

Lecture Notes in Morphogenesis

Series Editor: Alessandro Sarti

Jean Petitot

Elements of Neurogeometry

Functional Architectures of Vision



Springer

Lecture Notes in Morphogenesis

Series editor

Alessandro Sarti, CAMS Center for Mathematics, CNRS-EHESS, Paris, France
e-mail: alessandro.sarti@ehess.fr

More information about this series at <http://www.springer.com/series/11247>

Jean Petitot

Elements of Neurogeometry

Functional Architectures of Vision

 Springer

Jean Petitot
CAMS, EHESS
Paris
France

Translated by Stephen Lyle

ISSN 2195-1934 ISSN 2195-1942 (electronic)
Lecture Notes in Morphogenesis
ISBN 978-3-319-65589-5 ISBN 978-3-319-65591-8 (eBook)
DOI 10.1007/978-3-319-65591-8

Library of Congress Control Number: 2017950247

Translation from the French language edition: *Neurogéométrie de la vision* by Jean Petitot, © Les Éditions de l'École Polytechnique 2008. All Rights Reserved
© Springer International Publishing AG 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Of Interest for Neurogeometry by Jean Petitot

- o *Les Catastrophes de la Parole. De Roman Jakobson à René Thom*, Maloine, Paris, 1985.
- o *Morphogenesis of Meaning*, (trans. F. Manjali), Peter Lang, Bern, 2003.
- o *Neurogéométrie de la vision. Modèles mathématiques et physiques des architectures fonctionnelles*, Les Éditions de l'École Polytechnique, Distribution Ellipses, Paris, 2008.
- o *Cognitive Morphodynamics. Dynamical Morphological Models of Constituency in Perception and Syntax*, (with R. Doursat), Peter Lang, Bern, 2011.

- o (ed.) *Logos et Théorie des Catastrophes*, Colloque de Cerisy à partir de l'œuvre de René Thom, Éditions Patiño, Genève, 1988.
- o (ed., with F. Varela, J.-M. Roy and B. Pachoud) *Naturalizing Phenomenology. Issues in Contemporary Phenomenology and Cognitive Science*, Stanford University Press, 1999.
- o (ed. with J. Lorenceau), “Neurogeometry and Visual Perception”, Special issue of the *Journal of Physiology-Paris*, 97, 2003.
- o (ed., with M. Bitbol and P. Kerszberg) *Constituting Objectivity. Transcendental Perspectives on Modern Physics*, The Western Ontario Series in Philosophy of Science, vol. 74, Springer: Berlin, New York, 2009.
- o (ed. with A. Sarti and G. Citti), “Neuromathematics of vision”, Special Issue of the *Journal of Physiology-Paris*, 103, 1–2, 2009.

Contents

1 Preface	1
1.1 The Goal of This Work	1
1.2 An Outline of This Work	3
1.2.1 Outline of the First Volume	3
1.2.2 Some Remarks Concerning the Second Volume	7
1.2.3 Limits of This Investigation	9
1.3 History, Context, and Acknowledgements	10
References.....	19
2 Introduction	21
2.1 Origin of Space and Neurogeometry.....	21
2.1.1 Geometric, Physical, and Sensorimotor Conceptions of Space	21
2.1.2 The Neurogeometric Approach	24
2.2 Perceptual Geometry, Neurogeometry, and Gestalt Geometry. . . .	25
2.3 Geometry’s ‘Twofold Way’	27
2.4 Idealities and Material Processes	28
2.5 Mathematical Prerequisites and the Nature of Models	30
2.6 Mathematical Structures and Biophysical Data.....	32
2.7 Levels of Investigation: Micro, Meso, and Macro	33
2.8 The Context of Cognitive Science	34
2.9 Complex Systems and the Physics of the Mental	36
2.10 The Philosophical Problem of Cognitive Science	37
2.11 Some Examples	38
2.11.1 The Gestalt Concept of Good Continuation.....	38
2.11.2 Kanizsa’s Illusory Contours.....	39
2.11.3 Entoptic Phenomena	40
2.11.4 The Cut Locus.....	41
References.....	42

3	Receptive Fields and Profiles, and Wavelet Analysis	45
3.1	Structure of the Retino-Geniculo-Cortical Visual Pathways	45
3.2	Receptive Fields and Receptive Profiles	50
3.2.1	Structure of the Retina	50
3.2.2	Neurons and Action Potentials	51
3.2.3	Structure of the Photoreceptors	52
3.2.4	Ganglion Cells	58
3.2.5	Retinal Colour Coding Circuitry	59
3.2.6	General Receptive Fields and Neural Coding	62
3.3	Visual Neurons as Filters	73
3.3.1	Gabor Wavelets and Derivatives of Gaussians	73
3.3.2	Steerable Filters	76
3.3.3	Linearity Versus Nonlinearity	77
3.3.4	Visual Neurons as Convolution Operators	79
3.3.5	Fine Orientation Discrimination	84
3.4	Vision and Wavelets	85
3.4.1	Fourier, Gabor, and Wavelets	85
3.4.2	Wavelets and Group Representation	88
3.4.3	Wavelets and Discontinuities	89
3.4.4	Redundancy of Wavelets	89
3.4.5	Compression and Geometry	92
3.4.6	Matching Pursuit and Rank Coding	92
3.5	Feature Detectors	94
3.6	Receptive Profiles and Information Theory	95
3.6.1	Signal Decorrelation and Efficient Coding	95
3.6.2	Receptive Profiles and Natural Images	97
3.7	Signal Processing and Geometrical Formatting	101
3.8	Grid Cells and Place Cells	102
3.8.1	Spatial Navigation	102
3.8.2	Place Cells	103
3.8.3	Grid Cells	103
3.8.4	Head Direction Cells	107
3.8.5	Implementing the Tangent Bundle	108
	References	108
4	Functional Architecture I: The Pinwheels of V1	113
4.1	The Areas of the Visual Cortex	114
4.2	Hypercolumnar Structure of the V1 Area	120
4.3	V1 as a Mesoscopic Fibration	127
4.3.1	'Bridging Scales': The Mesoscopic Level	127
4.3.2	Fibrations and Engrafted Variables	127
4.3.3	Fibre Bundles	129
4.3.4	V1 as a Geometric Fibre Bundle	130
4.3.5	V1 as a 1-jet Fibre Bundle	132

4.3.6	Legendrian Lifts	135
4.3.7	Integrability Condition	135
4.3.8	$SE(2)$ Invariance of 1-jets	136
4.3.9	Generalizing the Model	137
4.3.10	Neurophysiology and Its Geometrical Idealization.	138
4.4	The Pinwheel Structure of $V1$	138
4.4.1	Observation of Pinwheels	138
4.4.2	Limitations of This Analysis	147
4.4.3	Functional Maps as Fields	147
4.4.4	Development of Pinwheels	151
4.4.5	Pinwheels and Evolution	156
4.4.6	End Points and Triple Points	158
4.4.7	Distortions and Defects in the Neighbourhood of the $V1/V2$ Boundary	162
4.5	Topological Universality of Pinwheels	165
4.6	Pinwheels as Phase Fields	170
4.6.1	Fields and Coordinates	170
4.6.2	Singularities of a Phase Field	171
4.6.3	Orientation and Iso-orientation Fields	173
4.6.4	Topological Charge and Index	174
4.6.5	Current, Vorticity, and Divergence	174
4.6.6	Helmholtz Equation	179
4.6.7	Illustration	180
4.6.8	Current Conservation	187
4.6.9	Critical Points	192
4.6.10	Mesogeometry and Microphysics	193
4.6.11	Statistics of Pinwheels as Phase Singularities	194
4.6.12	Pinwheels and Gaussian Fields	195
4.6.13	Evolution of Pinwheels as Phase Singularities	200
4.7	Pinwheel Singularities	201
4.7.1	Structure in the Vicinity of Singularities	201
4.7.2	The Problem of Resolution	210
4.7.3	Two-Photon Confocal Microscopy	212
4.8	Pinwheels and Blow-ups	214
4.8.1	The Geometric Concept of Blow-up	215
4.8.2	Blow-ups and Lines of Dislocations	217
4.8.3	From Blow-up to Fibre Bundle	217
4.8.4	Discrete Versus Continuous Models	219
4.9	Different Aspects of Pinwheels	220
4.9.1	Position–Orientation Independence and Local Triviality	220
4.9.2	Other Engrafted Variables	221
4.9.3	Spatial Frequency	223

4.9.4	Generality of Pinwheels.	239
4.10	Retinotopic Maps and Their Transversality.	240
4.10.1	Pinwheels and Ocular Dominance.	242
4.10.2	Independent Maps and Transversality Principle.	246
4.10.3	Binocularity.	250
4.10.4	Blobs and Colour	257
4.10.5	Functionality of Maps	259
4.11	Hemispheres and Callosal Connections.	259
4.12	Homogeneous and Inhomogeneous Qualities	263
4.12.1	Responses to Homogeneous Surfaces	263
4.12.2	Colour Processing	264
	References.	268
5	Functional Architectures II: Horizontal Connections and Contact Structure	275
5.1	From Pinwheels to Contact Geometry	275
5.1.1	Horizontal Intracortical Connections	275
5.1.2	Semi-local Structures.	276
5.1.3	Parallelism and Coaxiality.	278
5.2	Integration of Contours and Association Field	282
5.2.1	Some Experimental Facts	282
5.2.2	Pop-out, Perceptual Salience, and the Helmholtz Principle	285
5.2.3	Explanation in Terms of Association Fields	288
5.2.4	Confirmation by fMRI.	289
5.2.5	Relationship with the Horizontal Connections.	290
5.2.6	Discretization of the Contact Structure	292
5.2.7	Binding	293
5.2.8	Comparison with Other Data.	293
5.3	Some Effects of the Horizontal Connections.	295
5.3.1	Contextuality of the Receptive Fields	295
5.3.2	Line-Motion Illusion	296
5.4	Contact Structure	298
5.4.1	Integrability Condition and Contact Form	298
5.4.2	Contact Structure as a Cartan Connection	299
5.4.3	Non-integrability of the Contact Structure.	299
5.4.4	Polarized Heisenberg Group	300
5.4.5	Scale and Characteristic Vectors	301
5.4.6	Jets, Contact Geometry, and ‘Simplexity’	302
5.5	Illusory Contours as Sub-Riemannian Geodesics	303
5.6	Curvature Detectors and 2-jets	304
5.6.1	Data.	304
5.6.2	Curvature, 2-Jets, and Engel Structure.	305
5.6.3	‘Good Continuation’ and the Statistics of Natural Images.	306

5.7	Relationship with Wavelets.	307
5.8	Structure of the $V2$ Area	310
5.9	Colour and Area $V4$	313
5.9.1	Colour Constancy: Semir Zeki and René Thom	313
5.9.2	Objectivism Versus Subjectivism.	317
5.10	Motion and the MT Area ($V5$).	319
5.11	Models of Direction and Singularities of Functions	319
5.11.1	Projection of Singularity Lines from $V2$ to MT	320
5.11.2	Swindale’s Model	320
5.11.3	Singularities and Universal Unfoldings	323
5.11.4	Swindale Model (Continued).	327
5.12	Neural Morphogenesis and Its Genetic Control	331
5.12.1	Guidance of Axon Connections.	331
5.12.2	Transcription Factors and Homeoboxes.	332
5.12.3	Some Guidance Factors.	333
5.12.4	Neurogenesis of the Retina	334
5.12.5	Retinotopy and Neurogenesis of Visual Pathways.	337
5.12.6	Dynamical Models of Neural Guidance.	342
	References.	343
6	Transition to Volume II	347
6.1	Introduction	347
6.2	Geodesics of the \mathbb{V}_J Model.	348
6.3	The \mathbb{V}_S Model	354
6.4	Geodesics of the \mathbb{V}_S Model	356
6.5	Elastica Revisited	358
6.6	Sub-Riemannian Diffusion, Heat Kernel, and Non-commutative Harmonic Analysis	360
6.7	Confluence Between \mathbb{V}_J and \mathbb{V}_S Models.	363
6.8	Other Themes.	364
	References.	364
	Author Index.	367
	Subject Index.	375

About the Author

Jean Petitot was born in 1944, and he specialized in mathematical modelling in the cognitive sciences. Former student and teacher at the École Polytechnique, he is currently Professor at the Mathematics Center of the École des Hautes Études en Sciences Sociales in Paris. He is a member of the International Academy of Philosophy of Science. Having worked for several years on the theory of singularities in differential geometry, he was one of the first to become interested in René Thom's morphodynamic models of visual perception and phonetics in the 1970s.

He is the author of *Cognitive Morphodynamics* (Peter Lang, 2011), *Neurogéométrie de la vision* (École Polytechnique, Ellipses, 2008), *Physique du Sens* (Centre National de la Recherche Scientifique, 1992), *Morphogenèse du Sens* (Presses Universitaires de France, 1985; English transl. Peter Lang, 2004), five other books, and more than 300 papers. He is also co-editor of *Constituting Objectivity* (Springer, 2009), *Neurogeometry and Visual Perception* (J. of Physiology-Paris, 2003), *Au Nom du Sens*, a tribute to Umberto Eco (Grasset, 2000), *Naturalizing Phenomenology* (Stanford University Press, 1999), *Logos et Théorie des Catastrophes*, a tribute to René Thom (Patiño, 1988).

Keywords

Association field · Contact structure · Fibre bundle · Functional architecture
Gaussian field · Geodesic · Gestalt · Heisenberg group · Helmholtz equation
Frobenius integrability · Jet · Legendrian lift · Orientation (hyper)column
Parallel transport · Pinwheel · Receptive profile · Sub-Riemannian geometry
Transversality · Universal unfolding of singularity · Wavelet

Chapter 1

Preface

1.1 The Goal of This Work

This two-volume work synthesizes several years of research on the use of specific geometric structures in the neuroscience of vision. It is a new edition, greatly expanded and enriched. The first edition [1], based on my lectures *Introduction aux sciences cognitives* at the École Polytechnique in Paris, was published in French in 2008 by the *Éditions de l'École Polytechnique*.

In close relation with a comprehensive set of experimental data, the book presents several physico-mathematical models of the primary visual cortex—in particular, area V1, the first of the visual cortex areas—and suggests an original geometric model for its *functional architecture*, that is for the very specific organization of its neural connections.

Its purpose is to describe in detail certain geometric algorithms implemented by this functional architecture. Its primary concern is therefore the *internal* neural origin of *external* spatial representations, and this is why we coined the term *neurogeometry* to name its field of investigation. Neurogeometry thus deals with the ‘internal’ and ‘immanent’ geometric algorithms that allow the visual system to build the ‘external’ and ‘transcendent’ geometry of our surrounding world.

To the extent that the origin of spatial representations constitutes a major problem, not only scientifically but also philosophically, our work also has important epistemological consequences. We shall return to these in the second volume. Briefly, the idea is that the internal geometry ‘immanent’ in neural infrastructures corresponding to external ‘transcendent’ space can greatly clarify what is meant in philosophy by the ‘synthetic a priori’ character of this space, a character that has been flatly rejected for the past two centuries by nearly all philosophers of science, but which is nonetheless imposed by the results of contemporary neuroscience.

I have long argued that contemporary cognitive neuroscience supports many tenets of transcendental philosophy and that there are neurophysiological roots for what Kant called ‘pure intuitions’. Recently, some of the best neuroscientists have defended similar theses. For instance, in June 2014, John O’Keefe of the Institute

of Cognitive Neuroscience, University College, London, who was awarded the 2013 Horwitz prize ‘for significant advancements to the field of neuroscience’, gave a public lecture entitled “Immanuel Kant: Pioneer neuroscientist” [2]. The abstract was as follows:

In his *Critique of Pure Reason*, Kant argued that our concept of space was not derived from sensations arising from our interaction with the physical world but instead represented the a priori basis for our perception of the world in the first place. Extensive work in modern neuroscience has provided strong evidence in support of this position. We now know that there is an extensive network of brain areas in the temporal lobes dedicated to the construction of an allocentric space framework and that some parts of this network develop relatively independently of the animal’s experience. This map-like spatial representation is constructed from more primitive representations of places, directions and distances and allows the animal to know where it is in an environment and how to navigate to desired locations. In my talk, I will present the evidence for these more primitive representations and discuss how they may interact with each other to produce a Kantian map-like representation of space. In the latter part of my talk I will discuss how our understanding of these brain systems sheds light on some of the postulates of Euclidean geometry, one of the conceptual domains used by Kant to support his view of the synthetic a priori nature of our spatial representations.

In the same year (2014), John O’Keefe was awarded the Nobel prize for his discovery of ‘place cells’ in 1971, along with May-Britt and Edvard Moser for their discovery of ‘grid cells’ in 2005. We will discuss place cells and grid cells in Sect. 3.8.

In 2012, another great specialist of cognitive neuropsychology, Stanislas Dehaene, professor at the Collège de France in Paris and member of the French Académie des Sciences, proposed a ‘Kantian’ research program as part of the *Human Brain Project*, the most important European Future and Emerging Technologies (FET) flagship in neurosciences. With Elizabeth Brannon, he formulated this proposal in *Trends in Cognitive Sciences, Special Issue: Space, Time and Number* [3]. The title was ‘Space, time, and number: a Kantian research program’ with the following abstract:

What do the representations of space, time and number share that might justify their joint presence in a special issue? In his *Critique of Pure Reason*, Immanuel Kant famously argued that they provide ‘a priori intuitions’ that precede and structure how humans experience the environment. Indeed, these concepts are so basic to any understanding of the external world that it is hard to imagine how any animal species could survive without having mechanisms for spatial navigation, temporal orienting (e.g. time-stamped memories) and elementary numerical computations (e.g. choosing the food patch with the largest expected return). In the course of their evolution, humans and many other animal species might have internalized basic codes and operations that are isomorphic to the physical and arithmetic laws that govern the interaction of objects in the external world. The articles in this special issue all support this point of view: from grid cells to number neurons, the richness and variety of mechanisms by which animals and humans, including infants, can represent the dimensions of space, time and number is bewildering and suggests evolutionary processes and neural mechanisms by which Kantian intuitions might universally arise. [...] If Immanuel Kant were born today, he would probably be a cognitive neuroscientist!

1.2 An Outline of This Work

1.2.1 Outline of the First Volume

After a general introduction focusing on the origin of spatial representations, the structures of perception, cognitive sciences, and the way the neurogeometry of vision fits into them, we devote the first three chapters of this first volume to neurophysiological experimental data and a first stage of modelling.

We present many results, obtained by the best specialists in the field. This is absolutely necessary, for two reasons. First, to raise the reader's awareness of the extraordinary richness and neuronal complexity of the visual system. Second, to give a sufficiently detailed justification for the construction of the models and the choice of geometrical tools enabling their construction. Even though this experimental incursion will be rather broad, it must nevertheless remain somewhat superficial, simply because the subject is so vast. That said, it will not just be a scholarly compilation, because we will put the results in perspective with respect to the geometric approaches that we develop. Their presentation will be consolidated by our desire to formulate adequate models.

In Chap. 3, we discuss receptive fields and receptive profiles of visual neurons, starting with the photoreceptors and the ganglion cells of the retina and proceeding via the neurons of the lateral geniculate nucleus to those of $V1$. We explain how they act on the optical signal as filters, and we broach the problem of their linearity or nonlinearity. In the linear case, the effect they have on the signal falls to a large extent under the rule of what is known as 'wavelet analysis' in signal processing. We shall give an outline of this fundamental notion and then discuss how receptive profiles can be interpreted within the framework of information theory as a means of optimizing the processing of *natural* images, which have very particular statistical properties. To a certain extent, the geometric formatting of the signal by means of a certain kind of wavelet (essentially, partial derivatives of Gaussians¹) is a way of optimizing the *compression* of natural images.

In Chap. 4, we present a first set of experimental data on the *functional architecture* of area $V1$ and, in particular, on what is referred to as its 'pinwheel' structure. If we simplify their highly complicated activity, we may say that most of the so-called simple neurons in $V1$ detect positions and orientations² in the visual field, with those detecting the various orientations for *each* position being grouped together into functional micromodules that can be defined anatomically and are called *orientation hypercolumns* or pinwheels. In this sense, $V1$ implements a discrete approximation of the *fibration* $\pi : \mathbb{V} = R \times \mathbb{P}^1 \rightarrow R$, with the retinal plane R as its base space and

¹Named after Carl Friedrich Gauss.

²In this book, we use the term 'orientation' in the same way as neurophysiologists to refer to straight lines in the plane with angles specified modulo π . Orientations can thus have two senses with angles specified modulo 2π . Mathematicians speak rather of 'direction' in the first case and orientation in the second case. However, neurophysiologists reserve the term 'direction' for the direction of motion of an oriented segment.

the projective line \mathbb{P}^1 of orientations in the plane as its fiber. This fiber bundle \mathbb{V} has an abstract dimension of 3 (two degrees of freedom for position and one degree of freedom for orientation).

The geometric notion of fibration is a first key geometric concept that makes it possible to model the ‘vertical’ retino-geniculo-cortical functional architecture of $V1$. We then explain something that is well known to geometers: to use their language, this fibration is nothing other than the space of 1-jets of curves in the plane. The retinotopic map that projects curves from the visual field onto chains of activation of neurons in $V1$ is thus geometrically identified with a process called ‘Legendrian³ lifting’ of planar curves to the bundle of their 1-jets.

A substantial part of this chapter is devoted to the fact that the pinwheels of $V1$ implement a *dimensional reduction* of the bundle of 1-jets, this reduction being necessary in order to solve a key problem: the two-dimensional layers of $V1$ implement a fibration of abstract dimension (at least) 3. After reduction, $V1$ then appears to be a field of orientations in 2 dimensions (called ‘orientation maps’ by neurophysiologists), a field whose *singularities* are the centres of the pinwheels. We study these singularities, give their normal forms, and specify the distortions and defects of their networks. We present many accurate experimental data on this matter, as well as some hypotheses regarding their evolutionary origins, such as minimization of wiring.

One way to model these orientation maps is to treat them as *phase fields*, analogous to those encountered in optics, whose singularities (called in this case ‘dislocations’) have been thoroughly analyzed by specialists such as Michael Berry. These fields are superpositions of solutions to the *Helmholtz equation*, whose wave number depends in a precise manner on the mesh of the pinwheel lattice. They enable the construction of very interesting models, such as those proposed by Fred Wolf and Theo Geisel. We explain these models with a fair amount of mathematical detail because we will not return to them in the second volume.

However, in these models of phase fields, orientation selectivity must vanish at singularities. Yet many experimental results show that this is not the case. We thus present another model based on the geometric notion of *blow-up*. As the mesh of the pinwheel lattice tends to 0, these alternative discrete models converge to the bundle of 1-jets.

In this chapter, we also explain how the fibration that models the orientation variable interferes with other fibrations (other visual ‘maps’) that model other variables such as direction, ocular dominance, phase, spatial frequency, or colour. For spatial frequency, we will present the dipole model proposed by Daniel Bennequin. $V1$ therefore implements fibrations of rather high dimension $2 + m$, $m \geq 3$, in 2-dimensional layers. This leads to the very interesting problem of knowing how to express the *independence* of these different variables. A plausible hypothesis relies on a *transversality principle*: the transversality of the level curves of two independent maps is maximal (indeed they are nearly orthogonal) in the areas of the base plane where the two gradients are *jointly* strong.

³Named after Adrien-Marie Legendre.

We conclude Chap. 4 with data on two other aspects of neurobiology: (i) the relation between the cerebral hemispheres through *callosal connections* and (ii) the primary processing of *colour* in the ‘blobs’ of V1.

In Chap. 5, we enrich and complete the data assembled in Chap. 4 by studying the *second* component of the functional architecture of primary visual areas, namely what are called *cortico-cortical ‘horizontal’ connections*. For the orientation hypercolumns, these connections join together neurons that detect orientations at different points which are *parallel*, and not only parallel, but even approximately *aligned*. This ‘parallel transport’ reinforced by a rough coaxiality has been confirmed by numerous psychophysical experiments on what David Field, Anthony Hayes, and Robert Hess called the *association field*. We explain its principles and results.

Thus we show that, to a first approximation, parallel transport can be considered as an implementation of what is known in differential geometry as the *contact structure* of the fibre bundle \mathbb{V}_J of 1-jets of curves in the plane R . \mathbb{V}_J is isomorphic to a space \mathbb{R}^3 with coordinates (x, y, p) , where (x, y) are the spatial coordinates and p is the angular coordinate $\tan \theta$. Its contact structure lies at the heart of our ‘neurogeometric’ approach. It is invariant under the action of the group $SE(2)$ of isometries of the Euclidean plane. In fact, it is composed of the field \mathcal{K} of planes generated by the action of $SE(2)$ on the plane K_0 of points (x, p) through the origin. Since K_0 is the plane $y = 0$, or in other words the kernel of the differential 1-form dy , \mathcal{K} is the kernel of the differential 1-form $\omega = dy - p dx$ obtained by translating dy by the elements of $SE(2)$.

What is more, the contact structure is associated with a *non-commutative group structure* on \mathbb{V}_J , the action of $SE(2)$ then being identified with the left translations of \mathbb{V}_J on itself. This group is isomorphic to the ‘*polarised*’ *Heisenberg group*, which is well known to physicists. When endowed with its field of contact planes, this is a *nilpotent* group belonging to the class of what are called *Carnot groups*. We use this to transfer to neurogeometry an important and beautiful set of results from mathematical physics. (In a similar way, the discovery of the fact that the Wilson–Cowan–Hopfield equations of neural networks were analogous to the equations of spin glasses made possible a massive transfer of results from statistical physics to the theory of neural networks.)

The definition of an $SE(2)$ -invariant metric (i.e. invariant under left translations) on the contact planes thus defines what is called a *sub-Riemannian*⁴ geometry. We show that, in this first neurogeometric approach, the natural mathematical framework for modelling the functional architecture of V1 is provided by the contact geometry of \mathbb{V}_J and the sub-Riemannian geometry that is closely associated with it. This type of sub-Riemannian geometry (also called Carnot–Carathéodory geometry) has been extensively studied by many geometers. For our part, we were greatly inspired by Misha Gromov and, subsequently, specialists such as Andrei Agrachev, Richard Beals, André Bellaïche, Ugo Boscain, Giovanna Citti, Jean-Paul Gauthier, Bernard Gaveau, Peter Greiner, Richard Montgomery, Pierre Pansu, Yuri Sachkov, Alessandro Sarti, and Robert Strichartz.

⁴Named after Bernhard Riemann.

Using the sub-Riemannian geometry of the contact structure, one obtains a natural explanation for the enigmatic phenomenon whereby the visual system constructs long-range *illusory contours*. These phenomena of filling in gaps in sensory data by the integrative mechanisms of perception have been known since the beginnings of the Gestalt theory, and they have given rise to countless studies, particularly since the celebrated work by Gaetano Kanizsa. Together with our colleague Jacques Ninio at the École Normale Supérieure, we have carried out experiments on illusory contours *with curvature* in order to test what kinds of curves one observes. We were naturally led to *variational models*. At the end of the 1980s, David Mumford already proposed a first model, defined in the plane \mathbb{R}^2 of the visual field, which minimized a certain functional of curve length and curvature. We propose a model in the contact bundle \mathbb{V}_J that rests on the hypothesis that illusory contours are sub-Riemannian *geodesics*.

After this initial analysis of the $V1$ contact geometry and the sub-Riemannian geometry associated with it in the model \mathbb{V}_J , we will say a word about the ‘prolongation’ of this structure, called the *Engel structure*, which corresponds to the 2-jets of curves in the plane \mathbb{R}^2 . This prolongation must be taken into account if we want to report on the experimental data showing that, in the primary visual system, there exist not only orientation detectors (tangents and 1-jets), but also *curvature* detectors (osculating circles and 2-jets).

To conclude this chapter, we then point out a few properties of the functional architecture of areas $V2$, $V4$ (for colour), $V5$ or MT (for motion). We also discuss Swindale’s model for *directions*.

Finally, we briefly describe the genetic control of the *neural morphogenesis* of the functional architectures whose geometry we have analyzed. The neurogenesis of visual pathways and axon guidance are marvels of self-organization that make it possible to glimpse the origin of spatial representation.

To conclude this first volume, in the final chapter we describe the transition to the mathematical themes that will be developed in the second volume.

We first outline the explicit calculation of geodesics in the \mathbb{V}_J model using the tools of control theory. To do this, we adapt to the *polarised* Heisenberg group results obtained for the Heisenberg group \mathbb{H} by a long line of specialists such as, in addition to those already mentioned, Jean-Michel Bismut, Luca Capogna, Vladimir Gershkovich, John Mitchell, Anatoly Vershik, and others, following the groundbreaking work by pioneers such as Lars Hörmander, Elias Stein, and Gerald Folland. It is then a simple matter to calculate all the elements of the sub-Riemannian geometry of \mathbb{V}_J , including its geodesics, unit sphere, wave front, caustic, cut locus, conjugate points.

However, the most natural model is not \mathbb{V}_J , constructed on the space of 1-jets by the action of $SE(2)$, but rather \mathbb{V}_S , constructed on $SE(2)$ itself. Here, we make the standard transition from a fibre bundle with a structure group to the associated *principal* bundle. With Alessandro Sarti and Giovanna Citti, we stressed the importance of this second model. Now, the isometry group $SE(2)$ of the Euclidean plane, with its so-called shift-twist action on the bundle, *is not nilpotent*. Its ‘nilpotentization’, which defines its ‘tangent cone’ at the origin, is indeed isomorphic to the polarised Heisenberg group \mathbb{H} , but globally, it has a very different sub-Riemannian geometry.

We observed with some surprise that this aspect had never been explored up until then, even though $SE(2)$ is a well-known and elementary group. Andrei Agrachev and coworkers (in particular Yuri Sachkov, Jean-Paul Gauthier, Ugo Boscain, and Igor Moiseev) subsequently examined this geometry, and we shall summarize their investigation of its geodesics.

We shall also discuss the fact that, in the present context, the sub-Riemannian contact geometries model a functional architecture of connections between neurons which act as filters. This means that the natural mathematical framework for low-level visual perception is the one in which *non-commutative harmonic analysis* on the group $SE(2)$ is related to its sub-Riemannian geometry. Since harmonic analysis is carried out using a ‘mother’ wavelet translated by $SE(2)$, i.e. by what is known as a *coherent state*, it is thus the relationship between coherent states and sub-Riemannian geometry which turns out to be the natural mathematical framework for low-level vision.

We then explain how David Mumford’s variational model, based on the *elastica* functional, and its stochastic interpretation can be reformulated in this novel context. As shown by Giovanna Citti, Alessandro Sarti, Remco Duits, and Markus van Almsick, this leads to advection–diffusion algorithms described by a Fokker–Planck equation which can be calculated explicitly for the \mathbb{V}_J model (while the calculation in \mathbb{V}_S remains very complicated).

Such techniques belong to the general theory of the *heat kernel* for the *hypoelliptic Laplacians*⁵ of sub-Riemannian manifolds. There has been an enormous literature on this subject since Hörmander’s pioneering work. Here, we discuss in particular results by Giovanna Citti and Andrei Agrachev’s group on the hypoelliptic Laplacians and heat kernels of the \mathbb{V}_J and \mathbb{V}_S models. To do this, one must obtain the irreducible unitary representations of these groups, i.e. their *dual* space equipped with a Plancherel measure, and the Fourier transforms of the Laplacians. This brings in the *Hermite functions* in the case of \mathbb{V}_J and the *Mathieu functions* in the case of \mathbb{V}_S .

We then explain how one can interpolate between \mathbb{V}_J and \mathbb{V}_S using a *continuous* family of sub-Riemannian models.

1.2.2 Some Remarks Concerning the Second Volume

In the second volume, we shall return in detail to these topics and focus more closely on *physical* models of $V1$. In this first volume, we concentrate mainly on analyzing the *geometrical* structure of the functional architecture. However, this structure connects neurons which are dynamical units with internal states and which interact together via the couplings determined by the ‘synaptic weights’ of their connections. To take these physical aspects into account, we shall use Hodgkin–Huxley-type equations to describe the dynamics of the action potentials generated by individual

⁵Named after Pierre-Simon de Laplace.

neurons and Wilson–Cowan–Hopfield equations to describe *neural nets*, applying ‘spin glass’ type models to describe the statistical physics of such networks. We shall also discuss the theory of *synchronization* of oscillator networks developed since the pioneering work by Yoshiki Kuramoto, Hiroaki Daido, Bard Ermentrout, and Nancy Kopell, among others.

We shall dwell particularly on the remarkable work carried out by Bard Ermentrout, Jack Cowan, Paul Bressloff, and Martin Golubitsky on the dynamics of Hopfield networks whose synaptic weights *encode the functional architecture of V1*. These authors have shown that, under a self-excitation of *V1*, even in the absence of any stimulus, the ground state of the network can spontaneously *bifurcate* into states corresponding to perceptual patterns with well-defined geometry. Several such patterns have long been known in visual hallucinations.

We shall also present our work with Alessandro Sarti and Giovanna Citti. This shows how, by introducing a *scale parameter* (defined by the width of the neural filters), one can move naturally from the contact structure of the fibration \mathbb{V} to a *symplectic structure* defined on $\mathbb{V} \times \mathbb{R}$, which is the symplectized space. As in any symplectic structure, a structure of the kind found in geometrical optics can be specified, and when a shape bounded by a closed contour in the visual plane R is treated in this way, we obtain the *cut locus*—also known as the ‘skeleton’ or ‘generalized symmetry axis’—of the shape which, as we know from the fundamental work by Harry Blum and René Thom, plays a key role in the perceptual analysis of shapes.

We shall also explain the basics of *multiscale* differential geometry, or ‘scale-space analysis’, which is needed when we consider optical signal processing by neurons acting by convolution as filters. As an example, we discuss James Damon’s application of the Thom–Mather singularity theory to the multiscale context.

We shall also introduce *image segmentation algorithms*. There are two closely related classes. The first consists of models in which segmentation is achieved by applying *nonlinear anisotropic diffusion equations*. These homogenize regions where the signal varies little, but introduce clear discontinuities where the signal varies sharply. The second class contains the *variational models* introduced by David Mumford and Jayant Shah. Here, we discuss the difficult subject of the ‘Mumford conjecture’ on the one hand and also work by Alessandro Sarti and Giovanna Citti which shows that the Mumford–Shah model is a limit (in the variational sense) of an oscillator synchronization model.

Further, in the second volume, we shall tackle certain epistemological issues, from the phenomenology of perception in the sense of Husserl and Merleau-Ponty to the Kantian problem of transcendental aesthetics. We shall show that the possibility of deriving the global morphology and geometry of percepts from what might be called ‘neural mesophysics’ confirms a strictly naturalistic *emergentist* understanding of perceptual consciousness and invalidates the claims of many philosophers who deny the reducibility of consciousness to neural activity.

Finally, we shall show how neurogeometry and the related neural materialism can provide a *naturalistic foundation* for the most profound philosophical approaches to perception and space, namely the transcendental approaches. We shall discuss the

phenomenology of perception as developed by Husserl, stressing just how close it comes to the neurogeometric understanding put forward in the preceding chapters. We shall end by going back to the Kantian problem of *transcendental aesthetics* and the claim that space is essentially of a ‘synthetic a priori’ nature, showing that all the neurophysiological data and all the models we present do indeed point in this direction.

As we said at the outset, this claim is in our opinion the main philosophical message of this work, and we are profoundly convinced that it should have far-reaching consequences when we realize that, for over a hundred years now, the whole area of philosophy of mind has been built upon the demise of transcendental aesthetics and the promotion of purely logical and conceptual principles of mental activity. On the contrary, neurogeometry shows that space is indeed a ‘pure intuition’, that is a non-conceptual, antepredicative, and prejudicative *format* for sensory data, that this format is specified by the functional architectures of the primary visual areas, and that, insofar as the latter are products of evolution, they are innate and ontogenetically a priori as regards perceptual content and judgement.

The *formats*, i.e. the functional architectures, which are synthetic and a priori for the percepts of an adult organism, are *not* ‘judgements’. There are doubtless no synthetic a priori *judgements*, but there are synthetic a priori formats, ‘forms of intuition’ or ‘pure intuitions’.

In these two volumes, we shall aim to carry out for the neuroscience of vision what has been done in fundamental physics, and more exactly, in field theory, with the development of *gauge theory*; that is, we shall geometrize the functional architectures governing the interactions between neurons. Indeed, as was well understood by the founders of the Gestalt theory, these interactions operate like ‘fields’, even though they propagate along material connections in this case. The geometrization of neurophysics into neurogeometry comes with a bonus, the transfer of the methods of differential geometry to this area.

1.2.3 Limits of This Investigation

This work is cross-disciplinary in a rather special sense. Indeed, it does not deal with general methodological or epistemological problems common to several disciplines. What we have here is rather an ‘intrinsic’ cross-disciplinarity imposed by the subject under investigation. It would thus be better to say that it is ‘multispecialized’ or perhaps ‘polyscientific’. There is an analogy with the biathlon in sport, which brings together two rather different, even antagonistic, skills, with the result that even the best specialists in this ‘polysportive’ discipline could not be the best specialists in both cross-country skiing and rifle shooting. The same is true here. I have gone to considerable lengths to bring together the skills required for cognitive neuroscience, mathematical modelling, and the philosophy of space, but I am fully aware of the limitations of the result. In each of these fields, one could certainly do a lot better.

So the achievement of this work, if indeed there is one, would be to bring together these highly disparate disciplines.

My main hope is simply that this book will allow mathematicians, physicists, and philosophers and historians of science who are interested in the foundations of geometry—from Euclid to Riemann, Helmholtz, Klein, Poincaré, Hilbert, and Cartan—to realize just how greatly neuroscience is likely to revolutionize hundreds of years of human understanding.

1.3 History, Context, and Acknowledgements

The considerations described here would not have been possible without the collaboration of a great many colleagues in many different disciplines. I extend my warmest thanks to them.

My interest in the geometry of perception goes back a long way to my encounter with René Thom and his dynamical models of morphogenesis at the end of the 1960s. I had begun doing research at the *Centre de Mathématiques* just set up by Laurent Schwartz at the École Polytechnique and started working on the theory of singularities in algebraic geometry with Jean Giraud (a disciple of Alexandre Grothendieck). There I was fortunate enough to follow courses given by two leading authorities, Heisuke Hironaka and René Thom.⁶

René Thom used the tools of singularity theory (of which he was one of the main inventors, following Marston Morse and Hassler Whitney) to explain how morphologies and patterns can appear and develop in material substrates M . The key idea, and we shall return to this several times, was that, at every point a of M , the physical, chemical, or metabolic properties of the substrate are described by an attractor A_a of an ‘internal’ dynamics X_a and that the dynamics of neighbouring points are coupled. Then for some critical values a_c of a , bifurcations can happen, the attractor A_a being replaced abruptly by another attractor B_a . The subset K of the a_c can be very complex (fractal, Cantor, etc.), but in not too complex cases it *stratifies* M and breaks its homogeneity, whereupon this *symmetry breaking* generates a morphology. In that sense, any morphology is a segmentation of the qualities of a substrate by a set of qualitative discontinuities.

Thom’s models (see [4, 5] and my surveys [6–8]) constitute a broad extension of the pioneering *reaction–diffusion* models introduced in 1952 by Alan Turing [9]. The challenge is the same: how can ‘the chemical information contained in the genes’ be ‘converted into a geometrical form’. For Turing, the internal dynamics are systems of (nonlinear) differential equations modelling the chemical reactions between ‘morphogens’ inside the substrate, the spatial coupling is afforded by diffusion, and the cause of the ‘patternized’ morphologies is the breaking of homogeneity induced by ‘diffusion-driven instabilities’ (see my survey [10]). The same great biologists

⁶For more on the historical background of neurogeometry, the interested reader is referred to my ‘Landmarks for Neurogeometry’ [11].

inspired both Turing's and Thom's projects: Sir D'Arcy Wentworth Thompson and Conrad Hal Waddington.

Thom's models were called 'catastrophic' by Christopher Zeeman. I preferred the term 'morphodynamical'. Morphodynamical models belong to a mathematical universe which experienced an extraordinary development in the 1960s–1970s under the leadership of masters such as René Thom, Bernard Malgrange, John Mather, Christopher Zeeman, Vladimir Arnold, Stephen Smale, David Ruelle, David Mumford, John Milnor, Martin Golubitsky, Robert MacPherson, and many others. Their main tools were as follows:

- The theory of dynamical systems (for the study of internal dynamics), their attractors (possibly with strong ergodic properties), their structural stability properties, and their bifurcations.
- The theory of critical points of differentiable mappings (when the internal dynamics is gradient-like).
- The geometrical theory of jet spaces and their stratifications.
- The universal unfoldings of finite codimension singularities.

They aim at a mathematical understanding of morphogenesis and, going beyond biological morphogenesis, of morphological *structures* whatever their substrate may be, from physics to cognitive sciences (see my survey [12]).

Thom's models were inspired by a deep and rather universal mathematical 'philosophy' elaborated in the 1950s and 1960s (see, e.g., his two classic papers 'Les singularités des applications différentiables' in 1956 [13] and 'Ensembles et morphismes stratifiés' in 1968 [14]):

- The singularities of a space or of a map between spaces concentrate information about its *global* structure into *local* singular morphologies.
- To analyze a differentiable map $f : M \rightarrow N$ locally, it is efficient to look at its successive jets, these being defined up to a change of coordinates (a diffeomorphism) in M and in N . In general, jets of sufficiently high order are 'determined' and f can be reduced locally to an algebraic 'normal' form.
- Jet spaces $J^k(M, N)$ of successive order k are manifolds *stratified* by a stratification Σ_k whose strata of increasing codimension (of decreasing dimension) correspond to more and more singular singularities.
- One of Thom's most fundamental theorems is the *transversality theorem* in jet spaces. The k -jet $j^k(f)$ of f is a map $j^k(f) : M \rightarrow J^k(M, N)$, and the theorem says that, *generically*, $j^k(f)$ is transversal on Σ_k . A consequence is that f cannot generically have singularities of codimension $> \dim(M)$.

Very soon, having joined the École des Hautes Études en Sciences Sociales (EHESS) Center for Mathematics (CAMS) in 1971, I focused on the applications of Thomian morphodynamics to models in the cognitive sciences and semiolinguistics and also on their far-reaching epistemological consequences.

In order to build such models, it was crucial to maintain in close contact with pure mathematics. In those early days, I was able to do this thanks to the colleagues

I worked with in Laurent Schwartz's laboratory: Bernard Teissier (specializing in the study of singularities in algebraic geometry along the lines of Zariski, Hironaka, and Mumford), Alain Chenciner (an expert on singularity theory and the theory of dynamical systems who would later turn to difficult Poincaré-type problems in celestial mechanics), Jean-Pierre Bourguignon (an expert on Hamiltonian⁷ mechanics and future director of the Institut des Hautes Études Scientifiques, IHÉS), Jean-Marc Deshouillers (an arithmetician who later became interested in cognitive science), Lê Dũng Tráng (an algebraic geometer who would subsequently direct the Mathematics Department at the *Abdus Salam International Center for Theoretical Physics* in Trieste), and then younger colleagues like Daniel Bennequin (an expert, among other things, on symplectic geometry and contact geometry who has been working in neuroscience for a few years now, in close collaboration with Alain Berthoz at the Collège de France) and Marc Chaperon (specializing in singularity theory).

Over many years, Thom's weekly seminar on singularity theory at the IHÉS allowed me to remain fully up to date with the latest research in this area and to learn so much from presentations by so many renowned geometers. It was an honour to be able to attend lectures given by such great mathematicians and physicists, including David Ruelle, Stephen Smale, David Mumford, Misha Gromov, and Alain Connes.

A great deal of cross-disciplinary activity was generated around singularity theory and its applications in the 1970s. It started out within mathematics (including mathematical physics, e.g. caustics in optics, defects in ordered media, critical phenomena, and phase transitions) with major conferences such as the *Liverpool Singularities Symposium*, organized and published by C.T.C. Wall in 1971, or the Summer Schools at the Institut d'Études Scientifiques in Cargèse, organized in 1973 and 1975 by Frédéric Pham. It evolved in an extremely rich scientific context, bringing together several traditions. The morphodynamic models came into it, sometimes controversially, as, for example, with the dissipative structures in thermodynamics and chemistry studied by Ilya Prigogine, Grégoire Nicolis, and Isabelle Stengers in Brussels, the work on self-organization by Henri Atlan, Jean-Pierre Dupuy, and Francisco Varela, the synergetics of Hermann Haken and Scott Kelso, etc. This was a whole new paradigm in science, which set out to understand the emergence of morphology in physics, chemistry, and biology.

There were exciting philosophical consequences, too. The question of a dynamics of shapes had been neglected since the revolution brought about by Galileo and Newton in the seventeenth century, when it became possible to develop a mechanics based on forces. This was because such a dynamics would have involved certain finalistic Aristotelian concepts like entelechy. From a philosophical point of view, many great thinkers had been aware of this. To begin with, there was Leibniz, who spent his whole life reflecting on this problem, then Diderot (see his debate with d'Alembert), Kant (see his *Critique of Judgement*), Geoffroy Saint-Hilaire, Goethe (inventor of modern structural morphology), Brentano in psychology, Husserl with his phenomenology, Gestalt theory, and D'Arcy Thompson and Waddington (already mentioned). By showing how it was possible to construct a *mathematical* dynamics

⁷Named after William Rowan Hamilton.

of shapes which was compatible with physics and could be extended, by virtue of its transphenomenal nature, to psychology and the social sciences, René Thom made a philosophical breakthrough which did much to dissolve the traditional boundaries between the natural sciences and the sciences of the mind.

I thus decided to concentrate my research on this *unification* of the issues of shape and morphogenesis that would involve *all* the empirical disciplines. To get some idea of the scope and the possibilities, in 1982, I organized the Cerisy conference entitled *Logos and Catastrophe Theory* [15], a tribute to René Thom, which was attended by Christopher Zeeman and many other mathematicians, but also physicists such as David Ruelle and Michael Berry, biologists specializing in morphogenesis such as Yves Bouligand and Brian Goodwin, philosophers of science, semiolinguists, and experts on Aristotle.

It was in the 1980s that my research inspired by Thom in perceptual geometry began to connect with parallel studies by experts working on natural and computational vision and image processing. I would like to mention here those that most influenced me.

To begin with, in the mid-1980s, I discovered the work of William Hoffman, one of those pioneering the application of differential geometry and Lie group theory to vision, and Jan Koenderink (and later the group at the University of Utrecht, and in particular, Luc Florack) who studied the geometry of visual perception. For the first time, I was to meet experts in perceptual psychology using the methods and resources of differential geometry, and in particular, the theory of singularities. Among the fundamental achievements of these scholars, I would like to mention the following:

- The idea that the visual cortex is a ‘*geometric engine*’, implementing structures such as fibrations, jet spaces of order 1 and 2, and contact structures [16, 17].
- The key structural role of singularities in perceptual geometry [18].
- The need to *integrate* (in the mathematical sense) local neuronal detections into global geometric structures. If the brain is indeed a ‘*geometric engine*’, this is because groups of receptive fields of visual neurons can detect ‘local features’ such as edge orientations, crossovers, points of inflection, which can subsequently be integrated through the functional architectures connecting these detectors together in a very specific way.
- The key role played by *scale*. Perceptual geometry results from the integration of local detections by receptive fields which have a certain width and so occurs at a certain scale, i.e. with a certain resolution. Perceptual differential geometry must therefore be *multiscale*, while conventional differential geometry corresponds to the idealization of infinite resolution (scale 0). Koenderink and Witkin thus introduced the method of *scale-space analysis*. It consists in uniformly parametrizing all the relevant geometrical structures by a scale parameter σ by imposing a so-called causality constraint; i.e., when σ increases, the complexity of the geometrical structures *simplifies*.

In general, the latter constraint is expressed using a diffusion operator D , the simplest of which is the one corresponding to the heat equation (Gaussian blurring). However,

Gaussian blurring has a major disadvantage in image processing and computational vision, in that it does not respect image morphology. Indeed, as we have seen, the latter is dominated by the perceptual salience of qualitative discontinuities, and by definition, the isotropic diffusion induced by the heat equation smooths out these discontinuities. This is why a certain number of specialists based scale-space analysis on highly *anisotropic* nonlinear parabolic diffusion equations in which the image intensity gradient actually *inhibits* diffusion. Along discontinuities, this gradient is very large or even divergent, whence there is no diffusion: in fact, diffusion can only occur in a direction transverse to the discontinuity, whence morphology is preserved. These kinds of partial differential equation (PDE) are difficult to integrate numerically because they involve reverse diffusion and deconvolution which makes them unstable. A certain class of such equations has been studied by geometers such as Michael Gage, Richard Hamilton, Matthew Grayson, Lawrence Evans, and Joel Spruck under the headings of ‘curve shortening’, ‘flow by curvature’, and ‘heat flow on isometric immersions’.

But as the term suggests, neurogeometry is about the *neural* implementation of the algorithms of perceptual geometry and is thus based upon neuroscience. Regarding my own relationship with neuroscience, the decisive event came in 1990 when Michel Imbert started a *Diplôme d’études approfondies* (DEA, now a Masters) in cognitive science at the École des Hautes Études en Sciences Sociales, the University of Paris VI, and the École Polytechnique, given the explosive development of these new disciplines on the international level. Imbert, an eminent neurophysiologist specializing in vision, but also a maths and philosophy enthusiast, succeeded in the difficult task of institutionalizing this emergent interaction of disciplines and realizing a whole range of innovations in the fertile soil of neuroscience, psychology, linguistics, physico-mathematical modelling, logical and programming formalisms, and philosophy. Associated with the ‘Cerveau, Cognition, Comportement’ (3C) doctoral school of Paris VI (Jean-Francois Allilaire, Philippe Asher, Marie-Jo Besson, Alain Prochiantz, and Danièle Tritsch, among others), this DEA, which I had the honour of running after Michel Imbert, gave me a way of redirecting my research and teaching, at least in some respects, and also the possibility of working closely with some of the best specialists in neuroscience. In particular, I would like to thank Yves Frégnac, director of the Unité de Neurosciences Intégratives et Computationnelles (UNIC⁸) where Alain Destexhe⁹ also worked, Alain Berthoz, director of the Laboratoire de Physiologie de la Perception et de l’Action (LPPA) and member of the Collège de France, my colleagues at the Laboratoire de Neurosciences Cognitives et Imagerie Cérébrale (LENA, CNRS), the research institute of Jean Lorenceau (who was also a member of the LPPA and UNIC), and my late friend Francisco Varela.

Being a member of the Centre d’Analyse et de Mathématique Sociales (CAMS) of the EHESS, founded by Georges Th. Guilbaud and Marc Barbut, helped me to engage in this exciting scientific and institutional adventure. I would therefore like

⁸Under the aegis of the Centre National de la Recherche Scientifique (CNRS), and today called the Unité de Neurosciences, Information et Complexité.

⁹Today, director of the European Institute for Theoretical Neuroscience (EITN) in Paris.

to thank the following directors Pierre Rosenstiehl and Henri Berestycki.¹⁰ In 1999, a special issue of the CAMS publication *Mathématiques, Informatique et Sciences Humaines* carried the first review of this work [19].

I was also greatly assisted by my membership of the Centre de Recherche en Epistémologie Appliquée at the École Polytechnique (CREA),¹¹ which I joined as an associate researcher in 1986 when Daniel Andler set up a cognitive science team there. I already knew the CREA and its founding director Jean-Pierre Dupuy because I had taken part in debates with him since the 1970s about models of self-organization and morphogenesis (René Thom, Henri Atlan, Ilya Prigogine). The exceptional scientific profile of Jean-Pierre Dupuy played a decisive role in precipitating my move to neurocognitivism, and it was therefore a great honour to become his successor as director of the CREA in 2000. At the CREA, I met with a friend from student days, Paul Bourguine, who had since become an expert on, and enthusiastic advocate of complex systems, for which he was setting up the first institute. We created a branch of the CREA at the Palaiseau site of the École Polytechnique.

These connections allowed me to introduce cognitive science into the curriculum of the École Polytechnique, and from 2000, to fulfil a double teaching bill in the Department of Humanities and Social Sciences, organizing on the one hand a course entitled *Introduction to Cognitive Science*, and then, with Patrick Charnay of the Biology Department at the École Normale Supérieure, a seminar called *Brain and Cognition* to which we were able to invite many specialists each year.¹² As I mentioned at the outset, it was this course that gave rise to the first edition of this book.

Three other circumstances at the beginning of the 1990s were particularly favourable to this work at the interface between differential geometry, signal processing, and the neuroscience of vision. One of these was the remarkable development of *wavelet* algorithms since the end of the 1980s, by analysts such as Yves Meyer, Stéphane Mallat, and Ingrid Daubechies (see, e.g., *A Wavelet Tour of Signal Processing* by Mallat [20]). A connection was soon made with the pioneering way in which David Marr [21] had approached retinal signal processing by ganglion cells in the retina, and it became clear that the retina, the lateral geniculate nucleus (LGN), and the primary cortical layers were carrying out a kind of successive wavelet analysis of the optical signal, which led to the primary geometrical formatting of visual input.

A second point was the link established with David Mumford's work on vision and image processing, and the interest for these subjects shown by some of my mathematician colleagues, including in particular Bernard Teissier who, back in 1991, gave me the opportunity to present the first elements of the neurogeometry of vision at the Mathematics Department of the École Normale Supérieure (ENS).

¹⁰ Today, CAMS is directed by Jean-Pierre Nadal, an expert on statistical physics, neural networks, and complex systems at the ENS. We shall say more about him later in the book.

¹¹ At that time, the CREA was located in Paris at the former site of the École Polytechnique.

¹² Apart from the colleagues already mentioned and those invited by Patrick Charnay, I would like to thank among others Yves Burnod, Jean-Pierre Changeux, Stanislas Dehaene, Gérard Dreyfus, Emmanuel Dupoux, Jacques Droulez, Étienne Koechlin, Denis Le Bihan, Pascal Mamassian, Claude Meunier, Kevin O'Reagan, Khashayar Pakdaman, and Bernard Renault.

Many cross-disciplinary collaborations were then built up, leading to the formation of a working group which organized several meetings in the wonderful environment of the Fondation des Treilles, founded by Madame Anne Gruner-Schlumberger. The first conferences in 1993–94 brought together Bernard Teissier, Jean-Michel Morel, David Mumford, Gérard Toulouse, Stéphane Mallat, Yves Frégnac, Jean Lorraine, Olivier Faugeras, and Elie Bienenstock, among others. The 1998 conference was devoted to the theme *Methodology in Cognitive Science*. There were several other meetings, including in 1995 a conference entitled *Mathematics and Brain* at the Institut Henri Poincaré in Paris. The upshot was a fruitful interaction between differential geometry, the theory of singularities, morphological image processing, signal theory, and the neurophysiology and psychophysics of vision, which led Jean-Michel Morel, David Mumford, and Bernard Teissier in 1998 to organize a special term focusing on *Mathematical Questions in Signal and Image Processing* at the Centre Émile Borel of the Institut Henri Poincaré. And then there was the important conference organized in 2001 at the Mathematisches Forschungsinstitut Oberwolfach by David Mumford, Christoph von der Malsburg, and Jean-Michel Morel under the heading *Mathematical, Computational and Biological Study of Vision*.

With Giuseppe Longo, whom I met in the 1990s, and Bernard Teissier, we also set up a seminar on *Geometry and Cognition* at the ENS. This later interacted fruitfully with the *Philosophy and Mathematics* seminar we ran together under the direction of Pierre Cartier. I should also mention the ‘Neurogeometry’ project of the *Nouvelles Interactions des Mathématiques* (ACI-NIM), run with Daniel Bennequin and Bernard Teissier (Paris VII, Chevaleret), Jacques Droulez and Chantal Milleret (LPPA), Yves Frégnac (UNIC), and several young doctoral students and researchers, a working group that proved to be particularly stimulating.

All these interactions led to a new review of the principles of neurogeometry [22] which appeared in 2003 in the *Journal of Physiology-Paris*.

The research described in the present book then benefitted greatly from collaboration with two researchers at the University of Bologna, namely Alessandro Sarti (today working at CAMS, CNRS-EHESS, in France), who specialized in models of vision and image processing, and his colleague Giovanna Citti, expert on functional analysis, PDEs, and harmonic analysis, who were then working on the first geodesic model of illusory contours, which I suggested in 1998. There began a particularly fruitful collaboration in 2001, when A. Sarti returned from Berkeley where he had been working at the Department of Mathematics from 1997 to 2000 with James Sethian on Kanizsa-style illusory contours.

We embarked upon a deeper interpretation of the geodesic model in terms of sub-Riemannian geometry. This led them to organize three conferences: the first in 2004 at the University of Bologna, entitled *Mathematical Models of Visual Perception. From Neuroscience to Phenomenology*, the second in 2006 at the Centro Ennio De Giorgi of the Scuola Normale Superiore in Pisa, entitled *School on Neuromathematics of Vision*, and the third in 2009 back at the University of Bologna again, under the title *International Conference of Sub-Riemannian Geometry and Vision*. This was once again an opportunity to discuss with eminent practitioners, including Paul Bressloff, Jack Cowan, Guy David, Alain Destexhe, Olivier Faugeras, Yves Frégnac, Walter

Gerbino, Jan Koenderink, Jean Lorenceau, Lamberto Maffei, Marc Mézard, Jean-Michel Morel, Scott Pauls, Martin Reimann, James Sethian, Wolf Singer, and Steve Zucker.

My colleague H el ene Frankowska in the CREA, expert on partial differential equations (the Hamilton–Jacobi–Bellman equation) and control theory, was working with a post doc, Piernicola Bettiol, of Andrei Agrachev at the SISSA in Trieste, and introduced me to his group (and in particular, Jean-Paul Gauthier, Ugo Boscain, Yuri Sachkov, and Igor Moiseev). I owe them much for improving my understanding of sub-Riemannian geometry. In this regard, I would also like to thank Andr e Bella iche.

The ‘vertical’ part of the \mathbb{V}_J and \mathbb{V}_S models was greatly improved in 2009 by Faugeras and Chossat [23]. They had the idea that the hypercolumns of $V1$ should encode (on the scale specified by the size of the receptive profiles), not only jets like the orientation and the curvature, but also the ‘structure tensor’ of the signal.

An active cross-disciplinary community thus came into being. I would like to mention several publications which will attest to this: first, the double special issue (no. 97, 2–3) *Neurogeometry and Visual Perception* of the *Journal of Physiology-Paris* which I organized in 2003 with J. Lorenceau, on the suggestion of Y. Fr egnac; then the *Journ ee annuelle* 2006 organized by the Soci et e Math ematique de France (SMF) on the theme of *Mathematics and Vision*, with presentations by St ephane Mallat, Jean-Michel Morel, and myself; and finally, another double special issue (no. 103, 1–2) *Neuromathematics of Vision* of the *Journal of Physiology-Paris*, published in 2009 with A. Sarti and G. Citti.

In the fall of 2014, Andrei Agrachev, Davide Barilari, Ugo Boscain, Yacine Chitour, Fr ed eric Jean, Ludovic Rifford, and Mario Sigalotti organized a special term at the Centre  mile Borel of the Institut Henri Poincar e, entitled *Geometry, Analysis and Dynamics on Sub-Riemannian Manifolds*. Many of the geometers already mentioned were present, including Montgomery, Bryant, Ambrosio, Gauthier, Pansu, and Bella iche. With J.-P. Gauthier, G. Citti, and A. Sarti, we organized a workshop attended by Jack Cowan, Yves Fr egnac, Jean Lorenceau, Pascal Chossat, Olivier Faugeras, Fred Wolf, Steve Zucker, Daniel Bennequin, Remco Duits, St ephane Mallat, Xavier Pennec, Romain Veltz, Ugo Boscain, Dmitri Alekseevsky, Peter Michor, and Giovanni Bellettini.

Regarding the psychology of perception and in particular the Gestalt theory, I had the opportunity to be involved in two conferences organized jointly in 2001–2002, the first by Arturo Carsetti at Tor Vergata University in Rome, entitled *Seeing, Thinking and Knowing* [24], and the other by Liliana Albertazzi, and Roberto Poli of the Mitteleuropa Foundation at the University of Trento, entitled *The Legacy of Kanizsa in Cognitive Science*. In 2010, A. Carsetti also organized a workshop on *Models of the Mind*, devoted largely to neurogeometry.

As regards the philosophical aspect of the present book, in which Husserl’s phenomenology of perception plays a central role, I am deeply indebted to the group *Naturalizing Phenomenology*, founded with Jean-Michel Roy, Francisco Varela, and Bernard Pachoud. I had already long since brought together the Thomian morphodynamic approach to perception and Husserlian eidetic descriptions, in fact since the 1970s. I had built up a particularly fruitful collaboration with my friends Kevin

Mulligan and Barry Smith who, for their part, had rehabilitated Husserlian phenomenology in the context of analytic philosophy and the philosophy of mind. At the time responsible for cognitive science at the Husserl Archives in the École Normale Supérieure in Paris, Jean-Michel Roy was greatly interested in these ideas and, fuelled by his enthusiasm, a seminar was set up. In 1995, Jean-Michel Roy organized an important conference at the University of Bordeaux, which served as the starting point for the reference book *Naturalizing Phenomenology* published by Stanford University Press in 1999.

At the same time, Alain Berthoz and Jean-Luc Petit, an expert on Husserl, organized a whole series of remarkable meetings to discuss the surprising convergence of phenomenology and cognitive neuroscience.

Finally, regarding the links between Husserlian phenomenology and the Kantian tradition of transcendental philosophy, I would like to thank my colleagues Michel Bitbol and Pierre Kerszberg at the CREA. We published a book with Springer called *Constituting Objectivity, Transcendental Perspectives on Modern Physics*, which described the high degree of convergence with philosophers of physics like Michael Friedman, Thomas Ryckman, and Bas van Fraassen.

The cross-disciplinary research described in this book has found many other fine opportunities for presentation in high level academic contexts. I would like to mention a few of these which have for me played a special role, so that I may thank their organizers: the International Academy of Philosophy of Science (president Evandro Agazzi); the seminars run by Jacques Bouveresse [25], Alain Berthoz [26], Alain Prochiantz and Philippe Descola at the Collège de France; the conference *Dynamic approaches to cognition* at the École Normale Supérieure in Lyon (Jean-Michel Roy); the Valparaiso Complex Systems Institute (directed by Eric Goles); the Scuola Superiore di Studi Umanistici and the Istituto Italiano di Scienze Umane (Umberto Eco and Patrizia Violi); the Center for Semiotic Research at the University of Aarhus (directed by Per Aage Brandt); the Istituto Universitario di Studi Superiori in Pavia (Amedeo Conte, Franco Rositi, and Salvatore Veca); the Scuola Internazionale Superiori Studi Avanzati (SISSA) in Trieste (Andrei Agrachev); the Laboratoire de Mathématiques, Image et Applications at the University of La Rochelle (Guy Wallet, Michel Berthier, Eric Benoît, Aziz Hamdouni, see [27]); the Philosophy Department of the University of Quebec in Montréal (Jean-Guy Meunier); the *ERMITES 2010* conference at the École de Recherche Multimodale d'Information, Techniques and Sciences (Hervé Glotin); the special term entitled *Theoretical, Mathematical and Computational Neuroscience* at the Centre International de Rencontre Mathématiques (CIRM) in Marseille-Luminy and the two conferences *Spatio-temporal Evolution Equations and Neural Fields* and *Mathematical Models of Cognitive Architectures* organized there by Olivier Faugeras, Paul Bressloff, Stephen Coombes, and Viktor Jirsa; the conference *Géométrie sans frontières* at the University of Paris VII, organized in 2012 in honour of Daniel Bennequin; and the conference *Singular Landscapes* organized in 2015 in honour of Bernard Teissier.

Finally, I would like to thank the editorial team which made possible the publication and translation of this second edition by Springer: first Catriona Byrne, then Eva Hiripi, Jan-Philip Schmidt, and Alessandro Sarti who accepted the book in the

series *Lecture Notes in Morphogenesis*, but also Olga Chiarcos and Federica Corradi Dell’Acqua, and last but not least, my excellent translator Stephen Lyle.

The final text of this English edition was for the main part written during a fellowship in 2013–2014 at the *Humanities Center* of Stanford University. Here I must thank the Reverend Marta Sutton Weeks for her euergetism and also the director Caroline Winterer and the staff at the SHC for their warm hospitality in this magnificent academic context, which thoroughly deserves the name ‘home of the human experience’.

References

1. Petitot, J.: Neurogéométrie de la vision. Modèles mathématiques et physiques des architectures fonctionnelles. Les Éditions de l’École Polytechnique, Distribution Ellipses, Paris (2008)
2. O’Keefe, J.: Immanuel Kant: Pioneer neuroscientist, Public Lecture, Royal Institution of Great Britain, London (June 2, 2014)
3. Dehaene, S., Brannon, E.M.: Space, time, and number: A Kantian research program. *Trends Cogn. Sci. Spec.* **14**(12), 517–519 (2010)
4. Thom, R.: *Stabilité structurelle et Morphogénèse*. Benjamin, New York, Ediscience, Paris (1972)
5. Thom, R.: *Modèles mathématiques de la morphogénèse*. Bourgois, Paris (1980)
6. Petitot, J.: *Eléments de théorie des singularités*. http://jean.petitot.pagesperso-orange.fr/ArticlesPDF_new/Petitot_Sing.pdf (1982)
7. Petitot, J.: *Introduction aux phénomènes critiques*. http://jean.petitot.pagesperso-orange.fr/ArticlesPDF_new/Petitot_CritPh.pdf (1982)
8. Petitot, J.: *Forme*. In: *Encyclopedia Universalis*, vol. XI, pp. 712–728. Paris (1989)
9. Turing, A.M.: The chemical basis of morphogenesis. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* **237**(641), 37–72 (1952)
10. Petitot, J.: Complexity and self-organization in Turing, The Legacy of A.M. Turing. In: Agazzi, E. (ed.), *Proceedings of the International Academy for the Philosophy of the Sciences Conference*, September, 25–28, 2012. pp. 149–182. Franco Angeli, Milano (2013)
11. Petitot, J.: Landmarks for neurogeometry. In: Citti, G., Sarti, A. (eds.) *Neuromathematics of Vision*, pp. 1–85. Springer, Berlin, New York (2014)
12. Petitot, J.: *La sémiophysique: de la physique qualitative aux sciences cognitives*. In: Porte, M. (ed.) *Passion des Formes*, à René Thom, E.N.S. Éditions Fontenay-Saint Cloud, pp. 499–545 (1994)
13. Thom, R.: Les Singularités des applications différentiables. *Annales de l’Institut Fourier* **6**, 43–87 (1956)
14. Thom, R.: Ensembles et morphismes stratifiés. *Bull. Amer. Math. Soc.* **75**(2), 240–284 (1968)
15. Petitot, J. (ed.): [LTC] *Logos et Théorie des Catastrophes*. Colloque de Cerisy en l’honneur de René Thom, Editions Patiño, Genève (1989)
16. Koenderink, J.J., Van Doorn, A.J.: Representation of local geometry in the visual system. *Biol. Cybern.* **55**, 367–375 (1987)
17. Hoffman, W.C.: The visual cortex is a contact bundle. *Appl. Math. Comput.* **32**, 137–167 (1989)
18. Koenderink, J.J., Van Doorn, A.J.: The singularities of the visual mapping. *Biol. Cybern.* **25**, 51–59 (1976)
19. Petitot, J., (with a contribution of Tondut, Y.): Vers une neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. *Mathématiques, Informatique et Sciences Humaines* **145**, 5–101 (1999)
20. Mallat, S.: *A Wavelet Tour of Signal Processing*. Academic Press, New York (1998)

21. Marr, D.: *Vision*. W.H. Freeman, San Francisco (1982)
22. Petitot, J.: The neurogeometry of pinwheels as a sub-Riemannian contact structure. In: Petitot, J., Lorenceau, J. (eds.) *Neurogeometry and Visual Perception*. *J. Physiol. Paris* **97**(2–3), 265–309 (2003)
23. Chossat, P., Faugeras, O.: Hyperbolic planforms in relation to visual edges and textures perception. [arXiv:0907.0963v3](https://arxiv.org/abs/0907.0963v3)
24. Petitot, J.: Functional architecture of the visual cortex and variational models for Kanizsa modal subjective contours. In: Carsetti, A. (ed.) *Seeing, Thinking and Knowing*, pp. 55–69. Kluwer, (2004)
25. Petitot, J.: Neurogéométrie et phénoménologie de la perception. In: Bouveresse, J., Rosat, J.-J. (eds.) *Philosophie de la Perception*, pp. 53–76. Collège de France–Odile Jacob, Paris (2003)
26. Petitot, J.: La simplicité de la notion géométrique de jet. In: Berthoz, A., Petit, J.-L. (eds.) *Simplicité-Complexité*. Collège de France. OpenEdition Books, Paris (2014)
27. Petitot, J.: Towards a nonstandard model for neurogeometry. In: Benoît, E., Furter, J.-Ph. (eds.) *Des Nombres et des Mondes*, pp. 99–118. Hermann, Paris (2013)

Chapter 2

Introduction

Without a proper shape mathematics for biology, we are in the position that physics would have been in trying to develop mechanics without Euclidean geometry.

Harry Blum

2.1 Origin of Space and Neurogeometry

2.1.1 *Geometric, Physical, and Sensorimotor Conceptions of Space*

The origin and status of spatial representations is a long-standing question that has been much discussed in the history and philosophy of science. It has been approached from several different angles up to now, including those of mathematics, physics, physiology, and psychology.

1. From the mathematical point of view, starting out with the basic reference provided by Euclidean geometry, the concept of space has been gradually generalized: non-Euclidean geometries and geometries specified by their transformation groups, Riemannian geometry, differential geometry with Cartan connections, and others.
2. From the physical point of view, starting with the idea that space and time form an a priori background structure for physical phenomena, a physical genesis of space has been gradually built up: from general relativity, based on Riemannian geometry, and non-Abelian¹ gauge theories in quantum field theory which are based

¹Named after Niels Henrik Abel.

on Cartan geometry, to the admirable synthesis of quantum physics and geometry developed by Alain Connes under the name of ‘non-commutative geometry’.

3. As regards perceptual space or ‘perceived space’, from the profound insights of physiologists like Helmholtz, geometers like Poincaré, phenomenologists like Husserl, and psychologists like the proponents of *Gestalt theory* (Stumpf, Klüver, Kanizsa, etc.), up to contemporary studies on the physiology of perception and action, we have deepened our understanding of the way the classic Euclidean space derives from our sensorimotor relationship with the environment, where solid objects play a fundamental role.

All these developments have been either directly mathematical or else based on considerable progress in mathematics. This is obvious for level (1) above, since these are simply the most impressive developments of geometry, which followed on so quickly from Gauss to Riemann and Poincaré, then Weyl and Cartan, and today Alain Connes.²

The link with mathematics is no less obvious for level (2), which concerns physical space. Here, it is worth emphasizing the way the astonishing progress in the formalisms of fundamental physics can be considered as a ‘geometrization of physics’. This is not the place to go further into this vast subject. Let us just say that the geometrization process consists in identifying more and more geometrical structures and symmetry groups of physical theories in such a way as to understand the whole complexity and diversity of observed physical phenomena in an ever more synthetic way. This is indeed the main process of mathematization since, on the one hand, it ‘reduces’ more and more physical phenomena to a priori geometrical statements, while on the other hand, it ‘unfolds’ these a priori statements in a profusion of different models, exploiting to the full the characteristic ‘generativity’ of mathematics. In the words of Jean-Marie Souriau, one of the founders of geometric quantization, who made this point so clearly [7]:

Philosophically, [geometrization] means reducing physics to geometric symmetries in order to do a priori [i.e., ‘rational’] physics.

In other words, as Souriau puts it:

There is nothing more in physical theory than symmetry groups, except the mathematical construction which allows us to show that there is nothing more.

Regarding this point, the interested reader may consult [8] and [6] and the references therein.

Regarding level (3), the question of perceptual space, the connection with fundamental mathematical structures, is less obvious, and it is precisely here that the present book aims to bring new insights. In fact, as we know, perception and motricity are tightly linked from a functional point of view, and one of the main ideas developed from Helmholtz to Poincaré was to relate the geometry of external space to our sensorimotor relationship with *solid objects* in our environment.

²Among many other works, see, for example, the classic *Space, Time, Matter* by Weyl [1], the proceedings [2], the book by Toretti [3], the studies [4] and [5] by Thomas Ryckman, or our own review of non-commutative geometry [6].

Consider for example the way Hermann von Helmholtz responded to Bernhard Riemann's famous *Habilitationsvortrag* in 1854, viz. *Über die Hypothesen welche der Geometrie zu Grunde liegen* [9], by his no less famous reply [10] *Über die Tatsachen, die der Geometrie zu Grunde liegen*.³ Helmholtz suggested reducing the problem of perceptual space to a system of axioms specifying, not Riemann's infinitesimal metric elements, but rather the *transformations* of space which are observed experimentally to be the congruences (of free motions) between rigid bodies. An elegant examination of this so-called Riemann–Helmholtz problem, concerning the origin of the geometry of external space, can be found in Joël Merker's *Le problème de Riemann–Helmholtz–Lie* [11].

There are four axioms, and the first three are rather natural:

1. The points of the space E can be represented by the values of three coordinates, in such a way that transformations correspond to (smooth) variations of these coordinates. (For Euclidean \mathbb{R}^3 , this gives the six-dimensional group of transformations comprising the three translational degrees of freedom and the three rotational degrees of freedom.)
2. There exists a function $f(a, b)$ defined on $E \times E$ which is invariant under all transformations.
3. Any point in E can be carried to any other point of E by a transformation (transitivity).

The fourth axiom, called the *monodromy* axiom, is much less obvious:

4. If we choose two points a and b in a rigid body, then there is one remaining degree of freedom (rotations with axis ab in the case of Euclidean \mathbb{R}^3), and such a 'rotation' must move all the points and bring the body point by point back onto itself after one complete turn.

In volume III of their famous treatise *Theorie der Transformationsgruppen* [12], Sophus Lie and his disciple Friedrich Engel spelt out the above axioms, rectifying certain errors made by Helmholtz and classifying all the solutions, noting that Euclidean geometry was only one solution among others (see [13]). To do this, they used the theory of the groups and algebras now known by the name of Lie groups and Lie algebras, something we shall make constant use of throughout this book.

Regarding Henri Poincaré, some of his basic ideas about physical space and perceived space are discussed in Chaps. 4 and 5 of *Science and Hypothesis* [14], entitled *Space and Geometry* and *Experience and Geometry*, respectively. The principle of geometric conventionalism asserts that the geometry we apply in physics is conventional, i.e. neither true nor false, that its axioms are neither experimental (criticism of empiricism), nor synthetic a priori (criticism of the narrow idealist interpretation of Kantian apriorism), and that the same factual physical contents can be described within alternative geometrical frameworks. As a convention, a geometry provides a language for description and does not possess any experimental or empirical truth in itself. By introducing the thesis that the group concept is an a priori feature of our

³'Tatsachen' or 'facts' are contrasted with 'Hypothesen'.

understanding that thus ‘pre-exists in our minds, at least potentially’, Poincaré was putting forward a version of the transcendental ideality of space which is compatible with the existence of several different geometries and hence with the progress made in theoretical physics. Let us recall here the conclusion of *Space and Geometry*:

The object of geometry is the study of a particular ‘group’; but the general concept of group pre-exists in our minds, at least potentially. It is imposed on us not as a form of our sensibility, but as a form of our understanding; only, from among all possible groups, we must choose one that will be the standard, so to speak, to which we shall refer natural phenomena. Experience guides us in this choice, which it does not impose on us. It tells us not what is the truest, but what is the most convenient geometry.⁴

Poincaré expands on this idea in *Experience and Geometry*, where he explains that the principles of geometry are not experimental facts. A given physical fact can always be expressed by changing the convention represented by the geometrical framework and changing the laws of physics; e.g., one can keep Euclidean geometry but reject the principle that light rays follow geodesics. This point of view was already anticipated by Clifford: there is an equivalence between (1) physical causes of changes in a space thought of a priori as flat and (2) a non-trivial (curved) space geometry. Physical experiments are always carried out on bodies, never on space. Therefore, they cannot help us to decide upon the geometry.

Regarding perceived space, Poincaré considered that its geometry must come essentially from our fundamental sensorimotor experience of the motions of solid bodies (see, in particular, *Science and Method* [15]). This constitutes our notion of space and, by distinguishing between proprioceptive internal changes and external changes that may balance them, leads to the aprioricity of the group concept and to the idea that geometry is conventional.⁵

2.1.2 The Neurogeometric Approach

With this in mind, it should be said that there is (at least) a *fourth* way to inquire as to the origin of spatial representations. Until recently, it had only been the subject of a few bold and generally incorrect speculations, due to the lack of experimental evidence. This fourth way concerns the highly complex *neurophysiological processes* through which the geometrical structures of the external space are *constituted* as a result of the internal activities of our brains.

This can be tackled by considering at least two main lines of approach:

- Sensorimotor and locomotor positioning and navigation of an organism moving through space. For example, the now classic book *The Hippocampus as a Cognitive Map* (1978) by John O’Keefe⁶ and Lynn Nadel [17] has a long first chapter with a

⁴This is a clear reference to Kant’s opposition between sensibility and understanding.

⁵For the relations between Mathematics and Physics in Poincaré, see *La Valeur de la Science* [16].

⁶See Sect. 1.1 of the *Introduction*.

historical slant relating this new work on navigation to the philosophies of space propounded by philosophers and mathematicians such as Leibniz, Euler, Kant, Helmholtz, and Poincaré. Similar discussions can be found in the many works of neurophysiologists of perception and action like Alain Berthoz.

- The geometrical structuring of visual images. This will be our main subject in the present book.

We coined the term ‘neurogeometry’ to refer to this *neural* origin of perceived space. The aim in the present work is to take a first step in this direction, something made possible by the huge amount of new and fascinating experimental results now available thanks to new imaging techniques. As long as the brain remained, from an experimental point of view, a ‘black box’, there was no way of developing such an approach. What made this possible was thus that the brain became, at least to some extent, a ‘transparent box’.

Brain imaging techniques are here the equivalent of the new observational methods that are always found to underlie any scientific revolution. We shall show that their results can be modelled using sophisticated mathematics that corresponds in the deepest possible ways to mathematics already invented by certain outstanding geometers like those already mentioned, and in particular Lie and Cartan, when they set out to understand mathematically how the geometry of the external world (Euclidean or otherwise) could come about. We would thus like to insert a new page in the age-old story of the *foundations of geometry*. There will be two main aims:

- To provide models for a whole new set of neurophysiological data.
- To fit these models into modern developments in the foundations of geometry.

The analogy with the history of the theories of physics could be illuminating here. Just as the modern theories of fundamental physics (general relativity, gauge theories, Higgs field, etc.) have led to ever further geometrization of empirical physical phenomena which, in its turn, provides a better understanding of the physical genesis of space, so neurogeometry consists in a geometrization of empirical neural phenomena which, in its turn, provides a better understanding of the neural genesis of space. Our whole investigation will be based on this ‘*dialectic*’ *between the geometrization of internal neural dynamics and the neural foundations of external geometry*.

2.2 Perceptual Geometry, Neurogeometry, and Gestalt Geometry

Let us begin by giving a few points of reference and some clarifications:

1. Following on from the great geometers, phenomenologists, and psychologists who have turned their attention to our perception of forms, as discussed above, a certain number of eminent scholars have recently made considerable contributions to the geometry of visual perception. We may mention René Thom, who developed the first general dynamical theory of shapes, Jan Koenderink, who

applied Thom's theory of singularities to visual neurophysiology, the heirs to the Gestalt psychologists, and in particular, Gaetano Kanizsa, to whom we shall return at length, David Marr, who, at the end of the 1970s, brought a host of new insights into the problem of vision, and David Mumford (Fields medallist like René Thom), who completely revolutionized the area. When we talk about neurogeometry here, what we shall aim for is the *neural implementation* of the algorithms of this geometry, the problem being to understand how perceptual 'macrostructures' and their morphodynamics can *emerge* from the underlying neural 'microlevel'.

2. The aspect of neurophysiology that is relevant in this research is *functional neuroanatomy*. It is not concerned with the biochemical details of the individual neurons (ion channels, membrane potentials, etc.), but treats them rather as functional units, e.g. threshold automata in neural network models, connecting to form neuroanatomically specifiable populations. We shall say a few words about the 'micro' cellular level relevant to molecular biology, but most of what follows will concern a 'meso' functional level.
3. One characteristic of perception is that perceptual 'phenomenal consciousness' results from integration, in the neurophysiological sense, of the partial processing carried out by a great many different brain modules connected together in an extremely complicated way with a high level of feedback. Processing is highly modular (whence the very specific nature of pathologies), but consciousness is highly integrated. This means that models for specific areas are necessarily incomplete. Here, we shall be dealing mainly with the first area, known as V1 (or area 17 in cats), of the primary visual cortex. This does of course limit the discussion, but we shall see that much can already be said and that this provides a good example of what is meant by neurogeometry. Furthermore, despite being so restrictive, this case can also be considered as fundamental if we adopt David Mumford and Tai Sing Lee's 'high-resolution buffer hypothesis'. According to this, V1 takes part in any higher level processing which requires high resolutions (see Mumford [18] and Lee et al. [19]).
4. We stress that neurogeometry is about the *internal* geometry (already referred to here as 'immanent') of low-level vision, and not therefore the conventional 'transcendent' geometry of the perceived external 3D Euclidean space. It concerns a much more fundamental level, and to use the nice expression adopted by Misha Gromov to speak about sub-Riemannian geometry, it tries to understand perceived space *from within*.
5. In neurogeometry, anything that is not implemented neurally does not exist. This means that all the mathematical concepts used operationally in the models must have some material counterpart. There is a similar situation in computer science, where the software only works if it is compiled and realized materially in the physics of the hardware. It is not easy to implement this equivalence between *geometric idealities* and *neural materialism*. Indeed, on the one hand, trivial mathematical structures such as alignments, gluing of local charts, or direct products are implemented neurally in a very subtle way that is hard to study experimentally, and on the other hand, certain properties of the modelling structures will not

be implemented and so will have no empirical meaning. The reader should bear this crucial point in mind: when a set of empirical phenomena is modelled by mathematical structures of a certain kind, *only certain aspects* of these structures will be open to empirical interpretation.⁷

6. Furthermore, implementations can differ significantly depending on the species, and the same abstract functional structure can be achieved materially in different ways in the various layers of V1 (see Sect. 4.9.4 in Chap. 4.). We need therefore to carry out very careful interspecific comparative studies on rats, ferrets, tree shrews (tupaia), cats, macaques, humans, etc.
7. The neurons in V1 have small receptive fields and thus process information from the photoreceptors in a very *local* manner, i.e. localized in the visual field. The main problem is to know how these local data are organized into *global* structures such as lines, edges, surfaces, and shapes. This is a problem of ‘integration’ in the mathematical sense, and here, the concept of *functional architecture*—referring to the design of the connectivity of neurons within an area—proves to be crucial. The enigmatic phenomena studied by *Gestalt theory* relate to the fact that perception ‘integrates’ local data and ‘fills in the gaps’, if there are any. In this sense, neurogeometry could be qualified as *Gestalt geometry*.

2.3 Geometry’s ‘Twofold Way’

Let us stress once more that, in neurogeometry, there is a *twofold* relationship between the geometry and neurophysiology of vision. As we shall explain in detail, it is the functional architecture of the visual areas, the precise organization of their neural connections, which generates the geometric properties of perceptual space, i.e. the perceived 3D space in which the objects of the external world are situated. We may thus envisage a ‘neural → spatial genesis’ of the kind ‘functional architecture → geometric properties of external space’. But as we shall see later, there exist geometric models of the functional architectures themselves; that is, the latter implement well-defined *sui generis* geometrical structures. It is important to distinguish carefully between the two levels at which geometry enters the discussion. The whole purpose of this book would become incomprehensible if they were confused. As we have seen, to formulate the distinction, we may return to the classical philosophical opposition between *immanence* and *transcendence*. The geometry of functional architectures is immanent in perception, internal and local, and its global structure is obtained by integration and coherent association of local data. In contrast, the geometry of perceived space is transcendent in the sense that it concerns the outside world and is given to us immediately as global.

⁷This problem already arose with the invention of rational mechanics. In order to mathematize the physical motions of material points, one must choose a frame of reference. But neither absolute positions, nor absolute directions, nor absolute velocities have any physical meaning, hence the advent of Galilean relativity.

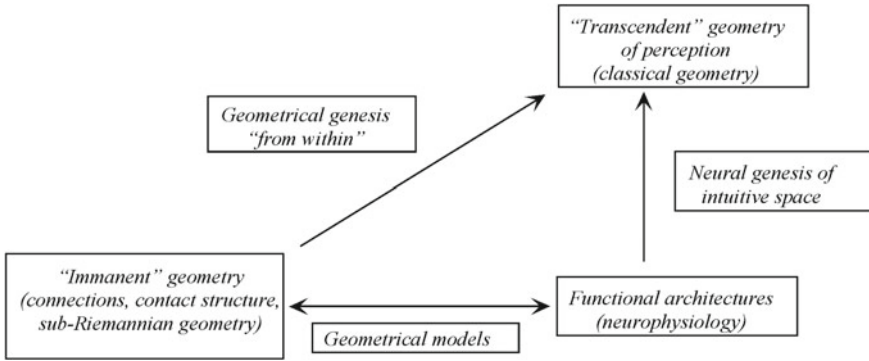


Fig. 2.1 Connections between ‘immanent’ geometry and ‘transcendent’ geometry

But it turns out that neurally implemented immanent geometry can itself be modelled using deep geometric structures already introduced by the geometers mentioned earlier, such as Elie Cartan, Hermann Weyl, René Thom, Alain Connes, and Misha Gromov, to understand the genesis of transcendent geometry. This implies that, once modelled in this way, the neural genesis of space can be *internalized* in the mathematics and thereby identified with a *mathematical* genesis of a macro and global geometry from a micro and local one, globalized by integration and coherent matching. This should come as no surprise, because the genesis of physical space occurs in exactly the same way: once physics has been mathematized, it is identified with the genesis of classical geometry from Riemannian geometry (in general relativity) or from the non-commutative geometry called ‘quantum’ or ‘spectral’ geometry in quantum field theory. The diagram in Fig. 2.1 explains this interaction between the different philosophical levels of understanding geometry.

2.4 Idealities and Material Processes

To clarify this key point, let us make an analogy. Although it differs with regard to content, the new direction provided by neurogeometry is methodologically speaking of the same kind as the one taken during the last century with the advent of the Turing machine, λ -calculus, and computers. This computational revolution took the symbols that underlie logical *idealities* and turned them into *material* operations. It explained how the dominant logical idealism and analytic apriorism expounded from Bolzano to Frege could be naturalized and even physicalized. In other words, it explained how logical ‘software’ could be implemented in physical ‘hardware’.

We are doing just the same here. The aim of the ‘neurogeometric’ approach is to obtain an explicit understanding of the material operations that underlie the geometric idealities of the synthetic a priori and to explain how some kind of geometric ‘software’ could be implemented in our neural ‘hardware’, hence the following analogy:

	Idealities	Type of a priori	Implementation
Logic	Logical idealities	Analytic	λ -calculus
Geometry	Spatial idealities	Synthetic	Neurogeometry

Let us say a little more about this analogy. In logic, we have what is known as the Curry–Howard correspondence which relates low-level machine language with the high-level language of logic. Low-level calculations are described, for example, by the λ -terms of a λ -calculus which describes the programs. In the simplest λ -calculus, the λ -terms (the programs) are constructed inductively by iterating two basic operations:

- the application MN of one λ -term M to another λ -term N ,
- the abstraction operation $\lambda x.M$ transforming the free occurrences of the variable x in M into places for other λ -terms.

The basic rule of λ -calculus (which corresponds to executing the programme described by the λ -term) is known as β -reduction. It consists in applying a λ -term $\lambda x.M$ to another λ -term N by substituting N in all the free occurrences of x in M , which can be written $(\lambda x.M)N \rightarrow_{\beta} M[x := N]$. The normalization of a λ -term is a sequence of β -reductions which stops at a β -irreducible λ -term. The normalizable λ -terms thus describe effective computations which stop and deliver a result. The fundamental link with logic comes from the *typing* of the λ -terms M into types μ (notation $M : \mu$). Intuitively, if $M : \mu$ is a λ -term of type μ , and if $x : \sigma$ is a variable of type σ , then the abstraction $\lambda x.M$ has the type $\sigma \rightarrow \mu$ of functions of source σ and target μ . Likewise, if $M : \sigma \rightarrow \tau$ is a λ -term with functional type $\sigma \rightarrow \tau$ and if $N : \sigma$ is of type σ , then MN is of type τ . In fact, these are the types which correspond to the formulas of a logic system: intuitionistic propositional logic. The Curry–Howard correspondence between programs and proofs is summarized here:

λ-calculus, programs	Logic, proofs
Low level	High level
Code	Expression
Compilation	Decompilation
Execution of the program	Theorem
Encoding	Typing
Instruction	Logic rule

It is this kind of correspondence that we shall describe in this book but yet with three fundamental differences:

1. The low-level calculations will be neural calculations and not programs written in a machine language.
2. The high-level structures will not be expressions of a logic system, but geometrical structures.

3. In contrast to computers (universal Turing machines), the neural hardware is *dedicated* to certain tasks, and its concrete physical activity is thus equivalent to the abstract ideal ‘calculation’ it carries out.

These points can be displayed as follows:

Neural ‘calculation’	Geometry
Low level	High level
Neural code	Geometric structures
Compilation	Decompilation
Neural activity	Geometric construction
Encoding	Typing
Instruction	Construction rule

2.5 Mathematical Prerequisites and the Nature of Models

By its very nature, the following will raise certain issues relating to didactic presentation, issues that might prove off-putting to some readers. Indeed, we shall use many mathematical concepts generally considered to be rather ‘advanced’: differential forms, connections, Lie groups, contact structures and symplectic structures, sub-Riemannian geometry, variational models, non-commutative harmonic analysis, and so on. We shall define these as we go along, assuming a basic understanding of differential and integral calculus, linear algebra, and elementary group theory. These are basic concepts that will be familiar to any science student and which are in any case easy to find in a good encyclopaedia.

Having said that, the reader may wonder quite rightly why such mathematics is relevant here. Our long experience as teacher and researcher in cognitive science has shown us that biologists and psychologists are often intrigued, even shocked, by the idea that non-trivial mathematical models (going beyond simple methods of data analysis) should be needed in their field of study.

A first source of suspicion comes from the idea that mathematics should only be applicable to intrinsically rational phenomena and that, insofar as evolution results from a ‘tinkering’ process, biological structures could not be intrinsically rational and so could not as a matter of principle be expressible in terms of mathematics. There are several possible answers to this. To begin with, there is no metaphysical reason why physical phenomena themselves should be intrinsically rational. It is rather because our efforts to express them mathematically have been so successful that they now appear a posteriori to be so rational. Secondly, what characterizes physical rationality expressed in this way is the existence of simple laws, from Kepler and Newton to superstring Lagrangians.⁸ But modelling goes well beyond what is governed by laws.

⁸Named after Joseph-Louis Lagrange.

For example, many differential equations can be applied to a whole range of different fields: Turing-type reaction–diffusion equations for morphogenetic processes, the Hodgkin–Huxley equation [20] for the propagation of action potentials, the spin glass equations of statistical physics for neural networks, the Lotka–Volterra equations for ecology, and so on. There is thus no deep reason why there should be any natural limit to the use of mathematical models.

Another argument often put forward is that if we make the hypothesis that algorithms are implemented neurally, this would mean that neurons ‘calculate’, which is impossible. But this argument is also mistaken. In Mechanics the planets do not ‘calculate’ their trajectories. The only thing we can say is that theories based on laws involving *global* interactions (as is the case with Newton’s universal law of gravitation) are problematic and that the interactions must be *localized* (something achieved by general relativity). However, in neuroscience, we can be sure of the locality of the interactions, because these interactions occur through *material connections* between neurons. What passes for a neural ‘calculation’ is essentially the propagation of activity along connections, and this is a ‘calculation’ because the connections are organized into highly specific *functional architectures*. In other words, it is the structure of the functional architectures—in a sense, the ‘design’ of the neural ‘hardware’—which amounts to a calculation.

A third argument is that even if we are convinced of the relevance of mathematical models in neurophysiology, we should at least seek out the simplest possible models and that we should in principle be suspicious of any complexity in this context. Once again, this is simply a prejudice and indeed constitutes another fallacy. To see this, we only need to return to the beginnings of differential and integral calculus and mathematical physics. To solve what seemed to be very simple problems, such as calculating the length of the arc of an ellipse, new functions had to be invented, viz. the elliptical functions, much more complicated than the trigonometric functions. Likewise in mechanics, to solve apparently very simple problems, such as the problem of a hanging chain, i.e. the shape of the curve adopted by a chain of uniform linear density when suspended by its two ends and subject to the force of gravity alone, the pendulum, or the shape of a uniform elastic rod when curved (elastica problem), mathematicians had to solve specific differential equations or variational problems, which were what they were and which turned out to involve astonishing internal complexity. Newton’s law of gravitation is expressed by an extremely simple second-order differential equation, but in most cases, when the relevant forces are fed in, it becomes a specific differential equation whose solutions have nothing simple about them at all. The complexity of the solutions often makes them quite inaccessible, as illustrated by the n -body problem.

The emergence of complexity is in fact perfectly commonplace, and we shall return to this in the second volume. It is often due to the fact that the integration of a differential equation involves *iterating* the infinitesimal generator of the equation. But the iteration of operations generally leads to a great deal of complexity, even if these operations are very simple. Fractals provide us with many examples.

2.6 Mathematical Structures and Biophysical Data

For our investigation of neurophysiology, we should like to return to the spirit of the pioneers of the seventeenth and eighteenth centuries, such as Euler and Lagrange, in their investigation of mechanics. Indeed, there is really no reason why ‘calculation’ of perceptual geometry by the visual cortex should be simpler and less subtle mathematically than the calculation of the arc of an ellipse, the hanging chain, the pendulum, or the elastic rod. Empirical phenomena have to be taken as they are. The important thing is to model them correctly, and it is perfectly understandable that, in order to do that, we must appeal to somewhat elaborate mathematics.

However, we are fully aware that we may convince *neither* the neurocognitivists *nor* the mathematicians, because we know from experience that the transition from ‘neither, nor’ to ‘both, and’ can be a difficult one. As soon as we leave the field of physics, whose practitioners have been making mathematical models of empirical reality for centuries, we find a ‘gap’, often even a ‘gulf’, and not only theoretical, but institutionalized, between mathematical structures and empirical observations (here neurophysiological). Experimenters tend to want to preserve the full complexity of the data they have acquired using highly sophisticated equipment and thus tend to prefer computer simulations rather than formal models which always simplify the data in order to extract structural properties. The computational programs ‘Blue Brain’ and ‘Human Brain’, to be discussed in Sect. 4.3.1 of Chap. 4, are good examples. And this mistrust on the part of experimental neuroscience will find little to counterbalance it from the mathematicians because, as one might imagine, many of these will only see in these models elementary special cases of structures they have long been perfectly familiar with, even though they may be considered insurmountably difficult to grasp by their neurocognitivist colleagues.

But we shall nevertheless take this risk, making the optimistic hypothesis that some readers will feel that, as far as the neuroscience of vision and neural genesis of perceived space are concerned, the gap between mathematics and experimentation is actually less difficult to negotiate than one might think.

In fact, we consider neurogeometry to be intrinsically cross-disciplinary, that is intrinsically involving many different disciplines, something forced upon us by the very nature of the phenomena it seeks to theorize, but with the long-term aim of becoming a discipline in its own right. Until now, the basis of neuromathematical projects has consisted above all of (ordinary or partial) differential equations for neural activity. Our purpose will be to introduce more abstract methods of differential geometry.

So let us stress that we shall therefore concentrate on *geometric* models. On the other hand, this will not prevent us from giving a glimpse of other methods when the opportunity arises. In this way, the reader will get a better idea of the wealth and diversity of neuroscience models.

2.7 Levels of Investigation: Micro, Meso, and Macro

Another potential problem here is the sheer breadth of the topics treated here. Of course, we shall focus on modelling the functional architecture of the primary visual areas and in particular $V1$. But despite the apparently rather limited nature of the subject, we shall nevertheless only discuss a very small part of it. It is easy to understand why. To begin with, we shall only be dealing with the so-called functional, integrative, and computational neurosciences, and apart from the discussion in Sect. 5.12 of Chap. 5 which we shall explain when the time comes, we shall not be concerned with any aspect of molecular biology or genetics. This said there are still three levels of investigation for the purpose at hand: those of *microneurophysiology*, *mesogeometry*, and *macrodynamics*. These will receive differing amounts of attention.

For instance, one of our basic experimental inputs (see Sect. 4.3 of Chap. 4) will be the fact that the single neurons in $V1$ detect a retinal position $a = (x, y)$ and a preferred orientation p at a , although naturally at a certain scale. The data (a, p) is called a *contact element* in differential geometry, and we shall thus consider the single neurons of $V1$ as filters extracting contact elements from the optical signal. But just this simple claim is the subject of a huge experimental effort. For example, one needs to compare the situations for different species and take into account the fact that, in these results, neurons are treated as linear filters acting on stimuli reduced to single bars (simulating the edge of an object) or systems of parallel bars in motion (drifting gratings), while it is clear that there are significant nonlinearities and also that natural stimuli may have very different structures.⁹ One must also take into account the fact that the imaging techniques used do not have sufficient spatial resolution to distinguish individual neurons,¹⁰ whence one is in fact dealing with *local averages* over small groups of neurons, and a piece of geometric data like a contact element (a, p) reflects an average of the underlying activity. The geometric quantity we refer to as a ‘contact element’ thus represents a *mesoscopic* entity when compared with the microscopic level of individual neurons.

One consequence of this choice of a mesoscopic level for neurogeometry is that what we shall call a ‘neuron’ will actually be a small patch of neurons, and we shall thus say little about true *elementary neural circuits*. There is an extensive literature on this subject and some sophisticated engineering, but we shall only refer to it from time to time.

It should also be noted that even a very high resolution would not remove the problem of levels. Indeed, the neural code is a *population coding*, where each elementary operation activates a large number of neurons. A ‘high-resolution’ neurogeometry that was truly microscopic would therefore have to be based on the tools of stochastic differential geometry, something pointed out by specialists such as David Mumford, Jack Cowan, and Daniel Bennequin. So let us stress once again that the neurogeometry developed here will idealize things by sticking to a mesoscopic level. The global

⁹See, for example, Marre [21].

¹⁰However, in Sect. 4.7.3 of Chap. 4, we shall discuss the latest methods of two-photon confocal microscopy, which can in fact distinguish individual neurons.

structures, processes, and dynamics that we shall study will thus be based on gluing together mesoscopic geometric elements.

All the various aspects of the microlevel are currently the subject of ever more highly specialized studies. What this means is that, while our neurogeometric meso-models are mathematically rather sophisticated, they concern only a very limited part of what contemporary neuroscience can teach us, and in a highly simplified way, so they are only a first step into this new field. What we would like to advocate in neuroscience is mainly the geometric framework, which seems relevant and natural for the mathematical modelling of functional architectures.

2.8 The Context of Cognitive Science

As we have rather briefly specified above, this book is about the problem of modelling in cognitive science, that is in the natural science of cognitive faculties and mental activities. Let us therefore say a few words about this context.¹¹ The cognitive sciences bring together all the various disciplines that tackle the question of human, animal, and artificial intelligence, starting with the underlying neurobiological substrate, its embodiment, that is the relationship between mental activity and the body apart from just the neural aspects, and its relationship also with the emotions, but going on to include its formal and mathematical structure (there are many different types of model in cognitive science), its computer simulations, and its linguistic, psychological, and social realizations.

The different areas of research in the cognitive sciences, specifically perception, action, reasoning, and language, are carried out with an endogenous, intrinsic, and unified ‘polyscientific’ approach, whose cross-disciplinary nature is imposed by the subjects of study and combines statistical physics, differential geometry, cognitive, computational, and integrative neuroscience, cognitive psychology, artificial intelligence, logic, linguistics, philosophy, and the social sciences. Biological evolution has produced an amazing biochemical machine, the brain, with intellectual, mental, and symbolic capacities. In a few tenths, or even hundredths of a second, this machine can recognize a complex visual shape, calculate the sequence of instructions required by the muscles to catch a ball in flight, or decode an acoustic message by identifying the words and their meaning. It includes a whole range of processing levels, from low-level peripheral sensory processing, such as retinal processing of an optical signal, to high-level central abstract symbolic processing, such as judgement and inference, or aesthetic assessment.

The aim of cognitive science is thus to explain mental phenomena—be they states, entities, structures, events, or processes—in a strictly naturalistic and causal way. These are the problems which, by definition, have long been studied by physiology and psychology. They have also been the subject of extensive and rigorous conceptual

¹¹There are many excellent introductory Websites to find out about cognitive science, such as *The MIT Encyclopedia of the Cognitive Sciences* [22].

analysis by philosophers from Aristotle to Descartes, Hume, Locke, Leibniz, Kant, and many others, who have reflected upon the nature of ‘ideas’, ‘human understanding’, and ‘mental faculties’. As a science of the ‘mind’, the cognitive sciences are thus by definition natural sciences, bringing with them a vast philosophical legacy. The novel aspect of the current scientific situation is, on the one hand, the remarkable harvest of results obtained over the past few decades in neuroscience and, in particular in brain imaging, and, on the other hand, the integration of theoretical work on cognition, not just in the natural sciences, statistical physics, and biochemistry, but also in the formal sciences of geometry, logic, and theoretical computing. What is more, insofar as the cognitive sciences also concern artificial cognition, they are now inseparable from information processing systems and methods for analyzing and synthesizing image and sound, not to mention artificial intelligence (AI) and robotics.

We therefore stress once more that the cross-disciplinary nature of cognitive science is *intrinsic* and *endogenous*: it is imposed by the very nature of the entities, structures, and mental processes it investigates. An ability such as the perception of objects in three-dimensional space on the basis of ‘pixellated’ two-dimensional retinal data can be studied on a formal level (to identify the mathematical and formal features of the problem of constituting objects bounded by edges and filled with perceived qualities), on a behavioural level (studying the computational procedures, i.e. processes of integration, recognition, inference, and interpretation), and on the level of the biological substrate (investigation of neurophysiological mechanisms). This ability thus involves several levels of integration in both space and time.

The cognitive sciences treat all these mental phenomena a priori as a broad class of *natural* phenomena. They do for the mental what biology has been doing for the living since the nineteenth century. Consequently, their status depends on the way we extend the concept of ‘nature’. If we understand ‘nature’ in the narrow (strictly physicalist) sense, this leads to a reductionist or ‘eliminativist’ understanding of the mental. But if ‘nature’ is taken in a broader sense, we arrive at an ‘emergentist’ understanding of the mental, e.g. emergence of macrostructures from microinteractions in complex systems, as in thermodynamics and sociology. But whichever option is chosen, the approach will be naturalistic and *monistic*, rejecting any Cartesian form of dualism between mind and body (two substances).

The term ‘natural sciences’ also includes mathematical modelling, computer simulation, and an experimental approach. Cognitive science has become a new frontier in the contemporary hard technosciences, with considerable technological spin-offs (neural networks, robotics, hybrid natural–artificial systems, and so on). The effect has been to completely break down the conventional boundaries between the physical and mathematical sciences, the biological sciences, and the social sciences. Thanks to what are now called the *convergent technologies*, the physical, the biological, and the mental come together into a unified understanding of *complexity* in nature.

This naturalization of all that is mental—and at the end of the day, that means also consciousness, intentionality, and meaning—brings with it formidable epistemological challenges, and it will thus be impossible to develop the cognitive sciences without facing up to a whole set of problems relating to the theory of knowledge.

2.9 Complex Systems and the Physics of the Mental

As ‘hard’ technosciences, the cognitive sciences are inextricably related to the study of complexity and derive from the intellectual environment that came into being in the 1940 to 50s, so admirably exemplified by exceptional scholars such as John von Neumann, Norbert Wiener, Warren McCulloch, and Walter Pitts. They belong to the movement that saw the joint emergence of the theories, techniques, and methods of computers, neural networks, cellular automata, information processing, and self-organizing, self-regulating complex systems.¹² After several decades of progress in constant interaction with neuroscience, cognitive psychology, linguistics, and certain approaches to economics, these activities are now mature enough to justify referring to them as a ‘science’.

This is part of a deep trend. There has been a gradual development of mathematical physics to treat the organizational complexity of material systems and the emergence of patterns and shapes, but also cognitive activities as ‘unphysical’ as conceptual categorization and learning. We began by understanding how shapes could ‘emerge’ and ‘self-organize’ in a stable manner on the macroscopic scale as causal consequences of complex interactions on the microscopic scale. Collective micro-physical phenomena, both cooperative and competitive, provide the causal origin of joint behaviour on a macroscopic level which can break the homogeneity of a substrate. The classic physical example is provided by *critical phenomena* like phase transitions. It was then realized that neural networks are the same kind of system, but in which emergent shapes and structures can be interpreted as cognitive processes.

If rather similar models crop up in rather disparate fields of empirical investigation, this is because complex systems possess certain relatively universal properties.¹³ By definition, these are large systems of interacting elementary units with emergent global macroscopic properties arising from cooperative or competitive collective interactions between these units. These systems contrast with classical deterministic mechanical systems in the following ways:

- They are singular and individuated, largely contingent, not concretely deterministic, even when they are ideally so: they are sensitive to tiny variations in their control parameters, a sensitivity that can induce divergence effects.
- They are historical products, resulting from processes of evolution and adaptation.

¹²For an introduction to this scientific revolution, the reader is referred to the reflections of Dupuy [23].

¹³This is actually a common theme throughout the history of science. From Galileo and Newton, we learnt that types of motion as apparently disparate as ‘sublunar’ ballistic motions and ‘superlunar’ celestial motions could be understood using a single mathematical theory, the universal theory of gravitation. This formal similarity between empirical areas hitherto considered to be ontologically incommensurable was once culturally traumatic, but it eventually became commonplace in science. The same can be said here. For example, the fact that neural networks carrying out cognitive operations of categorization are formally analogous to spin glasses may look quite bewildering, given the gulf that separates these two ontologically incommensurable areas. But it is already a scientific commonplace for the young generation.

- They are out of equilibrium and have an internal regulation that keeps them within their range of viability.

They have little to do with classical mechanistic determinism. They are analyzed using new physical and mathematical theories and a computational approach making heavy use of computer simulation. The role of nonlinear dynamical systems (attractors, structural stability properties, and bifurcations), chaos theory, fractals, statistical physics (renormalization group), self-organized criticality, algorithmic complexity, genetic algorithms, and cellular automata has become key to understanding their statistical and computational properties. In short, through the engineering of self-organized, non-hierarchical, distributed, and acentered artificial systems, we are beginning to be able to model and simulate reasonably well biological systems (immunological systems, neural networks, evolutionary processes), ecological systems, cognitive systems, social systems, and economic systems.

2.10 The Philosophical Problem of Cognitive Science

Cognitive science can be approached in a purely operational and instrumental way, but its development nevertheless raises many issues on the philosophical level because, as we have just seen, it questions the traditional dividing line between the science of nature and the science of mind. To be more specific here, let us return for a moment to certain epistemological basics.

In the *formalization* of the so-called exact sciences, there is a lot more than, on the one hand, the processing of empirical data using universally applicable methods such as statistics, factor analysis, principal component analysis, data mining and, on the other hand, the axiomatization of theoretical concepts. These two types of formalization also exist in the social sciences and involve general methods that are independent of the source of the data and the kinds of things to which they are applied.

But in the physical sciences, there is also modelling in a stronger sense which is of a quite different kind. For this modelling in the strong sense, methods are *specific* to the theoretical conceptualization of a particular kind of object and can be used to *reconstruct* the phenomena in some real field from its constitutive theoretical concepts. Mathematical physics is able to reconstruct the whole diversity of physical phenomena from its theoretical concepts. This completely changes the status and function of concepts. We no longer subsume empirical diversity by abstraction under the unity of theoretical categories and concepts. Rather, concepts are transformed into algorithms for reconstructing the diversity of phenomena. Put another way, conceptual analysis is converted into a *computational synthesis*.

At the present time, the ideal of a computational synthesis of phenomena has only really been achieved in physics, which is restricted to a very narrow and highly constrained region of empirical reality. Huge regions of phenomena have been left outside the reconstruction zone, even though a fair number of these regions have been studied in detail by many empirical and descriptive disciplines. Here, we may cite:

- The whole macroscopic organizational and morphological complexity of material systems.
- All cognitive operations, including categorization, inference, induction, learning.
- The whole semiotic and linguistic dimension of meaning.
- And in fact anything having to do with *phenomenality* itself as a process of phenomenalization of an underlying physical objectivity.

In other words, it is only by restricting phenomenal reality to its most elementary form (essentially, the trajectories of material bodies, fluids, particles, and fields) that we have been able to carry through the programme of reconstruction and computational synthesis. For the other classes of phenomena, this project has long come up against unsurmountable epistemological obstacles.

At this point, it was taken as self-evident that there was an unavoidable scission between phenomenology (being as it appears to us in the perceived world and the cognitive faculties that process it) and physics (the objective being of the material world). However, we may say that it is not so much self-evident as a straightforward prejudice. In any case, this disjunction transformed the perceived world into a world of subjective-relative appearances—mental projections—with no objective content and belonging to psychology. Beyond psychology, the most that could be attributed to these appearances in the way of objectivity was a logical form of objectivity to be found in the theories of meaning and mental contents, from Bolzano and Frege, Husserl and Russell, to contemporary analytical philosophy.

We may say that the current work aims to go beyond this scission by developing a *mathematical neurophysics of the phenomenology of the perceived world and common sense*. The neurogeometry of vision presented here will be one aspect of this.

2.11 Some Examples


To end this introduction, let us mention some of the most striking examples of perfectly intuitive but theoretically problematic perceptual features that we shall attempt to understand.

2.11.1 *The Gestalt Concept of Good Continuation*

Figure 2.2 shows small aligned segments against a background of random distractors.¹⁴ The alignment seems to jump out at us, and indeed, this is typical of what is known as a ‘pop out’ phenomenon. It results from ‘binding’ and integration of local information into a global structure. Psychophysical experiments have shown that it

¹⁴These are Gabor patches, not geometric segments, as we shall explain.

Fig. 2.2 Example of ‘good continuation’. From Hess et al. [24]




is indeed the global alignment that causes this effect. But what is the meaning of a global alignment on the neural level? For each of us conscious sentient beings, it is trivially and immediately obvious from the perceptual point of view. But each neuron only filters a tiny part of the visual field. There is no homunculus in the brain. There is no ‘ghost in the machine’, and the perceptual consciousness of a given individual is precisely the great mystery that we would like to explain. On the neural level, the Gestalt principle of ‘good continuation’, which asserts that alignment will be perceptually prominent, can thus be taken to identify a formidable problem.

2.11.2 *Kanizsa’s Illusory Contours*

Figure 2.3 shows an example of a still more spectacular phenomenon. The red sectors of the concentric grey rings specify the boundary conditions generating the illusory (or subjective) contours which constitute one of the most enigmatic manifestations of the Gestalt properties of completion of missing sensory data. Furthermore, a pink-tinted square emerges from this configuration (neon or watercolour effect), showing that not only does the visual system construct long-range contours that do not exist in the sensory stimulus, but these hallucinated contours can serve as the edges for a colour-spreading process that is just as much a hallucination.

The transition from local to global works over a very long range here on the neural length scale, and this is why these phenomena have always been considered so particularly enigmatic.

Fig. 2.3 Example of a Kanizsa style illusory contour with a neon effect



2.11.3 Entoptic Phenomena

Our third example is the even more surprising case of visual hallucinations in which there is absolutely no stimulus, while the percept is richly structured from the geometric standpoint. Some of these purely geometric hallucinations relating to what has been called ‘entoptic vision’ were already classified long ago by Heinrich Klüver, who first brought *Gestalt theory* to the USA. Figure 2.4 shows some examples of these visual patterns perceived under the influence of mezcal. It also shows some neurogeometric models with a remarkable empirical fit which are due in particular to Paul Bressloff, Jack Cowan, Martin Golubitsky (see Bressloff et al. [25]) and will be discussed further in the second volume.






Fig. 2.4 Left I, II, III, IV: Visual hallucinations observed by Klüver. Right a, b, c, d: Neurogeometric models for the Klüver data. See [25] and the second volume

2.11.4 The Cut Locus


Our last example concerns the cut locus of a figure, also called the generalized symmetry axis or ‘skeleton’. Following the psychologist Blum [26], Thom [27] always stressed its fundamental role in perception (see Fig. 2.5).

Once again, imaging can show us the neural reality of the construction of this inner skeleton, for which there is no trace whatever in the sensory input, the latter consisting merely of an outer contour. Figures 2.6 and 2.7, produced by David Mumford’s disciple Tai Sing Lee, illustrate the response of a population of simple V1 neurons, whose preferred orientation is vertical, to textures with edges specified by opposing orientations. Up to around 80–100 ms, the early response involves only the local orientation of the stimulus. Between 100 and 300 ms, the response concerns the overall perceptual structure and the cut locus appears. These experiments are rather delicate to carry out, and they are much debated, but the detection of cut loci seems to be well demonstrated experimentally.

Fig. 2.5 Example of a cut locus. From Kimia [28]



Texture strip



Spatiotemporal Response

Fig. 2.6 Response to a stimulus whose form is specified by opposing textural orientations. From Lee [29]




Fig. 2.7 Recording of the construction of the cut locus. From Lee [29]

All these examples share the fact that the geometry of the percept is *constructed*—Husserl would say ‘constituted’—from sensory data *which do not contain it*, whence it must originate somewhere else. Put another way, they all involve subjective Gestalts. This is indeed why we chose them, because, as claimed by Jancke et al. [30], these subjective global structures ‘reveal fundamental principles of cortical processing’, the kind of principles that interest us here.

The origins of visual perceptual geometry can be found in the functional architecture which implements an immanent geometry, and it is the latter that provides the focus of neurogeometry. So the time has come to get down to business, by presenting some neurophysiological data for the receptive profiles and receptive fields of the visual neurons.

References

1. Weyl, H.: Space-Time-Matter. Dover, New York (1922)
2. Boi, L., Flament, D., Salanskis, J.-M. (eds.): 1830–1930: A century of geometry. In: Epistemology, History, Mathematics. Springer, Berlin (1992)
3. Toretti, R.: Philosophy of Geometry from Riemann to Poincaré, 2nd edn. Reidel, Dordrecht (1984)
4. Ryckman, T.: The Reign of Relativity. Oxford University Press, Oxford (2005)
5. Ryckman, T.: Hermann Weyl and ‘first philosophy’: constituting gauge invariance. In: Bitbol, M., Kerszberg, P., Petitot, J. (eds.) Constituting Objectivity, Transcendental Perspectives on Modern Physics, pp. 279–298. Springer, Berlin (2009)
6. Petitot, J.: Noncommutative geometry and transcendental physics. In: Bitbol, M., Kerszberg, P., Petitot, J. (eds.) Constituting Objectivity. Transcendental Perspectives on Modern Physics, pp. 415–455. Springer, Berlin (2009)
7. Souriau, J.M.: Géométrie symplectique et physique mathématique (Colloques Internationaux du CNRS 237) Paris (1975)
8. Petitot, J.: Actuality of Transcendental Aesthetics for Modern Physics. In: Boi, L., Flament, D., Salanskis, J.-M. (eds.) 1830–1930: A Century of Geometry, pp. 273–304. Springer, Berlin (1992)
9. Riemann, B.: Über die Hypothesen, welche der Geometrie zu Grunde liegen (1854). In: Dedekind R., Weber, F. (eds.) Bernhard Riemanns Gesammelte Mathematische Werke, pp. 272–287. Druck und Verlag von B.G. Teubner, Leipzig (1892)

10. Helmholtz, H. von.: Über die Tatsachen, die der Geometrie zu Grunde liegen (1868). In: *Wissenschaftliche Abhandlungen*, vol. II, pp. 618–639. Barth, J.A., Leipzig (1883). English Edition in *Epistemological Writings*, pp. 39–58. Reidel, Dordrecht (1977)
11. Merker, J.: *Le problème de l'espace*. Sophus Lie, Friedrich Engel et le problème de Riemann-Helmholtz, Hermann, Paris (2010)
12. Engel, F., Lie, S.: *Theorie der Transformationsgruppen*, Erster-Zweiter-Dritter und Letzter Abschnitt, Teubner, Leipzig, Berlin, 1888, 1890, 1890. Reprinted by Chelsea Publishing Co, New York (1970)
13. Lie, S.: *Theorie der Transformationsgruppen*, I, *Mathematische Annalen*, **16**, 441–528 (1880). English Translation: Ackermann, M., Hermann, R. Sophus Lie's 1880 Transformation Group paper, *Mathematical Science Press*, Brookline, MA (1975)
14. Poincaré, H.: *La Science et l'Hypothèse*, Flammarion, Paris, 1902. Republished with a Preface by J. Vuillemin (1968)
15. Poincaré, H.: *Science et Méthode*, Flammarion, Paris, 1908. Republished in *Philosophia Scientiae*, Kimé, Paris (1999)
16. Poincaré, H.: *La Valeur de la Science*, Flammarion, Paris, 1905. Republished with a Preface by J. Vuillemin (1970)
17. O'Keefe, J., Nadel, L.: *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford (1978)
18. Mumford, D.: *Pattern theory: A unifying perspective*. In: Knill, D.C., Richards, W. (eds.) *Perception as Bayesian inference*, pp. 25–62. Cambridge University Press, Cambridge (1996)
19. Lee, T.S., Mumford, D., Romero, R., Lamme, V.A.F.: *The role of primary visual cortex in higher level vision*. *Vis. Res.* **38**, 2429–2454 (1998)
20. Hodgkin, A.L., Huxley, A.F.: *Currents carried out sodium and potassium ions through the membrane of the giant axon of Loligo*. *J. Physiol.* **117**, 500–544 (1952)
21. Marre, O.: *Irregular but reliable activity in recurrent cortical networks: A theoretical and intracellular study in the primary visual cortex*. PhD, Ecole Polytechnique (2008)
22. Wilson, R.A., Keil, F. (eds.): *The MIT Encyclopedia of the Cognitive Sciences*. MIT Press, Cambridge, MA
23. Dupuy, J.-P.: *Mechanization of the Mind: On the Origins of Cognitive Sciences*. Princeton University Press, Princeton (2000)
24. Hess, R.F., Hayes, A., Field, D.J.: *Contour integration and cortical processing*. In: Petitot, J., Lorenceau, J. (eds.) *Neurogeometry and Visual Perception*, pp. 105–119. *J. Physiol.* **97**, 2–3 (2003)
25. Bressloff, P., Cowan, J., Golubitsky, M., Thomas, P., Wiener, M.: *Geometric visual hallucinations. Euclidean symmetry and the functional architecture of striate cortex*. *Philos. Trans. R. Soc. Lond. B* **356**, 299–330 (2001)
26. Blum, H.: *Biological shape and visual science*. *J. Theoret. Biol.* **38**, 205–287 (1973)
27. Thom, R.: *Modèles mathématiques de la morphogenèse*. Bourgois, Paris (1980)
28. Kimia, B.B.: *On the role of medial geometry in human vision*. In: Petitot, J., Lorenceau, J. (eds.) *Neurogeometry and Visual Perception*, pp. 155–190. *J. Physiol.* **97**, 2–3 (2003)
29. Lee, T.S.: *Computations in the early visual cortex*. In: Petitot, J., Lorenceau, J. (eds.) *Neurogeometry and Visual Perception*, pp. 121–139. *J. Physiol.* **97**, 2–3 (2003)
30. Jancke, D., Chavane, F., Naaman, S., Grinvald, A.: *Imaging cortical correlates of illusion in early visual cortex*. *Nature* **428**, 423–426 (2004)

Chapter 3

Receptive Fields and Profiles, and Wavelet Analysis


In this chapter, we focus on the ‘ascending’ visual pathways, i.e. from the retina to the V1 area or striate cortex, the most occipital area, located around the calcarine fissure (or calcarine sulcus), and we shall very (too) succinctly summarize the way in which their neurons can be considered to a first approximation as linear filters acting on the optical signal.

3.1 Structure of the Retino-Geniculo-Cortical Visual Pathways

We begin with some orders of magnitude. The human brain contains some 10^{12} neurons and about 10 times as many glial cells. Neurons are cells measuring a few μm long (but up to a few hundred μm), with axons that can be very long and very complicated dendritic trees. There are around 10^4 synapses per neuron, making a total of around 10^{16} synapses, which are either electric (gap junctions) or chemical (emission of neurotransmitters across the synaptic gaps by specialized vesicles).

It would not be possible here to present the history of neuroscience, even in a rudimentary way. The interested reader may consult the remarkable *Traité du Cerveau* (2006) by Michel Imbert [1]. This explains how, back in the fourth century BC, Hippocrates declared the brain to be the seat of consciousness, and how in the third century BC, Herophilus of Chalcedon and Erasistratus of Chios undertook to describe it. It also explains more recent progress, from Descartes to the Enlightenment (Diderot, La Métrie, Helvétius) and the pioneering work of the nineteenth and early twentieth centuries: Paul Broca, the discovery of functional areas (Hughlin Jackson, Gustav Theodor Fritsch, Eduard Hitzig), the discovery of neurons (Camille Golgi and Santiago Ramón y Cajal), brain maps, measurement of nerve impulses and action potentials, synaptic transmission (Charles Sherrington), right up to the formation of the first neuroscience associations and international programmes in the 1960s.

Fig. 3.1 Structure of the retino-geniculo-cortical visual pathways. From [3]



Neither would it be possible here to present the fine structure of the visual system. A great deal of information is available from many reliable websites such as [2–8], but also the books by Imbert [1] and Buser and Imbert [9]. Figure 3.1 gives a rather general idea. The optic nerve fibres from the left and right hemiretinas (corresponding to the right and left visual hemifields, respectively) come together at the optic chiasm and lead to the left and right lateral geniculate nucleus, respectively, and from this thalamic relay, onto the left and right visual areas $V1$, respectively. Each visual hemifield is thus projected onto the contralateral hemisphere, the nerve fibres from the nasal hemiretina crossing the chiasm and those of the temporal hemiretina remaining on the ipsilateral side. Transmission of signals from the retina to the lateral geniculate nucleus takes from 2.5 to 10 ms, and transmission to $V1$ from 40 to 60 ms.

Figure 3.2 shows that, in mammals, apart from the retino-geniculo-cortical pathway, there is also a pathway to the superior colliculus which is a subcortical brain structure, a kind of first step, followed by the pulvinar, providing a route from the retina that does not lead to $V1$ (the striate cortex), but to the extrastriate cortex $V3$ and MT (see also Feldheim and O’Leary [11], and see the beginning of Chap. 4 for the definition of these areas).

The retino-geniculo-cortical pathways comprise two large classes of cells known, at least in primates, as parvocellular, and magnocellular neurons. To these, we must add the koniocellular neurons, which project their axons onto the cytochrome oxidase

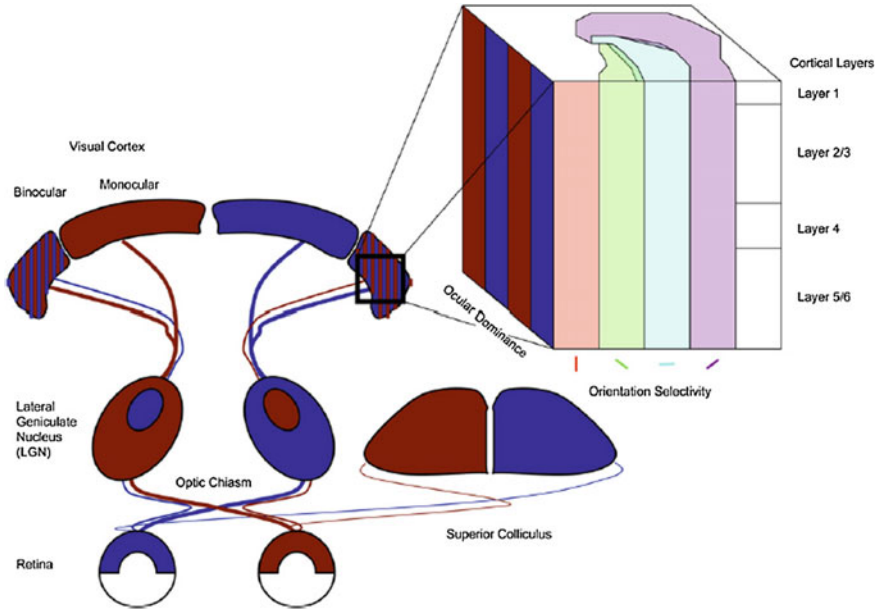


Fig. 3.2 A diagram showing the pathways from the retina to the visual cortex and the superior colliculus. From Hooks and Chen [10]

Table 3.1 Differences between the two classes of ganglion cells (GCs) of the retina, viz., *P/M* or *X/Y*

Parvocellular neurons (<i>X</i> cells)	Magnocellular neurons (<i>Y</i> cells)
Tonic cells: respond slowly (in ~ 30 ms) whenever the stimulus is present	Phasic cells: respond quickly (in ~ 10 ms) when the stimulus appears or disappears
Small	Large
Thin axons	Thick axons
Slow conduction rate (~ 6 m/s)	Fast conduction rate (~ 15 m/s)
Narrow centre, high centre-to-periphery gradient (see the next section)	Large centre, low centre-to-periphery gradient
High spatial resolution, low temporal resolution	Low spatial resolution, high temporal resolution
Sensitive to colour contrast	Insensitive to colour contrast
Analysis of spatial contrast (spatial analysis)	Detection of motion (temporal analysis)
Calculation of $\Delta G * I$ (see Sect. 3.3.4)	Calculation of $\partial(\Delta G * I)/\partial t$
Linear summation	Nonlinear summation

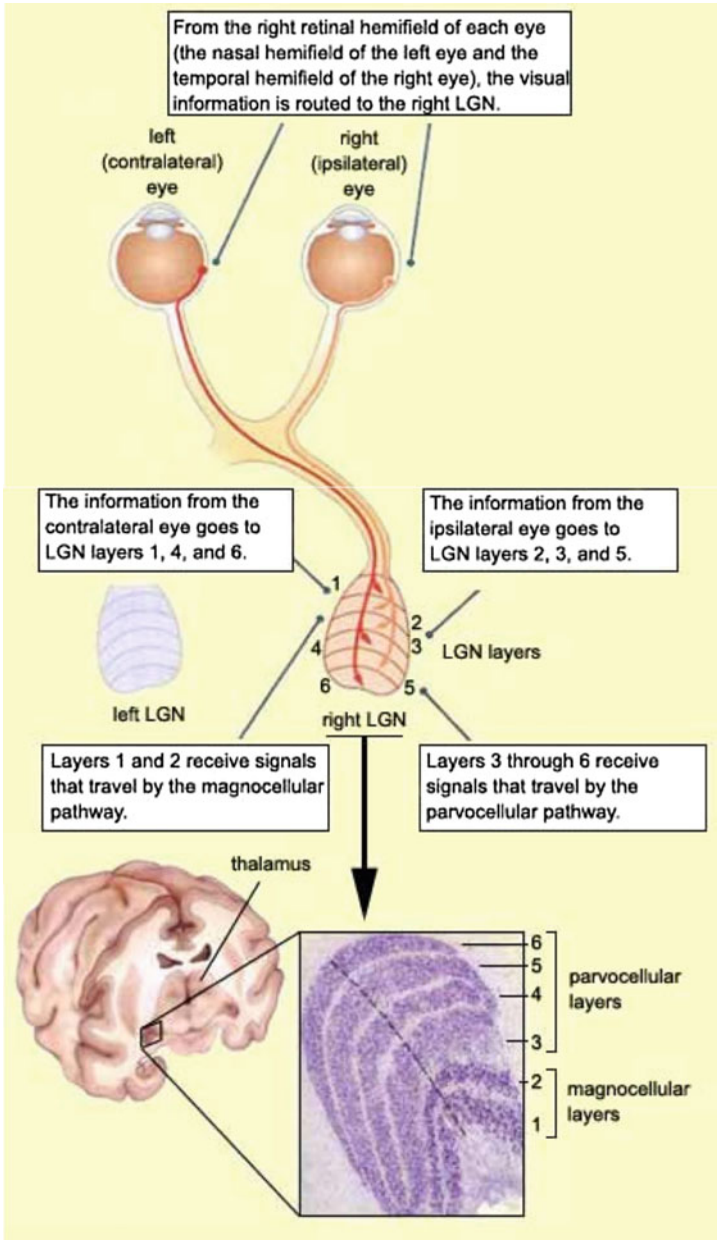


Fig. 3.3 Layered structure of the right lateral geniculate nucleus which receives the projections of the two right hemiretinas, i.e. the two left visual hemifields. We can see how ocular dominance is distributed: the left (contralateral) eye sends signals to layers 1, 4, and 6, and the right (ipsilateral) eye to layers 2, 3, and 5. We can also see how the parvocellular and magnocellular pathways are distributed: 1 and 2 for the magnocellular pathways, 4 and 3, 6 and 5 for the parvocellular pathways. From [3]




Fig. 3.4 From *V1*, the magnocellular and parvocellular pathways split into the ‘Where pathway’ and the ‘What pathway’. From *Wikipedia*, *Visual Cortex*

blobs processing colour in *V1*. In cats, the *P/M/K* cells correspond to the *X/Y/W* cells (see Shapley and Perry [12]), although there is still some debate about this homology. Table 3.1 shows the main differences between the two classes of ganglion cells (GCs) of the retina, viz., *P/M* or *X/Y*, corresponding to about 80% and 10%, respectively, of all the cells.

The parvo- and magnocellular pathways are finely distributed through the layers of the lateral geniculate nucleus, as shown in Fig. 3.3. For example, for the right LGN, the left (contralateral) eye sends signals to layers 1, 4, and 6 and the right (ipsilateral) eye to layers 2, 3, and 5, while the magnocellular pathways lead to layers 1 and 2 and the parvocellular pathways to layers 4 and 3, 6 and 5.

From *V1*, the magno- and parvocellular pathways split, one—called the dorsal stream—going up to the parietal areas for motion, close to those for motor functions, the other—the ventral stream—going to the temporal areas for shape recognition, close to those for language. Ungerleider and Mishkin [13] called the first (magno) the ‘Where’ pathway and the second (parvo) the ‘What’ pathway (see Fig. 3.4).¹ The ‘Where pathway’ or ‘dorsal stream’ goes from *V1* to *V2* and *V5 (MT)*, while the ‘What pathway’ or ‘ventral stream’ goes from *V1* to *V2*, *V4*, and *IT* (inferior temporal area) which are no longer retinotopic. We shall discuss these areas in Chap. 5.

¹Note that these are Aristotelian categories. We shall return to this philosophical aspect in the conclusion of the second volume.

3.2 Receptive Fields and Receptive Profiles

Let us now discuss some basic neurophysiological concepts and some related experiments, involving new experimental techniques that can observe the activity of the micromodules taking part in visual processes, even though the results are still hard to interpret given that the resolution is only a few hundred neurons, each containing hundreds or even thousands of synapses. The first concept required to understand how, even at the very lowest levels, the visual system formats the optical signal is the concept of *receptive field* (RF).

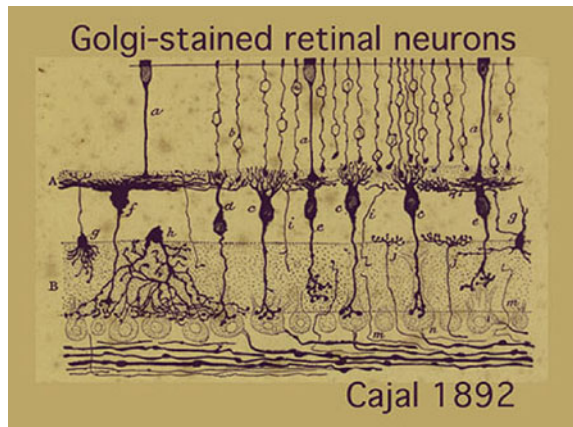
3.2.1 Structure of the Retina

From the earliest sensorial and peripheral level of the ganglion cells in the retina, which carry out what is known as *transduction* (neural coding) of the signal by measuring it and transforming it into neural information that can be exploited by the central nervous system, the signal is formatted geometrically. At the end of the nineteenth century, Camillo Golgi (1885), Ramón y Cajal (1892), and others were already beginning to understand the microanatomy of the retina (see Fig. 3.5).

With a thickness of around 0.25 mm (0.1–0.5 mm), apart from the Müller glial cells, the retina comprises five layers (3 + 2) and even as many as ten if we include the retinal pigment epithelium (RPE), the bounding membranes, and the layer of fibres of the optical nerve. The three main layers are as follows:

- (i) A layer containing photoreceptors. These constitute the ‘measurement device’ (in the quantum sense) which picks up the photon signal and pixelates it. These have dimensions of a few μm and number around 6×10^6 cones and 120×10^6 rods (about twenty times as many) in humans.

Fig. 3.5 One of the first cross sections of the retina, by Ramón y Cajal, in 1892 [5]



- (ii) A layer containing bipolar cells which connect the photoreceptors to the third layer.
- (iii) A layer containing ganglion cells (GCs), whose axons are the fibres of the optical nerve. There are of the order of 1.5×10^6 of these, and the compression ratio of photoreceptors \rightarrow GCs, what neurophysiologists call the ‘degree of convergence’ of the photoreceptors on the GCs, is thus of the order of 100. GCs can respond to few photons (see Barlow et al. [14]).

These three main layers are joined together by two ‘horizontal’ intermediate layers known as *plexiform layers*, one comprising the horizontal cells (the outer plexiform layer between photoreceptors and bipolar cells, involved mainly in spatial analysis of stimuli) and the other comprising amacrine cells (the inner plexiform layer between the bipolar cells and the ganglion cells, involved mainly in the temporal analysis of stimuli). Figures 3.6, 3.7 and 3.8 show the structure of the retina.

Through the complicated ‘vertical’ and ‘horizontal’ connectivity of these layers, each GC is joined to a precise domain on the retina (a disk with an area of $0.1\text{--}2\text{ mm}^2$ and a visual angle of $0.5\text{--}10^\circ$) which is called its receptive field (Fig. 3.9). This concept was introduced in 1938 by Hartline [15] (Nobel prize 1967).

3.2.2 Neurons and Action Potentials

Since we are concerned here with integrative and functional cognitive neuroscience, we shall say almost nothing about the fine structure of neurons. However, a certain minimum will be useful. Figure 3.10 shows a typical neuron (a pyramidal neuron in layer V of V1) with its soma, nucleus, dendrites, and axon.

The basic signals transmitted through neural networks are action potentials, or spikes, emitted by the neurons along their axons. Spikes are formed by a universal process involving the opening and closing of ion channels in cell membranes. The flow of sodium, potassium, calcium, and chlorine ions (Na^+ , K^+ , Ca^{++} , Cl^-) through their specific channels defines a rest potential (RP) $V \approx -75\text{ mV}$. The membrane is initially polarized. The respective equilibrium potentials (EP) of the ions are typically $EP_{\text{Na}^+} = +62\text{ mV}$, $EP_{\text{K}^+} = -80\text{ mV}$, $EP_{\text{Ca}^{++}} = +123\text{ mV}$, and $EP_{\text{Cl}^-} = -65\text{ mV}$. A depolarization increases the RP and, when it reaches the threshold -45 mV , the neuron discharges very quickly ($\sim 1\text{ ms}$), emitting a spike which propagates along the axon with a speed from 0.1 m/s up to many tens of m/s . The discontinuous nature of the spikes is due to a positive feedback of the depolarization on the opening of the Na^+ channels, which are very fast; the RP increases catastrophically to $+60\text{ mV}$ (roughly the EP of Na^+). The Na^+ channels are subsequently deactivated. But the sudden large depolarization opens the K^+ channels, which are slow, and the membrane is repolarized, returning to the basic RP value after a period of hyperpolarization (-85 mV , roughly the EP of K^+).




Fig. 3.6 The 5 layers of the retina. Lumière = Light, Couche des ganglionnaires = Ganglion cell layer, Couche granuleuse interne = Inner granular layer, Couche des photorécepteurs = Photoreceptor layer, Cellule ganglionnaire = Ganglion cell, Cellule amacrine = Amacrine cell, Cellule bipolaire = Bipolar cell, Cellule horizontale = Horizontal cell, Bâtonnet = Rod, Cône = Cone. From [2]

In the second volume, we shall return to the mathematical models (FitzHugh–Nagumo, etc.) of these processes, which derive from the fundamental model of Hodgkin-Huxley [16].

3.2.3 Structure of the Photoreceptors

The photon flux of the optical signal is converted into action potentials by the photoreceptors. This transduction reveals them to be quite remarkable quantum detectors. Figures 3.11 and 3.12 show the structure. There are two main types:





Fig. 3.7 Cross section of the retina. From [4]

Fig. 3.8 Cross section of the author’s retina obtained by Optical Coherence Tomography (OCT). Courtesy of Dr. David Sayag



- The rods are sensitive to low light intensities and black/white contrast. They have high sensitivity but low definition, and regenerate slowly.
- The cones are sensitive to high light intensities and colour. They have low sensitivity but high definition, and regenerate quickly.

Photoreceptors have two ‘segments’. The outer segment, where the signal is actually transduced, contains molecules of a photopigment called *rhodopsin*, inserted into its phospholipid membrane. It is constituted by a pile of around 2000 discs, each containing about 40 million rhodopsin molecules.




Fig. 3.9 Simplified diagram of the receptive field of a ganglion cell. Nerf optique = Optic nerve, Fibre optique ou axone = Axon, Cellule ganglionnaire = Ganglion cell, Cellule amacrine = Amacrine cell, Cellules bipolaires = Bipolar cells, Cellules horizontales = Horizontal cells, Photorécepteurs = Photoreceptors, Pourtour = Periphery, Centre = Centre, Champ rétinien = Retinal field. From [2]

The rhodopsin molecule, shown in Fig. 3.13, is a protein made from opsin and a molecule of the chromophore retinal, derived from vitamin A, at the 11-cis position. In fact, vitamin A in the form of the alcohol group retinol is oxidized to retinaldehyde or retinal by the enzyme dehydrogenase. This 11-cis retinal gives rhodopsin a very strong absorption spectrum for visible light in the range 400–600 nm, with a peak around 500 nm. It is the variation in the amino acid sequences of the opsins in the cones, a variation controlled by genes located on the *X* chromosome (OPN1, OPN = opsin), which is responsible for the spectral differences between the various pigments of the long/middle/short (*L/M/S*) cones. We shall return to this in Sect. 3.2.5. However, the 11-cis retinal chromophore remains unchanged.

A photoreceptor detects a photon through an isomerization that alters the configuration of the opsin by a *cis*–*trans* transition of the retinal, thereby triggering a cascade of enzyme activity. When the rhodopsin is activated in this way, it binds to transducin, and this in turn causes the enzyme phosphodiesterase to act on 3–5

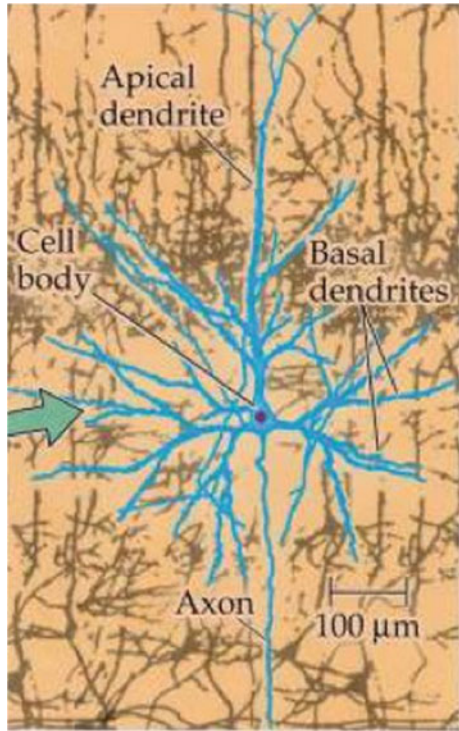


Fig. 3.10 Structure of a typical neuron with its cell body, nucleus, dendrites, and axon. From [6]




Fig. 3.11 Structure of the outer segment of a photoreceptor. From [5]




Fig. 3.12 Disk structure of a photoreceptor. From [5]




Fig. 3.13 Structure of rhodopsin. From [5]

cyclic guanosine monophosphate (cGMP). The latter controls the flow of cations into the outer segment. The Na^+ , Ca^{++} , and K^+ channels are open in darkness and the voltage across the photoreceptor membrane is -40 mV. Light causes the Na^+ and Ca^{++} channels of the outer segment to close, and this induces a hyperpolarization of the photoreceptor membrane, which itself induces a hyperpolarization of the OFF horizontal and bipolar cells and a depolarization of the ON horizontal and bipolar

cells. This is followed by a depolarization of the amacrine cells and the ganglion cells, and hence spike emission in the ON/OFF regions of their receptive profiles. We shall return to the definition of these ON/OFF regions in Sect. 3.2.6.

The process of disk formation or morphogenesis has been studied in detail. The discs are formed through the action of an intermediate filament (IF) protein called peripherin, at the proximal end (the base) of the outer segment, where it is joined to the inner segment by a connecting cilium (flagellum). The ciliary plasma membrane transforms into a closed disk from the outer edge: opsin is synthesized in the inner segment and transferred to the outer segment by the cilium. The growth rate is about $2.5 \mu\text{m}/\text{day}$ in humans. The reader may consult the paper by Arikawa et al. [17]. Figure 3.14 shows the development of cones and rods, and also their structure.

In humans, this morphogenesis can be observed non-invasively *in vivo* using a very high resolution OCT and adaptive optics techniques transforming the outer segment into a kind of ‘biological interferometer’ (see, for example, Jonnal et al. [19]).

At the distal end, the outer segment of a photoreceptor is in contact with the retinal pigment epithelium (see Fig. 3.7), a tissue containing some 5 million cells, each acting on several tens of photoreceptors. The pigment epithelium (PE) plays a

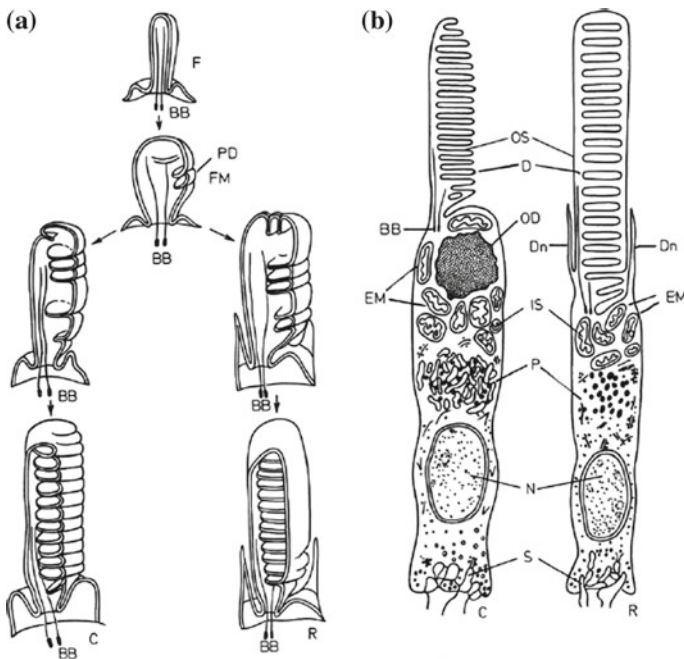



Fig. 3.14 **a** Development of cones and rods. **b** Their structure. C = cone, R = rod, F = flagellum (cilium), FM = folds of the ciliary membrane, BB = basal body, PD = presumptive disk, OS = outer segment, D = disk, OD = oil droplet, Dr = dendrite, EM = ellipsoid mitochondria, IS = inner segment, P = paraboloid, N = nucleus, S = synapse. From Ruben Adler and Pamela Raymond [18]

Fig. 3.15 Example of a ganglion cell, taken from Helga Kolb et al. [21]



fundamental role, regulating the subretinal medium on the microlevel and carrying out several functions, as an antioxidant, for the transport and deposit of retinoids derived from vitamin A, for phagocytosis of discs in the plasma membrane of the outer segments, and for regenerating the visual pigment rhodopsin. About 10% of the discs are renewed each day; old ones detach and are phagocytosed. Apoptosis of PE cells occurs in age-related macular degeneration (AMD).

On the genetic level, the development, differentiation, and regulation of photoreceptors are controlled by a network of transcription factors (proteins binding to specific DNA sequences and controlling the transcription of their genetic information into mRNA) centred on the *Crx* gene (cone-rod homeobox containing gene). This gene contains a homeobox (a DNA sequence involved in regulating morphogenetic processes) in the *Otx* family (orthodenticle homeobox gene) (see, for example, Henning et al. [20]). We shall return to genetic control in Sect. 3.2.5 of this chapter and Sect. 5.12.4 of Chap. 5.

3.2.4 Ganglion Cells

The circuitry of the retina is very subtle. There are 8 different kinds of cones, 3 or 4 kinds of horizontal cells, 9–11 kinds of bipolar cells, and over 23 kinds of amacrine cells. Regarding the ganglion cells, there are 15–20 kinds (very probably 18), with different physiological properties, different cortical targets, different dendritic trees in terms of size, shape, and stratification, and different kinds of connected bipolar and amacrine cells, with all these differences being controlled genetically. Figure 3.15, taken from the paper [21] by Helga Kolb et al., shows a 500 μm human GC.²

²The paper contains many other examples of retinal cells.

As already mentioned earlier, parvocellular (or X) GCs, which detect spatial contrasts, act as essentially linear filters, while magnocellular (or Y) GCs, which detect movements, are more nonlinear.

Moreover, the GCs have remarkable acuity. As pointed out by Barry Lee et al. [22, p. 2743]:

[...] a single cell can signal stimulus position within its receptive field to within a fraction of the centre diameter.

This paper describes physiological experiments on macaques and psychophysical experiments on humans. One of the protocols consists in flashing edges at 300 ms intervals, shifting them each time by an increment of $\Delta = 36$ arcsec to 1.5 arcmin. Discrimination of the order of 2 arcmin is then observed for GCs and receptive fields with centre diameters of 14 arcmin and 36 arcmin, respectively.

Regarding transcription factor networks controlling the morphological and functional diversity of the GCs, the reader may refer to the work of Badea et al. [23]. Many GCs express the genes *Brn3a* and *Brn3b* in their development. If *Brn3a* is not expressed, the dendritic tree is dramatically modified. If *Brn3b* is not expressed, many GCs become amacrine cells or horizontal cells and die; axon development, axon morphology, path-finding, and target selection become incoherent (see Sect. 5.12.4 of Chap. 5).

3.2.5 Retinal Colour Coding Circuitry

Here, we discuss briefly the genetic control of the cone \rightarrow GC retinal circuitry involved in colour processing. The photoreceptors are sensitive to wavelengths λ reflected by objects in a certain interval of a few hundred nanometres in the order infrared, red, orange, yellow, green, cyan, blue, violet, ultraviolet. The three kinds of cone of the human being, a trichromatic primate, are $L/M/S$ (many mammals are dichromatic or have only rods, while certain species of birds and fish are tetrachromatic). The class $L = \text{long} = \text{red}$ have an absorption peak near 565 nm and absorb over a range of about 500–600 nm, $M = \text{medium} = \text{green}$ have an absorption peak near 535 nm and absorb over a range of about 480–580 nm, and $S = \text{short} = \text{blue}$ have an absorption peak near 420 nm and absorb over a range of about 400–440 nm. We can distinguish some 10 million colours, with a discrimination of about 1 nm in the green and yellow and about 10 nm in the red and blue.

Since the pioneering work by Jeremy Nathans (see, for example, Nathans et al. [24]), genetic control has been pretty well understood. The reader may consult the paper [25] by Jay and Maureen Neitz. In humans, the opsin genes *OPNILW*, *OPNIMW*, and *OPNISW* (see above) controlling the difference between the long-wave-sensitive (LW), medium-wave-sensitive (MW), and short-wave-sensitive (SW) cones are located at position $Xq28$ on the X chromosome for the LW and MW cones and at position $7q32$ on chromosome 7 for the SW cones. The specialists think that the L/M difference may have arisen in primates from an allelic

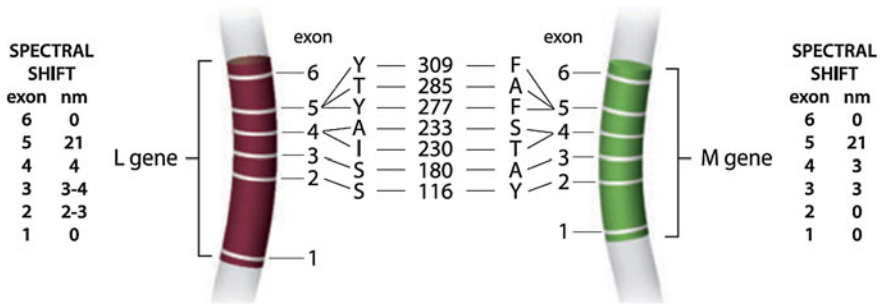


Fig. 3.16 The 6 exons of the genes *OPNILW* and *OPNIMW* distinguishing the *L/M* cones. These are the *white lines* separating the introns. The initials of the amino acids and the numbers indicating their codons are as follows: *Y* = tyrosine, *T* = threonine, *A* = alanine, *I* = isoleucine, *S* = serine, *F* = phenylalanine. For exon 5, *L* has *Y/277, T/285, Y/309*, while *M* has *F/277, A/285, F/309*. The spectral shifts induced by these differences in the amino acids are given in nm. From [25]

diversification of a single gene. Indeed, the *L* and *M* opsin genes each have 6 exons, the first and sixth being identical and the *L/M* spectral shift being explained essentially by the difference in exon 5. Figure 3.16 shows the 6 exons for *L/M* separating the introns. Once again, we shall return to these questions in Sect. 5.12.4 of Chap. 5.

Genetically controlled in this way, the retinal colour circuitry is rather subtle. There is a centre/surround opponency in the colour sensitive GCs corresponding to opposition of colours: these are the opponent cells. In dichromatic species, the basic opposition is *S* vs *L/M*, with *L* – *S*³ giving *Y* (yellow) and *S* – *L* giving *B* (blue).⁴ Trichromatic species evolved from dichromatic ones by separation of *L* and *M*. This gave *M* – (*S* + *L*) (*G* = green) cells and *L* – (*S* + *M*) (*Y* = yellow) cells, both ON and OFF, with specific circuitry. Figure 3.17, taken from Neitz [25], is a schematic view of the different classes of colour sensitive GCs.⁵

Figure 3.18, also taken from Neitz [25], shows how the various types of cones are connected to the GCs via the bipolar cells in a trichromatic primate. For example, in the first column on the left, we see how an *S* cone (blue) and an *M* cone (green) can be involved in an ON-centre *M* (green) GC with *M* – (*S* + *L*) (green/red) opponency and an OFF-centre (*S* + *L*) (red) GC with (*S* + *L*) – *M* (red/green) opponency.

This kind of work on the retinal circuitry confirms some very early scientific intuitions, and we can only admire the insight of those responsible for them. As early as 1802, Thomas Young put forward the idea of three kinds of retinal receptors. In 1866 Hermann von Helmholtz carried out psychophysical experiments which

³ *A* – *B* indicates a colour opposition between the centre and the periphery of the GC.

⁴ Some specialists like Nathans think that the blue/yellow (*B/Y*) opposition evolved from single-celled organisms in which it controlled the circadian (day/night) response by signalling large spectral changes in the sun's light.

⁵ Even though the names are the same, the colour labels of the cones should not be confused with those of the GCs. The mechanisms are not the same. The first refer to spectral absorption properties and the second to circuitry and a receptive field structure.

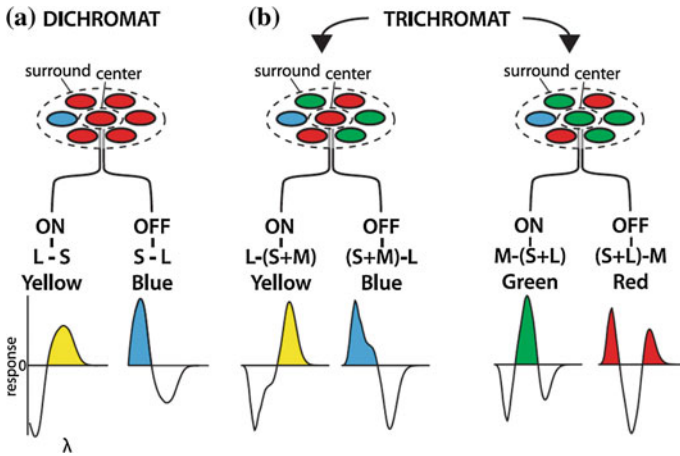


Fig. 3.17 Different classes of colour sensitive ganglion cells. **a** Dichromatic GCs. The basic opposition is S vs L , with $L - S$ giving Y and $S - L$ giving B . **b** Trichromatic GCs. Here L and M are separated, whence $L - (S + M)$ cells give Y , $(S + M) - L$ cells give B , $M - (S + L)$ cells give G , and $(S + L) - M$ cells give R . From Neitz [25]




Fig. 3.18 Retinal circuitry connecting the different types of cone to GCs (*diamonds*) via bipolar cells (*circles*) in a trichromatic primate. The various layers of the retina are indicated *on the right*. From Neitz [25]

supported this hypothesis and developed what has since become known as the Young–Helmholtz trichromatic theory. Later, in 1892, Ewald Hering introduced the idea that there are in fact 4 fundamental colours RGBY, forming two pairs of complementary colours R/G and B/Y (the opponent process theory). On the philosophical level, Goethe defended a physiological theory of colours in his famous *Farbenlehre* of 1810, thus opposing the purely physicalist theory of Newton. We shall return to this story in the second volume.

3.2.6 General Receptive Fields and Neural Coding

3.2.6.1 Receptive Fields and Profiles


Given a visual neuron belonging to the retino-geniculo-cortical pathways going from the retina to the cortex via the thalamic relay of the lateral geniculate nucleus, the simplest definition of its receptive field is the zone D of the retina to which it is connected through this highly complicated connectivity. So the neuron can respond to stimuli inside D by emitting spike trains. It can be shown that there are regions of the receptive field (RF)—said to be ON —which have positive excitatory response to pointlike light signals (effectively Dirac deltas, this is a case of impulse response). Other regions—said to be OFF —have negative inhibitory response. This leads to the concept of the receptive profile (RP) of a receptive field. The RP is a function $\varphi : D \rightarrow \mathbb{R}$ defined on the domain D of the RF and representing the response $\varphi(x, y)$ of the neuron (positive for ON and negative for OFF) to stimulation at the point (x, y) . It can be considered as the transfer function of the neuron treated as a filter. For GCs in the retina, the spatial $ON-OFF$ centre/surround antagonism was discovered by Stephen Kuffler [26]. The classic Fig. 3.19 shows a few examples of RFs.

In fact, the definition of the RF can vary enormously depending on the way we define the response of a visual neuron. The impulse responses in terms of spike trains lead to a rather narrow classical concept of RF called the *minimal discharge field* (MDF). But since a threshold of the membrane potential must be exceeded in order to trigger a spike (neurons act like threshold automata), a neuron may have many subthreshold responses for a given global contextual activity. As shown by Yves Frégnac (Frégnac [27], Frégnac and Shulz [28], Seriès et al. [29], but see also Maffei [30], Gilbert [31] and Lamme [32]), the classical concept of RF can thus be refined. In this section, we shall only consider the MDF, returning to the generalization in Chap. 4 and Sect. 5.3.1 of Chap. 5.

The spikes and subthreshold activity of the membranes determine the neural coding. One coding is based on the frequency of the spikes in the spike trains, referred to as *rate coding*. Temporal accuracy needs to be taken into account. For example, it has been shown (Butts et al. [33]) that the accuracy of the spike trains of cells in the visual system (retina, LGN, cortex) for white noise stimuli is of millisecond order. There is lower accuracy for natural images in which the spatial frequency spectrum is different, but the ratio of the relative accuracy to the time required to process the stimulus remains roughly constant.

However, there is another kind of coding called *rank coding*, introduced by Simon Thorpe and other specialists in this field [34]. The idea arose from experiments on ultrafast image categorization. For example, we can recognize the image of an animal when it is flashed for only 20 ms. When we calculate the transfer times involved, we conclude that only the first few spikes could possibly be involved in the coding, whence rate coding would not be plausible here. This leads to the idea of rank coding, which asserts that this kind of classification is carried out on the basis of the order of the first spikes. As explained by Thorpe [34]:

Fig. 3.19 Examples of receptive fields and profiles. *Upper neuron* in the lateral geniculate nucleus with *ON* centre and *OFF* surround. A light impulse in the centre activates the neuron (spike train), while one in the surround inhibits it. *Centre V1 neuron* sensitive to a specific orientation (vertical). *Lower neuron* sensitive to the direction of motion. It only responds significantly when the horizontal bars go down. From [6]



In an attempt to explain this sort of ultra-rapid processing I proposed a novel coding scheme that uses the order in which cells fire spikes, rather than firing rates, to encode information. It turns out that using such a code may allow us to recognize objects when as few as 1% of the neurons in the visual pathways have fired a spike.

3.2.6.2 Receptive Profiles of the Retina and the LGN

Here, we shall say a few words about the structure of the receptive profiles. Highly developed methods of electrophysiology (see, for example, DeAngelis et al. [35]) have been able to measure the *level curves* of the RPs of different visual neurons. Using sinusoidal gratings of light and dark bands and varying the widths, spatial frequency, orientation, speed, temporal frequency, and size of the bands, the position, size, and preferred orientation of the RF can be precisely specified. The next step is to use fast and random series of pointlike or bar-shaped flashes, e.g. lasting 50 ms, located on a grid, e.g. 20×20 , at intervals of about 100 ms to 1 s (white noise

analysis). This produces a few thousand spikes. The correlation between the input (flashes) and the output (action potentials), possibly with a time delay parameter, delivers the transfer function of the neuron. This is an astonishing experimental feat and the results are quite remarkable.

We shall consider two main types of RP. To begin with, it is a standard result of neurophysiology, already stressed by David Marr at the end of the 1970s, that the RFs of retinal ganglion cells have RPs that are Laplacians of Gaussians, viz., ΔG . The same is true for the cells of the lateral geniculate nucleus (see Figs. 3.20, 3.21 and 3.22).

3.2.6.3 Spatial Receptive Profiles of V1

We now move on from the retina and the LGN to V1. There are several classes of visual neurons in V1 and V2 (areas 17 and 18, respectively, for the cat):




Fig. 3.20 Receptive profile of an *ON*-centre cell in the lateral geniculate nucleus. *Left* Diagram of the RP with + (*ON*) and - (*OFF*) domains in *green* and *red*, respectively. *Right* Level curves of the RP. From DeAngelis et al. [35]

Fig. 3.21 The Laplacian of a Gaussian model ΔG of a ganglion cell and an *ON* cell of the lateral geniculate nucleus


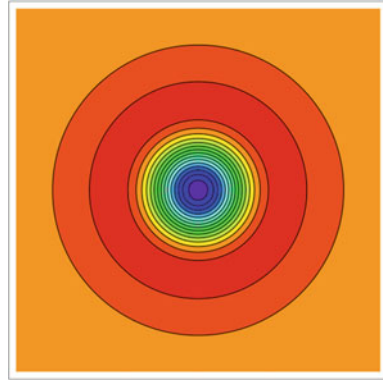


Fig. 3.22 Level curves of a Laplacian of a Gaussian ΔG , to be compared with the empirical results of Fig. 3.20



1. The so-called ‘simple’ cells have an anisotropic RF elongated in some preferred orientation. They are sensitive to the position and orientation of an edge or a grating at rest or in motion. They dominate in layer 4. Since they all have inputs from the LGN, Hubel and Wiesel conjectured that their RPs could be constructed by ‘converging’ the patterns of LGN receptive fields: if we consider a short row of *ON*-centre cells of the LGN whose RFs are aligned and if we couple it with a parallel row of *OFF*-centre cells, we obtain the characteristic anisotropy of simple cells. Other hypotheses, no longer based on converging thalamocortical afferents, but, e.g. local intracortical excitation, have been investigated (see, e.g. Somers et al. [36] and Troyer et al. [37]). Computer models of orientation selectivity using afferent and intracortical connections have also been built (see, e.g. Wörgötter and Koch [38]).
2. The so-called ‘complex’ cells (layers 2, 3, and 5) also respond to edges and gratings, but their RFs are not patterned into antagonistic *ON*–*OFF* zones and have no *preferred* orientation. They only respond to moving stimuli and are sensitive to the direction of motion. One way to construct them is to converge the RPs of simple cells with different preferred orientations.
3. The so-called ‘hypercomplex’ cells are also sensitive to orientation, movement, and direction, but only respond significantly to edges ending *inside* their RF. They thus have an ‘end-stopping’ structure (essential for detecting the ends of edges and corners). It was originally thought that these only occurred in *V2*. However, Geoffrey Henry, Bogdan Dreher, and Charles Gilbert (a doctoral student of Hubel and Wiesel) discovered them in *V1*. In fact, *V1* contains both simple and complex ‘end-stopped’ cells. Apart from the pattern of *ON* and *OFF* regions, their RF has end-zones which inhibit their activity when the stimulus is present there.

There are many computational models for these different classes of cells with appropriate circuitry. They can be as follows:

- (i) ‘Hierarchical’, building the RPs of a higher level by converging the RPs of lower levels (hierarchical model of convergence of LGN RFs to RFs of simple cells in layer 4, or the hierarchical model of convergence of simple cell RPs to complex cell RPs).

- (ii) ‘Parallel’, with complex cells built directly by convergence from patterns with *no* preferred orientation of LGN cells, and not from simple cells.
- (iii) ‘Recurrent’, taking into account the fact that the complex cells receive most of their input from other cortical layers, and thus building RPs from *intracortical* circuits.

The reader will find a survey in Martinez-Alonso [39] and a discussion of the relations between the RF properties and the underlying cortical circuitry in Gilbert [40]. However, as already mentioned, we shall not go into the technical details of these circuits, but focus rather on the structure of the RPs.

The receptive profiles of the simple V1 neurons most commonly encountered comprise an *ON* domain extended in one preferred orientation, with two smaller *OFF* domains on either side. A simple idealized model is provided by the second derivative of a Gaussian, viz., $\partial^2 G / \partial x^2$ (see Fig. 3.23). But since the experimental data cannot be particularly accurate, they can also be interpreted as Gabor wavelets (see e.g. [41]). Figure 3.24 compares the former, viz.,

$$\varphi(x, y) = \frac{\partial^2 G}{\partial x^2}, \quad G = \exp[-(x^2 + y^2)],$$

with the Gabor wavelet model

$$\varphi(x, y) = \exp(i2x) \exp[-(x^2 + y^2)] \quad (\text{real part}).$$

There are several neurophysiological schemes for the orientational selectivity of V1 neurons. They use different combinations of thalamic feedforward signals and intracortical feedback, together with different short-range excitatory and inhibitory connection effects (see, for example, Kayser et al. [42], Carandini et al. [43], Ben-Yishai et al. [44], McLaughlin et al. [45]). But even though the underlying mechanisms are




Fig. 3.23 Idealized standard receptive profile of a simple V1 neuron




Fig. 3.24 Two models of the standard receptive profile of a simple V1 neuron. *Upper* Second derivative of a Gaussian, viz., $\varphi(x, y) = \partial^2 G / \partial x^2$, with $G = \exp [- (x^2 + y^2)]$. *Lower* Gabor wavelet $\varphi(x, y) = \exp(i2x) \exp [- (x^2 + y^2)]$ (real part)

still debated, it is nevertheless clear that intracortical connections select a preferred orientation and eliminate the others. Moreover, this selectivity is acquired very early in development. In monkeys, orientation discrimination and selection (tuning) are innate, and in humans, although weak at birth, they are as good as in adults by the age of 3 months.

There are other RPs. For example, Fig. 3.25, produced by Gregory DeAngelis and coworkers at Berkeley, shows the RP of a simple V1 neuron comprising two symmetric *ON/OFF* regions extended in a preferred orientation, with smaller *OFF* and *ON* domains on either side. These two types of RPs (a central domain with two smaller antagonistic domains on either side, or two equivalent antagonistic central domains with two smaller domains on either side) were found right from the start by Hubel and Wiesel. Figures 3.26 and 3.27 show two models: a third derivative of a Gaussian, viz.,

$$\varphi(x, y) = \frac{\partial^3 G}{\partial x^3}, \quad G = \exp [- (x^2 + y^2)],$$

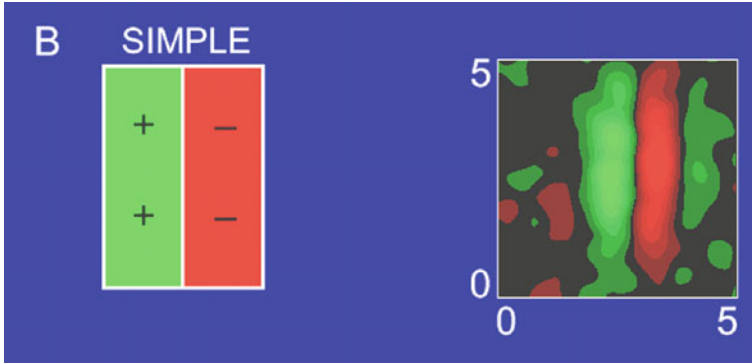


Fig. 3.25 Another receptive profile of a simple V1 neuron. *Left* Structure of the RP with + (ON) and - (OFF) lobes. *Right* Level curves. From DeAngelis et al. [35]




Fig. 3.26 Model of the RP of a simple V1 neuron using a third derivative $\partial^3 G / \partial x^3$ of a Gaussian

and a Gabor wavelet model

$$\varphi(x, y) = \exp(i4x) \exp [- (x^2 + y^2)] \quad (\text{imaginary part}) .$$

3.2.6.4 Spatiotemporal Receptive Profiles

If we introduce the correlation delay as a time variable t , we can also measure *spatiotemporal receptive profiles*. In these ‘reverse correlation’ techniques, we calculate the cross-correlation of the spike train of the neuron with a random sequence of bright and dark bars with an orientation corresponding to the preferred orientation of the neuron and flashed at different positions in the RF (typically, bars of length




Fig. 3.27 Model of the RP of a simple V1 neuron using a Gabor wavelet

15°, flashed for about twenty ms and at about thirty positions in the RF). Figure 3.28 exemplifies for LGN cells.

In general, the spatiotemporal RP has a *biphasic* time response, i.e. like the first derivative of a Gaussian, and a centre/surround spatial organization, i.e. like the second derivative of a Gaussian. The RP is then the third derivative of a Gaussian (see Figs. 3.29 and 3.30). When bars with sinusoidally modulated luminance are applied at different positions with a certain frequency and the cell response is recorded, one can measure its latency. When the latency is low (< 100 ms) and the first phase of the time response dominates, the cell is said to be *non-lagged*. When the latency is high (> 100 ms) and the second phase dominates, the cell is said to be *lagged*. The role of lagged and non-lagged cells is particularly important in the LGN, because it could explain how orientation selectivity in V1 is built up from the LGN (see Saul [47]).

The above example is a case of *space/time separability*, which means that the *ON/OFF* regions of the RP do not change position, only intensity, the latter being uniformly modulated by a temporal profile. In other words, the *spatiotemporal RP* $\varphi(x, y, t)$ is a product of the form $\varphi(x, y, t) = \varphi(x, y)\psi(t)$. As already mentioned, $\psi(t)$ is often biphasic, i.e. like the first derivative of a Gaussian, and profiles like the one above are obtained from $\partial^3 G / \partial x^2 \partial t$. However, there are also monophasic cases $\psi(t) = G$, and triphasic cases $\psi(t) = \partial^2 G / \partial t^2$, the latter given by fourth derivatives of a Gaussian, viz., $\partial^4 G / \partial x^2 \partial t^2$.

The fit between models and experimental data can be excellent. For example, there are neurons in which the spatial profile and the temporal profile are not separable. When there is *non-separability*, the *ON/OFF* regions move, the response of the surround is delayed relative to the response of the centre, and this means that the neuron can be motion-selective and detect edge speeds. Figure 3.31, taken from DeAngelis et al. [35], shows the temporal evolution of just such a non-separable receptive profile. We see here the way the receptive profile of a neuron of this kind evolves in time. Figure 3.32 shows a model using derivatives of a Gaussian. In fact, it uses a third derivative $\partial^3 G / \partial u^2 \partial v$, the (x, t) plane being obtained by rotating the (u, v) plane through $\pi/10$. The fit to the data is remarkable.

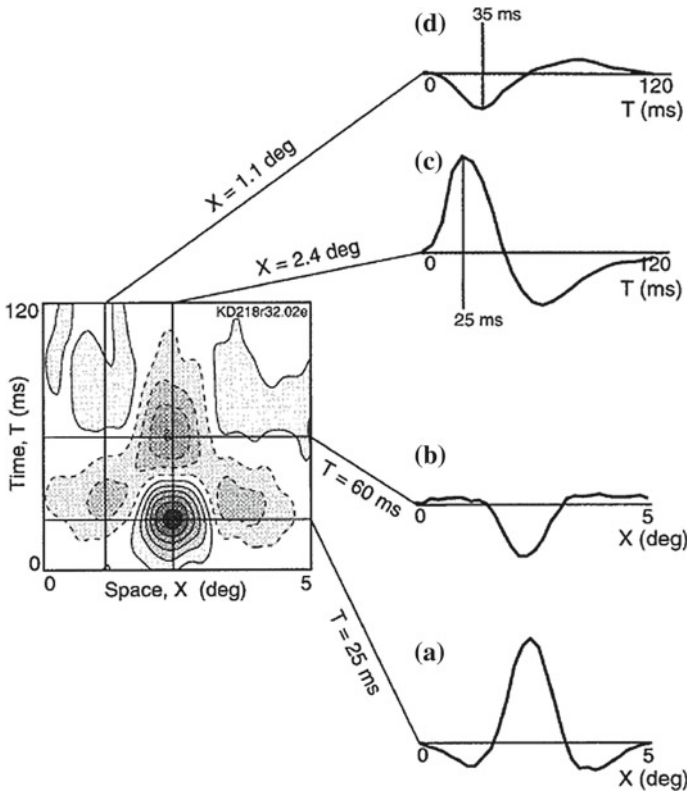


Fig. 3.28 Method for calculating a spatiotemporal receptive profile in the case of an LGN cell. Four cross sections of the RP are represented relative to the space dimension X and the time dimension T , two of these at constant T and two at constant X . From Cai et al. [46]




Fig. 3.29 Receptive profile of the form $\partial^3 G / \partial x^2 \partial t$ on the (x, t) plane, with y varying as a Gaussian




Fig. 3.30 Empirical level curves of a receptive profile of the form $\partial^3 G / \partial x^2 \partial t$ on the (x, t) plane. The coordinate x on the *horizontal axis* is spatial, and the *vertical axis* represents time, from 0 to 200 ms (time delay). From DeAngelis et al. [35]

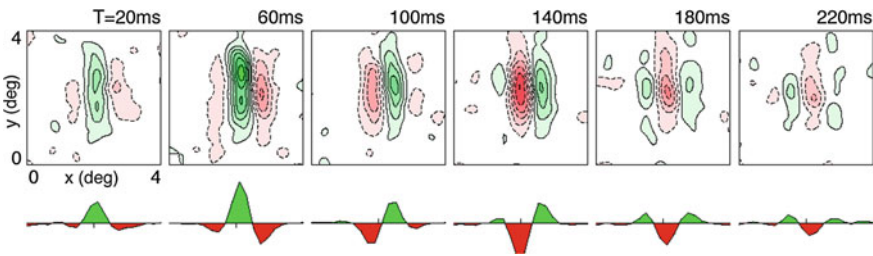


Fig. 3.31 Time evolution of a non-separable receptive profile $\varphi(x, y, t)$ for $t = 20\text{--}220$ ms. The spatial organization of the receptive profile changes, something that does not happen for separable receptive profiles. *Top* Spatial profile $\varphi_t(x, y) = \varphi(x, y, t)$. *Bottom* Cross section for $y = 0$. From DeAngelis et al. [35]

In humans, RPs cannot be measured using such invasive electrophysiological methods. However, other techniques have been developed since the end of the 1990s. They combine millimetre resolution functional magnetic resonance imaging (fMRI, introduced in 1990 by Seiji Ogawa, see Wandell and Winaver [48]) with the properties of the RFs on a scale of μm . For example, one can put checkerboard stimuli, such as moving bars, expanding rings, rotating wedges, at different points of the visual field and then measure the response. This is exemplified in Fig. 3.33, from Dumoulin and Wandell [49].

As the neurons in a voxel (a 3D pixel) respond to several positions, we are in fact studying the corresponding *population* pRF of RFs.⁶ Visual maps and the sizes of the pRFs are derived by integrating the data from different stimuli. This inference

⁶One voxel is typically $2.5 \times 2.5 \times 3 \text{ mm}^3$ in fMRI.

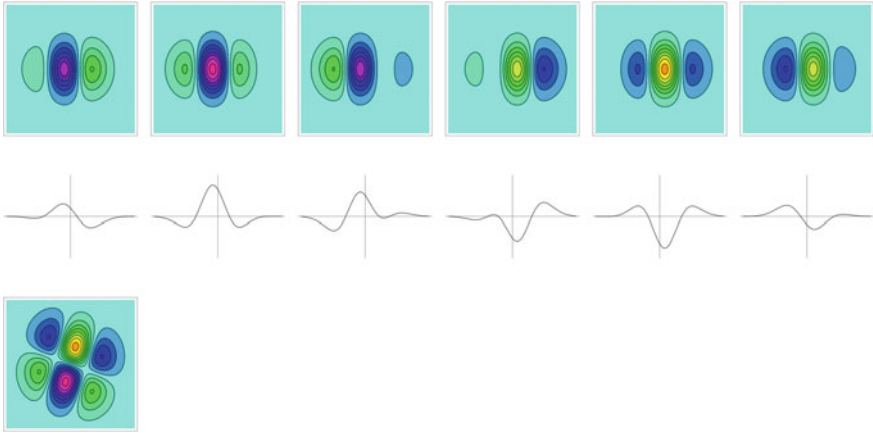


Fig. 3.32 Gaussian derivative model for the temporal evolution of the receptive profile $\varphi(x, y, t)$ in Fig. 3.31. Here we have taken $\varphi(x, y, t) = \partial^3 G(u, y, v) / \partial u^2 \partial v$ with the rotation $(u, v) = r_\theta(x, t)$ and $\theta = \pi/10$. *Upper* Evolution of the spatial profile $\varphi_r(x, y) = \varphi(x, y, t)$. *Centre* Cross sections for $y = 0$. *Lower* Receptive profile $\varphi(x, 0, t)$ on the plane (x, t) for $y = 0$




Fig. 3.33 **a–d** Checkerboard stimuli used in an experiment to determine RFs in humans by fMRI: rotating wedges, expanding rings, mean luminance periods, and moving bars. **e** Rotation cycles of the wedge in **(a)**: 6 cycles of 24 s per run and 8–16 runs per session. **f** The rotating wedge can be replaced by grey mean luminance blocks as in **(c)**. **g** Motion of the bar in **(d)**, with 4 orientations and 2 opposite directions per run. From [49]

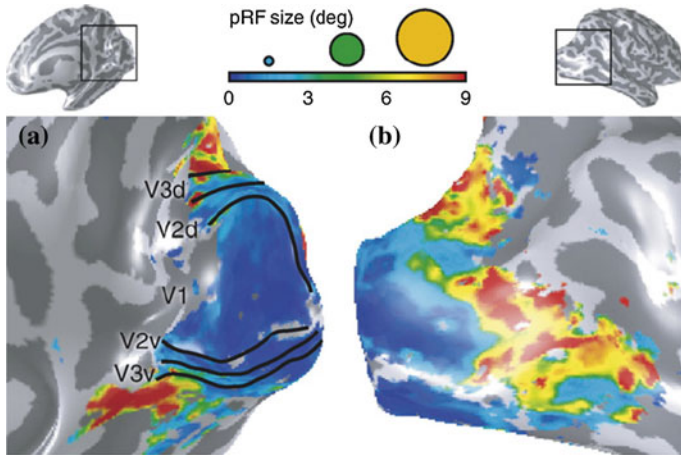


Fig. 3.34 Two views of a human visual cortex. **a** Medial. **b** Lateral. Colours code the sizes of the local RF populations. The size increases with the eccentricity and the transition from V1 through V2 to V3. From [49]

amounts to solving the inverse of the following direct problem: starting with the individual RFs, we consider a Gaussian population of such RFs centred at a point (x_0, y_0) with given width σ and calculate its response $r(t)$; we convolute $r(t)$ with a hemodynamic response function $h(t)$; we obtain a blood oxygenation level dependent (BOLD) signal $p(t)$ which we assume to be linearly related to the response $y(t)$ of the fMRI; we fit $p(t)$ to the data $y(t)$ and infer the sizes of the pRFs.

We then check that this non-invasive method agrees well with electrophysiological measurements on monkeys, and if so, we can draw conclusions for humans. Figure 3.34, also taken from [49], shows an estimate of the size of the pRFs. In V1–V3, it is $0.5\text{--}2^\circ$, and it increases with the eccentricity and the transition from V1 through V2 to V3.

3.3 Visual Neurons as Filters

3.3.1 Gabor Wavelets and Derivatives of Gaussians

There are many discussions about the exact form of the RPs. It was Daugman [50], then Judson Jones and Palmer [51] who first showed on the basis of detailed and accurate empirical data that receptive fields and profiles of visual neurons, e.g. the striated cortex, area 17 of the cat, could be approximated by Gabor filters, i.e. by trigonometric functions modulated by a Gaussian:

The Gabor function provides a useful and reasonably accurate description of most spatial aspects of simple receptive fields. [51, p. 1233]

The motivation for treating RPs as Gabor patches is twofold. To begin with, the ratio of the scale (the width of the Gaussian) to the spatial frequency, and also the phase factor, can be varied continuously. We then obtain wavelets associated with windowed Fourier transforms which are well suited to *harmonic analysis* on Lie groups (like the Heisenberg group), and these come into the models of functional architecture we shall be dealing with in later chapters.

But it is also illuminating to consider Gaussian derivative models. We still obtain the scale factor and the phase factor, but the spatial frequency is now determined by the order of differentiation and is thus highly constrained (see Fig. 3.35). In any case, these two classes of functions have good properties when we seek minimal solutions to the uncertainty principle. For a quantitative comparison between them, the reader may consult, for example, the paper by Bloom and Reed [52].

In 1D, the derivatives $d^n(e^{-x^2/2\sigma^2})/dx^n$ are given by the formula


$$\frac{d^n}{dx^n}(e^{-x^2/2\sigma^2}) = H_{n,\sigma}(x)e^{-x^2/2\sigma^2}, \quad H_{n,\sigma}(x) = \left(\frac{1}{\sigma\sqrt{2}}\right)^n H_n(x/\sigma\sqrt{2}),$$

where H_n is the n th Hermite polynomial.⁷ For $2\sigma^2 = 1$, we thus have

$$\frac{d^n}{dx^n}(e^{-x^2}) = H_n(x)e^{-x^2}.$$

The $H_n(x)$ are given by recurrence $H_0(x) = 1$, $H_1(x) = -2x$, and

Fig. 3.35 Successive derivatives of a Gaussian in 1D, showing how the spatial frequency is determined by the order of differentiation



⁷Readers wishing to experiment with derivatives of Gaussians are referred to the *Gaussian Derivative* package of *Mathematica*, by Bart M. ter Haar Romeny and Markus van Almstik. These techniques in which derivatives of Gaussians are used as filters have applications in image compression (JPEG). See, for example, the reference books by Mallat [53] or Morgan et al. [54].

$$H_n(x) = 2xH_{n-1}(x) - 2(n-1)H_{n-2}(x) .$$

In particular, $H_2(x) = -2(2x^2 + 1)$ and $H_3(x) = 4x(-2x^2 + 1)$. The $H_n(x)$ are polynomials of order n , even for even n and odd for odd n . The functions $H_n(x)e^{-x^2/2}$ are orthogonal⁸ in $L^2(\mathbb{R})$ with norm squared $2^n n! \sqrt{\pi}$. This gives the orthonormal system of *Hermite functions*

$$\varphi_n(x) = \frac{1}{\sqrt{2^n n! \sqrt{\pi}}} H_n(x) e^{-x^2/2} .$$

The $\varphi_n(x)$ satisfy the differential equation $\varphi_n''(x) + (2n + 1 - x^2)\varphi_n(x) = 0$ and are eigenfunctions of the Fourier transform with eigenvalue $(-i)^n$. Indeed, the generating function of the $H_n(x)$ is $\exp(2xt - t^2)$, since

$$e^{2xt-t^2} = \sum_{n=0}^{n=\infty} H_n(x) \frac{t^n}{n!} .$$

Multiplying by $e^{-x^2/2}$ and taking the Fourier transform F of the left-hand side, we obtain

$$\begin{aligned} e^{(2xt-t^2-x^2/2)} &= \sum_{n=0}^{n=\infty} H_n(x) e^{-x^2/2} \frac{t^n}{n!} , \\ F\left(e^{(2xt-t^2-x^2/2)}\right) &= \sum_{n=0}^{n=\infty} H_n(k) e^{-k^2/2} \frac{(-it)^n}{n!} = \sum_{n=0}^{n=\infty} F\left(H_n(x)e^{-x^2/2}\right) \frac{t^n}{n!} , \end{aligned}$$

and hence,

$$F\left(H_n(x)e^{-x^2/2}\right) = (-i)^n H_n(k) e^{-k^2/2} .$$

It was without doubt Richard Young who best analyzed the theoretical interest and empirical superiority of these models. The reader may refer to Young [55, 56], where, following Jan Koenderink, the author stresses that:

The initial stage of processing of receptive fields in the visual cortex approximates a ‘derivative analyzer’ that is capable of estimating the local spatial and temporal directional derivatives of the intensity profile in the visual environment.

We shall return to this in Sect. 3.3.4.

⁸For the definition of $L^2(\mathbb{R})$ see below Sect. 3.4. $H_n(x)e^{-x^2/2}$ can be in $L^2(\mathbb{R})$ because $e^{-x^2/2}$ decreases more quickly at infinity than any polynomial can increase there.

3.3.2 Steerable Filters

One advantage of derivatives of Gaussians is that they can provide *steerable filters*. Let us explain what this means. Suppose F is a filter with a preferred orientation and F^θ the filter derived from F by rotation through an angle θ . Then, F is said to be steerable if all the F^θ can be obtained by linear combinations, i.e. by interpolation, from a small number of them, viz., F^{θ_j} , $j = 1, \dots, S$:

$$F^\theta = \sum_{j=1}^{j=S} k_j(\theta) F^{\theta_j} .$$

This is a big advantage from the computational point of view since convolution is a linear operation and the filtering of a signal I by F^θ then amounts essentially to the filtering of I by the F^{θ_j} , since the interpolation coefficients $k_j(\theta)$ are independent of I .

For example, if $G = e^{-(x^2+y^2)}$ is a Gaussian, its first derivative G_1^θ in the direction θ is steerable. Indeed,

$$\frac{\partial G}{\partial x} = -2xe^{-(x^2+y^2)} = G_1^0, \quad \frac{\partial G}{\partial y} = -2ye^{-(x^2+y^2)} = G_1^{\pi/2},$$

and

$$G_1^\theta = \cos(\theta) G_1^0 + \sin(\theta) G_1^{\pi/2} .$$

Hence, if I is a signal,

$$G_1^\theta * I = \cos(\theta) G_1^0 * I + \sin(\theta) G_1^{\pi/2} * I .$$

Steerable functions are remarkable, and there are many theorems characterizing them. The interested reader may refer, for example, to the classic paper [57] by William Freeman and Edward Adelson. For example, if $F(x, y)$ written in polar coordinates $F(r, \varphi)$ has a Fourier expansion of the form

$$F(r, \varphi) = \sum_{n=-N}^{n=N} a_n(r) e^{in\varphi} ,$$

then F is steerable of the form $F^\theta(x, y) = \sum_{j=1}^{j=S} k_j(\theta) F^{\theta_j}(x, y)$ if and only if we have

$$\begin{pmatrix} 1 & \cdots & 1 \\ e^{i\theta_1} & \cdots & e^{i\theta_s} \\ \vdots & \ddots & \vdots \\ e^{iN\theta_1} & \cdots & e^{iN\theta_s} \end{pmatrix} \begin{pmatrix} k_1(\theta) \\ \vdots \\ k_s(\theta) \end{pmatrix} = \begin{pmatrix} 1 \\ e^{i\theta} \\ \vdots \\ e^{iN\theta} \end{pmatrix} ,$$

only the rows of the matrix corresponding to $a_n \neq 0$ being taken into account.

For example, if we return to the oriented first derivative G_1 of G , we have

$$G_1(r, \varphi) = -2re^{-r^2} \cos(\varphi) = -re^{-r^2} (e^{i\varphi} + e^{-i\varphi}) ,$$

whence $N = 1$, $a_0 = 0$, only a_{-1} and a_1 are nonzero, and $S = 2$. We must therefore solve the equation

$$(e^{i\theta_1} \ e^{i\theta_2}) \begin{pmatrix} k_1(\theta) \\ k_2(\theta) \end{pmatrix} = e^{i\theta} ,$$

$$k_1(\theta) e^{i\theta_1} + k_2(\theta) e^{i\theta_2} = e^{i\theta} = \cos(\theta) + i \sin(\theta) ,$$

giving $\theta_1 = 0$, $e^{i\theta_1} = 1$, and $k_1(\theta) = \cos(\theta)$, $\theta_2 = \pi/2$, $e^{i\theta_2} = i$, $k_2(\theta) = \sin(\theta)$.

3.3.3 Linearity Versus Nonlinearity

Visual neurons act on the optical signal like largely *linear* filters. This is remarkable since the connections and the spike inputs are highly nonlinear. The great specialist Robert Shapley wrote an interesting synopsis of this issue in 2009, in which he explains that the visual system reconstructs ‘a linearly filtered version of the visual world’ [58, p. 908]. He begins by observing that, under normal conditions, the vertebrate retinas are highly linear, from the photoreceptors to the GCs, because the synapses act linearly near the average level of the photon flux. For the rods, this average level is less than 1 photon, whence these actually act as photon counters. For the cones, the average level is several photons. If we apply a stimulus with illumination $E(t)$ given by an average illumination $E_0(t)$ modulated by a sinusoidal term with temporal frequency f (in cycles/s, i.e. in Hz) and wave number $\omega = 2\pi f$, so that $E(t) = E_0(t) + E_1(t) \cos(\omega t)$, a linear response will have the form

$$R(t) = L_0 E_0(t) + L_1(\omega) E_1(t) \cos(\omega t + \theta(\omega)) ,$$

and will thus contain only the first (fundamental) harmonic. A second-order response will have the form

$$R(t) = N_0 E_0(t)^2 + 2N_0 N_1(\omega) E_0(t) E_1(t) \cos(\omega t + \theta(\omega)) \\ + N_1(\omega)^2 E_1(t)^2 \cos(2(\omega t + \theta(\omega))) ,$$

and will thus exhibit a second harmonic with wave number 2ω , and so on. Experience shows that these higher harmonics are very weak. The photoreceptors thus respond in an essentially linear way.

For GCs, for example in cats, the animal is shown a grating with sinusoidally varying contrast, i.e. a signal of the form




Fig. 3.36 Response of an *X OFF* GC in a cat to a spatially and temporally modulated grating. The fundamental harmonic is significant, while the second is practically non-existent. Clearly, since the cell is *OFF*, the amplitude of the fundamental is maximal (minimal) when a dark (light) band occurs at the centre of the receptive field. From [58]


$$I(x, t) = I_0 + I_1 \sin(\varphi x + \xi) \cos(\omega t) ,$$

where $\varphi = 2\pi k$ is the wave number corresponding to the spatial frequency k of the grating (in cycles per degree) and ξ is the spatial phase. The response (spike train) is then measured to determine its modulation and the amplitude of its harmonics. Figure 3.36 shows that the *X* cells are linear, because the amplitude of the first harmonic is significant, while the second harmonic is practically non-existent. We already pointed out this linearity in Sect. 3.1. However, as can be seen from Fig. 3.37, the *Y* cells are nonlinear, with a significant second harmonic. We also pointed out this nonlinearity in Sect. 3.1.

Regarding the *V1* neurons, we observe that, for the same stimuli, the responses of simple cells are linear, with a weak second harmonic, and vary with the spatial phase, whereas those of complex cells have a dominant second harmonic that is insensitive to the spatial phase. To explain the linearity of the simple cells, Robert Shapley makes the hypothesis that it is caused by ‘the cancellation of nonlinear LGN excitation by cortico-cortical inhibition’ [58, p. 916], while the second harmonic of the LGN cells is eliminated by the second harmonic of the short-range inhibitory local cortico-cortical connexions. He concludes:

The apparent linearity of response time course and spatial summation are the result of a balance between nonlinear excitation and nonlinear inhibition. [58, p. 918]

Fig. 3.37 Response of a *Y OFF* GC in a cat to a spatially and temporally modulated grating. The fundamental harmonic not only is significant, but so also is the second, which is even constant, i.e., independent of the spatial phase. From [58]



3.3.4 Visual Neurons as Convolution Operators

In the context of the linear approximation, which we have just seen to be justified for the parvocellular GCs and simple cells in *V1*, the great advantage of receptive profile models based on derivatives of Gaussians is that one can do *multiscale differential geometry* on the optical signals, despite the fact that the data is noisy and highly non-differentiable. Indeed, let $I(x, y)$ be the intensity of the optical signal defined on the domain R of the retina, and $\varphi(x, y)$ the RP, centred on 0 (the centre of R), of an RF of a certain type of visual neuron. If the RF is centred on (x_0, y_0) , its RP is $\varphi(x - x_0, y - y_0)$ and, since the visual neuron acts to a first approximation as a linear filter on the optical signal, its response is the average of the signal weighted by φ , i.e. the integral

$$I_\varphi(x_0, y_0) = \int_D I(x', y')\varphi(x' - x_0, y' - y_0)dx'dy' .$$

This is due to the fact that the RP $\varphi(x - x_0, y - y_0)$ is by definition the impulse response to a Dirac delta $\delta(x - x_0, y - y_0)$ at (x_0, y_0) , that any signal $I(x, y)$ is a superposition (integral) of Dirac deltas, and that linearity implies that the response is the superposition of the responses. $I_\varphi(x_0, y_0)$ is the *measure* of the signal I at (x_0, y_0) .

For example, for a receptive profile going as the second derivative of a Gaussian, as in Fig. 3.23, we can consider the measure it provides of a stimulus in the form of a thin bar rotating and translating as in Fig. 3.38. Figure 3.39 shows that, during a rotation, the response has a maximum when the bar is aligned with the preferred orientation of the receptive field, then decreases and almost vanishes when the bar is perpendicular to this orientation. Fig. 3.40 for its part shows that, if the bar moves parallel to this preferred orientation, the response is maximal positive when the bar

is centred on the *ON* lobe and minimal negative when it is centred on the *OFF* lobes. All this is quite intuitive.

If a *field* of RFs with identical profiles now covers the retina *R* and measures the optical signal in parallel, then, in the continuous limit where there exists an RF for each position (x, y) in *R*, the response of this global field is the *convolution* of *I* with φ , viz.,

$$I_\varphi(x, y) = \int_D I(x', y')\varphi(x' - x, y' - y)dx'dy' = (I * \varphi)(x, y) .$$

As suggested by Luc Florack [59], a disciple of Jan Koenderink, a good way to see things is to treat the signal *I* (a noisy, hence ‘bad’ function) as a *distribution* in the sense of Laurent Schwartz, i.e. a continuous linear functional $\langle I | \varphi \rangle$ on a space of test functions (C^∞ functions with compact support, or fast decreasing in the case of tempered distributions), and treat the RPs $\varphi(x - x_0, y - y_0)$, which are well-localized regular functions, as classes of test functions *neurally wired* into the visual system. The measure of the signal *I* by a neuron with RP φ then delivers the representation $\langle I | \varphi \rangle$ in the sense of distributions.

Now, we know that the Dirac distribution δ is the basic operator for the differential calculus of distributions, since for any distribution *T*,

$$\delta * T = T , \quad \delta' * T = T' , \quad \delta^{(m)} * T = T^{(m)} ,$$

where $T^{(m)}$ is the *m*th derivative of *T*, and more generally,

$$D\delta * T = DT ,$$

for any differential operator *D* with constant coefficients. It thus suffices to know what δ becomes from a multiscale point of view in order to give meaning to the




Fig. 3.38 Thin bar rotating and translating in the receptive field of a simple V1 neuron




Fig. 3.39 During a rotation, the response of the simple V1 neuron has a maximum when the bar stimulus is aligned with the preferred orientation of the receptive field. It decreases and becomes negligible when the bar is perpendicular to this orientation




Fig. 3.40 When the bar stimulus moves parallel to the preferred orientation of the simple V1 neuron, the response is maximal positive when it is centred on the *ON* lobe and minimal negative when it is centred on the *OFF* lobes

concept of *scale space* (see Florack et al. [60]). In standard scale-space theories, we take the Gaussians

$$G_\sigma = G(x, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right).$$

This Gaussian kernel is the multiscale version of the identity operator. It tells us what happens to a point on the scale specified by the width σ of G_σ . As the convolution product $G_\sigma * G_\tau = G_{\sqrt{\sigma^2 + \tau^2}}$, the scale composition rule \dagger is $\sigma \dagger \tau = \sqrt{\sigma^2 + \tau^2}$. Put another way, it is the parameter σ^2 which is additive.

However, since the Gaussian is the heat kernel, it is natural to consider that the multiscale point of view consists in treating the signal I as the initial condition for a solution of the *heat equation*:

$$\left(\frac{\partial}{\partial s} \Delta \right) I = 0 \quad (\text{with } 2s = \sigma^2) .$$

This diffusion equation relates ‘pure geometry’ to its ‘physical’ counterpart (the multiscale aspect). It expresses the operational constraint of transformation of the signal *into a geometrical observable*. It replaces the infinitesimal level by a local multiscale level, with classical differential geometry corresponding to the ideal case of a zero scale (infinite resolution).

Given that, if D is a differential operator of order p , we have by definition

$$\langle DI | \varphi \rangle = (-1)^p \langle I | D\varphi \rangle ,$$

we see that the main effect of RPs given by partial derivatives of Gaussians is to *differentiate the signal in the sense of distributions*. More precisely, since a Gaussian G specifies a scale by its width, an RP of the form DG serves to apply the differential operator D at a certain scale. It is thus multiscale differential geometry that gets implemented neurally.

This is a really crucial point. If perception is geometrically structured, it is because, one way or another, the visual system is able to measure derivatives. However, noise in the signal induces random high-frequency components which render the idea of differentiation inapplicable. Therefore, the system must first eliminate these high frequencies by means of a low-pass filter, e.g. smoothing the signal I by convolving it with a Gaussian G , in order to be able to calculate $D(I * G)$. But since $D(I * G) = I * (DG)$, this is effectively what is done by fields of cells with RP DG . As observed by Hale [61]:

Convolution of a signal with a Gaussian derivative is equivalent to differentiating that signal before or after low-pass filtering. This combination of differentiation with Gaussian low-pass filtering is an efficient way to perform both operations simultaneously. It computes the derivative of the lower-frequency signal while attenuating higher-frequency noise.

At the end of the 1970s, the great vision specialist Marr [62] was the first to understand the functionality of the RFs produced by biological evolution. As already noted, neurophysiologists have long known empirically that retinal ganglion cells detect spatial contrasts (see Buser and Imbert [9]). But it was above all Marr (although we should also mention Ellen Hildreth) who introduced the idea that this was due to the fact that they carry out a very special kind of convolution, able to extract the qualitative discontinuities encoded in the signal, discontinuities that he called *zero-crossings*. Marr also asserted that the higher levels of visual processing were rooted in this first level of morphological organization of the retinal image, which he called the *2D primal sketch*. In fact, as we shall see in Sect. 3.4, he had discovered that the convolution of the signal with RPs of the form ΔG was effectively a *wavelet analysis*



Fig. 3.41 Gaussian blurring or smoothing $G * I$ of an image I by Gaussians of increasing width

of the signal, i.e. a spatially localized multiscale Fourier analysis able to extract the discontinuities.

Let us spell this out. If the RP $\varphi(x, y)$ is a derivative DG of a Gaussian, the classic formula $I * DG = D(I * G)$ provides a functional interpretation of the filtering $I * DG$: it is equivalent to $D(I * G)$, i.e. to application of the differential operator D to the signal I , smoothed (regularized) on the scale specified by G . Naturally, for scale 0, we recover the standard derivative DI , since $I * D\delta = D(I * \delta) = DI$.

For example, the ganglion cells in the retina or those in the lateral geniculate nucleus whose RFs are the Laplacian ΔG of a Gaussian calculate the Laplacian of the regularized signal, i.e. $\Delta(G * I) = \Delta G * I$. The convolution $G * I$ of I with

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-r^2/2\sigma^2},$$

where $r^2 = x^2 + y^2$, corresponds to the smoothing of I by *Gaussian blurring* (see Fig. 3.41).

The Laplacian $\Delta(G * I)$ of this convolution extracts the edges of $G * I$ (see Fig. 3.42). But since $\Delta(G * I) = \Delta G * I$, with

$$\Delta G(x, y) = -\frac{1}{\pi\sigma^4} \left(1 - \frac{r^2}{2\sigma^2}\right) e^{-r^2/2\sigma^2},$$

edge detection is equivalent⁹ to convolution of the image I with receptive profiles of the form ΔG . We stress the fact that this is an extraction of local edges, i.e. ‘pointlike’ on the scale of G , while our visual system ‘sees’ global edges in Fig. 3.42, because it automatically integrates these local edges. The local \rightarrow global integration is not carried out in the individual neurons of the retina, the LGN, or V1, but rather, as we shall see in Chaps. 4 and 5, through the functional architecture of V1.

⁹This algorithm can be found today in any image-processing software.




Fig. 3.42 Extraction of edges from an image by convolution with receptive profiles in the form of the Laplacian of a Gaussian




Fig. 3.43 Let $f(x)$ be a differentiable function on \mathbb{R} with a ‘discontinuity’ at x_0 , i.e. a sudden change in value. At x_0 , the first derivative $f'(x)$ has a peak (a Dirac distribution δ , if x_0 is a genuine discontinuity) and the second derivative $f''(x)$ has two peaks, one positive and the other negative, with a zero-crossing between them

The zero-crossing criterion says that, in this context and in one dimension for simplicity, a discontinuity corresponds to $\Delta G * I$ crossing zero, a crossing which occurs between two sharp peaks, one positive and one negative. But this is nothing other than the multiscale version of the well-known dipole structure of the second derivative δ'' of the Dirac distribution δ . Figure 3.43 illustrates the idea.

3.3.5 Fine Orientation Discrimination

Having said this, since the linear filtering just described is based on RFs with a certain width, it would not appear to explain the remarkably fine discrimination of orientations, much more precise than this width. Beaudot and Mullen [63] have put forward an explanation. Their hypothesis is that the fine discrimination in the

response of the visual system, based on population coding, should not be confused with that of the individual RFs, and in fact corresponds to the *maximal gradient* of the latter:

Despite their broad orientation tuning, single neurons in the primary visual cortex can reliably signal orientation differences of about 1 degree and it is the slopes of their tuning curves and response variability that determine the minimum orientation differences that can elicit a reliable response change. [63, p. 27]

3.4 Vision and Wavelets

3.4.1 Fourier, Gabor, and Wavelets

As indicated, the mathematical interpretation of the kind of algorithms proposed by Marr is now provided by the wavelet algorithm. Let us spell this out (in one dimension) according to the works of Meyer [64] and Mallat [65]. We begin with the conventional Fourier transform. Let $L^2(\mathbb{R})$ be the Hilbert space of square-integrable functions on \mathbb{R} , i.e. finite energy signals on \mathbb{R} , equipped with the scalar product¹⁰

$$\langle f | g \rangle = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(x) \bar{g}(x) dx .$$

Harmonic analysis provides a decomposition of each element in $f \in L^2(\mathbb{R})$ in terms of an orthonormal basis of trigonometric functions¹¹ $e^{i\omega x}$. The decomposition of $f(x)$ is the Fourier transform (FT) given by

$$\widehat{f}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(x) e^{-i\omega x} dx = \langle f(x) | e^{i\omega x} \rangle .$$

It can be shown that $\widehat{\widehat{f}} = f$, i.e., that the frequency analysis can give back f by synthesis:

$$f(x) = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} \widehat{f}(\omega) e^{i\omega x} d\omega = \langle \widehat{f}(\omega) | e^{-i\omega x} \rangle ,$$

and that the norms $\|f(x)\|$ and $\|\widehat{f}(\omega)\|$ are equal. In other words, the FT is an isometry between the Hilbert spaces $L^2(\mathbb{R})_x$ and $L^2(\mathbb{R})_\omega$, where the indices x and ω are the coordinates of the two spaces \mathbb{R} of positions and frequencies.

The FT is fundamental largely for the following reason. Consider a linear filter on $L^2(\mathbb{R})$, that is, a continuous linear operator T . The action of T on the

¹⁰The signals correspond to real-valued $f(x)$. However, it is convenient also to allow the $f(x)$ to be complex-valued. $\bar{f}(x)$ is then the complex conjugate of $f(x)$.

¹¹The $e^{i\omega x}$ are not actually in this Hilbert space, but we can ignore this well-known technical detail for our present purposes.

$f(x) \in L^2(\mathbb{R})$ (filtering) is calculated from its impulse response, that is, its action on the Dirac distribution $\delta(x)$: $\varphi(x) = T(\delta(x))$. Indeed, by definition of δ , $f(x) = \int_{\mathbb{R}} f(u)\delta(x-u)du$, i.e. $f = f * \delta = \delta * f$, whence δ is the identity for convolution, and by linearity and continuity

$$T(f(x)) = \int_{\mathbb{R}} f(u)T(\delta(x-u))du = \int_{\mathbb{R}} f(u)\varphi(x-u)du = f * \varphi.$$

We deduce immediately that the plane waves $e^{i\omega x}$ are eigenvectors of T . Indeed,

$$T(e^{i\omega x}) = \int_{\mathbb{R}} \varphi(u)e^{i\omega(x-u)}du = e^{i\omega x} \int_{\mathbb{R}} \varphi(u)e^{-i\omega u}du.$$

The eigenvalue of $e^{i\omega x}$ is therefore $\sqrt{2\pi}\widehat{\varphi}(\omega)$. It is thus straightforward to calculate the action of T . Since $\widehat{f * \varphi} = \widehat{f} \cdot \widehat{\varphi}$, T acts by multiplying the amplitudes $\widehat{f}(\omega)$ of the frequency components $e^{i\omega x}$ by $\sqrt{2\pi}\widehat{\varphi}(\omega)$.

The problem raised by the FT is that the spatial information it provides on cues relating to the position x is *delocalized* owing to the infinite range of the plane wave $e^{-i\omega x}$. In fact, it is determinate in the frequency space and indeterminate in position space (compare with the Heisenberg uncertainty relations in quantum mechanics). In order to obtain more localized information processing, the Anglo-Hungarian physicist Dennis Gabor (inventor of holography, Nobel prize in 1971) introduced the windowed Fourier transform (WFT) or Gabor transform (GT) in the 1940s, inspired by quantum mechanics:

$$Gf(\omega, u) = \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} f(x)g(x-u)e^{-i\omega x}dx = \langle f(x) | g_{\omega,u}(x) \rangle,$$

where g is a real spatial window, i.e. $\bar{g}(x) = g(x)$, which gets translated along the x axis. The $g_{\omega,u}(x) = e^{i\omega x}g(x-u)$ are frequency-space ‘atoms’, more localized in position than the FT but less in frequency, and $Gf(\omega, u)$ is a convolution product since

$$Gf(\omega, u) = \frac{1}{\sqrt{2\pi}} f * \tilde{g}_{\omega,u}, \quad \text{with } \tilde{g}_{\omega,u}(X) = g(X)e^{i\omega(X-u)}.$$


The GT is invertible and its inverse transformation is

$$f(x) = \frac{1}{2\pi} \int_{\mathbb{R}^2} Gf(\omega, u)g(x-u)e^{i\omega x}d\omega du.$$

It can be shown that $\|f\| = \|Gf\|/\sqrt{2\pi}$, which means that, up to a multiplicative constant, the GT is an isometry of the Hilbert spaces $L^2(\mathbb{R})_x$ and $L^2(\mathbb{R})_{\omega,u}$. The GT is generally *redundant*, so it is possible to restrict to *discrete* samples of u and ω .

The problem with the GT is that, even if it localizes the information, it will only do it at a *single* level of resolution. If σ_u is the standard deviation of the window

Fig. 3.44 Change of scale of a Laplacian of a Gaussian wavelet



$g(x)$, i.e.

$$\sigma_u^2 = \int_{\mathbb{R}} x^2 |g(x)|^2 dx ,$$

and if σ_ω is the standard deviation of its FT $\widehat{g}(\omega)$, i.e.

$$\sigma_\omega^2 = \int_{\mathbb{R}} \omega^2 |\widehat{g}(\omega)|^2 d\omega ,$$

then the signals f and \widehat{f} are both determined as well as possible with respect to the resolution cells

$$[u_0 - \sigma_u, u_0 + \sigma_u] \times [\omega_0 - \sigma_\omega, \omega_0 + \sigma_\omega] ,$$

but inside these cells, spatial information remains delocalized and diffuse. It is thus impossible to localize edges on a shorter scale than σ_u . In particular, if the signal is multiscale, e.g. fractal, it cannot be correctly analyzed.

This is why we need an FT that is not only localized but also, as for the retina, *multi-scale*, and this is what is achieved by *wavelets*. The problem is to find a decomposition of $L^2(\mathbb{R})$ using only a real function $\psi(x)$, i.e. $\overline{\psi(x)} = \psi(x)$, the *mother wavelet*, and its translates $\psi(x - u)$, and changes of scale $\psi_s(x) = \sqrt{s}\psi(sx)$ or $\psi_s(x) = \psi(x/s)/\sqrt{s}$. Yves Meyer, Ingrid Daubechies, and Stéphane Mallat showed that there exist functions ψ such that, for $j, k \in \mathbb{Z}$, the $\psi_{j,k}(x) = 2^{-j/2}\psi(2^{-j}(x - k))$ yield an orthonormal basis of $L^2(\mathbb{R})$. Figure 3.44 shows the example of a Laplacian of a Gaussian, which does indeed provide a system generating $L^2(\mathbb{R})$, although it exhibits redundancy and is not therefore an orthonormal basis. A typical example of a wavelet is precisely Marr’s wavelet ΔG .

We then obtain a wavelet transform (WT):

$$Wf(s, u) = \int_{\mathbb{R}} f(x)\psi_s(x - u)dx = \langle f(x) | \psi_{s,u}(x) \rangle .$$




Fig. 3.45 Laplacian of a Gaussian $\psi(x) = (1 - x^2)e^{-x^2/2}$ in one dimension and its Fourier transform $\widehat{\psi}(\omega) = \omega^2e^{-\omega^2/2}$. The condition (C) is satisfied because $\widehat{\psi}(0) = 0$ and $C_\psi = \int_{\mathbb{R}^+} \frac{|\widehat{\psi}(\omega)|^2}{\omega} d\omega$ is equal to $\int_{\mathbb{R}^+} \omega^3 e^{-\omega^2} d\omega = 1/2 < \infty$

It is well defined if an admissibility condition (C) is satisfied on the FT $\widehat{\psi}(\omega)$ of the mother wavelet $\psi(x)$

$$(C) : \widehat{\psi}(0) = 0, \quad C_\psi = \int_{\mathbb{R}^+} \frac{|\widehat{\psi}(\omega)|^2}{\omega} d\omega < \infty .$$

This expresses the fact that the FT of ψ is not only zero but also sufficiently slow-growing at the origin. This is the case, for example, for the Laplacian of a Gaussian Marr wavelet ΔG (see Fig. 3.45). For the WT, the resolution cells are therefore

$$\left[u_0 - \frac{\sigma_u}{s}, u_0 + \frac{\sigma_u}{s} \right] \times [\omega_0 - s\sigma_\omega, \omega_0 + s\sigma_\omega] .$$

A theorem due to Jean Morlet and Alex Grossmann says that, up to a constant multiple, W is an isometry of $L^2(\mathbb{R})_x$ on $L^2(\mathbb{R}^+ \times \mathbb{R})_{s,u}$. The inverse transform is given by the following formula, proven long ago by Calderon, well before wavelets had become relevant to signal and image processing:

$$f(x) = \frac{1}{C_\psi} \int_{\mathbb{R}^+} \int_{\mathbb{R}} Wf(s, u)\psi_s(x - u)ds du .$$

3.4.2 Wavelets and Group Representation

Here, we comment on something that will become essential in the second volume. The above wavelets are obtained by making transforms $\varphi_{s,u}(x) = \sqrt{s}\varphi(s(x - u))$ of a function $\varphi(x) \in L^2(\mathbb{R})$. Note to begin with that φ and $\varphi_{s,u}$ have the same norm in $L^2(\mathbb{R})$. Indeed, if we consider $\|\varphi_{s,u}\|^2 = \int_{\mathbb{R}} s |\varphi(s(x - u))|^2 dx$, we see

that by translating x by u , which will not change dx , then doing an enlargement (or homothety) $x' = sx$, we obtain $\int_{\mathbb{R}} |\varphi(x')|^2 dx'$, which is the square of the norm of φ .

Note then that, since $\varphi_{s,u}$ is obtained by acting on φ on the right by (s, u) , we define a linear operator on $L^2(\mathbb{R})$, and also that the composition of two such operators is given by the composition rule $(s, u) \cdot (s', u') = (ss', u + su')$. Now, this can be taken as the product in a group with identity $(1, 0)$ and inverse $(s, u)^{-1} = (1/s, -u/s)$, and this corresponds to the *affine group* $G = \mathbb{R}^+ \times \mathbb{R}$. Wavelets are thus associated with a unitary representation of the affine group G in the Hilbert space $L^2(\mathbb{R})$ which is given by

$$\begin{aligned} \rho : G &\rightarrow \text{Aut}(L^2(\mathbb{R})) \\ (s, u) &\mapsto \rho(s, u) : \varphi \mapsto \varphi_{s,u}. \end{aligned}$$

3.4.3 Wavelets and Discontinuities

One crucial point is that, as can be seen in Fig. 3.46, the amplitude of the WT is a privileged indicator of the *discontinuities* (singularities) encoded in the signal. These are only implicitly contained in the FT, but become *explicit*, and hence exploitable, in the WT. As stressed by Mallat [66, p. 9]:

The ability of the WT to characterize the type of local singularities is a major motivation for its application to detect the signal sharper variations.¹²

Figure 3.47 shows the quality of the reconstruction of the geometry of a signal by inverse WT. Only 10% of the information is conserved, and yet the reconstructed signal is almost indistinguishable from the original one.

3.4.4 Redundancy of Wavelets

There is generally a high level of redundancy in the WT (when there is none, we speak of orthogonal wavelets). This is expressed in terms of the *reproducing kernel*, viz.,

$$K(s, s', u, u') = \frac{1}{C_\psi} \int_{\mathbb{R}} \psi_s(x - u) \psi_{s'}(x - u') dx ,$$

¹²See also Mallat [53].




Fig. 3.46 Wavelet transform (WT) of a piecewise constant function using the Morlet wavelet $\psi(x) = (e^{2i\pi x} - e^{-k^2/2})e^{-2\pi^2 x^2/k^2}$. The *abscissa* is the space variable x , and the *ordinate* the scale k . The figure shows the modulus (*left*) and the phase (*right*) of the WT. Values are *colour* coded. We see that the coefficients essentially measure the discontinuities in the signal. Image by Marie Farge

by the formula

$$Wf(s', u') = \int_{\mathbb{R}^+} \int_{\mathbb{R}} Wf(s, u) K(s, s', u, u') ds du .$$

For the Marr wavelet $\psi = \Delta G$, the redundancy is expressed by the heat equation:

$$\frac{\partial W_s f}{\partial t} = \Delta W_s f \quad (t = s^2) .$$

Owing to the redundancy of the WT, one can discretize by sampling the variables u and ω . For example, by discretizing the scale s , we obtain the discrete WT

$$W_{2^j} f(x) = \langle f(x) | \psi_{2^j, u}(x) \rangle .$$

With the condition

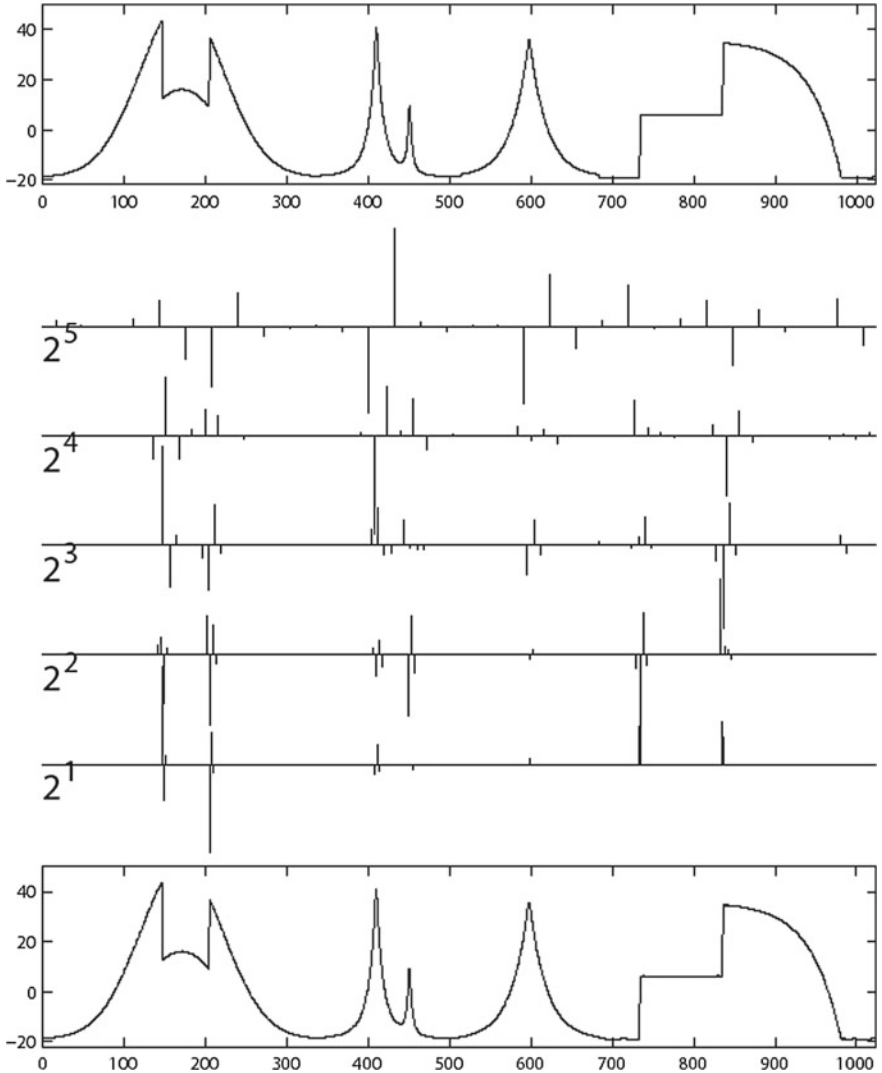


Fig. 3.47 Quality of the reconstruction of the geometry of a signal by the inverse WT. *Top* Original signal. *Centre* Coefficients of the dyadic WT for a Daubechies wavelet (a sort of Laplacian of a Gaussian without redundancy, giving an orthogonal basis) for five scales from 2^1 to 2^5 . *Bottom* Signal reconstructed by inverse WT from only 10% of the information. From Mallat and Peyré [67]

$$\sum_{j \in \mathbb{Z}} |\widehat{\psi}(2^j \omega)|^2 = 1, \text{ for all } \omega \neq 0,$$

we obtain

$$\|f\|^2 = \sum_{j \in \mathbb{Z}} \|W_{2^j} f\|^2.$$

In this case, we speak of dyadic wavelets.

In 2D, we introduce preferred orientations and use wavelets looking like the receptive profiles of simple V1 cells.

3.4.5 Compression and Geometry

Using transformation tools of this kind, it becomes possible to *compress* an image in an intrinsic way, that is, in conformity with its specific geometric structure. In fact, the image can be reconstructed from its multiscale edges (Marr's conjecture). The reconstruction can be extremely faithful because it is based on the morphological structure of the image. Only the finer details, such as textures, are smoothed out in this process.

More precisely, if we use the Laplacian ΔG of a Gaussian G (Marr) or the Laplacian $\Delta \vartheta$ of a regularizing function ϑ with compact support (Mallat), we can extract the qualitative discontinuities using the zero-crossing criterion of the last section. If we use the gradient of these functions, ∇G or $\nabla \vartheta$, as wavelet, we can extract qualitative discontinuities by selecting the maxima of the WT. Stéphane Mallat has given spectacular examples. Figure 3.48 shows the maxima of the modulus of the wavelet transform of a reference image (the face of the model Lena, the image-processing icon) on a certain scale. By selecting only a few scales and strongly thresholding the maxima so as to keep only the largest, thereby significantly compressing the image, the inverse transform nevertheless leads to excellent reconstructions (see Fig. 3.49).

The extraction of discontinuities is so important because those that are stable under large scale variations can be taken as *objective*, and in particular, as *edges* of external objects.

In short, regarding Marr's ideas and wavelet analysis, we can say that the compression of visual information, which is a constraint on information, becomes identified with a morphological analysis, i.e. a geometrical constraint. The morphological representation of images obtained in a bottom-up and data driven way by extraction of qualitative discontinuities provides a foundation for higher level symbolic representations. As stressed by Marr [62, p. 67]:

Zero-crossing provides a natural way of moving from an analogue or continuous representation like the two-dimensional image intensity values $I(x, y)$ to a discrete, symbolic representation.

3.4.6 Matching Pursuit and Rank Coding

In some of his work, Stéphane Mallat introduced maximal redundancy. He considered the *three-parameter* family of space–scale–frequency atoms:



Fig. 3.48 Wavelet analysis of an image (*top left*). *Top right* Modulus of the wavelet transform. *Bottom left* Maxima of the modulus. *Bottom right* Thresholded maxima. From Stéphane Mallat



Fig. 3.49 Reconstruction of the image in Fig. 3.48 by inverse wavelet transform. *Left* Original image. *Centre* Reconstruction from weakly thresholded maxima. *Right* Reconstruction from strongly thresholded maxima (compression by a factor of 32). The geometry of the image is perfectly preserved. Only the textures are slightly smoothed. From Stéphane Mallat

$$D = g_\gamma(x) = g_{s,u,\omega}(x) = \frac{1}{\sqrt{s}} g\left(\frac{x-u}{s}\right) e^{i\omega x},$$

where u = position, s = scale, and ω = frequency. He then iterated the adaptive algorithm known as *matching pursuit*, which involves seeking the g_γ that best approximates f :




Fig. 3.50 Detectors for points of inflection and ‘cross’ singularities in a checkerboard. *Left* Initial image with clear edges. *Centre* Image almost completely hidden by noise. *Right* Detector response (white = maximum response) for an appropriate smoothing scale. *Top* ‘Point of inflection’ singularity. *Bottom* ‘Cross’ singularity. From Florack [59]

$$\begin{cases} f = \langle f | g_{\gamma_0} \rangle g_{\gamma_0} + Rf, & \text{with } Rf \text{ orthogonal to } g_{\gamma_0}, \\ R^n f = \langle R^n f | g_{\gamma_n} \rangle g_{\gamma_n} + R^{n+1} f. \end{cases}$$

The method can be refined by introducing *wavelet packets*. Several wavelets are used in parallel to optimize the choice of decomposition basis to suit the intrinsic structure of the signal. The fitting criterion is then minimization of the information entropy.¹³

The matching pursuit algorithm is analogous to the one known as *rank coding*, which was introduced for neurophysiological purposes by Simon Thorpe (see Sect. 3.2.6). This further strengthens the parallel between the wavelet analysis of images and optical signal processing by visual neurons.

3.5 Feature Detectors

These methods can easily be used to construct detectors of local geometrical features that are more complicated than simple straight discontinuities. For example, a corner detector will pick up high curvatures in edges. The edges are detected by the norm of the gradient $|\nabla I|$ of the signal I . The curvature of an isophote (a level curve of I) is given by the divergence of the normalized gradient

$$\kappa = \text{div} \left(\frac{\nabla I}{|\nabla I|} \right) = \frac{1}{|\nabla I|} \left(\Delta I - \frac{H(\nabla I, \nabla I)}{|\nabla I|^2} \right),$$

¹³See, for example, Wickerhauser [68].

where H is the Hessian¹⁴ of I , i.e. the matrix of second derivatives of I with respect to x and y . Hence, a good corner detector is the invariant $\kappa |\nabla I|^3$. When used in a multiscale way, we obtain RPs that respond only to those places where the signal (smoothed by G) has a line of qualitative discontinuity which itself has a discontinuity in its tangent (Florack [59], Hamy [69]). Figure 3.50 by Luc Florack shows two examples of feature detectors: for points of inflection and for ‘cross’ singularities in a checkerboard pattern. We see that, even when the image is practically totally hidden in noise, the geometry remains detectable, provided that we use an appropriate smoothing scale.

3.6 Receptive Profiles and Information Theory

3.6.1 Signal Decorrelation and Efficient Coding

3.6.1.1 Models of Decorrelation

Several specialists, such as Joseph Atick and Jean-Pierre Nadal, have shown that the receptive profiles we have discussed can be recovered from extremely simple general hypotheses in the context of information theory. This can help us to understand how evolution might have selected them. They correspond to an information strategy and design principles that optimize the efficiency with which information is represented:

Efficiency of information representation in the nervous system potentially has evolutionary advantages. Atick [70, p. 213]

Efficient representations help us to understand the ‘visual vocabulary’ (geometric features) that can describe the environment in a compact way. Now, the statistics of *natural* images is rather particular and not at all Gaussian, because there are very strong correlations between the pixels, and because of the presence of edges. It is essential to take this into account because the different possible statistics of the inputs will affect the spike distribution. For example, Yves Frégnac and coworkers have studied four different statistics: drifting gratings, dense noise, natural images with eye movements, and gratings with eye movements. The variability of the spikes decreases with the complexity of these classes of stimuli, while the accuracy in their emission times increases.

The pixellated representation in the photoreceptors is thus fundamentally inefficient, because it does not take this into account. It must therefore be made more efficient. Back in the 1950s to 1960s, Fred Attneave and Horace Barlow suggested that, in order to be efficient, the neural coding must eliminate as far as possible the enormous redundancy in the inputs. Put another way, it must compress them. This is why, as explained by Field in [71], the statistical distributions of natural images must

¹⁴Named after Ludwig Otto Hesse.

play a determining role in the evolutionary explanation of the design of neural hardware. Ecological constraints, in the sense of James Gibson, and priors, in the sense of Thomas Bayes, are thereby imposed on neural information processes. Regarding the relations between neurons, the maximization of coding efficiency depends on the possibility of making the responses of the different neurons statistically *independent*, insofar as this can be done, in order that the joint probability of the responses become the product of the probabilities. However, this is a very hard problem to solve.

The guiding idea is thus to apply some principle for optimizing the information and suitably decorrelate the signal by eliminating higher and higher order correlations. The simplest example occurs when there is a linear filtering (convolution by a receptive profile) which compresses the signals $I(x)$ by first decorrelating the spatial autocorrelation $R(x_1, x_2) = \langle I(x_1)I(x_2) \rangle$.¹⁵ Due to homogeneity and isotropy, the autocorrelation only actually depends on $x = x_1 - x_2$, so $R(x_1, x_2)$ has the form $R(x_1, x_2) = R(x_1 - x_2) = R(x)$. The Fourier transform of $R(x)$ is the *power spectrum* of the signals:

$$\widehat{R}(\omega) = \int R(x)e^{-i\omega x} dx .$$

Now, an examination of natural image statistics shows that these have a power spectrum going as $1/|\omega|^2$ (Field's law). So let us consider the extreme case where $\widehat{R}(\omega) = 1/|\omega|^2$. It corresponds to the fact that the spatial autocorrelation is scale invariant, in the sense that $R(\alpha x) = \alpha R(x)$. Indeed, by inverse Fourier transform

$$R(x) = \int \frac{e^{i\omega x}}{|\omega|^2} d\omega ,$$

so by a change of variable $\omega = \lambda/\alpha$, we obtain

$$R(\alpha x) = \int \frac{e^{i\omega \alpha x}}{|\omega|^2} d\omega = \int \frac{\alpha^2 e^{i\lambda x}}{|\lambda|^2} \frac{d\lambda}{\alpha} = \alpha R(x) .$$

Then, let $\varphi(x)$ be the receptive profile of the given linear filters and let $T(I)(x) = (I * \varphi)(x)$ be the result when the signal I has been filtered. The decorrelation means that $\langle T(I)(x)T(I)(x') \rangle = \delta(x - x')$. Now, if we only take into account Field's law, we can assume that the images are generated by the Gaussian distribution with covariance matrix R . But then $T(I)$ is generated by the Gaussian with covariance matrix $T(R) = \varphi * R * \varphi'$, where $\varphi'(x) = \varphi(-x)$, and we would like to have $T(R) = \delta$. Fourier transformed, the equation becomes

$$\widehat{T(R)}(\omega) = \widehat{\varphi}(\omega)\widehat{R}(\omega)\widehat{\varphi'}(\omega) = 1 ,$$

whence $|\widehat{\varphi}(\omega)|^2 = 1/\widehat{R}(\omega)$, and since $\widehat{R}(\omega) = 1/|\omega|^2$, we therefore end up with $|\widehat{\varphi}(\omega)| = |\omega|$. Since $\langle T(I)(x)T(I)(x') \rangle = \delta(x - x')$ implies $|\widehat{T(I)}(\omega)|^2 = 1$ when

¹⁵ $\langle \rangle$ is the average over signals.

Fourier transformed, we see that the power spectrum of $T(I)$ is flat. The decorrelation thus leads to a ‘whitening’ of the signal, obtained by dividing each frequency component by the square root of its variance and taking the inverse Fourier transform.

While this method is well suited to spectra of the form $1/|\omega|^2$, it performs very poorly with noise, because it amplifies the high frequencies, precisely where it already dominates. The decorrelation must therefore be associated with a smoothing of the signal to suppress noise. To do this, we may take a filter satisfying

$$|\widehat{\varphi}(\omega)|^2 = \frac{\widehat{R}(\omega) + N^2}{\widehat{R}(\omega)^2},$$

where N^2 is the noise power. For $N \sim 0$, we recover the previous filter. However, for large N , we find $|\widehat{\varphi}(\omega)| \sim N$. To obtain $\varphi(x)$ from $|\widehat{\varphi}(\omega)|$ by inverse Fourier transform, we must fix the phase $\theta(\omega)$ of $\widehat{\varphi}(\omega)$. This is where it is useful to take the $\varphi(x)$ even, because we then have $\widehat{\varphi}(\omega) = \widehat{\varphi}(-\omega)$, which implies that $\theta(\omega) = 0$ or π .

This is how Atick showed that, if we wish to decorrelate the signal correctly, even when there is noise, we must use ganglion-cell-type receptive profiles of the form ΔG .

Jean-Pierre Nadal and coworkers [72] have gone further with this, not only by taking into account Field’s law, but also by explicitly considering multiscale *edges* in natural image statistics. They have shown that minimizing the redundancy of neural coding under the constraint of edge detection leads to oriented wavelets.

3.6.1.2 Empirical Decorrelation

It should be noted that decorrelation of neurons like retinal GCs can be tested experimentally with the methods of information theory. For example, the reader may consult the work by Nirenberg et al. [73] on the way GCs behave as independent encoders. The authors show that:

More than 90% of the information about the stimuli can be obtained from the cells when their correlated firing is ignored.

3.6.2 Receptive Profiles and Natural Images

More generally, we can take very large data bases of natural images and try to carry out an independent component analysis (ICA, a non-Gaussian version of the classic factor analysis into principal components, or PCA) of a kind that might be neurally implemented. We try to obtain components (the RFs of the neurons) which are both statistically independent and *sparse*, where sparsity means that, for a given image, most of the components have a very weak response, and only a very few

have any significant response. Sparse representations have many advantages. They provide a compromise between, at one extreme, representations in which almost all the neurons in the network would be involved in the processing of each stimulus and, at the other extreme, representations in which, for each stimulus, there would only be one neuron which responds selectively. They are made up of basic patterns (the RFs treated as ‘atoms’), these being adapted to the class of stimuli under consideration, natural images here. They can increase the storage capacity of the network associative memory without there being interference between the patterns of activity elicited by the different inputs; they give access to the structure of the stimuli because the components are the most relevant features; they allow the network to save energy; and they are easily learnt using Hebbian learning rules. It is useful to recall Hebb’s rule here: ‘neurons that fire together, wire together’. That is, when neurons are activated together by the same stimulus, their connections are strengthened.

The problem of finding optimal sparse representations for a class of stimuli is not an easy one. It can be formulated as follows. Let $\Sigma \subset \mathbb{R}^P$ be a class of stimuli I^k , $k = 1, \dots, S = \#(\Sigma)$ (where $\#E$ denotes the cardinality of the set E), these being images I made up of $P = p^2$ pixels and described by the vectors $I = (I_r)_{r=1, \dots, P}$ in the canonical basis of \mathbb{R}^P . Note that these are natural images or retinal inputs, and in the latter case, the pixels are ganglion cell RFs. We would like to find a ‘good dictionary’ Φ of atoms (of RFs) $\varphi_i \in \mathbb{R}^P$, $i = 1, \dots, N$, providing a ‘good’ decomposition of the $I \in \Sigma$ in the form

$$I = \sum_{i=1}^{i=N} s_i \varphi_i .$$


The representation of $I \in \Sigma \subset \mathbb{R}^P$ by $s = (s_i)_{i=1, \dots, N}$ is the code for I relative to the ‘dictionary’ Φ , which is in turn a $P \times N$ matrix whose columns are the $\varphi_i \in \mathbb{R}^P$. In the present case, i.e. the retina and V1, we have $N \gg P$, a situation referred to as ‘overcompleteness’, because the number of neurons in V1 is several hundred orders of magnitude greater than the number of ganglion cells. This means that the φ_i generate \mathbb{R}^P , but then are not at all linearly independent. It is just because of this that the codes s can be sparse. The space Σ is not at all a vector subspace of \mathbb{R}^P because a linear combination of natural images is not a natural image. It has a complicated form, and the atoms φ_i are a way to analyze it locally at many points using a kind of tangent structure.

So how can we find optimal sparse codings? The basic idea is to minimize an energy of the form

$$E(I, s) = \frac{1}{2} \|I - \Phi s\|^2 + \lambda \sum_{i=1}^{i=N} |s_i| ,$$

where the first term is the square of the Euclidean distance between the stimulus I and its coding Φs and the second term is the L^1 norm of the coding. Minimizing the

Fig. 3.51 Difference between a normal distribution (Gaussian) and a leptokurtic distribution. k is the excess kurtosis. The leptokurtic distribution has a more marked peak and fatter tail, but a thinner base. From [74]



first term ensures that the representation Φs will be a good approximation of I and minimizing the second term ensures that the coding is sparse. If we spell out this guiding idea algorithmically, we end up with a host of hard problems. The reader is referred to the papers by Bruno Olshausen and David Field [74] and Karol Gregor and Yann LeCun [75].

A good coding will reflect the statistical properties of the probability distributions of the pixels in the natural images. One of these properties is that the distributions are *non-Gaussian*, and more precisely, *leptokurtic*, i.e. more strongly peaked than Gaussians and with a fat tail or heavy tail, with most of the variance coming from either very weak or very strong extreme values, rather than fluctuations in intermediate values. A measure of this non-Gaussian character is the *kurtosis*, calculated from the fourth moment by the formula

$$K = \frac{\langle (I_r - \mu)^4 \rangle}{(\sigma^2)^2},$$

where $\mu = \langle I_r \rangle$ is the expected value of the pixels in the image and $\sigma^2 = \langle (I_r - \mu)^2 \rangle$ is the variance. It takes into account the fact that the edges represent sudden jumps, or ‘catastrophes’, compared with the pixel-to-pixel variations in the image. As the kurtosis is 3 for Gaussian distributions, we often consider the excess kurtosis $k = K - 3$, which is positive for leptokurtic distributions. Figure 3.51 shows the difference between a normal distribution (Gaussian) and a leptokurtic distribution. For the set Σ of natural images considered here, the kurtosis is the average kurtosis of the individual images.

As has been shown in a certain number of studies, such as those by Olshausen, Field, and Simoncelli [74, 76], the basis functions obtained by these methods look much like the RPs of simple V1 cells (see Figs. 3.52 and 3.53). We observe the orientation selectivity of the filters, that they are even, i.e. $\varphi(-x) = \varphi(x)$, with three lobes, or odd, i.e. $\varphi(-x) = -\varphi(x)$, with four lobes, their ON/OFF property, their spatial frequency selectivity, and their ability to detect endpoints.

Fig. 3.52 Basis functions obtained from large data bases of natural images. From [74, 76]





Fig. 3.53 Another example of basis functions obtained from large data bases of natural images



For further discussion of the models of RFs obtained by this method, the reader is referred to the paper [77] by Hyvärinen and Hoyer. The images $I(x, y)$ are treated as linear combinations of RPs $\varphi_i(x, y)$ with sparse responses s_i that have a non-Gaussian distribution, and this gives a discrete version of a convolution $I(x, y) = \sum_{i=1}^{i=N} s_i \varphi_i(x, y)$. The filters $\varphi_i(x, y)$ specify a discrete wavelet transformation which is invertible and can be used to reconstruct the s_i from $s_i = \langle w_i, I \rangle = \sum_{x,y} w_i(x, y) I(x, y)$. By whitening, i.e. elimination of second-order correlations, we may assume that the inverse filters $w_i(x, y)$ are orthonormal and the s_i statistically independent.

What is interesting is that, on top of this first layer of simple cells, the authors then introduce a second layer of more complicated cells which coordinate the simple cells locally, showing how their *functional architecture* also emerges as a principle of maximization of the ‘sparsity’ of the neural code. Because of their neighbourhood relations, the simple cells are no longer statistically independent and their responses s_i become correlated. We shall return to this point in Sect. 4.4.5 of Chap. 4, and in Sect. 5.6.3 of Chap. 5.

3.7 Signal Processing and Geometrical Formatting

We have just seen that one of the central theoretical problems of low-level vision is to understand how signal processing can also be geometrical formatting. Such formatting is necessary for the subsequent treatment of the optical signal by higher level routines, viz., cognitive, symbolic, and inferential. As we have seen the difficulty lies in the fact that the signal is not in itself a well structured geometrical object. It must therefore be transformed into a geometrical observable. For a signal, differentiation operations are not well posed problems. The algorithms of differential geometry cannot therefore be applied directly. To make differential geometry possible here, the price to pay is the introduction of a new degree of freedom, viz., scale.

Once we have thoroughly grasped the fact that, after transduction by the photoreceptors, the optical signal becomes a geometrical observable as a result of successive filtering operations which serve as measurement devices, and that the resulting measurements format it geometrically, we can understand the extent to which the objects of perception are *constructed* and not simply given. This construction—it would be better to say *constitution*—involves in an essential way, not only filtering, but also the structure of the network of connections between the neural filters, i.e. the cortical functional architectures.

Let us now move on to this problem, after just one last remark concerning other types of neurons.

3.8 Grid Cells and Place Cells

In the rest of the book, we shall focus the discussion mainly on visual neurons detecting geometrical features such as orientation or curvature. However, many other types of geometrical ‘calculations’ are carried out by neural modules. For example, we may cite *grid cells* and *place cells*, but also *head direction cells* in rodents, e.g. rats, which have been the subject of many studies.

3.8.1 Spatial Navigation

This kind of work comes under the general heading of *spatial navigation*, a particularly remarkable sensorimotor feat (i.e. perceptual and locomotor), clearly presented in the book *Le sens du mouvement* by Alain Berthoz [78]. For an animal, navigation involves localization and orientation in space, bringing together a considerable amount of geometric information regarding positions, displacement vectors, distances, and paths. To achieve this, it must be able to identify and monitor the positions and orientations (x, y, θ) of its body in a plane [or $(x, y, z, \theta, \varphi)$ in space]. There is an important distinction between *egocentric* and *allocentric* processing, which, in geometry, corresponds to the distinction between *moving frames* and *fixed frames*, introduced by Elie Cartan (see Table 3.2).

The aim is to understand the computational mechanisms involved in this processing and the way they are implemented by the neurons. The limbic system, and in particular the hippocampus for memory, the thalamus, the entorhinal cortex (parahippocampal gyrus), and the regions *CA1*, *CA2*, *CA3*, *CA4* (cornu ammonis regions), play a key role here. The animal can encode relative spatial positions by *integrating* linear and angular motions along the paths it follows from some initial position. Many experiments have been carried out on model animals such as rats and mice, which are able, starting from some point *A* and reaching some other point *B* along an extremely confusing random path, full of loops, to calculate and memorize the vector \overrightarrow{AB} . This *path integration* is quite astonishing. It makes it possible to glue different locations and orientations together into a single representation. The reader is referred to the 2006 review by McNaughton et al. [79].

Table 3.2 Egocentric and allocentric frames

Egocentric processing	Allocentric processing
Fixed observer	Moving observer
Moving external space	Fixed external space
Objects are targets of action	Subject reorients relative to landmarks
Cartan geometry: moving frame	Cartan geometry: fixed frame

3.8.2 *Place Cells*

In 1971, John O’Keefe and John Dostrovsky made the important discovery of the so-called *place cells* (PC), a major step forward which deeply changed many studies on perception and earned J. O’Keefe the Nobel prize in 2014, along with May-Britt and Edvard Moser, as already mentioned in Sect. 1.1 of the Preface. The first results were brought together in the book *The Hippocampus as a Cognitive Map*, published by O’Keefe and Nadel in 1978 [80] (already discussed in the *Introduction*).

The exceptional navigation skills of rats are well known. They can find their way through any labyrinth. (But note that, in these species of rodent, the column structure of the visual cortex is not very pronounced.) So these rats or mice are released into a circular arena D , bounded by a little wall, perfectly rotationally symmetric and without any kind of landmark that would allow the animals to pick out particular positions. The animal is then left to navigate freely in this homogeneous and isotropic space. After a certain time, it is noted that certain pyramidal cells in the CA1 of CA3 regions of the hippocampus (especially the dorsal hippocampus) are activated and fire, emitting several dozen spikes/s, when the animal’s head passes specific places, and that each of these cells codes a well-defined place, whence the name *place cells*. This coding, which is in fact a population coding, constitutes an *allocentric location system*. It is stable if the environment remains fixed and it is maintained in darkness. The places coded by the PCs are distributed fairly homogeneously in the arena D . The reader may consult the review by Muller [81], and also Muller et al. [82].

The question as to whether the PCs associate some preferred direction with which the animal crosses the given place is not easy to answer. Clearly, if the arena D is anisotropic, e.g. star-shaped, the distribution of the angles θ is no longer homogeneous and a directionality appears. When D is isotropic, it seems that most of the PCs, at least those coding positions far from the edge of D , are not directional. But some are, and they may perhaps be coupled with the head direction cells (see below). The reader is referred to the study by Muller, Bostock, and Taube [82].

3.8.3 *Grid Cells*

The *grid cells* were discovered in 2005 by Edvard and May-Britt Moser (who received the 2014 Nobel prize with John O’Keefe) in the pre- and post-subiculum, the entorhinal cortex (EC), or rather the dorsal EC and the dorsocaudal medial EC, the dorsal hippocampus, and the posterior parietal cortex of the rat. These are cells that fire when the animal, confined within an arena D and moving freely along long and repeated random paths, goes over the vertices of a rather regular triangular lattice Γ of positions in D . The mesh of the lattice is of the order of 25–30 cm in an arena D measuring 2 m across, and it increases along the dorsal–ventral axis. The grid cells supplement the place cells, which fire when the animal passes at a precise point of D .

They constitute topographic maps in which neighbouring grid cells have the same orientation and the same mesh, but different vertex positions (phases).

The last point is important. The grid cells are *mutually coherent*. This means that, idealizing and taking $D = \mathbb{R}^2$, they encode triangular lattices Γ_a which are copies of a basic lattice Γ_0 , translated by vectors a within the mesh of Γ_0 . To a first approximation, we may say that they implement the *quotient* D/Γ (see Fig. 3.54).

The grid cells form an original coordinate system for localization, orientation, and distance, allowing the animal to ‘calculate’ distances travelled and update its position. They are checked by means of landmarks, but once established, remain stable. The reader will find a discussion in the 2005 paper [83] by Hafting et al. Figure 3.55 shows the recording of a grid cell over 30 min. The ‘grid’ of positions is clearly visible.

Figure 3.56 by Doeller et al. [84] shows the relationship between the grid (triangular lattice Γ of positions) of a grid cell and the *preferred crossing direction* θ

Fig. 3.54 Coherence of the grid cells. The lattices encoded by the different cells are translated relative to one another by vectors lying within the mesh of the lattice. Here, we show a single hexagon of the lattice





Fig. 3.55 Activity of a grid cell monitored for 30 min. *Left* Random walk of a rat in a circular arena and activation clusters (*red*) of the cell. *Centre* Rate map showing the firing rate with peaks (*red*) and minima (*dark blue*). *Right* Spatial autocorrelation of the rate map, from $r = -1$ (*blue*), through $r = 0$ (*green*), to $r = 1$ (*red*). The triangular lattice of positions is clearly visible. From Hafting et al. [83]




Fig. 3.56 Relationship between the triangular lattice Γ of positions of a grid cell (*left*) and the preferred crossing direction θ associated with the positions (*right*). The direction $\theta = 32^\circ$ (red line in the left-hand image) is aligned with one of the axes of Γ

associated with the positions. The major experimental result is that θ is *aligned* with one of the axes of Γ .

Christian Doeller, Caswell Barry, and Neil Burgess used certain properties of the grid cells to show in [84] that they also exist in humans. The methodological problem is that, since invasive methods like electrodes are not possible, one must turn to non-invasive techniques such as fMRI, applying them to *virtual* navigation paradigms, allowing the subject to navigate around large spaces while remaining within the laboratory. The idea is (i) to identify fMRI signals in rats which are *correlated* with characteristic signals of the grid cells recorded using electrodes, and then (ii) to identify fMRI signals of the same kind in humans. This is no simple matter, because fMRI signals are macroscopic and result from averaging over thousands of individual neurons.

The authors use the coherence of the grid cells. Insofar as the axes of the lattices of the different cells are the same, for a given subject, the preferred directions are aligned along these axes, and the firing rate increases with the speed of navigation, one would expect a modulation of the fMRI signals analogous to the one in Fig. 3.57 (upper), that is, a modulation that increases with the speed and has a rotation symmetry of order 6. As the authors put it:

Entorhinal grid cells form a coherent population in which the common effects of orientation and speed of movement could produce a macroscopic signal visible with fMRI. [84, p. 657]

And this is precisely what is observed experimentally, as can be seen from Fig. 3.57 (lower).

Grid cells may be more fundamental than place cells. Indeed, as can be seen from Fig. 3.58 taken from McNaughton et al. [79], by summing over several grid cells on different scales and thresholding the result, we can obtain a place cell which activates only when the animal crosses some specific position.

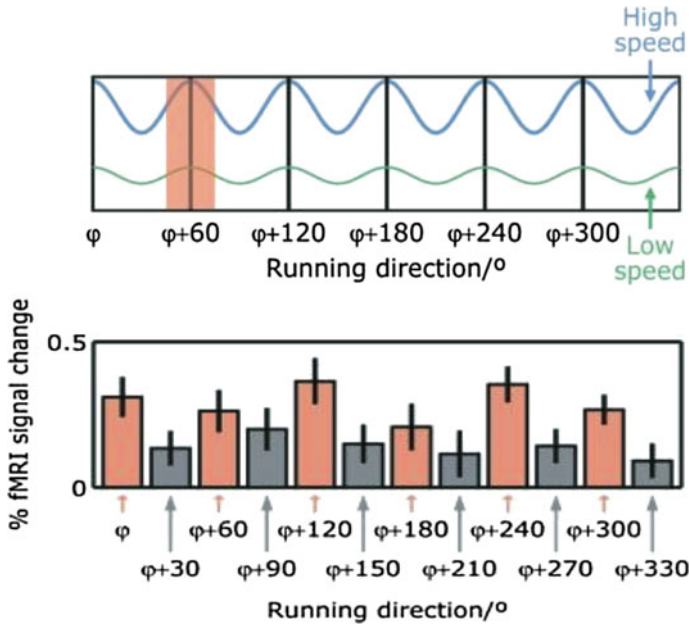


Fig. 3.57 Modulation of fMRI signals depending on the direction and speed of motion. *Upper* What we would anticipate: rotation symmetry of order 6 and an increase with speed. *Lower* What we observe. From Doeller et al. [84]




Fig. 3.58 By summing over several grid cells on different scales and thresholding the result, we obtain a place cell. From McNaughton et al. [79]

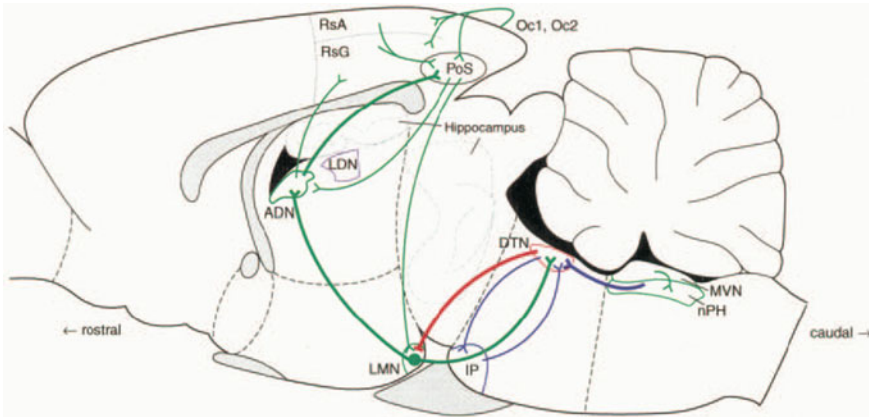


Fig. 3.59 HDC circuit in rats. The neural regions shown are as follows: LDN = laterodorsal thalamic nucleus, MVN = medial vestibular nucleus, nPH = nucleus prepositus, DTN = dorsal tegmental nucleus, LMN = lateral mammillary nucleus, IP = interpeduncular nucleus, ADN = anterodorsal thalamic nucleus, PoS = post-subiculum, RsA = agranulus retrosplenial cortex, RsG = granular retrosplenial cortex, Oc1 = primary visual area, Oc2 = secondary visual area. Signals regarding head movements enter the MVN from vestibular afferents. They are then conveyed to the nPH, and on to the DTN. The latter makes an inhibitory projection to the LMN, which makes ipsilateral excitatory projections back to the DTN. Likewise between the DTN and the IP. The LMN also projects bilaterally to the ADN. For its part, the latter projects to the PoS and the RsG. Cues from visual landmarks enter the HDCs via projections of the visual areas Oc1 and Oc2 to the retrosplenial cortex and the PoS, then from there to the subcortical structures of the LMN and the ADN. We thus have the projections nPH → DTN → LMN → ADN → Rs → PoS, and feedback PoS → LMN. From [85]

3.8.4 Head Direction Cells

Navigation involves locomotor systems which specify the displacement vector of the body and the direction of vision. There are also cells coding this direction. They were discovered by James Ranck Jr in the post-subiculum (PoS) of the rat and fire only when the animal’s head points in a certain direction (the direction relative to the surroundings, rather than relative to the body), independently of the motion and the position of the body. Then, Jeffrey Taube, a doctoral student of Ranck, extended the research on these *head direction cells* (HDC). There are many in the PoS (25% of some 200,000 cells, so around 50,000) and in the anterodorsal thalamic nucleus (ADN) (60% of around 20,000 cells, so around 12,000). They are controlled by information both from the surroundings (visual, auditory, and somatic cues) and from within the animal (vestibular, proprioceptive, and motor cues) regarding the animal’s proper motion. However, once established, they remain stable. Figure 3.59 from the 2003 paper [85] by Taube and Bassett shows the remarkable complexity of the HDC circuit in the rat.

3.8.5 Implementing the Tangent Bundle

Equipped with these different kinds of cells, the animal can develop its navigation skills using cognitive maps of its position, direction, and speed in the surrounding space. To simplify, on the 2D surface of the arena D , it can ‘calculate’, for its body and its head which guides the motion, its position $a \in D$, its direction of motion $\theta \in [2\pi]$, and its speed v . This means that its (sensorimotor, locomotor, cerebral) organism is able to ‘calculate’ the tangent vectors $\vec{v} \in T_a D$ and thereby implement the tangent bundle TD of D made up of the tangent planes $T_a D$. With the computational mechanisms of path integration, it can ‘calculate’ the vector between the ends of a complicate path, and from there, once the grid, place, and head direction cells have been established, optimize its movements and take shortcuts to get back to some position that is relevant to it, e.g., because there is a food reward there. The animal thus proves itself to be an ‘expert’ in differential and even Riemannian geometry ... As noted by Robert Muller:

Examples of interesting navigational capacities are the production of geodesic paths in unobstructed space, taking detours and taking shortcuts. [81, p. 821]

References

1. Imbert, M.: *Traité du Cerveau*. Odile Jacob, Paris (2006)
2. Bioinformatics. <http://www.bioinformatics.org/oeil-couleur/dossier/retine.html>
3. Brain. <http://thebrain.mcgill.ca>
4. Thalamus. <http://neurosci.wustl.edu>
5. Vision Web. <http://webvision.med.utah.edu>
6. Vision. Accessible at http://jeanpetitot.com/Woodruff-Pak_vision.ppt
7. Visual Cortex. <http://www.vision.ee.ethz.ch/en/>
8. Visual system. <http://brain.phgy.queensu.ca/pare/assets/Higher%20Processing%20handout.pdf> Processing handout.pdf
9. Buser, P., Imbert, M.: *Vision*. Hermann, Paris (1987)
10. Hooks, B.M., Chen, C.: Critical periods in the visual system: changing views for a model of experience-dependence plasticity. *Neuron* **56**, 312–326 (2007)
11. Feldheim, D.A., O’Leary, D.D.: Visual map development: bidirectional signaling, bifunctional guidance, molecules and competition. *Cold Spring Harb. Perspect. Biol.* **2**(11):a001768. doi:10.1101/cshperspect.a001768 (2010)
12. Shapley, R., Perry, V.H.: Cat and monkey retinal ganglion cells and their visual functional roles. *Trends Neurosci.* **9**, 229–235 (1986)
13. Ungerleider, L.G., Mishkin, M.: Two cortical visual systems. In: Ingle, D.J., Goodale, M.A., Mansfield, R.J.W. (eds.) *Analysis of Visual Behavior*, pp. 549–586. MIT Press, Cambridge, MA (1982)
14. Barlow, H.B., Levick, W.R., Yoon, M.: Responses to single quanta of light in retinal ganglion cells of the cat. *Vis. Res. Suppl.* **3**(11), 87–101 (1971)
15. Hartline, H.K.: The receptive fields of optic nerve fibers. *Am. J. Physiol.* **130**, 690–699 (1940)
16. Hodgkin, A.L., Huxley, A.F.: Currents carried out sodium and potassium ions through the membrane of the giant axon of *Loligo*. *J. Physiol.* **117**, 500–544 (1952)
17. Arikawa, K., Molday, L.L., Molday, R.S., Williams, D.S.: Localization of peripherin/rds in the disk membranes of cone and rod photoreceptors: relationship to disk membrane morphogenesis and retinal degeneration. *J. Cell Biol.* **116**(3), 659–667 (1992)

18. Adler, R., Raymond, P.A.: Have we achieved a unified model of photoreceptor cell fate specification in vertebrates? *Brain Res.* **1192**, 134–150 (2008)
19. Jonnal, R.S., Besecker, J.R., Derby, J.C., Kocaoglu, O.P., Cense, B., Gao, W., Wang, Q., Miller, D.T.: Imaging outer segment renewal in living human cone photoreceptors. *Opt. Express* **18**(5), 5257–5270 (2010)
20. Hennig, A.K., Peng, G.-H., Chen, S.: Regulation of photoreceptor gene expression by Crx-associated transcription factor network. *Brain Res.* **1192**, 114–133 (2008)
21. Kolb, H., Linberg, K.A., Fisher, S.K.: Neurons of the human retina: A Golgi study. *J. Comp. Neurobiol.* **318**, 147–187 (1992)
22. Lee, B.B., Wehrhahn, C., Westheimer, G., Kremers, J.: The spatial precision of macaque ganglion cell responses in relation to vernier acuity of human observers. *Vis. Res.* **35**(19), 2743–2758 (1995)
23. Badea, T.C., Cahill, H., Ecker, J., Hattar, S., Nathans, J.: Distinct roles of transcription factors *Brn3a* and *Brn3b* in controlling the development, morphology, and function of retinal ganglion cells. *Neuron* **61**, 852–864 (2009)
24. Nathans, J., Thomas, D., Hogness, D.S.: Molecular genetics of human color vision: the genes encoding blue, green, and red pigments. *Science* **232**, 193–202 (1986)
25. Neitz, J., Neitz, M.: The genetics of normal and defective color vision. *Vis. Res.* **51**, 633–651 (2011)
26. Kuffler, S.W.: Discharge patterns and functional organisation of mammalian retina. *J. Neurophysiol.* **16**, 37–68 (1953)
27. Frégnac, Y., Bringuier, V., Chavane, F., Glaeser, L., Lorenceau, J.: An intracellular study of space and time representation in primary visual cortical receptive fields. *J. Physiol.* **90**, 189–197 (1996)
28. Frégnac, Y., Shulz, D.: Activity-dependent regulation of receptive field properties of cat area 17 by supervised Hebbian learning. *J. Neurobiol.* **41**(1), 69–82 (1999)
29. Seriès, P., Lorenceau, J., Frégnac, Y.: The ‘silent’ surround of *V1* receptive fields: theory and experiments. *J. Physiol. Paris* **97**(4–6), 453–474 (2004)
30. Maffei, M., Fiorentini, A.: The unresponsive regions of visual cortical receptive fields. *Vis. Res.* **16**, 1131–1139 (1976)
31. Gilbert, C.D.: Horizontal integration and cortical dynamics. *Neuron* **9**, 1–13 (1992)
32. Lamme, V.A.F., Super, H., Spekreijse, H.: Feedforward, horizontal and feedback processing in the visual cortex. *Curr. Opin. Neurobiol.* **8**, 529–535 (1998)
33. Butts, D.A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N.A., Alonso, J.-M., Stanley, G.B.: Temporal precision in the neural code and the timescales of natural vision. *Nature* **449**, 92–95 (2007)
34. Thorpe, S.J., Guyonneau, R., Guilbaud, N., Allegraud, J.M., Vanrullen, R.: SpikeNet: real-time visual processing with one spike per neuron. *Neurocomputing* **58–60**, 857–864 (2003)
35. DeAngelis, G.C., Ohzawa, I., Freeman, R.D.: Receptive-field dynamics in the central visual pathways. *Trends Neurosci.* **18**(10), 451–458 (1995)
36. Somers, D.C., Nelson, S.B., Sur, M.: An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* **15**, 5448–5465 (1995)
37. Troyer, T.W., Krukowski, A.E., Priebe, N.J., Miller, K.D.: Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *J. Neurosci.* **18**, 5908–5927 (1998)
38. Wörgötter, F., Koch, C.: A detailed model of the primary visual pathway in the cat: comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *J. Neurosci.* **11**, 1959–1979 (1991)
39. Martinez, L.M., Alonso, J.-M.: Complex receptive fields in primary visual cortex. *Neurosci.* **9**(5), 317–331 (2003)
40. Gilbert, C.D.: Circuitry, architecture and functional dynamics of visual cortex. In: Bock, G.R., Goode, J.A. (eds.) *Higher Order Processing in The Visual System*, Ciba Foundation Symposium, vol. 184, pp. 35–62. Wiley (1994)
41. Marcelja, S.: Mathematical description of the response of simple cortical cells. *J. Opt. Soc. Am.* **70**, 1297–1300 (1980)

42. Kayser, A., Priebe, N.J., Miller, K.D.: Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J. Neurophysiol.* **85**, 2130–2149 (2001)
43. Carandini, M., Ringach, D.L.: Predictions of a recurrent model of orientation selectivity. *Vis. Res.* **37**, 3061–3071 (1997)
44. Ben-Yishai, R., Bar-Or, R.L., Sompolinsky, H.: Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci.* **92**, 3844–3848 (1995)
45. McLaughlin, D., Shapley, R., Shelley, M., Wielaard, D.J.: A neuronal network model of macaque primary visual cortex (*V1*): orientation selectivity and dynamics in the input layer 4C α . *Proc. Natl. Acad. Sci.* **97**(14), 8087–8092 (2000)
46. Cai, D., DeAngelis, G.C., Freeman, R.D.: Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *J. Neurophysiol.* **78**, 1045–1061 (1997)
47. Saul, A.B.: Lagged cells. *Neurosignals* **16**, 209–225 (2008)
48. Wandell, B.A., Winaver, J.: Imaging retinotopic maps in the human brain. *Vis. Res.* **51**, 718–737 (2011)
49. Dumoulin, S.O., Wandell, B.A.: Population receptive field estimates in human visual cortex. *NeuroImage* **39**, 647–660 (2008)
50. Daugman, J.G.: Uncertainty relation for resolution in space, spatial frequency and orientation optimized by two dimensional visual cortical filters. *J. Opt. Soc. Am.* **2**(7), 1160–1169 (1985)
51. Jones, J.P., Palmer, L.A.: An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.* **58**, 1233–1258 (1987)
52. Bloom, J.A., Reed, T.R.: An uncertainty analysis of some real functions for image processing applications. In: *International Conference on Image Processing*, vol. 3, IEEE, pp. 670–672 (1997)
53. Mallat, S.: *A Wavelet Tour of Signal Processing*. Academic Press, New York (1998)
54. Morgan, A.P., Watson, L.T., Young, R.A.: A Gaussian derivative based version of JPEG for image compression and decompression. *IEEE Trans. Image Process.* **7**(9), 1311–1320 (1998)
55. Young, R.A.: The Gaussian derivative model for spatial vision. *Spat. Vis.* **2**(4), 273–293 (1987)
56. Young, R.A., Lesperance, R.M., Meyer, W.W.: The Gaussian derivative model for spatio-temporal vision, I. *Spat. Vis.* **14**(3–4), 261–319 (2001)
57. Freeman, W.T., Adelson, E.H.: The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.* **13**(9), 891–906 (1991)
58. Shapley, R.: Linear and nonlinear systems analysis of the visual system: why does it seem so linear? *Vis. Res.* **49**, 907–921 (2009)
59. Florack, L.M.J.: *The Syntactical Structure of Scalar Images*. PhD Thesis, University of Utrecht (1993)
60. Florack, L.M.J., ter Haar Romeny, B.M., Koenderink, J.J., Viergever, M.A.: Scale and the differential structure of images. *Image Vis. Comput.* **10**(6), 376–388 (1992)
61. Hale, D.: *Recursive Gaussian filters*, Center for Wave Phenomena. Report 546, Colorado School of Mines (2006)
62. Marr, D.: *Vision*. W.H. Freeman, San Francisco (1982)
63. Beaudot, W.H.A., Mullen, K.T.: Orientation selectivity in luminance and color vision assessed using 2D band-pass filtered spatial noise. *Vis. Res.* **45**, 687–695 (2005)
64. Meyer, Y.: Ondelettes, filtres miroirs en quadrature et traitement numérique de l’image. *Gazette des Mathématiciens* **40**, 31–42 (1989)
65. Mallat, S.: Multifrequency channel decompositions of images and wavelet models. *IEEE Trans. Acoust. Speech Signal Process.* **37**(12), 2091–2110 (1989)
66. Mallat, S., Zhong, S.: Complete signal representation with multiscale edges. Technical Report no. 483, Department of Computer Sciences, New York University (1989)
67. Mallat, S., Peyré, G.: Traitements géométriques des images par bandelettes. *Journée annuelle de la SMF, Mathématiques et Vision*, 39–67 (24 June 2006)
68. Wickerhauser, M.V.: *Lectures on Wavelet Packet Algorithms*. Technical Report, Department of Mathematics, Washington University (1991)
69. Hamy, H.: *Méthodes géométriques multi-échelles en vision computationnelle*. PhD Thesis, École Polytechnique, Paris (1997)

70. Atick, J.: Could information theory provide an ecological theory of sensory processing? *Network* **3**, 213–251 (1992)
71. Field, D.J.: Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4**(12), 2379–2394 (1987)
72. Turiel, A., Nadal, J.-P., Parga, N.: Orientational minimal redundancy wavelets: from edge detection to perception. *Vis. Res.* **43**(9), 1061–1079 (2003)
73. Nirenberg, S., Carcieri, S.M., Jacobs, A.L., Latham, P.E.: Retinal ganglion cells act largely as independent encoders. *Nature* **411**(7), 698–701 (2001)
74. Olshausen, B.A., Field, D.J.: Sparse coding of sensory inputs. *Cur. Opin. Neurobiol.* **14**, 481–487 (2004)
75. Gregor, K., LeCun, Y.: Learning fast approximations of sparse coding. <http://www.icml2010.org/papers/449.pdf>
76. Simoncelli, E.P., Olshausen, B.A.: Natural image statistics and neural representation. *Ann. Rev. Neurosci.* **24**, 1193–1216 (2001)
77. Hyvärinen, A., Hoyer, P.O.: A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vis. Res.* **41**(18), 2413–2423 (2001)
78. Berthoz, A.: *Le sens du mouvement*. Odile Jacob, Paris (1997)
79. McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.-B.: Path integration and the neural basis of the ‘cognitive map’. *Nat. Rev. Neurosci.* **7**, 663–678 (2006)
80. O’Keefe, J., Nadel, L.: *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford (1978)
81. Muller, R.U.: A quarter of a century of place cells. *Neuron* **17**, 979–990 (1996)
82. Muller, R.U., Bostock, E., Taube, J.S., Kubie, J.L.: On the directional firing properties of hippocampal place cells. *J. Neurosci.* **14**(12), 7235–7251 (1994)
83. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**(11), 801–806 (2005)
84. Doeller, C.F., Barry, C., Burgess, N.: Evidence for grid cells in a human memory network. *Nature* **463**, 657–661 (2010)
85. Taube, J.S., Bassett, J.P.: Persistent neural activity in head direction cells. *Cerebral Cortex* **13**, 1162–1172 (2003)

Chapter 4

Functional Architecture I: The Pinwheels of V1

We have just seen in the previous chapter how the mechanisms used by visual neurons to filter the optical signal can implement localized differential data. But this is still insufficient; the visual cortex can enable the perception of forms only because it can pass from local to global and *integrate* local differential data into global geometric forms. This extraordinary performance has intrigued all vision specialists, even well before the famous psychophysical experiments of the Gestalt theory at the beginning of the last century.

In this chapter and the next, we will explore, among other things, the neural infrastructures that underlie the simplest of the Gestalt principles, that of so-called good continuation, which is associated with the ideal geometric notion of a line. To do so, we must introduce the *functional architecture* of the visual areas. It is indeed the very particularity and global coherence of this architecture that create the geometry we want to model. We should be clear that the remarkable geometric properties of perception result from very specific properties of the generating neural dynamics and that this specificity requires equally specific functional architectures. Architectures with connections that are either too local or global, either too isotropic or random, *could not* create such geometric structures.

As we have said, we will give particular importance to V1, the first of the primary visual areas. Such a restriction might seem too drastic, insofar as the later areas like V2 or V4 have many ‘top-down’ feedback connections to V1. But V1 is already very important, more important than one might think. Let us reiterate that we are adopting Mumford and Lee’s high-resolution buffer hypothesis (Lee et al. [1]), according to which V1 is not a simple ‘bottom-up early module’, but participates in all visual processes requiring fine resolution, whence its functional architecture is essential for the totality of the visual system. As William Beaudot and Kathy Mullen argue in [2, p. 688]:

All higher aspects of form perception rely on this early orientation-selective processing stage.

The literature in this area is immense. We therefore run into an obvious problem of presentation. We have tried to resolve this problem by presenting, although certainly in too brief a manner, enough data to give the reader some idea of the variety, the richness, and the difficulty of the questions. This research requires innumerable experimental feats, so we must be aware of the fact that, despite the mass of remarkable experimental results already available, all of the data presented here is ‘work in progress’, subject to debate.

The functional architecture includes two main components, that is two main classes of connections. The first is the class of retino-geniculo-cortical ‘vertical’ connections, which will be studied in this chapter. The second is the class of cortico-cortical ‘horizontal’ connections, which will be discussed in the next chapter.

4.1 The Areas of the Visual Cortex

In Sect. 3.1 of Chap. 3, we very briefly mentioned the visual pathways and areas. Here, we give a few additional clarifications. In humans, there are about fifty such areas in the cortex, totalling an area of about 2 500 cm². The thickness of the cortex varies between 2 and 3 mm, which makes a volume of about 625 cm³. Estimating the average number of neurons per mm³ to be from 20 000–50 000 yields on the order of 10 to 30 billion neurons and 60 to 240 × 10¹² synapses, although these estimates vary slightly among specialists.

We have seen in Sect. 3.1 of Chap. 3 (see Fig. 3.1) the general structure of the retino-geniculo-cortical pathways. Area V1 corresponds to area 17 in the classification of Korbinian Brodmann, represented in Fig. 4.1. A large part is located in the calcarine fissure, with the retina’s ‘fovea → periphery’ gradient corresponding to the ‘occipital pole → anterior fissure’ gradient with a magnification of the fovea (see Fig. 4.12 below). It includes about 100 million neurons (recall that the optical nerve contains about a million fibres). Figures 4.2 and 4.3 show the localization of the visual areas around V1. Figure 4.4 shows a planar flattening of the structure of the visual areas.

Area V2 consists of a ventral part and a dorsal part in each hemisphere; as we shall see in Sect. 5.8 of Chap. 5, it plays an important role in the detection of illusory contours and the determination of an edge of a figure in front of a background. It is the same for V3. V2 projects onto the dorsal V3, which in turn projects onto the parietal cortex. The ventral V3 projects onto the inferotemporal (*IT*) cortex. Area V4, the third ventral area after V2 and V3, receives direct projections from V1 and V2. As we shall see, it is important for colour processing and recognition of simple forms. Area V5, also called *MT* for ‘medio-temporal’, is sensitive not only to orientations but also to directions, and it is essential for detecting local movement (e.g. the motion of points) as well as for controlling eye movements. When V5 is damaged, vision is altered by giving the subject the impression of perceiving static images. It is equally sensitive to binocular disparity and therefore plays a part in stereopsis. It receives




Fig. 4.1 The Brodmann areas (1909). ●1, 2, 3 Primary sensory cortex. Postcentral gyrus. *Sensitivity*. ●4 Primary motor cortex. Precentral gyrus. *Motricity*. ●5 Posterior parietal association area. Superior parietal gyrus. *Stereognosis*. ●6 Premotor cortex and supplementary motor area. Precentral gyrus and adjacent rostral cortex. *Programming of movements*. ●7 Posterior parietal association area. Superior parietal gyrus. *Visuomotor coordination, perception*. ●8 Frontal oculomotor field. Superior and middle frontal gyrus, internal face. *Saccades*. ●9–12 Prefrontal association cortex. Superior and middle frontal gyrus, internal face. *Cognitive areas, programming of movements*. ●13–16 Vegetative areas. Insular cortex. ●17 Primary visual area. Calcarine fissure. *Vision*. ●18 Secondary visual area. Around area 17. ●19 Tertiary visual area. Around area 18. ●20 Inferotemporal visual area. Inferior temporal gyrus. *Recognition of forms*. ●21 Inferotemporal visual area. Middle temporal gyrus. *Recognition of forms*. ●22 Association auditory area. Superior temporal gyrus. *Hearing*. ●23–27 Limbic association cortex. Subcallosal, cingulate, retrosplenial, parahippocampal cortex. *Emotion, memory*. ●28 Olfactory cortex, limbic association cortex. Parahippocampal gyrus. *Smell, emotions*. ●29–33 Limbic association cortex. Cingulate and retrosplenial gyrus. *Emotions*. ●34–36 Olfactory cortex, limbic association cortex. Parahippocampal gyrus. *Scents, emotions*. ●37 Parietal-temporal-occipital association cortex, middle temporal visual area. Middle and inferior temporal gyrus (T-O junction). *Perception, vision, reading, language*. ●38 Olfactory cortex, limbic association cortex. Temporal pole. *Scents, emotions*. ●39 T-P-O association cortex. Temporo-parieto-occipital junction (angular gyrus). *Perception, vision, reading, language*. ●40 T-P-O association cortex. Temporo-parieto-occipital junction (supramarginal gyrus). *Perception, vision, reading, language*. ●41–42 Primary auditory cortex. Heschl’s gyrus and superior temporal gyrus. *Hearing*. ●43 Olfactory cortex. Insular cortex, frontal parietal operculum. ●44 Broca’s area, lateral premotor cortex. Inferior frontal gyrus (frontal operculum). *Language, planning of movement*. ●45 Prefrontal association cortex. Inferior frontal gyrus. ●46 Prefrontal association cortex. Middle frontal gyrus. ●47 Inferior frontal gyrus. Taken from Hasboun [3]




Fig. 4.2 Localization of visual areas around V1. Taken from [4]




Fig. 4.3 Three-dimensional view of the localization of visual areas. Taken from [5]




Fig. 4.4 Planar flattening of the structure of the visual areas. Taken from [6]

many connections from V1, but also from V2 and the dorsal V3. All of these areas possess numerous feedback connections coming back down towards the LGN.

The dorsal and ventral pathways that we have encountered are, respectively:

1. The magnocellular pathway dealing with spatial localization and movement (the ‘Where’ pathway of Ungerleider and Mishkin [7]), viz.

LGN (magno) → V1 (4C α) → V1 (4B) → V2 (thick stripes) → MT (V5)

which leads, after the parietal cortex, all the way to the frontal cortex, where the frontal oculomotor field is located (among others). The latter directs the gaze.

2. The parvocellular pathway dealing with forms and colour (the ‘What’ pathway of Ungerleider and Mishkin [7]), viz.




Fig. 4.5 Visual areas in humans (see text for acronyms). The prefix h stands for ‘human’, while o, +, – denote the fovea and the superior and inferior visual fields, respectively. Taken from [8]




Fig. 4.6 fMRI of the retinotopic projection of a visual hemifield onto the corresponding visual hemisphere. The concentric circles of varying eccentricity and the rays of varying orientation on the half-disc target are colour-coded. Taken from [8]




Fig. 4.7 Position of V1 (and V2, V3) in the cortex. The *upper edge* corresponds to the inferior vertical semi-meridian, and the *lower edge* to the superior vertical semi-meridian. The horizontal semi-meridian corresponding to the represented hemisphere is situated along the median fissure. Note that the scale of the eccentricities (radii of the circles of the retinal half-target) is logarithmic. Taken from [11]

LGN (parvo) $\rightarrow V1$ ($4C\beta$) $\rightarrow V1$ (2/3 blobs)
 $\rightarrow V2$ (thin stripes) $\rightarrow V3 \rightarrow V4$ and IT

With fMRI methods, already mentioned in Sect. 3.2.6.4 of Chap. 3 (phase-encoded retinotopy on travelling wave) which use checkered dynamic stimuli (expanding rings and rotating wedges) and record the induced cortical activity waves, we can get a better idea of what these areas do in a human brain. In Fig. 4.5, taken from the article [8] by Brian Wandell and Serge Dumoulin, we see the medial occipital areas ($V1$, $V2$, $V3$), lateral occipital areas ($LO-1$ and $LO-2$, see [9], $hMT+$, corresponding in humans to the macaque’s area $MT = V5$), the ventral occipital areas ($hV4$, $VO-1$, $VO-2$), the dorsal occipital areas ($V3A$, $V3B$), and the posterior occipital areas (from $IPS-0 = V7$ to $IPS-4$, $IPS =$ intraparietal sulcus).

Figure 4.6 is an fMRI version of the classic Fig. 4.12 (see below) first obtained by Tootell et al. [10] in the macaque using lethal methods. More detail is shown in Fig. 4.7, which is a more accurate representation of the position of $V1$ and its boundaries, showing how the two hemispherical $V1$ areas are glued along the vertical meridian by the callosal connections.

The definition of the visual areas is often a delicate matter. To give just one example, it remains an open debate as to whether, in humans, an area $V8$ should be

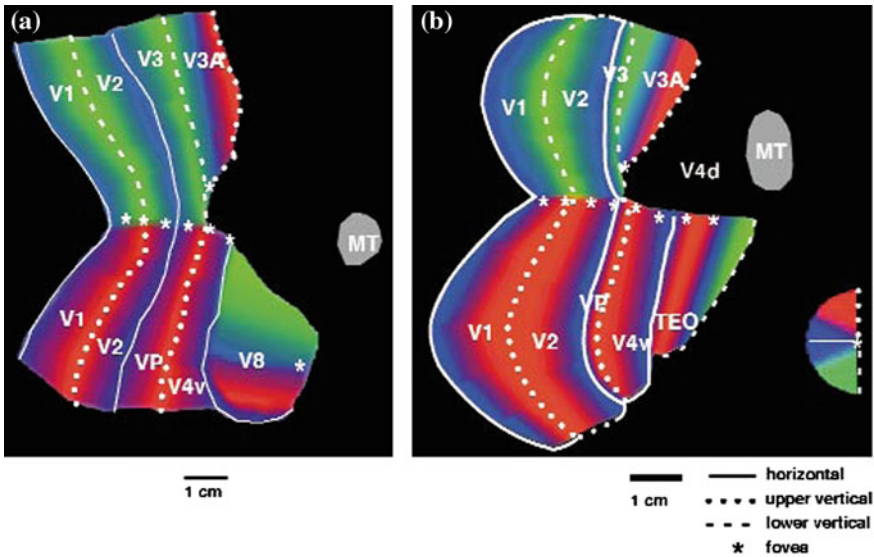


Fig. 4.8 Comparison of the polar angle retinotopy in humans (a) and in macaques (b). The stimuli are radial gratings rotating at low spatial frequency (the colour code of the angles is given *bottom right*). The boundaries of visual areas $V1$ to $V4$ can thus be established, and the mechanism for representing angles. The orthogonality of the angular distributions between area $V8$ in humans and area TEO in macaques is clearly visible. $V4v$ is the ventral $V4$ and $V4d$ is the dorsal $V4$. Taken from [13, Fig. 1]

introduced—possessing its own retinotopic representation—that would correspond to the area *TEO* (posterior inferotemporal cortex) in macaques (see [12]), or whether it would be more properly considered as a sub-area of *V4*, the problem being that, with respect to the retinotopy of the already well-identified *V4*, the retinotopy of this hypothetical area would lie at an angle of 90° (see Fig. 4.8 of [13]).

All these areas develop and stabilize with great plasticity through (i) critical periods, (ii) exposure to an enormous amount of stimuli, and (iii) the spontaneous cortical activity that induces neuronal waves (see, e.g., Hooks and Chen [14]).

4.2 Hypercolumnar Structure of the V1 Area

In *V1*, the density of neurons is on the order of $2 \times 10^5/\text{mm}^2$, while the receptive fields are on the order of a few degrees. Neurophysiological studies have made it possible to distinguish three types of structures in *V1*. These are layered, retinotopic, and (hyper)columnar, respectively.

Layered Structure

This is about 1.8 mm thick and composed of 6 ‘horizontal’ layers, i.e., parallel to the surface of the cortex. The most important for us is layer 4 (traditionally called ‘granular’) and more specifically the sub-layer *4C*. Most of the fibres coming from the lateral geniculate nucleus project onto the latter, with sub-layer *4C α* receiving in particular the magnocellular projections and sub-layer *4C β* the parvocellular projections (see Figs. 4.9 and 4.10). Layer 4 relays these inputs towards the ‘supragranular’ layers 2 and 3 and also projects in the other direction, onto the ‘infragranular’ layers, connecting them to the thalamus (layer 6) and the spinal cord (layer 5). Figure 4.11 shows for the macaque a small part of this complex network of projections containing several loops (see also Sect. 5.10 of Chap. 5). This structure does depend on the species, but we shall not go into the details here, apart from a brief note in Sect. 4.9.4, despite the importance of this observation.

Fig. 4.9 Layered structure of *V1* and layer 4. Taken from [6]




Fig. 4.10 Projection of the lateral geniculate nucleus onto layer 4C of V1. Taken from [6]





Fig. 4.11 Part of the network of projections between the layers of V1 in the macaque. Taken from [15]



Retinotopic Structure

Retinotopy refers to the fact that the projections (in the neurophysiological sense) of the retina onto the cortical layers are in fact mappings (in the mathematical sense) preserving the retinal topography. A typical example is the logarithmic conformal mapping between the retina and the 4C sub-layer of layer 4 where, as we have just seen, it is mainly fibres from the lateral geniculate nucleus that project. If R is the plane of the retina and M the cortical layer, the retinotopy is described by a mapping $\chi : R \rightarrow M$ which is a non-isometric isomorphism for a certain level of geometric structure that is not as rigid as the metric level (see Fig. 4.12).¹

There are several models for this retinotopic mapping. The first was a monopole model $\text{Log}(z + a)$, where z is a complex variable varying in R , but better is a dipole model like

¹When we don't need to distinguish between R and M , we shall set $R = M$ and $\chi = Id$.

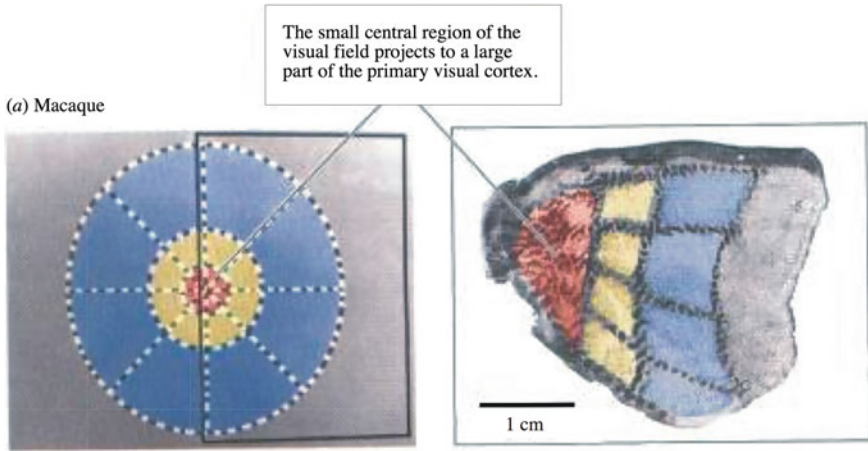



Fig. 4.12 Retinotopic projection of the retina onto layer 4C of area V1. The retinal hemitarget is transformed by a diffeomorphism which is not an isometry but a conformal map. From [10]. See also [16]

Fig. 4.13 Mathematical model of the retinotopic projection of the retina onto V1 using a conformal map that is a logarithm of a homography



$$\text{Log} \frac{z + 0.333}{z + 6.66} .$$

This agrees quite well with empirical data [17], as shown in Fig. 4.13. A still better model is a wedge-dipole model common to the three areas V1, V2, and V3, such as

$$\text{Log} \frac{w(z) + a}{w(z) + b} ,$$

illustrated in Fig. 4.14.




Fig. 4.14 Model of the form $\text{Log}[(w(z) + a)/(w(z) + b)]$ for V1, V2, and V3. *Left* Areas V1, V2, and V3. *Right* Fit with the model. From Balasubramanian et al. [18]

Column and Hypercolumn Structure

This was the great discovery of Nobel prizewinners David Hubel and Torsten Wiesel in the early 1960s, following up work by their PhD supervisor Stephen Kuffler on retinal GCs (see Sect. 3.2.6.1). This was preceded in the late 1950s by the work of Vernon Mountcastle on the somatosensory cortex of the cat, and after Hubel and Wiesel, it was also found in the motor cortex and the auditory cortex. As we saw in Sect. 3.2.6 of Chap. 3, there are ‘simple neurons’ in the V1 area (as opposed to the ‘complex’ and ‘hypercomplex neurons’ discussed in Sect. 3.2.6.3) which are sensitive to orientation, ocular dominance, and colour. These are the ones of interest to us here.

We have already considered the structure of their receptive fields and receptive profiles. We shall focus particularly on those having the form of a second derivative of a Gaussian. If to begin with we simplify the situation as far as possible by not taking into account either the scale (the resolution and the spatial frequency) or the phase, we may say that these neurons detect pairs (a, p) of retinal positions a and orientations p at a . In Sect. 2.7 of the *Introduction*, these were referred to as mesoscopic *contact elements* of the visual plane. Indeed, as we said in Sect. 3.3.4 of Chap. 3, if we consider an edge crossing the receptive field, the response of the neuron will be maximal when the edge is aligned with the preferred orientation.

Using the methods to record responses to appropriate stimuli discussed in the last chapter (oriented bars crossing the RF of the neurons, etc.), it can be shown that, perpendicular to the surface of the cortex, the retinal position a and the preferred orientation p remain roughly constant. This ‘vertical’ redundancy—which specifies a population coding of the position—defines *orientation columns* of about 20 μm . As shown by DeAngelis [19], it is the *phase variation* that dominates in the columns:

Spatial phase is the single parameter that accounts for most of the difference between receptive fields of nearby neurons.




Fig. 4.15 Classic experiments by Hubel and Wiesel which led to the discovery of the orientation hypercolumns in V1. Taken from [24]

Moreover (see Sect. 3.3.5 of Chap. 3), the population coding gives the system a higher resolution than single neurons. The reader is referred to Beaudot and Mullen [2, 20], who suggest a mechanism to explain this, and also Snippe and Koenderink [21], and Ringach [22]. The latter studies the dependence of the resolution on the Fourier spectrum of the tuning curves of the neurons.

On the other hand, along lines parallel to the surface of the cortex, the preferred orientation p varies in steps of about 10° . To a first approximation, it can be considered to vary monotonically along line segments. In fact, this is not really the case, and for subtle topological reasons which we shall explain in Sect. 4.4.3 when we comment on Braitenberg's paper [23], the variation of p is not necessarily monotonic. However, we shall suppose here that it is. A 'horizontal' grouping of columns whose orientations vary over a range of π defines an *orientation hypercolumn* which is a broad neural micromodule measuring between $200\ \mu\text{m}$ and $1\ \text{mm}$ (see Figs. 4.15, 4.16 and 4.17).

The idea of a column just discussed clearly needs to be fleshed out, and a great deal of work has been devoted to this, dealing with problems such as establishing the link between anatomical and functional definitions, investigating the genuinely columnar nature of different ways of processing the stimuli (i.e. asking to what extent they cross the layers of V1), analyzing the strengthening of the edges of the columns by lateral inhibition. Among the various reviews available on these matters, the reader is referred to the one by Lund et al. [25].

We shall discuss different species of mammals, such as the tree shrew (tupaia), cats, primates, and humans, where the concept of orientation column is of key importance. This is quite restrictive because, as we saw in Sect. 3.8 of Chap. 3, in other species of mammals, and in particular rodents like the rat, this idea is less relevant.





Fig. 4.16 Hypercolumn of V1 from Hubel and Wiesel

Fig. 4.17 Classic diagram of hypercolumns. From [6]



There we encounter other types of spatial processing, as carried out, for example, by the grid and place cells (see Sect. 3.8 of Chap. 3). There is orientation selectivity, but it is *dispersed* throughout V1; i.e., it is not brought together anatomically in columns and hypercolumns by a functional architecture.

However, even restricting to species for which this idea is relevant, there is nevertheless a certain diversity in the structure of V1. The general structure, concerning retinotopy, orientation selectivity, spatial frequency tuning, etc., remains the same, but the fine structure varies. As noted by Stephen Van Hooser:

There is considerable diversity in the abundance of different cell classes, laminar organization, functional architecture, and functional connectivity. [26]

Figure 4.18 compares several different species: the macaque (primate, diurnal, frontally placed eyes), the cat (carnivorous, crepuscular, frontally placed eyes), the tree shrew or tupaia (Scandentia, diurnal, laterally placed eyes), and the gray squirrel (rodent, diurnal, laterally placed eyes).² We find that there is indeed orientation selectivity everywhere, but that it is not always present in all layers, as in the cat or the

²Van Hooser’s paper also discusses the rat (rodent, nocturnal, laterally placed eyes), the night monkey, also known as the owl monkey or *douroucoulis* (New World primate, nocturnal, frontally placed eyes), the ferret, etc.

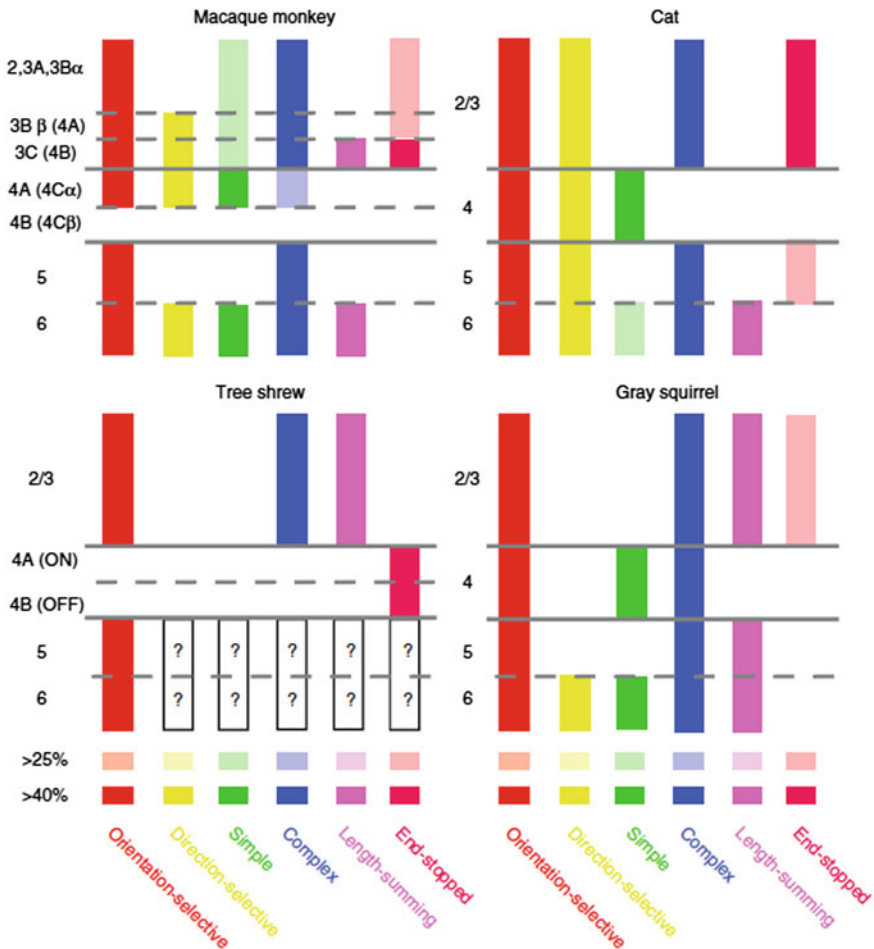


Fig. 4.18 Similarities and differences in a certain number of geometric features for different species of mammals. Taken from Van Hooser [26]

squirrel. Direction selectivity occurs throughout all layers in the cat (and also the ferret), but less in primates and scarcely at all in the tree shrew. The hypercomplex (end-stopped) cells occur in layers 2 and 3, but not in layer 4, except in the tree shrew, where they are in fact present only in layer 4. The cells responding most strongly when the bars on the gratings are longer (length-summing cells) are rare in carnivores, etc.

We thus see just how important the hypercolumn structure of V1 is and how crucial it is to find a suitable mathematical structure to model it. We shall show that, on the *mesoscopic* scale that is our own, the right concept, which is altogether natural, is the geometric concept of a *fibration*, so well known to mathematicians.

4.3 V1 as a Mesoscopic Fibration

4.3.1 ‘Bridging Scales’: The Mesoscopic Level

We should stress, as we have already done in Sect. 2.7 of the *Introduction*, that the geometric structures we have introduced refer to the *mesoscopic* scale. Indeed, realistic ‘micro’ simulations of an orientation column are already unbelievably complex and have featured among some of the biggest computational projects in the world.

Launched in 2005 by Henry Markram at the École Polytechnique Fédérale de Lausanne (EPFL), the *Blue Brain Project* (BBP) aimed to simulate a rat cortical column containing some 10,000 neurons and 30 million synapses. It had access to a computational power (supplied by IBM) of 20 teraflops (one teraflop is 10^{12} operations per second). Since 2013, the *Human Brain Project* (rival of the American project *Brain*) has picked up where Blue Brain left off, using supercomputers of up to 10^6 teraflops. The HBP aims to simulate complete brain areas, and in particular, the visual cortex. Bringing together many universities and research institutes, it has been chosen as one of the two Future and Emerging Technologies (FET) Flagships of the European Union and will receive a total of a billion euros over 10 years. It explicitly exploits the possibilities of the information and communication technologies (ICT) and involves six platforms: neuroinformatics, brain simulation, high-performance computing, medical informatics, neuromorphic computing, and neurorobotics. Mathematical models do not play a major role because, as noted by those who designed the project, they are just ‘toy models’, drastically simplifying the biophysical data to make them amenable to mathematical analysis.

However, certain modelling problems have been identified, in particular regarding the four themes:

- bridging scales,
- synaptic plasticity, learning, and memory,
- large-scale models,
- principles of brain computation.

An internal debate has sprung up between massively computational approaches and more model-oriented and structural approaches. Our own model is an element in this debate.

4.3.2 Fibrations and Engrafted Variables

On the mesoscopic level, through the hypercolumn functional architecture of V1, a (discretized) copy of the space P of directions p in the plane is associated (in a retinotopic and anatomically observable way) with each retinal position $a \in R$. As a consequence, there is a *neural implementation* of the projection $\pi : R \times P \rightarrow R$

of the Cartesian product $R \times P$ onto its first factor R , a projection which is in fact a (trivial) fibration with base the retinal space R and fibre the manifold P .

In these mesoscopic models, a column is thus viewed as a simple contact element (a, p) . This may look like an unacceptable reduction with regard to the computational projects just mentioned. However, we shall set out to show here that, despite the enormous reduction in complexity, the functional organization of these elements leads to a significant geometric complexity. This means that the complexity of a realistic ‘micro’ model of a functional architecture only becomes manageable by using highly elaborate multiscale bridges.

According to the materialist principle that ‘a structure only exists if it is implemented’, we must inquire into the neural materiality of the mesoscopic projection $\pi : R \times P \rightarrow R$ over and above its formal ideality. Now, if the projection π is mathematically (ideally) trivial, it is not at all so from the neurophysiological (material) point of view:

- (i) The receptive fields of the ganglion cells and cortical neurons, or rather the cortical columns, are very small *local charts* which overlap and can be glued together on the overlap.
- (ii) We shall see in Sect. 4.9.1 that it is no easy matter to test the direct product structure experimentally, i.e. the independence of the position and orientation variables. There, too, there are significant differences between species, as we saw above: for some species, it is a single layer of V1 which implements the product $R \times P$, while for others, several layers are involved.
- (iii) The projection π (in the geometrical sense) is implemented by the fine circuitry of the retino-geniculo-cortical pathways, which project the retina onto V1 (in the neurophysiological sense).³

The geometric structure of the product of R as base space with a space of ‘secondary’ variables such as orientation, ocular dominance, direction of motion was well expressed by David Hubel when he spoke of ‘engrafted variables’:

What the cortex does is map not just two but many variables on its two-dimensional surface. It does so by selecting as the basic parameters the two variables that specify the visual field coordinates (distance out and up or down from the fovea), and on this map it engrafts other variables, such as orientation and eye preference, by finer subdivisions. [27, p. 131]

A hypercolumn can thus be modelled as the Cartesian product of the RFs with the space of secondary variables which are ‘grafted on’, so to speak. The overlaps of the RFs are then interpreted geometrically as a *gluing together* of these local models, and the projection π is for its part implemented in the vertical connections:

- from the retina to the hypercolumns,
- within the hypercolumns themselves.

³In the geometric models of neural functional architectures, there are many problems of terminology. Lexical items such as ‘fibre’, ‘projection’, ‘connection’ are used in different ways by mathematicians and neurophysiologists. In general, the meaning should be clear from the context.

This is the idealized model of $V1$ proposed by Hubel, leading to the fundamental geometrical concept of a *fibration* or *fibre bundle*.

4.3.3 Fibre Bundles

The key idea of a fibration or fibre bundle was developed by mathematicians, and more recently by theoretical physicists, for deep reasons. The problem was to associate with each point in some base space M an entity of a certain type F , such as a scalar, a vector, a covector, a tensor, an exterior form, a direction, a phase, or a quantum number, which depends smoothly on this point. An obvious solution for modelling such a *field* on M would be to use maps $\varphi : M \rightarrow F$. But in many cases, we must take into account the fact that the *whole set* F of possible values of such maps φ is associated with each point of M . To take a concrete technical example, in each pixel of a computer screen, all grey levels (1 byte) or colours RGB (3 bytes) are represented.

Intuitively, a fibre bundle comprises a base space M (a differentiable manifold) and copies of a manifold F called the fibre ‘above’ each point of M . Globally, the space E of the fibre bundle, with the fibres glued together, is not necessarily a Cartesian product $M \times F$. It results from gluing together several Cartesian products $U_i \times F$ defined on local open domains U_i of M . Up to now, this *local triviality* has only been of interest to geometers in cases where the base space M is not a globally trivial space like \mathbb{R}^n but rather a manifold that may not be simply connected and thus may have a non-trivial homology. In our case, the fibre bundles are globally trivial, but their local structure is imposed by neurophysiology (the receptive fields).

By definition, a fibre bundle is a 4-uple (E, M, F, π) such that:

1. E , M , and F are differentiable manifolds, called the total space, the base space, and the fibre of the fibre bundle, respectively.
2. $\pi : E \rightarrow M$ is a surjective differentiable map called the projection of the fibre bundle.
3. The inverse images $E_x = \pi^{-1}(x)$ ($x \in M$) are isomorphic to F and $E_x \cong F$ is called the fibre at the point x (see Fig. 4.19).
4. For any point $x \in M$, there is a neighbourhood U of x such that $\pi^{-1}(U)$ is diffeomorphic to the Cartesian product $U \times F$ equipped with its canonical projection $\pi : U \times F \rightarrow U$, $(x, q) \mapsto x$ (local triviality, see Fig. 4.20).

A *section* of a fibre bundle is a differentiable map that lifts the projection π , associating an element of the fibre E_x with each point x of the base space M . If $s : M \rightarrow E$ is a section, we thus have $\pi \circ s = Id_M$ (see Fig. 4.21). Sections can be defined simply locally on open sets $U \subset M$. In the case of a globally trivial fibre bundle $\pi : E = M \times F \rightarrow M$, a section above U is nothing other than a map $s : U \rightarrow F$.

Fig. 4.19 Fibre bundle with base space M , fibre F , and total space E . Above each point x of M , the fibre $\pi^{-1}(x) = E_x$ is isomorphic to F






Fig. 4.20 Local triviality of a fibre bundle. For any point x in M , there is a neighbourhood U of x whose inverse image $\pi^{-1}(U) = E_U$ is the direct product $U \times F$, where π is the projection onto the first factor



4.3.4 V1 as a Geometric Fibre Bundle

If we idealize the functional architecture of the retino-geniculo-cortical pathway mathematically on the mesoscopic scale, the retinotopic and hypercolumn structures of V1 can be naturally modelled by the bundle $\pi : V \rightarrow R$ associating a copy P_a of the space P of directions in the plane with each point a of the retina R . The total space V of these copies P_a of P , glued together by local coordinate changes in the base space R , is a fibre bundle. We shall see in the next section that this is in fact

Fig. 4.21 A section of a fibre bundle defined on an open set U of M associates a value $s(x)$ in the fibre E_x above x with each point x of U



the *contact fibre bundle* CR of R , in other words, the projectivization of the tangent bundle TR of R . The points of V , that is, the pairs (a, p) comprising a point a of R and an orientation p at a , are the contact elements of R , already mentioned several times (see Fig. 4.22).




Fig. 4.22 The fibre bundle $E = V$ with base space the retinal plane $M = R$ (represented by a line to simplify) and fibre the projective line \mathbb{P}^1 of directions in M . **a** Elements of the fibre above a are represented by rotating horizontal line segments viewed in perspective. **b** Elements of the fibre above a are represented as points (the coordinate in the fibre encodes the angle θ of the direction p)

Through this functional architecture, a (discretized) copy of the orientation space P is associated retinotopically with each retinal position a . There thus exists a neural implementation of the structure $\pi : R \times P \rightarrow R$, the set of feed-forward projections (in the neurophysiological sense) of the retino-geniculo-cortical pathways implementing the projection π (in the geometric sense).

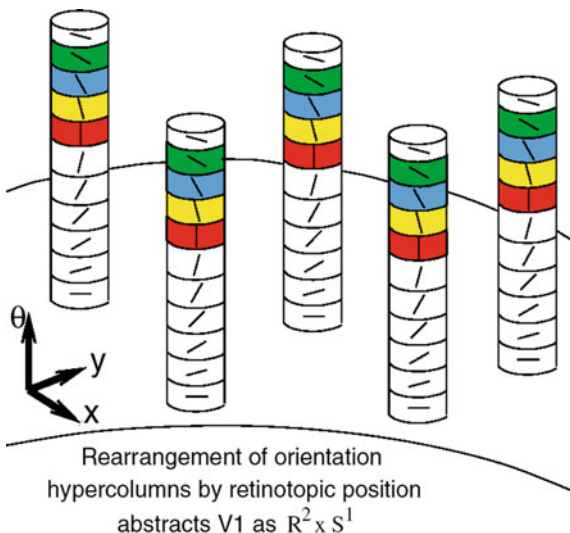
4.3.5 V1 as a 1-jet Fibre Bundle

Few specialists grasp the fundamental importance of an abstract geometrical model with dimension equal to the number of degrees of freedom of the given empirical structure, in this case 3. This theoretical requirement leads us to introduce the fibre bundle $\pi : V = R \times P \rightarrow R$, and this can be interpreted in different ways:

- (i) As $\pi : R \times \mathbb{P}^1 \rightarrow R$ if the fibre P is identified with the projective line of orientations in the plane specified modulo π (P is topologically a circle).
- (ii) As $\pi : R \times \mathbb{S}^1 \rightarrow R$, where \mathbb{S}^1 is the unit circle, if the fibre P is described by an angular coordinate θ modulo 2π , taking into account sense as well as orientation.
- (iii) As $\pi : R \times \mathbb{R} \rightarrow R$ if the fibre P is described by the tangent $p = \tan \theta$.

An excellent example of interpretation (i) is provided by the work of Steve Zucker and Ohad Ben-Shahar. In their paper [28], we find Fig. 4.23.

Fig. 4.23 Fibre bundle $\pi : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$ introduced by Ben-Shahar and Zucker [28] to model the abstract structure of the orientation hypercolumns of V1



However, in interpretation (iii), V1 is identified with what is known as the *1-jet bundle* of curves in R . The idea of jets generalizes the classical notion of Taylor⁴ expansion, thereby giving it an intrinsic geometric meaning, i.e., independent of coordinates. Suppose that, in a certain coordinate system (x, y) of R , a smooth curve γ is the graph $\{x, f(x)\}$ of a real-valued function f on \mathbb{R} . The first-order jet of f at x , denoted by $j^1 f(x)$, is characterized by 3 arguments: the coordinate x , the value $y = f(x)$ of f at x , and the value $p = f'(x)$ of the derivative of f at x , i.e., the slope of the tangent to the graph of f at the point $a = (x, f(x))$ of R . So if we identify R with a domain of \mathbb{R}^2 , a 1-jet is just a pair $c = (a, p)$, i.e. a contact element. Conversely, with each contact element $c = (a, p)$, we may associate the set of regular functions f whose graph is tangent to c at a . Then, $J^1 R$ will denote the fibre bundle with base space R , usually denoted $J^1(\mathbb{R}, \mathbb{R})$, of 1-jets of curves in R .

These 1-jets are feature detectors specialized in the detection of tangents. The fact that the V1 area can be ideally identified with $J^1 R$ in the case of ‘simple neurons’ explains why it is functionally essential for *contour integrations*. In the 2D manifold R , the determination of the direction p tangent to a contour γ at a point a requires one to compare the values of γ in the neighbourhood of this point. However, the neural system can access this *local* geometric information directly in the *pointlike format* of a simple numerical value, provided that it calculates in the jet space V1, a three-dimensional space. This saves having to carry out a local computation that would be costly in terms of wiring.

Jan Koenderink in [29] stressed the importance of the jet concept for theories of vision. Without jets, it is hard to understand how the visual system could extract geometrical features like the tangent or the curvature of a curve at a point:

Geometrical features become multilocal objects, i.e., in order to compute [boundary or curvature] the processor would have to look at different positions simultaneously, whereas in the case of jets it could establish a format that provides the information by addressing a single location. Routines accessing a single location may aptly be called *point processors*, those accessing multiple locations *array processors*. The difference is crucial in the sense that point processors need no geometrical expertise at all, whereas array processors do (e.g. they have to know the environment or neighbours of a given location). Koenderink [29, p. 374]

This is indeed the key point: V1 must carry out geometrical tasks, but without being availed of any ‘geometrical know-how’. However, this is possible *only if the geometry is neurally hard-wired* and, as we shall see, this is precisely the role of a functional architecture.

We can already say to some extent why functional architectures are so important. In the brain, there are what are essentially *temporal correlations*. So how can temporal coherence transform itself into spatial morphologies, Gestalten, or patterns? As we shall see in detail in the second volume, the key phenomenon is *synchronization*. So how can synchronizations be transformed into patterns? To do this, we require highly constrained connections and hence a functional architecture.

⁴Named after Brook Taylor.

The deepest point, and without doubt the most difficult to understand for a non-mathematician, is that the introduction of a further independent variable p provides a way to replace ‘geometrical know-how’ by a functional architecture. Let us say once again that neurons are point processors (on the scale specified by the size of their RFs) and as a consequence can only measure a quantity at a point. But in order to do differential geometry with point processors rather than local processors, the only way is to add supplementary variables evaluating partial derivatives of appropriate degree. This is why Jan Koenderink stresses the fact that the (hyper)columns implement jet spaces neurobiologically:

The modules (like ‘cortical columns’ in the physiological domain) of the sensorium are local approximations (*Nth order jets*) of the retinal illuminance that can be addressed as a *single datum* by the point processors. [29, p. 374]

Let us specify how $V1$ can be interpreted as a neural realization of the space of 1-jets of curves in R . At each point a in M , consider not the tangent space T_aR , but the set of its hyperplanes (its vector subspaces of codimension 1 and hence its straight lines), denoted by C_aR . This is isomorphic to the projective space \mathbb{P}^1 . The total space in which these fibres are glued together is called the *contact fibre bundle* of R , denoted by CR .⁵

CR is almost the space of 1-jets $J^1(\mathbb{R}, \mathbb{R}) = J^1R$ associated with the choice of coordinates (x, y) . To see this, we only have to interpret the coordinate on its fibres C_aR in terms of T_aR . Equipping R with local coordinates (x, y) at a and the tangent plane T_aR with the natural coordinates (ξ, η) associated with it in the basis $(\partial/\partial x, \partial/\partial y)$, then on an open set not containing the ‘vertical’ straight line $\xi = 0$, a local coordinate on C_aR is $p = \eta/\xi$, and in the neighbourhood of $\xi = 0$, we may take the coordinate $p = \xi/\eta$. An element c of CR is thus attributed the coordinates $(x, y, p) = (a, p)$. It is straightforward to check that changes of chart associated with these natural coordinates on CR are diffeomorphisms. Hence, CR is a 3D differentiable manifold isomorphic to $V = R \times \mathbb{P}^1$.

The difference between CR and J^1R is that the fibre of J^1R is not the whole of \mathbb{P}^1 but the \mathbb{R} given by the values of $\tan \theta$, the angle θ being specified modulo π , measured given a choice of x -axis and never taken equal to $\pi/2$. To obtain the fibre \mathbb{P}^1 of CR , we must compactify \mathbb{R} by adding a point at infinity. CR is the compactification at infinity of J^1R , and its fibre corresponds to that of J^1R via the stereographic projection $\mathbb{P}^1 \rightarrow \mathbb{R}, \theta \rightarrow \tan \theta$. In the language of algebraic geometry, J^1R is the open affine subset of CR complementary to the section at infinity, the choice of this section corresponding to the choice of an x -axis in the plane R .

Specifying a section s of the fibre bundle $\pi : V = CR \rightarrow R$ above a subset U of R means associating an element $s(a)$ of the fibre C_aR above a , i.e., an orientation, with each point a of U . The sections are thus fields of pairs $(a, p) = (\text{position}, \text{orientation})$. A fundamental special case is the sections restricted to differentiable

⁵We could distinguish between the retinal plane R and the cortical layer M (the base space of V) to which it projects. However, to simplify, we shall not do so, considering the retinotopic map $\chi : R \rightarrow M$ as the identity.

curves γ in R . They are obtained by choosing, above each point a of γ , the orientation of the *tangent* to γ at a .

4.3.6 Legendrian Lifts

By the implicit function theorem, any differentiable curve γ is locally the graph of a function $f(x)$, except at points with a vertical tangent. Let $j^1\gamma(a)$ be its 1-jet. It can be identified with the 1-jet $j^1f(x)$ comprising the abscissa x of a , the value $y = f(x)$ of f at x , and the value $p = f'(x)$ of its derivative there. Rather than considering the *plane* equipped with coordinates (x, y) and calculating $y' = dy/dx$ —as we have seen, this requires knowing not only the value $y = f(x)$ of f at x , but also the values of f in a neighbourhood of x —we work in the space with three dimensions spanned by coordinates (x, y, p) , imposing the constraint $y' = p$. This very profound idea goes back to William Hamilton who, by introducing the conjugate momenta p_i of the position variables q_i of a mechanical system as independent variables, replaced the Lagrangian formulation of mechanics by what is now known as the Hamiltonian formulation.

Quite generally, if γ is a differentiable curve in R , parametrized by the equations $x(s)$ and $y(s)$, the 1-jet $j^1\gamma(a(s))$ of γ at $a(s) = (x(s), y(s))$ is the contact element $(a(s), p(s))$, where $p(s) = y'(s)/x'(s)$ is the slope of the tangent to γ at $a(s)$.

The image of $j^1\gamma$ is called the *Legendrian lift* of γ . The Legendrian lifts of curves γ in R no longer represent these curves as sets of points in R , but in a dual sense, i.e., what is called *projective duality*, as envelopes of their tangents. It is remarkable that biological evolution should have created two neurophysiological structures, the retina and the V1 area, in order to implement projective duality for contours.

4.3.7 Integrability Condition

We can thus associate a Legendrian lift $\Gamma = j^1\gamma$ with any smooth curve γ in R . However, these lifts Γ are rather specific so we need to characterize them carefully. Indeed, let $\Gamma = v(s) = (a(s), p(s)) = (x(s), y(s), p(s))$ be any skew curve in V . Its projection $a(s) = (x(s), y(s))$ is indeed a curve γ in R . However, there is no reason why Γ should be the Legendrian lift $j^1\gamma$ of its projection γ . This is only the case if $p(s) = p_{a(s)}(s)$. Put another way, a curve Γ defined locally by equations $y = f(x)$, $p = g(x)$ is the lift of a curve γ in V if and only if $g(x) = f'(x)$, i.e. $p = y'$. This condition is called an *integrability condition*. It is crucial, and we shall return to this at length.

4.3.8 $SE(2)$ Invariance of 1-jets

To simplify, set $R = \mathbb{R}^2$. The structure of $J^1\mathbb{R}^2$ is invariant under the action of the Euclidean group $SE(2) = \mathbb{R}^2 \rtimes SO(2)$ of isometries of the plane which is the semi-direct product \rtimes of the group of translations \mathbb{R}^2 and the group of rotations $SO(2)$. Generally speaking, if G is a group and if H is a subgroup of G which operates on another, normal subgroup N of G , then G is the semi-direct product $N \rtimes H$ if its product law \circ is $(n', h') \circ (n, h) = (n'h'(n), h'h)$.

Let (q, r_θ) be an element of $SE(2)$, where q is a point in \mathbb{R}^2 and r_θ the rotation through angle θ . Then, (q, r_θ) acts on the points a of \mathbb{R}^2 according to

$$(q, r_\theta)(a) = q + r_\theta(a) .$$

If (q, r_θ) and (s, r_φ) are 2 elements of $SE(2)$, their (non-commutative) product is given by

$$(s, r_\varphi) \circ (q, r_\theta) = (s + r_\varphi(q), r_{\varphi+\theta}) .$$

The product is non-commutative because $(q, r_\theta) \circ (s, r_\varphi) = (q + r_\theta(s), r_{\theta+\varphi})$. Naturally, $r_{\varphi+\theta} = r_{\theta+\varphi}$, but $s + r_\varphi(q) \neq q + r_\theta(s)$ (see Fig. 4.24).

The rotation r_θ acts on the fibre bundle $J^1\mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$r_\theta(a, \psi) = (r_\theta(a), \psi + \theta) ,$$

where ψ is the angular coordinate in the fibre. This further action on the fibre ensures that the alignment of the preferred orientations is itself also $SE(2)$ -invariant (see




Fig. 4.24 Non-commutativity of the Euclidean group $SE(2)$. The vectors $s + r_\varphi(q)$ and $q + r_\theta(s)$ are completely different

Fig. 4.25 $SE(2)$ invariance of the jet space




Fig. 4.25). This makes the action interesting from the point of view of group representation theory.

Regarding the group $G = SE(2)$, it is interesting to note a general phenomenon that is quite fundamental here. By defining $SE(2)$ as the group of isometries in the plane \mathbb{R}^2 , we assume that \mathbb{R}^2 is given and we consider G only as a consequence. But this approach can be reversed. Indeed, \mathbb{R}^2 and the action of G can be reconstructed from the group structure of G . For if $H = SO(2) \simeq \mathbb{S}^1$ is the compact commutative subgroup of rotations around 0, its conjugates gHg^{-1} for $g = (q, r_\theta)$ in G give the rotation subgroups about different points q . We may then identify \mathbb{R}^2 with the quotient G/H and the action of G on \mathbb{R}^2 with the quotient of the action of G on itself.

4.3.9 Generalizing the Model

In accordance with Hubel’s idea of engrafted variables, this model can be extended to other characteristic variables of the visual signal which are represented in the hypercolumns. To do this, we consider spaces in which the new variables vary, viz. the interval $[0, 1]$ for the level of ocular dominance, the projective plane \mathbb{P}^2 for the colour, the circle \mathbb{S}^1 for directions of motion (the reader is referred to, e.g. Zhang and Wu [30] or Weliky et al. [31]). Then, $V1$ will be modelled by a fibre bundle with base space R and fibre the Cartesian product of the spaces of secondary variables.

Such generalizations might also lead us to consider k -jets of order $k > 1$, as we shall see in Sect. 5.6 of Chap. 5 when we discuss *curvature detectors*.

4.3.10 Neurophysiology and Its Geometrical Idealization

There are significant differences between the neurophysiological data and the geometrical idealization of the fibre bundle. Here, we mention three of these:

- (i) To begin with, the RFs introduce a resolution scale so we require a *multiscale* theory of fibre bundles. Moreover, the RFs are adaptive and modulated by stimuli.
- (ii) Then, there is a significant *redundancy* in the columns. Indeed, a ‘point’ (a, p) of the fibre bundle actually corresponds to a whole column. As we have seen, this so-called population coding is essential for adaptive capacity and refining the resolution. Among other things, it allows *oscillatory responses* (the columns can become oscillators through Hopf bifurcation) and hence synchronizing effects through phase-locking.
- (iii) Finally, there is a fundamental *dimensional* constraint. From an abstract point of view, the fibre bundle $\pi : R \times P \rightarrow R$ has dimension 3, i.e. 2 degrees of freedom for the retinal position $a = (x, y)$ and 1 degree of freedom for the orientation p , whereas the cortical layers are essentially 2D. There is therefore a problem of ‘dimensional collapse’. The visual systems produced by evolution that interest us here have solved this problem through the fascinating structure of ‘pinwheels’ in the V1 area, and these have received much attention since the pioneering work of Tobias Bonhöffner, Gary Blasdel, and Amiram Grinvald.

4.4 The Pinwheel Structure of V1

The model of V1 as a 1-jet space is a *continuous* model which is in fact the limit of a *discrete* model defined on a lattice in R . We shall return at length to the idea of continuous models in the second volume. But in the present chapter, we shall begin by gathering together some experimental data regarding discrete models.

4.4.1 Observation of Pinwheels

4.4.1.1 Functional Orientation Maps

Fundamental experiments made possible by recent progress in brain imaging have shown that the hypercolumns are arranged geometrically in little wheels called *pinwheels*. The observed cortical layer is covered with a lattice of singular points (about 1 200 μm apart in the cat and about 600 μm apart in primates): the centres of local pinwheels which join up to form a global structure. The imaging method used here was developed at the beginning of the 1990s, by Bonhöffner and Grinvald [32] among others, and it is referred to as *in vivo* optical imaging based on activity-dependent

intrinsic signals. It exploits the fact that the metabolic activity of the nerve tissue changes its optical properties, whence it can acquire images of the activity of the surface cortical layers. More precisely, it exploits the differential absorption of oxyhemoglobin or deoxyhemoglobin, or of dyes whose fluorescence indicates local depolarization of neurons.

The method has profoundly transformed the observation of neural activity. Previously, there were only multielectrode methods for recording the activity of a few individual neurons,⁶ or post-mortem visualization of cortical activity using 2-deoxyglucose maps. Both methods were drastically inadequate. In vivo optical imaging made it possible to visualize functional organization; in other words, it made the cerebral black box ‘transparent’. As pointed out by Ohki and Reid [33]:

Optical imaging revolutionized the study of functional architecture by showing the overall geometry of functional maps.

However, the experimental challenge is enormous. First, the signal-to-noise ratio is very small, viz. $\sim 10^{-3}$, because the background noise is huge. The intrinsic signal comes from the hemodynamic properties of the cortical tissue and thus from an area of vascular metabolism greater than the area of the activated neurons. Then, millions of neurons are connected together, each with hundreds or even thousands of synapses, and this imaging operates on a *mesoscale* defined by averaging (in Sect. 4.7.3, we shall discuss the methods of two-photon confocal microscopy on the microlevel). One ‘neuron’ is in fact a cortical position at which a bunch of neurons is located. Moreover, the in vivo optical imaging methods with their good mesoscopic spatial resolution (50 μm) actually have a rather poor temporal resolution, because the intrinsic signals are slow, so they can only analyze slow intrinsic changes in the optical properties of the cortical layer. Other methods are required to visualize the cortical dynamics, such as voltage-sensitive dyes, which colour the active cells. The dye molecules bind to the neural membranes and act:

[...] as molecular transducers that transform changes in membrane potential into optical signals. [34]

As the changes in the membrane potential are correlated with millisecond changes in the absorbed or emitted fluorescence, we may thus obtain temporal resolutions of millisecond order.

For a discussion of these new techniques for investigating mammalian brains, the reader is referred to the review [34], presented by their inventor Amiram Grinvald. Figure 4.26 shows the accuracy that can be reached in the topography of the V1 area of the owl monkey (*douroucoulis*). Such maps became available from the end of the 1980s. Figure 4.27 produced by Blasdel and Salama in 1986 already shows pinwheels in the V1 area of the macaque.

⁶The transition from recordings of a few isolated neurons to a visualization of the overall activity of a piece of brain area is analogous to the leap forward in meteorology when recordings made by weather balloons were replaced by satellite imaging. No need for further comment.

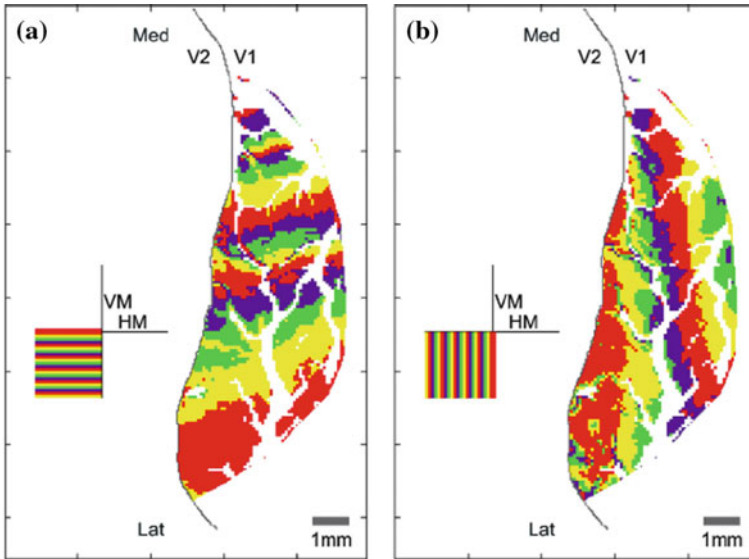


Fig. 4.26 Topography of the V1 area of the New World owl monkey as measured by in vivo optical imaging. A target made of two gratings, one horizontal and the other vertical, is projected onto V1. *VM* vertical meridian, *HM* horizontal meridian. From Grinvald [34]




Fig. 4.27 Pinwheels in the V1 area of the macaque. From Blasdel and Salama [35]

According to a standard experimental protocol, the animal is shown high-contrast gratings between 20 and 80 times. These are made of black bands (e.g. 6.25°) alternating with white bands (e.g. 1.25°), with several different orientations (e.g. 8) and an angular speed of say $22.5^\circ/s$. A cranial window is opened above V1, and the cortex is illuminated with orange light (605 nm). The orientation maps have very low amplitude relative to the light intensity of the recorded cortical images. We thus subtract the average intensity of the responses for all orientations, which is known as the cocktail blank. Then, depending on the orientation of the gratings, differential absorption patterns are observed, due to local spatial non-uniformities in the ratio of deoxyhemoglobin to oxyhemoglobin. We subsequently sum the images of V1 activity obtained for different gratings and construct differential maps. These are normalized by dividing the relative deviations from the average of each pixel by the global average deviation, and low-frequency noise is also eliminated.

We thus obtain *functional maps* like those in Figs. 4.28 and 4.29 obtained by Crair et al. [36] and Bosking et al. [37], respectively, which concern the 2 and 3 layers of a tree shrew (tupaia): LGN \rightarrow layer 4 \rightarrow (strictly feed-forward) \rightarrow layers 2 and 3. (The tree shrew looks like a primate but with the difference that there is




Fig. 4.28 Method for producing an orientation map of V1. Preferred orientations are conventionally colour-coded. Taken from Crair et al. [36]




Fig. 4.29 V1 of a tree shrew (tupaia). The different orientations are colour-coded. *Right* Zoom on examples of regular points and singular points of opposite chirality, as explained in Sect. 4.4.1.2. From Bosking et al. [37]

no fovea/periphery distinction, and this simplifies observation. For a survey of the functional organization of his visual cortex, see Fitzpatrick [38] and Lund et al.

[39].) The orientations are colour-coded, and the iso-orientation lines are thus the monochromatic lines.

As emphasized by Yves Frégnac et al. [40], these mesoscale optical functional imaging methods can be supplemented by microscale recordings of individual neurons, either extracellular and simultaneous recordings of several cells using multielectrodes able to measure correlations in the activities of the various neurons involved, or intracellular recordings that can measure both spiking activity and sub-threshold activity, and by inverse analysis can be used to reconstruct the afferent network of the given neuron (functional synaptic imaging). Figure 4.30 from [40] shows the relationship between these three methods schematically.

Optical imaging data helps us to understand what is meant by the population coding of a stimulus. Figure 4.31, which we shall consider again in Sect. 4.9.1, shows the neurons in V1 which are activated by a long vertical bar. We see that it is a rather thick band, also very patchy, and nothing like a line. As pointed out by Ulf Eysel [41, p. 641]:

A continuous line across the whole visual field would be cortically depicted in a patchy discontinuous fashion.

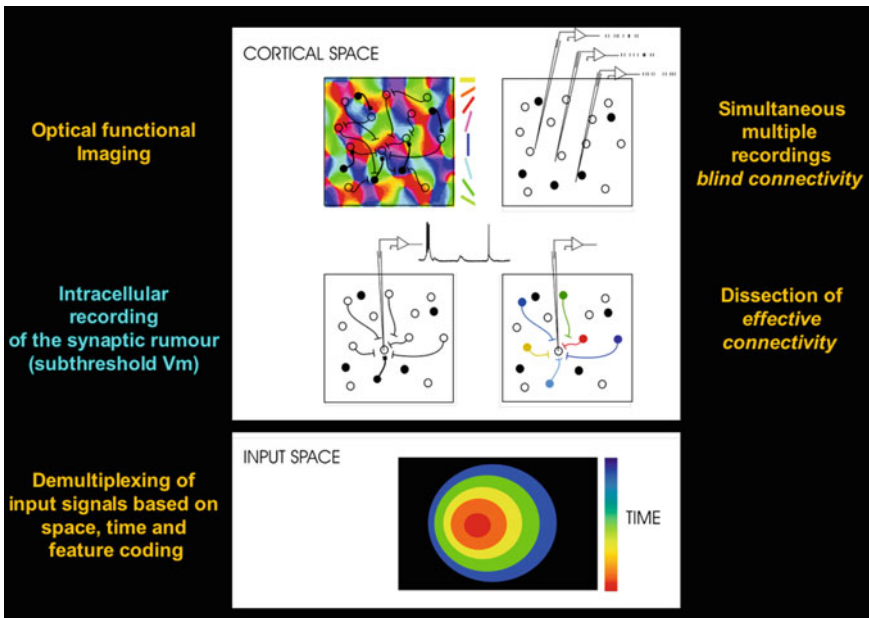


Fig. 4.30 Three methods for microlevel recording of the activity of individual neurons. Simultaneous extracellular recording of several cells, intracellular recording, and reconstruction of the afferent network of a neuron by functional synaptic imaging. Taken from Frégnac et al. [40]




Fig. 4.31 Neurons of V1 activated by a long vertical bar. Note that we obtain a thick and rather patchy band. From [42]

4.4.1.2 Periodicity, Retinotopy, and Singularities

We return here to Fig. 4.29. Note to begin with that the lattice has a kind of *characteristic length*—a mesh or *wave number*. This periodicity can be measured accurately by taking the map of the pinwheels, translating it by $t = (u, v)$, and calculating the correlation between the two maps. The autocorrelation obviously has a primary peak at $t = (0, 0)$ (the map correlates perfectly with itself) and secondary peaks giving the periodicity. Figure 4.32 is taken from the paper [43] by McLoughlin and Schiessl and gives the example of the marmoset monkey.

We shall see in Sect. 4.6.6 that, under a Fourier transform, an orientation map can be interpreted as a superposition of plane wave solutions of the Helmholtz equation and that the mesh of the pinwheel lattice comes from the fact that these plane waves have almost the same wave number. More precisely, let $\Phi(z) = e^{i\varphi(z)}$ be the orientation map [or $\Phi(z) = e^{i2\varphi(z)}$, if we prefer to work modulo 2π and not take into account the chirality of the pinwheels]. The autocorrelation function of $\Phi(z)$ is given by $C(u) = \int_{\mathbb{C}} \Phi^*(z) \Phi(z - u) dz$, where $\Phi^*(z)$ is the complex conjugate of $\Phi(z)$. Isotropy implies that $C(u) = C(r)$ depends only on the modulus $r = |u|$ of u . So let us consider the Fourier transform $P(k)$ of $C(r)$ (known as the power spectrum), viz.

$$P(k) = \int_{\mathbb{R}^+} e^{-ikr} C(r) dr .$$

$P(k)$ is maximal on a ring of radius $k_0 = 2\pi/\Lambda_0$, and Λ_0 specifies the mesh of the pinwheel lattice. This structure in $P(k)$ arises in the following way:




Fig. 4.32 Autocorrelation with its translates of the pinwheel map of the marmoset monkey V1 area. It is colour-coded from 0 to 1. At the *centre*, the *red* point means that the autocorrelation is 1 for $t = (0, 0)$. The first *dark blue* ring around the centre corresponds to the distance at which the pinwheels of opposite chirality are superposed. Then, the first *light blue* ring corresponds to the basic period. The average distance between pinwheels is $575 \mu\text{m}$ in this example. Taken from [43]

- (i) Local triviality outside singularities indicates a small local variation in orientations, and this low local correlation implies that there are no Fourier components with high spatial frequencies $k \gg k_0$, because these would induce large local variations in the orientations everywhere.
- (ii) Uniformity, i.e. the fact that all orientations are represented in equal amounts, implies that there are no Fourier components with low spatial frequencies $k \ll k_0$, because these would induce long-range non-uniformities.

Figure 4.33, produced by Niebur and Wörgötter [44], shows an orientation map for a macaque (area 18) and the power spectrum concentrated on a ring of average radius $k_0 = 2\pi/\Lambda_0$. As already mentioned, we shall return in Sect. 4.6.6 to the limiting case in which the power spectrum is totally concentrated on a *circle* of radius k_0 .

Note that in Fig. 4.29, the orientations (represented by colours) are distributed in a globally homogeneous way. If the sectors of different colours seem to be distributed inhomogeneously in the pinwheels, this is because of the perceptual structure of the continuous wheel of colours; in contrast with what happens for discretized wheels, the names of the colours correspond to sectors of a different angular width (see Fig. 4.34).

Note also that there are three classes of points (they are represented in the idealized diagram of Fig. 4.35):




Fig. 4.33 Periodicity of an orientation map (*right*). The power spectrum (Fourier transform of the autocorrelation function of the map) is concentrated on a ring of average radius $k_0 = 2\pi/\Lambda_0$, and Λ_0 specifies the periodicity of the map. From Niebur and Wörgötter [44]

Fig. 4.34 In the standard continuous colour wheel, the names of the colours correspond in fact to angular sectors of different angles. The primary colours in additive colour mixing, viz. *blue*, *green*, and *red*, occupy large sectors, whereas their complementary colours, which are the primary colours in subtractive colour mixing, viz. *yellow*, *magenta*, *cyan*, each occupy narrow sectors





Fig. 4.35 Idealized ‘crystalline’ pinwheel structure on a regular square lattice of singularities. There are three classes of points: regular, singular (pinwheel), and saddle points



- (i) Regular points where the orientation field is locally trivial in the sense that the iso-orientation lines are approximately parallel.
- (ii) Singular points at the centres of the pinwheels where all the orientations converge. They have a positive or negative chirality depending on whether, when we move around the centre, the orientations rotate in the same sense or not. They have opposite chiralities when adjacent.
- (iii) Saddle points at the centres of the lattice cells, points where the iso-orientation lines bifurcate: two neighbouring iso-orientation lines leave the same singular point but end up at two opposite singular points.

Xu et al. [45] studied analogous structures for the prosimian primate⁷ known as the bush baby (galago), which is a monkey ancestor with an analogous V1 to the monkey primates, but very different V2. Figure 4.36, taken from [5], shows the isochromatic lines more clearly, these being the iso-orientation field lines of V1. Figure 4.37, produced by Hongbo Yu et al. [46], shows the receptive fields corresponding to positions along a line segment in V1. We see the superposition of receptive fields for neighbouring positions in a pinwheel and their slow displacement when we move from one pinwheel to another. This explains the pioneering experiments by Hubel and Wiesel.

4.4.2 *Limitations of This Analysis*

As noted by Amit Basole et al. [47], optical imaging methods are very powerful, but the result must be interpreted cautiously because they depend among other things on the type of inputs used. For example, they often use gratings which cross the receptive fields (RFs), so we can only test the motion selectivity of the neurons in the orthogonal direction (the aperture problem). If the input patterns are changed by using shorter segments than the RFs in order to be able to test other directions of motion, we obtain the same pinwheel structure for the orthogonal motions, but a *distortion* of this structure for oblique motions.

4.4.3 *Functional Maps as Fields*

It is interesting to note that the pinwheel lattices with their iso-orientation lines look like ‘field models’. The singularities c_i with their chiralities are analogous to positive and negative charges producing field lines in \mathbb{R}^2 . This kind of field structure was introduced early on by Valentino Braitenberg, in fact in 1979, in a paper entitled *Geometry of orientation columns in the visual cortex* [23], well before the development of in vivo optical imaging methods. Starting with Hubel and Wiesel’s results

⁷There are two classes of primates: on the one hand, monkeys and humans, and on the other, the prosimians.




Fig. 4.36 The isochromatic *lines* are the iso-orientation lines of V1. From [5]

on monkeys and cats in 1962, and their discoveries (i) that the preferred orientation of neurons depends smoothly on the tangential penetrations of the electrodes, (ii) that there is chirality, and (iii) that the latter can reverse along a penetration, the Braitenbergs tried to build orientation fields. Through quite remarkable abduction, they came to the following conclusion:




Fig. 4.37 Receptive fields of simple V1 neurons at different positions along a line segment. We see the superposition of the receptive fields for neighbouring positions in a pinwheel and their slow displacement when we move from one pinwheel to another. From Yu et al. [46]


We believe that the most natural explanation of the facts observed would be in terms of orientations arranged with circular symmetry around centres, either radially or along concentric circles.

They clearly anticipated by clever abduction the pinwheel structure discovered later experimentally. Figure 4.38 of [23] shows how a straight line penetration L encounters the pinwheels and can reverse the chirality when the singularities lie on either side of L .

Following the work by the Braitenbergs, other specialists like William Baxter and Bruce Dow carried out further investigations of these inferences. In their paper [48] on pinwheels in the macaque, they explicitly constructed the flow of the orientation field and introduced singularities of a different chirality and different topological index, i.e., singularities of index ± 1 giving sources and sinks and those of index $\pm 1/2$ giving end points and triple points, which we shall discuss shortly. Figure 4.39 gives an example.

It would be interesting to reconstruct more or less explicit models of orientation fields that can be observed empirically in the pinwheels of different species. To do this, we could use the relationship between the fields of straight lines provided by the orientation fields and vector fields carried by these straight lines. Away from singularities where the vectors vanish or diverge, a vector field automatically induces an orientation field, simply by taking the orientations of the vectors. Supposing that it is meaningful to consider the singularities with their chirality as positive and

Fig. 4.38 The Braitenbergs' theoretical reconstruction by abduction of the pinwheel structure of V1. We observe the orientations along the straight line tangential penetration L of an electrode. The pinwheels are reconstructed to explain the data, and in particular the reversal of chirality when the singularities lie on either side of L . From Braitenberg [23]



negative charges, we could then consider the field lines of an associated ‘physical’ field deriving from a potential.

Consider, for example, the field of Fig. 4.40 produced by Geoffrey Goodhill in [49]. An arrangement of charges ± 1 is constructed by applying the rule that two neighbouring singularities always have opposite charges, together with several supplementary terms to take into account edge effects. By inspection, we find that the ‘physical’ analogy is roughly correct. In Fig. 4.41, the field lines are the orange lines from the sources (+) towards the sinks (–), while the lines orthogonal to the field line are equipotentials, i.e. lines of equal potential, for the potential whose gradient gives the field.

To give another example, consider the fragment shown in Fig. 4.42 of the image in Fig. 4.97 that we will comment on in Sect. 4.9.3. Figure 4.43 shows the physical analogy with the field induced by the $+1$ and -1 charges located at singular points and with a certain number of equipotential curves which are by definition orthogonal to the field.

The equipotential line of mean level is a separatrix between the influence zones of charges $+1$ and -1 (see Fig. 4.44).

For a regular ‘crystalline’ pinwheel lattice with square unit cell like the one in Fig. 4.35 of Sect. 4.4, we obtain Figs. 4.45 and 4.46, whose idealized geometry is shown in Fig. 4.47.

In Figs. 4.48 and 4.49, we show how the ideal crystal model can be distorted.

With the notion of phase field, we shall discuss a deeper way to understand pinwheels as ‘physical fields’ in Sect. 4.6.




Fig. 4.39 Four examples of abstract 'crystal' models of pinwheels put forward by William Baxter and Bruce Dow. *White and black squares* represent singularities with topological indices $+1$ and -1 , respectively, while *white and black triangles* represent singularities with topological indices $+1/2$ and $-1/2$, respectively. Taken from [48]

4.4.4 Development of Pinwheels

There is some experimental data concerning the evolution of orientation maps and their pinwheels while the V1 area is developing. For example, in the mid-1990s, Barbara Chapman, Michael Stryker, and Tobias Bonhöffer [51] adapted the optical imaging method to carry out what they called *chronic* optical imaging of intrinsic signals. The aim was to study the emergence and development of the pinwheel structure. (They were working on the ferret, whose visual system is similar to the cat's.) They showed that, starting from a barely structured initial orientation map,



Fig. 4.40 Example of an orientation field reconstructed from neural network models, with its pinwheels and field lines. The *black lines* are the boundaries of ocular dominance domains, discussed in Sect. 4.10.1. From Goodhill [49]

Fig. 4.41 ‘Physical’ field with charges $+1$ and -1 distributed like the pinwheels in Fig. 4.40 (ordinate $-y$). Field lines are orange lines going from the sources ($+$) to the sinks ($-$). Lines orthogonal to the field lines are equipotentials, i.e. lines of equal potential, for the potential whose gradient gives the field. Plot drawn with *Mathematica* [50]




Fig. 4.42 Fragment of the image in Fig. 4.97



Fig. 4.43 Physical analogue of the field in Fig. 4.42 induced by +1 and -1 charges located at singular points. A certain number of equipotential curves orthogonal to the field are also shown




Fig. 4.44 If we take the mean level equipotential of the field in Fig. 4.43, we get a separatrix between the influence zones of charges $+1$ and -1 . The small white discs around the singular points are due to the fact that the divergences of the potential at these points are clipped




Fig. 4.45 Field of the regular ‘crystalline’ pinwheel lattice with square unit cell in Fig. 4.35 of Sect. 4.4




Fig. 4.46 Equipotential separating the influence zones of $+1$ and -1 charges of the field in Fig. 4.45




Fig. 4.47 $+/-$ dominance domains in the 'crystal' model idealized in Fig. 4.35




Fig. 4.48 Idealized crystalline geometric model of Fig. 4.47, interpreted using a physical analogy in which the pinwheel singularities generate a field. Gradient lines (corresponding to iso-orientation lines) and equipotentials are shown





Fig. 4.49 Distortion of the physical analogy in Fig. 4.48



an organized map gradually comes into being and eventually becomes remarkably stable.

We shall return to pinwheel morphogenesis and the associated learning processes in Sect. 4.7.1.2, but this time in terms of models.

4.4.5 Pinwheels and Evolution

Regarding the pinwheel structure of V1, as for any other biological structure, we may wonder how it evolved, and what exactly was being optimized as a result of some evolutionary pressure. We already discussed this question in Sect. 3.6 of Chap. 3, in relation to the shape of the receptive profiles of ganglion cells and LGN and V1 neurons.

4.4.5.1 Minimizing the Wiring

Alexei Koulov and Dmitri Chklovskii [52] investigated the hypothesis that the pinwheel structure *minimizes the total length L* of intracortical connections, while maintaining a good representation of the properties of the stimuli. To do this, they con-

sidered connection functions $c(\theta)$ giving the number of connections between a neuron \mathcal{A} with preferred orientation 0 and neurons with orientations $\theta \in [-\pi/2, \pi/2]$. The functions $c(\theta)$ considered have the form $c + G(\theta)$, where c is a constant favouring a number of connections independent of θ , and $G(\theta)$ is a Gaussian with mean zero favouring connections with neurons having the same orientation as \mathcal{A} . In the limit, $G(\theta)$ gives the Dirac distribution at 0. These two models c and $G(\theta)$ compete to minimize L , and when the Gaussian G is narrow enough, pinwheels appear.

Seen from this point of view, the advantage of the pinwheels is that a neuron does not have to go further than the nearest pinwheel in order to connect itself to a neuron of any given orientation. Figure 4.50a of [52] shows an example function $c(\theta)$ and Fig. 4.50b the pinwheel map obtained by minimizing L . The lattice contains $N = 50 \times 50 = 2500$ neurons. We consider the set \mathfrak{M} of 2500×2500 matrices $M = (M_{ij})_{i,j=1,\dots,N}$ of connections between neurons $i = (x_i, y_i)$ and $j = (x_j, y_j)$ which satisfy the constraint $c(\theta)$ with $M_{ij} = 1$ if the neurons i and j are connected and 0 otherwise. We have $L = \sum_{i,j=1}^N d_{ij} M_{ij}$, where d_{ij} is the distance between i and j , and we minimize L in \mathfrak{M} , which is difficult, using simulated annealing algorithms.⁸




Fig. 4.50 Minimizing the total length L of intracortical connections for the connection function $c(\theta)$ shown in (a). **b** Pinwheel map for a square lattice of $N = 50 \times 50 = 2500$ neurons. We consider 2500×2500 matrices of connections $M = (M_{ij})_{i,j=1,\dots,N}$ between neurons $i = (x_i, y_i)$ and $j = (x_j, y_j)$ which satisfy the constraint $c(\theta)$ with $M_{ij} = 1$ if neurons i and j are connected and 0 otherwise. We have $L = \sum_{i,j=1}^N d_{ij} M_{ij}$, where d_{ij} is the distance between i and j . We minimize L using simulated annealing algorithms. From Koulikov and Chklovskii [52]

⁸For an introduction to simulated annealing, see, for example, the Bourbaki lecture by Robert Azencott [53].

4.4.5.2 Column Structure and Sparse Representations

For their part, Karol Gregor, Arthur Szlam, and Yann LeCun [54] related the column and hypercolumn structure of V1 to the sparse code representation of statistical regularities in natural images. They used the model we discussed in Sect. 3.6.2 of Chap. 3. Natural images I with $P = p^2$ pixels form a subset $\Sigma \subset \mathbb{R}^P$ and are thus described by vectors $I = (I_r)_{r=1, \dots, P}$ in the canonical basis of \mathbb{R}^P . We seek representations $I = \sum_{i=1}^{i=N} s_i \varphi_i$ using filters (RPs of neurons) $\varphi_i \in \mathbb{R}^P, i = 1, \dots, N$. The latter constitute a ‘dictionary’ Φ which is a $P \times N$ matrix and we wish these representations to be sparse; that is, for each image I , only a limited number of filters respond, most of the s_i being zero. As we have seen, the basic idea is to minimize an energy of the form

$$E(I, s) = \frac{1}{2} \|I - \Phi s\|^2 + \lambda \sum_{i=1}^{i=N} |s_i| .$$

But we can also try to *structure* the lattice of filters by introducing, rather like Hyvriinen and Hoyer [55], a set of lateral neural connections between the filters and by penalizing the simultaneous activity of two neurons by means of *inhibitory connections*. If $U = \{(i_1, j_1), \dots, (i_k, j_k)\}$ is the set of these connections, we cannot have both $s_i \neq 0$ and $s_j \neq 0$ in the coding of an image I when $(i, j) \in U$.

In order to impose this constraint, we introduce a quadratic term $s^T W s$ into the energy E , where W is a matrix of weights for the inhibitory connections. If U is given, learning will reinforce W . However, if U has to be learnt by exposure to natural images, then learning will weaken the weights of the connections (i, j) between filters that are often activated together. The authors of [54] show that the lattice of filters gets structured into columns and hypercolumns. We shall return to this point in the second volume.

4.4.6 End Points and Triple Points

All the pinwheels observed experimentally have topological index ± 1 . However, singularities with topological index $\pm 1/2$ enter the scene when we look at the way the orientation itself varies. Figure 4.51 was produced by Shmuel. We have included the orientation field lines in the neighbourhood of two singularities of opposite chirality.

We see that dextrorotatory and levorotatory pinwheels are associated with the two types of generic singularity in the orientation fields in the plane. This is due to the fact that, when the ray rotates through an angle θ about the centre of the pinwheel, the associated orientation rotates through $\theta/2$. Hence, two diametrically opposite rays correspond to orthogonal orientations.




Fig. 4.51 Orientation maps and pinwheels in the V1 area of a tree shrew. Here, we see the relation between pinwheels (colours) and the preferred orientations. We have represented the orientation field lines in the neighbourhood of two singularities with opposite chirality. Adapted from Shmuel [56]. Copyright (2000) National Academy of Sciences, USA

If the orientation ψ_θ associated with the ray at angle θ is

$$\psi_\theta^+ = \alpha + \frac{\theta}{2} = \frac{\varphi_\theta^+}{2},$$

the two directions will be the same for $\psi_{\theta_0}^+ = \alpha + \theta_0/2 = \theta_0$, i.e. for $\theta_0 = 2\alpha$, and then $\varphi_{\theta_0}^+ = 4\alpha$. As α is specified modulo π , there is only one solution and we obtain the local model in Fig. 4.52 (an end point). If on the other hand the orientation ψ_θ associated with the ray at angle θ is

Fig. 4.52 End point singularity





Fig. 4.53 Triple point singularity



$$\psi_{\theta}^{-} = \alpha - \frac{\theta}{2} = \frac{\varphi_{\theta}^{-}}{2},$$

the two directions will be the same for $\psi_{\theta_0}^{-} = \alpha - \theta_0/2 = \theta_0$, i.e., for $\theta_0 = 2\alpha/3$, and then $\varphi_{\theta}^{+} = 4\alpha/3$. There are three solutions, and we obtain the local model shown in Fig. 4.53 (triple point).

It is very easy to calculate the integral curves of the orientation field in the neighbourhood of these singularities, for example for $\alpha = 0$. Adopting polar coordinates (ρ, θ) , we have $x = \rho \cos(\theta)$ and $y = \rho \sin(\theta)$, whence

$$\begin{cases} dx = \cos(\theta) d\rho - \rho \sin(\theta) d\theta, \\ dy = \sin(\theta) d\rho + \rho \cos(\theta) d\theta. \end{cases}$$

For end points, the constraint is $dy/dx = \tan(\theta/2)$, i.e.

$$\sin\left(\frac{\theta}{2}\right) dx - \cos\left(\frac{\theta}{2}\right) dy = 0,$$

or

$$\begin{aligned} \sin\left(\frac{\theta}{2}\right) [\cos(\theta) d\rho - \rho \sin(\theta) d\theta] - \cos\left(\frac{\theta}{2}\right) [\sin(\theta) d\rho + \rho \cos(\theta) d\theta] \\ = \left[\sin\left(\frac{\theta}{2}\right) \cos(\theta) - \cos\left(\frac{\theta}{2}\right) \sin(\theta) \right] d\rho \\ - r \left[\sin\left(\frac{\theta}{2}\right) \sin(\theta) + \cos\left(\frac{\theta}{2}\right) \cos(\theta) \right] d\theta = 0. \end{aligned}$$

But

$$\cos(\theta) = \cos^2\left(\frac{\theta}{2}\right) - \sin^2\left(\frac{\theta}{2}\right), \quad \sin(\theta) = 2 \cos\left(\frac{\theta}{2}\right) \sin\left(\frac{\theta}{2}\right),$$

and

$$\cos^2\left(\frac{\theta}{2}\right) + \sin^2\left(\frac{\theta}{2}\right) = 1,$$

so we obtain the differential equation

$$\sin\left(\frac{\theta}{2}\right) d\rho + \rho \cos\left(\frac{\theta}{2}\right) d\theta = 0,$$

and hence,

$$\sin\left(\frac{\theta}{2}\right) \frac{d\rho}{d\theta} + \rho \cos\left(\frac{\theta}{2}\right) = 0.$$

The solutions are

$$\rho = \frac{\rho_\pi}{\sin^2(\theta/2)} = 2 \frac{\rho_\pi}{1 - \cos(\theta)},$$

where the constant of integration ρ_π is the value $\rho(\pi)$. This is a parabola. Its axis of symmetry is the x -axis, it has vertical tangent at $(\rho = \rho_\pi, \theta = \pi)$, and its branches tend to infinity for $\theta \rightarrow 0$. Figure 4.54 shows the case $\rho_\pi = 1$.




Fig. 4.54 Integral curves of the orientation field in the neighbourhood of the pinwheels. *Left* End point singularity. *Right* Triple point singularity

For the triple points, the constraint is $dy/dx = -\tan(\theta/2)$ and an analogous calculation leads to

$$\sin\left(\frac{\theta}{2}\right) \left[3 \cos^2\left(\frac{\theta}{2}\right) - \sin^2\left(\frac{\theta}{2}\right) \right] d\rho \\ + \rho \cos\left(\frac{\theta}{2}\right) \left[3 \sin^2\left(\frac{\theta}{2}\right) - \cos^2\left(\frac{\theta}{2}\right) \right] d\theta = 0,$$

which has solutions

$$\rho = \frac{\rho_\pi}{\sin^{2/3}(3\theta/2)}.$$

The symmetry $\theta \rightarrow \theta + 2\pi/3$ is obvious, since it changes $3\theta/2$ to $3\theta/2 + \pi$ and thus leaves $\sin^{2/3}(3\theta/2)$ invariant. Figure 4.54 shows one of the trajectories in the case $\rho_\pi = 1$.

When $\alpha \neq 0$, the solutions become

$$\rho = \frac{\rho_{\pi+2\alpha}}{\sin^2\left(\frac{\theta}{2} - \alpha\right)}, \quad \rho = \frac{\rho_{(\pi+2\alpha)/3}}{\sin^{2/3}\left(\frac{3\theta}{2} - \alpha\right)},$$

respectively, and the orientation fields rotate through α . However, it should be noted that the fields of the $\varphi = 2\psi$ depend in a more subtle way on α . The singularity of the field φ^+ varies from a node with central symmetry ($\alpha = 0$), first to stable foci, then to a centre ($\alpha = \pi/4$), then to unstable foci, whereas the singularity of the field φ^- gives different kinds of saddle point. Figure 4.55 (obtained using Mathematica's StreamPlot) shows the field lines of φ^+ , φ^- , ψ^+ , and ψ^- for four values of the parameter α : 0, $\pi/8$, $\pi/4$, and $3\pi/8$. These different prototype models of field lines all correspond to a prototype pinwheel like the one in Fig. 4.56.

4.4.7 Distortions and Defects in the Neighbourhood of the V1/V2 Boundary

Given the chirality of the pinwheels, the most standard V1 structure is a distortion of a square lattice (the 'crystal' structure in Fig. 4.35). However, like any distortion of a regular lattice, it may contain defects inducing, for example triangular or pentagonal domains (see Fig. 4.57).

The distortion becomes significant at the boundary between V1 and V2 because, as has been shown by Ohki et al. [57] among others, pinwheels of the same chirality are aligned along this boundary (see Fig. 4.58). Their field lines and level sets can be studied and simulated in detail, as shown in Fig. 4.59. Distortion is manifested by the fact that the V1/V2 boundary resembles a 'cliff', as can be seen from the tightly bunched parallel level sets.




Fig. 4.55 Left to right columns Field lines of the fields φ^+ , φ^- , ψ^+ , ψ^- (see text) for the 4 values (top to bottom) 0 , $\pi/8$, $\pi/4$, and $3\pi/8$ of the parameter α

Fig. 4.56 Typical pinwheel with $\psi_\theta = \theta/2$




Fig. 4.57 Part of Fig. 4.36 showing the distortion relative to the 'crystal' lattice in Fig. 4.35 and the appearance of defects in the ordered structure. We have not indicated the connections to pinwheels outside the frame





Fig. 4.58 Arrangement of pinwheels near the boundary between V1 and V2. It is highly distorted because it includes alignments of several pinwheels with the same chirality. *Small white and black discs* represent singularities with levorotatory and dextrorotatory chiralities, respectively. From Ohki et al. [57]




Fig. 4.59 Physical field model of the distortion of the pinwheel lattice at the V1/V2 boundary. This boundary resembles a ‘cliff’, as can be seen from the tightly spaced bundles of parallel level sets

4.5 Topological Universality of Pinwheels

We shall see in Sect. 4.9.4 that the pinwheel structure is widespread in visual systems. Note, however, that this concrete empirical omnipresence is matched by an abstract *topological universality* which helps to explain it. In a very elegant paper [58], Daniel Bennequin and coworkers Alberto Romagnoni, Jérôme Ribot, and Jonathan Touboul provided a topological demonstration. The point of this theorem is to show just how far apparently very general and barely restrictive conditions can effectively amount to drastic constraints. It is one of the great advantages of mathematics to be able to identify this kind of ‘Platonic’ constraint.

The idea is to consider the way cells parametrized by $z = (x, y)$ in a domain $\Omega \subset \mathbb{C}$ of the visual plane identified with the plane $\mathbb{C} = \mathbb{R}^2$ can encode orientations (modulo π) $\psi \in \mathbb{S}_\pi^1$ (where \mathbb{S}_π^1 is $[0, \pi]$, with π identified with 0), i.e., a *periodic quantity*. Such a coding can be described by a mapping $f : \Omega \rightarrow \mathbb{S}_\pi^1$ which may have singularities, i.e. points where it is not defined because multivalued there.

Since the idea is to model hypercolumn functional architectures, we assume that f is *local*, i.e. defined apart from a few possible singular points on a domain Ω which is a small disc centred on 0 and having a circle as boundary $\Gamma = \partial\Omega$. But at the same time, since the notion of size is relative to a *scale*, we seek local models which would ideally be *scale invariant*. This means that these are models centred on a singularity and with a qualitative structure that does not vary when we zoom in or out.

The authors begin by defining a rather general class \mathcal{F} of sufficiently smooth maps f satisfying natural conditions. They then assume that f is *exhaustive*, in the sense that it codes all orientations, i.e. it is *surjective*. Clearly, the standard model of the pinwheel, viz.

$$f_0 : \Delta \rightarrow \mathbb{S}_\pi^1, \quad z = \rho e^{i\theta} \mapsto \psi = \frac{\theta}{2},$$

defined over the unit disc Δ , must belong to \mathcal{F} . Now, f_0 has several properties:

- (i) f_0 is defined away from 0, with 0 being a singularity at which it is multivalued (the ‘image’ of 0 is the whole of \mathbb{S}_π^1).
- (ii) Away from 0, f_0 is regular (smooth, infinitely differentiable).
- (iii) The level sets $C_\psi = f_0^{-1}(\psi)$ on which f_0 is constant are rays and hence smooth curves joining the boundary Γ to the singular point 0.

The fact that the level sets C_ψ are curves is to be expected. Indeed, suppose that $f : \Omega \rightarrow \mathbb{S}_\pi^1$ is smooth (apart from singularities) and *generic*, i.e. with a topology that remains qualitatively the same when f is slightly deformed. Then, the level sets C_ψ must be sub-manifolds of Ω with codimension $c = \dim(\mathbb{S}_\pi^1) = 1$ so that there exists in Ω a degree of freedom transverse to the C_ψ , corresponding to changes in ψ . Hence, since Ω has dimension $n = 2$, the level sets C_ψ must be sub-manifolds of dimension $n - c = 2 - 1 = 1$.⁹ As curves, the C_ψ can have singularities (which are not necessarily singularities of f , where f is not defined). Generically, these are end points or triple points.

The authors then introduce a *minimal complexity* constraint known as ‘parsimony’, which stipulates that the topological redundancy, that is the number of connected components of the level sets C_ψ , should be *minimal*, i.e., equal to unity (as for the rays of the standard pinwheel). The class \mathcal{F} defined in this way is thus the class of maps $f : \Omega \rightarrow \mathbb{S}_\pi^1$ which are smooth away from singularities in the interior of Ω , surjective (exhaustivity constraint), and whose level sets C_ψ are smooth connected curves (parsimony constraint). The authors then prove the following theorem:

Theorem *Topological universality of pinwheels. The elements $f : \Omega \rightarrow \mathbb{S}_\pi^1$ of \mathcal{F} have the pinwheel topology, i.e., an isolated singularity in the interior of Ω with the C_ψ joining the different points of the boundary $\Gamma = \partial\Omega$ to this singularity.*

The proof is based on purely topological arguments. Here, we give a heuristic account. We use a fundamental topological property of any circle, e.g. the standard circle \mathbb{S}^1 : if we remove any point or interval, it will not be disconnected. So if we remove an arc $J = [\theta_1, \theta_2]$ containing $(\theta_1 + \theta_2)/2$, we can still move from θ_1 to θ_2 by following the other arc containing $\pi + (\theta_1 + \theta_2)/2$ [the point opposite to $(\theta_1 + \theta_2)/2$ on \mathbb{S}^1]. This is not the case with $\mathbb{R} = (-\infty, +\infty)$: if we remove an interval $I = [x_1, x_2]$, we disconnect \mathbb{R} into two disjoint intervals $(-\infty, x_1)$ and $(x_2, +\infty)$ (see Fig. 4.60).

Let $\psi_0 \in f(\Gamma)$ be an orientation coded by a point on the boundary Γ . By hypothesis, $C_{\psi_0} = C$ is connected and intersects Γ . Now, consider the complement $\Omega - C$ of C in Ω and let c be the number of its connected components (connected by arcs¹⁰).

⁹For more on such questions of dimensions, the reader could consult my 1982 review [59] and references therein.

¹⁰Given the assumptions of continuity and smoothness, we need consider only connectedness by arcs.





Fig. 4.60 A circle remains connected if we remove an arc $[\theta_1, \theta_2]$, whereas a straight line becomes disconnected if we remove an interval $[x_1, x_2]$

Fig. 4.61 Level curves of f cannot have triple points decomposing Ω into three connected components. Symbols $<$ and $>$ indicate $f < \psi_0$ or $f > \psi_0$, respectively



It can be shown that we cannot have $c \geq 3$ (see Fig. 4.61). Indeed, suppose that C has a triple point t and consider a small enough open neighbourhood T of t and a small enough open neighbourhood V of ψ_0 . Let $W = f^{-1}(V) \cap T$. Then, W is made up of points close to t , where f takes values close to ψ_0 . Let $W_i, i = 1, 2, 3$, be the three connected components of $W - C$. We can then arrange for f to be $> \psi_0$ or $< \psi_0$ in each W_i . As this makes two possibilities for three components, there must therefore be two components with the same property. Suppose, for example, that $f < \psi_0$ on W_1 and W_2 and $f > \psi_0$ on W_3 . This means that (i) the part C_{12} of C which is the shared boundary of W_1 and W_2 is a curve along which f is transversally maximal, and (ii) the parts C_{23} and C_{31} of C which are shared boundaries of W_2 and W_3 and of W_3 and W_1 are curves along which f is transversally monotonic. Figure 4.62 shows that, in this case, the level lines C_ψ for $\psi \neq \psi_0, \psi \in V$ are no longer connected, which is forbidden by parsimony.




Fig. 4.62 *Left* Structure of level lines in the neighbourhood of a triple point of C . *Right* A level line close to C (*continuous line*) which is disconnected and another level line close to C (*dashed line*) which is connected

But neither is it possible to have $c = 2$, $\Omega - C$ being decomposed into two disjoint connected components U_1 and U_2 on which we would have $f > \psi_0$ and $f < \psi_0$, respectively (see Fig. 4.63). Let $I_1 \subset \mathbb{S}_\pi^1$ be the values taken by f on $U_1 \cup C$. Since $U_1 \cup C$ is compact and connected and f is continuous, the image I_1 is compact and connected¹¹ and is thus an interval $I_1 = [\psi_0, \psi_1]$ (by continuity, we can have $\psi_1 > \pi$). Likewise, the image of $U_2 \cup C$ is an interval $I_2 = [\psi_2, \psi_0]$ (by continuity, we can have $\psi_2 < 0$). If $\psi_1 - \psi_2 < \pi$, then part of \mathbb{S}_π^1 is not covered by f , so f is not surjective, a situation excluded by hypothesis. If, on the other hand, $\psi_1 - \psi_2 \geq \pi$, then there is at least one value $\psi \neq \psi_0$ for which the level curve C_ψ intersects both U_1 and U_2 , which are disjoint (see Fig. 4.64). These two pieces of C_ψ cannot be connected by C_ψ because, for this to happen, C_ψ would have to cross the boundaries of U_1 and U_2 , and hence C_{ψ_0} . But this is impossible because $\psi \neq \psi_0$, whence C_ψ and C_{ψ_0} are disjoint. Therefore, C_ψ is not connected and the topological redundancy is ≥ 2 , which is forbidden by hypothesis.

So the connected level line $C = C_{\psi_0}$ does not disconnect Ω . It can be shown that it cannot be entirely contained within the boundary Γ and that it enters into Ω and ends at a singular point s . We then choose two points z_1 and z_2 in $\Omega - C$ and connect them by two paths: one of them $\gamma : z_1 \rightarrow z_2$ crossing C and the other $\gamma' : z_1 \rightarrow z_2$ going round C , in such a way as to give the situation in Fig. 4.65.

Let $\psi_1 = f(z_1)$ and $\psi_2 = f(z_2)$. Since $z_1, z_2 \notin C$, we have $\psi_1, \psi_2 \neq \psi_0$. Since γ and γ' are compact and connected, their images are intervals $[\psi_1, \psi_2]$ of \mathbb{S}_π^1 . Since γ crosses C , its image is the interval $[\psi_1, \psi_2]$ containing ψ_0 , and since γ' goes around C , its image is the interval $[\psi_1, \psi_2]$ not containing ψ_0 . Hence, the

¹¹Here, the authors use basic theorems of general topology going back to Bolzano, Weierstrass, Heine, Borel, and Lebesgue. For separated topological spaces, the continuous image of a compact (connected) set is compact (connected).




Fig. 4.63 Level curves of f cannot decompose Ω into two connected components. Symbols $<$ and $>$ indicate $f < \psi_0$ or $f > \psi_0$, respectively





Fig. 4.64 *Left* If $(\psi_1 - \psi_0) - (\psi_2 - \psi_0) = \psi_1 - \psi_2 < \pi$, the interval (ψ_1, ψ_2) (bottom) is not covered by f . *Right* If, on the other hand, $(\psi_1 - \psi_0) - (\psi_2 - \psi_0) = \psi_1 - \psi_2 \geq \pi$, the interval (ψ_2, ψ_1) (bottom) is covered twice (example given by angle ψ)

image of $\gamma \cup \gamma'$ under f is the whole circle \mathbb{S}^1_π . But the configuration in Fig. 4.65 can be shrunk towards s . Thus, all the level lines C_ψ must converge towards the point s .

Finally, it can be shown that f is already surjective on the boundary Γ . Indeed, $f|_\Gamma$ is a continuous map of a circle onto a circle. If it is not surjective, it then takes each of its values generically at least twice (counted with multiplicity). This means that there are at least two level curves C_ψ for each value ψ in the range, which contradicts the assumption of minimal topological redundancy (equal to unity). Therefore, f is

Fig. 4.65 The level line C does not disconnect Ω , enters Ω , and ends at a singular point s . z_1 and z_2 are two points in $\Omega - C$ joined by two paths, one of them $\gamma : z_1 \rightarrow z_2$ (red) crossing C and the other $\gamma' : z_1 \rightarrow z_2$ (green) going around C



surjective. In fact, by the same argument of minimal topological redundancy, it is also injective.

4.6 Pinwheels as Phase Fields

Having discussed the topological universality of pinwheels as local models of hypercolumns, we shall now turn to their global lattice structure. On a certain scale, functional orientation maps assign an orientation $\psi(a) \bmod \pi$ to each point a of the cortical surface of V1. For simplicity, we shall treat this surface as a plane \mathbb{R}^2 with coordinates (x, y) , identified with a complex plane \mathbb{C} with coordinate z . By treating V1 in Sect. 4.3.5 as an implementation of the space of 1-jets of curves in \mathbb{R}^2 , we represented ψ by its tangent $\tan(\psi)$. It was associated with a function $e^{i\varphi}$, where φ is the phase $\varphi = 2\psi$, defined mod 2π .

These maps are thus *phase fields* in which the pinwheels are singularities. This kind of field turns up in many types of physical phenomenon, and in particular in *optical* and *liquid crystal* structures. There is a vast literature about this, and especially about the singularities. In this section, we shall import a certain number of the associated formalisms and models to neurogeometry.

We do this with a certain amount of mathematical detail because we shall not return to this model in the second volume (where we shall in fact develop an alternative model in considerable mathematical detail).

4.6.1 Fields and Coordinates

Since we shall be considering several spaces and several fields, we begin by defining the notation. The cortical layer is modelled by \mathbb{R}^2 . If $a = (x, y)$ is a point of \mathbb{R}^2 ,

it will often be useful to consider it as a complex number $z = x + iy$ using the standard \mathbb{R} -linear isomorphism between \mathbb{C} and \mathbb{R}^2 . To study the structure of the field in the neighbourhood of a singular point a_0 , it will often be useful to transport the coordinate system on \mathbb{R}^2 to a_0 by a translation carrying 0 to a_0 and an appropriate rotation of the axes, and then to take polar coordinates (ρ, θ) around $a_0 = 0$, i.e. $z = \rho e^{i\theta}$, $x = \rho \cos(\theta)$, $y = \rho \sin(\theta)$. A phase field assigning the phase $\varphi(a)$ to each point $a \in \mathbb{R}^2$ is thus a map $\Phi : \mathbb{R}^2 \rightarrow \mathbb{S}^1$, $\Phi(a) = e^{i\varphi(a)}$, in other words a section of the fibre bundle $\pi : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$, with singularities of a particular type in places where the phase $\varphi(a)$ is not defined. Clearly, as $\varphi(a)$ is a function defined on \mathbb{R}^2 , if it is differentiable, it can have singularities in the classical sense, i.e. *critical points* where the gradient $\nabla\varphi = 0$. Generically, there are three types: extrema, i.e. maxima or minima, and saddle points. Experimentally, there do not seem to be any extrema (where the isochromatic curves would be concentric circles locally), whereas there are many saddle points occupying the ‘centres’ of the cells defined by the pinwheels.

The experimental data show that the orientation maps $\psi(a) = \varphi(a)/2$ can also have lines of discontinuity called fractures, across which the orientation jumps from one value to another. If there are no fractures, $\psi(a)$ is smooth away from singular points, where it is not specified. In the local pinwheel models involving end point and triple point singularities, discussed in Sect. 4.4.6, we took $\psi = \alpha \pm \theta/2$.

It is often possible and natural to associate an *amplitude*, that is a modulus $r(a)$, with the phase of a phase field. For functional orientation maps, Fred Wolf and Theo Geisel suggested introducing the strength of the orientation selectivity, i.e. the width of the tuning curve. We shall return to this key idea in Sect. 4.7.1. Under such an assumption, the phase field Φ becomes the ‘phase part’ of a *complex scalar field* $Z : \mathbb{C} \rightarrow \mathbb{C}$, $z = \rho e^{i\theta} \mapsto r(z) e^{i\varphi(z)}$, also denoted $r(a) e^{i\varphi(a)}$. In Cartesian coordinates, $Z(a)$ will be written $Z(a) = X(a) + iY(a)$, where X and Y are two real functions of the variables (x, y) .

The map Z can be viewed as a section of the bundle $\pi_1 : \mathbb{R}^2 \times \mathbb{C} \rightarrow \mathbb{R}^2$ whose fibre is no longer \mathbb{S}^1 but \mathbb{C} . The group $E(2)$ acts on π_1 as it acts on π : in the standard way in the base space \mathbb{R}^2 and by rotation of the phase in the fibre.

4.6.2 Singularities of a Phase Field

Given a phase field $Z(a)$ —which may be denoted $Z(z)$, $X(a) + iY(a)$, $X(z) + iY(z)$, $r(a) e^{i\varphi(a)}$, or $r(z) e^{i\varphi(z)}$, depending on the context—the geometry can be analyzed using the standard tools of vector analysis, viz. gradients ∇ , divergences, curls, Laplacians Δ , and so on.

The singularities of Z are points where the phase φ is not defined, but where Z is for its part *well defined*, although with a particular value, 0 or ∞ . They correspond to *zeros and poles* of Z . If $|Z|$ is bounded, then they can only be zeros. We go from zeros to poles by allowing $|Z|$ to diverge and taking for the space of values of Z , not only the complex plane \mathbb{C} , but its completion $\widehat{\mathbb{C}}$ by a point at infinity (called the

Riemann sphere or the complex projective line). The reciprocal

$$\frac{1}{Z} = \frac{1}{r(a)} e^{-i\varphi(a)}$$

of Z then has a pole at a if Z has a zero at a . For a zero, the simplest local model is $\pm z$ if Z is a *holomorphic* function, i.e., a smooth function of the *complex* variable z , hence *differentiable with respect to z* (and not only with respect to x and y separately, see below), and for a pole, it is $\pm 1/z$ if Z is a meromorphic function (see below). However, there is no reason a priori why Z should be holomorphic. In the neighbourhood of a point a_0 of the base space \mathbb{R}^2 taken as the origin 0 , we then have to first order

$$Z(x, y) \approx X(0) + x \frac{\partial X}{\partial x}(0) + y \frac{\partial X}{\partial y}(0) + i \left[Y(0) + x \frac{\partial Y}{\partial x}(0) + y \frac{\partial Y}{\partial y}(0) \right],$$

whence

$$Z(a) \approx Z(0) + a \cdot \nabla_0 X + ia \cdot \nabla_0 Y,$$

where $\nabla_0 X$ is the value of the gradient of X at 0 , $\nabla X = (\partial X/\partial x, \partial X/\partial y)$, and the same for Y . Hence,

$$|Z(a) - Z(0)|^2 \approx R^2 = (a \cdot \nabla_0 X)^2 + (a \cdot \nabla_0 Y)^2,$$

and the level lines $R = \text{const.}$ are *ellipses* to this order of approximation. They are only *circles* if $Z(a)$ can be written as a function $Z(z)$ of z , in other words if on the one hand $x\partial X/\partial x + iy\partial Y/\partial y$ is proportional to z , which requires $\partial X/\partial x = \partial Y/\partial y$, and if on the other hand $y\partial X/\partial y + ix\partial Y/\partial x$ is proportional to iz , which requires $\partial X/\partial y = -\partial Y/\partial x$. These fundamental conditions, known as the *Cauchy–Riemann equations*, express the fact that the gradients ∇X and ∇Y of Z are *orthogonal*. They characterize holomorphic functions.

Now in all the situations where we analyze fields, the *singularities* play a determining structural role and contain most of the essential morphological information. As noted by Berry [60], the eminent specialist in optics (whether it be geometrical, wave, or quantum optics):

Each singularity is a window to a deeper theory.

Regarding lines of phase singularities for 3D waves,¹² Berry also stresses that [61, p. 724]:

Wave vortex lines can be regarded as a skeleton, characterizing and supporting the full structure of the wave.

¹²Phase singularities are generically points in 2D and lines in 3D because they are specified by two conditions and so have codimension 2 (see below).

This is a general ‘philosophy’ which geometers have clearly established since René Thom (see Sect. 1.3 of the Preface). We shall take this as a starting point, and all the more so in that, for a long time now, we have been following Michael Berry’s work in optics very closely, and in particular his work on *caustics*, because these represent one of the main applications of the theory of singularities in fundamental physics. More recent applications of these studies in phase fields are now also used in the theory of pinwheels by Fred Wolf and Theo Geisel (we shall return to this), Bennequin [62], Afgoustidis [63, 64], and Giovanna Citti and Sarti [65].

Many of the concepts used in the theory of singularities, such as genericity, codimension, bifurcations, unfolding, or normal forms, have thus turned out to be highly relevant in neurogeometry. The reader will find much more about this in our surveys [59] and [66] on the theory of singularities and critical phenomena, and in particular, further reading (but see also our *Landmarks* [67]).

Let us assume that the field Z is smooth outside singular points, where the phase φ is indeterminate and where Z vanishes. Since $Z = X + iY$, these points are intersections between curves with equations $X = 0$ and $Y = 0$. The condition $X = 0$ corresponds to $r \cos(\varphi) = 0$, i.e. $\varphi = \pi/2 \bmod \pi$ if $r \neq 0$, and $Y = 0$ corresponds to $r \sin(\varphi) = 0$, i.e. $\varphi = 0 \bmod \pi$ if $r \neq 0$. If $X = Y = 0$, we necessarily have $r = 0$ because the two conditions on φ are incompatible. Generically, the curves $X = 0$ and $Y = 0$ cross transversally at isolated points. This means that the points which satisfy both conditions have codimension 2 and, as the surrounding space \mathbb{R}^2 is 2D, are isolated points (whereas in a 3D surrounding space, they would be lines). In the theory of mesophases (liquid crystals), these phase singularities are called *dislocations*. The same term could be used here to say that the pinwheels are dislocations in the orientation field implemented in $V1$.

Note that these dislocations $Z = 0$ are *invariant* under gauge transformation $Z \rightarrow e^{i\sigma} Z$ in the target space and under change of differentiable coordinates $(x, y) \rightarrow (\xi, \eta)$ in the source space.

4.6.3 Orientation and Iso-orientation Fields

There are several different fields here. The field Z is the phase field $\varphi(a) = 2\psi(a)$. It has *isophase lines*, called wave fronts by analogy with optics. As isophase means iso-orientation, they are represented by isochromatic lines in the pinwheel maps.

In addition, there is the orientation field $\psi(a) = \varphi(a)/2$. As such, it defines a foliation of the plane \mathbb{R}^2 by its integral curves. With the end point and triple point models in Sect. 4.4.6, we identified the local geometry of these foliations at the singular points. Let us denote this field by $W(a) = s(a)e^{i\psi(a)}$, assuming that meaning can be attributed to the amplitude $s(a)$. W also has field lines and isophase lines. However, in contrast to the phase φ , which is defined modulo 2π , the angle ψ is only defined modulo π . For pinwheels without distortion, φ can in fact be identified with $2\alpha \pm \theta$ and ψ with $\alpha \pm \theta/2$. In this case, we have $e^{i\varphi} = (e^{i\psi})^2$. The simplest

way to satisfy this condition is to take $Z = W^2$. Since $W = \sqrt{Z}$, there is a Riemann cut, because when φ changes by 2π , which leaves Z invariant, ψ changes by π and W changes sign, the square root function changing determination.

4.6.4 Topological Charge and Index

Let a_0 be a singularity of Z taken as origin. The *topological charge* of this singularity is defined to be $q = \oint_{\gamma} d\varphi/2\pi$, where $\oint_{\gamma} d\varphi$ denotes the integral of the differential 1-form $d\varphi$, i.e. the change in φ around a small closed path γ going once around a_0 in the right-handed sense. Since a_0 is an isolated singularity, there are such paths γ enclosing only this singularity and it can be shown that the integral is independent of the choice of γ . From the Euclidean structure of \mathbb{R}^2 , the differential 1-form

$$d\varphi = \frac{\partial\varphi}{\partial x}dx + \frac{\partial\varphi}{\partial y}dy$$

corresponds to the vector gradient $\nabla\varphi$ with components $\nabla\varphi = (\partial\varphi/\partial x, \partial\varphi/\partial y)$, and if the differentials are interpreted as in the past as infinitesimal variations, we have $d\varphi = \nabla\varphi \cdot da$ (scalar product) and $\oint_{\gamma} d\varphi = \oint_{\gamma} \nabla\varphi \cdot da$ becomes what is known as the circulation of the gradient field $\nabla\varphi$ around the path γ . The topological charge $q = \oint_{\gamma} \nabla\varphi \cdot da/2\pi$ can then be interpreted as the *topological index* of the field $\nabla\varphi$.

For the field Z , the phase φ varies as $\pm\theta$ and the index is ± 1 . However, for the field W , the orientation ψ varies as $\pm\theta/2$ and the index is $\pm 1/2$. We already noted this for the pinwheels in Sects. 4.4.3 and 4.4.6.

In the vicinity of a pinwheel, the isophase lines $\varphi = \text{const.}$, i.e. the wave fronts, are rays of the pinwheel. Along these wave fronts, $d\varphi = 0$ and hence $\nabla\varphi \cdot da = 0$, which means that the field $\nabla\varphi$ is orthogonal to the rays and that its trajectories are therefore locally qualitatively like concentric circles centred on the singular point. Quite generally, the trajectories of $\nabla\varphi$ are orthogonal to the wave fronts.

4.6.5 Current, Vorticity, and Divergence

To understand the behaviour of $\nabla\varphi$ in the neighbourhood of a singularity, the best thing is to use polar coordinates. Since $x = \rho \cos(\theta)$ and $y = \rho \sin(\theta)$, differentiating we obtain

$$\begin{cases} dx = \cos(\theta) d\rho - \rho \sin(\theta) d\theta, \\ dy = \sin(\theta) d\rho + \rho \cos(\theta) d\theta, \\ d\rho = \cos(\theta) dx + \sin(\theta) dy, \\ \rho d\theta = -\sin(\theta) dx + \cos(\theta) dy. \end{cases}$$

Then, since

$$d\varphi = \frac{\partial\varphi}{\partial x}dx + \frac{\partial\varphi}{\partial y}dy ,$$

we have

$$\begin{cases} \frac{\partial\varphi}{\partial x} = \cos(\theta) \frac{\partial\varphi}{\partial\rho} - \frac{\sin(\theta)}{\rho} \frac{\partial\varphi}{\partial\theta} , \\ \frac{\partial\varphi}{\partial y} = \sin(\theta) \frac{\partial\varphi}{\partial\rho} + \frac{\cos(\theta)}{\rho} \frac{\partial\varphi}{\partial\theta} . \end{cases}$$

At the singular point $\rho = 0$, the gradient $\nabla\varphi$ is not defined and diverges.

To regularize this situation, physicists usually consider the *current* \mathcal{J} of the field, which is the vector pointing in the direction of the gradient $\nabla\varphi$ when it does not vanish, defined by

$$\mathcal{J} = r^2\nabla\varphi .$$

Note that if $Z = X + iY$, then

$$\mathcal{J} = X\nabla Y - Y\nabla X ,$$

and hence that \mathcal{J} is well defined, even at singular points of the phase φ of Z . Indeed, since $X = r \cos(\varphi)$ and $Y = r \sin(\varphi)$,

$$\left\{ \begin{array}{l} \nabla X = \left[\cos(\varphi) \frac{\partial r}{\partial x} - r \sin(\varphi) \frac{\partial\varphi}{\partial x}, \cos(\varphi) \frac{\partial r}{\partial y} - r \sin(\varphi) \frac{\partial\varphi}{\partial y} \right] , \\ \nabla Y = \left[\sin(\varphi) \frac{\partial r}{\partial x} + r \cos(\varphi) \frac{\partial\varphi}{\partial x}, \sin(\varphi) \frac{\partial r}{\partial y} + r \cos(\varphi) \frac{\partial\varphi}{\partial y} \right] , \\ X\nabla Y = \left[r \cos(\varphi) \sin(\varphi) \frac{\partial r}{\partial x} + r^2 \cos^2(\varphi) \frac{\partial\varphi}{\partial x}, \right. \\ \qquad \qquad \qquad \left. r \cos(\varphi) \sin(\varphi) \frac{\partial r}{\partial y} + r^2 \cos^2(\varphi) \frac{\partial\varphi}{\partial y} \right] , \\ Y\nabla X = \left[r \sin(\varphi) \cos(\varphi) \frac{\partial r}{\partial x} - r^2 \sin^2(\varphi) \frac{\partial\varphi}{\partial x}, \right. \\ \qquad \qquad \qquad \left. r \sin(\varphi) \cos(\varphi) \frac{\partial r}{\partial y} - r^2 \sin^2(\varphi) \frac{\partial\varphi}{\partial y} \right] , \\ X\nabla Y - Y\nabla X = \left[r^2 \frac{\partial\varphi}{\partial x}, r^2 \frac{\partial\varphi}{\partial y} \right] = r^2\nabla\varphi . \end{array} \right.$$

Note also that, in terms of the complex conjugate values Z and \bar{Z} , the current \mathcal{J} can be written

$$\mathcal{J} = \text{Im}(\bar{Z}\nabla Z) .$$

Indeed,

$$\bar{Z}\nabla Z = (X - iY)(\nabla X + i\nabla Y) = X\nabla X + Y\nabla Y + i(X\nabla Y - Y\nabla X) .$$

We shall see an example of a current in Sect. 4.6.7.

Another vector, in fact a pseudovector, used by physicists is the *vorticity* Ω of the current \mathcal{J} , i.e. its *curl*, up to a factor. By definition,

$$\Omega = \frac{1}{2}\nabla \times \mathcal{J} = \nabla X \times \nabla Y ,$$

where the symbol \times stands for the *exterior product* of two vectors in \mathbb{R}^2 . If

$$u = (u_x, u_y) = u_x e_x + u_y e_y , \quad v = (v_x, v_y) = v_x e_x + v_y e_y ,$$

are two vectors in \mathbb{R}^2 , where e_x and e_y are unit vectors associated with the x - and y -axis, the exterior product $u \times v$ is a vector of magnitude

$$\det \begin{pmatrix} u_x & u_y \\ v_x & v_y \end{pmatrix} = u_x v_y - u_y v_x = \omega$$

along an axis orthogonal to \mathbb{R}^2 , with unit vector e_3 , such that the frame $\{e_x, e_y, e_3\}$ is right-handed. It is the area of the parallelogram constructed from u and v and oriented normal to the plane they lie in. For $\nabla X \times \nabla Y$,

$$\omega = \det \begin{pmatrix} \partial X/\partial x & \partial X/\partial y \\ \partial Y/\partial x & \partial Y/\partial y \end{pmatrix}$$

is the determinant of the Jacobian¹³ of Z considered as a map from \mathbb{R}^2 into \mathbb{R}^2 . In fact, the right interpretation of $\Omega = \omega e_3$ is as a differential 2-form. It can be checked that $\Omega = \nabla \times \mathcal{J}/2$, the components \mathcal{J}_x and \mathcal{J}_y of \mathcal{J} being

$$\mathcal{J}_x = X \frac{\partial Y}{\partial x} - Y \frac{\partial X}{\partial x} , \quad \mathcal{J}_y = X \frac{\partial Y}{\partial y} - Y \frac{\partial X}{\partial y} .$$

¹³Named after Carl Gustav Jacob Jacobi.

Indeed,

$$\begin{aligned}
 \Omega &= \frac{1}{2} \nabla \times \mathcal{J} = \frac{1}{2} \left(\frac{\partial \mathcal{J}_y}{\partial x} - \frac{\partial \mathcal{J}_x}{\partial y} \right) e_3 \\
 &= \frac{1}{2} \left(\frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} + X \frac{\partial^2 Y}{\partial x \partial y} - \frac{\partial Y}{\partial x} \frac{\partial X}{\partial y} - Y \frac{\partial^2 X}{\partial y \partial x} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x} - X \frac{\partial^2 Y}{\partial x \partial y} \right. \\
 &\quad \left. + \frac{\partial Y}{\partial y} \frac{\partial X}{\partial x} + Y \frac{\partial^2 X}{\partial y \partial x} \right) e_3 \\
 &= \det \begin{pmatrix} \partial X / \partial x & \partial X / \partial y \\ \partial Y / \partial x & \partial Y / \partial y \end{pmatrix} e_3 = \omega e_3 .
 \end{aligned}$$

Note that when $\Omega = 0$, either $\nabla X = 0$ or $\nabla Y = 0$ (isolated points with codimension 2), or the real gradients ∇X and ∇Y have the same orientation (lines of codimension 1), whence $\nabla Y = \alpha \nabla X$ for $\alpha \in \mathbb{R}$. The complex gradient $\nabla Z = \nabla X + i \nabla Y \in \mathbb{R}^2 \oplus i \mathbb{R}^2$ is thus the real vector ∇X multiplied by a factor $1 + i\alpha$. The condition $\Omega = 0$ saying that ∇X and ∇Y are parallel is the opposite of the Cauchy–Riemann equations saying that ∇X and ∇Y are orthogonal. We shall also see an example of vorticity in Sect. 4.6.7.

Note also that, in terms of the values of the complex conjugates Z and \bar{Z} , the vorticity Ω can be written

$$\Omega = \frac{1}{2} \text{Im} (\nabla \bar{Z} \times \nabla Z) .$$

Indeed,

$$\begin{aligned}
 \nabla \bar{Z} \times \nabla Z &= \frac{\partial \bar{Z}}{\partial x} \frac{\partial Z}{\partial y} - \frac{\partial \bar{Z}}{\partial y} \frac{\partial Z}{\partial x} \\
 &= \left(\frac{\partial X}{\partial x} - i \frac{\partial Y}{\partial x} \right) \left(\frac{\partial X}{\partial y} + i \frac{\partial Y}{\partial y} \right) - \left(\frac{\partial X}{\partial y} - i \frac{\partial Y}{\partial y} \right) \left(\frac{\partial X}{\partial x} + i \frac{\partial Y}{\partial x} \right) \\
 &= 2i \left(\frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x} \right) .
 \end{aligned}$$

Points where $\Omega = 0$ are invariant under the gauge transformation $Z \rightarrow e^{i\sigma} Z$ (multiplication by $e^{i\sigma}$ is effectively a rotation in the Z plane so parallel vectors are mapped to parallel vectors) and coordinate changes $(x, y) \rightarrow (\xi, \eta)$ [because ω becomes $\omega' = \omega \det(\text{Jac})$, where Jac is the Jacobian of the change of variables and, since by definition $\det(\text{Jac}) \neq 0$, $\omega = 0 \iff \omega' = 0$].

The vorticity of \mathcal{J} is not generally trivial. This is not the case for $\nabla \varphi$ because the curl of a gradient $\nabla \times \nabla f$ is always zero due to the fact that, for any vector u , we have $u \times u = 0$. Away from pinwheels where it is not defined, the field $\nabla \varphi$ is thus curl-free.

The idea of vorticity can be used to resolve a slight difficulty regarding the difference between the topological charge and index of a singularity. For the charge, the right-handed orientation of a small closed path γ about a singular point is the right-handed orientation of \mathbb{R}^2 . The charges of the field Z (whose trajectories are isophase lines, i.e. wave fronts) are then ± 1 . For the index, we often orient the γ with respect to Ω and then the index is always $+1$.

We saw above that, to first order, in the neighbourhood of a point a_0 taken as origin 0, the modulus of Z is given by

$$|Z(a) - Z(0)|^2 = R^2 = (a \cdot \nabla_0 X)^2 + (a \cdot \nabla_0 Y)^2 .$$

The current \mathcal{J} is given to first order by the 2-column vector

$$\begin{aligned} \mathcal{J}(a) &= \mathcal{J}(x, y) \\ &\approx \begin{pmatrix} \left[X(0) + x \frac{\partial X}{\partial x}(0) + y \frac{\partial X}{\partial y}(0) \right] \left[\frac{\partial Y}{\partial x}(0) + x \frac{\partial^2 Y}{\partial x^2}(0) + y \frac{\partial^2 Y}{\partial x \partial y}(0) \right] \\ - \left[Y(0) + x \frac{\partial Y}{\partial x}(0) + y \frac{\partial Y}{\partial y}(0) \right] \left[\frac{\partial X}{\partial x}(0) + x \frac{\partial^2 X}{\partial x^2}(0) + y \frac{\partial^2 X}{\partial x \partial y}(0) \right] \\ \left[X(0) + x \frac{\partial X}{\partial x}(0) + y \frac{\partial X}{\partial y}(0) \right] \left[\frac{\partial Y}{\partial y}(0) + x \frac{\partial^2 Y}{\partial x \partial y}(0) + y \frac{\partial^2 Y}{\partial y^2}(0) \right] \\ - \left[Y(0) + x \frac{\partial Y}{\partial x}(0) + y \frac{\partial Y}{\partial y}(0) \right] \left[\frac{\partial X}{\partial y}(0) + x \frac{\partial^2 X}{\partial x \partial y}(0) + y \frac{\partial^2 X}{\partial y^2}(0) \right] \end{pmatrix} \\ &= \begin{pmatrix} X(0) \frac{\partial Y}{\partial x}(0) - Y(0) \frac{\partial X}{\partial x}(0) + x \left[X(0) \frac{\partial^2 Y}{\partial x^2}(0) - Y(0) \frac{\partial^2 X}{\partial x^2}(0) \right] \\ + y \left[X(0) \frac{\partial^2 Y}{\partial x \partial y}(0) - Y(0) \frac{\partial^2 X}{\partial x \partial y}(0) + \frac{\partial X}{\partial y}(0) \frac{\partial Y}{\partial x}(0) - \frac{\partial X}{\partial x}(0) \frac{\partial Y}{\partial y}(0) \right] \\ X(0) \frac{\partial Y}{\partial y}(0) - Y(0) \frac{\partial X}{\partial y}(0) + y \left[X(0) \frac{\partial^2 Y}{\partial y^2}(0) - Y(0) \frac{\partial^2 X}{\partial y^2}(0) \right] \\ + x \left[X(0) \frac{\partial^2 Y}{\partial x \partial y}(0) - Y(0) \frac{\partial^2 X}{\partial x \partial y}(0) - \frac{\partial X}{\partial y}(0) \frac{\partial Y}{\partial x}(0) + \frac{\partial X}{\partial x}(0) \frac{\partial Y}{\partial y}(0) \right] \end{pmatrix} . \end{aligned}$$

If $X(0) = 0$ and $Y(0) = 0$ (pinwheel, i.e., dislocation), then these formulas simplify enormously to give

$$\begin{aligned} \mathcal{J}(x, y) &\approx \begin{pmatrix} y \left[\frac{\partial X}{\partial y}(0) \frac{\partial Y}{\partial x}(0) - \frac{\partial X}{\partial x}(0) \frac{\partial Y}{\partial y}(0) \right] \\ x \left[-\frac{\partial X}{\partial y}(0) \frac{\partial Y}{\partial x}(0) + \frac{\partial X}{\partial x}(0) \frac{\partial Y}{\partial y}(0) \right] \end{pmatrix} \\ &= (a \cdot \nabla_0 X) \nabla_0 Y - (a \cdot \nabla_0 Y) \nabla_0 X = \Omega_0 \times a = \omega_0 \begin{pmatrix} -y \\ x \end{pmatrix}. \end{aligned}$$

We can thus evaluate $|\mathcal{J}| = r^2 |\nabla\varphi| \approx |\omega| \rho$ in the vicinity of dislocations where $\omega \neq 0$. However, locally, φ is constant on rays from such a singular point and $\nabla\varphi$ is orthogonal to the rays, and in polar coordinates

$$\nabla\varphi = \frac{\partial\varphi}{\partial\rho} e_\rho + \frac{1}{\rho} \frac{\partial\varphi}{\partial\theta} e_\theta,$$

where e_ρ is the unit vector along the radius at a and e_θ is the unit vector orthogonal to e_ρ , i.e. e_ρ rotated through $+\pi/2$, so we have

$$\nabla\varphi \approx \frac{1}{\rho} \frac{\partial\varphi}{\partial\theta} e_\theta,$$

and hence,

$$r^2 \left| \frac{\partial\varphi}{\partial\theta} \right| \approx \rho^2 |\omega|.$$

This tells us that, whereas r is locally constant on the *ellipses*

$$r^2 = (a \cdot \nabla X)^2 + (a \cdot \nabla Y)^2,$$

the quantity $r^2 |\partial\varphi/\partial\theta|$ is constant on the *circles* $\rho = \text{const}$. As noted by Mark Richard Dennis, this is a kind of Kepler's law for $r^2 |\partial\varphi/\partial\theta|$, which is analogous to the angular momentum [68, p. 41]:

Equal area vectors of the core anisotropy ellipse [$r^2 = \text{const}$.] are swept out in equal intervals of phase.

Note that the eccentricity of the ellipses measures the anisotropy of the vorticity. As we have seen, there is only isotropy (the ellipses are only circles) if the Cauchy–Riemann equations are satisfied.

4.6.6 Helmholtz Equation

When we discussed the experimental results for pinwheels, we saw that these appear when we superpose maps showing the response to different orientations. We also

saw that the pinwheel lattice has a *characteristic length* (see Sect. 4.4.1). These two empirical facts suggest considering the field Z as a superposition of simpler fields with a characteristic mesh. Moreover, on the mathematical level, any field can be considered through its Fourier transform as a superposition of *plane waves*. The latter are the simplest fields with a characteristic length. They have the form $Ae^{i\kappa \cdot a}$, where A is a complex amplitude $Ee^{i\phi}$ and $\kappa = (\kappa_x, \kappa_y)$ is a vector called the *wave vector*, whose magnitude $k = |\kappa|$, called the *wave number*, is analogous to a momentum and associated with a wavelength $\Lambda = 2\pi/k$ (the smaller the wavelength, the greater the wave number). When they evolve in time, their prototype is $Ae^{i(\kappa \cdot a - \omega t)}$, where ω is an angular frequency associated with a frequency $\nu = \omega/2\pi$ and a period $T = 1/\nu = 2\pi/\omega$.

It is straightforward to check that the plane waves $U = Ae^{i\kappa \cdot a}$ satisfy a fundamental equation called the *Helmholtz equation*, viz. $\Delta U + k^2 U = 0$. Indeed, $\kappa \cdot a = x\kappa_x + y\kappa_y$, whence

$$\Delta U = \frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} = -A\kappa_x^2 e^{i\kappa \cdot a} - A\kappa_y^2 e^{i\kappa \cdot a} = -k^2 U .$$

As the Helmholtz equation is linear, any linear superposition of solutions with different values of κ but the *same* magnitude k is also a solution. It is thus natural to assume that the field Z satisfies the Helmholtz equation for a certain wave number k :

$$\Delta Z + k^2 Z = 0 .$$

Let \mathcal{H}_k be the space of solutions C^∞ of the Helmholtz equation with wave number k . It can be shown that \mathcal{H}_k is $SE(2)$ -invariant; that is, if $Z(a) \in \mathcal{H}_k$ and if $g \in G = SE(2)$, then $gZ(a) = Z(g^{-1}(a))$ is also an element of \mathcal{H}_k . The action of G on \mathcal{H}_k therefore defines a representation of G , and this has the property of *irreducibility* in the usual sense that there is no *closed* subspace which is G -invariant.

Figure 4.66, recomputed from the data of Michael Berry's work [69] on optical currents, shows a superposition of 10 plane waves with the same value of k . We see how closely this phase field resembles our pinwheels, with its isophase lines, i.e. iso-orientation lines, its orthogonal gradient lines, and its saddle points.

4.6.7 Illustration

It is interesting to consider Michael Berry's example, which is a superposition of 10 plane waves, viz.

$$Z = \sum_{j=1}^{j=10} E_j \exp \left[i \left[\phi_j + 2\pi x \cos(\alpha_j) + 2\pi y \sin(\alpha_j) \right] \right] = re^{i\varphi} ,$$

Fig. 4.66 Superposition in the rectangle $[0, 1] \times [0, 1]$ of 10 plane waves with the same wave number $k = 2\pi$. The wave vectors $\kappa_j = (2\pi \cos(\alpha_j), 2\pi \sin(\alpha_j))$ are given below in Table 4.1. Recomputed from the data of Berry [69]




Table 4.1 Values used in Berry’s example [69]

	1	2	3	4	5	6	7	8	9	10
α_j	5.971	2.666	0.939	4.629	1.023	1.537	2.710	3.273	4.356	5.032
ϕ_j	3.846	0.777	5.008	2.916	6.274	4.344	2.411	5.688	1.734	0.214
E_j	0.337	0.015	0.762	0.785	0.625	0.442	0.688	0.065	0.064	0.035

with the same wave number $k = 2\pi$ (and hence wavelength $\Lambda = 2\pi/k = 1$) and wave vectors $\kappa_j = (2\pi \cos(\alpha_j), 2\pi \sin(\alpha_j))$. The angles α_j are chosen randomly in $[0, 2\pi]$, the phase shift ϕ_j randomly in $[0, 2\pi]$, and the amplitudes E_j randomly in $[0, 1]$. The values used by Berry are shown in Table 4.1. If Φ_j are the phases $\phi_j + 2\pi x \cos(\alpha_j) + 2\pi y \sin(\alpha_j)$, we get an expression $Z = \sum_{j=1}^{10} E_j \exp(i\Phi_j)$ that is easy to calculate with. Figure 4.67 shows the pinwheels of Z . The white lines are cuts where φ jumps by 2π , due to the fact that φ takes values in \mathbb{S}^1 but is represented as having values in \mathbb{R} . The coordinates of the 5 pinwheels are

$$\{0.528545, 0.942654\}, \{0.988124, 0.811337\}, \{0.433271, 0.516137\},$$

$$\{0.761954, 0.258734\}, \{0.0838329, 0.0359263\}.$$

They are given by the intersections of the level lines $X = 0, Y = 0$. Figure 4.68 represents the lines $X = 0$ in red and the lines $Y = 0$ in blue.

Figure 4.69 shows the structure of the phase field Z on the square $x, y \in [0, 1]$. The first line shows the modulus r of Z (i) with its level lines and (ii) as a function

Fig. 4.67 Pinwheels of the phase field Z for $x, y \in [0, 1]$. White lines are cuts where φ jumps by 2π





Fig. 4.68 Lines $X = 0$ (red) and $Y = 0$ (blue) of the phase field Z






Fig. 4.69 Structure of the phase field Z for $x, y \in [0, 1]$. *First line* modulus r of Z . (i) Level lines. (ii) The function $r(x, y)$ from two perspectives, the second giving a good view of the singular points (dislocations) where $r = 0$ and also the maxima of r . *Second line* argument φ of Z . (i) Level lines. (ii) Function $\varphi(x, y)$ from two perspectives, the second giving a good view of the cuts where φ jumps by 2π

$r(a) = r(x, y)$ from two perspectives, the second giving a good view of the singular points (dislocations), where $r = 0$, and the maxima of r . The second line shows the argument φ of Z (i) with its level lines (we recover those in Fig. 4.66) and (ii) as a function $\varphi(a) = \varphi(x, y)$ from two perspectives, the second giving a good view of the cuts, where φ jumps by 2π .

Figure 4.70 shows more pinwheels of $Z(x, y \in [0, 3])$. The white cuts represent $\varphi = 0 = 2\pi$. Note that there are 29 pinwheels in an area of $3^2 = 9$, giving a density d equal to $29/9 \sim 3.2$. We shall explain in Sect. 4.6.11 a formula giving $d = \pi/\Lambda^2$, which means $d = \pi$ in our case, since $\Lambda = 1$. We see that the approximation is excellent. Concerning Fig. 4.71, it represents the phase φ for partial sums of Z .

In Fig. 4.72, we also show the orientation lines, that is the field lines of W , and the field lines of Z (not to be confused with the isophase lines). We do indeed see examples of the typical singularities illustrated in Fig. 4.55 of Sect. 4.4.6.

It is straightforward to calculate the current

$$\mathcal{J} = r^2 \nabla \varphi = X \nabla Y - Y \nabla X = \text{Im}(\bar{Z} \nabla Z) .$$

Fig. 4.70 Pinwheels of Z for $x, y \in [0, 3]$. *White cuts* represent $\varphi = 0 = 2\pi$





Fig. 4.71 Phase φ for partial sums of Z




Fig. 4.72 Field lines of W and Z for the example of Fig. 4.69

Indeed,

$$\left\{ \begin{array}{l} X = \sum_{k=1}^{k=10} E_k \cos(\Phi_k) , \\ Y = \sum_{k=1}^{k=10} E_k \sin(\Phi_k) , \\ \frac{\partial X}{\partial x} = - \sum_{j=1}^{j=10} 2\pi \cos(\alpha_j) E_j \sin(\Phi_j) , \\ \frac{\partial X}{\partial y} = - \sum_{j=1}^{j=10} 2\pi \sin(\alpha_j) E_j \sin(\Phi_j) , \\ \frac{\partial Y}{\partial x} = \sum_{j=1}^{j=10} 2\pi \cos(\alpha_j) E_j \cos(\Phi_j) , \\ \frac{\partial Y}{\partial y} = \sum_{j=1}^{j=10} 2\pi \sin(\alpha_j) E_j \cos(\Phi_j) , \end{array} \right.$$

and hence the components \mathcal{I}_x and \mathcal{I}_y of \mathcal{I} are

$$\mathcal{I}_x = \sum_{j,k=1}^{j,k=10} (\mathcal{I}_x)_{j;k} , \quad \mathcal{I}_y = \sum_{j,k=1}^{j,k=10} (\mathcal{I}_y)_{j;k} ,$$

with

$$\left\{ \begin{array}{l} (\mathcal{I}_x)_{j;k} = 2\pi E_j E_k \cos(\alpha_j) \cos \left[\phi_j + 2\pi [x \cos(\alpha_j) + y \sin(\alpha_j)] \right] \\ \quad \times \cos \left[\phi_k + 2\pi [x \cos(\alpha_k) + y \sin(\alpha_k)] \right] \\ \quad + 2\pi E_j E_k \cos(\alpha_j) \sin \left[\phi_j + 2\pi (x \cos(\alpha_j) + y \sin(\alpha_j)) \right] \\ \quad \times \sin \left[\phi_k + 2\pi [x \cos(\alpha_k) + y \sin(\alpha_k)] \right], \\ (\mathcal{I}_y)_{j;k} = 2\pi E_j E_k \sin(\alpha_j) \cos \left[\phi_j + 2\pi [x \cos(\alpha_j) + y \sin(\alpha_j)] \right] \\ \quad \times \cos \left[\phi_k + 2\pi [x \cos(\alpha_k) + y \sin(\alpha_k)] \right] \\ \quad + 2\pi E_j E_k \sin(\alpha_j) \sin \left[\phi_j + 2\pi (x \cos(\alpha_j) + y \sin(\alpha_j)) \right] \\ \quad \times \sin \left[\phi_k + 2\pi [x \cos(\alpha_k) + y \sin(\alpha_k)] \right]. \end{array} \right.$$

Figure 4.66 already displayed above recovers the structure of the phase field Z with its isophase lines in orange and the current lines orthogonal to them, since the current \mathcal{I} is parallel to the gradient $\nabla\varphi$ of φ and hence orthogonal to the lines $\varphi = \text{const}$.

The vorticity $\Omega = \nabla \times \mathcal{I} / 2 = \nabla X \times \nabla Y$ is shown in Fig. 4.73 by the graph of the function

$$\omega(x, y) = \frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x},$$

for $x, y \in [0, 1]$. Figure 4.74 shows the position of the lines $\Omega = 0$ for $x, y \in [0, 3]$. We see that adjacent pinwheels do indeed have opposite chirality, since they belong to regions where Ω has opposite sign.

We know that, when $\Omega = 0$, either $\nabla X = 0$ or $\nabla Y = 0$ (isolated points of codimension 2), or the real gradients ∇X and ∇Y have the same orientation, i.e. $\nabla Y = \alpha \nabla X$ (see Sect. 4.6.5). Figure 4.75 shows the curves $\partial X / \partial x = \partial_x X = 0$ and $\partial_y X = 0$ in red and $\partial_x Y = 0$ and $\partial_y Y = 0$ in blue [the $\partial_x X$, etc., are denoted by $\tilde{X}(x)$, etc., in the label]. Note that points where $\nabla X = 0$ (intersection of red lines at top left) and $\nabla Y = 0$ (intersection of blue lines at top right) do indeed lie on the lines $\Omega = 0$. Likewise, the points where ∇X and ∇Y are both vertical ($\partial_x X = 0$ and $\partial_x Y = 0$) or both horizontal ($\partial_y X = 0$ and $\partial_y Y = 0$) also lie on the lines $\Omega = 0$ (intersections of red and blue lines at bottom left and bottom right).

Figure 4.76 shows the curves $\partial_x X = \alpha \partial_x Y$ and $\partial_y X = \alpha \partial_y Y$ (and their intersections, where $\nabla X = \alpha \nabla Y$) for different values of α . We see that these points do indeed lie on the lines $\Omega = 0$. Figure 4.77 shows the lines $\Omega = 0$ for $x, y \in [0, 1]$. It is interesting to look at the details. The typical situation is when two adjacent pinwheels are ‘close’, a line $\Omega = 0$ intersecting the isophase line that connects them




Fig. 4.73 Vorticity Ω of the phase field Z for $x, y \in [0, 1]$, i.e. the graph of the function $\omega(x, y) = \frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x}$

the most directly in a strongly transverse manner somewhere near the middle. This typical situation can be highly distorted for ‘distant’ adjacent pinwheels.

4.6.8 Current Conservation

Assuming that Z solves the Helmholtz equation, we consider the divergence of the current \mathcal{J} given by

$$\begin{aligned} \operatorname{div}(\mathcal{J}) &= \frac{\partial \mathcal{J}_x}{\partial x} + \frac{\partial \mathcal{J}_y}{\partial y} = \frac{\partial}{\partial x} \left(r^2 \frac{\partial \varphi}{\partial x} \right) + \frac{\partial}{\partial y} \left(r^2 \frac{\partial \varphi}{\partial y} \right) \\ &= r^2 \frac{\partial^2 \varphi}{\partial x^2} + 2r \frac{\partial r}{\partial x} \frac{\partial \varphi}{\partial x} + r^2 \frac{\partial^2 \varphi}{\partial y^2} + 2r \frac{\partial r}{\partial y} \frac{\partial \varphi}{\partial y} \\ &= r^2 \Delta \varphi + 2r \nabla r \cdot \nabla \varphi, \end{aligned}$$

where $\nabla r \cdot \nabla \varphi$ is the scalar product and the Laplacian operator $\Delta \varphi$ appears because it is by definition the divergence of the gradient. The Laplacian ΔZ of $Z = r e^{i\varphi}$ is given by




Fig. 4.74 Position of the vorticity lines $\Omega = 0$ for the phase field Z when $x, y \in [0, 3]$

$$\begin{aligned} \Delta Z &= \frac{\partial^2 Z}{\partial x^2} + \frac{\partial^2 Z}{\partial y^2} = \frac{\partial^2 Z}{\partial \rho^2} + \frac{1}{\rho} \frac{\partial Z}{\partial \rho} + \frac{1}{\rho^2} \frac{\partial^2 Z}{\partial \theta^2} \\ &= e^{i\varphi} \left[\Delta r - r |\nabla \varphi|^2 + i (r \Delta \varphi + 2 \nabla r \cdot \nabla \varphi) \right], \end{aligned}$$

so if $\Delta Z + k^2 Z = 0$, we must have

$$\begin{cases} \Delta r + r(k^2 - |\nabla \varphi|^2) = 0, \\ r \Delta \varphi + 2 \nabla r \cdot \nabla \varphi = 0. \end{cases}$$

The second equation expresses the fact that the divergence of the current is zero, that is $\operatorname{div}(\mathcal{J}) = 0$. This is a *conservation law*. It implies that \mathcal{J} can be written in the form $e_3 \times \nabla S = (-\partial S / \partial y, \partial S / \partial x)$. Indeed, since \mathbb{R}^2 is simply connected, a necessary and sufficient condition for this is $\partial^2 S / \partial x \partial y = \partial^2 S / \partial y \partial x$, i.e., $\partial \mathcal{J} / \partial x = -\partial \mathcal{J} / \partial y$, which says precisely $\operatorname{div}(\mathcal{J}) = 0$. We then have

$$\omega = \frac{1}{2} \left(\frac{\partial \mathcal{J}_y}{\partial x} - \frac{\partial \mathcal{J}_x}{\partial y} \right) = \frac{1}{2} \left[\frac{\partial}{\partial x} \left(\frac{\partial S}{\partial x} \right) + \frac{\partial}{\partial y} \left(\frac{\partial S}{\partial y} \right) \right] = \frac{1}{2} \Delta S.$$

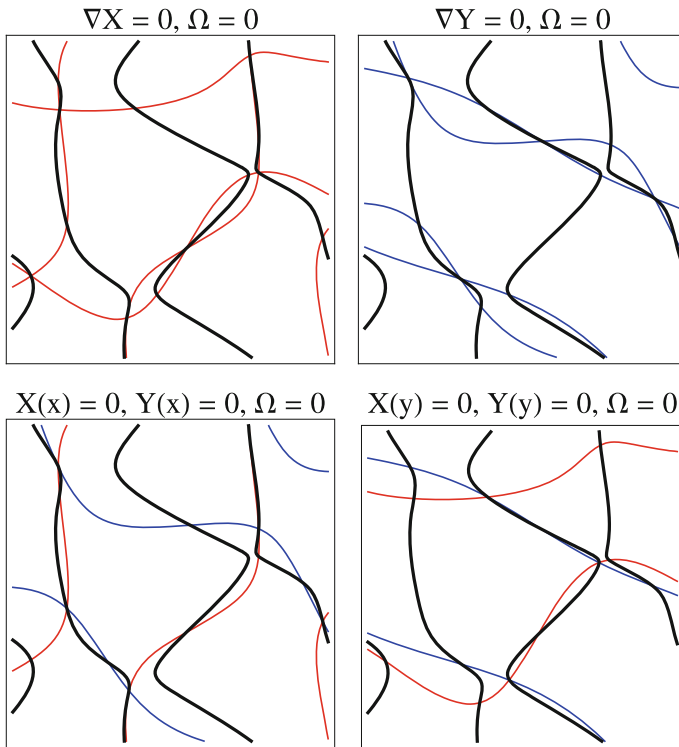


Fig. 4.75 Curves $\partial_x X = 0$ and $\partial_y X = 0$ (red) and $\partial_x Y = 0$ and $\partial_y Y = 0$ (blue). Top left points where $\nabla X = 0$ (intersection of red lines). Top right points where $\nabla Y = 0$ (intersection of blue lines). Bottom left points where ∇X and ∇Y are both vertical [intersections of red and blue lines $\partial_x X = X(x) = 0$ and $\partial_x Y = Y(x) = 0$]. Bottom right points where ∇X and ∇Y are both horizontal [intersections of red and blue lines $\partial_y X = X(y) = 0$ and $\partial_y Y = Y(y) = 0$]. All these points lie on the lines $\Omega = 0$

Regarding the first equation, it turns up everywhere in optics in a form where k is part of the phase, i.e. where $Z = re^{ik\varphi}$. It then becomes $\Delta r = k^2 r (|\nabla\varphi|^2 - 1)$. When $k \rightarrow \infty$, which corresponds to the geometrical optics approximation, it becomes the well-known *eikonal equation* $|\nabla\varphi|^2 = 1$, which expresses the fact that ‘light rays’, i.e. the trajectories of the field gradient $\nabla\varphi$, go around the singularity at a constant rate, while the wave fronts $\varphi = \text{const.}$ are the rays coming from the singularity (leading to some confusion between the two meanings of the word ‘ray’). To see this, we expand r asymptotically in powers of k , viz.

$$r(a) \sim k^\mu \sum_{n=0}^{n=\infty} k^{-n} r_n(a) ,$$




Fig. 4.76 Four examples of curves $\partial_x X = \alpha \partial_x Y$ and $\partial_y X = \alpha \partial_y Y$ (and their intersections where $\nabla X = \alpha \nabla Y$) for different values of α . *Top left* $\alpha = -1$ (blue). *Top right* $\alpha = -1/2$ (green). *Bottom left* $\alpha = 1/2$ (orange). *Bottom right* $\alpha = 1$ (red). We see that all points where $\nabla X = \alpha \nabla Y$ do indeed lie on the lines $\Omega = 0$

with $r_0(a) \neq 0$. When $k \rightarrow \infty$, we have $r(a) \rightarrow k^\mu r_0(a)$ and we deduce that $r_0(1 - |\nabla\phi|^2) = 0$, and hence $|\nabla\phi|^2 = 1$.

The quantity S can be calculated explicitly for superpositions of plane waves. Indeed, S is a sum of integrals like

$$\int 2\pi E_j E_k \cos(\alpha_j) \cos\left[\phi_j + 2\pi\left[x \cos(\alpha_j) + y' \sin(\alpha_j)\right]\right] \times \cos\left[\phi_k + 2\pi\left[x \cos(\alpha_k) + y' \sin(\alpha_k)\right]\right] dy',$$

which gives




Fig. 4.77 Lines of vorticity $\Omega = 0$ for the phase field Z when $x, y \in [0, 1]$

$$\frac{1}{2} E_j E_k \cos(\alpha_j) \left\{ \frac{\sin \left[2\pi \left[x \cos(\alpha_j) + y \sin(\alpha_j) - x \cos(\alpha_k) - y \sin(\alpha_k) \right] + \phi_j - \phi_k \right]}{\sin(\alpha_j) - \sin(\alpha_k)} \right. \\ + \frac{\sin \left[2\pi \left[y \sin(\alpha_j) + y \sin(\alpha_k) + x \cos(\alpha_j) + x \cos(\alpha_k) \right] + \phi_j + \phi_k \right]}{\sin(\alpha_j) + \sin(\alpha_k)} \\ - \frac{\sin \left[2\pi x \left[\cos(\alpha_j) + \cos(\alpha_k) \right] + \phi_j + \phi_k \right]}{\sin(\alpha_j) + \sin(\alpha_k)} \\ \left. - \frac{\sin \left[2\pi x \left[\cos(\alpha_j) - \cos(\alpha_k) \right] + \phi_j - \phi_k \right]}{\sin(\alpha_j) - \sin(\alpha_k)} \right\},$$

for $k \neq j$ and

$$\frac{1}{4} E_j^2 \cot(\alpha_j) \left[\sin \left[2\phi_j + 4\pi x \cos(\alpha_j) + 4\pi y \sin(\alpha_j) \right] \right. \\ \left. - \sin \left[2\phi_j + 4\pi x \cos(\alpha_j) + 4\pi y \sin(\alpha_j) \right] \right],$$

for $k = j$.




Fig. 4.78 Flow of the current \mathcal{J} shown with the level lines of S

Figure 4.78 shows the flow of the current \mathcal{J} with the level lines of S . We note that the dislocations do indeed occur in regions of highest vorticity (> 0 or < 0), but this does not mean that they coincide with the extrema of S . The definition of the latter involves second derivatives of the field Z :

$$\begin{cases} \frac{\partial \omega}{\partial x} = \frac{\partial Y}{\partial y} \frac{\partial^2 X}{\partial x^2} + \frac{\partial X}{\partial x} \frac{\partial^2 Y}{\partial x \partial y} - \frac{\partial Y}{\partial x} \frac{\partial^2 X}{\partial x \partial y} - \frac{\partial X}{\partial y} \frac{\partial^2 Y}{\partial x^2}, \\ \frac{\partial \omega}{\partial y} = \frac{\partial Y}{\partial y} \frac{\partial^2 X}{\partial x \partial y} + \frac{\partial X}{\partial x} \frac{\partial^2 Y}{\partial y^2} - \frac{\partial Y}{\partial x} \frac{\partial^2 X}{\partial y^2} - \frac{\partial X}{\partial y} \frac{\partial^2 Y}{\partial x \partial y}. \end{cases}$$

4.6.9 Critical Points

Apart from dislocation singularities where the phase φ is not defined, there are other geometrically interesting points giving structure to the field Z . As already noted, these are the *critical points* of φ where the gradient $\nabla\varphi$ vanishes, and hence also the current $\mathcal{J} = r^2\nabla\varphi$. At such a point a_0 , the first term in the Taylor expansion of φ

is second order. If δa is an increment in a , we have to second order $\varphi(a_0 + \delta a) \approx \varphi(a_0) + \delta a^T H_\varphi \delta a$, where δa is treated as a column vector, δa^T is its transpose, and H_φ is the symmetric 2×2 matrix of second partial derivatives of φ at a_0 , a matrix known as the *Hessian* of φ :

$$H_\varphi = \begin{pmatrix} \partial^2 \varphi / \partial x^2 & \partial^2 \varphi / \partial x \partial y \\ \partial^2 \varphi / \partial x \partial y & \partial^2 \varphi / \partial y^2 \end{pmatrix}.$$

The eigenvalues of the Hessian H_φ determine the type of critical point. They are real because of the symmetry of H_φ , generically distinct, and nonzero.¹⁴ If they are both strictly positive, then $\delta a^T H_\varphi \delta a > 0$ and a_0 is a minimum of φ . If they are both strictly negative, a_0 is a maximum. If they have opposite signs, then a_0 is a saddle point of φ . Pinwheel maps would not appear to have orientation maxima or minima. However, we have seen that there are many saddle points, mainly located at the centre of the cells of the pinwheel lattice.

There can also be critical points in the *amplitude* r^2 , where $\nabla r^2 = 0$. Since $r^2 = X^2 + Y^2$, we have $\nabla r^2 = 2(X\nabla X + Y\nabla Y)$. At a singularity that is a zero of Z , $X = Y = 0$ and the point is critical. However, there are other critical points where $X, Y \neq 0$. At these points, ∇X and ∇Y are necessarily parallel, so the vorticity $\Omega = \nabla \times \mathcal{J} / 2 = \nabla X \times \nabla Y$ is necessarily zero. We thus see that the critical points in the intensity include dislocations and the points where $\Omega = 0$. The critical points in the amplitude are invariant under gauge transformations and coordinate transformations.

4.6.10 Mesogeometry and Microphysics

The optical analogy which suggests treating orientation maps as phase fields and pinwheels as dislocation singularities in such fields is also useful for understanding the relations between different levels. In optics, there are three levels: geometric, wave, and quantum. In our analogy, the geometric level corresponds to the mesogeometric level involving the contact structures, symplectic structures, and sub-Riemannian structures that we shall discuss here in detail.

The wave level corresponds to what we have just been doing, identifying pinwheels with singularities in the phase fields. However, as noted by Berry [69], wave optics is an average over microphysical interactions described by quantum optics. In particular, the optical current is an energy flow whose trajectories are the level lines of S , in a certain sense a momentum density, giving the classical force on a small particle placed at a :

¹⁴Let $H_\varphi = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$. The eigenvalues are solutions of the quadratic equation $\text{Det}(H_\varphi - \lambda I) = 0$, which can be written $\lambda^2 - \lambda(a+c) + (ac - b^2) = 0$. The discriminant $(a+c)^2 - 4(ac - b^2) = (a-c)^2 + 4b^2$ is always non-negative, and the roots are therefore real.

The current gives the time-averaged force on small particles.

For its part, the phase gradient $\nabla\varphi$ gives the momentum imparted to the particle by impacts from individual photons. And as the probability of these impacts is r^2 , the average momentum is indeed $\mathcal{J} = r^2\nabla\varphi$.

We thus recover the assumption (see Sect. 4.3.2 of Chap. 4) that there is a micro-physics of elementary events for which the mesogeometry of the orientation maps is a kind of morphological skeleton. It is presumably the action potentials that play the role of the tiny particles.

4.6.11 Statistics of Pinwheels as Phase Singularities

The pinwheel maps like phase fields can come in many forms. It is thus interesting to carry out a *statistical investigation* on the basis of certain simplifying assumptions. Such studies have already been carried out in optics, in particular by Michael Berry and Mark Richard Dennis (see, e.g. [70, 71]). This is a topical subject bringing together work by Wolf and Geisel [72, 73], studies by Daniel Bennequin and coworkers, and also recent work by Citti, Sarti, and one of their doctoral students, Davide Barbieri [74].

In his thesis [68], Dennis gives precise results for superpositions of plane waves

$$Z = \sum_{\kappa} A_{\kappa} e^{i\kappa \cdot a}$$

with complex amplitudes $A_{\kappa} = E_{\kappa} e^{i\phi_{\kappa}}$, in particular in the *isotropic case*, i.e. where there is rotation invariance, so that the E_{κ} have a distribution depending only on the magnitude $k = |\kappa|$ of the wave vectors (the wave number) and where the spatial phases ϕ_{κ} are random variables uniformly distributed on $[0, 2\pi]$. If the sampling of the κ in the given sums Z is fine enough, we can consider that the statistics of the components X and Y of Z and their partial derivatives are circular Gaussian distributions, which makes the calculations more accessible. In particular, we define the energy spectrum by

$$\frac{1}{2} \sum_{\kappa} E_{\kappa}^2 = \int \mathcal{E}(\kappa)^2 d\kappa,$$

and the radial energy spectrum by

$$\frac{1}{2} \sum_{\kappa} E_{\kappa}^2 = \int \frac{\mathcal{R}(k)^2}{2\pi k} dk.$$

A further simplification comes by considering ‘monochromatic’ waves, with the same wave number k , where the wave vector κ thus varies over a circle of radius k .

In this case, $\mathcal{R}(u)$ becomes the Dirac delta $\delta(u - k)$. This hypothesis corresponds to the fact that Z is a solution of the Helmholtz equation.

We can then calculate the average density d of phase dislocations. As these are defined by the conditions $X = 0, Y = 0$, it will be given by the average of $\delta(X) \delta(Y)$ with respect to the measure $dXdY$. Relative to the measure $dx dy$, we need to insert the Jacobian of $Z(x, y) = X(x, y) + iY(x, y)$, i.e.

$$|\omega| = |\nabla X \wedge \nabla Y| = \left| \frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x} \right|.$$

We must therefore calculate the average

$$\left\langle \delta(X) \delta(Y) \left| \frac{\partial X}{\partial x} \frac{\partial Y}{\partial y} - \frac{\partial X}{\partial y} \frac{\partial Y}{\partial x} \right| \right\rangle.$$

Assuming that X, Y , and their partial derivatives are independent Gaussian random variables, we can do this using the integrals

$$\delta(u) = \frac{1}{2\pi} \int e^{itu} dt, \quad |u| = -\frac{1}{\pi} \int \frac{\partial}{\partial s} (e^{isu}) \frac{ds}{s},$$

taking the latter in the sense of the Cauchy principal value. We thus obtain

$$d = \frac{K}{4\pi}, \quad K = \int_0^\infty k^2 \mathcal{R}(k) dk = \langle k^2 \rangle_{\mathcal{R}} \text{ for the measure } \mathcal{R}(k) dk.$$

As noted by Michael Berry, the fact that Z is a superposition of waves barely comes into the calculation, and [70, p. 2076]:

The results apply to any complex scalar random function. [...] The geometry thus revealed is extraordinarily complicated and occasionally counterintuitive.

The wave number k is proportional to the reciprocal of a wavelength $\Lambda = 2\pi/k$, whence $\Lambda^2 = 4\pi^2/k^2$ and $k^2/4\pi = \pi/\Lambda^2$. Consequently, the density of singularities d is the average $\langle \pi/\Lambda^2 \rangle_{\mathcal{R}}$. In Sect. 4.7.1.2, we shall once again come across this term π/Λ^2 , also found by Fred Wolf and Theo Geisel.

4.6.12 Pinwheels and Gaussian Fields

Quite generally, the orientation maps can be treated as *random sections* of the fibre bundle $\mathbb{R}^2 \times P \rightarrow \mathbb{R}^2$ satisfying a set of constraints explaining their pinwheel geometrical structure. The problem then is to calculate their dislocation distribution in statistical terms. At each point a of the base space \mathbb{R}^2 , we thus consider a random variable \mathcal{L}_a , which defines a random field \mathcal{L} whose orientation maps $Z(a)$ are

samples. To simplify, we generally assume that the field \mathcal{Z} is *Gaussian* (RGF), i.e. that the \mathcal{Z}_a are Gaussians with mean $m_a = \mathbb{E} \{ \mathcal{Z}_a \}$, where \mathbb{E} is the expectation, and variance $\sigma_a^2 = \mathbb{E} \{ (\mathcal{Z}_a - m_a)^2 \}$, and that all the joint distributions $\sum_i \alpha_i \mathcal{Z}_{a_i}$ for a finite number of points a_i are also Gaussian. Moreover, it is natural to assume that the distribution of the \mathcal{Z}_a is *SE* (2)-invariant [the distribution of the \mathcal{Z}_a and obviously not the samples $Z(a)$]. Translation invariance is known as *stationarity* and rotation invariance as *isotropy*. For an introduction to random Gaussian fields, the reader could consult, for example, Petter Abrahamsen's review [75].

In the paper [63] on Gaussian models of pinwheels, Alexandre Afgoustidis carried out a numerical calculation of the *variance* of the number $N(J, \theta)$ of neurons with preferred orientation θ along a line segment J in the plane V1. His calculation used a formula due to Cramer and Leadbetter which contains oscillating integrals that are rather difficult to calculate. The main result is that the variance is minimal when the power spectrum is concentrated on a circle.

While preparing the paper, the author had an interesting discussion with one of the referees (I was aware of this, being a referee myself). For the initial orientation field, the assumption of an RGF model is plausible. However, we may wonder whether it remains so for the stable orientation fields obtained by learning. Indeed, the role of long-range horizontal connections becomes crucial and changes the statistical properties of the fields. In his thesis [64], Afgoustidis discusses this point with Fred Wolf.

The random variables \mathcal{Z}_a cannot be decorrelated (i.e. independent, since for Gaussian random variables, independence and decorrelation are equivalent), because otherwise there would only be Gaussian noise and no geometrical structure. Furthermore, the very definition of a *continuum* of independent Gaussian random variables raises some tricky questions. What characterizes the field \mathcal{Z} is the *correlation function* $C(a, b) = \mathbb{E} \{ (\mathcal{Z}_a - m_a) (\mathcal{Z}_b - m_b) \}$. Dividing by the variances, we obtain the normalized correlation function $\Gamma(a, b) = C(a, b) / \sigma_a \sigma_b$. Stationarity implies that $C(a, b) = C(a - b)$, and isotropy implies in addition that $C(a, b) = C(a - b) = C(\|a - b\|) = C(r)$. The means are all equal, i.e. $m_a = m$, and so are the variances, i.e., $\sigma_a^2 = \sigma^2 = C(0)$, and $\Gamma(r) = C(r) / C(0)$.

The correlation functions are rather special, being symmetric and non-negative definite:

- (i) In the case of stationarity, a theorem due to Salomon Bochner tells us that they admit a spectral representation which is a generalized Fourier transform. This implies that $\Gamma(a) = \int_{\mathbb{R}^2} e^{i\langle a, \kappa \rangle} dF(\kappa)$ ¹⁵, where F is a bounded non-negative measure on the space of wave vectors κ dual to the positions a . We have $F(\mathbb{R}^2) = \Gamma(0) = 1$. In polar coordinates (k, α) for κ , we have

$$\Gamma(a) = \int_{\mathbb{R}^+ \times \mathbb{S}^1} e^{ik[x \cos(\alpha) + y \sin(\alpha)]} d\alpha dP(k) ,$$

for a measure $P(k)$ on \mathbb{R}^+ .

¹⁵ $\langle \cdot, \cdot \rangle$ is the natural pairing between vectors and dual covectors. It can also be expressed by a dot.

- (ii) If the measure F is smooth enough relative to the Lebesgue measure $d\kappa$, then it has a spectral density $f(\kappa)$ and the generalized FT in Bochner’s theorem reduces to a Fourier transform

$$\Gamma(a) = \int_{\mathbb{R}^2} e^{i(a,\kappa)} f(\kappa) d\kappa ,$$

with inverse transform

$$f(\kappa) = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} e^{-i(a,\kappa)} \Gamma(a) da .$$

- (iii) If there is also isotropy, then $\Gamma(r) = \int_0^\infty J_0(kr) kf(k) dk$ with $k = \|\kappa\|$, where J_0 is the Bessel function.
- (iv) If in addition, we consider the solutions of the Helmholtz equation with wave number k_0 , then $f(k)$ is proportional to $\delta(k - k_0)$ and $\Gamma(r)$ is proportional to $J_0(kr) k_0$.

In this context, the formula for the statistics of the dislocations in the fields $Z(a)$ sampling the random Gaussian field \mathcal{Z} is a special case of a fundamental formula called the *Kac–Rice formula* (from Kac and Rice) [76]. We would like to calculate the average $d = \mathbb{E} \{ \# \{ a \in T : Z(a) = 0 \} \}$ (recall that # indicates the cardinality of a set) of the number of zeros of Z in a unit square T . Let N_T be this number. We thus have $d = \mathbb{E} \{ N_T \}$. The Rice formula tells us that

$$d = \int_T \mathbb{E} \{ |\det(\text{Jac}(\mathcal{Z}_a))| : \mathcal{Z}_a = 0 \} p_{\mathcal{Z}_a}(0) da ,$$

where Jac is the Jacobian and $p_{\mathcal{Z}_a}$ the density of \mathcal{Z}_a . The calculation done by Dennis [68] involved working out Gaussian integrals. Let \mathcal{X}_a and \mathcal{Y}_a be the components of \mathcal{Z}_a and $\mathcal{J}_a = |\det(\text{Jac}(\mathcal{Z}_a))|$. It can be shown that the 6 random variables

$$\mathcal{X}_a , \quad \mathcal{Y}_a , \quad \left(\frac{\partial \mathcal{X}}{\partial x} \right)_a , \quad \left(\frac{\partial \mathcal{X}}{\partial y} \right)_a , \quad \left(\frac{\partial \mathcal{Y}}{\partial x} \right)_a , \quad \left(\frac{\partial \mathcal{Y}}{\partial y} \right)_a ,$$

are independent Gaussian variables, each with a distribution of the form

$$\frac{1}{\sqrt{2\pi}\sigma} e^{-\xi^2/2\sigma^2} .$$

As we have seen, we then have to evaluate an integral of the form

$$\int \delta(X) \delta(Y) J_p d(X) d(Y) d\left(\frac{\partial X}{\partial x}\right) d\left(\frac{\partial X}{\partial y}\right) d\left(\frac{\partial Y}{\partial x}\right) d\left(\frac{\partial Y}{\partial y}\right) ,$$

where p is the product of the distributions. The first two variables \mathcal{X}_a and \mathcal{Y}_a have variance 1, which introduces a factor $(1/\sqrt{2\pi})^2 = 1/2\pi$ in the integral, and the other four have variance $\sigma^2 = K/2$, which introduces a factor $(1/\sqrt{2\pi}\sigma)^4 = (1/2\pi\sigma^2)^2 = 1/(\pi K)^2$. The condition $\mathcal{Z}_a = 0$ amounts to putting $\delta(X)\delta(Y)$ in the integral, but $\int \delta(X) e^{-X^2/2} = 1$ and the same for Y . It thus remains to evaluate

$$\frac{1}{2\pi} \frac{1}{(\pi K)^2} \int J e^{-(\|\nabla X\|^2 + \|\nabla Y\|^2)/K} d\left(\frac{\partial X}{\partial x}\right) d\left(\frac{\partial X}{\partial y}\right) d\left(\frac{\partial Y}{\partial x}\right) d\left(\frac{\partial Y}{\partial y}\right).$$

If we change to polar coordinates by writing $\nabla X = R_X e^{i\psi_X}$ and $\nabla Y = R_Y e^{i\psi_Y}$, the integral becomes

$$\frac{1}{2\pi} \frac{1}{(\pi K)^2} \int_{R_X=0}^{R_X=\infty} \int_{R_Y=0}^{R_Y=\infty} \int_{\psi_X=0}^{\psi_X=2\pi} \int_{\psi_Y=0}^{\psi_Y=2\pi} (R_X)^2 (R_Y)^2 |\sin(\psi_Y - \psi_X)| e^{-(R_X^2 + R_Y^2)/K} dR_X dR_Y d\psi_X d\psi_Y.$$

The integral of the sine gives 8π and the integrals of R_X and R_Y each give $K^{3/2}\sqrt{\pi}/4$. So finally,

$$d = \frac{1}{2\pi} \frac{1}{(\pi K)^2} 8\pi \frac{1}{16} K^3 \pi = \frac{K}{4\pi}.$$

When the wave vectors κ are concentrated on a circle of radius $k_0 = 2\pi/\Lambda$ in the Fourier space, the Gaussian random variables

$$\left(\frac{\partial \mathcal{X}}{\partial x}\right)_a, \quad \left(\frac{\partial \mathcal{X}}{\partial y}\right)_a, \quad \left(\frac{\partial \mathcal{Y}}{\partial x}\right)_a, \quad \left(\frac{\partial \mathcal{Y}}{\partial y}\right)_a,$$

with their distribution

$$\frac{1}{\sqrt{2\pi}\sigma} e^{-\xi^2/2\sigma^2},$$

which has variance $\sigma^2 = K/2$ satisfy $K = k_0^2$, and we thus obtain, as stated above,

$$d = \frac{K}{4\pi} = \frac{k_0^2}{4\pi} = \left(\frac{2\pi}{\Lambda}\right)^2 \frac{1}{4\pi} = \frac{\pi}{\Lambda^2}.$$

These statistical calculations, which only give one particularly simple example of the connection between statistics and geometry, are especially interesting from a theoretical point of view (and not only numerically) for the following reason. In their 2005 reference book *Random Fields and Geometry* [77], Robert Adler and Jonathan

Taylor studied in great depth the generalizations of the Kac–Rice formula for random fields \mathcal{F}_a defined on a base space M and with values in \mathbb{R}^k .

Let us take, for example, $k = 1$. One of the main problems, and an extremely difficult one, is to calculate $\mathbb{P}\{\sup_{a \in M} \mathcal{F}_a \geq u\}$ for large u . This so-called probability of excursion in the interval $[u, \infty)$ is well approximated by $\mathbb{E}\{\chi(A_{[u, \infty)})\}$, where quite generally, if D is a domain of \mathbb{R}^k , $A_D := \{a \in M : \mathcal{F}_a \in D\}$, and where χ is the *Euler–Poincaré characteristic*. Assuming Gaussianity, stationarity, isotropy, and smoothness of the correlation functions C , we can obtain explicit but complicated formulas for the $\mathbb{E}\{\chi(A_D)\}$. What is interesting is that the proof of these formulas requires all the fundamental tools of the Morse–Whitney–Thom ‘philosophy’ discussed in Sect. 1.3 of the *Preface*.

To begin with, in order to be able to handle enough cases, we assume that M is a manifold with boundary of dimension N , equipped with a ‘good’ stratification $M = \cup_{k=0}^{k=N} \partial_k M$, called a Whitney stratification, satisfying Whitney’s properties A and B . Here, k is the dimension of the strata making up $\partial_k M$, with $\partial_N M = \overset{\circ}{M}$, where $\overset{\circ}{M}$ is the interior of M , and $\partial_0 M = \{\text{vertices of } M\}$. We assume that the smoothness properties of C imply that the samples F of the field \mathcal{F} are *Morse functions* on M , using the generalization of Morse theory to stratified manifolds due to Robert MacPherson. In addition, the field \mathcal{F} defines a *natural metric* $d_{\mathcal{F}}(a, b)$ on the base space M through the formula

$$d_{\mathcal{F}}^2(a, b) = \mathbb{E}\{\|\mathcal{F}_a - \mathcal{F}_b\|^2\}.$$

We can thus also make use of the resources of Riemannian geometry, such as curvature tensor, Levi-Civita connection, covariant derivative, Lipschitz–Killing curvatures, and so on.

For $k = 1$, we thus apply to $A_{[u, \infty)}$ the formulas relating Morse theory to the Euler–Poincaré characteristic. If the sample F of \mathcal{F} is Morse and if u is a regular, i.e. non-critical, value of F , then $A_{[u, \infty)}$ is a sub-manifold with boundary of M that is ‘well stratified’ by the intersection strata $\overset{\circ}{A}_{[u, \infty)} \cap \partial_k M$ and $\partial A_{[u, \infty)} \cap \partial_k M$. F is not necessarily Morse on $A_{[u, \infty)}$, but it can be approximated by a Morse function \tilde{F} whose critical points correspond to the critical points of F situated above u .

The Rice–Kac formula corresponds to a rectangle $M = T$ of \mathbb{R}^N and to \mathcal{F} with values in \mathbb{R}^N . Let J be the Jacobian of a sample F (an $N \times N$ matrix). Let N_u be the number of points of T for which $F(a) = u \in \mathbb{R}^N$. Then, the formula tells us that

$$\mathbb{E}\{N_u\} = \int_T \mathbb{E}\{|\det(J)| : F(a) = u\} p_a(u) dt.$$

In our case $N = 2$, $\mathcal{F} = \mathcal{L}$, $u = 0$, and T is a unit square.

4.6.13 Evolution of Pinwheels as Phase Singularities

We can also study the temporal evolution of the pinwheels, that is, the singularities in the field Z , e.g. by assuming that Z is a superposition of plane waves $Z = \sum_{\kappa} A_{\kappa} e^{i(\kappa \cdot a - \omega t)}$, where, as we saw in Sect. 4.6.6, ω is an angular frequency associated with the frequency $\nu = \omega/2\pi$ and the period $T = 1/\nu = 2\pi/\omega$, the corresponding wavelength being $\lambda = cT = c/\nu = 2\pi c/\omega$, where c is the speed of propagation.

Topological accidents can occur during this evolution, in particular when a singularity $Z = 0$ also satisfies $\Omega = \omega e_3 = 0$. We saw in Sect. 4.6.5 that, when $\Omega = 0$, the real ∇X et ∇Y have the same orientation, $\nabla Y = \alpha \nabla X$, $\alpha \in \mathbb{R}$, and hence that the complex gradient ∇Z is the real vector ∇X multiplied by a factor $1 + i\alpha$. As explained by Michael Berry and Mark Richard Dennis in [71], by a gauge transformation and a coordinate change, we can reduce to the case in which the field $Z(a)$ has the form

$$Z(a) = i\beta y + \frac{1}{2} a^T H a, \quad \beta \in \mathbb{R},$$

where the complex Hessian

$$H = \begin{pmatrix} H_{11} & H_{12} \\ H_{12} & H_{22} \end{pmatrix}$$

is a symmetric 2×2 complex matrix. We do indeed have $Z(0) = 0$, and since $\nabla X(0) = 0$, $\Omega(0) = 0$.

This unstable normal form can be unfolded using a time variable t , giving the dynamical model

$$\begin{aligned} Z(a, t) &= t + i\beta y + \frac{1}{2} a^T H a, \\ X(a, t) &= t + \frac{1}{2} [x^2 \operatorname{Re}(H_{11}) + 2xy \operatorname{Re}(H_{12}) + y^2 \operatorname{Re}(H_{22})], \\ Y(a, t) &= \beta y + \frac{1}{2} [x^2 \operatorname{Im}(H_{11}) + 2xy \operatorname{Im}(H_{12}) + y^2 \operatorname{Im}(H_{22})]. \end{aligned}$$

For the singularities $Z = 0$, the equation $Y = 0$ gives, for the lowest order terms, $y \sim -x^2 \operatorname{Im}(H_{11})/2\beta$ and the equation $X = 0$ gives $x^2 \operatorname{Re}(H_{11}) \sim -2t$. If $\operatorname{Re}(H_{11}) > 0$, then for $t < 0$, there are two solutions with abscissa values $x = \pm [-2t/\operatorname{Re}(H_{11})]^{1/2}$ on the parabola $y = -x^2 \operatorname{Im}(H_{11})/2\beta$ which coalesce and disappear for $t \geq 0$: two singularities with opposite chiralities $x\beta \operatorname{Re}(H_{11})$ annihilate in what is known as a fold catastrophe. If on the other hand $\operatorname{Re}(H_{11}) < 0$, then when t becomes positive, two singularities of opposite chirality appear.

We shall return to these pinwheel bifurcations in more detail in Sect. 4.7.1.2.

4.7 Pinwheel Singularities

We have seen that pinwheels are centred on singularities. In this section, we shall discuss some of the experimental data which supports the hypothesis that the functional orientation maps can be treated as phase fields.

4.7.1 Structure in the Vicinity of Singularities

4.7.1.1 Pinwheels as Zeros of the Orientation Field

If we enter the cortex ‘vertically’ at a regular point, we come across the columns of Hubel and Wiesel with their redundancy and their population coding. However, when we enter at a singular point, we come across neurons of every orientation. It is tricky but crucial to understand this phenomenon correctly. Orientation selectivity is highly variable in V1, and clearly, on the level of the mesoscopic resolution of the images presented so far, it is necessarily rather low or even zero on average at singular points, since all orientations are present in a roughly uniform manner. In fact, we may assume that the meso- or ‘coarse-grained’ level is one where the preferred orientation corresponds to what is known as an *order parameter* in statistical physics, and that the models of this level are analogous to the so-called *mean field* models.

As an example of *mesoscopic models*, let us return to those proposed by Wolf and Geisel [72] to understand the *learning process* behind orientational selectivity (see Sect. 4.6.1). They are phase field models in the sense of Sect. 4.6. The field of preferred orientations $\Phi(a) = e^{i\varphi(a)}$ of simple V1 neurons can be treated as a section of the fibre bundle $\pi : R \times P \rightarrow R$ defined outside the lattice L of the pinwheel centres of R . Identifying R with \mathbb{R}^2 and \mathbb{R}^2 with \mathbb{C} , Wolf and Geisel suggested modelling the columns by a continuous *complex field* $Z(a) = r(a)e^{i\varphi(a)} = r(a)\Phi(a)$ on the complex variable $a = \rho e^{i\theta}$, where the spatial phase $\varphi(a)$ thus encodes the orientation preference $\psi = \varphi/2$ and where the modulus $r(a) = |Z(a)|$ encodes the strength of selectivity, i.e. the width of the response curve, which is Gaussian. The complex value Z is thus a section of the fibre bundle $\pi : \mathbb{C} \times \mathbb{D} \rightarrow \mathbb{C}$, which has base space \mathbb{C} , and if the maximal selectivity is normalized to unity, its fibre is the unit disc \mathbb{D} . As we have seen, the singular points (pinwheel centres) then correspond to the *zeros* of Z . To obtain a naive and simple illustration of the behaviour of the field Z in the neighbourhood of a zero, we take $\varphi = \theta$ and $r = \rho/2$. Above the points a of a small circle C_ρ with centre 0 and radius ρ fixed in the base space \mathbb{C} , we have in the fibre the circles $\Sigma_{\rho/2}$ of fibres \mathbb{D}_a of radius $\rho/2$, and the inverse image of C_ρ under π is thus the torus $\pi : C_\rho \times \Sigma_{\rho/2} \rightarrow C_\rho$. The lift of C_ρ on this torus is the curve Γ_ρ (see Fig. 4.79) given parametrically by

$$\Gamma_\rho = \left(\frac{\rho}{2} \sin(\theta), \rho \left[1 - \frac{1}{2} \cos(\theta) \right] \cos(\theta), \rho \left[1 - \frac{1}{2} \cos(\theta) \right] \sin(\theta) \right).$$

Fig. 4.79 The torus
 $\pi : C_\rho \times \Sigma_{\rho/2} \rightarrow C_\rho$ and
 the lift Γ_ρ of C_ρ (see text)





Fig. 4.80 When $\rho \rightarrow 0$, the
 section $\Gamma_\rho \rightarrow 0$



As the orientation selectivity vanishes at 0, when $\rho \rightarrow 0$, the section Γ_ρ also tends to 0 and the projection π is locally a diffeomorphism (see Fig. 4.80).

4.7.1.2 Pinwheel Morphogenesis and Learning

One of the advantages of this approach is that it provides a good theory of the learning process. The paper [78] by Nicholas Swindale contains a summary of these development models. As the author puts it:

The repetitive stochastic patterns of eye dominance and orientation preference [are of an] intriguing nature. Many aspects of their development seem likely to be dependent upon both spontaneous and visually driven patterns of neural activity. [...] Remarkably simple models, based on Hebbian synaptic plasticity, intracortical interactions and competitive interactions between cells and growing axons, have been able to explain much of the phenomenology.

A good example of such models is the one developed by James Bednar and Risto Miikkulainen in [79, 80] using self-organized map (SOM) models (see Sect. 4.10.1.4), which are¹⁶:

[...] networks of simple artificial neurons with initially unspecific connections that are modified by Hebbian learning and homeostatic plasticity.

They show quite clearly how the double constraints imposed on the one hand by the geometry and statistics of *external stimuli* and, on the other, by the geometry and statistics of the *internal structure* of the cortical areas lead in a self-organized way to pinwheel maps that depend only weakly on the initial synaptic weights of the neurons, but heavily on the way the stimuli arrive. Figure 4.81 gives an example.

The dynamical self-organization models proposed by Fred Wolf, Theo Geisel, Matthias Kashube, and Michael Schnabel in [72, 81, 82] are models of pattern formation, starting from an unstructured initial state, through *Turing instabilities*. They




Fig. 4.81 Pinwheel map built from GCAL self-organized maps (gain control, adaptive, laterally connected SOMs). Starting from random synaptic weights, the neural network is subjected to a flow of stimuli of a certain type. The weights are modified by Hebbian learning and homeostatic plasticity, gradually producing a pinwheel structure which stabilizes. It can be shown that such maps depend little on the initial synaptic weights of the neurons, but a great deal on the flow of stimuli. In the last figure, *white lines* indicate the inhibitory lateral connections of a neuron. Figure adapted from Bednar [79]

¹⁶For Hebb's law, see Sect. 3.6.2.

adopt the development model first introduced by Nicholas Swindale and describe the morphogenesis of the field by a partial differential equation (PDE) of the form

$$\frac{\partial Z(a, t)}{\partial t} = F(Z(a, t)) + \eta(a, t) ,$$

where F is a linear operator and η a stochastic term representing the intrinsic fluctuations associated with this activity. This kind of dynamics can be induced by a Hebbian learning process, with $F(Z(a))$ the average of $Z(a)$ for rapidly changing stimuli A with a certain probability distribution.

The authors also take into account ocular dominance and distortion of retinotopy (see Figs. 4.12 and 4.13 in Sect. 4.2). They show that, starting from an initial state in which the field $Z(a, 0)$ is weak (with little orientational selectivity), there is first a period of proliferation of pinwheels, with a growth in selectivity and the appearance of a characteristic wavelength Λ , the pinwheel density¹⁷ becoming greater than π/Λ^2 . Then, the nonlinearities and long-range lateral interactions stabilize the process and the number of pinwheels through their displacements, collisions, and annihilations of pairs of pinwheels of opposite chirality (see Fig. 4.82). The reader could consult Ha Youn Lee et al. [83].

More precisely, the standard technique for analyzing such a PDE is to expand $F(Z)$ in a series

$$F(Z) = L(Z) + F_2(Z, \bar{Z}) + F_3(Z, Z, \bar{Z}) + \dots ,$$

where $L(Z)$ is linear, $F_2(Z, \bar{Z})$ is bilinear, $F_3(Z, Z, \bar{Z})$ is trilinear, and so on.¹⁸ In the given model, it can be shown that the linearized $L(Z(a, t))$ whose spectral analysis is to be carried out is of fourth order and can be written $L = \mu - (k_c^2 + \Delta)^2$, where μ is a bifurcation parameter such that, for $\mu < 0$, the homogeneous base state $Z \equiv 0$ is stable, while for $\mu > 0$, it becomes unstable and bifurcates towards




Fig. 4.82 Annihilation of pinwheels of opposite chirality simplifies and stabilizes the pinwheel geometry. From Wolf and Geisel [72]

¹⁷We have already encountered this density π/Λ^2 in Sect. 4.6.11.

¹⁸We have to use the two variables Z and \bar{Z} because the functions F_j are not necessarily analytic functions depending only on Z .

patterns with characteristic wavelength $\Lambda = 2\pi/k_c$ for critical eigenvalues $\lambda(k_c)$ ($k_c \neq 0$) of the Fourier representation of L given by $\mu = \lambda(k_c)$.¹⁹ To third order, the dynamical PDE becomes analogous to the equation that Jack Swift and Pierre Hohenberg introduced into the study of convection in hydrodynamics (see [84]).

It is important to note that the PDE must be *equivariant* under the action of the group of symmetries of the plane $SE(2)$, and even $E(2)$. We shall return at length in the second volume on the structure of the group $SE(2)$ and the importance of symmetries. Let us just say for the moment that they impose strong constraints on the form of the functional F . Indeed, F must be equivariant under the action of $SE(2)$, acting by translations $a \rightarrow a + b$ and rotations $a \rightarrow R_\theta(a)$ on the base plane $R = \mathbb{C}$ of the a , and also by rotations $Z \rightarrow e^{i\psi}Z$ in the fibres. Now recall that, if $Z(a)$ is a function of a variable $a \in R$ and if T is an element of a group G of transformations of R , T acts on $Z(a)$ by $\widetilde{T}(Z(a)) = Z(T^{-1}(a))$. In addition, we say that a function or functional $F(Z)$ on which G operates is G -equivariant if $T(F(Z)) = F(\widetilde{T}(Z))$, i.e. if F commutes with the actions of G on Z and on the space of values of F . Consequently, we must have

$$\widetilde{b}(F(Z(a))) = F(\widetilde{b}(Z(a))) = F(Z(a - b)) ,$$

for any translation \widetilde{b} by a vector b in the plane R ,

$$\widetilde{R}_\theta(F(Z(a))) = F(\widetilde{R}_\theta(Z(a))) = F(Z(R_{-\theta}(a))) ,$$

for any rotation \widetilde{R}_θ through angle θ of the plane R , and $F(e^{i\psi}Z) = e^{i\psi}F(Z)$ for any rotation of Z . Note that the last symmetry implies that $F(0) = 0$, since if $Z = 0$, then $e^{i\psi}Z = 0$ for any ψ , whence $e^{i\psi}F(0) = 0$ for all ψ .

When F is expanded in the neighbourhood of the equilibrium state $Z \equiv 0$, its linear part L commutes with the symmetries, and the Fourier representation is therefore diagonal with eigenvalues $\lambda(k_c)$ which depend only on the magnitude k_c . The authors show that the second-order term F_2 can be neglected and that we may choose a third-order term F_3 depending on $|Z|^2 Z$ and long-range lateral connections. We shall see in the second volume that this cubic normal form $|Z|^2 Z$ often arises in problems with $E(2)$ symmetry.

The analysis of patterns $Z(a)$, i.e. pinwheel maps, that can emerge by bifurcation thus reduces to a classic problem that has been widely studied in physics, the problem of *symmetry breaking*. If a modulus (wave number) k is selected by a bifurcation at μ , we may consider superpositions of eigenstates (Fourier modes) with eigenvectors (wave vectors) κ_j with magnitude $|\kappa_j| = k$. We thus obtain *planforms*:

¹⁹In the expansion of L , the powers of partial derivatives like $(\partial/\partial x)^n$ mean repeated differentiation $\partial^n/\partial x^n$. The linear PDE $\partial Z/\partial t = L(Z)$ for $L = \mu$ describes exponential damping towards 0 for $\mu < 0$ (stability) and exponential growth for $\mu > 0$ (instability). The PDE $\partial Z/\partial t = L(Z)$ for $L = -\Delta$ is a diffusion equation.

$$Z(a) = \sum_{j=0}^{j=n-1} A_j e^{i\kappa_j \cdot a}, \quad |\kappa_j| = k, \quad \kappa_j = k \left(\cos\left(\frac{2\pi j}{n}\right), \sin\left(\frac{2\pi j}{n}\right) \right).$$

The symmetries impose a precise structure on the amplitudes A_j . By expressing the fact that they are solutions of the PDE satisfied by Z , we show that they must satisfy differential equations of the form

$$\frac{dA_i}{dt} = A_i - \sum_{j=1}^{j=n} g_{ij} |A_j|^2 A_i - \sum_{j=1}^{j=n} f_{ij} A_i A_{j^*} \overline{A_{i^*}},$$

where j^* , the mode antiparallel to the mode j , is defined by $\kappa_{j^*} = -\kappa_j$, and where the coefficients g_{ij} and f_{ij} satisfy the properties

$$\begin{cases} g_{ij} = \left(1 - \frac{1}{2}\delta_{ij}\right) g(|\alpha_i - \alpha_j|), \\ f_{ij} = (1 - \delta_{ij} - \delta_{i^*j}) f(|\alpha_i - \alpha_j|). \end{cases}$$

Here, δ_{ij} is the Kronecker symbol ($\delta_{ii} = 1$ and $\delta_{ij} = 0$ if $i \neq j$), α_i is the angle of mode i , and $g(\alpha)$ and $f(\alpha)$ are π -periodic angle functions calculated from F . A refined version of this model can be found in the paper Schnabel et al. [85].

Figure 4.83, taken from [81, 82], shows the planforms for $n = 1, 2, 3, 5, 15$, with the positions of the wave vectors κ_j and the density of the pinwheels as a function of n . An example of a pinwheel map constructed as a linear combination of planforms is also shown.

4.7.1.3 Pinwheels as Genuine Singularities

These models based on the parallel between pinwheels and dislocations are quite remarkable. However, two points should be emphasized:

1. The model assumes that the orientation selectivity vanishes on the dislocations located at the centres of the pinwheels.
2. No meaning can be attached to the limit as the mesh of the pinwheel lattice tends to 0, since it then gives a field Z that is identically zero.

It thus requires closer examination because it may be that the selectivity remains good at singular points on the *micro* level of individual neurons and only vanishes at the *meso* level due to averaging. As noted by Jonathan Polimeni et al. when commenting on a paper by Pedro Maldonado et al. [86] which we shall discuss in a moment [87, p. 4158]:

The optical recording signal is a population average, and so individual neurons in the pinwheel centres—even if they were strongly tuned to orientation—would average to a weak population response at the pinwheel centres imaged through optical recording.

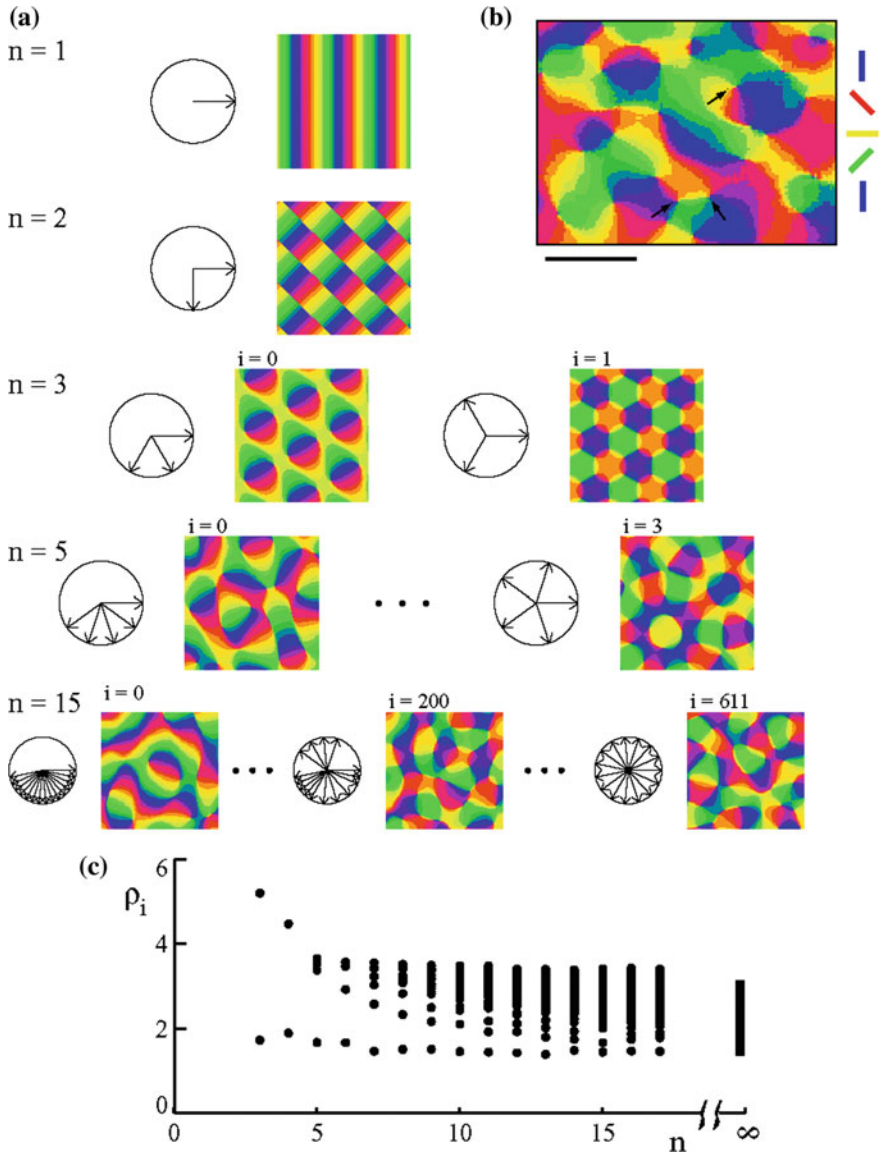


Fig. 4.83 Model by Wolf and Geisel. **a** Planforms for $n = 1, 2, 3, 5, 15$, with the position of the wave vectors κ_j . **b** Example of a pinwheel map built as a linear combination of planforms. **c** Diagram giving the density ρ of the pinwheels as a function of n : for a given n , each point is the density of one of the patterns corresponding to n . From [82]

The difference is clear on the level of the geometry of the models. Let us come back to Fig. 4.79 and suppose, for example, that the selectivity r is everywhere maximal, viz. $r = 1$, even at the origin. Then, the lift of C_ρ is the curve Γ_ρ given by

$$\Gamma_\rho = \left(\frac{\rho}{2} \sin(\theta), \left(1 - \frac{1}{2}\rho \cos(\theta)\right) \cos(\theta), \left(1 - \frac{1}{2}\rho \cos(\theta)\right) \sin(\theta) \right),$$

and when $\rho \rightarrow 0$, the section Γ_ρ no longer tends to 0 but to the ‘vertical’ unit circle $\Gamma_0 = (0, \cos(\theta), \sin(\theta))$, and the projection π is not at all a local diffeomorphism at 0. It is such away from 0, even if it turns out to be a highly twisted diffeomorphism, but at 0 itself, it has an ‘exceptional fibre’ Γ_0 of dimension 1 (see Fig. 4.84).


It is thus well worth carrying out precise experiments on the structure of the orientation field in the neighbourhood of singularities. They are difficult to do but rich in results. By combining imaging methods with methods for intracellular recording of spikes triggered by the synaptic inputs, but also recording of membrane potentials of single neurons, Pedro Maldonado, Imke Gödecke, Charles Gray, and Tobias Bonhöffer were able to analyze the fine structure of the orientation maps at singularities, observing that [86, p. 969]:

Orientation columns contain sharply tuned neurons of different orientation preference lying in close proximity.

Put another way, column redundancy seems to disappear at singular points.

For further discussion of these results, the reader is referred to David McLaughlin et al. [88] and Michael Shelley et al. [89], who model the cortical processes generating orientation selectivity in the $4C\alpha$ layer of the macaque V1 area. They start from the microlevel with a network of 16000 integrate-and-fire neurons and reduce it at the meso level to what they refer to as ‘a spatially coarse-grained system for firing rates of

Fig. 4.84 When $\rho \rightarrow 0$, the section Γ_ρ no longer tends to 0 but to the ‘vertical’ unit circle Γ_0 (see text)



neuronal subpopulations' [89, p. 97]. They then introduce a plausible connectivity for the 'vertical' connections coming from the LGN and for the 'horizontal' intracortical inhibitory and excitatory connections (which we shall return to at length in Sect. 5.1 of Chap. 5). In the second volume, we shall discuss cortical models comprising a huge number of differential equations describing the membrane potentials of each neuron in the network, along with coarse-grained mean field methods which allow us to study them on the meso level. What is important here is the result from these authors, formulated here for an orientation hypercolumn, i.e. a pinwheel, including isotropic internal inhibitory and excitatory connections, but no anisotropic long-range lateral connections with other hypercolumns:

This analysis showed it is an interaction between the pinwheel structure of the preferred orientation mapping and the isotropic architecture that produces greater orientation selectivity near pinwheel centres [89, p. 121].

In other words, these models show that the pinwheels are 'true' singularities where all the orientations are present, *each with strong selectivity*.

Many experimental results confirm this idea. For example, using moving gratings as stimuli, James Schummers et al. have shown that:

Neurons near pinwheel centres have sub-threshold responses to all stimulus orientations but spike responses to only a narrow range of orientations. [90, p. 969]

They assess the selectivity of spike responses and sub-threshold responses as a function of position relative to the pinwheels. Far from the pinwheels:

[Cells] show a strong membrane depolarization response only for a limited range of stimulus orientation, and this selectivity is reflected in their spike responses. [90, p. 970]

At the centre of a pinwheel, on the other hand, only the spike response is selective and the membrane is highly depolarized for all orientations (see Fig. 4.85).

These results show that, at singular points, all orientations are indeed present, but using a novel solution with regard to the connectivity of the relevant neural microcircuits:

These examples indicate that both simple and complex cells located near pinwheel centres receive synaptic inputs over a broad range of stimulus orientations, although not all of these inputs are represented in the spike outputs. [90, p. 971]

Schummers [91] obtained further results with Jorge Mariño (on cats). The local neural circuits are very different at regular points and at singular points. If we look at the total conductance g , we observe that it is strongly peaked at a regular point and flatter at a singular point. It is the different ways in which it is distributed between inhibitory and excitatory conductances g_i/g_e and the interaction between the excitatory and inhibitory connections which show us how, 'despite the diversity of local environments', there can nevertheless be 'a sharp orientation tuning at all locations in the orientation map' [91, p. 195]. With highly localized injections of retrograde tracer (where 'localized' refers to a region $<100 \mu\text{m}$), the difference between excitatory and inhibitory can be clearly distinguished. Figure 4.86 shows

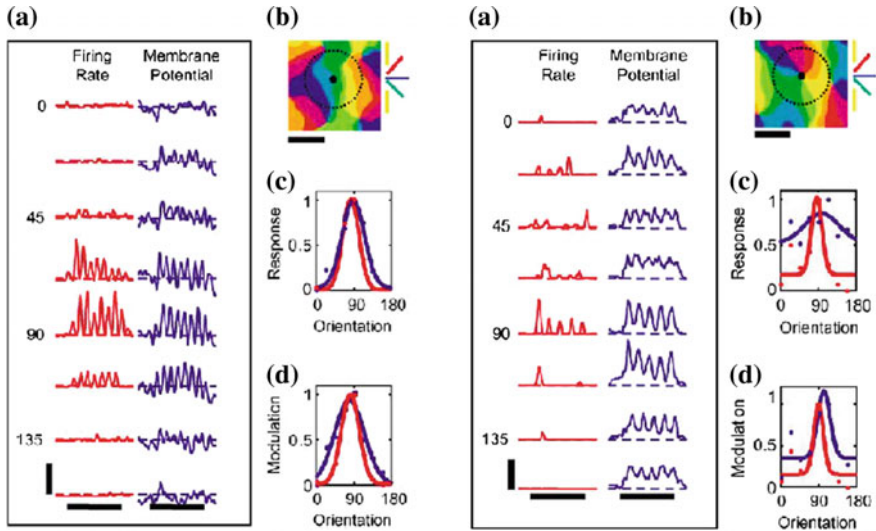


Fig. 4.85 Responses to stimuli in the form of moving periodic gratings with different orientations. *Left* Far from singularities. Scales 8 spikes/s, 10 mV, 2 s. *Right* At the centre of a pinwheel. Scales 3 spikes/s, 8 mV, 2 s. **a** Comparison between spike emissions (*red*) and the membrane potential (*blue*). **c** The two curves representing the average amplitude of the responses for different orientations. **d** Modulation of responses. From Schummers [90]

an injection site with the inhibitory and excitatory connections. We observe that inhibition is more localized and more isotropic, whereas excitation propagates over greater distances, but in a highly *anisotropic* way, clustering in regions of the same colour and hence of the *same orientation* as the injection site. We shall return to this crucial point, the key to the functional architecture of V1, in Sect. 5.1 of Chap. 5.

These measurements confirm that the difference between regular and singular points is not reflected in the level of the spikes (there is good orientation selectivity in both cases), but at the level of the sub-threshold signals. This is because:

[...] as a result of the appropriate inhibitory balance at orthogonal orientations, which keeps the membrane potential below threshold, it is not reflected in the spike responses.

4.7.2 The Problem of Resolution

As already mentioned, it is important to note that the question of *resolution* plays an important role in the definition of the pinwheels. In an optical device, there is always a resolution transforming points (Dirac measures) into Gaussians of a certain width σ . In the optical imaging of intrinsic signals *in vivo*, there are systematic statistical errors. As noted by Jonathan Polimeni et al., these are due to:




Fig. 4.86 The injection site is the *small central circle*. The *outer circle C* is at $250\ \mu\text{m}$ from the *intermediate circle* (see the $250\ \mu\text{m}$ scale bar in the *bottom left corner*). *White dots* represent *inhibitory connections* and *black dots* *excitatory connections*. We observe that inhibition is more localized, bounded by the circle *C*, and more isotropic, whereas excitation propagates well outside the circle *C*, but in a highly anisotropic way, clustering in regions of the same colour and hence of the same orientation as the injection site. From Mariño and Schummers [91]

[...] photon scatter and absorption in brain tissue combined with the blurring introduced by the optics of the imaging system. [87, p. 4158]

In this sense:

Optical recording, as it has been used to date, has insufficient spatial resolution to accurately locate pinwheel centers. [87, p. 4158]

Fortunately, as we shall soon see, there are now much higher resolution methods which provide a way around these difficulties.

However, it remains important to take into account the fact that the functional orientation maps are obtained by superposing single-orientation maps and are thus vector-valued images. They have a resolution corresponding to a width σ of about $250\ \mu\text{m}$. Figure 4.87, taken from [87], shows what happens to a pinwheel map when it is convolved with a Gaussian kernel of width $250\ \mu\text{m}$. Obviously, we observe that the map changes significantly. In other words, good *mesoscopic* models must always be *multiscale* models.




Fig. 4.87 Transformation of a pinwheel map by convolution with a Gaussian kernel of width $250\ \mu\text{m}$. *White pinwheels* are (+), *black pinwheels* are (−), and *continuous* and *dashed lines* are the zero-crossings of the map. The *scale bar* represents 1 mm. We note how much the map changes. **a** Initial map with 124 pinwheels. **b** Transformed map with only 93 pinwheels. From Polimeni et al. [87]. Copyright (2005) National Academy of Sciences, USA

4.7.3 Two-Photon Confocal Microscopy

The imaging techniques mentioned up to now are not accurate enough. However, recently, new techniques like *in vivo* imaging by *two-photon confocal microscopy* have given us functional maps with a resolution down to the level of single neurons. Kenichi Ohki and coworkers [92] have shown that, in cats, pinwheels defined on the mesolevel remain highly ordered on the microlevel. Consequently:

Pinwheel centers truly represent singularities in the cortical map.

The idea of the method is to inject calcium-sensitive indicators (Oregon Green BAPTA-1 acetoxymethyl ester) which label a few thousand neurons in regions of $300\text{--}600\ \mu\text{m}$. We simultaneously measure calcium signals triggered by visual stimuli in hundreds of neurons at different depths (from 130 to $290\ \mu\text{m}$, in steps of $20\ \mu\text{m}$), and we find pinwheels with the same orientation wheel (see Fig. 4.88):

This demonstrates the columnar structure of the orientation map at a very fine spatial scale.

Whence the problem of connectivity implementing the fine selection of orientations near singularities; several hypotheses have been put forward by the authors about the dendritic tree near the centre C (a few tens of μm) in an iso-orientation domain D (see Fig. 4.89):




Fig. 4.88 Two-photon optical imaging can simultaneously measure calcium signals triggered by visual stimuli in hundreds of neurons at different depths (from 130 to 290 μm in steps of 20 μm). We find analogous pinwheels with the same orientation wheel at these different depths. From Ohki et al. [92]

- (a) Unbalanced dendritic tree towards D .
- (b) Symmetric dendritic tree, but excitatory inputs unbalanced towards D .
- (c) Symmetric dendritic tree, symmetric excitatory inputs, but local and within D (good segregation in the vicinity of C).
- (d) Symmetric dendritic tree, symmetric excitatory inputs, integrated over a wide dendritic area.




Fig. 4.89 Four dendritic tree structures in the vicinity of a singularity, as proposed by Kenichi Ohki et al. [92] (see text)

4.8 Pinwheels and Blow-ups

The results of Schummers and Maldonado lead to two competing idealized models. In both cases, all the orientations are indeed present with good selectivity in the neighbourhood of singularities, but at the singularity itself (insofar as this has a precise meaning on the meso level), either the orientation selectivity is sub-threshold (Schummers), or there is a loss of column redundancy (Maldonado). Further experiments will be needed to elucidate this, but we may nevertheless make a suggestion for Maldonado's results. The theoretical problem we have to solve concerns the dimensional collapse of an ideally 3D structure onto 2D neural layers.

4.8.1 The Geometric Concept of Blow-up

So how can the concrete 2D structure of pinwheels fit with the abstract 3D structure of the fibre bundle? To understand this, we shall use the geometric operation known as ‘blowing up’ of singularities. The intuitive idea, which needs to be carefully tested, is that the pinwheels could be like local blow-ups of singular points, and the pinwheel lattice like a gluing together of these local blow-ups to give a discrete approximation of the projection $\pi : R \times \mathbb{P}^1 \rightarrow R$. In this picture, π would correspond to a limiting situation where in some sense all the points of R would blow up in parallel.

In algebraic geometry, the blow-up of a variety such as the plane $M = \mathbb{R}^2$, at a point, e.g. the origin $O = (0, 0)$, is defined in the following way. Let $a = (x, y) \neq (0, 0)$ be a point of \mathbb{R}^2 . We associate the direction p of Oa with this point and thus specify a map δ :

$$\begin{aligned} \delta : \mathbb{R}^2 - \{O\} &\longrightarrow \mathbb{P}^1 \\ a = (x, y) &\longmapsto \delta(a) = p = y/x \end{aligned}$$

The graph of δ is a helicoidal surface H in the 3D fibre bundle $V = \mathbb{R}^2 \times \mathbb{P}^1$, and the topological closure of H in V is a helicoid \overline{H} with $\pi^{-1}(O) = \Delta \cong \mathbb{P}^1$. This copy of \mathbb{P}^1 above O is an exceptional fibre which, for historical and technical reasons, is called the *exceptional divisor* of the blow-up. The restriction to \overline{H} of the projection $\pi : \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$ is an isomorphism of H on $\mathbb{R}^2 - \{O\}$. If d is the straight line generated by Oa in \mathbb{R}^2 , the closure of the inverse image $\pi^{-1}(d - \{O\})$ is made up of points $(\lambda a, \delta(a) = p)$ of $V = \mathbb{R}^2 \times \mathbb{P}^1$, i.e., by the straight line d' at height $\delta(a) = p = y/x$. When the straight line d rotates in the plane \mathbb{R}^2 , d' also rotates, but while translating through Δ , whence the helicoidal motion.

Since the inverse image of the point O under π is the projective line $\Delta = \mathbb{P}^1$, π is not at all an isomorphism at O , but a projection collapsing a 1D fibre to a zero-dimensional point. In this sense, the blow-up of the plane at a point generates a geometric structure which is somehow ‘intermediate’ between the 2D plane and the 3D fibre bundle. It is the fibre bundle V above O and the (twisted) plane \mathbb{R}^2 away from O . One might say that the blow-up $\pi : \overline{H} \rightarrow \mathbb{R}^2$ unfolds the orientation wheel centred on $O = (0, 0)$ in a third dimension (see Figs. 4.90 and 4.91).

This construction can be understood as an interpretation of *polar coordinates* in terms of a *field of directions*. Indeed, consider the fibre bundle $\pi_1 : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$. On the plane with one point removed $\mathbb{R}^2 - \{O\}$, the argument $\theta(a) \in [0, 2\pi]$ of a point a is well defined and we may thus consider the section Θ_1 of π_1 defined by $\Theta_1 : a \rightarrow \Theta_1(a) = (a, e^{i\theta(a)})$. The fibre bundle $\pi : \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$ is the quotient of the fibre bundle $\pi_1 : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$ obtained by identifying θ with $\theta + \pi$ (i.e., by identifying $e^{i\theta}$ with $-e^{i\theta}$) and Θ_1 lifts to π_1 the section of π defined by $\Theta : a \rightarrow \Theta(a) = (a, e^{i\theta(a)})$, where $\theta(a) \in [0, \pi]$ is now considered modulo π . $\Theta_1(a)$ is constant on the rays $\theta = \text{const.}$, and when it is lifted from $\mathbb{R}^2 \times \mathbb{P}^1$ to $\mathbb{R}^2 \times \mathbb{S}^1$, the surface H of $\mathbb{R}^2 \times \mathbb{P}^1$ becomes the image of Θ_1 .

Fig. 4.90 Blow-up of the plane at a point a . The directions at a are unfolded in a third dimension





Fig. 4.91 When the third dimension collapses, the blow-up becomes a pinwheel



The concept of blow-up was introduced under the name of ‘quadratic transformation’ at the end of the nineteenth century by specialists of projective algebraic geometry. It represents the simplest case of what are known as birational transformations, and it is fundamental for desingularizing singular curves. If a curve γ in \mathbb{R}^2 has a singular point at O where several branches intersect with different tangents, then by lifting γ to \bar{H} , we obtain a curve $\Gamma = \pi^{-1}(\gamma)$ with various branches at different heights, thereby eliminating the intersections.

We can now *localize* this algebraic model and even consider an *infinitesimal version*, where we restrict to points $a = (dx, dy)$ infinitely close to the blow-up point $O = (0, 0)$. To do this, we take what is known as the ‘germ’ of the structure in the neighbourhood of (Δ, O) . In the local model, which is no longer algebraic but differential, we have $p = dy/dx$ and the surface H is thus included in the kernel of the differential form $\omega = dy - p dx$ defined on $V = \mathbb{R}^2 \times \mathbb{P}^1$ (we shall return to this key point in Sect. 5.4.1 of Chap. 5). Conversely, the algebraic model can be considered as the ‘tangent’ structure to the local model, where infinitesimal segments are replaced by tangent vectors.

The blow-up model can be used to understand orientation singularities. We can neglect the redundancy of the cortical columns outside the singularities insofar as, at regular points, simple $V1$ neurons mainly detect the same pair (a, p) . However, we have seen that, according to Maldonado et al. [86], this is not the case at singular points, where all orientations are in fact present in the column. The fact that the orientation associated with a ray of a pinwheel selects this orientation at the centre of the pinwheel can then be expressed by saying that the operation yielding the topological closure \overline{H} of H is implemented neurally and that the singular point is thereby blown up throughout the thickness of the cortical layer.

When the pinwheel is modelled in this way, it corresponds to the section Ψ of the projection $\pi : (\mathbb{R}^2 - \{O\}) \times \mathbb{P}^1 \rightarrow (\mathbb{R}^2 - \{O\})$ given by $\Psi(a) = (a, e^{i[\alpha + \theta(a)/2]})$, where α is the orientation encoded by the ray $\theta(a) = 0$. We thus require a double rotation of θ through angle π (hence two round trips of \mathbb{P}^1) to get back to the same ray, and this is why, as explained above, two diametrically opposite rays at angles θ and $\theta + \pi$ correspond to orthogonal orientations $\alpha + \theta/2$ and $\alpha + \theta/2 + \pi/2$. This means that the pinwheels implement what is known in geometry as a *spin structure*. The section Ψ is not the section $\Theta : a \rightarrow \Theta(a) = (a, e^{i\theta(a)})$, $\theta(a) \in [0, \pi]$, of the bundle $\pi : (\mathbb{R}^2 - \{O\}) \times \mathbb{P}^1 \rightarrow \mathbb{R}^2 - \{O\}$, but rather a field of orientations of the same order, and this justifies modelling it by a blow-up.

4.8.2 Blow-ups and Lines of Dislocations

A link can be made between these blow-ups and the singularities of the phase fields discussed in Sect. 4.6. On the one hand, if we rotate the x -axis by θ_0 without changing the origin in \mathbb{P}^1 , the helicoid H is translated by θ_0 to the helicoid H_{θ_0} and the blow-up can be considered as the set of all the H_{θ_0} . On the other hand, we can consider the field of phases $Z = ae^{-ip} = \rho e^{i(\theta-p)}$ in $\mathbb{R}^2 \times \mathbb{P}^1$, where p varies in the fibre \mathbb{P}^1 . Surfaces of constant phase are then the helicoids H_{θ_0} and the fibre \mathbb{P}^1 above 0 is a line of dislocations.

4.8.3 From Blow-up to Fibre Bundle

To construct a fibre bundle model for a global pinwheel structure from the local blow-up model, one must blow up all the points of an (irregular) lattice L of the plane, the lattice of singular points, *in parallel*. This is not possible in the framework of algebraic geometry, because the local structures determine the global structures, and to iterate the algebraic model for the blow-up of a point, one would have to immerse the successive blow-ups in spaces of ever higher dimension. However, it can be done in the context of differential geometry by gluing together the local models of different points in the lattice. We thus obtain a model for the pinwheel structure (see Figs. 4.92 and 4.93).





Fig. 4.92 Simultaneous blow-up of a lattice of points

Fig. 4.93 When the third dimension collapses, a simultaneous blow-up of a lattice of points yields a lattice of pinwheels



In a field model like the one in Sect. 4.4.3, the field can be lifted from \mathbb{R}^2 to $V = \mathbb{R}^2 \times \mathbb{P}^1$ by blowing up the singularities c_i and lifting the field lines like curves: if a field line γ goes from a source c_1 with angle θ_1 to reach a sink c_2 with angle θ_2 , γ lifts to a curve Γ in V starting at height θ_1 in the fibre above c_1 and arriving at height θ_2 in the fibre above c_2 . Only the fibres above c_1 are involved, so everything happens as though the c_i had been blown up in parallel. The structures tangent to this multiple blow-up in the neighbourhoods of the fibres $\Delta_i \rightarrow c_i$ are all isomorphic to the local algebraic model.

We could therefore consider that, when the mesh of the lattice L tends to 0, the limit of this multiple blow-up gives back the bundle $\pi : \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$, gluing back the infinitesimal models at all the points of \mathbb{R}^2 . In this sense, the pinwheel structure

could effectively be interpreted as a discrete approximation of the bundle π , and conversely, π could be considered as the simultaneous blow-up of all the points in the plane. A good way could be to use a non-standard model²⁰ $(\mathbb{R}^*)^2$ of \mathbb{R}^2 where, around each standard point $a = (x, y)$, there would be a ‘halo’ or a ‘monad’ of infinitesimals $\mu(a) = \{(x + dx, y + dy)\}$. In the blow-up, the exceptional fibre Δ^* would then be $(\mathbb{P}^1)^*$ and dy/dx a non-standard real number $p^* = p + dp$, equivalent to $p \in \mathbb{P}^1$. The field lines would all lie within the monads, and so would be non-standard, and to first order, there would only remain the segments $a \rightarrow (a + da)$. The standard part of the structure at $\Delta^* \rightarrow a$ would thus be the tangent structure. In this way, we would obtain not only the bundle $\pi : \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$, but also the infinitesimal structure defined on it by the differential 1-form $\omega = dy - p dx$, i.e. the *contact structure*, to which we shall return.

The advantage of a non-standard model is that it provides an intuitive picture of the characteristic dimensional reduction of the pinwheel structure. We take the bundle $\pi : \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$, compactify the fibres²¹ to make them infinitesimal, then project them into the monads $\mu(a)$. I used to think that the idea of blow-up with the exceptional fibre made infinitesimal and projected back onto the base plane was original. However, reading the correspondence between Pierre Deligne, Bernard Malgrange, and Jean-Pierre Ramis on ‘irregular singularities’, I found a letter dated 7 January 1986 on the singularities of analytic functions and Gevrey classes, where Deligne [94] introduced the concept of ‘thick points’. The idea is to replace a point $a \in \mathbb{C}$, say $a = 0$, by a small disc D with boundary Δ , consider the space $\tilde{\mathbb{C}} = \mathbb{C}^* \cup D$, union of $\mathbb{C}^* = \mathbb{C} - \{0\}$ and D , and equip it with the topology of the blow-up of 0 in \mathbb{C} along $\mathbb{C}^* \cup \Delta$. Moreover, in his last paper on Gevrey classes (edited by Jean-Pierre Ramis [95]), Martinet [96] used this construction with the discs D that are infinitesimal in the sense of non-standard analysis.²²

We could thus say that, in the continuous limit, a pinwheel lattice model amounts to treating the points of the plane as infinitesimal ‘thick points’ in the sense of Deligne and Martinet, with the standard part of such a structure giving back the bundle $\pi : V = \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$.

4.8.4 Discrete Versus Continuous Models

We thus see that there are two complementary ways to model pinwheels and one which combines both:

1. We can go to the continuous limit and work in $\pi : V = \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$. This is what we shall do at great length in the second volume.

²⁰For a didactic introduction to non-standard analysis, see, for example, Petitot [93] and the references therein.

²¹Rather as in the Kaluza–Klein field theories of physics.

²²I thank Guy Wallet and Michel Berthier for this reference.

2. Alternatively, we do not go to the limit, but keep the finite mesh size of the lattice L and consider the orientation field $\Psi(a)$ of simple V1 neurons as a section of the bundle π defined on the open set $\mathbb{R}^2 - L$. This is the approach we shall pursue in the following sections of this chapter.
3. However, it is also possible to adopt a mixed approach and introduce continuous models that get discretized by breaking their symmetry groups. We shall also pursue this idea in the second volume.

4.9 Different Aspects of Pinwheels

To illustrate our models with experimental data, we shall discuss here some other aspects of the pinwheel maps.

4.9.1 Position–Orientation Independence and Local Triviality

The pinwheel structure is a good example of the way neurophysiology requires us to rethink the most basic geometric structures used to model it. In the above, we have assumed that the direct product structure $U \times \mathbb{P}^1$ ($U \subset R$) in the fibre bundle $\pi : U \times \mathbb{P}^1 \rightarrow U$ (local triviality) raises no particular problems. However, it assumes an independence between the position and orientation variables which is not at all obvious from the neurophysiological point of view (see, e.g. Das, Gilbert [97]) and must be carefully checked. This is what has been done by Bosking et al. [42] by analyzing the pattern of neural activity elicited by a long line crossing the visual field.

Figure 4.94, part of which appears as Fig. 4.31 in Sect. 4.4.1, shows the following:

- (a) The band of neurons in V1 which are activated by a long line located at a precise (vertical) position x (scale 1 mm).
- (b) The way this band is situated within the population of V1 neurons responding to the same vertical orientation but at different positions.


By carefully analyzing the way these bands and the response peaks change as the stimulus moves (10 positions x at intervals of 1°),²³ the authors have shown that the maps of the positions of the stimulus and the orientations (pinwheels) are essentially independent. Figure 4.95 shows the relationship between a $1^\circ \times 1^\circ$ lattice of vertical and horizontal positions of the bar and the pinwheel structure. From a scale of $4^\circ \times 4^\circ$ the coverage becomes uniform.

In short, Bosking shows that:

The map of visual space in V1 is orderly at a fine scale and has uniform coverage of position and orientation without local relationships in the mapping of these features.

²³Lengths are measured in degrees of the visual field.

Fig. 4.94 Neurons in V1 activated by a long line located in a precise vertical position. From Bosking et al. [42]



In geometrical terms, this means that the Cartesian product structure (the triviality) of the bundle $\pi : R \times \mathbb{P}^1 \rightarrow R$ is indeed neurally implemented.

4.9.2 Other Engrafted Variables

Other variables are engrafted in Hubel's sense in the pinwheel structure, for example the direction of motion, phase, spatial frequency, and ocular dominance.

4.9.2.1 Direction of Motion

Simple V1 neurons are orientation selective and detect edges. However, they are often also sensitive to the direction of motion of the edges, and in general, this is optimal in the orthogonal direction. This is the case, for example, in the cat, as shown in Fig. 4.18 taken from Van Hooser [26]. In [98], Ohki et al. confirmed this on the level of individual cells in area 18 of the cat (see Fig. 4.96).

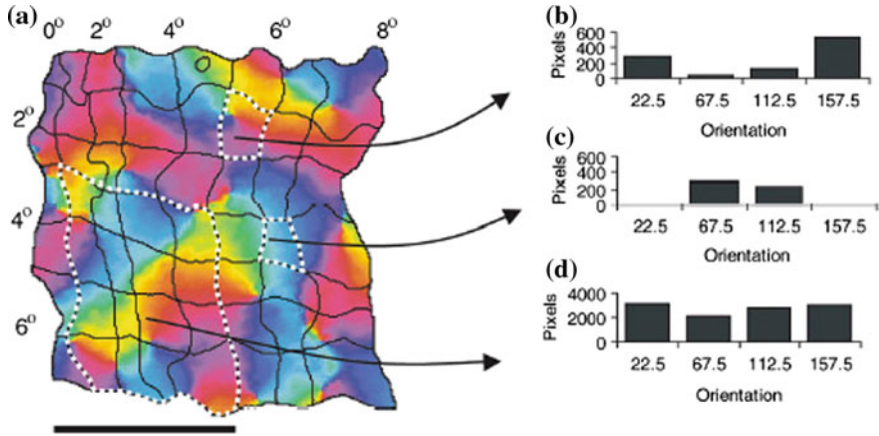


Fig. 4.95 Maps of the positions of the stimulus and the pinwheels. The coordinate system (x, y) reconstructed from the iso-azimuth contours of a $1^\circ \times 1^\circ$ lattice of vertical and horizontal positions of the bar is independent of the pinwheel structure. The tiny $1^\circ \times 1^\circ$ domains clearly have dominant orientations [histograms (b) and (c)]. But from a scale of $4^\circ \times 4^\circ$, the coverage is uniform [histogram (d)]. From Bosking et al. [42]




Fig. 4.96 Direction selectivity in cats. The colours *red* and *green* encode the two directions orthogonal to a given orientation. From Ohki et al. [98]

In Sect. 5.11 of Chap. 5, we shall return to the problem of direction in the context of other visual areas. We shall also discuss this a little more in Sect. 4.10.1.4.

4.9.2.2 Phase

Concerning the variation of the *phase* in a given column, DeAngelis et al. [19] compared the *spatiotemporal* RF ($X, Y = \text{space}$ and $T = \text{time} = \text{correlation delay}$) of two neighbouring cells in the same column. The stimuli were small randomly flashed bars (length 1.5° , 40 ms flashes) with the preferred orientations of the two cells, these being recorded simultaneously. The authors measured the cross-correlation between the sequence of stimuli and the response (spike trains) with different correlation delays. They observed that the visuotopy, orientations, and spatial frequencies are the same, but not the phases.

4.9.3 Spatial Frequency

Here, we attribute a little more importance to another engrafted variable, viz. the spatial frequency, where interesting new models have recently been developed.

4.9.3.1 Some Preparatory Notes

Note that pinwheel rays leave room for a further parameter. Recent work has conjectured that this could be the *spatial frequency* (SF, see, e.g., DeAngelis et al. [19] and Bressloff-Cowan [99]). The experiments are not easy to carry out. Care must be taken to use stimuli that only select a single SF, and the results are still controversial.

Figure 4.97 due to Hübener et al. [100] shows the edges of domains of low spatial frequency. Statistically, the pinwheels tend to be located towards the centre of the frequency domains and the iso-orientation lines are rather strongly *transverse*, sometimes almost orthogonal, to the edges. This is a strong transversality condition.

In their work [101], Naoum Issa, Christopher Trepel, and Michael Stryker made a detailed analysis of this spatial frequency (SF) distribution in cats (see Fig. 4.98). Their results are slightly different, which shows that this is only the beginning of such investigations. They have suggested that, for this parameter, there may also be a column structure comprising columns of about 0.7 mm with ends close to the pinwheels. As they note:

The organization of cortical maps permits nearly all combinations of orientation and SF preference to be represented in V1. [101, p. 8504]

Put another way, it would seem that there is a neural implementation of the bundle $\pi : \mathbb{R}^2 \times P \times F \rightarrow \mathbb{R}^2$, where F is the interval of the SFs observed (roughly from




Fig. 4.97 Edges of domains of low spatial frequency (*gray*). From Hübener et al. [100]

0.2 to 1.8 cycle/degree). In fact, the binarization of the SFs into high and low frequencies may be correlated on the one hand with the segregation of the parvocellular (X) and magnocellular (Y) pathways and on the other with the chirality of the pinwheels. It often happens that two adjacent pinwheels (hence of opposite chirality) are in frequency columns of opposite colours (hence high and low frequencies).

We see immediately the kind of problems that such results raise for modelling. We may for example pair together the orientation θ with the SF f in a fibre space H and thus consider a field with values in H defined on the cortical surface, i.e. a section of the bundle $\mathbb{R}^2 \times H$. But how should H be defined? If θ and f are treated as independent variables, then $H = P \times F$. But we may also pair θ and f . This is what is done by Paul Bressloff and Jack Cowan in [99] by introducing a local ‘spherical’ model in which they glue together a disc (pinwheel) of maximal f with a disc (pinwheel) of minimal f along their boundaries. The common boundary is the equator, f_{\min} and f_{\max} are the poles, and a point of the sphere is thus a pair (θ, f) located at a distance from the singular point (the centre of the sphere) defined by f .

However, these results from Issa et al. [101] were taken up by other specialists who came to different conclusions. For example, Sirovich and Uglyesich [102] detected a bias due to vascular artefacts in the hemodynamic measurements and, when they eliminated it, they concluded that there was no column organization in the SFs; for them, the SF was more like a parameter *distorting* the orientation map. Figure 4.99 shows, for two pinwheels, the difference between the maps obtained with high SFs (left) and low SFs (right). See also Born and Tootell [103].





Fig. 4.98 **a** Orientation columns. **b** Spatial frequency columns. High frequencies are coded in *blue* and low frequencies in *red*. Note that they correspond well to the two chiralities of the pinwheels, i.e. two adjacent pinwheels (hence of opposite chirality) are in frequency columns of opposite colours. The *scale bar* is 1 mm. From Issa et al. [101]

Fig. 4.99 Spatial frequency (SF) dependence of orientation maps. For two pinwheels (*top and bottom*), we observe the difference between the map obtained with high SFs (*left*) and low SFs (*right*). From Sirovich and Uglesich [102]. Copyright (2004) National Academy of Sciences, USA



These critical results were themselves repeated with autofluorescence imaging techniques using oxidation and reduction of flavoproteins in the metabolism of mitochondria. These techniques can correct for the vascular artefact pointed out by Sirovich, and it would seem that there is in fact no column organization (see Issa et al. [104]).

There are many other studies of this issue. We may cite for example the results of Zhu et al. [105] on the correlated variation of the OR and the SF.

4.9.3.2 First Dipole Model

If we accept this idea, the basic module is no longer the isolated pinwheel as a hypercolumn, but a *pair* of adjacent pinwheels with opposite chiralities. Near the singularities, i.e. the pinwheel centres (PCs), only a small part of the SFs (low or high) would be represented. A natural model would then be a *dipole–dipinwheel model*.

In Sect. 4.4.3, we discussed the analogy between the functional pinwheel maps and the topological charge fields encountered in physics. A pinwheel is locally like the field produced by a $+$ or $-$ charge. A pair of pinwheels with opposite chiralities thus corresponds to the field produced by a pair of opposite charges ($+$, $-$). Such a field source is known in physics as a dipole. It will be useful to dwell on this for a moment.

In electrostatics, a dipole is made by placing opposite charges $-q$ and $+q$, e.g. -1 and $+1$, at two points A and B . Let d be the distance between A and B and place A at $(-d/2, 0)$ and B at $(d/2, 0)$, as in Fig. 4.100. Let (r, θ) be the polar coordinates of M . Up to a multiplicative constant which we ignore, we have the potential $V = 1/BM - 1/AM$. Since

Fig. 4.100 Dipole structure. Opposite charges $-q$ and $+q$ are placed at two points A and B and we calculate the potential V at different points M . The figure shows several equipotentials. The values of V range from $-\infty$ (black) at A to $+\infty$ (white) at B . The field is $-\text{grad}(V)$, and its field lines are orthogonal to the equipotentials





Fig. 4.101 Dipole field with equipotentials in *light blue* and field lines in *orange*



$$\begin{cases} BM^2 = r^2 + \frac{d^2}{4} - dr \cos(\theta) , \\ AM^2 = r^2 + \frac{d^2}{4} + dr \cos(\theta) , \end{cases}$$

we obtain the explicit formula

$$V(M) = \frac{1}{r} \left[\frac{1}{\sqrt{1 + \frac{d^2}{4r^2} - \frac{d}{r} \cos(\theta)}} - \frac{1}{\sqrt{1 + \frac{d^2}{4r^2} + \frac{d}{r} \cos(\theta)}} \right] .$$


In polar coordinates, the field $E = -\text{grad}(V)$ has components

$$E_r = -\frac{\partial V}{\partial r} , \quad E_\theta = -\frac{1}{r} \frac{\partial V}{\partial \theta} .$$

The + charge is attractive and the – charge repulsive. This is illustrated in Fig. 4.101, which is a special case of the field configurations of pinwheel lattices considered in Sect. 4.4.3.

In this dipole–dipinwheel model of an orientation–spatial frequency module comprising two adjacent pinwheels of opposite chirality, the field lines are iso-orientation curves going from low SFs centred on one of the pinwheels to high SFs centred on the other pinwheel, while the equipotentials are level lines of the SF. These two families of curves are orthogonal, i.e. maximally transverse, thus satisfying the Hübener hypothesis. Locally, in the neighbourhood of each of the PCs, we have a classical

Fig. 4.102 Equipotentials and field lines of an intrinsically 2D dipole. If we compare with the field of Fig. 4.101 with potential going as $1/r$, we see that the geometries are qualitatively similar, but that the directions of the field lines are reversed, so $+$ charges are repulsive and $-$ charges are attractive



orthogonal model made up of rays for the orientations and nested circles for the SFs, the latter varying smoothly near the PC.

Even though what interests us in these physical analogies is the geometry of the fields, not the physics they actually represent, we should nevertheless note that potentials of the form $1/r$ are related to forces of the form $1/r^2$ which depend intrinsically on the 3D nature of the space. Since we are working on phenomena that are intrinsically 2D, we should use 2D physical analogies where the forces go as $1/r$ and the potentials as $\text{Log}(r)$. Such situations have been investigated, but they are rather strange, because positive charges become repulsive: $1/r$ decreases when r increases, whereas $\text{Log}(r)$ increases. However, this does not change much in our case since we consider pairs of opposite charges, and as shown in Fig. 4.102, the geometry of the field barely changes qualitatively.

In fact, we may even consider forces of the form $1/r^\sigma$, where the exponent σ is a parameter to be determined experimentally.

4.9.3.3 Second Dipole Model

A certain number of more refined experiments using functional imaging techniques with better resolution have shown that the standard dipole–dipinwheel orthogonal model is not corroborated experimentally, in ways that go beyond the criticisms made by Issa [101]:

1. There are *both maxima and minima* of the SF in the vicinity of many PCs.
2. The strong transversality condition of Hübener is not satisfied at all.

In a paper in 2016, Ribot et al. [106] (see also [107]) analyzed in great detail the fine structure of the functional maps of the SFs in area 17 of the cat. They used sophisticated high-resolution imaging methods, essential from the methodological point of view if we are to test the possibility of large variations in the SF in the vicinity of a PC. They discovered that the SF generally has *two extrema*, a minimum and a maximum, in the vicinity of a given PC. This is a functional organization because it means that all the ORs and all the FSs can be represented at the same time within *the same* hypercolumn. However, it must then have a completely different topology to that of the standard orthogonal circular model.

Figure 4.103 shows the results of the data analysis carried out by Jérôme Ribot. The orders of magnitude are: (i) 1 mm for the size of a hypercolumn, (ii) 300 μm for the extent of the neighbourhood of the PC analyzed (so this remains well within the hypercolumn), (iii) 150 μm for the extent of the little neighbourhood of the PC where measurement errors remain too great. Three examples are shown. The two extrema are perfectly clear.

Hence the idea of keeping the pinwheel–dipole model, but *dissociating* the size d of the dipole from the distance between two adjacent pinwheels. We thus assume that the dipole size d is *small* compared with the size of the pinwheel as hypercolumn module. In the limit, we would then have an *infinitesimal* dipole, a situation that is easy to model.

Indeed, if the dimension d of the dipole is small enough relative to the distances r where the field is measured, then by expanding up to a few terms in d/r , we may use the approximations

$$\frac{1}{BM} = \frac{1}{r} \left[1 + \frac{d}{2r} \cos(\theta) \right], \quad \frac{1}{AM} = \frac{1}{r} \left[1 - \frac{d}{2r} \cos(\theta) \right], \quad V(M) = \frac{d}{r^2} \cos(\theta).$$

The equipotentials are then given by the equation $r^2 = k \cos(\theta)$. The components of the field in polar coordinates are




Fig. 4.103 Level lines of the SFs and their extrema in the neighbourhood of three PCs. They lie within a single hypercolumn. *Black pixels* are those where measurements are not accurate enough. From Ribot et al. [106, Fig. 5]

$$E_r = -\frac{\partial V}{\partial r} = \frac{2d}{r^3} \cos(\theta) , \quad E_\theta = -\frac{1}{r} \frac{\partial V}{\partial \theta} = \frac{d}{r^3} \sin(\theta) .$$

The equations of the field lines are therefore

$$\frac{dr}{r} = \frac{E_r}{E_\theta} d\theta = 2 \frac{\cos(\theta) d\theta}{\sin(\theta)} ,$$

that is $\text{Log}(r) = 2\text{Log}(\sin(\theta)) + c$, whence $r = k \sin^2(\theta)$.

For the intrinsically 2D dipole, the formulas change. In the limit, we obtain an approximation of the potential $V = -(d/r) \cos(\theta)$, equipotentials with equation $r = 2R \cos(\theta)$, and field lines with equation $r = 2R \sin(\theta)$. The first constitute a family of circles of radius R centred on the x -axis and with tangent the y -axis at 0, since

$$x = r \cos(\theta) , \quad y = r \sin(\theta) , \quad r^2 = x^2 + y^2 , \quad (x - R)^2 + y^2 = R^2 ,$$

whence $r^2 = 2Rx = 2Rr \cos(\theta)$, i.e. $r = 2R \cos(\theta)$. The second constitute the orthogonal family of circles centred on the y -axis and with tangent the x -axis at 0 (see Fig. 4.104). Figure 4.105 shows the potential $V(M)$ in the neighbourhood of 0.




Fig. 4.104 Equipotentials of an intrinsically 2D infinitesimal dipole. This is a family of circles centred on the x -axis and with tangent the y -axis at 0




Fig. 4.105 Potential $V(M)$ of the infinitesimal dipole in Fig. 4.104. We see how V diverges to $+\infty$ at 0^+ and $-\infty$ at 0^-

4.9.3.4 Topological Universality of Dipoles

In their elegant paper [58], already cited in Sect. 4.5, Alberto Romagnoni, Jérôme Ribot, Daniel Bennequin, and Jonathan Touboul showed that just as the topological structure of the pinwheel is universal for the ‘exhaustive coding’ and the ‘parsimonious coding’ of periodic quantities, so the topological structure of the dipole is universal for the ‘exhaustive coding’ and the ‘parsimonious coding’ of non-periodic quantities.

As for the pinwheels, the idea is to consider the way the cells parametrized by $z = (x, y)$ in a small disc $\Omega \subset \mathbb{C}$ (with boundary $\Gamma = \partial\Omega$) of the visual plane (identified with the plane $\mathbb{C} = \mathbb{R}^2$) can encode spatial frequencies $\nu \in \mathbb{F}$ in a *local* and, ideally, *scale invariant* way, where \mathbb{F} is the allowed spatial frequency domain. As we are only interested in topological structure, we can take $\mathbb{F} = \mathbb{R}$, introducing saturation if necessary. Such a coding can thus be described by a map $g : \Omega \rightarrow \mathbb{R}$ which may have singularities, i.e. points where it is not defined because it is multivalued or divergent.

We use the standard topological properties of \mathbb{R} :

- \mathbb{R} is *not compact*, i.e. its points escape to infinity.
- In contrast to \mathbb{S}^1 , if we remove a point or a finite interval from \mathbb{R} , it disconnects into two connected components.
- *Intermediate value theorem* (which fails disastrously in \mathbb{S}^1): if $h : I \rightarrow [h_{\min}, h_{\max}]$ is continuous for an interval I and if $\nu \in]h_{\min}, h_{\max}[$ is an intermediate value, then there exists $x \in I$ with image $h(x) = \nu$.

As for the pinwheels, the authors define a rather general class \mathcal{G} of sufficiently smooth maps g satisfying natural constraints. They assume to begin with that g is continuous away from its singularities and ‘exhaustive’, i.e. it codes all the SFs. In other words, it is *surjective*. This implies that g must have at least one singularity. Indeed, if g




Fig. 4.106 A connected level line $C = C_{v_0} = g^{-1}(v_0)$ of a surjective g cannot be entirely contained within the boundary Γ of Ω unless $g \geq v_0$ or $g \leq v_0$ everywhere in Ω . But then g is not surjective. So C must enter Ω . As usual, (\gtrsim) means $z \gtrsim v_0$

had no singularities, it would be a continuous map on a compact set Ω and its image $g(\Omega)$ would therefore be compact, hence contained in a closed *bounded* interval which could not then be equal to \mathbb{R} . Since we work locally and topologically, we shall assume that g has just one singularity and that it is at 0. We write $\Omega^* = \Omega - \{0\}$.

The authors then introduce three assumptions to define \mathcal{G} :

- H1 Smoothness condition: the map $g : \Omega^* \rightarrow \mathbb{R}$ is smooth (infinitely differentiable) and its level lines $C_v = g^{-1}(v)$ are smooth curves (as we saw in Sect. 4.5, this is a generic property [59]).
- H2 Exhaustivity condition: since the model must have scale-invariant properties, g must be exhaustive (surjective) in *any* neighbourhood of 0.
- H3 Minimal complexity or parsimony condition: topological redundancy must be minimal on *all* scales. There are arbitrarily small discs Ω' centred on 0 where the topological redundancy is minimal.

Concerning assumption H3, the authors use a lemma according to which *there is no* continuous surjective $g : \Omega^* \rightarrow \mathbb{R}$ with topological redundancy 1. Indeed, let $v_0 \in g(\Gamma)$ be a value of g taken on the boundary Γ of Ω and consider the level line $C = C_{v_0} = g^{-1}(v_0)$. To begin with, C cannot be completely included in the boundary Γ (see Fig. 4.106). Indeed, let $z_1 \neq z_2 \in \mathring{\Omega}^*$ be two points in the interior $\mathring{\Omega}^*$ of Ω^* and assume that $g(z_1) < v_0$ and $g(z_2) > v_0$. Let $\gamma : z_1 \rightarrow z_2$ be a continuous path joining z_1 and z_2 in $\mathring{\Omega}^*$. Then, γ is an embedding of an interval I in Ω^* and $h = g \circ \gamma : I \rightarrow \mathbb{R}$ is a continuous map, so according to the intermediate value theorem, there is a point z of γ where $g(z) = v_0$ and hence $z \in C \cap \mathring{\Omega}^*$. But this is impossible since this intersection is empty by the assumption about C . Hence, on $\mathring{\Omega}^*$, either $g \leq v_0$ everywhere or $g \geq v_0$ everywhere and g is not surjective, contradicting H1.

Thus, C enters the interior $\mathring{\Omega}^*$ of Ω^* . If we assume that the topological redundancy is 1, then C is *connected*. If C stops in $\mathring{\Omega}^*$, then we can find a topological disc $\Omega' \subset \mathring{\Omega}$ around 0 (possibly highly twisted) which is decomposed into two parts by C , i.e.,





Fig. 4.107 If the connected level line $C = C_{v_0} = g^{-1}(v_0)$ of g starts from the boundary Γ of Ω and enters Ω without reaching 0 , then C decomposes a neighbourhood Ω' of 0 into two subdomains and g is not surjective on the new $\Omega = \Omega'$, because $g \leq v_M$ everywhere on Ω^* . As usual, (\geq) means $g \geq v_0$

Fig. 4.108 If the connected level line $C = C_{v_0} = g^{-1}(v_0)$ starting from the boundary Γ of Ω reaches the singular point 0 , then either $g \geq v_0$ or $g \leq v_0$ everywhere on Ω^* and g is not surjective



C joins two points on the boundary Γ' . Replacing Ω by Ω' , we may thus assume that C decomposes Ω into two closed domains Ω_1 and Ω_2 with a common boundary C (see Fig. 4.107). Assume that $0 \notin \Omega_1$, whence $0 \in \Omega_2 - C$. By the intermediate value theorem once again, we have $g > v_0$ everywhere on $\Omega_1 - C$ and $g < v_0$ everywhere on $\Omega_2^* - C$, or the opposite. Assume the first case. Since Ω_1 is compact and g is defined on the whole of Ω_1 , g reaches a maximum $v_M \geq v_0$ on Ω_1 and hence $g \leq v_M$ everywhere on Ω^* . Consequently, g cannot be surjective.

So if C is connected, C must reach the singular point 0 . But then $\Omega^* - C$ is *connected*, and g is continuous on $\Omega^* - C$ and does not take the value v_0 . As its image $g(\Omega^* - C)$ is connected, being the continuous image of a connected set, then from the topological properties of \mathbb{R} , either $g \geq v_0$, or $g \leq v_0$ everywhere on Ω^* , so g cannot be surjective (see Fig. 4.108).

In short, level curves C_v with points on the boundary Γ necessarily have at least *two* connected components and the topological redundancy of g is ≥ 2 . However, $C_v \subset \overset{\circ}{\Omega}^*$ lying completely in the interior of Ω^* can be connected.

Clearly, the standard model of the 2D infinitesimal dipole $g_0 : \Delta^* \rightarrow \mathbb{R}$ defined on the unit disc $\Delta^* = \Delta - \{0\}$ with the origin removed,

$$z = \rho e^{i\theta} \mapsto v = \frac{\cos(\theta)}{r},$$

belongs to \mathcal{G} :

- (i) g_0 is well defined away from 0, 0 being a singularity where it is completely undefined, since θ is not defined and $1/r$ diverges.
- (ii) Away from 0, g_0 is smooth (infinitely differentiable).
- (iii) The level sets $C_v = g_0^{-1}(v)$, where g_0 is constant, are smooth curves.
- (iv) g_0 is surjective on any neighbourhood of 0 (and even on any small circle surrounding 0).
- (v) The topological redundancy of g_0 is 2, and it is therefore minimal.

Note that (see Fig. 4.104), insofar as the level lines constitute a bundle of circles (with the origin 0 removed), small enough circles of radius $\leq 1/2$ lie within Δ and are connected level lines, whereas circles of radius $> 1/2$ stop at Γ and are level lines with two connected components. The authors refer to the configuration of internal circles as a ‘bouquet’ of circles. They prove the following theorem:

Theorem *Topological universality of dipoles. The elements $g : \Omega^* \rightarrow \mathbb{R}$ of \mathcal{G} , i.e., satisfying hypotheses H1, H2, and H3 on all scales, have the dipole topology.*

As for the theorem on the topological universality of pinwheels, the ideas used in the proof are purely topological. We can give a heuristic insight (incomplete and not rigorous). Topologically, the connected components C of the level lines C_v of $g \in \mathcal{G}$ can be of 4 types: (i') a closed loop not going through 0 and not going around zero, (i'') a closed loop not going through 0 but going around 0, (ii) a path from Γ to Γ , (iii) a path connecting Γ to 0, and (iv) a loop going through 0 (see Fig. 4.109).

As g is continuous on the boundary Γ and the boundary is compact, $g(\Gamma)$ is a compact subset K of \mathbb{R} and hence bounded. As g has its values in K for paths of




Fig. 4.109 The four types of connected components of level lines for a $g \in \mathcal{G}$. *i* Loops in the interior of Ω and not going through 0 but either encircling it or not. *ii* Paths connecting two points on the boundary Γ . *iii* Paths connecting a point on the boundary Γ with the singularity 0. *iv* Loops in the interior of Ω and passing through 0

types (ii) and (iii) g , there must necessarily be level lines of types (i) or (iv). Let C be of type (i'), i.e. $C \subset \mathring{\Omega}^*$, with 0 outside C . Let Λ be the closed topological disc with boundary C and $\Omega^{*'} a small disc centred on 0 and contained in $\Omega^* - \Lambda$. From the scale invariance hypothesis, we can localize in $\Omega^{*'}$ and Λ only adds redundancy. We can thus eliminate the case (i'), and also case (ii) for the same reason (we can remove the connected component of $\Omega - C$ which does not contain 0).$

There remain the cases (i''), (iii), and (iv). Let C be of type (i'') and $\Lambda \ni 0$ the topological disc with C as boundary; since C is a level curve, according to the smoothness assumptions, no level line in the interior of Λ^* can meet C . There can thus only be C' of types (i'') and (iv) in Λ^* . However, there cannot only be C' of type (i'') because they would be 'concentric', i.e. nested, and localizing even more closely around 0, g would no longer be surjective, thus violating H2. So there is at least one C' of type (iv) in the interior of Λ^* .

Let C be a level curve C_ν of type (iv) on which $g = \nu$. Then, C is the boundary of a disc Λ and $0 \in C = \partial\Lambda$. Since 0 is a singularity of g , the latter can diverge as we approach 0 in Λ^* , but it can diverge only to either $+\infty$ or $-\infty$ and not both. There will thus be curves C^+ and C^- bounding regions Λ^+ and Λ^- . Consider the case where g diverges to $+\infty$ and let $\mathring{\Lambda}_m^+$ be the interior of the Λ^+ of the largest component of type (iv) diverging to $+\infty$. We obtain a lobe of concentric regions Λ^+ on which g tends to ν when the Λ^+ expands to Λ_m^+ and tends to $+\infty$ when Λ^+ contracts to 0. If ν_m^+ is the value of g on the boundary C_m^+ of Λ_m^+ , g takes the values $[\nu_m^+, +\infty]$ on Λ_m^+ .

The complement $\Omega - \mathring{\Lambda}_m^+$ is topologically a closed ring on which g must be able to diverge to $-\infty$ (because this divergence is not realized in $\mathring{\Lambda}_m^+$ by construction). This requires curves C of type (iv) in the interior of $\Omega - \mathring{\Lambda}_m^+$ and bounding regions Λ^- . We thus have, as for the standard dipole (see Fig. 4.104), a lobe of concentric regions Λ^+ up to Λ_m^+ and a lobe of concentric regions Λ^- expanding out to Λ_m^- . If ν_m^- is the value of g on the boundary C_m^- of Λ_m^- , then g takes the values $[-\infty, \nu_m^-]$ on Λ_m^- .

From the minimal complexity hypothesis, i.e. that the topological redundancy is equal to 2, there cannot be more than two lobes of this kind. Indeed, consider a level line of type (iv). If we have restricted it to a small enough disc Ω' around 0, we obtain two connected components (see Fig. 4.110i). The values of g on two lobes cannot be the same because, for small enough Ω' , that would give four connected components, thereby violating hypothesis H3 (see Fig. 4.110ii). Suppose then that there are three lobes. Then, g will take different values on each of them. Suppose that $\nu_1 < \nu_3 < \nu_2$ with $g(a) = \nu_1$ for a point a of the lobe L_1 , $g(b) = \nu_2$ for a point b of the lobe L_2 , and $g(c) = \nu_3$ for a point c of the lobe L_3 . Let $\gamma : a \rightarrow b$ be a path passing outside the lobe L_3 (see Fig. 4.110iii). According to the intermediate value theorem, there is a point d on γ , where g takes the value ν_3 . But then there is a connected component of C_{ν_3} outside L_3 and reaching 0. Since the part of C_{ν_3} inside L_3 itself has two connected components if Ω is small enough, C_{ν_3} has a topological redundancy ≥ 3 , thereby violating the hypothesis H3.





Fig. 4.110 Diagrams illustrating the proof of the topological universality of the dipole

Fig. 4.111 Illustration of the last part of the proof of the topological universality of the dipole



We thus end up with a ‘bouquet’ comprising a lobe Λ_m^+ of concentric regions Λ^+ on which g takes the values $[\nu_m^+, +\infty]$ and a lobe Λ_m^- of concentric regions Λ^- on which g takes the values $[-\infty, \nu_m^-]$. On the respective boundaries C^+ and C^- of Λ_m^+ and Λ_m^- , g takes the values ν_m^+ and ν_m^- , respectively. It remains to consider the values in the interval (ν_m^-, ν_m^+) . The complement of the union $\Lambda_m^+ \cup \Lambda_m^-$ of the two lobes in Ω comprises two disjoint open sets U_1 and U_2 . Let γ_1 be a path from C^+ to C^- in U_1 and γ_2 a path from C^+ to C^- in U_2 . Using the intermediate value theorem once again, if $\nu \in (\nu_m^-, \nu_m^+)$, g takes the value ν on both γ_1 and γ_2 . The level line C_ν then comprises two connected components of type (iii) joining Γ to 0 (see Fig. 4.111).

4.9.3.5 Pinwheel–Dipole Model

Daniel Bennequin and coworkers thus arrive at a model of the orientation–spatial frequency module of the form (f, g) with $f(z) = \theta/2$ and $g(z) = \cos(\theta)/r^\alpha$, with

$z = re^{i\theta}$ in the unit disc with the centre removed Δ^* and α an exponent used to fit the model to the data. Here, 0 is a singularity of both f and g :

[This model] is the unique scale-invariant topology achieving minimal representation of a pair (θ, ν) for the circular variable θ and the variable ν taking values in an open interval. [58, p. 9]

In this model, the orthogonality of the level lines of f and g is no longer satisfied at all. Indeed, let $z_0 = r_0 e^{i\theta_0}$ be a point of Δ^* . The level line F_0 of f corresponding to $f(z_0) = \theta_0/2 = \text{const.}$ is the ray through z_0 . The level line G_0 of g corresponding to $g(z_0) = \cos(\theta_0)/r_0^\alpha = \text{const.}$ has equation $\cos(\theta)/r^\alpha = \cos(\theta_0)/r_0^\alpha = C$, i.e. $\cos(\theta) - Cr^\alpha = 0$. Considering the symmetries of the configuration, we can restrict to $\theta \in [0, \pi/2]$.

Suppose to begin with that $\alpha = 1$. Then, G_0 has equation $r = 2R \cos(\theta)$, with $2R = r_0/\cos(\theta_0)$ so the tangent at z satisfies $dr + 2R \sin(\theta) d\theta = 0$. Using

$$x = r \cos(\theta) = 2R \cos^2(\theta) \quad , \quad y = r \sin(\theta) = 2R \cos(\theta) \sin(\theta) \quad ,$$

we obtain

$$\begin{cases} \frac{dx}{d\theta} = -4R \cos(\theta) \sin(\theta) = -2R \sin(2\theta) \quad , \\ \frac{dy}{d\theta} = 2R [\cos^2(\theta) - \sin^2(\theta)] = 2R \cos(2\theta) \quad , \\ \frac{dy}{dx} = -\frac{1}{\tan(2\theta)} = \tan\left(\frac{\pi}{2} + 2\theta\right) \quad . \end{cases}$$

Hence, the tangents to F_0 and G_0 at z_0 make angles θ_0 and $\varphi_0 = 2\theta_0 + \pi/2$, so the angle between them is $\psi_0 = \varphi_0 - \theta_0 = \theta_0 + \pi/2$. In particular, they are orthogonal for $\theta_0 = 0$, i.e. at the point $(2R, 0)$, and become tangents when θ_0 tends to $\pi/2$. The normalized distribution $P(\psi)$ of the angles ψ is thus uniform and equal to $P(\psi) = 2/\pi$, so that $\int_{\pi/2}^\pi P(\psi) d\psi = 1$.

For $\alpha \neq 1$, an exact calculation is still possible, but more complicated. We write G_0 in the form $\cos(\theta) = Cr^\alpha$, with $C = \cos(\theta_0)/r_0^\alpha$. So along C , we have $\sin(\theta) d\theta + C\alpha r^{\alpha-1} dr = 0$, i.e.,

$$r \frac{d\theta}{dr} = -C \frac{\alpha r^\alpha}{\sin(\theta)} \quad .$$

The tangent is thus

$$\begin{aligned} \frac{dy}{dx} &= \frac{\sin(\theta) dr + r \cos(\theta) d\theta}{\cos(\theta) dr - r \sin(\theta) d\theta} = \frac{\sin(\theta) - C \frac{\alpha r^\alpha \cos(\theta)}{\sin(\theta)}}{\cos(\theta) + C\alpha r^\alpha} \\ &= \frac{\sin^2(\theta) - C\alpha r^\alpha \cos(\theta)}{\sin(\theta) \cos(\theta) + C\alpha r^\alpha \sin(\theta)} \quad . \end{aligned}$$

With $\theta = \theta_0$ and $C = \cos(\theta_0)/r_0^\alpha$, we thus obtain

$$\begin{aligned} \tan(\varphi_0) &= \tan(\psi_0 + \theta_0) = \frac{\tan(\psi_0) + \tan(\theta_0)}{1 - \tan(\psi_0)\tan(\theta_0)} = \frac{dy}{dx} \\ &= \frac{\sin^2(\theta_0) - \frac{\cos(\theta_0)}{r_0^\alpha} \alpha r_0^\alpha \cos(\theta_0)}{\sin(\theta_0)\cos(\theta_0) + \frac{\cos(\theta_0)}{r_0^\alpha} \alpha r_0^\alpha \sin(\theta_0)} = \frac{\tan^2(\theta_0) - \alpha}{(1 + \alpha)\tan(\theta_0)}. \end{aligned}$$

We conclude that $\tan(\psi_0) = -\alpha/\tan(\theta_0)$, so the angular distribution is

$$d\theta = \frac{\alpha}{1 + (\alpha^2 - 1)\cos^2(\psi)} d\psi = P_1(\psi) d\psi.$$

Normalizing $P_1(\psi)$ to $P(\psi)$ so that $\int_{\pi/2}^\pi P(\psi) d\psi = 1$, and since ψ varies from $\pi/2$ to π when θ varies from 0 to $\pi/2$ and $\int_0^{\pi/2} d\theta = \pi/2$, we obtain the distribution

$$P(\psi) = \frac{2\alpha}{\pi[1 + (\alpha^2 - 1)\cos^2(\psi)]}.$$

Figure 4.112 shows several distributions $P(\psi)$ for different values of the exponent α .

The high-resolution experimental data for the cat reported in [58] and [106] show that there is a certain advantage for parallelism over orthogonality due to saturation blocking the dipole divergence at the PC. This can be modelled by taking an exponent $\alpha < 1$. The data in Fig. 4.113 from [58] fit particularly well with the value $\alpha = 0.73$.





Fig. 4.112 Examples of the angular distribution $P(\psi)$ for different values of α . The angle ψ varies from 0 to $\pi/2 \sim 1.57$. *Black* $\alpha = 1$ and $P(\psi) = 2/\pi \sim 0.64$. *Green* $\alpha = 0.7$. *Red* $\alpha = 0.5$. *Blue* $\alpha = 1.5$

Fig. 4.113 Observed distribution of the angles between the level lines of the ORs and the SFs. From [58]



4.9.3.6 Orientation–Spatial Frequency Dependence

In the classic orthogonal circular model centred on a PC where the SF diverges, for example, to $-\infty$ and grows to, let’s say, ν_0 , there is for each value of $\nu \in (-\infty, \nu_0)$ a full circle of orientations $\omega = \theta/2$. Thus, Ω^* has a *direct product* structure $\mathbb{S}_\pi^1 \times (-\infty, \nu_0)$ which is a subset of the direct product $\mathbb{S}_\pi^1 \times \mathbb{R}$. Combined with orthogonality, i.e. maximal transversality, this shows that, in this coding of the (ω, ν) , the variables ω and ν are *independent*. We shall return to the connection between transversality and independence in Sect. 4.10.2. However, note that, in the pinwheel–dipole model, ω and ν are *not* coded independently.

Indeed, in Ω^* , r varies from 0 to 1 and hence, for θ fixed, $\nu = \cos(\theta)/r$ varies between $\pm\infty$ and $\cos(\theta)$, \pm being the sign of $\cos(\theta)$. Figure 4.114 shows the set of pairs (ω, ν) represented. They are distributed throughout the full direct product $\mathbb{S}_\pi^1 \times \mathbb{R}$, but without covering it and without having a direct product structure. This agrees with results like those of Tani et al. [108], showing that the ORs and SFs develop *interdependently*, at least in cats.

We thus see that the two coding strategies of the dipole–dipinwheel and pinwheel–dipole models are very different.

4.9.4 Generality of Pinwheels

Since it is so functional, the pinwheel structure is extremely general. It is not only found in V1, but also in the other areas of the primary visual cortex. As shown, e.g., by Xiangmin Xu et al. [109] for the New World owl monkey (douroucoulis), as we go from V1 to V2 and then V3, the pinwheel structure remains, but the mesh of the singularity lattice increases with the size of the receptive fields.

Furthermore, the structure of V1 depends on the species. However, the existence of pinwheels is a robust interspecies phenomenon. Here, we give just one example in Fig. 4.115, produced by Liu and Pettigrew [110], which compares orientation maps




Fig. 4.114 Image of a regular lattice of $(\theta, r) \in [0, 2\pi] \times [0, 1]$ in the cylinder $\mathbb{S}_\pi^1 \times \mathbb{R}$ of the $(\omega = \theta/2, \nu = \cos(\theta)/r)$ (represented by the band $[0, \pi] \times \mathbb{R}$). The divergence of ν is bounded by ± 7 . The image of $[0, 2\pi] \times [0, 1]$ is clearly visible in $\mathbb{S}_\pi^1 \times \mathbb{R}$

of the V1 and V2 areas in the cat and the marmoset monkey with the equivalent in the tawny owl. We also observe a pinwheel structure in the tawny owl with an interpatch distance of about 0.9 mm.

4.10 Retinotopic Maps and Their Transversality

With ever more sophisticated experimental techniques, specialists have been able to build up several retinotopic maps corresponding to different geometric features and in several different species. We have already discussed orientation, direction, phase, and spatial frequency. Now, we shall go on to consider ocular dominance, colour, temporal frequency, and motion. With so many different issues, we may ask the following questions:

1. Are the geometric features studied with the kind of stimuli generally used, viz. gratings, etc., sufficient to describe the cortical activity induced by *complex and natural* visual stimuli?
2. What are the *relations* between the maps? Are the features independent of one another, or not?




Fig. 4.115 Pinwheels in the V1 and V2 areas of the cat, (a) and (b), respectively, and the marmoset (c) and (d), and equivalent areas in the visual cortex of the tawny owl (e) and (f). From Liu and Pettigrew [110]

3. What is the mechanism of the *dimensional reduction* which implements abstract $(2 + N)$ -dimensional fibre bundles in the 2D cortical layers (where N is the number of features)?

Regarding the first question, it seems that the feature detections are adequate if we take into account, for each feature, the preferred values, intensities, and variances, i.e. the complete tuning curves, and any possible dependences between the features (see, e.g. Issa et al. [104]).

Regarding the second question, we have seen several examples, and we shall discuss several more.

For the third question, it should be emphasized that the maps and their relations are the solutions to a problem with two opposing constraints. One concerns the uniformity of coverage of the features in the 2D cortical layer. For example, in the standard circular model of the SFs, when a pinwheel lies at the centre of a domain of high spatial frequency, all orientations are automatically combined locally with this SF (see Sect. 4.9.3.6). The other is a continuity constraint: the features are continuous functions of the retinotopic position, although with the possibility of singularities.

4.10.1 Pinwheels and Ocular Dominance

4.10.1.1 Some Experimental Data

We have seen the relations between the pinwheels and the spatial frequencies. There are also quite remarkable relations between the pinwheel structure of V1 and its organization into ocular dominance domains (ODDs), i.e. dominance of the left or right eye, or more precisely, ipsilateral or contralateral dominance relative to the given hemisphere. The ocular dominance bands measure about 1 mm in the monkey and the cat and about 2 mm in humans. As shown by Hübener et al. [100], the iso-orientation lines are essentially *transverse*, and even almost orthogonal, to the boundaries of the ODDs (see Figs. 4.116 and 4.117):

Many iso-orientation lines cross the borders between ocular dominance domains close to right angles, and the pinwheel centers are preferentially located in the middle of these ocular dominance domains. [100]

Xiangmin Xu et al. [45] found the same relationship in the bush baby or galago, a small nocturnal African primate, already encountered in Sect. 4.4.1. The histogram of the angles of intersection between the iso-orientation lines and the boundaries of the ODDs shows that transversality is statistically well represented: the angles between $3\pi/8$ and $\pi/2$ represent more than 50%. Moreover, as clearly shown by Crair et al. [36], ocular dominance peaks are situated very close to pinwheel centres in the middle of the ODDs (see Fig. 4.118).

One can also compare the ocular dominance domains with the spatial frequency domains (see, e.g., Hübener et al. [100]).




Fig. 4.116 Relations between pinwheels and ocular dominance domains. Several iso-orientation lines cross the boundaries of the ODDs almost at *right angles*. From Hübener et al. [100]





Fig. 4.117 Relations between pinwheels and ocular dominance domains in the macaque. From Obermayer, Blasdel [111]

Fig. 4.118 Pinwheel centres (green stars) and ocular dominance peaks (red stars) are neighbours and located on ODD ridges. From Crair et al. [36]



4.10.1.2 Fields and Equipotentials

As for the pinwheels, the geometrical configuration of the OD peaks and the ODDs can be made intuitive using a physical field deriving from a potential V (see Sect. 4.4.3), the peaks corresponding to the extrema of V . Then, the ODD boundaries can be identified with the level lines of the potential segregating the influence zones of the maxima and the minima. If the peaks of OD would be close to the pinwheels (but it is only partially the case), then the two OR and OD fields would be close and, as in a field equipotentials are orthogonal to field lines, this could explain the strong transversality between the iso-orientation lines of the OR field and the boundaries of the ODDs.

4.10.1.3 An Elastic Net Model

It is easy enough to simulate these field structures with orientation columns and ODDs using computational models. For example, Carreiro-Perpiñán and Goodhill [112] used elastic net models which minimize the total length of wiring while satisfying a compromise between the uniformity U and the continuity C of the cortical representation of the various features of the stimuli.²⁴ To do this, they minimize an energy $E = U + \beta C/2$. We work in the space of features

$$\left\{ \begin{array}{l} \text{position } a = (x, y), \text{ ocular dominance, orientation (angle and selectivity)} \\ \text{in polar coordinates} \end{array} \right\} ,$$

corresponding to a lattice of $N = N_x \times N_y \times N_{\text{OD}} \times N_{\text{OR}} \times 1$ stimuli in the space

$$T = [0, 1] \times [0, 1] \times [-\ell, \ell] \times [-\pi/2, \pi/2] \times [0, \rho] .$$

Let y_m be the centre of the receptive field of neuron m , $m = 1, \dots, M$, in the space T and x_n the vector of stimulus n , $n = 1, \dots, N$. The uniformity U is defined by

²⁴We have already encountered this problem of minimizing the wiring in Sect. 4.4.5.1.




Fig. 4.119 Field obtained by minimizing the energy E (see text) for $p = 1$ and $\beta = 10$. From [112]

$$U(y_1, \dots, y_M, K) = -K \sum_{n=1}^{n=N} \log \left[\sum_{m=1}^{m=M} \exp \left(-\frac{1}{2} \left\| \frac{x_n - y_m}{K} \right\|^2 \right) \right],$$

where K fixes the size of the receptive fields. The continuity C is defined by

$$C(y_1, \dots, y_M) = \sum_{m=1}^{m=M} \|f(y_m)\|^2,$$

where f is a linear combination of the neurons neighbouring neuron m which approximates an order p differential operator. Figure 4.119, already shown in Sect. 4.4.3 (see Fig. 4.40), illustrates the result of minimizing for $p = 1$ and $\beta = 10$. Figure 4.120 gives the result for $p = 3$. We see that $p = 1$ looks like the empirical maps.

4.10.1.4 A LISSOM Model

Other models of joint self-organization of maps of orientation, direction of motion, and ocular dominance can be found in the work of Bednar and Miikkulainen [113, 114]. They are based on the laterally interconnected synergetically self-organizing map (LISSOM) model. Their 2005 book entitled *Computational Maps in the Visual Cortex* [80], written with Yoonsuck Choe and Joseph Sirosh is particularly interesting:

[It] presents a unified computational approach to understanding the structure, development and function of the visual cortex.




Fig. 4.120 Field obtained by minimizing the energy E (see text) for $p = 3$ and $\beta = 10$. From [112]

For example, Fig. 4.121 shows the result of joint learning of orientation (OR, colours), ocular dominance (OD, boundaries), and direction of motion (DR, arrows) in a LISSOM model on the basis of a flux of simple stimuli. Note that the DRs are reasonably orthogonal to the ORs and that in each orientation patch there are two sub-patches with opposite directions (orthogonal to the preferred orientation of the patch). This synthesis provides a good model for the cat. We shall return to direction maps in Sect. 5.11 of Chap. 5.

4.10.2 Independent Maps and Transversality Principle

We see that, after the spatial frequency, the ocular dominance is another parameter that is implemented in the 2D neural layers. With the orientation, an abstract 3D structure collapsed to two dimensions. Now, we have an abstract 5D structure. Such a drastic reduction in dimension obviously raises questions about how the *independence* of the parameters can be represented in two dimensions. It seems that the solution discovered by evolution was to maximize a *transversality condition*: the boundaries of the frequency domains and those of the ODDs are strongly transverse to the lines of the iso-orientation field. Clearly, if there is more than one extra parameter, the transversality cannot be strong everywhere, but there can nevertheless be optimization of two opposing constraints on transversality.

Nicholas Swindale explicitly addressed the question of how to understand the optimization of the interactions between different maps in his paper [115] entitled *How many maps are there in visual cortex?* (see also [116]). To simplify, Swindale started with N binary variables (e.g. binarizing continuous variables like the orien-




Fig. 4.121 Orientation maps (*colours*), ocular dominance (*boundaries*), and direction of motion (*arrows*) in a LISSOM model. The DRs are strongly transverse to the ORs. In each orientation patch there are two sub-patches with opposite directions (orthogonal to the preferred orientation of the patch: for example, *vertical arrows* in the *red regions* correspond to the horizontal orientation, etc.). From [80, Fig. 5.29]

tation, ocular dominance, or spatial frequency), regions where a variable is constant corresponding to the different values of the features of the given stimuli. He then asked about the maximal N ensuring good coding efficiency, the effect of introducing a new map on the other maps, and the information provided by a map about the presence of other maps. To answer these questions, he used the well-known algorithm called self-organizing feature maps due to Kohonen, which is analogous to the algorithms used in the previous section.

For this, we take a retinal grid (i, j) , $i, j = 0, 1, \dots, M$ ($M = 150$) and associate with the neuron (i, j) the vector $w_{ij} = (x, y, t_1, \dots, t_N)$ defined as the set of central values of its receptive field and the features t_k it detects. This is therefore a discrete and generalized version of what we introduced at the outset when we said that a simple V1 neuron codes a contact element $(a = (x, y), p)$. In other words, we work in the fibre bundle $\pi_N : V = R \times T \rightarrow R$, where T is a fibre of dimension N . For



Fig. 4.122 Swindale's model for 6 maps. From Swindale [115]

their part, the stimuli are also coded by vectors v of V , $v_s = (x_s, y_s, u_1, \dots, u_N)$ and we take enough of them (2.5×10^6) to ensure that the learning is significant. We select random sequences of them and, at each step v , the w_{ij} are incremented by Δw_{ij} applying the following standard rule: if (i_0, j_0) is the neuron in the network which is closest to the v , then $\Delta w_{ij} = \varepsilon h(r) (v - w_{ij})$, where ε is the learning rate ($\varepsilon = 0.01$), r is the distance between (i_0, j_0) and (i, j) , and $h(r) = \exp(-r^2/2\sigma_c)$ is the learning Gaussian ($\sigma_c = 2.5$). To minimize, we then carry out simulated annealing, the algorithm already discussed in Sect. 4.4.5.1 when minimizing the total length of wiring in V1. For example, Fig. 4.122 shows a piece for $N = 6$. There are $2^N = 64$ feature values, each coded by a colour.

As the variables are binary, each map comprises blobs and stripes whose intersections encode the structural relations between the maps. The morphology of the individual maps does not change much qualitatively when other maps are added, but their structural relations change much more. Figure 4.123 shows the boundaries for $N = 6$. We observe strong transversality relations.

Hongbo Yu et al. [46] also examined the relations between the three maps, viz. orientation, ocular dominance, and spatial frequency, but using continuous variables and focusing on the way strong transversality codes the independence of the associated variables. They considered the *gradients* of the variables and showed first that the gradients are maximal in disjoint regions and then that the transversality is *maximal* when the two gradients are jointly high enough:

Two features are mapped orthogonally in their high-gradient overlap regions. [46, p. 277]

Figure 4.124 shows the orientation field lines (the level lines of the 'orientation' variable) and the OD level lines. The pinwheel centres and the ODD boundaries



Fig. 4.123 Boundaries of the domains of the $N = 6$ binary variables in Fig. 4.122. From Swindale [115]


avoid one another since the pinwheel centres are located on the ridges of the ODDs. The grey zones are those where the two gradients are jointly highest. We observe that transversality is indeed very strong; in fact, we almost have orthogonality.

One interesting result obtained by Swindale in [115] concerns the retinotopic map. Figure 4.125 shows the projection of the cortical network on the retinal space for the case $N = 7$. Note that there are many folds and pleats. Two remarks are in order:

1. Firstly, there should be no surprise in finding folds and pleats. Indeed, according to the well-known theorem due to Hassler Whitney and generalized by René Thom, the only singularities possible in a generic differentiable map between 2D manifolds are fold lines and isolated pleat points.
2. Secondly, these singularities are not those of the conformal retinotopic map (a complex logarithm, see Sect. 4.2) between the retinal positions and the cortical positions. They are induced by what happens in the fibres of the bundle $\pi_N : V = R \times T \rightarrow R$ and show that the values of the features $w_{ij} = (x, y, t_1, \dots, t_N)$ do not generate a single-valued section of π_N , but rather a multivalued section whose image is a surface in V that is not everywhere transverse to the fibres.²⁵

²⁵For an adequate treatment of this point, one must introduce the rather technical geometric notion of a Lagrangian sub-manifold. We shall say a little more about this in the second volume. Here we only make elementary remarks about the geometry.

Fig. 4.124 Iso-orientation lines and ODD level lines. Grey regions are those where the two gradients are at their highest. We observe that transversality is very strong there. From Yu et al. [46]



4.10.3 Binocularity

4.10.3.1 Ocular Dominance and Binocular Disparity

The ocular dominance maps and ODDs of the contralateral and ipsilateral monocular cells must be studied in relation to *binocular disparity* (BD) of the binocular cells. Among others, Prakash Kara and Jamie Boyd have studied the functional architecture of the BD and its relationship with OD in a piece of the V2 area (18) of the cat [117]. Using two-photon confocal microscopy, method already discussed in Sect. 4.7.3 in the context of Ohki's work on pinwheels and which can be used to measure the activity of several hundred individual neurons in layer 2/3 (a region of about 300 μm), Kara and Boyd were able to show that there is a BD selectivity map and that OD and BD are *mutually independent* at the level of individual neurons. Once again, this independence is manifested through strong transversality properties: in regions where the gradients of the two variables are both high, they are orthogonal.

The stimuli used are gratings with the preferred direction of motion and orientation of the cell being measured. Either one of the eyes is stimulated by such a grating




Fig. 4.125 Projection of the cortical network on the retinal space in the Swindale $N = 7$ model. From Swindale [115]

and the other by a uniformly grey stimulus, or both eyes are stimulated by such gratings, with the same orientation and direction, but out of phase (8 values of the phase difference are used). Figure 4.126 shows the responses of several cells to the right eye (R) and the left eye (L) and to the phase of the BD. Figure 4.127 shows maps of the BD and the OD, confirming that, in regions where both gradients are high, the level lines are close to orthogonal.

4.10.3.2 Bistability and Binocular Rivalry

Since we have been discussing binocularity, let us say a word about the phenomena of *binocular rivalry* and *bistable images*, important subjects when studying the relationship between high and low levels of visual cognition.

Ambiguous bistable images like the Necker cube, which can be perceived in two different ways, are well known. When we stare at the image for long enough, we observe regular switches between the conflicting interpretations. It is thus natural to assume that there are two competing populations of neurons implementing the two interpretations, the dominant population selecting the perceived interpretation at a given time and inhibiting the other. When the inhibition falls sufficiently or the cues

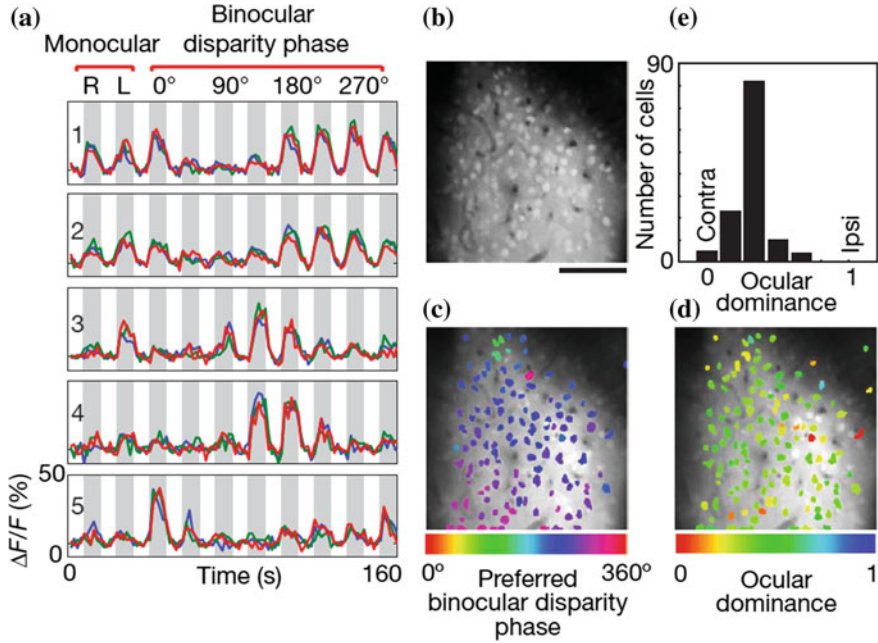


Fig. 4.126 Responses of 5 cells to monocular stimuli and the phase of the binocular disparity. For example, cell 1 responds to L and R and to phase differences of $0, \pi, 5\pi/4, 3\pi/2,$ and $7\pi/4$, but not to phase differences of $\pi/4, \pi/2,$ and $3\pi/4$. **c** gives the BD, **d** the OD, and **e** the histogram of the OD. From Kara and Boyd [117]




Fig. 4.127 Maps of the BD and the OD. In regions where both gradients are high, level lines are close to orthogonal. From Kara and Boyd [117]

favouring the other interpretation evolve sufficiently, a switch can occur. Here, we should ask at what level in the elaboration of percepts an interpretation is selected. Is it a basically low-level mechanism or rather some high-level cognitive mechanism involving knowledge? In actual fact, it almost certainly involves both because any interpretation involves *inference*.

Spontaneous bifurcation phenomena are well known in every field. We shall meet some fine examples in the second volume. They can be modelled using the following guiding idea introduced by René Thom in [118] and [119], an idea already discussed in Sect. 1.3 of the *Preface*. Let S be a system satisfying the following hypotheses:

1. Inside S , there is an internal process X which specifies the internal states that S can occupy.
2. The internal process X specifies all the internal states of S .
3. There is a selection criterion I which selects the current state among the various possible internal states on the basis of certain criteria specific to the system, and which may vary significantly.
4. Finally, the system S is controlled, with good regularity properties, by a certain number of control parameters. These vary in a space W called the *external space* (or control space, or substrate space) of S , to distinguish it from the internal space of S . The internal process X is thus a process X_w depending on w .

If \mathfrak{X} is the space of possible internal processes, the system S will then be described by the field $\sigma : W \rightarrow \mathfrak{X}$ associating the corresponding internal process X_w with $w \in W$, and also by the selection criterion I . As w varies, the actual state will vary and, for certain critical values of w , it may bifurcate towards another internal state.

The best known model of this kind is known as the cusp model. The states of the system are represented by the minima of a potential function $f_w : \mathbb{R} \rightarrow \mathbb{R}$ parametrized by a parameter w varying in a 2D external space W , and the internal dynamics X_w are the gradient dynamics of the f_w . Depending on the value of w , f_w has either one minimum or two minima A and B separated by a threshold, and when there are two minima, these are in competition. Hence, there are two kinds of possible bifurcation corresponding to the exceptional values of w varying along three lines of W : along two of these lines, A or B disappear, and along the third, A and B compete and balance each other (see Fig. 4.128).

We shall return to this model in Sects. 5.9.2 and 5.11.3 of Chap. 5. It was used to model bistable images in the 1970s by Christopher Zeeman, the minima of f_w corresponding to attractors of the relevant neural dynamics.

Binocular rivalry does not concern the conflict between two possible interpretations of a given stimulus, but something deeper and more enigmatic, namely the conflict between two different stimuli. We consider two stimuli A and B , e.g. a chessboard and a face, and we present one to each eye using appropriate apparatus. The two stimuli are thus both present in $V1$. However, the subject does not perceive a superposition $A + B$, but in each case only one image, the two images alternating regularly. In other words, there are successive intervals of time in which the subject perceives only A or B (exclusivity and uniqueness), and between these intervals, short transition periods during which the stimulus that is actually present, say A , is destabilized and bifurcates towards the other stimulus B . Temporal series of such spontaneous shifts have been studied in detail.

There are roughly speaking two classes of theory to explain the phenomena of binocular rivalry: on the one hand, low-level theories which say that the suppression




Fig. 4.128 Universal unfolding f_w of the cusp singularity. The parameter w varies in a 2D external space W and the internal dynamics is the gradient descent $-\text{grad}(f_w)$. Depending on the value of w , f_w has either one minimum or two competing minima separated by a threshold. There are two kinds of catastrophe along three strata in W : either one of the minima bifurcates by collision with the maximum, or the two minima compete and balance each other. This figure will be taken up again in Fig. 5.39 and explained further in Sect. 5.11.3 of Chap. 5

of one of the competing stimuli already begins in the V1 area, or even before V1, and on the other hand, high-level theories which say that the suppression is made through cognitive inferences. According to the low-level hypothesis known as interocular competition, rivalry arises because the two images cannot be merged by binocular stereopsis and this mismatch blocks one or other of the monocular pathways. But the two kinds of theory are probably both partly correct. Indeed, fMRI studies show that high-level non-visual areas (frontal and parietal) are active during bifurcations. This

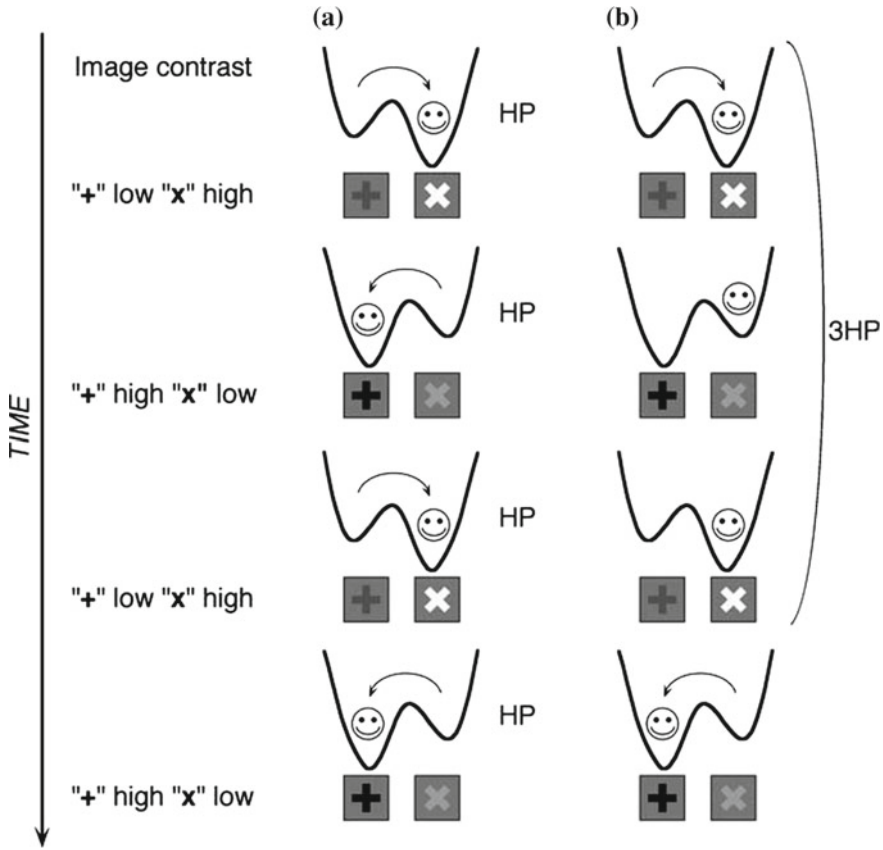



Fig. 4.129 Cusp model for binocular rivalry. The two percepts are symbolized by + and \times and the depths of their respective potential wells are alternately modulated by the image contrast. Noise provides the means to overcome the threshold (the maximum). Stochastic resonance is manifested by the fact that the distribution of the dominance periods has resonance peaks for the values $(2k + 1) HP$, where HP is the half period of contrast modulation. **a** Peak for HP . **b** Peak for $3HP$. From Kim et al. [121]

is therefore a complex phenomenon. For a good summary, the reader is referred to the paper [120] by Randolph Blake and Hugh Wilson.

The cusp dynamical model turns up in certain studies of binocular rivalry. For example, Yee-Joon Kim et al. [121] use it to apply the methods of stochastic resonance. In Fig. 4.129, the two percepts are symbolized by + and \times and the depths of their respective potential wells are alternately modulated by the *image contrast*. Noise is introduced to overcome the threshold represented by the maximum separating the two minima. Stochastic resonance is manifested by the fact that the distribution of the dominance periods, i.e. the time series of the shifts, has resonance peaks for the values $(2k + 1) HP$, where HP is the half period of contrast modulation.

Fig. 4.130 Initial fast field (h_1, h_2) with two nullclines $\partial h_1/\partial t = 0$ (red) and $\partial h_2/\partial t = 0$ (blue) intersecting at three fixed points (two attractive nodes and one partially repulsive saddle point). Recomputed from de Jong [122]



In his 2008 master's thesis entitled *The Dynamics of Visual Rivalry* [122], Ties Marijn de Jong investigated a simple but explicit model of bifurcation with slow/fast dynamics. The competing stimuli were orthogonal gratings. By averaging the dynamics of single neurons, we obtain, for the two populations $i = 1, 2$:

- The local activity fields h_i of the percepts (average membrane potentials), these being the fast variables.
- Slow adaptation variables a_i expressing the fact that the activity of the dominant percept gradually weakens.

Introducing a scale change parameter $\tau \gg 1$ between slow and fast times and averaging the standard neural equations, the author produced the following model:

$$\begin{cases} \tau \frac{\partial h_i}{\partial t} = X_i - (1 + a_i) h_i - \gamma \sigma(h_j) , \\ \frac{\partial a_i}{\partial t} = -a_i + \alpha \sigma(h_i) , \end{cases}$$

where α and γ are constants arising from the underlying neural interactions, the X_i are the initial forces of the two stimuli as inputs, and where $\sigma(h)$ is a sigmoid curve, i.e. a C^∞ approximation to the Heaviside step function, equal to zero for $h < 0$ and unity for $h \geq 0$. As the mean firing rate is $\sigma(h)$, $h_i < 0$ implies that the stimulus i is suppressed by the perceptual system.

If we begin with a symmetric rivalry, that is two stimuli of the same strength and equal adaptation coefficients, and examine the field (h_1, h_2) , a classic bifurcation scenario comes to light (see Fig. 4.130). The slow manifolds are the nullclines $\partial h_i/\partial t = 0$ with a fixed. At the beginning, they intersect at three points which are equilibrium points since the field vanishes there. Two of these points are attractive (nodes), on either side of a third which is partially repulsive (a saddle point). For each

attractor, one of the h_i is strictly positive while the other (h_j) is strictly negative, implying that only the stimulus i is perceived. If we begin in a state with no activity, close to 0, and if we activate the system with the stimuli, the fast dynamics projects the system onto one of the attractors. From this point, the slow dynamics modifies the field and adaptation breaks the initial symmetry $a_1 = a_2$. It displaces the nullclines, thereby leading to a bifurcation by the merging of the initial attractor and the saddle point. The system then jumps into the other attractor. We shall return to this kind of fast/slow dynamics and bifurcations in the second volume.

The most profound causes of perceptual bistabilities are still debated. For example, in [123], David Leopold and Nikos Logothetis criticize the favoured hypothesis according to which the bifurcations (spontaneous reversals) result from antagonistic connections. In their view:

[Alternations] reflect responses to active, programmed events initiated by brain areas that integrate sensory and non-sensory information to coordinate a diversity of behaviours. [123, p. 254]

4.10.4 Blobs and Colour

We should also say a word about ‘blobs’ of cytochrome oxidase (CO) in $V1$ and $V2$ which are sensitive to *colour* and already process it in $V1$. By imaging, we obtain colour-selective response maps in those regions which, in $V1$, are centred on the ODDs, but which are in fact barely selective to orientation and which do not overlap with the central regions of the pinwheels. In $V2$, the colour regions coincide with the thin stripes and the orientation regions with the pale and thick stripes. There is therefore a kind of functional segregation of colour and orientation corresponding in part (this point is debated) to the distinction between blobs and interblob regions.

The structure of the blob map has been studied by many specialists, such as Haidong Lu and Anna Roe [124], using optical imaging methods on the macaque. In primates, the blobs in $V1$ are found mainly in layers 2 and 3, but also in layers 1, 4B, 5, and 6. They measure in the range 150–250 μm , and there are about $5/\text{mm}^2$. They tend to be centred on the ODD axes with an average period of about 350 μm (see Polimeni et al. [87]). Figure 4.131 shows their distribution. It is valid for a monkey like the macaque, but in a prosimian primate like the bush baby, the constraints on the positions of the pinwheels and the blobs within the ODDs are less stringent (see Xiangmin Xu et al. [45]). Finally, Fig. 4.132 shows a diagram of $V1$ with pinwheels, ODDs, and blobs.

We shall return to the band structure of $V2$ in Sect. 5.8 of Chap. 5.





Fig. 4.131 Distribution of blobs of cytochrome oxidase in the ODDs of a macaque. *Red*: left eye. *Green*: right eye. They are centred preferentially on the ODD axes. *Arrows* indicate blobs furthest from the centres. *Scale bar*: 1 mm. From Lu and Roe [124]

Fig. 4.132 Diagram of V1 with pinwheels, ODDs, and blobs. From [24]



4.10.5 *Functionality of Maps*

Here, we see just how rich the structure of $V1$ is. All the various ‘secondary’ variables, engrafted in the sense of Hubel onto the primary spatial variables, are processed by micromodules measuring $400\ \mu\text{m} \times 800\ \mu\text{m}$ and containing some 60 000 neurons. With their functional architectures, they constitute very precise maps that are just as precisely correlated with one another. This is how, according to Yu et al. [46], dimensional reduction can:

[...] smoothly map several (more than two) response properties onto a two-dimensional cortical surface.

It is also how it can solve the fundamental functional problem of ensuring the uniformity and continuity of features with respect to their positions by the design of the neural hardware. To minimize the complexity of this structure and make a ‘simplex’ structure in the sense of Alain Berthoz (who generalized what we shall say about 1-jets in Sect. 5.4.6 of Chap. 5), evolution grouped together functionally related neurons locally through the spatial extent of the cortical layers.

4.11 Hemispheres and Callosal Connections

Here, we shall say something about the way the two halves of $V1$ in the two hemispheres are connected up. This provides a remarkable example of a geometrical *gluing* process. It is the *callosal connections* of the corpus callosum that do this. The corpus callosum is the biggest bundle of nerve fibres in mammals. In humans, it comprises some 200 million axons.²⁶ The region of the visual field located close to the vertical meridian, called the transition zone (TZ) or visual midline (VM), is projected onto the two parts of $V1$ in the vicinity of the $V1$ – $V2$ boundary.

The structure of the gluing map is fascinating. Figure 4.133 shows this map for the cat (areas 17/18). The sub-zones *A*, *B*, *C*, *D*, *E* of a hemisphere are connected to the zones with the same label in the other hemisphere: zones within the TZ are connected to zones outside it and conversely.

If we cut through the fibres in the optical chiasm coming from the nasal hemiretina, while keeping those from the temporal hemiretina (split-chiasm preparation), the right visual hemifield projects onto the left $V1$ area via the left eye and the activity of the right $V1$ area becomes entirely due to the callosal connections (see Fig. 4.134). Depending on which eye is stimulated, we can thus activate either the geniculocortical pathway or the transcallosal pathway.

The problem is to understand the distribution of the callosal connections in relation to the pinwheel structure in $V1$. This is a very delicate matter. An experiment by William Bosking (Bosking et al. [126]) shows (see Fig. 4.135) that, for the tree shrew

²⁶Recall that the optic nerve contains about 1.5 million axons, so less than a hundredth of the number.




Fig. 4.133 Structure of the callosal gluing map between the two hemispherical parts of V1 and V2 in the cat. Zones A, B, C, D, E of one hemisphere are connected to zones with the same label in the other hemisphere. From Rochefort [125]




Fig. 4.134 If we cut through the optical chiasm, the right visual hemifield projects onto the left V1 area via the left eye and the activity of the right V1 area is entirely due to the callosal connections. From Rochefort [125]




Fig. 4.135 Callosal connections in the tupaia (tree shrew). The two *red/green* sites of $V1_L$ in **a** project onto $V1_R$ regardless of orientations as shown in **b** and **c**. In $V1_R$, we find a rather uniform distribution of activated neurons: they are distributed across several regions of different preferred orientation (*colour*). From Bosking et al. [126]

(tupaia), if we inject rhodamine into a small region of $V1_L$ with vertical preferred orientation (red circle in a black region) and fluorescein into another small region with horizontal preferred orientation (green circle in a neighbouring white region), the callosal projections onto $V1_R$ do not exhibit any orientation specificity:

Callosal connections appear to terminate without regard for the map of orientation preference, showing little sign of the orientation-specific modular and axial specificity that is characteristic of long-range horizontal connections. (Bosking et al. [126])

For the cat, the situation seems to be different. The classic work by Olavarria [127] showed that the distribution of retrogradely labelled callosal cells is polarized by the ocular dominance domains:

Callosal cells correlate preferentially with contralateral ODCs (ocular dominance columns) within the 17/18 transition zone (TZ), and with ipsilateral ODCs in regions of areas 17 and 18 located outside the TZ.

Milleret and Rochefort [128, 129] made further investigations using in vivo optical imaging of intrinsic signals coupled with a method for reconstructing and labelling the callosal axons. They showed that the callosal connections *do conserve* orientation selectivity. More precisely, using the split-chiasm preparation with an injection site in the left hemisphere, they were able to reconstruct the distribution, in layers II, III, and IV (upper part) of the transition zone (ZT) of the right hemisphere, of the synaptic buttons of several labelled callosal axons. They thus observed that these axons

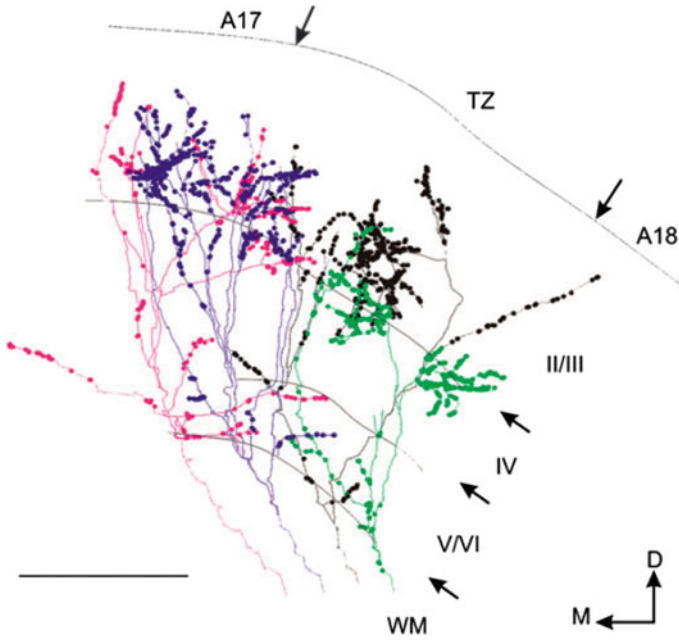


Fig. 4.136 Front view of the reconstruction of four callosal axons in cats (two neighbouring injection sites and two axons per site) and their clusters of synaptic buttons in layers II, III, and IV (*upper part*) of the transition zone TZ between area 17 and area 18. WM white matter, D dorsal, M medial. Scale bar 500 μm . From Rochefort et al. [129]




Fig. 4.137 Conservation of orientation by transcallosal connections in the cat. *Left* Injection site in the left hemisphere. *Right* Reconstruction of the distribution of the synaptic buttons of a labelled callosal axon in the right hemisphere. The axon projects into isochromatic zones (*blue*) of the same colour as the injection site. From Rochefort [125]

project into isochromatic zones relative to the colour of the injection site, thereby showing that the orientation is conserved. Figure 4.136 shows four reconstructed axons. Figure 4.137 shows the conservation of orientation by the transcallosal connections.

4.12 Homogeneous and Inhomogeneous Qualities

4.12.1 Responses to Homogeneous Surfaces

Up to now the basic geometric feature we have studied is the orientation of contour elements. However, if we consider in the most qualitative phenomenological way a natural image made up of identifiable objects, we see immediately that there are two fundamentally different kinds of local structure:

1. Locally *homogeneous* points near which features vary only slightly and in a continuous way.
2. Locally *heterogeneous* points near which certain features reveal discontinuities.

This dichotomy between homogeneous and heterogeneous is absolutely fundamental and very general, going back at least to Aristotle, who, himself referring to Anaxagoras, made a distinction in biology between homoeomeric parts ('made up' of similar [omios] parts [meros]) and anhomoeomeric parts.²⁷ As we shall see in the second volume, this distinction is of key importance in Husserl's phenomenology of perception. It also underpins René Thom's morphological models. Thom referred to the locally homogeneous points of the substrate of a form as its *regular points* and the locally inhomogeneous points as its *singular points*. The latter are the phenomenological manifestation of symmetry breaking in the underlying processes, and this is why the theory of singularities, the theory of bifurcations, and the theory of symmetry breaking formed the backbone of all his models of natural morphology, not just in physics, chemistry, and biology, but also in psychology and the social sciences.

These models are based on Thom's guiding idea as discussed in Sect. 4.10.3.2. Here, we apply to the case where the external space W is that of the visual field and where the attractors of the internal dynamics X_w implement sensory qualities. In this case, when the internal dynamics X_w undergoes a bifurcation, the point w is singular.²⁸

The dichotomy between regular and singular points was also important in the visual theories of Grossberg [131], which assume that there are two fundamental systems of visual perception:

²⁷ *On generation and corruption*, I, 5, 312b.

²⁸ All the details can be found in René Thom's two books [118] and [119]. For a didactic introduction, see [130].

1. The *boundary contour system* (BCS), which detects, enhances, and completes edges using a ‘spatially long-range cooperative process’.
2. The *featural contour system* (FCS), which fills in the regions bounded by the BCS with qualities (‘featural filling-in’): ‘These filling-in processes lead to visible percepts of colour-and-form-in-depth at the final stage of the FCS’.

According to Grossberg [131, p. 35]:

Boundary contours activate a boundary completion process that synthesizes the boundaries that define perceptual domains. Feature contours activate a diffusion filling-in process that spreads featural qualities, such as brightness or color, across these perceptual domains.

It is thus interesting to know whether the dichotomy between homogeneous-regular and heterogeneous-singular is relevant in the coding of features in V1 or V2.²⁹ Toshiki Tani et al. [132] reached this conclusion by exhibiting in the 17/18 areas of the cat a map of neurons that detect the *interior* of homogeneous domains and respond to spots which completely cover their receptive fields. These neurons are located in the V2 area near the boundary with the V1 area. The authors conclude that, in the primary visual cortex, there is not only a representation of contours but also a representation of *surfaces*. They stress that these regions of V2 are small and centred on the pinwheels where all the orientations are present.

As these zones for local detection of homogeneous surfaces are located close to the boundary between V2 and V1, which is the transition zone (TZ) of the vertical meridian where the corpus callosum comes into play, as we have seen, the authors suggest that these neurons may transmit information about the interior of homogeneous zones from one hemifield and one hemisphere to the other:

One possible function of the surface-responsive regions may be to link visual information about large surfaces extending across both the right and left hemifields. [132, p. 1123]

We shall now specify the problem of regular and singular points by returning to the processing of *colour*.

4.12.2 Colour Processing

4.12.2.1 First Steps: Goethe, Helmholtz, Hering

We discussed colour processing in the retina in Sect. 3.2.5 of Chap. 3, and we have just been talking about blobs in V1 and V2. Let us now say a little more about this.

The trichromatic theory developed by Thomas Young and Hermann von Helmholtz in the first half of the nineteenth century involved a remarkable anticipation of the

²⁹The pinwheel structure is based on the dichotomy between regular and singular points, but this concerns field lines of the orientation field. Here we are talking about something quite different.

way the retina actually processes colour. However, it was clearly inadequate. Goethe had already recorded decades of phenomenological observation of colour perception in his *Farbenlehre* of 1810. He was interested in the phenomenon of colour as an experienced quality and sought to describe this experience in as precise a way as possible, since accurate description was for him essential. From the details of these observations (which are experimental in the scientific sense and also in the subjective sense of experience), he hoped to infer more explicative underlying mechanisms. In particular, he noted that the persistent image of a coloured figure has another colour: pieces of white paper on a yellow wall become tinted with purple, and when we remove an orange sheet of paper from a white wall, this induces a blue image and an orange background. As he notes in § 60:

These phenomena are of the greatest importance, because they point to laws of vision.

Indeed, they point to the antagonisms between complementary colours, namely yellow/purple (Y/R+B), blue/orange (B/Y+R), and red/green (R/B+Y), where R = red, G = green, B = blue, and Y = yellow.

Later, in 1892, the great Austrian physiologist Ewald Hering (1834–1918) developed an antagonistic theory, known as the opponent process theory, in a simplified tetrachromatic form, assuming that there are in fact four fundamental colours RGBY, organized in two pairs of complementary colours R/G and B/Y. With this theory, he was able to solve the problem of yellow, which seemed to be an elementary colour rather than a mixture. Hence, the idea that there were indeed three RGB detectors, but four simple primary physiological colours.

Hering's theory of pairs of complementary colours anticipated in a quite remarkable way what we know today about colour processing in the cortex, between the LGN and V1.

4.12.2.2 Chromatic Opponency: Single Versus Double

In the LGN, colour-tuned cells are essentially parvocellular 'single-opponent' cells detecting R/G contrasts, i.e. L/M.³⁰ Magnocellular cells are insensitive to colour. However, there are also cells in the koniocellular pathway, discovered in 1994, which detect B/Y contrasts, i.e. S/(L+M). In V1, the colour cells located in blobs are of two kinds. The first calculate *chromatic contrasts* R/G, i.e., L/M, and B/Y, i.e. S/(L+M). The second calculate *spatial contrasts* between the complementary colours R/G and B/Y. These are therefore doubly antagonistic (double-opponent cells) which calculate opponency in both the external space and the chromatic space. They can be denoted $\pm L / \mp M$. In other words, these are orientation cells that also detect chromatic contrasts. They are essential because, as emphasized by Shapley and Hawken [133, p. 701]:

³⁰Recall (see Sect. 3.2.5 of Chap. 3) that the three kinds of cones in humans are L/M/S, L red, M green, S blue.

Color, form and motion are inextricably linked as properties of objects in visual perception and in the visual cortex.

So from the lowest levels of the V1 and the V2 areas, there is a functional entanglement between the spatiality of perceived scenes and the colours of objects making them up. The brain must reconstruct by inference from the colour the objective reflectance of the surfaces perceived and this independently of a host of extremely variable factors such as the illumination of sources, indirect irradiation, angles of incidence, and reflection. This *inverse problem* is exceedingly hard to solve. The *direct* problem is this: for each position $a = (x, y)$ of the visual field, given the reflectances $\rho(a, \lambda)$ and an illumination spectrum $\sigma(a, \lambda)$, to calculate the spectrum $\gamma(a, \lambda) = \sigma(a, \lambda) \rho(a, \lambda)$, where $\gamma(a, \lambda)$ is itself encoded locally by the excitation of the L/M/S cones at a . The inverse problem is not well posed, like most inverse problems, and can only be solved if the system has priors at its disposal, i.e. priors in the Bayesian sense, regarding the $\sigma(a, \lambda)$. An examination of these difficulties can be found in the paper by Foster [134].

In [133], Robert Shapley and Michael Hawken study the receptive profiles $\varphi(a, \lambda)$ of colour-sensitive V1 neurons which, in the linear case, act by convolution on the signal $I(a, \lambda)$. For a trichromatic species like the macaque or humans, they have the form

$$\varphi(a, \lambda) = \alpha_L L(\lambda) r_L(a) + \alpha_M M(\lambda) r_M(a) + \alpha_S S(\lambda) r_S(a) ,$$

where the α are coefficients, the $r(a)$ are spatial receptive profiles, and $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ are the spectral responses of the different kinds of cone. For example, the opponent cells R/G, i.e. L/M, correspond to $\alpha_S = 0$ and $\alpha_L = -\alpha_M$, while the B/Y, i.e. S/(L+M) correspond to $\alpha_S = -(\alpha_L + \alpha_M)$. Figure 4.138, produced by Elizabeth Johnson in [135] and reproduced in [133], represents the receptive profile of a double-opponent cell R/G, where the ON and OFF regions have positive and negative heights, respectively. We see the level lines corresponding to L-cones and M-cones. Figure 4.139 shows schematically single- and double-opponent cell models with oriented or circular spatial profile. Note in particular the orientation cells $L \pm /M \mp$.

The difference between single- and double-opponent cells is crucial from the theoretical point of view, but also for the purposes of modelling. Indeed, single-opponent cells respond to *interior* regions of continuously varying colour domains (with no detectable spatial contrast), whereas double-opponent cells respond to the edges of such regions where the colour undergoes a *qualitative discontinuity*. This confirms neurophysiologically the phenomenological models suggested by René Thom at the end of the 1960s [118, 119], models which we have already related to Stephen Grossberg's work in [131].




Fig. 4.138 **a** Complete receptive profile of a double-opponent cell R/G of V1. This is the superposition of an L profile and an M profile. The ON and OFF regions (where an increase or a decrease in luminosity bring about a greater response) have positive and negative heights, respectively. **b** and **c** Level lines of the L-cones and the M-cones obtained by the technique of inverse correlation. The scale of growth goes from *blue* (low) to *red* (high). At the position marked with a *star*, the L map is OFF (*blue*) and the M map is ON (*red*), and at the position marked with a *circle*, the opposite. From Johnson [135] and Shapley and Hawken [133]

4.12.2.3 Double Opponency and Natural Images

As pointed out in Sect. 3.6.2 of Chap. 3, the receptive profiles (RPs) can be deduced from the statistical properties of natural images using the techniques of independent component analysis (ICA), insofar as they minimize redundancy in the neural representations of sense data. Likewise for the RPs of colour processing neurons. For example, Dharmesh Tailor, Leif Finkel, and Gershon Buchsbaum showed in [136] that independent spatiochromatic filters extracted from databases of natural colour images look much like those of double-opponent colour cells observed, for example, in the macaque V1 area. For statistical correlations between colour and orientation in natural images, see also [137].

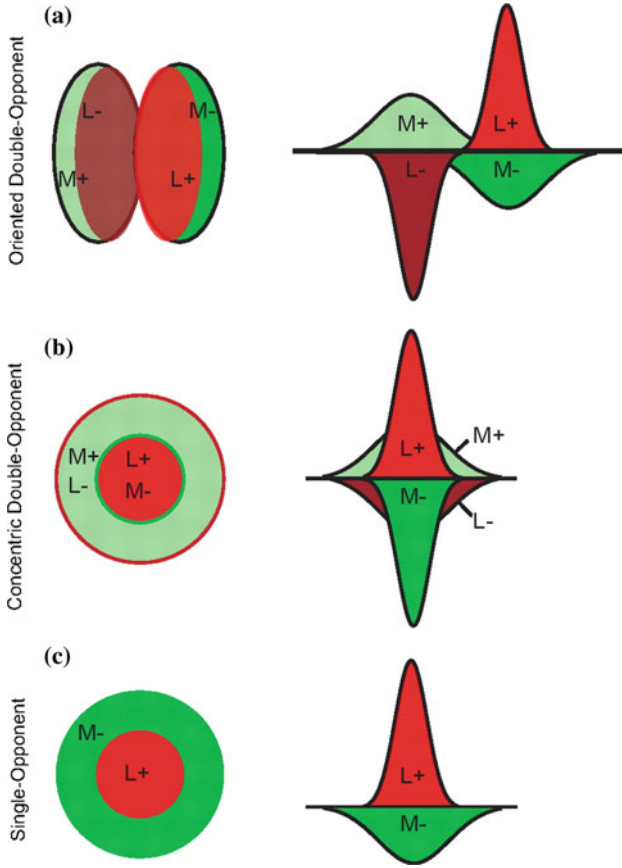


Fig. 4.139 Diagrams of single- and double-opponent L/M cells with oriented or circular spatial profile. From Johnson [135]

References

1. Lee, T.S., Mumford, D., Romero, R., Lamme, V.A.F.: The role of primary visual cortex in higher level vision. *Vision Res.* **38**, 2429–2454 (1998)
2. Beaudot, W.H.A., Mullen, K.T.: Orientation selectivity in luminance and color vision assessed using 2D band-pass filtered spatial noise. *Vision Res.* **45**, 687–695 (2005)
3. Hasboun, D.: neur@nat. <http://www.chups.jussieu.fr/ext/neuranat/>
4. Brain. <http://thebrain.mcgill.ca>
5. Visual Cortex. <http://www.vision.ee.ethz.ch/en/>
6. Vision Web. <http://webvision.med.utah.edu>
7. Ungerleider, L.G., Mishkin, M.: Two cortical visual systems. In: Ingle, D.J., Goodale, M.A., Mansfield, R.J.W. (eds). *Analysis of Visual Behavior*, MIT Press, pp. 549–586. Cambridge, MA (1982)
8. Wandell, B.A., Dumoulin, S.O., Brewer, A.A.: Visual field maps in human cortex. *Neuron* **56**, 366–383 (2007)

9. Larsson, J., Heeger, D.J.: Two retinotopic visual areas in human lateral occipital cortex. *J. Neurosci.* **26**(51), 13128–13142 (2006)
10. Tootell, R.B., Switkes, E., Silverman, M.S., Hamilton, S.L.: Functional anatomy of macaque striate cortex. II retinotopic organization. *J. Neurosci.* **8**(5), 1531–1568 (1988)
11. Dumoulin, S.O., Wandell, B.A.: Population receptive field estimates in human visual cortex. *NeuroImage* **39**, 647–660 (2008)
12. Brewer, A.A., Press, W.A., Logothetis, N.K., Wandell, B.A.: Visual areas in macaque cortex measured using functional magnetic resonance imaging. *J. Neurosci.* **22**(23), 10416–10426 (2002)
13. Hadjikhani, N., Liu, A.K., Dale, A.D., Cavanagh, P., Tootell, R.B.H.: Retinotopy and color sensitivity in human visual cortical area V8. *Nat. Neurosci.* **1**(3), 235–241 (1998)
14. Hooks, B.M., Chen, C.: Critical periods in the visual system: changing views for a model of experience-dependence plasticity. *Neuron* **56**, 312–326 (2007)
15. Gur, M., Snodderly, D.M.: Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. *J. Physiol.* **585**(2), 383–400 (2007)
16. Vision. Accessible at http://jeanpetitot.com/Woodruff-Pak_vision.ppt
17. Alexander, D.M., Sheridan, P., Bourke, P.D., Konstantatos, O., Wright, J.J.: Global and local symmetry of the primary visual cortex: derivation of orientation preference. <http://paulbourke.net/papers/visualneuro/>
18. Balasubramanian, M., Polimeni, J., Schwartz, E.L.: The V1-V2-V3 complex: quasi conformal dipole maps in primate striate and extra-striate cortex. *Neural Networks* **15**(10), 1157–1163 (2002)
19. DeAngelis, G.C., Ghose, G.M., Ohzawa, I., Freeman, R.D.: Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J. Neurosci.* **19**(9), 4046–4064 (1999)
20. Beaudot, W.H.A., Mullen, K.T.: Orientation discrimination in human vision: psychophysics and modeling. *Vision Res.* **46**, 26–46 (2006)
21. Snippe, H.P., Koenderink, J.J.: Discrimination thresholds for channel-coded systems. *Biol. Cybern.* **66**, 543–551 (1992)
22. Ringach, D.L.: Population coding under normalization. *Vision Res.* **50**, 2223–2232 (2010)
23. Braitenberg, V., Braitenberg, C.: Geometry of orientation columns in the visual cortex. *Biol. Cybern.* **33**, 179–186 (1979)
24. Visual system. <http://brain.phgy.queensu.ca/pare/assets/Higher%20Processing%20handout.pdf>
25. Lund, J.S., Angelucci, A., Bressloff, P.C.: Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cereb. Cortex* **13**(1), 15–24 (2003)
26. Van Hooser, S.D.: Similarity and diversity in visual cortex: is there a unifying theory of cortical computation? *Neuroscientist* **13**(6), 639–656 (2007)
27. Hubel, D.H.: *Eye, Brain and Vision*. Scientific American Library, W.H. Freeman & Co, New York (1988)
28. Ben-Shahar, O., Zucker, S.: Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Comput.* **16**(3), 445–476 (2004)
29. Koenderink, J.J., Van Doorn, A.J.: Representation of local geometry in the visual system. *Biol. Cybern.* **55**, 367–375 (1987)
30. Zhang, J., Wu, S.: Structure of visual perception. *Proc. Natl. Acad. Sci. USA* **87**, 7819–7823 (1990)
31. Weliky, M., Bosking, W., Fitzpatrick, D.: A systematic map of direction preference in primary visual cortex. *Nature* **379**, 725–728 (1996)
32. Bonhöffer, T., Grinvald, A.: Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* **353**, 429–431 (1991)
33. Ohki, K., Reid, R.C.: Specificity and randomness in the visual cortex. *Curr. Opin. Neurobiol.* **17**(4), 401–407 (2007)

34. Grinvald, A., Shoham, D., Shmuel, A., Glaser, D., Vanzetta, I., Shtoyerman, E., Slovlin, H., Sterkin, A., Wijnbergen, C., Hildesheim, R., Arieli, A.: In-vivo optical imaging of cortical architecture and dynamics, Technical Report GC-AG/99-6, The Weizmann Institute of Science (2001)
35. Blasdel, G.G., Salama, G.: Voltage sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* **321**, 579–585 (1986)
36. Crair, M.C., Ruthazer, E.S., Gillespie, D.C., Stryker, M.P.: Ocular dominance peaks at pinwheel center singularities of the orientation map in cat visual cortex. *J. Neurophysiol.* **77**, 3381–3385 (1997)
37. Bosking, W.H., Zhang, Y., Schofield, B., Fitzpatrick, D.: Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J. Neurosci.* **17**(6), 2112–2127 (1997)
38. Fitzpatrick, D.: The functional organization of local circuits in visual cortex: insights from the study of tree shrew striate cortex. *Cereb. Cortex* **6**, 329–341 (1996)
39. Lund, J.S., Fitzpatrick, D., Humphrey, A.L.: The striate visual cortex of the tree shrew. In: Jones, E.G., Peters, A. (eds). *Cerebral Cortex*, pp. 157–205. Plenum, New York (1985)
40. Frégnac, Y., Baudot, P., Chavane, F., Lorenceau, J., Marre, O., Monier, C., Pananceau, M., Carelli, P., Sadoc, G.: Multiscale functional imaging in V1 and cortical correlates of apparent motion. In: Masson, G.S., Ilg, U.J. (eds). *Dynamics of Visual Motion Processing. Neuronal, Behavioral, and Computational Approaches*, pp. 73–94. Springer, Berlin, New York (2010)
41. Eysel, U.: Turning a corner in vision research. *Nature* **399**, 641–644 (1999)
42. Bosking, W.H., Crowley, J.C., Fitzpatrick, D.: Spatial coding of position and orientation in primary visual cortex. *Nat. Neurosci.* **5**(9), 874–882 (2002)
43. McLoughlin, N., Schiessl, I.: Orientation selectivity in the common marmoset (*Callithrix jacchus*): The periodicity of orientation columns in V1 and V2. *NeuroImage* **31**, 76–85 (2006)
44. Niebur, E., Wörgötter, F.: Design principles of columnar organization in visual cortex. *Neural Comput.* **6**, 602–614 (1994)
45. Xu, X., Bosking, W.H., White, L.E., Fitzpatrick, D., Casagrande, V.A.: Functional organization of visual cortex in the prosimian bush baby revealed by optical imaging of intrinsic signals. *J. Neurophysiol.* **94**, 2748–2762 (2005)
46. Yu, H., Farley, B.J., Jin, D.Z., Sur, M.: The coordinated mapping of visual space and response features in visual cortex. *Neuron* **47**, 267–280 (2005)
47. Basole, A., White, L.E., Fitzpatrick, D.: Mapping multiple features in the population response of visual cortex. *Lett. Nat.* **423**(26), 986–990 (2003)
48. Baxter, W.T., Dow, B.M.: Horizontal organization of orientation-sensitive cells in primate visual cortex. *Biol. Cybern.* **61**, 171–182 (1989)
49. Goodhill, G.J.: Contributions of theoretical modeling to the understanding of neural map development. *Neuron* **56**, 301–311 (2007)
50. Nöckel, J.: Gradient field plots in Mathematica. <http://pages.uoregon.edu/noeckel/computernotes/Mathematica/fieldPlots.html>
51. Chapman, B., Stryker, M.P., Bonhöffer, T.: Development of orientation preference maps in ferret primary visual cortex. *J. Neurosci.* **16**(20), 6443–6453 (1996)
52. Koulakov, A.A., Chklovskii, D.B.: Orientation preference patterns in mammalian visual cortex: a wire length minimization approach. *Neuron* **29**, 519–527 (2001)
53. Azencott, R.: Simulated Annealing, Séminaire Bourbaki 697, Astérisque 161–162, Paris (1988)
54. Gregor, K., Szlam, A., LeCun, Y.: Structured sparse coding via lateral inhibition. <http://books.nips.cc/nips24.html>
55. Hyvärinen, A., Hoyer, P.O.: A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Res.* **41**(18), 2413–2423 (2001)
56. Shmuel, A., Grinvald, A.: Coexistence of linear zones and pinwheels within orientation maps in cat visual cortex. *Proc. Natl. Acad. Sci.* **97**(10), 5568–5573 (2000)
57. Ohki, K., Matsuda, Y., Ajima, A., Kim, D.-S., Tanaka, S.: Arrangement of orientation pinwheel centers around area 17/18 transition zone in cat visual cortex. *Cereb. Cortex* **10**, 593–601 (2000)

58. Romagnoni, A., Ribot, J., Bennequin, D., Touboul, J.: Parsimony, exhaustivity and balanced detection in neocortex. *PLoS Comput. Biol.* **11**(11), e1004623
59. Petitot, J.: *Eléments de théorie des singularités*. http://jean.petitot.pagesperso-orange.fr/ArticlesPDF_new/Petitot_Sing.pdf (1982)
60. Berry, M.V.: Much ado about nothing: optical dislocation lines (phase singularities, zeros, vortices ...). In: Soskin, M.S. (ed.) *Proceedings of International Conference on Singular Optics*, vol. 3487, pp. 1–5. SPIE (Society of Photographic Instrumentation Engineers) (1998)
61. Berry, M.V., Dennis, M.R.: Reconnections of wave vortex lines. *Eur. J. Physics* **33**, 723–731 (2012)
62. Bennequin, D.: Remarks on invariance in the primary visual systems of mammals. In: Citti, G., Sarti, A. (eds) *Neuromathematics of Vision*, pp. 243–333. Springer, Berlin (2014)
63. Afgoustidis, A.: Monochromaticity of orientation maps in V1 implies minimum variance for hypercolumn size. *J. Math. Neurosci.* **5**, 10 (2015)
64. Afgoustidis, A.: *Représentations de groupes de Lie et fonctionnement géométrique du cerveau*, Thèse (D. Bennequin dir.), Université Paris VII (2016)
65. Citti, G., Sarti, A.: From functional architectures to percepts: a neuromathematical approach. In: Citti, G., Sarti, A. (eds) *Neuromathematics of Vision*, pp. 131–171. Springer (2014)
66. Petitot, J.: *Introduction aux phénomènes critiques*. http://jean.petitot.pagesperso-orange.fr/ArticlesPDF_new/Petitot_CritPh.pdf (1982)
67. Petitot, J.: Landmarks for neurogeometry. In: Citti, G., Sarti, A. (eds) *Neuromathematics of Vision*, pp. 1–85. Springer, Berlin (2014)
68. Dennis, M.R.: *Topological Singularities in Wave Fields*, PhD Thesis, H.H. Wills Laboratory, University of Bristol (2001)
69. Berry, M.V.: Optical currents. *J. Opt. A: Pure Appl. Opt.* **11**(9), 094001 (2009)
70. Berry, M.V., Dennis, M.R.: Phase singularities in isotropic random waves. *Proc. Roy. Soc. London A* **456**, 2059–2079 (2000)
71. Berry, M.V., Dennis, M.R.: Topological events on wave dislocation lines: birth and death of loops, and reconnection. *J. Physics A: Math. Theor.* **40**, 65–74 (2007)
72. Wolf, F., Geisel, T.: Spontaneous pinwheel annihilation during visual development. *Nature* **395**, 73–78 (1998)
73. Wolf, F., Geisel, T.: Universality in visual cortical pattern formation. In: Petitot, J., Lorenceau, J. (eds) *Neurogeometry and Visual Perception*, pp. 253–264. *J. Physiol. Paris* **97**, 2–3 (2003)
74. Barbieri, D., Citti, G., Sanguinetti, G., Sarti, A.: An uncertainty principle underlying the functional architecture of V1. *J. Physiol. Paris* **106**(5–6), 183–193 (2012)
75. Abrahamsen, P.: *A Review of Gaussian Random Fields and Correlation Functions*. http://publications.nr.no/917_Rapport.pdf (1997)
76. Azañs, J.-M., León, J.R., Wschebor, M.: Rice formulae and Gaussian waves. *Bernoulli* **17**(1), 170–193 (2011)
77. Adler, R.J., Taylor, J.E.: *Random Fields and Geometry*. Springer, Berlin (2007)
78. Swindale, N.V.: The development of topography in the visual cortex: a review of models. *Netw. Comput. Neural Syst.* **7**(2), 161–247 (1996)
79. Bednar, J.A.: Hebbian learning of the statistical and geometrical structure of visual inputs. In: Citti, G., Sarti, A. (eds) *Neuromathematics of Vision*, pp. 335–366. Springer, Berlin (2014)
80. Miikkulainen, R., Bednar, J.A., Choe, Y., Sirosh, J.: *Computational Maps in the Visual Cortex*. Springer, New York (2005)
81. Wolf, F.: Symmetry, multistability, and long-range interaction in brain development. *Phys. Rev. Lett.* **95**, 208701, 1–4 (2005). Erratum: 103, 20 (2009)
82. Kaschube, M., Schnabel, M., Wolf, F.: Self-organization and the selection of pinwheel density in visual cortical development. [arXiv:0801.3651v1](https://arxiv.org/abs/0801.3651v1) (2008)
83. Lee, H.Y., Yahyanejad, M., Kardar, M.: Symmetry considerations and development of pinwheels in visual maps. [arXiv:cond-mat/0312539v1](https://arxiv.org/abs/cond-mat/0312539v1)
84. Swift, J.B., Hohenberg, P.C.: Hydrodynamic fluctuations at the convective instability. *Phys. Rev. A* **15**, 319–328 (1977)

85. Schnabel, M., Kaschube, M., Wolf, F.: Pinwheel stability, pattern selection and the geometry of visual space. [arXiv:0801.3832v2](https://arxiv.org/abs/0801.3832v2) (2008)
86. Maldonado, P.E., Gödecke, I., Gray, C.M., Bonhöffer, T.: Orientation selectivity in pinwheel centers in cat striate cortex. *Science* **276**, 1551–1555 (1997)
87. Polimeni, J.R., Granquist-Fraser, D., Wood, R.J., Schwartz, E.L.: Physical limits to spatial resolution of optical recording: clarifying the spatial structure of cortical hypercolumns. *Proc. Natl. Acad. Sci.* **102**(11), 4158–4163 (2005)
88. McLaughlin, D., Shapley, R., Shelley, M., Wielaard, D.J.: A neuronal network model of macaque primary visual cortex (V1): orientation selectivity and dynamics in the input layer 4C α . *Proc. Natl. Acad. Sci.* **97**(14), 8087–8092 (2000)
89. Shelley, M., McLaughlin, D.: Coarse-grained reduction and analysis of a network model of cortical response: I. drifting grating stimuli. *J. Comput. Neurosci.* **12**, 97–122 (2002)
90. Schummers, J., Mariño, J., Sur, M.: Synaptic integration by V1 neurons depends on location within the orientation map. *Neuron* **36**, 969–978 (2002)
91. Mariño, J., Schummers, J., Lyon, D.C., Schwabe, L., Beck, O., Wiesing, P., Obermayer, K., Sur, M.: Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat. Neurosci.* **8**(2), 194–201 (2005)
92. Ohki, K., Chung, S., Kara, P., Hübener, M., Bonhöffer, T., Reid, R.C.: Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* **442**, 925–928 (2006)
93. Petitot, J.: Rappels sur l'Analyse non standard. In: Salanskis, J.-M., Barreau, H. (eds). *La Mathématique non standard*, pp. 187–209. Éditions du CNRS, Paris (1989)
94. Deligne, P., Malgrange, B., Ramis, J.-P.: Singularités irrégulières. Correspondance et documents, Société Mathématique de France (2007)
95. Ramis, J.-P.: Les derniers travaux de Jean Martinet. *Annales de l'Institut Fourier* **42**(1–2), 15–47 (1992)
96. Martinet, J.: Introduction à la théorie de Cauchy sauvage (unpublished). See Ramis above
97. Das, A., Gilbert, C.D.: Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* **387**, 594–598 (1997)
98. Ohki, K., Chung, S., Ch'ng, Y.H., Kara, P., Reid, R.C.: Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* **433**, 597–603 (2005)
99. Bressloff, P., Cowan, J.: The functional geometry of local and horizontal connections in a model of V1. In: Petitot, J., Lorenceau, J. (eds). *Neurogeometry and Visual Perception*, pp. 221–236. *J. Physiol. Paris* **97**, 2–3 (2003)
100. Hübener, M., Shoham, D., Grinvald, A., Bonhöffer, T.: Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.* **17**, 9270–9284 (1997)
101. Issa, N.P., Trepel, C., Stryker, M.P.: Spatial frequency maps in cat visual cortex. *J. Neurosci.* **22**, 8504–8514 (2000)
102. Sirovich, L., Uglesich, R.: The organization of orientation and spatial frequency in primary visual cortex. *Proc. Natl. Acad. Sci.* **101**(48), 16941–16946 (2004)
103. Born, R.T., Tootell, R.B.: Spatial frequency tuning of single units in macaque supragranular striate cortex. *Proc. Natl. Acad. Sci.* **88**, 7066–7070 (1990)
104. Issa, N.P., Rosenberg, A., Husson, T.R.: Models and measurements of functional maps in V1. *J. Neurophysiol.* **99**, 2745–2754 (2008)
105. Zhu, W., Xing, D., Shelley, S., Shapley, R.: Correlation between spatial frequency and orientation selectivity in V1 cortex: implications of a network model. *Vision Res.* **50**, 2261–2273 (2010)
106. Ribot, J., Romagnoni, A., Milleret, C., Bennequin, D., Touboul, J.: Pinwheel-dipole configuration in cat early visual cortex. *NeuroImage* **128**, 63–73 (2016)
107. Ribot, J., Aushana, Y., Bui-Quoc, E., Milleret, C.: Organization and origin of spatial frequency maps in cat visual cortex. *J. Neurosci.* **33**, 13326–13343 (2013)
108. Tani, T., Ribot, J., O'Hashi, K., Tanaka, S.: Parallel development of orientation maps and spatial frequency selectivity in cat visual cortex. *Eur. J. Neurosci.* **35**, 44–55 (2012)
109. Xu, X., Bosking, W.H., Sáty, G., Stefansic, J., Shima, D., Casagrande, V.: Functional organization of visual cortex in the owl monkey. *J. Neurosci.* **24**(28), 6237–6247 (2004)

110. Liu, G.B., Pettigrew, J.D.: Orientation mosaic in barn owl's visual Wulst revealed by optical imaging: comparison with cat and monkey striate and extra-striate areas. *Brain Res.* **961**, 153–158 (2003)
111. Obermayer, K., Blasdel, G.G.: Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci.* **13**, 4114–4129 (1993)
112. Carreira-Perpiñán, M.A., Goodhill, G.J.: Influence of lateral connections on the structure of cortical maps. *J. Neurophysiol.* **92**, 2947–2959 (2004)
113. Bednar, J.A., Miikkulainen, R.: Joint maps for orientation, eye, and direction preference in a self-organizing model of V1. *Neurocomputing* **69**, 1272–1276 (2006)
114. Miikkulainen, R., Bednar, J.A., Choe, Y., Sirosh, J.: A Self-organizing neural network model of the primary visual cortex. In: *Proceedings of the Fifth International Conference on Neural Information Processing (ICONIP'98)* (1998)
115. Swindale, N.V.: How many maps are there in visual cortex? *Cereb. Cortex* **7**, 633–643 (2000)
116. Swindale, N.V.: How different feature spaces may be represented in cortical maps? *Network* **15**, 217–242 (2004)
117. Kara, P., Boyd, J.D.: A micro-architecture for binocular disparity and ocular dominance in visual cortex. *Nat. Lett.* **458**, 627–631 (2009)
118. Thom, R.: *Stabilité structurelle et Morphogenèse*. Benjamin, New York, Ediscience, Paris (1972)
119. Thom, R.: *Modèles mathématiques de la morphogenèse*. Bourgois, Paris (1980)
120. Blake, R., Wilson, H.: Binocular vision. *Vision Res.* **51**, 754–770 (2011)
121. Kim, Y.-J., Grabowecy, M., Suzuki, S.: Stochastic resonance in binocular rivalry. *Vision Res.* **46**, 392–406 (2006)
122. De Jong, T.M.: *The Dynamics of Visual Rivalry*. University of Utrecht, Thesis (2008)
123. Leopold, D.A., Logothetis, N.K.: Multistable phenomena: changing views in perception. *Trends Cogn. Sci.* **3**(7), 254–264 (1999)
124. Lu, H.D., Roe, A.W.: Functional organization of color domains in V1 and V2 of macaque monkey revealed by optical imaging. *Cereb. Cortex* **18**(3), 516–533 (2008)
125. Rochefort, N.L.: *Functional Specificity of Callosal Connections in the Cat Visual Cortex*. PhD Thesis, Ruhr-Universität, Bochum and Université de Paris VI, Paris (2007)
126. Bosking, W.H., Kretz, R., Pucak, M.L., Fitzpatrick, D.: Functional specificity of callosal connections in tree shrew striate cortex. *J. Neurosci.* **20**(6), 2346–2359 (2000)
127. Olavarria, J.F.: Callosal connections correlate preferentially with ipsilateral cortical domains in cat areas 17 and 18, and with contralateral domains in the 17/18 transition zone. *J. Comp. Neurol.* **433**(4), 441–457 (2001)
128. Rochefort, N.L., Buzás, P., Kisvárdy, Z.F., Eysel, U.T., Milleret, C.: Layout of transcallosal activity in cat visual cortex revealed by optical imaging. *NeuroImage* **36**(3), 1804–1821 (2007)
129. Rochefort, N.L., Buzás, P., Quenech'du, N., Koza, A., Eysel, U.T., Milleret, C., Kisvárdy, Z.F.: Functional selectivity of interhemispheric connections in cat visual cortex. *Cereb. Cortex* **19**, 2451–2465 (2009)
130. Petitot, J.: *Physique du Sens*. Éditions du CNRS, Paris (1992)
131. Grossberg, S.: *Neural Networks and Natural Intelligence*. MIT Press, Cambridge, MA (1988)
132. Tani, T., Yokoi, I., Ito, M., Tanaka, S., Komatsu, H.: Functional organization of the cat visual cortex in relation to the representation of a uniform surface. *J. Neurophysiol.* **89**, 1112–1125 (2003)
133. Shapley, R., Hawken, M.J.: Color in the cortex: single- and double-opponents cells. *Vision Res.* **51**, 701–717 (2011)
134. Foster, D.H.: Color constancy. *Vision Res.* **51**, 674–700 (2011)
135. Johnson, E.N., Hawken, M.J., Shapley, R.: The orientation selectivity of color-responsive neurons in macaque V1. *J. Neurosci.* **28**, 8096–8106 (2008)
136. Taylor, D.R., Finkel, L.H., Buchsbaum, G.: Color-opponent receptive fields derived from independent component analysis of natural images. *Vision Res.* **40**, 2671–2676 (2000)
137. Cecchi, G.A., Rao, A.R., Xiao, Y., Kaplan, E.: Statistics of natural scenes and cortical color processing. *J. Vision* **10**(11), 1–13 (2010)

Chapter 5

Functional Architectures II: Horizontal Connections and Contact Structure

In this chapter, we shall investigate what we consider to be the most important geometrical property of the functional architecture of V1, namely the connectivity between the various hypercolumns which will help us to understand how V1 can *integrate* local data into global configurations. To end this first volume, devoted to neurophysiology, we shall then say a word about the genetic control of functional architectures.

5.1 From Pinwheels to Contact Geometry

5.1.1 Horizontal Intracortical Connections

So far we have been focusing mainly on the ‘vertical’ connections of the retinogeniculo-cortical pathway defining the fibre bundle $\pi : M \times P \rightarrow P$ with base space $M = R$ the 2D visual plane and fibre $P = \mathbb{P}^1$, $P = \mathbb{R}$, or $P = \mathbb{S}^1$, the 1D space of orientations. However, this structure is not sufficient for V1 to integrate contours. Indeed, for this, there must be a way to *compare* orientations (and hence fibres P_a and P_b) above *different* points a and b of R . This process is indeed neurally implemented, through what neurophysiologists call ‘horizontal’ intracortical connections, one of the great experimental discoveries of the 1980s. As we shall see below, the horizontal connections for the orientation are long range (up to 6–8 mm), excitatory, slow (around 0.2 m/s), and distributed patchily and anisotropically.




Fig. 5.1 Arborescence of an inhibitory cell of layer 2/3. The soma is shown in *white* and the axon distribution in *black*. The axon branches project anisotropically in domains of the same colour (*green–yellow*, parallelism) and highly transverse colours, including orthogonal ones (*violet–red*). The *arrow* shows a singular point and the *asterisk* a regular point. From Eysel [1]

5.1.2 *Semi-local Structures*

Before introducing the main structure, let us say a word about semi-local structures which already operate between the various hypercolumns.

5.1.2.1 **Cross-Over Detectors**

Short-range intracortical connections join neurons belonging to the same hypercolumn. They are essentially inhibitory and implement the local triviality of the fibre bundle $\pi : R \times P \rightarrow P$. There are also medium-range inhibitory connections (up to a few hundred μm) which connect neurons belonging to neighbouring hypercolumns and which code the different orientations enabling the detection of angles, corners, and T junctions. For example, Fig. 5.1, taken from the paper [1] by Ulf Eysel, shows the arborescence of an inhibitory cell of layer 2/3. We observe that the axon projects in an extremely anisotropic way in neighbouring domains of the same colour (parallelism) and of colours coding highly transverse orientations, including orthogonal ones.

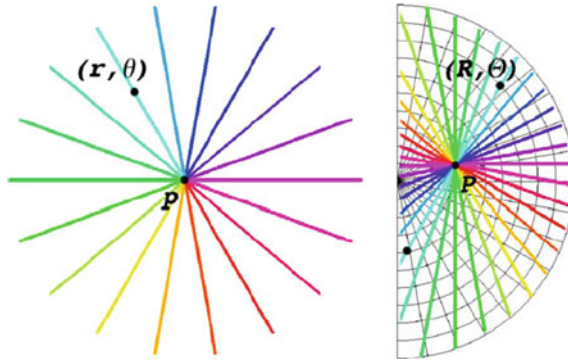


Fig. 5.2 The local–global map proposed by David Alexander and co-workers. *Left* A column C of polar coordinates (r, θ) in a pinwheel of centre p . The column C receives inputs from the point (R, Θ) of the visual field line issuing at angle Θ from the point P which projects retinotopically onto p . This point is thus represented at (r, θ) on the map. Note that the orientations (light blue) of the points (R, Θ) and $(R, \Theta + \pi)$ are represented by the same (r, θ) and that we do indeed have $\Theta = \alpha \pm \theta/2$, as in Sect. 4.4.6. From Alexander et al. [2]

5.1.2.2 Alexander’s Local–Global Map

Concerning the transition from local to global, we would also like to mention the work by David Alexander, Paul Bourke, Otto Konstandatos, Phil Sheridan, and James Wright in [2]. These authors developed the idea that the pinwheels are local maps of a few hundred μm arranged visuotopically, maps centred on singular points, with the saddle points (see Sect. 4.4.1) serving as transitions between them. They explain that these maps encode the geometric feature of orientation, corresponding to the fibre bundle structure $\pi : V = R \times P \rightarrow R$ which we introduced for the contact elements $(a, p) = (\text{position}, \text{orientation})$. They also stress that the fibre P is not the circle \mathbb{S}^1 , but the projective line \mathbb{P}^1 (spin structure) and that there is therefore:

[...] a doubling of the visuotopic angles in the pinwheel. [2, p. 860]


The great originality of their work lies in the way they treat p as a *non-retinotopic feature*, generalizing the bundle π and identifying each pinwheel with a domain D , in fact an open disc \mathbb{D} , which is:

[...] a remapping of V ’s global retinotopic map. [2, p. 860]

This leads to their conjecture of a local–global map which we may interpret in our geometric language as the fibre bundle $\pi_A : R \times D \rightarrow R$, where the long-range, anisotropic, and patchy horizontal connections of the supragranular layers serve to identify the different fibres D_a and D_b . Figure 5.2 illustrates the model.

The model essentially involves describing the tangent bundle TR in polar coordinates and quotienting modulo π , thus keeping 2D fibres. Our model, on the other hand, consists in projectivizing TR , which reduces the fibres to one dimension, and introducing a further structure in order to go from local to global.

Fig. 5.3 Short-range horizontal connections of two neurons in the upper layers of V1. These are distributed isotropically. From [5]. Copyright (1993) National Academy of Sciences, USA



5.1.3 Parallelism and Coaxiality

As discussed above, it is essential to be able to *compare* orientations, and hence fibres P_a and P_b , above different points a and b of R . This is a process known as *parallel transport*, and it is implemented neurally by long-range horizontal intracortical connections. The fundamental property of these connections is indeed to *link cells that have more or less the same preferred orientation in distant hypercolumns*. To detect them, we can measure the correlations between cells belonging to different hypercolumns: we compare the orientations of the cells encountered in a cortical penetration with the orientation of a single reference cell, and by establishing cross-correlograms, we then observe that cells with similar orientations are strongly correlated, i.e. there is a peak in the correlogram, while cells with sufficiently different orientations are decorrelated (see, e.g. Ts'o et al. [3]).

We can also use *in vivo* optical imaging methods (see e.g. [4]). From the beginning of the 1990s, it thus became possible to visualize horizontal connections by labelling them with anterograde markers like biocytin. In this way, long-range horizontal connections of 2–5 mm were identified in cats and macaques, joining hypercolumns with the same ocular dominance. They are very different from short-range connections (up to 400 μm) within hypercolumns, which are isotropic and homogeneous and not connected to specific functional domains. Figure 5.3, taken from the paper [5] by Rafael Malach et al., shows the short-range horizontal connections of two neurons in the upper layers of V1. We see that they are distributed very isotropically and cross several isochromatic regions: therefore they do not preserve orientation.

On the other hand, Fig. 5.4, due to Bosking et al. [6], shows how biocytin injected locally in a region of around 100 μm of layer 2/3 of the V1 area of a tupaia (tree shrew), a given orientation region coded by blue–green, diffuses along the horizontal connections in an anisotropic, selective, and patchy way. As in the previous figure, short-range diffusion is isotropic and occurs through inhibitory intrahypercolumn connections. But long-range diffusion is in fact highly anisotropic, occurring through




Fig. 5.4 Diffusion of the biocytin marker along horizontal connections of layer 2/3 in the V1 area of a tupaia (tree shrew). The distribution is anisotropic and patchy, concentrated locally in domains of the same orientation as the injection site *top left* and mainly along the *top-left to bottom-right diagonal*. From Bosking et al. [6]

excitatory interhypercolumn connections, and it is restricted to domains with roughly the *same* orientation (of the same colour) as the injection site. Previous analogous results, using horseradish peroxidase diffusion, were already obtained in the early 1980s, before the invention of imaging methods, by Mitchison and Crick [7] (see also Rockland and Lund [8, 9]).

We have already encountered the property of parallel transport in Sect. 4.7.1.3 of Chap. 4 with the work of Jorge Mariño and James Schummers. It can be measured by statistical analysis of the differences in orientation of the marked synaptic buttons and those of the injection site: the curve is sharply peaked around 0. Note that the marked synaptic buttons also cluster along a top-left to bottom-right diagonal. We shall give an interpretation of this key fact.

When these horizontal intracortical connections were discovered, some thought they might violate retinotopy, whereas in fact they strengthen it, by guaranteeing its large-scale coherence (see Grinvald et al. [10]). Without them, neighbouring hypercolumns would remain independent, and the retinotopy would lose any *immanent* reality for the system.

The fact that horizontal intracortical connections link neurons with similar orientations in different hypercolumns indicates that the neural system can find out whether, for $b \neq a$, an orientation p at a is roughly the same as an orientation q at b . In other words, while the ‘vertical’ retino-geniculo-cortical connections confer an immanent internal sense on the relations between pairs (a, p) and (a, q) (different orientations p and q at the same point a), the ‘horizontal’ intracortical connections




Fig. 5.5 Horizontal intracortical connections allow the system to compare in an immanent way the orientations in two different hypercolumns corresponding to two different retinal locations a and b (highly schematic)




Fig. 5.6 While the ‘vertical’ retino-geniculo-cortical connections endow the relations between pairs (a, p) and (a, q) (different orientations p and q at the same point a) with an immanent internal sense, the ‘horizontal’ intracortical connections do so for the relations between (a, p) and (b, p) (same orientation p at different points a and b)

do so for the relations between pairs (a, p) and (b, p) (same orientation p at different points a and b) (see Figs. 5.5 and 5.6).

What is more, and this is absolutely crucial for the geometry of percepts, it can be shown that the intracortical connections preferentially link not only parallel contact elements (a, p) and (b, p) but in particular *coaxial* contact elements, i.e. pairs such that the common orientation p is the orientation of the axis ab (see Figs. 5.7 and 5.8). This can be seen intuitively by noting that the diagonal where marked synaptic buttons cluster is roughly parallel to the $V1-V2$ boundary representing the vertical meridian of the visual space. We may deduce that this overall direction is the quasi-




Fig. 5.7 Intracortical connections preferentially link neurons detecting pairs (a, p) and (b, p) which not only have the same orientation, but are also coaxial, where p is the orientation of the axis ab




Fig. 5.8 In Fig. 5.4, the clustering of synaptic buttons connected to the region where the marker was injected occurs along a global diagonal. The preferred orientation p of the injection site is 80° (top-right corner). The direction of the global elongation of the clusters (long grey strip) rotates with p . We know that the direction of the vertical meridian of the visual field (90°) is parallel to the $V1-V2$ boundary. We see that it has rotated a little. Precise measurements in Bosking et al. [6] show that this corresponds to the phenomenon of coaxiality

vertical direction marked in the upper right corner of Fig. 5.4. Now this direction is encoded by a colour which is precisely that of the marked regions.

To sum up, as explained by Bosking [6]:

The system of long-range horizontal connections can be summarized as preferentially linking neurons with co-oriented, co-axially aligned receptive fields.

As we shall see at some length and with considerable mathematical detail in the second volume, it is basically this functional architecture that models the geometric notion of contact structure. However, let us first see how it is confirmed by psycho-physical observations.

5.2 Integration of Contours and Association Field

The intracortical connections of V1 play a key role in the explanation of the first intriguing perceptual phenomenon mentioned in the *Introduction*, namely the integration of contours. ‘What is a line on the neural level?’ From a psychophysical point of view, one of the great breakthroughs was made by David Field, Anthony Hayes, and Robert Hess who, in their classic paper of 1993 [11], introduced the idea of *association field*.

5.2.1 Some Experimental Facts

The basic idea is simple. If we consider a randomly distributed ensemble of short line segments, our perception will derive no particular structure (see Fig. 5.9). But colinearity of segments enhances the response and improves visual sensitivity (see Kapadia et al. [12, 13]).

The experimental procedure of Field, Hayes, and Hess involves briefly presenting subjects with a grid of 256 oriented elements (a, p), e.g. they use Gabor patches, that is, oriented bandpass elements in the spatial domain which, by selecting a spatial frequency, single out only those cells corresponding to a well-defined scale. The size of the receptive field of a cell that selects for the wavelength λ is typically 2λ .

Half the time the grid of Gabor patches contains elements aligned along a smooth path γ , the other elements being oriented randomly (see Fig. 5.10). The rest of the time, all the elements are oriented randomly. The subject’s task is to decide whether he/she detects the alignment γ in the grid (two-alternative forced choice method). The results show that subjects do indeed perceive the alignment if the elements (a, p) are aligned ‘tangentially’ to γ and if the change in the slope between two consecutive elements is not too great ($\leq 30^\circ$). This perceptual salience results from a characteristic *pop-out* phenomenon.

Fig. 5.9 In a randomly distributed ensemble of short line segments, we perceive no particular structure





Fig. 5.10 The path (*top-left corner*) is embedded in a background of random distractors. In this ensemble of Gabor patches (a, p), the subject observes the pop-out of a group of elements (a_i, p_i) if these are suitably aligned. From Field et al. [11]



A key point is that the elements of the grid are too far apart to belong to a single receptive field in V1 and that:

It is clear that this ‘association field’ covers a considerably wider area than would be covered by the receptive field of a mammalian cortical cell [11, p. 185].

Now, subjects spontaneously group these elements together. There must therefore be an automatic mechanism that makes a connection between several receptive fields. This is *low-level integration*. As the authors explain:

Recent computational studies have suggested that a useful segregation process for real scenes may be based on local (rather than global) integration. [...] The general theme of these algorithms is that the points along the length of a curved edge can be linked together according to a set of local rules that allow the edge to be seen as a whole, even though different components of the edge are detected by independent mechanisms [11, p. 174].

So here is an excellent formulation of the problem of how to make the transition from local to global: there are *local neurophysiological binding rules* which bring about a *global perceptual organization*. The fact that the grouping phenomenon is local and not global in the sense of being associated with large receptive fields turns out to be crucial:


In our stimuli, there does not exist any ‘global’ feature that allows the path to be segregated from the background. It is not possible to segregate the path by filtering along any particular dimension. Our results imply that the path segregation is based on local processes which group features locally [11, p. 191].

If the grouping is purely local, there must necessarily be an *integration mechanism* that is not a straightforward filtering: as the feature detectors are point processors (on a certain scale), how can their measurements be globalized?

The experiments also show two other fundamental phenomena:

- If the change of slope between consecutive elements is too great, the subject will no longer perceive any alignment (see Fig. 5.11).

Fig. 5.11 If the change in the slope between consecutive elements (a_i, p_i) is too great, the subject no longer perceives any alignment. From Field et al. [11]




- Likewise, if the orientation of the elements is not tangent but transverse, e.g. orthogonal, to the curve γ followed by their centres (see Fig. 5.12).

There is then no pop-out effect, and the paths are only detected in a cognitive and inferential way.

These experiments relate to the Gestalt principle of *good continuation*. And the more Gestalt principles are involved, the stronger the pop-out effect. For example, Fig. 5.13 due to Kovács and Julesz [14] shows that if the alignment of the segments closes in accordance with the Gestalt law of closure, then pop-out increases significantly.

Fig. 5.12 If the orientation p_i of the elements (a_i, p_i) is not tangent but transverse, e.g. orthogonal, to the 'curve' γ followed by their centres a_i , then the subject no longer perceives the alignment. From Field et al. [11]






Fig. 5.13 Adding the Gestalt law of closure to the Gestalt principle of good continuation strengthens perceptual pop-out. From Kovács and Julesz [14]. Copyright (1993) National Academy of Sciences, USA

5.2.2 *Pop-Out, Perceptual Salience, and the Helmholtz Principle*

The psychophysical experiments of Field, Hayes, and Hess are based on the phenomenon of perceptual salience known as pop-out. We shall give a neurophysiological interpretation in the next section. But this can also be characterized in probabilistic terms expressing the novelty and rarity of certain stimuli.

For example, Sajda and Han [15] introduced the following idea. We consider a pinwheel map S_k with N possible orientations $\theta = \theta_i, i = 1, \dots, N$; that is, we decompose each pinwheel into N angular sectors each corresponding to a filter. The response space of a pinwheel is \mathbb{R}^N , where a response is a vector $f(\theta)$. To calculate $f(\theta)$, the signal is convolved with the filter associated with θ and we take into account the functional architecture connecting sectors of the same orientation in different pinwheels.

We thus consider a distribution P of orientations corresponding to a particular class of stimuli (e.g. textures made up of segments), and we define the salience of the pinwheel S_k by $S_k = -\log(p(f_k | P))$, where $p(f_k | P)$ is the conditional probability of $f_k(\theta)$ relative to the distribution P . This formula tells us that the rarer the function $f_k(\theta)$ in the given class of stimuli, i.e. the smaller the conditional probability, the more salient it becomes.¹ In the response space \mathbb{R}^N of a pinwheel,

¹Recall that $\log(0) = -\infty$, so if $p \in [0, 1]$ is small, then $-\log(p)$ is a large positive number.

P defines clusters where the responses to the given class of stimuli accumulate, so saying that a response is salient amounts to saying that it is far from these clusters.

The idea that ‘salient’ is equivalent to ‘rare’ is sometimes called the *Helmholtz principle*. It was applied by Jean-Michel Morel to the detection of alignments in images, which is equivalent to the problem addressed by Field, Hayes, and Hess. The basic idea is that a large deviation from a statistically generic situation will lead to perceptual salience. The less probable a configuration, the more salient it will be. The advantage with Morel’s approach is that it can define ‘maximal meaningful events’ without using priors as is usually done in the conventional Bayesian approach. In a standard Bayesian model, we take an image I , introduce a class of models satisfying prior conditions, and try to infer abductively the best model M for I that satisfies these constraints. According to Bayes’ law, the conditional probability $P(M | I)$ of the abduction leading to M from the data I is given by

$$P(M | I) = \frac{P(I | M)P(M)}{P(I)} .$$

The conditional probability $P(I | M)$ of I when M is given corresponds to the *direct problem* of obtaining the image I from the model M , and it is usually straightforward to calculate. However, the conditional probability $P(M | I)$ of the model M when the image I is given corresponds to the *inverse problem*, and it is generally very hard to calculate; it requires us to know the priors $P(M)$. The good model M that we seek under the hypothesis I corresponds to its maximization.

If for example M is given by a Gibbs distribution $P(M) = \exp(V(M)) / Z$, where Z is a partition function, as in statistical physics, and $V(M)$ is an ‘energy’, and if $P(M | I)$ is given by a Gaussian $C \exp(-D(M, I))$, where D is a ‘distance’ between M and I , then the optimal M is given by minimizing the functional $D + V$.

Bayesian models are classic top-down techniques of template matching. They are not well suited to a probabilistic interpretation of salience, where no priors are known and where we are not trying to find a model giving the best possible interpretation of I , but only to extract the salient features of I . In this case, it is better to use a probabilistic version of the opposition between the generic and non-generic situations. This is what Jean-Michel Morel did with Agnès Desolneux and Lionel Moisan for the example of alignment in [16]. His basic assumption was as follows:

The main idea is that a meaningful event is an event that, according to probabilistic estimates, should not happen in the image and therefore is significant.

Here, ‘meaningful’ means ‘not generic’. Consider, for example, a black square of 10×10 pixels on a white background in a frame of 100×100 pixels. If the colour of the pixels is chosen at random, the probability of such a configuration is

$$90 \times 90 \times \left(\frac{1}{2}\right)^{100} \times \left(\frac{1}{2}\right)^{900} ,$$

a probability of the order of 10^{-300} which is zero for all practical purposes.²

For the alignment property, we take a grid G of $N \times N$ sites and at each site $x = (i, j)$ we consider a small segment centred on the point and with a direction d_x which is a random variable with uniform probability distribution $p = 1/n$, i.e. we consider n equiprobable directions that can join the points together. Let $S = (x_1, \dots, x_r)$ be a segment with r points, D the direction of S , and d_i the directions of the points x_i of S . We calculate the maximal number $k(r)$ of aligned points which must be observed for the alignment event to be perceptually salient. Let X_i be the Boolean³ random variable with value 1 if d_i is aligned with D and 0 otherwise. We have $P(X_i = 1) = p$ and $P(X_i = 0) = 1 - p$. The random variable representing the number of points x_i of S with orientation D is thus the sum S_r of the X_i . Since the X_i are independent random variables, we calculate $P(S_r = k)$ using the binomial law and define the ε -significance by a significance threshold $k(r)$:

$$k(r) = \min \left\{ k : P(S_r \geq k) \leq \frac{\varepsilon}{N^4} \right\},$$

where N^4 is the number of oriented segments (of arbitrary length r) in the $N \times N$ grid. The calculation gives precise estimates of the threshold $k(r)$. For example, for $N = 512$ and $p = 1/16$, the minimal length of a 1-significance segment is 9. As N tends ‘rather quickly’ to ∞ , we obtain the estimate

$$k(r) \approx pr + \sqrt{2p(p-1)r \log \left(\frac{N^4}{\varepsilon} \right)}.$$

We then define the maximally salient segments S , which are those producing perceptual pop-out effects.

This kind of calculation can be generalized. The alignment of segments randomly oriented along a path (with curvature) is not generic in the space of random segment configurations, whence it is salient. This is the result obtained by Field, Hayes, and Hess. However, if we now consider such paths, the constitution of a closed path is not generic in the space of paths, so the alignment of randomly oriented segments along a closed path is in a certain sense doubly non-generic. It is therefore doubly salient. This is the result obtained by Kovács and Julesz.

² 90×90 is the number of ways to position the little square, $1/2$ is the probability of black or white for each pixel, $(1/2)^{100}$ is the probability that the 100 pixels of the small square are black, and $(1/2)^{900}$ is the probability that the 900 pixels remaining in the large square are white.

³Named after George Boole.




Fig. 5.14 An association field. The elements are pairs $(a, p) = (\text{position}, \text{orientation})$. Two elements (a_1, p_1) and (a_2, p_2) can be connected (*thick lines*) if we can interpolate a curve γ that is not too curved between positions a_1 and a_2 , and which is tangent to p_1 and p_2 at a_1 and a_2 , respectively. Otherwise the two elements are not connectable (*thin lines*)




Fig. 5.15 ‘Field lines’ of the association field

5.2.3 Explanation in Terms of Association Fields

By measuring the changes in detection rate for different spatial positions and relative orientations of the elements forming a contour, Field, Hayes, and Hess were able to conclude that the tendency to perceive the elements $c_i = (a_i, p_i)$ as being aligned followed from the existence around each element, of a region in which other elements tend to be perceived as grouped. This region, referred to as the *association field*, is defined by *joint constraints* on the position and orientation. The form of the association field is described in Figs. 5.14 and 5.15.

Field, Hayes, and Hess give a remarkable interpretation of the deep geometrical nature of the association field. To begin with, the association is not simply:

[...] a general spread of activation, linking together all types of features within the field [11, p. 185].

It manifests a correlation between position and orientation:

Elements are associated according to joint constraints of position and orientation [11, p. 187].

This is a key point. The pop-out phenomenon comes about because the elements are aligned in such a way as to be tangent to the curve described by their centres:

There is a *unique link* between the relative positions of the elements and their relative orientations. [...] *The orientation of the elements is locked to the orientation of the path*; a *smooth curve* passing through the long axis can be drawn between any two successive elements [11, p. 181].




Fig. 5.16 Co-activation of two coaxial cells (a, p) and (b, q) (with $p = q = ab$) in turn co-activates intermediate coaxial cells

The process that underlies this segregation shows a specific relation between position and orientation. Similar orientations are grouped together only when the alignment falls along particular axes [11, p. 189].

This is a discrete formulation of the integrability constraint. Contact elements such as $c_i = (a_i, p_i)$ embedded in a background of random distractors generate a perceptually salient curve if and only if the orientations p_i are tangent to a curve γ interpolating between the positions a_i . This is due to the fact that the co-activation of simple neurons with approximately coaxial preferred orientations, detecting pairs (a, p) and (b, q) such that b is roughly aligned with a in the direction p and q is close to p , spontaneously co-activates in its turn the intermediate cells through the horizontal connections (see Fig. 5.16).

When the distance between the positions a_i tends to 0, the ‘joint constraints’ and the ‘unique link’ between the positions a and the orientations p become precisely the integrability constraint $p = f'(x)$ explained in Sect. 4.3.7 of Chap. 4.

5.2.4 Confirmation by fMRI

In the above psychophysical experiments, we must refer to the responses made by subjects, and hence their lived experience of perceptual pop-out. But if we wish to obtain neurophysiological measurements, we cannot use local measurement methods, precisely because the phenomenon under investigation is a process integrating local orientations into global contours. Fortunately, the emergence of a Gestalt alignment under the good continuation constraint can be tested using fMRI methods. For example, Zoe Kourtzi et al. [17] used an adaptive paradigm to show that there is indeed a selectivity for collinear contours in the primary visual areas. The idea of the adaptive protocol is to present random configurations over a long period (150 ms), during which the response of the visual system gradually decreases, then to measure the rebound effect brought about by a sudden change leading either to a collinear configuration, or to another random configuration. As pointed out by the authors:

Stronger rebound effects for conditions with collinear than random patterns would indicate visual areas responsive to global visual configurations [17, p. 334].

Fig. 5.17 shows that this is indeed the case for the monkey V1 area.

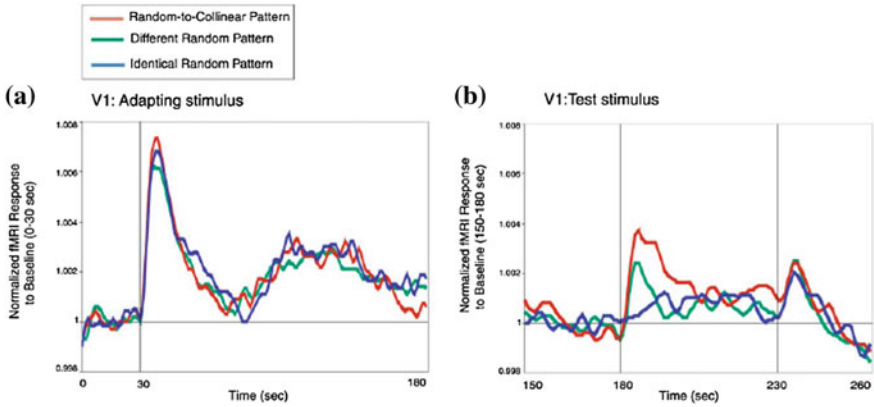


Fig. 5.17 **a** fMRI response at the onset of adaptation at time $t = 30$ ms. We see the decrease in the response over the interval [40, 180] ms. **b** Another stimulus is then presented and the elicited response is measured up to 280 ms. The rebound effect is clear, and we observe that it is much stronger for a transition to a collinear pattern. From Kourtzi et al. [17]

5.2.5 Relationship with the Horizontal Connections

In their paper, Field, Hayes, and Hess put forward several physiological speculations about the way the association field might be implemented by the horizontal connections. These were confirmed by Jean Lorenceau, Peggy Seriès, and Sébastien Georges in Yves Frégnac's laboratory [18, 19]. They measured the apparent speed of rapidly moving sequences of oriented Gabor patches (the speed-up illusion). The subject is presented with a sequence of Gabor patches which pass by at a rate of $64^\circ/s$ (5 images 17 ms apart) along a vertical straight line at a certain speed. The patches can be oriented in the direction of motion (vertical collinearity) or in the direction orthogonal to the motion (horizontal parallelism). The sequences are themselves separated by intervals of 500 ms. Figure 5.18 shows this protocol. We observe that the apparent speed is greater (overestimated) in the case of collinearity and smaller (underestimated) in the case of parallelism.

The authors conjecture that the over- and underestimation of the apparent speed is a field association effect and that the latter results from the propagation of activity along the long-range horizontal intracortical connections between neurons with similar preferred orientations:

Long-range horizontal connections in V1 elicit differential latency modulations in response to apparent motion sequences, whose read-out at an MT stage results in a perceptual speed bias [18, p. 2757].

This conjecture was confirmed by the fact that, in the case of collinearity, the increase in apparent speed measured using these *psychophysical methods* turns out to be essentially the same (around 0.2 m/s) as the speed of propagation in the horizontal

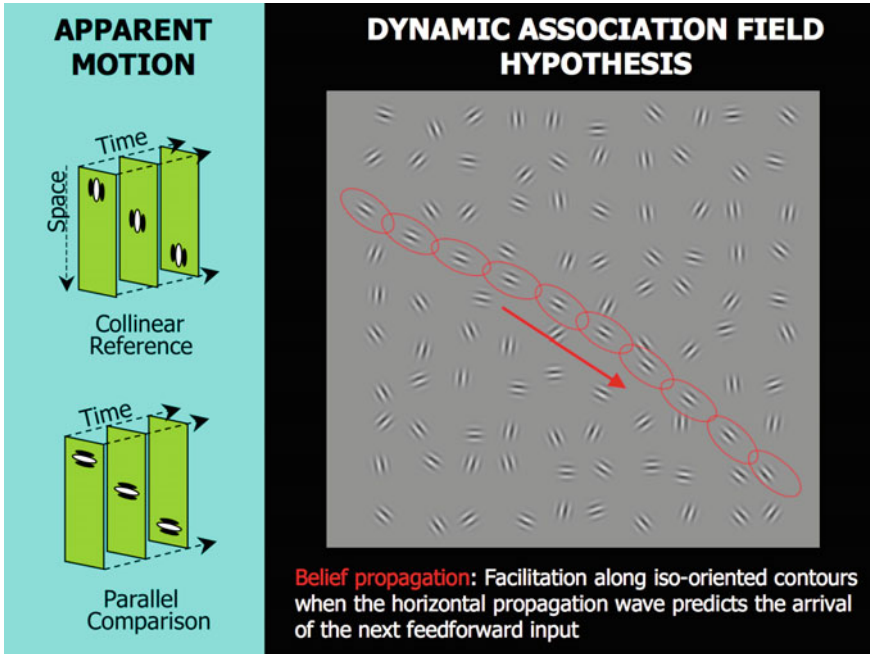


Fig. 5.18 Sequence of Gabor patches along a vertical straight line. The patches are oriented either in the direction of motion (vertical collinearity), or in the direction perpendicular to the motion (horizontal parallelism). The apparent speed is greater in the collinear case and smaller for parallelism. From [20]

intracortical connections as measured using *electrophysiological methods*. In [21], the authors explain that:

Spreading activity through [long-range horizontal] connections evoked by a first stimulus may modulate the dynamics—and in particular the latency—of the neuronal responses to a second stimulus, presented from a few milliseconds to a few tens of milliseconds later, at neighbouring positions in visual field.

David Alais, Jean Lorenceau, et al. [22] used the method known as *binocular rivalry* (already discussed in Sect. 4.10.3.2 of Chap. 4) to improve these results. The idea is to send stimuli comprising two Gabor patches (a, p) and (b, q) with $p = \overline{ab}$ and q roughly aligned with p , i.e. collinear stimuli, to one of the visual fields and random noise around a and b to the other visual field. The subject must say when the pair of contact elements appears or disappears, the hypothesis being that the structure of the time series encodes the strength of the connection between the two patches. This can be used to construct an association field that confirms those obtained by the other methods.

In another paper [23], J. Lorenceau and co-workers used magnetoencephalography (MEG) to confirm this increase in the apparent vertical speed. The imaging data

show that a wave does indeed propagate along the horizontal intracortical connections, and more precisely, that there is an alignment mechanism in the spike trains ‘which synchronizes the neural activity produced by a contour’. This alignment shortens the response latency and induces a phase advance along the contours. This would indeed appear to explain the phenomenon of pop-out and perceptual salience observed psychophysically.

5.2.6 Discretization of the Contact Structure

Our graduate student at the École Polytechnique, Yannick Tondut, showed long ago (in 1997, see [24]) that the association field can indeed be interpreted as a *discretization* of the contact structure \mathcal{K} of the contact fibre bundle V . As this structure idealizes the limit of a pinwheel arrangement for a scale tending to 0, it is natural to discretize.


We fix a spatial discretization step Δs and discretize also the orientations with a step Δp . We consider the curves Γ in $V = R \times P$ which are locally defined by the equations $y = f(x)$, $p = g(x)$. As we saw in Sect. 4.3.7 of Chap. 4, Γ lifts a curve γ of R with equation $y = f(x)$ if and only if the integrability condition $p = g(x) = f'(x) = dy/dx$ is satisfied. To discretize this equation, we consider two consecutive points⁴ A and B on γ which satisfy the metric relation $d(A, B) = \Delta s$, where d is the Euclidean distance in the visual field. If θ is the angle between the vector AB and the x axis, we have (see Fig. 5.19) $x_B - x_A = \Delta s \cos \theta$ and $y_B - y_A = \Delta s \sin \theta$. In this discretized framework, the tangent to γ at A can be approximated by the straight line AB , and the derivative $f'(x)$ by the ratio $(y_B - y_A)/(x_B - x_A) = \tan(\theta)$. The equation corresponding to the contact structure at A then takes the discrete form $|p_A - \tan(\theta)| \leq \Delta p$. Symmetrically, we also have $|p_B - \tan(\theta)| \leq \Delta p$. Adding the two inequalities, we obtain the symmetric form $|p_A - \tan(\theta)| + |p_B - \tan(\theta)| \leq 2\Delta p$.

If we represent graphically the oriented elements that can follow a given element, we obtain the desired result: two elements that come one after the other satisfy the constraints of the association field (see Fig. 5.19). The threshold $2\Delta p$ is roughly $\tan(\pi/6)$. We may thus conclude that the association field is indeed a discrete version of the integrability constraint. Note that one consequence of the discretization of the contact structure is a limitation of the curvature that the visual system can accept at a given scale.

It is straightforward to explain the third experiment of Field, Hayes, and Hess (see Fig. 5.12) in terms of the contact structure: the curve Γ in V is *not integrable*; that is, it does not satisfy the integrability condition. Figure 5.20 (top left) shows a curve in J^1R which is a Legendrian lift. Figure 5.20 (top right) corresponds to the situation in the third experiment, where we have added a constant p_0 to $p = f'(x)$ (an angle $\pi/2$, whence the orientations p become orthogonal to the curve γ described by the

⁴In the rest of this section, A and B will denote the ends of curves.

Fig. 5.19 Discretization of the integrability condition. Courtesy of Yannick Tondut



positions a). Figure 5.20 (bottom) shows two examples of non-integrable curves in V , the first (left) because p is constant, whereas f' is not, the second (right) because p varies more quickly than f' . In each of the three last cases, the pop-out of the curve γ described by the positions a is impossible because the integrability condition, i.e. the joint constraint on position and orientation, is not satisfied.

5.2.7 Binding

The pop-out of curves that are generated by approximately aligned contact elements $c_i = (a_i, p_i)$ is a typical Gestalt phenomenon resulting from *binding* between the activities of the relevant neurons. These *synchronize* their firing by means of the horizontal connections. The temporal coherence of the correlated firings binds the features they encode, and this explains why elements aligned along a contour are perceived as a *single whole*. Put another way, *the integrability constraint is a geometrical binding condition*. As stressed by Tai Sing Lee in [25], one might think that feedback from $V2$ is necessary for binding, but in agreement with the high-resolution buffer hypothesis discussed in Sect. 2.2 of the *Introduction* and the beginning of Chap. 4, it is nevertheless the underlying geometry of $V1$ that turns out to be essential.

5.2.8 Comparison with Other Data

As noted by Yannick Tondut, this model of the association field can also explain other results of Polat, Sagi, Gilbert, or Westheimer which showed that the facilitation zone induced by an oriented element (a, p) in fact implements the *contact plane* at the point corresponding to the contact fibre bundle.

Polat and Sagi [26] considered facilitation and suppression effects in target detection tasks. Using Gabor patches, they studied the detection of a low-contrast stimulus (the target) when it is flanked by two high-contrast stimuli (masks) aligned with it. By varying the distance between the target and the masks, they observed a facilitation effect for distances between 2λ and 10λ , where λ is the wavelength of the stimuli.




Fig. 5.20 Association field as integrability condition. *Top left* The integrability condition is satisfied. In the other cases it is not. *Top right* We add a constant angle to the tangent [$p = f'(x) + p_0$]. *Bottom left* p is constant, whereas f' is not. *Bottom-right* p rotates more quickly than f'

Insofar as the typical size of a receptive field tuned to a wavelength λ is 2λ , the effect extended well beyond this size. They also compared different configurations of relative positions and orientations of the stimuli: facilitation occurs in particular when the orientation of three stimuli coincides with that of their alignment.

The experiments of Gilbert et al. [27] confirmed these results. Using simple oriented bars as stimuli, these authors observed a facilitation in the detection of a low-contrast target when it is aligned with an analogous high-contrast bar. Two of their results are particularly interesting:

1. A very small deviation from collinearity that is not accompanied by an appropriate displacement of the position very quickly reduces facilitation.
2. On the other hand, facilitation persists when the orientation of the relevant bar is modified in such a way as to preserve the continuity of the path determined by the two bars, but it then decreases with the curvature of the path and vanishes beyond a deflection of 30° .

Another interesting comparison can be made with the work of Steve Zucker (Zucker et al. [28], Parent and Zucker [29]). To detect the curves, Zucker used an estimate of their geometrical coherence based on the compatibility of the neighbouring tangents. This compatibility is measured either directly from the difference in the orientations, or by evaluating the local curvature using specialized detectors. Zucker's model is rich and goes beyond the association field. However, the compatibility criterion is very similar to that of the association field.

It would also be extremely interesting to study the relationships between the association field and the cooperation/competition processes used by Grossberg and Mingolla [30]. Between the aligned elements, these authors introduced a cooperation materialized by oriented dipole cells with large receptive fields, able to detect approximate alignments. The weight of an element in the RF of a dipole cell depends on its relative position and orientation. The ensemble of weights defines a *cooperation field*, which is strikingly similar to the association field of Field, Hayes, and Hess.

5.3 Some Effects of the Horizontal Connections

5.3.1 Contextuality of the Receptive Fields

We thus discover the role of the horizontal intracortical connections and incidentally also the top-down projections from V1 to the lateral geniculate nucleus. They play an important part in making the functional architecture of the V1 RFs flexible and plastic. The RFs are in fact adaptive, their responses being modulated by the stimuli. The notion of RF must therefore be extended. In Sect. 3.2.6 of Chap. 3, we saw that the minimal discharge field (MDF) was only part of the actual RF, the part where highly localized stimuli elicit spike responses. However, the horizontal connections induce strong *contextual* effects which are important, not only for integrating contours, but for many other global phenomena identified a long time ago by the Gestalt theory, such as the perception of surfaces, segmentation, figure/background segregation. To go 'beyond the receptive field in the primary visual cortex', the reader is referred to the study [31] by David Fitzpatrick. Activity waves propagate:

[...] in a wave-like manner via polysynaptic pathways (Wright et al. [32, p. 2704]).

This explains how neurons can be driven in a purely contextual way without any stimulus affecting their RF.

Figure 5.21 due to Yves Frégnac et al. [20] shows schematically the way the MDF extends, thanks to the propagation of cortical waves along the horizontal connections, to a synaptic integration field (SIF).

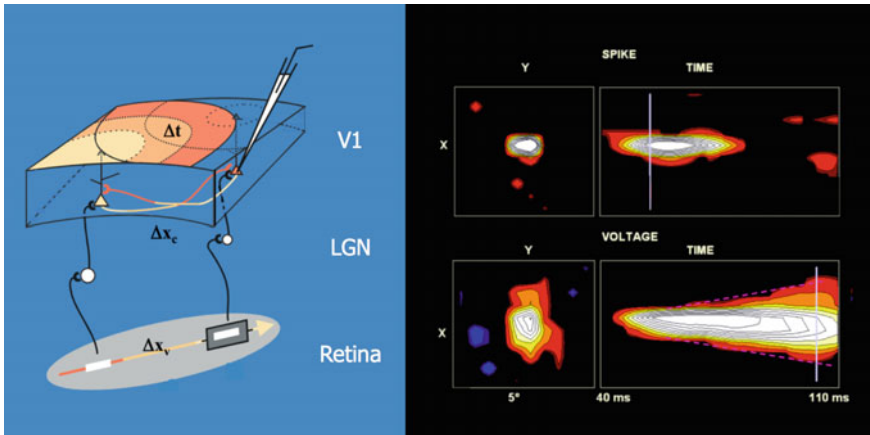


Fig. 5.21 *Left* Δx_v is the eccentricity of the distal retinal stimulus (white segment on the left) relative to the stimulus located in the MDF (white segment on the right in the grey rectangle). Δx_c is the cortical distance between the neurons of V1 activated by the two stimuli. Δx_t is the latency time induced by propagation along the horizontal connection. *Right* At the top we see the minimal discharge field (MDF) defined by spike responses in the (x, y) plane and its temporal evolution in the (x, t) space. At the bottom, we see the synaptic integration field (SIF) defined by the sub-threshold activity. The SIF is much more extensive than the MDF. From Frégnac et al. [20]

5.3.2 Line-Motion Illusion

As we shall see, the horizontal connections can considerably deepen our understanding of perceptual phenomena. We shall return to this at length. But for now, we shall consider the simple example provided by Okihide Hikosaka's so-called line-motion illusion (LMI). This is the perception of an apparent motion that is not present in the stimuli, which are in fact motionless. The experimental protocol is straightforward: a little square of side L is flashed to the subject, followed a few milliseconds later by the flash of a long rectangle of width L which extends the square. The subject then perceives the static rectangle as a dynamic motion of the small square. This apparent motion, which depends on V1, and possibly also on V2 and MT, is only illusory relative to the stimulus. Indeed, the experiment shows that the dynamics of the cortical activity in V1 and V2 is *the same* as that occurring for a real motion of the stimulus.

Figure 5.22 due to Jancke et al. [33] concerns the primary visual areas 17 and 18 of an anaesthetized cat and shows the temporal evolution (time on the horizontal axis) of the membrane potentials using voltage-sensitive dye imaging (VSD). In this technique, fluorescence indicates local depolarization of neurons and hence marks active cells (see Sect. 4.4.1 of Chap. 4). This real-time imaging method is highly sensitive to sub-threshold activity of neural membranes. The authors used four stimuli: (c) a flashed motionless square, (d) a flashed motionless bar, (e) the moving square, (f) the motionless square and bar flashed successively. The images show the sub-threshold




Fig. 5.22 Cortical dynamics of the LMI in an anesthetized cat, measured using the VSD technique. There are four stimuli: (c) a flashed motionless square, (d) a flashed long motionless bar, e the moving square, (f) the motionless square and the motionless bar flashed in succession. The area under consideration is shown *top right* at the end of line (c). At the beginnings of lines (c) and (f), the *dashed yellow contours* delimit the cortical area activated by the square and the bar. The spike emission regions are those with high membrane potentials (*red regions* bounded by a black contour), and regions of subthreshold activity are those in *yellow and light blue* surrounded by a white contour. The duration of stimulation is indicated *at the bottom* of the lines. *Green vertical bars* in line (e) indicate the position of the moving square. From [33]

activity of the neurons in the relevant part of the area [top right at the end of line (c)]. At the beginning of lines (c) and (f), the dashed yellow lines delimit the cortical area activated by the square and the bar. The spike emission regions are those with high membrane potentials (red regions bounded by a black contour) and the regions of sub-threshold activity are shown in yellow and light blue surrounded by a white contour. The duration of stimulation is indicated at the bottom of the lines. Vertical green bars in line (e) indicate the position of the moving square. The important result here is that the cortical dynamics (e) and (f) are essentially the same because the activity propagates along horizontal connections, thereby explaining the LMI. As noted by the authors:

These findings demonstrate the effects of spatio-temporal patterns of sub-threshold synaptic potentials on cortical processing and the shaping of perception [33].

In [34], Aaditya Rangan, David Cai, and David McLaughlin proposed a model of the LMI in the context of a model of V1 containing horizontal connections in which removal of inhibition allows lateral spreading of the excitation of orientation cells. The fit with the experimental data is excellent.

5.4 Contact Structure

In the last chapter (Sects. 4.3 and 4.8, and in particular Sect. 4.8.4), we saw how the ‘vertical’ (retino-geniculo-cortical) part of the functional architecture of V1 can be modelled as a *continuous* limit of a pinwheel structure, in terms of the 1-jet bundle of curves in \mathbb{R}^2 . We shall now show how the ‘horizontal’ cortico-cortical part of this functional architecture (discovered in the late 1980s, see [35, 36]) can be modelled in terms of the contact structure.

5.4.1 Integrability Condition and Contact Form

Let us investigate more closely the integrability condition defined in Sect. 4.3.7 of Chap. 4. Recall that we are working with the model

$$\pi_J : J^1(\mathbb{R}, \mathbb{R}) = \mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2,$$

and that, if γ is a smooth parametrized curve $a(s) = (x(s), y(s))$ in the base plane \mathbb{R}^2 [with $x'(s) \neq 0$], it can be lifted to \mathbb{V}_J using the 1-jet map $j^1\gamma(a(s))$ which associates the contact element $(a(s), p_{a(s)})$ with $a(s) = (x(s), y(s))$, where $p_{a(s)} = y'(s)/x'(s)$ is the slope of the tangent to γ at $a(s)$ (if there are vertical tangents, we have to use the compactification of \mathbb{V}_J). Now if $\Gamma = v(s) = (a(s), p(s)) = (x(s), y(s), p(s))$ is a *general skew curve* in \mathbb{V}_J , its projection $a(s) = (x(s), y(s))$ is a curve γ in \mathbb{R}^2 , but Γ is the Legendrian lift of γ if and only if $p(s) = p_{a(s)}(s)$. This integrability condition is the *geometric interpretation of the functional architecture of V1 and the association field*.

The integrability condition can be formulated in a more interesting way. Let $t = (a, p; \alpha, \pi) = (x, y, p; \xi, \eta, \pi)$ be tangent vectors⁵ to \mathbb{V}_J at $v = (a, p) = (x, y, p)$. Along γ (we suppose x is the independent variable), $t = (x, y, p; 1, y', p')$ and the integrability condition $p = y'$ means that we have in fact $t = (x, y, p; 1, p, p')$. It is straightforward to check that this is equivalent to t being in the *kernel* of the differential 1-form $\omega_J = dy - p dx$, where $\omega_J = 0$ simply means that $p = dy/dx$. Indeed, to compute the value of a 1-form ϖ on a tangent vector $t = (\xi, \eta, \pi)$ at (x, y, p) , we apply the rules $dx(t) = \xi$, $dy(t) = \eta$, and $dp(t) = \pi$. If $\varpi(t) = \sum \varpi_i t_i$, where t_i and ϖ_i are the components of t and ϖ with respect to the bases of $T\mathbb{V}_J$ and $T^*\mathbb{V}_J$ associated with the coordinates (x, y, p) , we get $\omega_J(t) = -p \times 1 + 1 \times p + 0 \times p' = -p + p = 0$, since $\omega_J = -p dx + 1 \times dy + 0 \times dp$ and dx (resp. dy, dp) applied to $(1, p, p')$ selects the first (resp. second, third) component 1 (resp. p, p'). Note that, if $p = y'$, the ‘vertical’ component $\pi = p'$ of the tangent vector t in the direction of the p -axis is the *curvature* of the projection γ at a . Indeed, $p = y'$ implies $p' = y''$ and therefore $\pi = p' = y''$.

⁵The reader should not confuse π the projection $\pi : R \times P \rightarrow R$ and π the component of a tangent vector t .

The 1-form ω_J is called the *contact form*, and its kernel is the field \mathcal{K} of planes K_v , called the *contact planes*, with equation $-p\xi + \eta = 0$. The tangent vectors

$$X_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y} = (\xi = 1, \eta = p, \pi = 0), \quad X_2 = \frac{\partial}{\partial p} = (\xi = 0, \eta = 0, \pi = 1),$$

are obvious generators. Now we can express the integrability condition purely geometrically: a curve Γ in \mathbb{V}_J is the Legendrian lift $\tilde{\gamma}$ of its projection γ if and only if it is everywhere tangent to the field \mathcal{K} of contact planes, i.e. if and only if it is an *integral curve* of \mathcal{K} .

5.4.2 Contact Structure as a Cartan Connection

The contact structure has several interesting properties. First, if we no longer consider the projection $\pi_J : \mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ on the plane (x, y) , but the projection $\pi'_J : \mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ on the plane (x, p) , we can show how the 1-form ω_J defines a *connection* in the sense of Elie Cartan, the contact planes K_v becoming what are called the ‘horizontal’ planes. This representation deepens the projective duality which we have already mentioned (Sect. 4.3.6 of Chap. 4). Instead of taking the (x, y) plane as base plane and the axis of the tangent p calculated from the *derivative* $p = dy/dx$ as fibre, we take the (x, p) plane as base plane and the y axis as fibre, the curves γ being given now as functions $p = g(x)$, i.e. as envelopes of their tangents. The y coordinate is then calculated from the *integral* $y = \int y' dx = \int p dx$. A general theorem says that the *curvature* $d\omega_J$ of the connection 1-form ω_J has to be a symplectic form on the new base plane. This is clear, since $d\omega_J = dx \wedge dp$ is the standard symplectic form on the (x, p) plane.

5.4.3 Non-integrability of the Contact Structure

The contact structure \mathcal{K} is the field of planes $K_v \subset T_v\mathbb{V}_J$ defined by the equations $\eta = p\xi$ parametrized by p . As the Legendrian lifts are its integral curves, there are many 1D integrals. But there are nevertheless *no* 2D integrals, no *surfaces* S of \mathbb{V}_J which are tangent to K_v at every point $v \in S$, i.e. such that $T_v S = K_v$. This is due to the fact that the field K_v spins too rapidly with p to be integrable: K_v is the ‘vertical’ plane above the ‘horizontal’ line of slope p and, when p varies along the fibre \mathbb{R}_a above a , it rotates with p .

More precisely, the non-integrability of \mathcal{K} —called *non-holonomy*—results from the violation of the *Frobenius integrability condition*, which says that a 1-form ϖ admits integral surfaces if and only if the wedge product $\varpi \wedge d\varpi = 0$, i.e., $d\varpi(t, t') = 0$ for all tangent vectors t and t' such that $\varpi(t) = \varpi(t') = 0$. Now, for $\omega_J = dy - p dx$,

from the calculus of differential forms,⁶

$$\begin{aligned} d\omega_J &= - \left(\frac{\partial p}{\partial x} dx \wedge dx + \frac{\partial p}{\partial y} dy \wedge dx + \frac{\partial p}{\partial p} dp \wedge dx \right) + d^2y - pd^2x \\ &= -dp \wedge dx = dx \wedge dp, \end{aligned}$$

whence

$$\omega_J \wedge d\omega_J = (-pdx + dy) \wedge dx \wedge dp = dy \wedge dx \wedge dp = -dx \wedge dy \wedge dp.$$

But this 3-form is a volume form of \mathbb{V}_J and vanishes nowhere. By the way, for the basis

$$X_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y} = (1, p, 0), \quad X_2 = \frac{\partial}{\partial p} = (0, 0, 1),$$

of K_v , we have

$$[X_1, X_2] = -X_3 = -\frac{\partial}{\partial y} = (0, -1, 0), \quad X_3 = (0, 1, 0) \notin K_v,$$

since $\omega_J(X_3) = 1 \neq 0$.

5.4.4 Polarized Heisenberg Group

A key point concerning the contact structure of \mathbb{V}_J is that it is left-invariant for a non-commutative *Lie group structure* which is isomorphic to the Heisenberg group and called the *polarized Heisenberg group*. The product is given by

$$(x, y, p) \cdot (x', y', p') = (x + x', y + y' + px', p + p').$$

It is easy to check that this product is associative, that the origin $(0, 0, 0)$ of \mathbb{V}_J is its neutral element, and that the inverse of $v = (x, y, p)$ is $v^{-1} = (-x, -y + px, -p)$. Due to the asymmetry of the coupling term px' , the product is non-commutative. \mathbb{V}_J is a semi-direct product $\mathbb{V}_J = \mathbb{R}^2 \rtimes \mathbb{R}$. The base plane \mathbb{R}^2 ($p = 0$) of the fibration $\pi_J : \mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$ is the commutative subgroup of translations, and the centre Z of \mathbb{V}_J is the y -axis. Indeed, $v' = (x', y', p')$ commutes with all $v \in \mathbb{V}_J$ if and only if for every $v = (x, y, p)$ we have $px' = p'x$, which implies $x' = p' = 0$.

If $t = (\xi, \eta, \pi)$ are vectors in the Lie algebra $\mathcal{V}_J = T_0\mathbb{V}_J$ of \mathbb{V}_J , \mathcal{V}_J has Lie bracket

$$[t, t'] = [(\xi, \eta, \pi), (\xi', \eta', \pi')] = (0, \xi'\pi - \xi\pi', 0)$$

⁶If ω_1 and ω_2 are two 1-forms, $\omega_2 \wedge \omega_1 = -\omega_1 \wedge \omega_2$ (and hence $\omega \wedge \omega = -\omega \wedge \omega = 0$), and if $\omega = df$ is the differential of a function $f(x, y, z)$, then $d\omega = d^2f = 0$.

and is generated as a Lie algebra by the basis of K_v :

$$X_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial x} = (1, p, 0), \quad X_2 = \frac{\partial}{\partial p} = (0, 0, 1),$$

at $v = 0$. Indeed, at 0, $X_1 = (1, 0, 0)$, $X_2 = (0, 0, 1)$ and $[X_1, X_2] = (0, -1, 0) = -X_3$ (the other brackets being zero). The fundamental fact that the basis $\{X_1, X_2\}$ of the distribution \mathcal{K} is *bracket generating*, i.e. Lie-generates the *whole* tangent bundle $T\mathbb{V}_J$, is called the *Hörmander condition*. It is the key property for generalizing our very simple contact structure \mathbb{V}_J to higher dimensions and general manifolds. Moreover, this group is *nilpotent* of step 2, which means that all brackets of the form $[t, [u, v]]$ vanish.

5.4.5 Scale and Characteristic Vectors

It should be emphasized that the definition of the contact structure using the 1-form ω_J contains more information than the definition using its kernel, i.e. the distribution \mathcal{K} of contact planes K_v . Indeed, the 1-forms ω_J and $\alpha\omega_J$ ($\alpha \neq 0 \in \mathbb{R}$) have the same kernel and define the same distribution. The supplementary information encoded in ω_J is the numerical value of ω_J on the ‘characteristic’ tangent vector field (called Reeb field) X_3 transverse to K_v .

Computations in \mathbb{V}_J become very easy if we use the matrix representation

$$v = (x, y, p) = \begin{pmatrix} 1 & p & y \\ 0 & 1 & x \\ 0 & 0 & 1 \end{pmatrix}$$

and

$$t = (\xi, \eta, \pi) = \begin{pmatrix} 0 & \pi & \eta \\ 0 & 0 & \xi \\ 0 & 0 & 0 \end{pmatrix}.$$

Indeed, the product in \mathbb{V}_J becomes the matrix product $v \cdot v'$ and the Lie product in \mathcal{V}_J becomes the commutator $[t, t'] = tt' - t't$. Using this trick, it is easy to see that the contact structure is left-invariant. The *left translation* L_v of \mathbb{V}_J is defined by $L_v(v') = vv'$ and is a diffeomorphism of \mathbb{V}_J whose tangent map at 0 is the linear map

$$\begin{aligned} T_0L_v : \mathcal{V}_J = T_0\mathbb{V}_J &\longrightarrow T_v\mathbb{V}_J \\ t = (\xi, \eta, \pi) &\longmapsto T_0L_v(t) = (\xi, \eta + p\xi, \pi) \end{aligned}$$

The matrix of T_0L_v is

$$T_0L_v = \begin{pmatrix} 1 & 0 & 0 \\ p & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

This shows that the basis $\{\partial/\partial x, \partial/\partial y, \partial/\partial p\}$ of the tangent bundle $T\mathbb{V}_J$ associated with the coordinates $\{x, y, p\}$ is *not* left-invariant. It is the source of non-holonomy. To get a left-invariant basis, we must translate via L_v the basis $\{\partial/\partial x, \partial/\partial y, \partial/\partial p\}_0$ at 0. We get the basis $\{\partial/\partial x + p\partial/\partial y, \partial/\partial y, \partial/\partial p\}$, that is $\{X_1, X_3, X_2\}$.

Now let t be a vector of the contact plane K_0 at 0. Since $\eta = p\xi$ and $p = 0$, we have $\eta = 0$. Its translate $T_0L_v(t)$ is therefore $(\xi, p\xi, \pi)$, and since $\eta = p\xi$, $T_0L_v(t)$ is an element of the contact plane K_v and the contact structure $\mathcal{K} = \{K_v\}$ is nothing else than the left-invariant field of planes left-translated from K_0 . In fact, the 1-form ω_j is itself left-invariant and left translates $\omega_{j,0} = dy$.

5.4.6 Jets, Contact Geometry, and ‘Simplicity’

To give a deeper appreciation of the functional role of the contact structure, we refer to works by Alain Berthoz. In [37], Berthoz introduced the idea of *simplicity* to qualify the novel solutions found by biological evolution to handle the unthinkable complexity of the real world. What he called the *detour principle* of simplicity involves

[...] decomposing complex problems into simpler sub-problems using specialized modules, then putting everything back together again [37, p. 22].

He gives the example of composite mechanical variables, as used in bioinspired robotics:

Instead of the simple variable which [the roboticist] wishes to control, [he can use] a mixture of variables combining position, speed, and acceleration [37, p. 31].

He particularly stresses the fact that ‘geometry is one of the tools of simplicity’ because:

The use of geometry and hence of space to organize neural activity leads to remarkable simplifications with regard to cerebral processing, flexibility, and adaptability [37, p. 166].

We showed (see [38]) that the notion of jet provides a particularly good example of simplicity. The complex problem is to calculate derivatives with respect to the position variables (x, y) and/or integrate them. As pointed out by Jan Koenderink, it is too complex to be carried out by point processors like neurons (see Sect. 4.3.5 in Chap. 4). The functional architecture of V1 is a simplex solution for this complex problem which has been discovered by evolution and which uses only simple operations like taking values at points:

1. We add an independent variable p : simple neurons with orientation selectivity.

2. The values of the new variable are organized in specialized modules, i.e. contact elements (a, p) recording position and orientation: orientation hypercolumns.
3. The process simply involves measuring point values of (a, p) .
4. To reconstruct, we introduce a functional architecture which ensures that taking point values of the three variables (x, y, p) is indeed equivalent to differentiating and/or integrating with respect to the initial variables (x, y) .

But it is precisely this simplicity that formalizes the idea of a jet, the functional architecture corresponding to a specific geometric structure, namely *contact geometry* on the jet space.

5.5 Illusory Contours as Sub-Riemannian Geodesics

In this neurogeometrical framework, we can interpret in a *principled way* the variational process giving rise to the Kanizsa illusory contours constituting the second example of strange phenomena evoked in the *Introduction* (Sect. 2.11, Fig. 2.3). In a Kanizsa figure, the pacmen⁷ define two contact elements (a, p) and (b, q) , and an illusory contour interpolating between (a, p) and (b, q) is a skew curve Γ in \mathbb{V}_J from (a, p) to (b, q) which is at the same time:

1. A Legendrian lift of a curve γ in the base plane \mathbb{R}^2 , i.e. an integral curve of the contact structure, with integrability condition $p(x) = y'(x)$.
2. A curve that is ‘as straight as possible’, as already emphasized by Shimon Ullman in 1972 [39] when he introduced the idea of a variational model ‘minimizing total curvature’.

The simplest way to satisfy these two requirements is to model illusory contours using *geodesics* for a natural metric because, since the variation of p measures the curvature κ of γ , we minimize both the length and the curvature of the projection γ at the same time. But, due to condition (1), the metric *only has to be defined on the distribution* \mathcal{H} of contact planes K_γ . In our first synthesis of neurogeometry, in 1999, entitled *Vers une neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux* [40], geodesics were computed by solving the Euler–Lagrange equations with the Lagrange multiplier expressing the integrability condition. And these equations were reformulated in the framework developed by Robert Bryant and Phillip Griffiths for variational models on Lie groups (see [41]). Later on, the search for minima was naturally interpreted in terms of a suitable metric defined on the contact distribution \mathcal{H} . Such a metric is said to be *sub-Riemannian*, so the modelling of illusory contours was embedded in the mathematical context of *left-invariant sub-Riemannian metrics on nilpotent Lie groups*.

⁷The term ‘pacman’ comes from a famous Japanese video game in which a disc with a variable angular sector representing a stylized mouth (the ‘pac-man’) has to find its way through a maze while ‘eating pac-dots’ and at the same time avoiding enemies.

So, in a rigorous sense, Kanizsa illusory contours are *geodesics of a sub-Riemannian geometry defined on the contact structure of the fibre bundle of 1-jets of planar curves*.

5.6 Curvature Detectors and 2-Jets

5.6.1 Data

Regarding the geometric features that are derivatives, or better, jets of curves, we focused on 1-jets, that is, tangents and first-order derivatives. We saw in Sect. 4.3.5 of Chap. 4 and Sect. 5.4.6 that there are two fundamentally different implementations of a tangent calculation:

1. One uses what Jan Koenderink called array processors, and for a curve with equation $y = f(x)$, corresponds to the standard school algorithm


$$f'(x) = \lim_{\Delta x \rightarrow 0} \frac{f(x + \Delta x) - f(x)}{\Delta x}.$$

2. The other is a ‘simplex’ implementation, precisely the one carried out by V1, using purely *point processors* which are tangent detectors. As already explained, it consists in including the supplementary variable p as independent variable, modules (columns) measuring the different values of p , and a functional architecture (the horizontal intracortical connections) ensuring that the contact geometry of the 1-jets is indeed implemented.

However, there is a certain amount of experimental evidence to support the idea that there are also *curvature calculations*, i.e. carrying out *second derivatives* on curves. For example, Ohad Ben-Shahar and Steve Zucker [42, 43] went back to the data of Bosking et al. [6] and explained that collinearity does not exhaust their information content. Insofar as the process of integrating local orientations into global curves uses the *variance* of the orientation, a certain curvature is processed by the functional architecture. The question is therefore whether this is done in a classic way using array processors which calculate second derivatives or in a simplex way using point processors which are curvature detectors, introducing the curvature κ as a new independent variable related by a specific functional architecture to the previous variables p and $a = (x, y)$.

The idea put forward by Zucker and Ben-Shahar is that curvature is implemented in V1 in the second way: there are point curvature detectors connected by a functional architecture that implements the space of 2-jets of curves. This means that there is detection of the *osculating circles* of curves at their regular points and ‘a curvature-based field of long-range horizontal connections’ [42] which defines a relation, not of

Fig. 5.23 An osculating circle detector. From Ben-Shahar et Zucker [42]



collinearity, but of *cocircularity* between neighbouring orientation cells. Figure 5.23 shows an example.

In continuous neurogeometric models, the functional architecture of the space of 2-jets which extends the contact structure of the space of 1-jets corresponds to what is known as the *Engel structure*.

5.6.2 Curvature, 2-Jets, and Engel Structure

In addition to x, y, p or x, y, θ , we have to add a *fourth* independent variable κ and a supplementary 1-form which forces its interpretation as a curvature. This is what Sarti, Citti, and I did together. We work in $\mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R}$ with the contact 1-form $\omega_J = dy - p dx$ and the non-holonomic basis of the contact planes

$$\left\{ X_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y}, X_2 = \frac{\partial}{\partial p} \right\},$$

the third basis tangent vector X_3 being given by the Lie bracket

$$[X_1, X_2] = -X_3 = -\frac{\partial}{\partial y}.$$

We introduce a fourth variable κ and working in the space $\tilde{\mathbb{V}}_J = \mathbb{R}^2 \times \mathbb{R} \times \mathbb{R}$, we express the idea that the natural interpretation of κ is associated with the second derivative $f''(x)$ for curves of equation $y = f(x)$. The space $\tilde{\mathbb{V}}_J$ of the $\{x, y, p, \kappa\}$ is the space of 2-jets $J^2(\mathbb{R}, \mathbb{R})$ and its canonical structure, the Engel structure,⁸ is the Pfaff system comprising the two 1-forms ω_J and $\tau_J = dp - \kappa dx$.

The kernel of τ_J in $\tilde{\mathbb{V}}_J$ is generated by the three tangent vectors

⁸Friedrich Engel was one of the main disciples and collaborators of Sophus Lie.

$$X_1^\kappa = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y} + \kappa \frac{\partial}{\partial p} = X_1 + \kappa X_2, \quad X_3 = \frac{\partial}{\partial y}, \quad X_4 = \frac{\partial}{\partial \kappa},$$

whereas the kernel of ω_J extended to $\widetilde{\mathbb{V}}_J$ is generated by X_1^κ , X_2 , and X_4 . The distribution of planes is now $\text{Span}\{X_1^\kappa, X_4\}$, and it generates the whole Lie algebra because

$$[X_1^\kappa, X_4] = -X_2 = -\frac{\partial}{\partial p}, \quad [[X_1^\kappa, X_4], X_1^\kappa] = -X_3 = -\frac{\partial}{\partial y}.$$

5.6.3 ‘Good Continuation’ and the Statistics of Natural Images

In Sect. 3.6 of Chap. 3 and Sect. 4.4.5 of Chap. 4, we saw that the receptive profiles of simple V1 neurons and also their ‘vertical’ (hyper)column organization could be reconstructed from the rather special statistical properties of natural images by making the assumption that the neural coding optimizes from the information standpoint, assuming sparse coding.

We can thus ask whether the ‘horizontal’ functional architecture of V1 implementing the 1-jets (tangents), 2-jets (osculating circles), and the Gestalt principle of good continuation as reflected by the structure of the association field cannot also be reconstructed from the statistics of natural images. This does indeed seem to be the case. For example, in [44, p. 1935], Mariano Sigman et al. investigated ‘the geometric regularities of oriented elements [...] present in an ensemble of visual scenes.’ They demonstrated ‘strong long-range correlations’ and showed that these are well explained by ‘a very simple geometric rule, cocircularity’, i.e. the transition from 1-jets to 2-jets in our approach, this being a ‘natural extension of collinearity to the plane’:

Our findings provide an underlying statistical principle for the establishment of form and for the Gestalt idea of good continuation. [...] The geometry of the pattern of interactions in primary visual cortex parallels the interactions of oriented segments in natural scenes [44, p. 1939].

The experimental method involves measuring the correlation between the orientations of the edges between an arbitrary origin a and another point b for a corpus of $N = 4\,000$ natural images comprising $1\,536 \times 1\,024$ pixels. Filters are 7×7 pixels, the angular resolution 1 arcmin/pixel, and we use 16 orientations in $[0, \pi]$. In pixels, the coordinate differences Δx and Δy between a and b vary within $[-256, 256]$. In order to treat edge orientations (which are cliffs in the level lines of the image) on the same footing as line orientations (which are peaks and valleys), we use typical even and odd filters $G_2(a, \theta)$ and $G_3(a, \theta)$ which are second and third derivatives of Gaussians (more precisely, what are known as quadrature filters with a phase offset of $\pi/2$) obtained by translating and rotating $G_2(0, 0)$ and $G_3(0, 0)$ (see Sect. 3.3.1 of Chap. 3), and we write the orientation energy as $E(a, \theta) = G_2^2(a, \theta) + G_3^2(a, \theta)$

with a discretization of θ , e.g. in 16 values.⁹ For each point a of each image i , we seek θ_{\max} maximizing $E_i(a, \theta)$ and we say that an orientation is detected at a if $E_i(a, \theta_{\max}) \geq T$, where T is a perceptually tested threshold. The correlation function for all the images is given by the following average of integrals:

$$C(\Delta x, \Delta y, \varphi, \psi) = \frac{1}{N} \sum_{i=1}^{i=N} \iint E_i(a, \varphi) E_i(a + \Delta a, \psi) da .$$

Figure 5.24 shows the results, and in particular, the correlation $C(a, \psi)$ between an orientation $\varphi = 0$ at 0 and an orientation ψ at a , for $\psi = 0, \pi/8, \pi/4, 3\pi/8$, and $\pi/2$. To simplify, we can also consider histograms giving the number of times an orientation φ at 0 coexists with an orientation ψ at a . In each vignette of the figure, the correlation is represented for ψ constant (ψ is the second orientation at the bottom of the vignette). Red is the maximum and dark blue the minimum. The vignette Fig. 5.24f represents solutions to the cocircularity problem for orientations at 20° (red) and 40° (green lines at 30° and 120°), with 0° corresponding to the vertical. This is based on an elementary geometrical observation: if Γ is the circle of radius R and centred at $(R, 0)$, hence with equation $r = 2R \cos(\theta)$ and such that the y axis is tangent to it at 0, and if $c = (r, \theta)$ is a point of Γ , the tangent to Γ at c makes an angle $\phi = 2\theta + \pi/2$ (modulo π) with the x axis (see Fig. 5.25). Indeed,

$$\left\{ \begin{array}{l} x = r \cos(\theta) = 2R \cos^2(\theta) , \\ y = r \sin(\theta) = 2R \sin(\theta) \cos(\theta) , \\ y' = 2R[\cos^2(\theta) - \sin^2(\theta)] = 2R \cos(2\theta) , \\ x' = -4R \cos(\theta) \sin(\theta) = -2R \sin(2\theta) , \\ \tan(\phi) = y'/x' = \tan(2\theta + \pi/2) . \end{array} \right.$$

5.7 Relationship with Wavelets

In Sect. 3.4 of Chap. 3, we saw that the ganglion cells in the retina and the cells of the LGN carry out a wavelet analysis of the optical signal. Such a link between natural vision and image compression was introduced by David Marr and extended by Stéphane Mallat, among others. Standard wavelets in the form of a Laplacian of a Gaussian are not able to follow the contours of objects. In order to extend the idea, we can first introduce oriented wavelets resembling the RPs of simple V1 cells. However, by taking biomimetics a little further, and observing the functional architecture for edge detection in V1, Mallat introduced the idea of bandlets, along

⁹There are some technical problems in the definition of the energy, but these are well known in signal theory. To obtain good results, the even/odd filters must be refined, for example, taking the *Hilbert transform* of G_2 instead of G_3 .

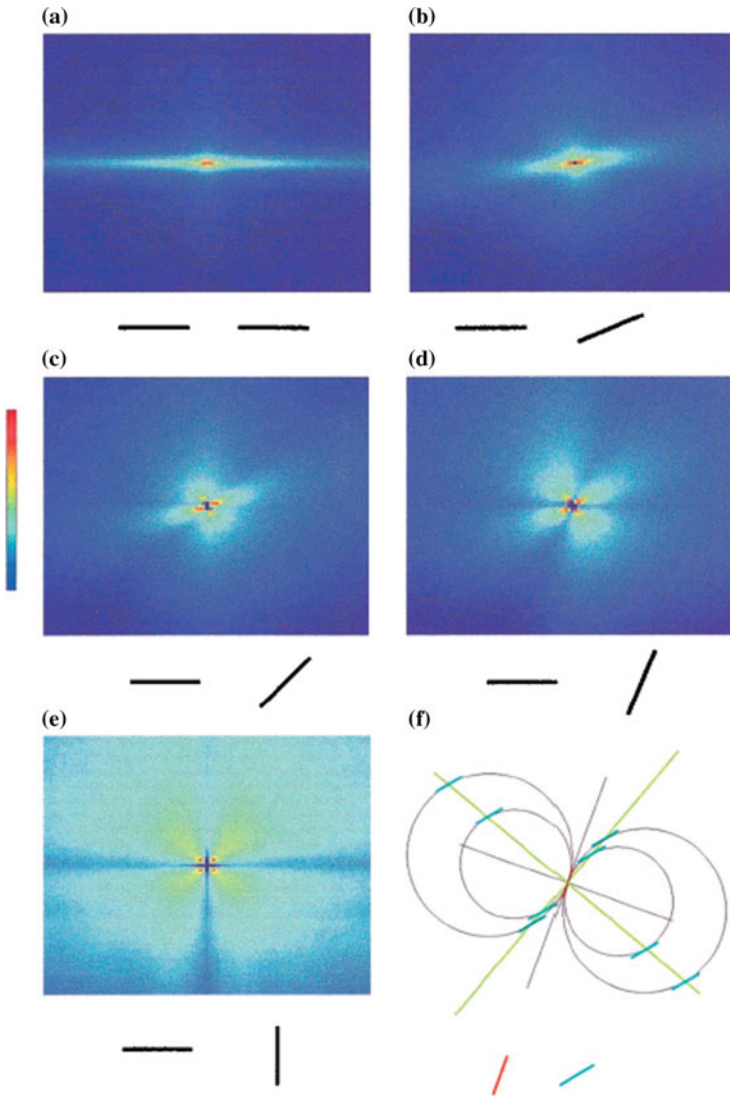


Fig. 5.24 Correlation $C(a, \psi)$ between an orientation $\varphi = 0$ at 0 and an orientation ψ at a in a corpus of 4 000 natural images. Each vignette represents the correlation for constant ψ , where ψ is the second orientation shown at the bottom of the vignette. Red is the maximum and dark blue the minimum. In (a), strict collinearity is represented by a very long horizontal lobe: there is only correlation with $(a, 0)$ when a lies on the horizontal axis. As the orientation ψ rotates, this collinearity lobe rotates and shortens and a second lobe appears. For $\psi = \pi/2$ in (e), the two lobes become equivalent and symmetric, with an orientation of their symmetry axes at $\pi/4$ and $3\pi/4$. **f** Solutions to the cocircularity problem for orientations 20° (red) and 40° (green lines at 30° and 120°), with 0° corresponding to the vertical. From Sigman et al. [44]. Copyright (2001) National Academy of Sciences, USA

Fig. 5.25 If Γ is the circle with equation $r = 2R \cos(\theta)$, so with radius R , centre at $(R, 0)$, and tangent to the y axis at 0 , and if $c = (r, \theta)$ is a point of Γ , the tangent to Γ at c makes an angle of $\phi = 2\theta + \pi/2$ (modulo π) with the x axis

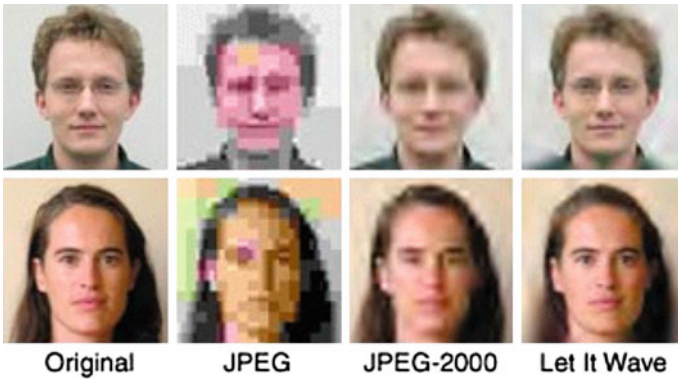



Fig. 5.26 Comparing Stéphane Mallat’s bandlet method with JPEG and JPEG-2000. The gain in efficiency is spectacular. See the home page of the start-up company *Let it wave*, and also Mallat and Peyré [45]

with a new algorithm. The compression properties are remarkable since a human face can be coded with only 500 bytes (see Fig. 5.26).

As stressed by Mallat and Peyré [45]:

Wavelet bases [in the form of Laplacians of Gaussians] take advantage of isotropic regularity on neighbourhoods of varying size. The existence of geometric regularity can be interpreted as an anisotropic regularity. Although discontinuous at a point on a contour, the image can nevertheless be differentiable in the direction of the tangent to this curve. The bandlet transform exploits this kind of anisotropic regularity by constructing orthogonal vectors which extend in the direction of maximal regularity.

The first bandlets were constructed by Erwan Le Pennec [46]. The idea was to relate the fact that wavelet coefficients are maximal near the edges with the Gestalt principle of good continuation along the edges.

The analogy between neurophysiology (the retina and the $V1$ area) and image compression (wavelets and bandlets) is therefore profound, and we come back to the idea that the constraint of compression may be the evolutionary origin of the geometry.

5.8 Structure of the $V2$ Area

Having explained the second part of the functional architecture of $V1$ comprising the horizontal connections, let us take a moment to discuss further clues regarding the structure of the other areas in the visual cortex, starting with area $V2$. We have already mentioned $V2$ on many occasions, in particular in Sects. 3.1 and 3.2.6 of Chap. 3 and Sects. 4.1, 4.4.7, 4.9.4, 4.10.4, and 4.12.1 of Chap. 4. $V2$ receives inputs directly from $V1$, and also from the thalamus, via the pulvinar nuclei among others. It also gets feedback from the dorsal stream (the ‘Where’ pathway) and the ventral stream (the ‘What’ pathway).

When discussing the blobs in $V1$ in Sect. 4.10.4 of Chap. 4, we explained how $V2$ decomposes into ‘stripes’. For example, in the macaque, $V2$ comprises 14 sets of stripes which are well-defined functional regions. We saw that they can be classified in three groups: thin, pale (or intermediate), and thick. They are between 1 and 2 mm wide and a complete series thin/pale/thick measures about 4 mm. Each type of stripe tends to process a particular quality: colour and luminosity for the thin ones, orientation, direction, and depth (binocular disparity and stereopsis) for the thick ones, and orientation and shape for the pale ones. $V2$ is an essentially binocular area, and this is why a good demarcation criterion between $V1$ and $V2$ is the end of the ocular dominance domains of $V1$.

In $V2$, we encounter the same kinds of neurogeometric issues as in $V1$, regarding the relations between the different retinotopic maps. For example, Ts’o et al. [47] studied the iso-orientation and iso-disparity lines in a thick band for the two geometric features orientation and binocular disparity. They observed their transversality relations. There is an excellent discussion of this structure in the chapter by Roe [48] in *The Primate Visual System*.

In the standard view, the blobs of $V1$ connect to the thin stripes of $V2$ (colour), the interblob regions of layers 2/3 of $V1$ to the pale stripes of $V2$, and layer 4B of $V1$ to the thick stripes of $V2$. However, the actual situation seems to be more complex, for example, with the interblobs also connecting to the thick stripes.

Regarding the receptive fields, it can be shown that, by taking natural images as inputs and using plausible models, the $V2$ neurons have *two* preferred orientations, the first $p \in P$ and the second $s \in S$, and this allows them to detect corners, junctions, and crossings. The orientation domains are larger than those in $V1$, and each domain, for a given primary orientation p , is structured in hypercolumns for the secondary orientation s (see, e.g. Sit and Miikkulainen [49]). An idealized model would thus be a bundle with 2D fibre, but also implementing a 1D sub-bundle:

$$P \times S \times \mathbb{R}^2 \longrightarrow P \times \mathbb{R}^2 \longrightarrow \mathbb{R}^2, \quad (p, s, a) \longmapsto (p, a) \longmapsto a.$$

Regarding direction, the situation is rather subtle. There is no direction map in the V1 area of the macaque (while there is in the cat, the ferret, and carnivores, and we shall return to this in Sect. 5.11), but there is in the thick and pale stripes of V2 (and not in the thin stripes). As direction is linked to motion, the latter is therefore processed in V2. Moreover, the thick stripes project onto the middle temporal area MT (V5). In fact, to be more precise, in V1, there are neurons sensitive to different directions but the point is that they are not segregated by a functional architecture, whereas they are in V2. This is the same phenomenon as the absence of orientation columns in rodents that was discussed in Sect. 3.8 of Chap. 3 and in Sect. 4.2 of Chap. 4. Figure 5.27 from Lu et al. [50] shows the pinwheels that form the directional domains of V2, still for the macaque, while Fig. 5.28 shows the pinwheels with segments indicating the preferred orientation and arrows indicating the preferred direction, the length of the arrow being proportional to the selectivity.

Area V2 plays a very important role in the geometrical constitution of shapes because it ‘calculates’ local geometrical features involving several orientations such as occlusions of contours, crossings, or complex illusory contours, especially in the thin and pale stripes.

In particular, neurons in V2 detect *virtual* orientations which are *orthogonal* at ends of segments (see Figs. 5.29 and 5.30 due to Esther Peterhans and Rüdiger von der Heydt [51]). This can explain subtle phenomena such as the illusory contours produced in the comb effect shown in Fig. 5.31, by assuming that the responses of V1 to those segments that are genuinely contained in the stimulus (the teeth of the comb) are inhibited at the end points by inversion of the response, with V2 activating the orthogonal direction.

This is a good example of inferred contours in the sense of Anna Roe:

Inferred contours are not defined by luminance contrast but rather by more global features that are perceived only by grouping multiple cues across space. They include occluded contours, texture element borders, and have been referred to as higher-order contours, illusory contours, cognitive contours, and anomalous contours [48, p. 121].

Illusory contours like the one in the comb effect have been studied in the cat, among other animals, by Sheth et al. [53]. Figure 5.32 shows the orientation maps for real objective orientations (Fig. 5.32c) and illusory subjective orientations (Fig. 5.32d). Pinwheels appear for the difference between Fig. 5.32c and d, and the histogram of this difference has a sharp peak at angle 0 which shows that the V2 neurons respond strongly to subjective contours.

Fig. 5.33a and b from [53] shows, as we have seen, that the situation is different in V1. The peak of the difference between Fig. 5.33a and b lies at $\pi/2$, which corresponds to real orientations.

Thanks to the treatment of complex boundaries in the thick stripes and surface qualities in the thin stripes, V2 also plays a role in ‘calculating’ fundamental segregation relations between a figure and its background in a perceived image, a calculation




Fig. 5.27 Pinwheel map forming the directional domains in the macaque V2 area (*scale bar 1 mm*). In the upper figure, each panel represents the difference between the direction indicated by the *arrow* and the opposite direction. From Lu et al. [50]

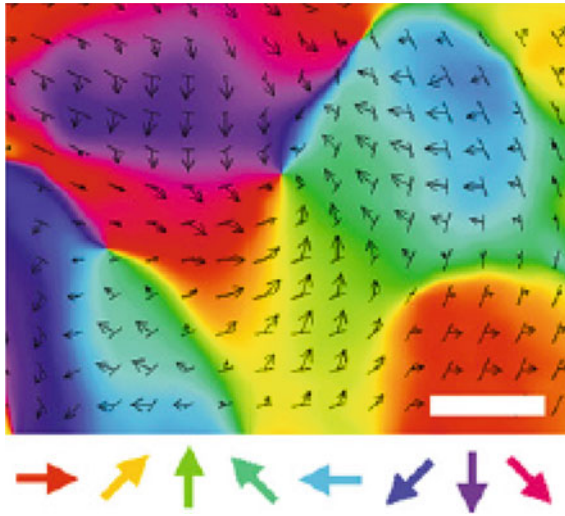


Fig. 5.28 Pinwheels with *segments* indicating the preferred orientation and *arrows* indicating the preferred direction. The length of the *arrow* is proportional to the selectivity. From Lu et al. [50]

that constitutes one of the great challenges in the theory of perception. As noted by Anna Roe, the problem is to determine:

[...] how boundaries and surfaces relate to each other, thereby forming the initial stage of figure/ground segregation [48, p. 127].

The problem of attributing edges to one of the two regions that they separate (border ownership) cannot be solved by V1 alone. However, V2 responds to edges belonging to a surface because its receptive fields respond, not only to pieces of lines but also to segments bounding one of the two half-spaces they separate. Background/figure segregation then results from a *global* and highly *contextual* calculation that can integrate these local edge-assigning data.

5.9 Colour and Area V4

5.9.1 Colour Constancy: Semir Zeki and René Thom

After areas V1 and V2, information about colour is processed in particular in area V4, as shown by Semir Zeki, then in the infratemporal (IT) area, where the colours are integrated with the treatment of shapes in the ventral pathway. Since the late 1970s, Semir Zeki has done much work on the way the *non-retinotopic* V4 area produces a perception of *colour constancy* over quite extensive spatial regions, while these colours are in fact clearly variable with regard to the spectral composition of




Fig. 5.29 In this figure and the following, the two discs represent the stimulus (bar or grating) with orientations of 0° and 90° . The rows of windows represent the responses of the neuron to the stimulus gradually rotated through 0° – 180° , with time on the *horizontal axis* and the angle on the *vertical axis*. *Left* A neuron of V1 responding to a preferred orientation p . *Right* When the stimulus is a comb whose teeth have the orientation p , the neuron responds to this orientation and does not respond to the virtual edge with the orthogonal orientation p^\perp . Redrawn from Peterhans and von der Heydt [51] and von der Heydt and Peterhans [52]




Fig. 5.30 *Left* Neuron in V2 responding to a preferred orientation p . *Right* When the stimulus is a comb whose teeth have orientation p , the neuron does not respond to this orientation, but responds rather to the virtual edge with the orthogonal orientation p^\perp . Redrawn from Peterhans and von der Heydt [51] and von der Heydt and Peterhans [52]

the stimuli (see his celebrated book *A Vision of the Brain* [54]). On the one hand, we only have to observe the pixels of a digital image to convince ourselves of this extreme variability, and on the other hand, we only have to recall how colour can be a predicate in statements like ‘the ball is red’ to convince ourselves of its constancy. The




Fig. 5.31 Comb effect involving the V2 area




Fig. 5.32 **c** Orientation maps in the cat for real objective orientations. **d** Orientation maps for illusory subjective orientations. There are pinwheels for the difference between **(c)** and **(d)**. The histogram of this difference has a sharp peak at 0, showing that the V2 neurons respond strongly to the subjective contours. From Sheth et al. [53]





Fig. 5.33 **a** and **b** Still for the cat, orientation maps of the V1 area for objective and subjective orientations, respectively. In contrast to the situation in V2 (Fig. 5.32), the peak in the difference between **(a)** and **(b)** lies at $\pi/2$, which corresponds to real orientations. From Sheth et al. [53]

Fig. 5.34 Old world monkey. Projections of the CO blobs of $V1$ onto the thin stripes of $V2$ and the interblobs onto the pale and thick stripes of $V2$, followed by projections of the thin and pale stripes of $V2$ to $V4$ and the thick stripes of $V2$ onto $V5$ (MT). From Shapley and Hawken [57]



assumption was that the retinotopic neurons of $V1$ would be ‘wavelength-only’, i.e. sensitive to true wavelengths λ , even if they include the processing described in the last section, whereas the non-retinotopic neurons of $V4$ would be ‘colour-only’, i.e. contextually sensitive to a given colour, even when that colour does not necessarily correspond to the actual wavelengths. Colour constancy over homogeneous regions bounded by edges would thus involve $V4$ through highly contextual non-retinotopic effects.

Semir Zeki’s idea can be formulated in terms of fibre bundles, as proposed by René Thom independently in [55, p. 248, Chap. 6] (see our paper [56]). Let $R = \mathbb{R}^2$ be the visual field. We consider the bundle $\pi_\Lambda : \mathbb{R}^2 \times \Lambda \rightarrow \mathbb{R}^2$ with base space \mathbb{R}^2 and fibre the colour space Λ , e.g. the hue circle or the 3D hue–saturation–lightness (HSL) space. The objective colour is then a generally highly non-smooth section σ of this bundle. However, the colour space Λ is categorized by the perceived colours C_i , $i = 1, \dots, n$. The problem is then to approximate σ by a much smoother section Σ , but one which can have qualitative discontinuities. This is a problem of *segmentation*. We shall encounter another example of such a problem in the second volume. Colour constancy corresponds to regions W_k of \mathbb{R}^2 whose image $\sigma(W_k)$ under the section σ is entirely contained in a single chromatic category C_i .

Figure 5.34, also due to Shapley and Hawken [57], shows projections of the CO blobs of $V1$ onto the thin stripes of $V2$ and the interblobs onto the pale and thick stripes of $V2$, along with projections of the thin and pale stripes of $V2$ to $V4$ and the thick stripes of $V2$ onto the motion area $V5$ (MT), all this for the case of the old world monkey.

Area $V4$ would thus appear to be specialized in colour processing, as claimed by Semir Zeki, but this process seems to begin in the $V1$ area. This issue is still debated. Moreover, the functional organization of $V4$ is rather complex with colour regions separated from orientation regions (see, e.g. Tanigawa [58]).

5.9.2 Objectivism Versus Subjectivism

The processing of colour provides an interesting example of the way contemporary neurophysiology can revitalize, and often solve, philosophical problems which have a long and hotly disputed history. There are no end of conceptual reflections on colour, even if we only consider those produced in the last century, as exemplified by Husserl and Wittgenstein. But the results we have brought together here can throw light on one of the most intriguing problems for these authors, namely the objective or subjective nature of colour.

This problem was very clearly spelt out in a target article by Francisco Varela, Evan Thompson, and Adrian Palacios [59], published in 1992 in *Behavioural and Brain Science* under the title *Ways of colouring: Comparative colour vision as a case study in cognitive science*. The authors discuss the question of colour in the framework of an enactive theory of perception which aims to go beyond conventional antinomies: computational objectivism versus neurophysiological subjectivism, external objectivity versus internal cognitive processing, detection of objective properties versus construction of enacted properties, heteronomous input–output systems versus autonomous self-organized systems. Their aim is ‘to offer a new empirical and philosophical perspective on colour vision’ which is experiential, i.e. not objectivist, and ecological in the sense of Gibson, i.e. not subjectivist.

Let S be a coloured surface made up of points s . There is a map $R : S \rightarrow \Delta, s \mapsto R(s)$ from S to the space Δ of activities of colour receptors (the three types of cone). However, due to post-retinal processing which involves the chromatic oppositions R/G and B/Y, we have a composite map $C(s) = MR(s)$, where M is the post-retinal transformation. The problem then is to go beyond the classic antinomy of objectivism versus subjectivism as formulated in the following way by the authors:

1. *Thesis*: ‘The distal world can be specified independently of the animal; it casts images on the perceptual system whose task is to recover the world appropriately from them’.
2. *Antithesis*: ‘The perceptual system projects its own world and the apparent reality of this world is merely a reflection of internal laws of the system’.

In our response to this target article (reprinted in [60]), we put forward a solution to this dilemma based on Thom’s idea, as discussed in Sect. 4.10.3.2 of Chap. 4 on binocular rivalry and in the previous section on homogeneous surfaces. From the objectivist computational standpoint, the cognitive task of the visual system is to reconstruct external 3D visual scenes E from their highly ambiguous 2D retinal projections. As we have seen, the geometric information about E is encoded in the qualitative discontinuities of the optical signal as the apparent contours of objects and provides a geometric configuration (W, K) , where W is the spatial extent of E and K a set of qualitative discontinuities delimiting the objects and structuring their shape. So what of sensory qualities like colour, which fill up the different regions of W bounded by K ? The objectivist approach makes the following assumptions:

1. At each point w of W , there is a well-defined objective value $R(w)$ of the reflectance of surfaces in E .
2. $R(w)$ varies continuously as a function of w , except along the discontinuities of K .
3. The perceptual quality of colour $C(w)$ encodes $R(w)$.

However, this is not plausible. Indeed, as we have seen, cortical processing of retinal transduction converts the reflectance $R(w)$ into the perceptual quality of colour $C(w)$ in a much more complex manner than a simple coding. Let σ be the relation between $R(w) \in \mathfrak{R}$, where \mathfrak{R} is the space of reflectances, and $C(w) \in \mathfrak{Q}$, where \mathfrak{Q} is the space of colours. The above antinomy comes from a misinterpretation of the relation σ :

1. For the objectivist, $C(w)$ is a simple coding of $R(w)$, and σ can effectively be ignored because it is an isomorphism between \mathfrak{R} and \mathfrak{Q} .
2. For the subjectivist, on the other hand, $C(w)$ cannot be reduced to $R(w)$, and σ is an incommensurability relation, resulting in an explanatory gap.

The solution is that the colour space \mathfrak{Q} is like a space of *internal states* of the perceptual system \mathfrak{S} , whereas the values $R(w)$ are like external stimuli. There is thus no direct correspondence $R(w) \rightarrow C(w)$, but only the realization of an internal state $C(w)$ cued by a stimulus $R(w)$. Then $C(w)$ is the dynamical response of the system \mathfrak{S} to the stimulus $R(w)$, and the colour opposition property is a characteristic feature of this.

The perceptual internal states are attractors of internal dynamics X defined on the internal space \mathfrak{Q} . Clearly, \mathfrak{Q} and X are highly dependent on the species, whereas R is objective. The discontinuous maps $C(w)$ along K thus result from complex dynamical processes. Here we have a field of internal dynamics X_w controlled by W , and $C(w)$ is the attractor of X_w which is selected by the stimulus $R(w)$. The latter is a *control parameter* for X_w and $W \times \mathfrak{R}$ is a control space. In this kind of morphodynamic model, the qualitative discontinuities comprising K are interpreted as dynamical bifurcation events: when w crosses K , the attractor $C(w)$ selected by $R(w)$ bifurcates towards another attractor, and this induces a qualitative discontinuity (see Sect. 4.10.3.2 of Chap. 4).

The relation σ between $R(w)$ and $C(w)$ is therefore nothing like a simple map between spaces. It results from the composition of at least four processes:

1. An objective application $W \rightarrow \mathfrak{R}$, $w \mapsto R(w)$.
2. An internal dynamical field $\sigma : W \rightarrow \mathfrak{X}$, $w \mapsto X_w$, embedding the extension W of the visual scene E in the (infinite-dimensional) function space \mathfrak{X} of the internal dynamics X .
3. A selection process which, depending on the objective input $R(w)$ and the general selection rules, selects an attractor A_w of X_w .
4. A phenomenological presentation of A_w as a quality $C(w)$.

We understand the sense in which these morphodynamical models bring together computational objectivism and neurophysiological subjectivism. There is no conflict

between the thesis that the colour system detects the reflectance and the antithesis that it constructs a specific response to the reflectance. As a control parameter, $R(w)$ selects an attractor $C(w)$ and this is in fact a construction. However, insofar as $C(w)$ is causally controlled by $R(w)$, $C(w)$ is cued by $R(w)$ and thus detects $R(w)$. The antinomy arises due to the confusion between the stimulus $R(w)$ and the attractor $C(w)$ which it selects, that is, between the external control of the internal dynamics X_w and the internal state A_w it determines.

5.10 Motion and the MT Area (V5)

Located in the superior temporal sulcus (STS) and investigated by Zeki and Dubner in the 1970s, then Barlow and Albright in the 1980s, then Van Essen and Orban in the 1990s, the middle temporal (MT) area, also known as V5, specializes in analyzing motion. For example, Fig. 4.11 from Gur and Snodderly [61] in Sect. 4.2 of Chap. 4 shows three projections from V1 to MT in the macaque. One goes from layer $4C\alpha$ of V1 (where the magnocellular pathways of the LGN are preferentially projected) to the MT via the spin stellate neurons of layer $4B$ of V1. Another goes from layer 6 of V1 and another from layer $4Cm$ (m middle, located in V1 between layers $4C\alpha$ and $4C\beta$) to MT via layer $3B$ of V1 which projects onto the thick stripes of V2, which in turn project onto MT. The excellent review [62] by Richard Born and David Bradley discusses the other pathways to MT through the superior colliculus and the thalamic nucleus of the pulvinar.

For their part, the receptive profiles are larger than those of V1, integrating more information and processing five features of the stimuli: (1) retinal position (MT is a retinotopic area), (2) direction of motion, (3) speed, (4) binocular disparity, (5) size of the stimulus. The link with binocular disparity, i.e. the way the horizontal shift between the two retinal images of a single object is an indication of depth, is important. The study [63] by Gregory DeAngelis and William Newsome presents a schematic view of the functional architecture of MT dealing with the direction of motion and binocular disparity.

5.11 Models of Direction and Singularities of Functions

Before discussing the genetic control of neurogenesis, let us return for a moment to the geometric feature of *direction* (of motion) which complements the geometric feature of orientation already analyzed at length. During development, selective responses to direction are established after selective responses to orientation and require a rich visual experience. In carnivorous species like cat, there is a functional architecture devoted to direction in V1. It does not occur in the V1 area of species like the macaque where, as we saw in Sect. 5.8, it is found in V2.

5.11.1 Projection of Singularity Lines from V2 to MT

For these primates, the reader may consult the paper [64], in which Shigeru Tanaka and Hiroyuki Shinbata simulate the projection of *singularity lines* or *fractures* from V2 onto MT. They show that the directions vary continuously and that the responses are unidirectional, except along the fractures, which act like lines of discontinuity across which the direction jumps by π and where the responses are thus bidirectional. Moreover, all directions are present at the ends of the fractures. These lines come about because orientations are defined modulo π , while directions are defined modulo 2π . Fractures result from balancing between these two geometrical features.

In their model, Tanaka and Shinbata considered MT as a network of neurons with lateral connections $U_{j,j'}$ with synaptic weights in the form of Laplacians of Gaussians joining neurons at positions j and j' . These connexions are modulated by synaptic inputs $\sigma_{j,\theta}$ from V1 and V2, with $\sigma_{j,\theta} = 1$ if there is a connection between the neuron j of MT and a neuron of V1/V2 coding the orientation θ . Let $\Gamma_{\theta,\theta'}$ be the correlation between the activities of two neurons of V1/V2 projecting onto j and j' , respectively, and coding the orientations θ and θ' , respectively. A key feature of the model is to take a sum of two Fourier components for Γ , the first specified modulo 2π (directions) and the second modulo π (orientations):

$$\Gamma_{\theta,\theta'} = c_1 \cos(\theta - \theta') + c_2 \cos(2(\theta - \theta')) .$$

Then by applying standard learning methods, we obtain direction and orientation maps. The latter depend on the parameter $r = c_2/c_1$, which measures the strength of the direction c_1 with respect to that of the orientation c_2 . For $r \sim 0$, the only singularities are points. For $r = 1$, however, we obtain Fig. 5.35, in which arrows indicate the preferred directions and the fractures along which the direction jumps by π and reversals are clearly visible. For large enough r , the directions become random. Figure 5.36 shows a 1D section of Fig. 5.35. We see the jumps in the direction across these lines of discontinuity.

5.11.2 Swindale's Model

Here we pause for a moment to consider the work by Nicholas Swindale, Amiram Grinvald, and Amir Shmuel on area 18 of the cat [65]. The preference curves have one peak for the orientation and two for the direction. Figure 5.37 from [65] shows the orientation pinwheel map (Fig. 5.37a), the width and height of the preference peak (Fig. 5.37c, e), the direction pinwheel map (Fig. 5.37b), the strength of selectivity (Fig. 5.37d), and the relations between the fractures (black lines) of the direction map and the singularities of the orientation pinwheels with their chiralities $+/-$ (green/red) (Fig. 5.37f).

Fig. 5.35 Model of preferred directions and orientations in area MT when these are balanced. Arrows indicate the preferred directions. The fractures along which the direction jumps by π are clearly visible. From [64]





Fig. 5.36 1D section of Fig. 5.35. The jumps in direction across the lines of discontinuity, i.e. the fractures, are clearly visible. From [64]



These results raise several interesting questions about how to model this. When we gave a geometric treatment of the orientation maps in the geometrical framework of fibre bundles, we considered the set S of pinwheels s_i , and with each point a of $\mathbb{R}^2 - S$, i.e. with each *regular point*, we associated the preferred orientation p_a of the orientation column corresponding retinotopically with this position. The orientation map thus corresponds to a *section* $\sigma : \mathbb{R}^2 - S \rightarrow V = \mathbb{R}^2 \times \mathbb{P}^1$ of the bundle $\pi : V = \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$ defined on $\mathbb{R}^2 - S$. As we saw in Sect. 4.8.3 of Chap. 4, the closure of the image $\sigma(\mathbb{R}^2 - S)$ of σ in V is defined on the whole of \mathbb{R}^2 , but it is no longer a section because above the s_i there is not just one point of the fibre $\mathbb{P}^1_{s_i}$, but the whole fibre $\mathbb{P}^1_{s_i}$. To see what is happening here, we have to *blow up* the singularities s_i and lift each field line γ of the orientation field which leaves s_1 at angle θ_1 to arrive at s_2 at angle θ_2 in the curve Γ of V starting at height θ_1 in the fibre $\mathbb{P}^1_{s_1}$ above s_1 to arrive at height θ_2 in the fibre $\mathbb{P}^1_{s_2}$ above s_2 . It is the continuous limit of this model, in which all points become centres of infinitesimal pinwheels, that we have used as an idealization of $V1$.

However, let us remain in a discrete model with a finite and somewhat limited number of pinwheels and consider the directions θ rather than orientations p . The directions are organized according to a functional architecture made up of direction columns. In the fibre bundle framework, we go from the fibre \mathbb{P}^1 to the fibre \mathbb{S}^1 , considering the bundle $\delta : \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$ and representing the direction field by a section τ of δ above $\mathbb{R}^2 - S$, where S is the singular locus of fractures and direction

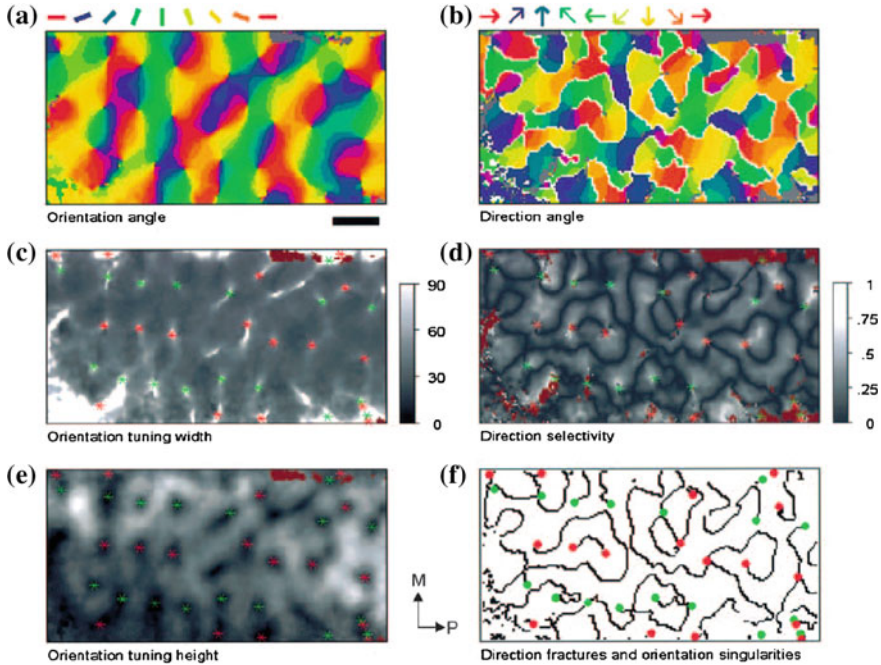


Fig. 5.37 Area 18 of the cat. **a** Pinwheel orientation map. **c** and **e** Width and height of the preference peak. **b** Pinwheel direction map. **d** Strength of direction selectivity. **f** Relations between fractures (black lines) of the direction map and singularities of the orientation pinwheels with their chiralities \pm (green/red). M and P denote the medial and posterior axes of the cortex, respectively. Scale bar 1 mm. From Swindale et al. [65]

pinwheels. But now we would like to take into account the fact that the preferred direction θ_a at $a \in \mathbb{R}^2 - S$ is the maximum of a Gaussian function $F_a(\theta)$ defined on the whole fibre \mathbb{S}_a^1 . The field is therefore now defined by the family of functions $F_a(\theta) : \mathbb{S}^1 \rightarrow \mathbb{R}$, parametrized by $\mathbb{R}^2 - S$. The section $\tau : \mathbb{R}^2 - S \rightarrow \mathbb{R}^2 \times \mathbb{S}^1$ is then found by taking the maximum of $F_a(\theta)$ at each position a , i.e. the point where $dF_a(\theta)/d\theta = 0$, which is unique if F_a is Gaussian. These are called *critical points* of F_a .

Let us suppose then, as in the paper by Swindale et al. [65], that the probability functions F_a can be more complicated than Gaussians and can have several maxima, for example, two. There will thus be competition between these two maxima, i.e. two preferred directions θ_1 and θ_2 , and when a varies, there could be a reversal of dominance between θ_1 and θ_2 . As the fibre is the compact manifold \mathbb{S}^1 rather than \mathbb{R} , the two maxima are separated by two minima and if the critical points of $F_a(\theta)$ are set above the regular points a , we obtain a surface with four sheets, two preferred in the sense of having locally maximal selectivity, corresponding to the two maxima, and the two others having locally minimal selectivity, corresponding to the minima.

If one of the maxima collides with one of the minima and disappears, then there remains only one preferred sheet.

5.11.3 *Singularities and Universal Unfoldings*

This brings us naturally to a mathematically well-known context, namely unfoldings of singularities of differentiable functions, already discussed in Sect. 4.10.3.2 of Chap. 4 in the context of binocular rivalry and in Sect. 5.9.2 when discussing the objectivity and/or subjectivity of colours. This highly technical theory was developed in the 1960s by great geometers like Hassler Whitney, René Thom, Vladimir Arnold, Bernard Malgrange, and John Mather. For a review of these results, the reader is referred to our 1982 compilation [66].

The general framework of the theory can be described in a little more detail as follows. We wish to analyze the structure of differentiable functions $f : M \rightarrow \mathbb{R}$ defined on a differentiable manifold M which we call the *internal space*, in our case $M = \mathbb{S}^1$. Let \mathfrak{F} be their functional space. It is an infinite-dimensional vector space over \mathbb{R} , because there are an infinite number of degrees of freedom to slightly deform f . We introduce a natural topology \mathcal{T} on \mathfrak{F} , known as the Whitney topology. When M is compact, which is our case, this is simply the uniform convergence of functions and all their partial derivatives. When M is not compact, it controls in addition the behaviour of functions at infinity. We consider the action on \mathfrak{F} of the group $G = \text{Diff}(M) \times \text{Diff}(\mathbb{R})$ of diffeomorphisms of M and \mathbb{R} , i.e., differentiable coordinate changes of the source and target. This specifies an equivalence on the f , and we wish to study the structure of the f up to this equivalence. To begin with, we introduce the key idea of *structural stability*: by definition f is structurally stable if any function g close enough to f in \mathfrak{F} in the sense of the topology \mathcal{T} is equivalent to f , in other words if the differentiability type of f does not change under small deformations. Two fundamental theorems can then be proven. First, *Morse theorem*, which says that f is structurally stable if and only if:

1. All its critical points, i.e., points where the gradient of f vanishes, are non-degenerate minima, saddle points, and maxima (non-degenerate means they do not arise from the merging of several critical points).
2. All its values at critical points, called critical values, are distinct.

So according to Morse theorem, there are only two causes of instability: (i) degeneracy of critical points, which leads to what are known as *bifurcation catastrophes*, and (ii) equality of critical values, which leads to what are known as *conflict catastrophes*.

By recurrence, we can attribute a sort of degree of instability c to the functions f : a structurally stable function will be of degree 0, the coincidence of two critical values will increase c by 1, as will the merging of one critical point with a non-degenerate critical point. For small values of c , it is easy to see what is happening. But c can be infinite and lose all intuitive meaning. For example, a constant function on an open set of M is infinitely unstable.

The second fundamental theorem concerns *universal unfolding* of local singularities with finite codimension. The exact theory is highly technical, but fortunately it is also fairly intuitive for small values of the degree of instability c . We assume that M is reduced to a small neighbourhood U (and even to an infinitesimal neighbourhood) of a point where f is singular. Let \mathfrak{F}_U be the space of the g on U . We consider $f \in \mathfrak{F}_U$ and let \tilde{f} be its orbit under the action of the group $G_U = \text{Diff}(U) \times \text{Diff}(\mathbb{R})$. If f is structurally stable, \tilde{f} contains an *open neighbourhood* of f , that is, *all functions close enough to f* . However, if f is unstable, this is not the case and locally \tilde{f} is a strict subspace of \mathfrak{F}_U whose ‘tangent space’ $T_f \tilde{f}$ at f can be defined. When the degree of instability is not too high, $T_f \tilde{f}$ admits supplementary spaces W at f in \mathfrak{F}_U (this notion can be rigorously defined) and these have a finite dimension equal to c . The difference in dimension between \mathfrak{F}_U and \tilde{f} is therefore c , even though both dimensions are infinite, and this is why c is called the *codimension* of f . If we parametrize the points $a \in W$ by local coordinates (a_1, \dots, a_c) in such a way that $(0, \dots, 0)$ corresponds to f and the c basis vectors $(1, 0, \dots, 0)$, $(0, \dots, 1, \dots, 0)$ (with the 1 in the i th place), $(0, 0, \dots, 1)$ correspond to the functions h_i , $i = 1, \dots, c$, then all functions g close enough to f in W can be written in the form

$$g = f + \sum_{i=1}^c a_i h_i .$$

This means that the h_i form a minimal system of functions which are not equivalent to f , which are linearly independent, and which can generate with \tilde{f} a whole neighbourhood of f in \mathfrak{F}_U , in such a way that each function close enough to f can be written in the form $f' + \sum_{i=1}^{i=c} a_i h_i$, with $f' \in \tilde{f}$ equivalent to f . We see that such a W parametrizes a *family* f_a of functions such that $f_0 = f$. This family is called a *deformation* or an *unfolding* of f . The space W is called the *external space* of this unfolding. The universal unfolding theorem asserts that all the supplementary spaces W constructed in this way are *equivalent* in a well-defined sense and that they classify all the possibilities for stabilizing f by small deformations. They are called *universal unfoldings*. When the codimension c is large enough, things get more complicated and refinements are needed.

Returning to the global case, e.g. compact M , if f is structurally stable except at an isolated singular point, we may then unfold this local singularity by deforming f in one of its neighbourhoods U without changing it qualitatively elsewhere.

We shall use this notion of universal unfolding in the case of a 1D internal space $M = \mathbb{S}^1$, with codimension ≤ 2 since, in our case, the parameters a of the families of functions $F_a(\theta)$ under consideration vary in the external space $W = \mathbb{R}^2$, which is 2D. In most applications, the functions f are a sort of energy function with the form of a potential well with several minima at the bottom of the well. Our case is analogous because the maximization of a tuning curve $F_a(\theta)$ with one or more peaks is equivalent to the minimization of $f_a(\theta) = 1/F_a(\theta)$ or $f_a(\theta) = -\log(F_a(\theta))$ (as for the information associated with a probability distribution) which has the form of a potential well as shown in Fig. 5.38.




Fig. 5.38 Maximization of a tuning curve $F(\theta)$ is equivalent to minimization of an associated potential well $f(\theta)$

We shall use the best known of the universal unfoldings f_w , that of the singularity known as the *cusp*, corresponding to the merging of three non-degenerate critical points, two minima and one maximum which separates them. We have already seen this in Sect. 4.10.3.2 of Chap. 4. The codimension c of f_0 is 2, so the external space W is 2D. One of the parameters, called the splitting factor, controls the fact that the potential has just one or three non-degenerate critical points. A second parameter, called the bias factor, controls the dominance of one of the minima over the other in the region of W where f_w has three critical points. Figure 5.39 (repeat of Fig. 4.128 in Sect. 4.10.3.2) shows this universal unfolding. We see that W is partitioned into 3 regions by a cusp curve with two branches $K_{b,l}$ and $K_{b,r}$ and a symmetry axis K_c :

1. The exterior of the cusp, a region where f_w has only one minimum, which is non-degenerate.
2. The interior of the cusp where f_w has two non-degenerate minima separated by a non-degenerate maximum, this region being itself partitioned into two sub-regions by the symmetry axis of the cusp:
 - 2.1. In the left-hand region lying between $K_{b,l}$ and K_c , it is the right-hand minimum which dominates.
 - 2.2. In the right-hand region lying between K_c and $K_{b,r}$, on the other hand, it is the left-hand minimum which dominates.

Along K_c , the two minima are at the same height, and these special values of w therefore correspond to a conflict catastrophe (of codimension 1). Along the branches $K_{b,l}$ and $K_{b,r}$ of the cusp, one of the minima merges with the maximum at a point of inflection and these special values of w therefore correspond to bifurcation catastrophes (of codimension 1).

The decomposition of W into regions is called a *stratification*, and the sub-manifolds $K_{b,l}$, $K_{b,r}$, K_c , and $\{0\}$ of W are called strata. Note the relation between the dimension d_k of the strata S_k and the codimension c_k of the associated singularities f_w for $w \in S_k$: indeed, we have $c_k = c - d_k$. The function f_0 has codimension 2 and corresponds to the cusp point which does indeed have dimension $c - 2 = 0$ and the strata $K_{b,l}$, $K_{b,r}$, K_c corresponding to singularities of codimension 1 do indeed have dimension $c - 1 = 1$.

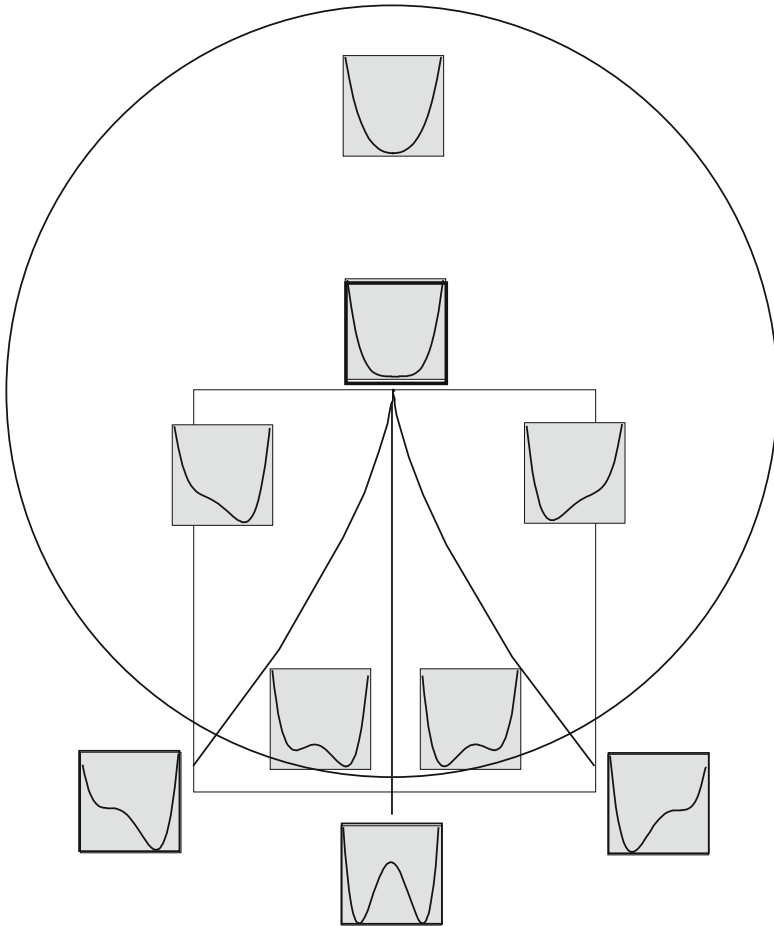


Fig. 5.39 Repeat of Fig. 4.128 from Sect. 4.10.3.2 of Chap. 4. Universal unfolding of the cusp singularity f_0 (merging of three non-degenerate critical points, two minima and one maximum) of codimension $c = 2$. The external space W is 2D. It is partitioned into 3 regions by a cusp curve with two branches $K_{b,l}$ and $K_{b,r}$ and a symmetry axis K_c . Exterior to the cusp, f_w has only one non-degenerate minimum. In the interior of the cusp, f_w has two non-degenerate minima separated by a non-degenerate maximum. Between $K_{b,l}$ and K_c , it is the right-hand minimum that dominates, and between K_c and $K_{b,r}$, it is the left-hand minimum

A fundamental consequence of this formula is the ‘transitivity’ of the universal unfoldings: if we consider a 1D section in W transverse to K_c , it is a universal unfolding of the singularity of codimension 1 (conflict between the two minima) that f_w exhibits along K_c . Likewise for $K_{b,l}$ and $K_{b,r}$ (where the singularities of codimension 1 are points of inflection). Figure 5.40 shows such a transverse section transverse with the three singularities I , J , and I' corresponding to $K_{b,l}$, K_c , and $K_{b,r}$, respectively.




Fig. 5.40 Transitivity of universal unfoldings. We consider a 1D section T in W transverse to $K_{b,1}$, K_c , and $K_{b,r}$ at the points I, J, I' . Locally at I, J, I' and for $t \in T$, (T, f_t) is a universal unfolding of the singularity f_I, f_J , or $f_{I'}$ of codimension 1




Fig. 5.41 Universal unfolding of the conflict catastrophe in the vicinity of J (see Fig. 5.40) for the probability distributions

Figure 5.41 shows the universal unfolding in the vicinity of J for the probability distributions.

5.11.4 Swindale Model (Continued)

Let us come back to the results of Swindale et al. [65]. We consider tuning curves $F_a(\theta)$ with one or two peaks parametrized by $a \in \mathbb{R}^2 - S$. As the internal space is the circle \mathbb{S}^1 , in addition to the minimum separating the maxima in the above images, there is also a minimum corresponding to the joining of the tails of the distributions F_a . Figure 5.42 shows the variation of the two-peak tuning curve $F_a(\theta)$ when a follows a path transverse to a fracture. We see that the discontinuity in the preferred direction does indeed correspond to a conflict catastrophe like the one described in




Fig. 5.42 Variation of the two-peak direction tuning curve $F_a(\theta)$ when a follows a path transverse to a fracture. The discontinuity in the preferred direction corresponds to a conflict catastrophe like the one described in Fig. 5.41. The line of fracture (*white*) corresponds to the conflict stratum between the two maxima. The x axis goes from 0 to 2π , and the vertical mark on the x axis indicates the average of the two peaks. It is located on the direction given by the minimum when a is on the conflict stratum. From Swindale et al. [65]




Fig. 5.43 End point of a fracture in the direction map which falls on a pinwheel. The wheel of eight directions is 0 (red), $\pi/4$ (violet), $\pi/2$ (blue), $3\pi/4$ (blue-green), π (green), $5\pi/4$ (green-yellow), $3\pi/2$ (yellow), $7\pi/4$ (orange), 2π (red). There is a discontinuity of π (green/red) in crossing the fracture. At the pinwheel, the distribution $F_0(\theta)$ becomes flat. As we move around the pinwheel, there is a continuum red \rightarrow violet \rightarrow blue \rightarrow blue-green \rightarrow green, and only half the direction wheel ($0, \pi/4, \pi/2, 3\pi/4, \pi$) is relevant. When a maximum occurs on the other half, it has no effect. This would appear to be a cusp singularity in the orientation map. From [65]

Fig. 5.41, the fracture corresponding to the conflict stratum K_c . We also see that when we follow the fracture towards the cusp, we move gradually from a red/green-yellow gap to a red/green, then orange/green, then yellow/green gap, and that the maxima thus shift towards the green while at the same time moving closer together, and this does indeed conform qualitatively to the cusp geometry despite the lack of accurate data.

Figure 5.43 is even more interesting. Indeed, it shows what happens when a fracture in the direction map stops on a pinwheel of the orientation map. The wheel of eight directions is 0 (red), $\pi/4$ (violet), $\pi/2$ (blue), $3\pi/4$ (blue-green), π (green), $5\pi/4$ (green-yellow), $3\pi/2$ (yellow), $7\pi/4$ (orange), and 2π (red). The discontinuity of π (green/red) in crossing the fracture is clearly visible, and we also see that,




Fig. 5.44 Superposition of the orientation pinwheel map and fractures in the direction map of Fig. 5.37

when we go around the pinwheel, there is a continuum red \rightarrow violet \rightarrow blue \rightarrow blue-green \rightarrow green. Note also that, at the pinwheel, the curve $F_0(\theta)$ becomes very flat. The cusp model thus seems to be correct here. However, the curves $F_a(\theta)$ have two maxima. But we made the assumption that only half of the direction wheel was relevant, i.e. $0, \pi/4, \pi/2, 3\pi/4, \pi$, and that the maximum on the other side was not involved.

Let us say a little more about these data. Figure 5.44 shows a superposition of the orientation pinwheel map and the fractures in the direction map of Fig. 5.37. When we compare the two maps more carefully, we observe a certain number of things:

1. To begin with the singular points which are simultaneously orientation pinwheels and the ends of fracture lines do indeed have the same chirality, as they should. The chirality tells us whether, when we move round a singular point in a certain sense, the orientations or the directions coded by colours rotate in the same sense or not.
2. The second thing we observe is that the discontinuities are mainly red/green, i.e. $0, \pi$, and when we go around the singular point we go through either blue ($\pi/2$) or yellow ($3\pi/2$).
3. The singularities R/G(B) and R/G(Y) can have both chiralities.
4. There seem to be *triple points*, e.g. B/G/Y.

More accurate experimental data will be necessary, but it seems that what we have here are cusp singularities in the direction map.

5.12 Neural Morphogenesis and Its Genetic Control

In this first volume focusing on neurophysiology, we have brought together experimental data relating to the visual system, and in particular the functional architecture of its primary cortical areas (early vision), and above all V1. We have begun to explain what geometric framework presents itself most naturally for modelling such architectures by focusing on stabilized adults of species like the cat, the macaque, or the human being. Even though we have at times mentioned the questions of development and learning, we have said hardly anything about the *morphogenesis* of these extraordinary neural systems. This is why, even though it goes beyond our own area of expertise, we would like to say a word here about this morphogenesis and its genetic control, and in particular about axon guidance (or axon pathfinding), which allows special kinds of connectivities to set up through systems of ‘vertical’ retino-geniculo-cortical and ‘horizontal’ intracortical connections. These are fascinating processes of great complexity, in which molecular factors and cell interactions set up the functional geometry of the neural hardware, an internal and ‘immanent’ geometry on the basis of which, as explained in the *Introduction*, the external and ‘transcendent’ geometry of perceived space comes into being.

The sources for this section are studies by Tudor Badea, Lisa Baye, Keely Bumsted-O’Brian, Alain Chedotal, David Feldheim, John Flanagan, Alex Kolodkin, Brian Link, Liqun Luo, Rodrigo Martins, Todd McLaughlin, Dennis O’Leary, Alain Prochiantz [67], Benjamin Reese, Linda Richards, Dorothea Schulte, Gabriel Scicolone, and Marc Tessier-Lavigne.

5.12.1 Guidance of Axon Connections

The establishment of axon connections is an intrinsically difficult and extremely diverse problem because there is such a wide variety of different neurons and projections. What are the specific molecular mechanisms that guide axons? How do axons find their way and carry out the series of connections allowing them to reach their target? In the 1990s, our understanding of these processes developed considerably. The purification of several proteins in the extracellular matrix and membranes brought to light many molecules like the semaphorins (1993), the netrins (1994), the ephrins (1995), and the slits. During development, axons are guided by these membrane or diffusible proteins, accompanied by morphogens, growth factors, cell-adhesion molecules (CAMS), and immunoglobulins, which stimulate or inhibit and orient growth spatially. Since then, amazing progress has been made in our understanding of the way neural positioning and morphological differentiation of cells are controlled within this system.

As noted by Chedotal and Richards [68, p. 2]:

[There is a] preferred growth of developing axons along preexisting axonal tracts.

There are pioneer axons which open up new paths to be followed later by other axons:

The incredible complexity of the mammalian brain, and the targeting and growth of axons over long distances, require a unique strategy for enabling brain wiring to occur during development [68, p. 5].

Many mechanisms are involved: guidance of axons and dendrites, branching, target recognition, synaptogenesis, degeneration (apoptosis) of neurons and their axons, regeneration of axons, pruning of dendritic trees, etc. (see, e.g. the synopsis by Kolodkin and Tessier-Lavigne [69]).

Axon guidance is based on the fact that the growth cone of the axons detects molecular signals called *guidance cues* or *guidance factors*, either fixed or diffusible in the extracellular environment. These provide directionality instructions by exerting attractive or repulsive influences. In its membrane, the growth cone contains receptors which recognize these guidance cues and interpret the signals chemotropically. The activated receptors have effects on the cytoskeleton, and if the growth cone is sensitive to concentration gradients of such morphogens, the effect will be asymmetric and the growth of the axon will thus be redirected.

There are three main kinds of guidance cues:

1. *Adhesion factors* like laminin or fibronectin in the extracellular matrix, or the cadherins and immunoglobulin factors.
2. Attractive and repulsive *tropism factors* transforming the cytoskeleton, like the netrins and the semaphorins.
3. *Modulation factors* modifying the sensitivity of the growth cones to certain other factors, like the neurotrophins.

5.12.2 *Transcription Factors and Homeoboxes*

Recall what was said very briefly in Sect. 3.2.3 of Chap. 3. On the genetic level, development, differentiation, and regulation are controlled by a network of transcription factors, i.e. proteins which bind to specific DNA sequences to control the transcription of their genetic information into mRNA. They are key to the regulation of gene expression. They can be activators or inhibitors, i.e. activate or inhibit the recruitment of the RNA polymerase enzyme which carries out the transcription. There are several thousand of these in a genome like the human genome. About 10% of genes are involved in coding them. As it often happens that genes are flanked by several binding sites for different transcription factors which work together, the combinatoric effects are highly complex. Transcription factors contain DNA-binding domains (DBD) attaching to specific DNA sequences of the genes whose expression they regulate.

Homeoboxes are DNA sequences present in certain genes called *hox genes*, which are involved in the regulation of morphogenetic processes. A homeobox encodes a protein homeodomain which, once expressed, can bind to the DNA in a specific way. The *hox genes* thus encode transcription factors which control the transcription of cascades of other genes.

5.12.3 Some Guidance Factors

For the following sections, the interested reader may consult the *GeneCards* [70] of the *Human Gene Compendium* at the Weizmann Institute (www.genecards.org). We shall say a few words about certain guidance factors. Receptors bind in a specific way to ligands, molecules they recognize:

1. The *semaphorins* are membrane proteins which act as guidance molecules on the axon growth cone. In fact they act as short-range inhibitory signals through receptors like *plexins*, *neuropilins*, and *integrins*. Each semaphorin is characterized by the expression of a specific region of about 500 amino acids (its sema domain). There are eight classes. For example, in humans, their genes are SEMA3A/.../G, SEMA4A/.../G, SEMA5A/B, SEMA6A/.../D, and SEMA7A. For example, the integrins are receptors for SEMA7 which mediate the attachment of a cell to the surrounding tissue by binding the cell surface to extracellular components like fibronectin, vitronectin, collagen, or laminin.
2. The *netrins* are trophic axon guidance factors with the double function of attracting or repelling axons. Their receptors are DCC for attraction and Unc5 for repulsion. For example, Netrin-1 activates DCC receptors which open Ca^{++} channels, producing a CaCnLs calcium current, and triggering the cAMP signalling pathway (depending on cyclic adenosine 3',5'-monophosphate). When the Ca^{++} ions enter the endoplasmic reticulum of the cells, this induces an attraction of the growth cone. On the other hand, activation of Unc5 activates cyclic guanosine monophosphate (cGMP), which decreases the CaCnLs current and induces a repulsion of the growth cone.
3. The *ephrins* and their transmembrane receptors *Eph* form the largest family of receptor tyrosin kinases (RTK). Their extracellular domain interacts with the ephrin ligands. Activation, through their ligands, of the *Eph* receptors located on the axon growth cones, is essential for the signal transduction mechanisms via their kinase domain. It deflects growth to regions of high ligand concentration. The ephrins-A bind to the cell membrane by glycosylphosphatidylinositol (GPI) bonds. The ephrins-B bind to the membrane by a transmembrane domain containing a small cytoplasmic motif called *PDZ*. Humans have eight ephrins, including 5 ephrins-A interacting with nine *EphA* receptors and three ephrins-B interacting with five *EphB* receptors. Others are found in other vertebrates.
4. The *slits* are proteins repelling growth cones via their *Robo* (Roundabout) receptors.
5. The *cell-adhesion molecules* (CAMS) are membrane proteins mediating adhesion between axons. They are essential for the fasciculation of fibres and include two subgroups, the *IgSF-CAMS* (*Ig* = immunoglobulins) and the *cadherins*.




Fig. 5.45 **a** The cells are attached to the two membranes, viz., apical (retinal pigment epithelium RPE, *top*) and basal (*bottom*), by junction and adhesion proteins, and their nuclei can be positioned at different heights (or phases). The extracellular adenosine triphosphate (ATP, supplying energy) produced by the RPE plays an important role in interkinetic nuclear migration. The Notch-Delta signalling pathway is based on the Notch gene and its transmembrane receptors and ligands (Delta among others). The Sonic Hedgehog Shh is a protein coded by the *hh* gene involved as ligand in the Hedgehog signalling pathway. It is a morphogen diffusing along a concentration gradient and whose concentration acts on the expression of certain genes. The glia-derived factor (GDF) is secreted by the ganglion cells and plays an important role in synaptogenesis and controlling the timing of progenitor competence. Lysophosphatidic acid (LPA) is an important extracellular regulator for the motion of nuclei. **b** The cells position themselves with different phases *M*, *S*, *G1*, *G2*. For example, some *G2* cells (*red*) whose nuclei are close to the basal membrane will have a specific cell fate. From Baye and Link [72]

5.12.4 Neurogenesis of the Retina

We begin with the transcription control of retinal neurogenesis. In this system, the development of cell lines from their multipotent progenitors can be monitored very closely. Two mechanisms are particularly important for movement between the apical retinal pigment epithelium (RPE) and the basal membrane of the retina. One of these is interkinetic nuclear migration (INM), which is an oscillating migration, first apical-basal then basal-apical, of the neuroepithelial cell nuclei which are attached to these two bounding surfaces of the retina, the different positions (or phases) corresponding to the different cell types. The *M*-phase occurs at the apical positions near the RPE, and the *S*-phase at the basal positions. The other mechanism is nuclear translocation in which, at the end of the cell cycle, the nucleus of the progenitor migrates into an appropriate laminar position, and the soma then retracts and detaches itself from the bounding surfaces. An excellent review of genetic control can be found in the paper *Development of the retina and optic pathway* by Reese [71]. Figure 5.45 due to Baye and Link [72] shows how the cells are attached to the two membranes by junction and adhesion proteins and how they position themselves with their different phases.

During neurogenesis, the ganglion cells first differentiate and develop the axons which will form the optic nerve. Then come the horizontals, the cones, the amacrine,




Fig. 5.46 General schematic of retinal neurogenesis. *Left* A multipotent progenitor attaches to the apical (RPE) and basal retinal membranes. Lower horizontal arrows and oblique arrows from bottom to top represent cell divisions increasing the number of progenitors. Lateral inhibition and gene expression processes (thick vertical arrows from top to bottom) lead to cell differentiation (small upper horizontal arrows). The different kinds of cells appear successively and gradually modify the cell environment via signal pathways (thick coloured diagonal arrows from top to bottom) which trigger the production of subsequent cell types. From Reese [71]

the rods, the bipolars, and the glials (Müller cells), while the plexiform layers are layers of connections. These seven kinds of cell appear in a characteristic order. Figure 5.46 due to Reese [71] gives a schematic view of the dynamics of cell production.

The proliferation of progenitors is precisely controlled by the transcription factors (Sect. 5.12.2) *Hes1* and *Hes5* of *bHLH*-type (basic helix-loop-helix, for the role of *bHLH* factors in the regulation of retinal cell fate specification, see [73]), and neurotransmitters like acetylcholine, noradrenaline, dopamine, and serotonin (or 5-hydroxytryptamine). The reader is referred, for example, to the paper by Martins and Pearson, R.A. [74]. More generally, the *Hes* genes in mammals code for transcriptional repressors *bHLH*. There are seven classes.

Regarding the genetic control of retinal cell differentiation, we have already mentioned this in Sects. 3.2.4 and 3.2.5 of Chap. 3 when discussing the cone opsins (the *OPN1LW* gene, or long-wave-sensitive opsin 1, etc.) and ganglion cells with transcription factors (i) *Brn3a* for the dendritic tree (*Brn3a* = brain-specific homeobox/*POU* domain protein 3A, coded in humans by the gene *POU4F1* = *POU* domain, class 4, transcription factor 1, where the acronym *POU* comes from the three transcription factors pituitary-specific Pit-1, Octamer, and Unc-86), (ii) *Brn3b* and *Isl1* (insulin gene enhancer) for differentiation into ganglions rather than amacrine or horizontals, and (iii) *Math5* for cell fate. Likewise, e.g. in mice (see Li et al. [75]), *Foxn4* (forkhead box protein N4) controls the amacrine and horizontals by activating the expression of the *Math3* factors (which encode a *bHLH*), *NeuroD1* (*bHLH* factor of neurogenic differentiation 1, i.e. $\beta 2$ beta-cell E-box transactivator 2, regulating the insulin gene), *Prox1* (encoding the Prospero homeobox protein 1), and *Ptf1a* (pancreas-specific transcription factor *bHLH*-type 1a), with *Prox1* differentiating between them.

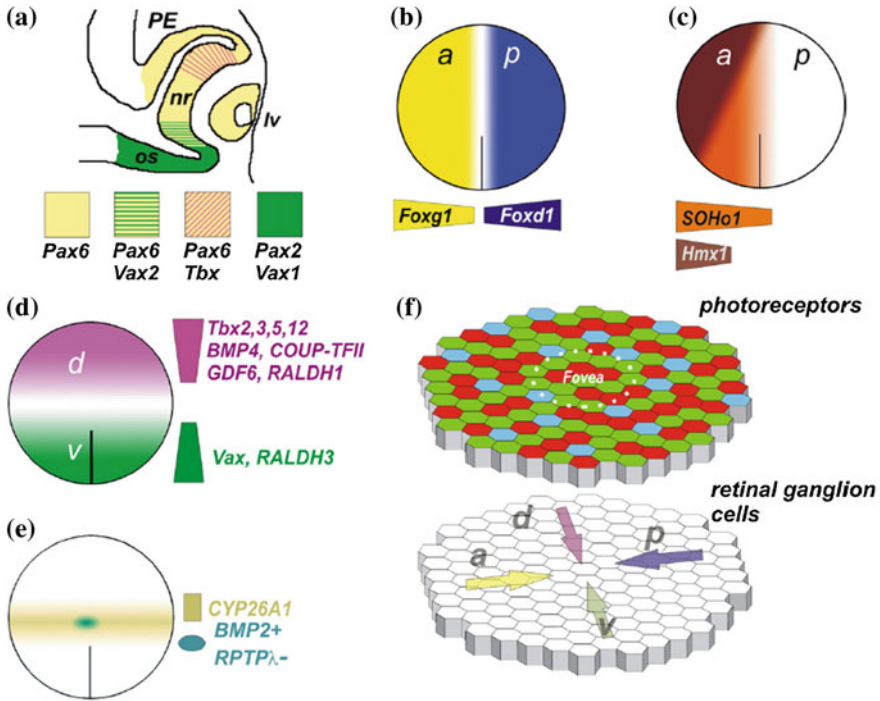


Fig. 5.47 Gene expression in the vertebrate retina during embryogenesis. **a** At the beginning of embryogenesis, the future optic nerve (*os*, optic stalk) and the neural retina (*nr*) compartmentalize by expressing, for the former, the transcription factor *Pax2* and the homeodomain *Vax1*, and for the latter, *Pax6* (also expressed by the pigment epithelium *PE* and the optical vesicle *lv*), *Pax6* and *Vax2* giving the ventral retina and *Pax6* and *Tbx* giving the dorsal retina. **b** During formation of the anterior–posterior axis, the anterior part *a* expresses *Foxg1* and the posterior part *p* expresses *Foxd1*, the medial line (*white*) expressing neither of these two genes. **c** At a later stage, after invagination of the optic vesicle, the patterning of the *a/p* axis is brought about by expression of the homeotic genes *SOHo1* and *Hmx1*. **d** Formation of the dorsal–ventral axis *d/v* of the retina. The dorsal retina expresses the transcription factors *Tbx*, *COUP-TFII*, *GDF6*, and *RALDH1* (coding an enzyme retinaldehyde dehydrogenase), while the ventral retina expresses the homeotic genes *cVax/mVax2* and the gene of the enzyme *RALDH3*. **e** A horizontal band expresses the retinoic acid metabolizing the enzyme *CYP26A1*. In chickens, the expression of *BMP2* dominates and the expression of *RPTPλ* is absent. **f** Photoreceptor and ganglion layers. From Schulte and Bumsted-O’Brien [76]

Likewise, the transcription factor *Otx2* (orthodenticle homeobox) controls photoreceptors, and *Crx* (cone-rod homeobox) is crucial for their differentiation. *Rorβ* (retinoid-related orphan receptor-beta) leads to rods via *Nrl* (neural retina-specific leucine zipper protein) and *Nr2e3* (nuclear receptor subfamily 2, group E, member 3). The *M*-cones express *Trβ2* (in mice). For the bipolars, the genes are *bHLH Mash1* and *Math3*, and the homeobox genes *Chx10* and *Isl1*.

It is also important to study the genetic control of migration which allows cells to reach their laminar destinations, but also the genetic control of the growth of

dendritic trees which overlap more and more during development (the overlap factor goes from 1 to 10). This control of the patterns of dendritic branching is crucial from the functional point of view.

Figure 5.47 due to Dorothea Schulte and Keely Bumsted-O'Brien [76] is a schematic view of gene expression in the retina of vertebrates during embryogenesis. At the beginning of this process, the future optic nerve and the neural retina compartmentalize by expressing: (i) the former, the transcription factor *Pax2* (paired box homeotic gene) and the homeodomain *Vax1* (ventral anterior homeobox); (ii) the latter, *Pax6* (also expressed by the pigment epithelium *PE* and the optic vesicle *lv*), *Pax6*, and *Vax2* giving the ventral retina and *Pax6* and *Tbx* giving the dorsal retina. During formation of the anterior–posterior axis, the anterior part *a* expresses *Foxg1* (forkhead box *G1*) and the posterior part *p* expresses *Foxd1* (forkhead box *D1*), the medial line expressing neither of these two genes. At a later stage, following invagination of the optical vesicle, the patterning of the *a/p* axis is brought about by expression of the homeotic genes *SOH1* (sensory organ homeobox, especially in chickens) and *Hmx1*. During formation of the dorsal–ventral axis *d/v* of the retina, the dorsal retina expresses the transcription factors, *Tbx*, *COUP-TF11* (*COUP* = chicken ovalbumin upstream promoter, in chickens, and *NR2F2* = nuclear receptor subfamily 2, group *F*, member 2 in humans), *GDF6* (growth differentiation factor), *RALDH1* (coding the enzyme retinaldehyde dehydrogenase), while the ventral retina expresses the homeotic genes *cVax/mVax2* (*c* = chicken, *m* = mouse) and the gene for the enzyme *RALDH3*. In Fig. 5.47e, a horizontal band expresses the gene *CYP26A1* encoding the monooxygenase protein Cytochrome P450 26A1 (Family 26, Subfamily A, Member 1 of a superfamily of enzymes) which participates in metabolizing the retinoic acid-4-hydroxylase. In chickens, the expression of *BMP2*, bone morphogenetic protein, dominates, and the expression of *RPTPλ* (receptor-type protein tyrosine phosphatase) is absent.

5.12.5 Retinotopy and Neurogenesis of Visual Pathways

We shall now say a few words about the development of the optical pathways, which is very precise and has extremely tightly controlled geometry. The selective guidance controlling the formation of the optic chiasm and the projections onto the superior colliculus and the lateral geniculate nucleus (LGN) is particularly remarkable. It is made possible by the presence along the optical tract of what David Feldheim and Dennis O'Leary [77, p. 3] call a 'graded distribution of topographic guidance molecules along its two axes', each ganglion cell having

[...] a unique profile of receptors for those molecules that would result in a position-dependent, differential response to them by retinal ganglion cells axons.

We have seen that there are growth cones at the end of the axons. Through the mechanisms for signalling spatial position information carried by attractive or repulsive proteins, these growth cones allow the axons to reach their targets and group to-

gether into bundles. The function of molecular guidance control is to allow cells to carry out *transduction* of extracellular signals through a series of changes in their cell morphology and their cytoskeleton. The ganglion cells of the ventral–temporal retina project onto the ipsilateral optic tract and so do not cross the optic chiasm. The region of the chiasm exerts a repulsion through cells expressing the *ephrin-B2* ligand decoded by the *EphB1* receptors on the growth cones of the axons in the temporal retina, these receptors being restricted for the ventral–temporal retina. This expression is controlled (in the mouse) by the transcription factor *Zic2*, a gene which is itself regulated by the *Foxd1* gene.

But perhaps one of the most fascinating morphogenetic processes concerns the control of *spatial positioning in the retinotopic maps*. This requires extremely accurate targeting, and axon guidance is not enough here. A more precise genetic control is needed, along with further refinement of this control through the activity arising from interactions between axons that have become juxtaposed by guidance. The *spontaneous* activity of neural networks turns out to be crucial, because it generates waves that correlate neighbouring dendritic trees. This process exemplifies in a quite spectacular way these links between positional information and genetic expression which occurs during embryogenesis, and it has led to heated debate between geneticists and proponents of a ‘structuralist’ approach such as Conrad Hal Waddington, Alan Turing, René Thom, Brian Goodwin, and Lewis Wolpert.¹⁰

A key mechanism is axon *chemotaxis*, which occurs through *gradients* of chemo-attractive and chemorepulsive molecules like the ephrins (see above) with their *Eph* receptors, the semaphorins for the axons of the corpus callosum, *Netrin 1* for the thalamocortical pathways, *slits* for the optic chiasm and the corpus callosum. The idea of chemoaffinity introduced by Roger Sperry in the 1950s asserts that there is an address system, i.e. molecular tags, distributed in complementary gradients on the axons and their targets, analogous to a lock and key, and which determine the specificity of the axon connections in topographic maps. We only began to understand the exact molecular nature of these gradients in the 1990s, with the discovery of the *Eph/ephrins* which specify a refined coarse-grained topographic projection, as already noted, through waves of spontaneous activity. The reader is referred to Todd McLaughlin et al. [79].

The acronym *Eph* comes from erythropoietin-producing human hepatocellular carcinoma cell line, and *ephrin* from Eph family receptor interacting proteins. There are ten *EphA* and six *EphB* in vertebrates. The *EphA* bind with six *ephrin-A* ligands and the *EphB* with three *ephrin-B* ligands. Figure 5.48 due to Scicolone et al. [80] is a schematic representation of the structural domains of the *EphA* and the *ephrin-A*, showing how these *EphA* and *ephrin-A* bind via the ligand-binding domain (LBD) by interdigitating their end amine group, which produces the associated two-way signalling. The interested reader is referred to the primer [69] (already cited) by Alex Kolodkin and Marc Tessier-Lavigne entitled *Mechanisms and Molecules of Neuronal Wiring*.

¹⁰For an introduction to structuralism in biology from Waddington to Thom, see [78].




Fig. 5.48 **a** Schematic view of the structural domains of the *EphA* and *ephrin-A*. **b** Binding *EphA/ephrin-A*. *LBD* = ligand-binding domain, *CysRD* = cysteine-rich domain, *FLR* = fibronectin-like repeats, *JD* = juxtamembrane domain, *TPoS* = tyrosine phosphorylation sites, *TyRK* = tyrosine kinase domain, *SAM* = sterile α motif, *PDZ-BD* = post-synaptic density protein/Drosophila disc large tumour suppressor/zonula occludens-1 protein (*PDZ*)-binding domain, *GPI* = glycosylphosphatidylinositol, *P* = tyrosine phosphorylated sites. From Scicolone et al. [80]

The formation of positional identities has received much attention, for example, in the optic tectum of non-mammalian vertebrates such as the chicken or the frog, or the mammalian superior colliculus and LGN. *ephrin-A* ligands (*ephrin-A2/A5/A6* in the chicken, *ephrin-A6* in the mouse) have been found in the neural membrane. These are expressed in an increasing rostral–caudal gradient and are detected via a decreasing temporal–nasal counter-gradient (i.e. opposite gradient) of the associated *EphA* receptors (*EphA3* in the chicken, *EphA5* in the mouse). *ephrin-A* and *EphA* act bifunctionally (as both attractors and repellers) and bidirectionally. There are also orthogonal tectal and retinal gradients mediated by *ephrin-B1/B2* and their receptors *EphB1/B2*, which control the geometry of the projection of the dorsal–ventral retina on the lateromedial axis of the tectum. This defines the local address system of the targets.

Figure 5.49 from Reese [71] shows this gradient-controlled dynamics in a schematic way for the case of a retina \rightarrow tectum projection (the retina \rightarrow colliculus projection is similar). It shows how the increasing naso \rightarrow temporal (*N* \rightarrow *T*) *EphA* and dorso \rightarrow ventral (*D* \rightarrow *V*) *EphB* gradients of the ganglion cell layer, and also the opposite retinal gradients of *ephrin-A* and *-B*, allow the axon growth cone




Fig. 5.49 The projection retina \rightarrow tectum. Increasing naso \nearrow temporal ($N \nearrow T$) *EphA* and dorso \nearrow ventral ($D \nearrow V$) *EphB* gradients of the ganglion cell layer allows the axon growth cone of these cells to reach a precise position in the tectum defined by increasing complementary gradients of rostro-caudal ($R \nearrow C$) *ephrin-A* and lateromedial ($L \nearrow M$) *ephrin-B*. This leads to a retinotopic map $N \rightarrow C$, $T \rightarrow R$, $D \rightarrow L$, $V \rightarrow M$. From Reese [71]

of these cells to reach a precise position in the tectum (or the colliculus) defined by the increasing complementary gradients of rostro-caudal ($R \rightarrow C$, or anterior-posterior for the colliculus) *ephrin-A* and lateromedial ($L \rightarrow M$, or ventral-dorsal for the colliculus) *ephrin-B* with opposite *EphA* and *EphB* gradients. This establishes a retinotopic map $N \rightarrow C$, $T \rightarrow R$, $D \rightarrow L$, $V \rightarrow M$. For example, *ephrin-A2* will be attracting at low concentration in R and repulsive at high concentration in C .

Figure 5.50 from Luo and Flanagan [81] gives further details concerning Fig. 5.49. We find the previous ephrin map, but also the role of the *Wnt3* (wingless-type MMTV integration site family, member 3) gradient $V \nearrow D$ with its retinal receptors *Ryk* (receptor tyrosine kinase) with gradient $D \nearrow V$, a tyrosine kinase receptor mediating repulsion, and *Fz* (frizzled) mediating attraction for low concentrations of *Wnt3*. *En-2* (engrailed) is a homeotic transcription factor for the guidance which attracts nasal axons and repels temporal axons. Its retinal receptors have not yet been clearly identified. The *En* proteins are expressed by an increasing $A \nearrow P$ gradient and regulate cell fate. The figure also shows the mechanism whereby the retinotopic map is produced by guidance. The region N of the retina projects onto the region P of the tectum, while the region T projects onto region A , this map being defined very precisely by the concentrations of *ephrin-A* and *-B* and *Wnt3*. Low concentrations of *ephrin-A* induce attractive positive effects (in green), while high concentrations induce repulsive negative effects (in red). If we consider a point on the axis $N \rightarrow T$, it will project onto the axis $A \rightarrow P$, and it is the *balance* between these two opposing effects that specifies the exact target position on $A \rightarrow P$. Likewise for the axis $D \rightarrow V$. The *EphA/B* and *ephrin-A/-B* concentrations thus specify a map $(x, y) \rightarrow (f(x, y), g(x, y))$ from genuine coordinate axes.

Finally, the figure shows (for the mouse) how the projection $N/T \rightarrow A/P$ is constituted from arborization of axons and transforms when the genetic control is altered in some way. For the mutant *ephrin-A2/-A3/-A5* knockout mice, the N and T axons have an adequate arborization but project onto *several* positions. In knockout

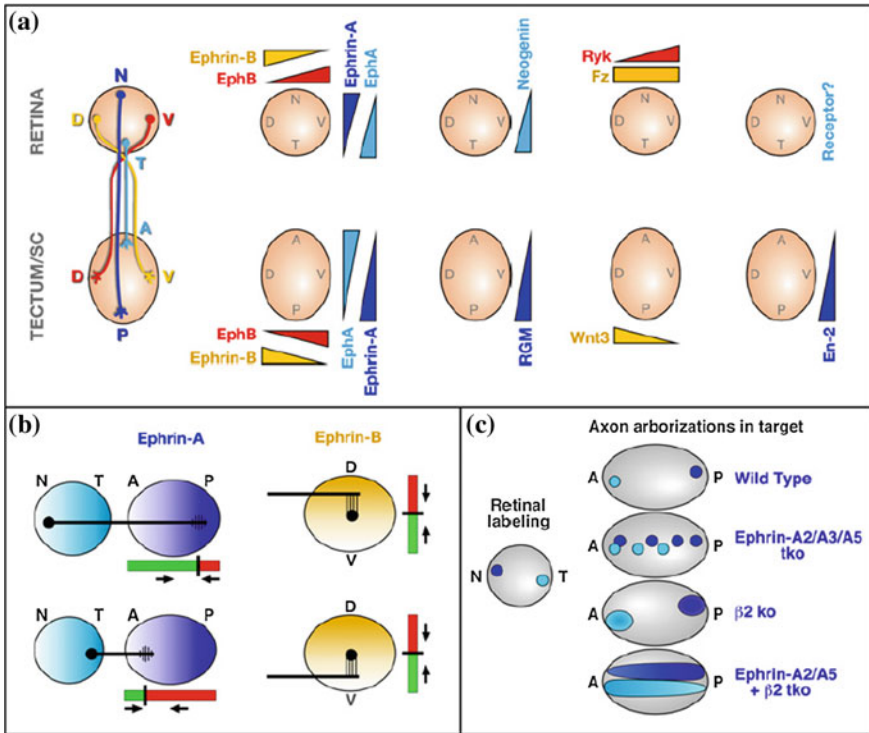


Fig. 5.50 **a** Map of the ephrins in Fig. 5.49. **b** Role of the *Wnt3* $V \nearrow D$ gradient with its retinal receptors *Ryk* $D \nearrow V$ and *Fz*. *En-2* attracts the nasal axons and repels the temporal axons. In the mechanism underlying the formation of the retinotopic map, region *N* of the retina projects onto region *P*, and region *T* onto region *A*, this map being very precisely specified by the concentrations of *ephrin-A* and *-B* and *Wnt3*. Low concentrations of *ephrin-A* induce attractive positive effects (green), while high concentrations induce repulsive negative effects (red). A point on the $N \rightarrow T$ axis projects onto the $A \rightarrow P$ axis, and it is the balance between these two opposing effects which specifies the exact target position on $A \rightarrow P$. Likewise for the $D \rightarrow V$ axis. **c** For the mouse, the way the projection $N/T \rightarrow A/P$ is built up by arborization of the axons and transformed when the genetic control is altered in some way. *Line 1* Adequate natural projection (wild type). *Line 2* For *ephrin-A2/-A3/-A5* knockout mutants, the axons *N* and *T* have adequate arborization but project onto several positions. *Line 3* In knockout mice for the $\beta 2$ subunit of the acetylcholine nicotinic receptor, activities of the ganglion cells become decorrelated and the arborizations are well positioned, but more diffuse. *Line 4* When there are mutations for both *ephrin-A2/-A5* and $\beta 2$, diffusion is extreme and covers the whole axis. From Luo and Flanagan [81]

mice for the β -2 subunit of the acetylcholine nicotinic receptor, activities of the ganglion cells become decorrelated and the arborizations are well positioned but much more spread out. When there are mutations both for *ephrin-A2/-A5* and β -2, diffusion is extreme and covers the whole axis.

As summed up by Benjamin Reese [71, p. 624]:

Complementary gradients across both retinal and tectal axes clearly participate in retinotopic map formation by providing retinal growth cones with the means to decipher positional identity upon the tectal surface.

David Feldheim and Dennis O’Leary also note [77, p. 18]:

Along each mapping axis dual gradients of axon guidance activity combined with axon competition and correlated neural activity are used to specify axonal branch points, guide the branches, and refine their terminations.

In a word, we may say once again that molecular addresses are able to realize a *coordinate system biochemically*. Liqun Luo and John Flanagan stress this point [81, p. 289]:

Positional information gradients provide coordinate systems that can be interpreted to regulate essentially any cell function, depending on the receptor systems used to decode this positional information.

5.12.6 Dynamical Models of Neural Guidance

We have seen just how complex and precise the genetic control of cell motility and directional guidance can be. On the basis of the available experimental results, we can begin to build dynamical models. The first precise models of morphogenesis go back to the famous paper by Alan Turing in 1952 entitled *The Chemical Basis of Morphogenesis* [82], which introduced the *reaction–diffusion equations*. Since then, these have seen considerable development. The central idea is that:

A system of chemical substances, called morphogens, reacting together and diffusing through a tissue, is adequate to account for the main phenomena of morphogenesis.

The following year (1953), Turing formulated this in a striking way in his last paper:

It was suggested in Turing (1952) that this might be the main means by which the chemical information contained in the genes was converted into a geometrical form.

As noted by one of Turing’s colleagues, the botanist Claude Wilson Wardlaw [83, p. 40]:

A localized accumulation of gene-determined substances may be an essential prior condition [for cell differentiation].

The guidance models discussed here are diffusion models, such as those of Anna Cai, Kerry Landman, K. and Barry Hughes, B. [84]. To simplify, we assume that there is only one spatial coordinate x . Let $u(x, t)$ be the cell population density at x at time t . The simplest equation is a pure diffusion equation $\partial u / \partial t = D \partial^2 u / \partial x^2$, where D is a diffusion constant. If there are interactions between the cells, their motility will depend on their density u and the coefficient D will depend on u . Then we have a more subtle diffusion equation of the form

$$\frac{\partial u}{\partial t} = \frac{\partial}{\partial x} \left(D(u) \frac{\partial u}{\partial x} \right) = D(u) \frac{\partial^2 u}{\partial x^2} + \frac{\partial D}{\partial u} \left(\frac{\partial u}{\partial x} \right)^2 .$$

If we add a concentration $s(x, t)$ of morphogens with a signalling function, we must then introduce assumptions about the way the cells detect the signal molecules. If the cells detect s locally near their position x , the equation might take the form

$$\frac{\partial u}{\partial t} = \frac{\partial}{\partial x} \left(D(u, s) \frac{\partial u}{\partial x} \right) = D(u, s) \frac{\partial^2 u}{\partial x^2} + \frac{\partial D}{\partial u} \left(\frac{\partial u}{\partial x} \right)^2 + \frac{\partial D}{\partial s} \frac{\partial s}{\partial x} \frac{\partial u}{\partial x} ,$$

where the diffusion coefficient D now also depends on s . Another local model corresponds to the case where the cells detect s only at their exact position x , through a function $I(s)$ of the concentration $s(x, t)$:

$$\frac{\partial u}{\partial t} = \frac{\partial}{\partial x} \left[D(u) \frac{\partial}{\partial x} [I(s)u] \right] = \frac{\partial}{\partial x} \left[D(u) I(s) \frac{\partial u}{\partial x} + D(u) \frac{\partial I}{\partial s} \frac{\partial s}{\partial x} u \right] .$$

A further refined model would include chemotaxis with the diffusion, whereby the cells could detect changes in s in the vicinity of their position. We introduce a chemotactic sensitivity $\chi(s)$ with the rule that if $\chi(s) < 0$, the cells climb the gradient $\nabla(s) = \partial s / \partial x$ (attraction), but if $\chi(s) > 0$, the cells move down the gradient $\nabla(s) = \partial s / \partial x$ (repulsion). A range of numerical simulations of these different models can be found in [84].

References

1. Eysel, U.: Turning a corner in vision research. *Nature* **399**, 641–644 (1999)
2. Alexander, D.M., Bourke, P.D., Sheridan, P., Konstantatos, O., Wright, J.J.: Intrinsic connections in tree shrew V1 imply a global to local mapping. *Vision Res.* **44**, 857–876 (2004)
3. Ts'o, D.Y., Gilbert, C.D., Wiesel, T.N.: Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* **6**(4), 1160–1170 (1986)
4. Das, A., Gilbert, C.D.: Long range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex. *Nature* **375**, 780–784 (1995)
5. Malach, R., Amir, Y., Harel, M., Grinvald, A.: Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc. Nat. Acad. Sci.* **90**, 10469–10473 (1993)
6. Bosking, W.H., Zhang, Y., Schofield, B., Fitzpatrick, D.: Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J. Neurosci.* **17**(6), 2112–2127 (1997)
7. Mitchison, G., Crick, F.: Long axons within the striate cortex: their distribution, orientation, and patterns of connections. *Proc. Nat. Acad. Sci.* **79**, 3661–3665 (1982)
8. Rockland, K.S., Lund, J.S.: Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science* **215**, 532–534 (1982)
9. Rockland, K.S., Lund, J.S.: Intrinsic laminar lattice connections in the primate visual cortex. *J. Comp. Neurol.* **216**(3), 303–318 (1993)

10. Grinvald, A., Lieke, E.E., Frostig, R.D., Hildesheim, R.: Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *J. Neurosci.* **14**(5), 2545–2568 (1994)
11. Field, D.J., Hayes, A., Hess, R.F.: Contour integration by the human visual system: evidence for a local ‘association field’. *Vision Res.* **33**(2), 173–193 (1993)
12. Kapadia, M.K., Ito, M., Gilbert, C.D., Westheimer, G.: Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron* **15**, 843–856 (1995)
13. Kapadia, M.K., Westheimer, G., Gilbert, C.D.: Dynamics of spatial summation in primary visual cortex of alert monkey. *Proc. Nat. Acad. Sci.* **96**(21), 12073–12078 (1999)
14. Kovács, I., Julesz, B.: A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation. *Proc. Nat. Acad. Sci.* **90**, 7495–7497 (1993)
15. Sajda, P., Han, F.: Perceptual salience as novelty detection in cortical pinwheel space. In: *Proceedings of the First International IEEE EMBS Conference on Neural Engineering* (2003)
16. Desolneux, A., Moisan, L., Morel, J.-M.: *From Gestalt Theory to Image Analysis: A Probabilistic Approach*, *Interdisciplinary Applied Mathematics*, vol. 34. Springer, Heidelberg (2008)
17. Kourtzi, Z., Tolias, A.S., Altmann, C.F., Augath, M., Logothetis, N.K.: Integration of local features into global shapes: monkey and human fMRI studies. *Neuron* **37**, 333–346 (2003)
18. Georges, S., Seriès, P., Frégnac, Y., Lorenceau, J.: Orientation dependent modulation of apparent speed: psychophysical evidence. *Vision Res.* **42**, 2757–2772 (2002)
19. Georges, S., Seriès, P., Lorenceau, J.: Contrast dependence of high-speed apparent motion. <http://www.perceptionweb.com/ecvp00/0300.html> (2000)
20. Frégnac, Y., Baudot, P., Chavane, F., Lorenceau, J., Marre, O., Monier, C., Pananceau, M., Carelli, P., Sadoc, G.: Multiscale functional imaging in V1 and cortical correlates of apparent motion. In: Masson, G.S. Ilg, U.J. (eds.) *Dynamics of Visual Motion Processing*. Neuronal, Behavioral, and Computational Approaches, pp. 73–94. Springer, Berlin (2010)
21. Seriès, P., Georges, S., Lorenceau, J., Frégnac, Y.: Orientation dependent modulation of apparent speed: a model based on the dynamics of feed-forward and horizontal connectivity in V1 cortex. *Vision Res.* **42**, 2781–2797 (2002)
22. Alais, D., Lorenceau, J., Arrighi, R., Cass, J.: Contour interactions between pairs of Gabors engaged in binocular rivalry reveal a map of the association field. *Vision Res.* **46**, 1473–1487 (2006)
23. Paradis, A.-L., Morel, S., Seriès, P., Lorenceau, J.: Speeding up the brain: when spatial facilitation translates into latency shortening. *Frontiers Human Neurosci.* **6**, 330 (2012)
24. Petitot, J., Tondut, Y.: *Géométrie de contact et champ d’association dans le cortex visuel*. CREA reports # 9725, École Polytechnique, Paris (1997)
25. Lee, T.S., Nguyen, M.: Dynamics of subjective contour formation in the early visual cortex. *Proc. Nat. Acad. Sci.* **98**(4), 1907–1911 (2001)
26. Polat, U., Sagi, D.: Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiment. *Vision Res.* **33**(7), 993–999 (1993)
27. Gilbert, C.D., Das, A., Ito, M., Kapadia, M., Westheimer, G.: Spatial integration and cortical dynamics. *Proc. Nat. Acad. Sci.* **93**, 615–622 (1996)
28. Zucker, S.W., David, C., Dobbins, A., Iverson, L.: The organization of curve detection: coarse tangent fields and fine spline covering. In: *Proceedings of the COST-13 Workshop (European Cooperation in Science and Technology)*, Bonas, France, 1988. Republished in Simon, J.C. (ed.) *From Pixels to Features*, pp. 75–90. North Holland, New York (1989)
29. Parent, P., Zucker, S.W.: Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(8), 823–839 (1989)
30. Grossberg, S., Mingolla, E.: Neural dynamics of form perception: boundary completion, illusory figures and neon color spreading. *Psychol. Rev.* **92**, 173–211 (1985)
31. Fitzpatrick, D.: Seeing beyond the receptive field in primary visual cortex. *Curr. Opin. Neurobiol.* **10**, 438–443 (2000)
32. Wright, J.J., Alexander, D.M., Bourke, P.D.: Contribution of lateral interactions in V1 to organization of response properties. *Vision Res.* **46**, 2703–2720 (2006)

33. Jancke, D., Chavane, F., Naaman, S., Grinvald, A.: Imaging cortical correlates of illusion in early visual cortex. *Nature* **428**, 423–426 (2004)
34. Rangan, A.V., Cai, D., McLaughlin, D.W.: Modeling the spatiotemporal cortical activity associated with the line-motion illusion in primary visual cortex. *Proc. Nat. Acad. Sci.* **102**(52), 18793–18800 (2005)
35. Gilbert, C.D., Wiesel, T.N.: Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J. Neurosci.* **9**(7), 2432–2442 (1989)
36. Gilbert, C.D.: Horizontal integration and cortical dynamics. *Neuron* **9**, 1–13 (1992)
37. Berthoz, A.: *La simplicité*. Odile Jacob, Paris (2009)
38. Petitot, J.: La simplicité de la notion géométrique de jet. In: Berthoz, A., Petit, J.-L. (eds.) *Simplicité-Complexité*. Coll'ège de France, OpenEdition Books, Paris (2014)
39. Ullman, S.: Filling in the gaps: the shape of subjective contours and a model for their generation. *Biol. Cybern.* **25**, 1–6 (1976)
40. Petitot, J., (with a contribution of Tondut, Y.): Vers une Neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. *Mathématiques, Informatique et Sciences Humaines* **145**, 5–101 (1999)
41. Bryant, R., Griffiths, P.: Reduction for constrained variational problems and $\int \kappa^2 ds/2$. *Am. J. Math.* **108**, 525–570 (1986)
42. Ben-Shahar, O., Zucker, S.: Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Comput.* **16**(3), 445–476 (2004)
43. Ben-Shahar, O., Huggins, P.S., Izo, T., Zucker, S.: Cortical connections and early visual function: intra- and inter-columnar processing. In: Petitot, J., Lorenceau, J. (eds.) *Neurogeometry and Visual Perception*. *J. Physiol. Paris* **97**, 191–208 (2003)
44. Sigman, M., Cecchi, G.A., Gilbert, C.D., Magnasco, M.O.: On a common circle: natural scenes and Gestalt rules. *Proc. Nat. Acad. Sci.* **98**(4), 1935–1940 (2001)
45. Mallat, S., Peyré, G.: Traitements géométriques des images par bandelettes. In: *Journée annuelle de la SMF, 24 June 2006, Mathématiques et Vision*, pp. 39–67 (2006)
46. Le Pennec, E., Mallat, S.: Sparse geometrical image representation with bandelets. In: *IEEE Transactions on Image Processing* (2003)
47. Ts'o, D.Y., Zarella, M., Burkitt, G.: Whither the hypercolumn? *J. Physiol.* **587**(12), 2791–2805 (2009)
48. Roe, A.W.: Modular complexity of area V2 in the macaque monkey. In: Kaas, J.H., Collins, C.E. (eds.) *The Primate Visual System*, Chap. 5, pp. 109–138. CRC Press LLC, Boca Raton, FL (2003)
49. Sit, Y.F., Miikkulainen, R.: Computational prediction on the receptive fields and organization of V2 for shape processing. *Neural Comput.* **21**(3), 762–785 (2009)
50. Lu, H.D., Chen, G., Tanigawa, H., Roe, A.W.: A motion direction map in macaque V2. *Neuron* **68**, 1–12 (2010)
51. Peterhans, E., von der Heydt, R.: Subjective contours: bridging the gap between psychophysics and psychology. *Trends Neurosci.* **14**(3), 112–119 (1991)
52. von der Heydt, R., Peterhans, E.: Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *J. Neurosci.* **9**(5), 1731–1748 (1989)
53. Sheth, B.R., Sharma, J., Rao, S.C., Sur, M.: Orientation maps of subjective contours in visual cortex. *Nature* **274**, 2110–2115 (1996)
54. Zeki, S.: *A Vision of the Brain*. Wiley-Blackwell, Oxford (1993)
55. Thom, R.: *Esquisse d'une Sémiophysique*. InterEditions, Paris (1988)
56. Petitot, J.: 'Le hiatus entre le logique et le morphologique'. *Prédication et perception*. In: Wildgen, W., Brandt, P.A. (eds.) *Semiosis and Catastrophes*. René Thom's Semiotic Heritage, pp. 141–166. Peter Lang, Bern (2010)
57. Shapley, R., Hawken, M.J.: Color in the cortex: single- and double-opponents cells. *Vision Res.* **51**, 701–717 (2011)
58. Tanigawa, H., Lu, H.D., Roe, A.W.: Functional organization for color and orientation in macaque V4. *Nature Neurosci.* **13**, 1542–1548 (2010)

59. Thompson, E., Palacios, A., Varela, F.: Ways of coloring: comparative color vision as a case study in cognitive science. *Behav. Brain Sci.* **15**, 1–74 (1992)
60. Petitot, J.: Morphodynamical enaction: the case of color. *Biol. Res.* (In: Bacigalupo, J., Palacios, A.G. (eds). *A Tribute to Francisco Varela*) **36**(1), 107–112 (2003)
61. Gur, M., Snodderly, D.M.: Direction selectivity in V1 of alert monkeys: evidence for parallel pathways for motion processing. *J. Physiol.* **585**(2), 383–400 (2007)
62. Born, R.T., Bradley, D.C.: Structure and function of visual area MT. *Annu. Rev. Neurosci.* **28**, 157–189 (2005)
63. DeAngelis, G.C., Newsome, W.T.: Organization of disparity-selective neurons in macaque area MT. *J. Neurosci.* **19**(4), 1398–1415 (1999)
64. Tanaka, S., Shinbata, H.: Mathematical model for self-organization of direction columns in the primate middle temporal area. *Bio. Cybern.* **70**, 227–234 (1994)
65. Swindale, N.V., Grinvald, A., Shmuel, A.: The spatial pattern of response magnitude and selectivity for orientation and direction in cat visual cortex. *Cereb. Cortex* **13**(3), 225–235 (2003)
66. Petitot, J.: Éléments de théorie des singularités. http://jean.petitot.pagesperso-orange.fr/ArticlesPDF_new/Petitot_Sing.pdf (1982)
67. Prochiantz, A.: *Machine-esprit*. Odile Jacob, Paris (2001)
68. Chedotal, A., Richards, L.J.: Wiring the Brain: The Biology of Neuronal Guidance. In: Cold Spring Harbor Perspectives in Biology. <http://cshperspectives.cshlp.org> (2010)
69. Kolodkin, A.L., Tessier-Lavigne, M.: Mechanisms and molecules of neuronal wiring: a primer. In: Cold Spring Harbor Perspectives in Biology. <http://cshperspectives.cshlp.org> (2011)
70. GeneCards: Human Gene Compendium. Weizmann Institute. www.genecards.org
71. Reese, B.E.: Development of the retina and optic pathway. *Vision Res.* **51**, 613–632 (2011)
72. Baye, L.M., Link, B.A.: Nuclear migration during retinal development. *Brain Res.* **2008**, 29–36 (1992)
73. Ohsawa, R., Kageyama, R.: Regulation of retinal cell fate specification by multiple transcription factors. *Brain Res.* **1192**, 90–98 (2008)
74. Martins, R.A.P., Pearson, R.A.: Control of cell proliferation by neurotransmitters in the developing vertebrate retina. *Brain Res.* **1192**, 37–60 (2008)
75. Li, S., Mo, Z., Yang, X., Price, S.M., Shen, M.M., Xiang, M.: *Foxn4* controls the genesis of amacrine and horizontal cells by retinal progenitors. *Neuron* **43**, 795–807 (2004)
76. Schulte, D., Bumsted-O’Brien, K.M.: Molecular mechanisms of vertebrate retina development: Implications for ganglion cell and photoreceptor patterning. *Brain Res.* **1192**, 151–164 (2008)
77. Feldheim, D.A., O’Leary, D.D.: Visual map development: bidirectional signaling, bifunctional guidance, molecules and competition. *Cold Spring Harb. Perspect. Biol.* <http://cshperspectives.cshlp.org> (2010)
78. Petitot, J.: *Morphogenèse du Sens*, Presses Universitaires de France, Paris, 1985. English Translation: Manjali, F. *Morphogenesis of Meaning*. Bern, Peter Lang (2003)
79. McLaughlin, T., Torborg, C.L., Feller, M.B., O’Leary, D.D.M.: Retinotopic map refinement requires spontaneous retinal wave during a brief period of development. *Neuron* **40**, 1147–1160 (2003)
80. Scicolone, G., Ortalli, A.L., Carri, N.G.: Key roles of Ephs and ephrins in retinotectal topographic map formation. *Brain Res. Bull.* **79**, 227–247 (2009)
81. Luo, L., Flanagan, J.G.: Development of continuous and discrete neural maps. *Neuron* **56**, 284–300 (2007)
82. Turing, A.M.: The chemical basis of morphogenesis. *Philos. Trans. R. Soc. B, Bio. Sci.* **237**(641), 37–72 (1952)
83. Turing, A.M., Wardlaw, C.W.: A diffusion reaction theory of morphogenesis in plants. *New Phytol.* **52**, 40–47. Also in *Collected Works of A.M. Turing*, P.T. Saunders, Amsterdam **1953**, 37–47 (1952)
84. Cai, A.Q., Landman, K.A., Hughes, B.D.: Modelling directional guidance and motility regulation in cell migration. *Bull. Math. Biol.* **68**, 25–52 (2006)

Chapter 6

Transition to Volume II

6.1 Introduction

In this first volume, we have focused on experimental data and the basics of neurogeometric modelling. We have discussed receptive fields and profiles of visual neurons in relation to wavelet theory, and we have interpreted them as a way of optimizing the compression of natural images. We have also described the two components of the functional architecture of the $V1$ area and other retinotopic areas, and in particular, the pinwheel structure of $V1$. We have seen that this ‘vertical’ retinogeniculo-cortical functional architecture of $V1$ implements a discrete approximation of the fibre bundle of 1-jets of curves in the visual plane. We have analyzed $V1$ as a 2D orientation field (orientation map) whose singularities are the centres of the pinwheels. More precisely, we have shown how to treat it as a phase field which is a superposition of solutions of the Helmholtz equation. Then, with reference to the experimental data, we have presented the alternative model of blow-ups which converge to the 1-jet bundle when the mesh of the pinwheel lattice tends to 0.

We have also discussed the relationship between the orientation maps and other maps, such as those for direction, ocular dominance, phase, spatial frequency, and colour. Their independence can be expressed in the form of a transversality principle.

We then investigated the second component of the functional architecture of $V1$ and the other areas of the visual cortex, namely the ‘horizontal’ intracortical connections underlying the association field. It is this component that forms the basis for our neurogeometric models. It implements the contact structure of the fibre bundle \mathbb{V}_J which is isomorphic to the polarized Heisenberg group and which is defined by the field of kernel planes of the differential 1-form $\omega = dy - p dx$ invariant under left translations. We use a sub-Riemannian metric on the contact structure to define geodesics that can serve as models for long-range illusory contours.

Finally, we discussed some properties of the functional architecture of areas $V2$, $V4$ (for colour), and $V5$ or MT (for motion) and reviewed the genetic control of their morphogenesis (neurogenesis and axon guidance).

To conclude this first volume, we shall now give a brief preview of the themes to be covered in the second volume. This is just an outline and the necessary mathematics will be developed there.

6.2 Geodesics of the \mathbb{V}_J Model

We have seen in Sect. 5.5 of Chap. 5 how to interpret illusory contours as sub-Riemannian geodesics of the \mathbb{V}_J model. Let us now give an idea of the computation of such geodesics and emphasize that, as the metric is *sub*-Riemannian, they are very different from Riemannian ones, even at the infinitesimal level.

As the distribution \mathcal{K} of the contact planes K_v is bracket generating and satisfies the Hörmander condition, a celebrated theorem due to Wei-Liang Chow tells us that every pair of points (v, v') of \mathbb{V}_J can be connected by an integral curve of \mathcal{K} . If \mathcal{K} is endowed with a sub-Riemannian metric, we can compute the length of such integral curves and look for geodesics, which are integral curves of minimal length. The problem of computing geodesics is quite difficult to solve. The study by Roger Brockett [1] entitled *Control Theory and Singular Riemannian Geometry* is a classic reference from 1981. Other excellent references are the books by Richard Montgomery [2] and Robert Strichartz [3], and also the papers by Zhong Ge [4] and Ursula Hammenstädt [5]. One of the main difficulties is that, in contrast to the Riemannian case, there can exist ‘abnormal’ geodesics, that is, geodesics which do not satisfy the differential equations canonically associated with the geodesic variational problem. Fortunately, we shall not encounter this ‘abnormality’ since our models, even though non-trivial, remain rather elementary.

Richard Beals, Bernard Gaveau, and Peter Greiner, who solved the geodesic problem for the (non-polarized) Heisenberg group with *explicit* formulas, emphasized [6, p. 634] “how complicated a control problem can become, even in the simplest situation”. This was a new mathematical result since, in 1977, Bernard Gaveau claimed [7, p. 114] that the variational problem of minimizing “the energy of a curve in the base manifold under the Lagrange condition that its lifting is given in the fibre bundle” seemed “not yet [to have been] studied”.

We will adapt Beals, Gaveau, and Greiner’s computations to the *polarized* Heisenberg group $\mathbb{V}_J = J^1(\mathbb{R}, \mathbb{R})$ (see Sect. 5.4.4 of Chap. 5) with coordinates $(x, y, p = \tan(\theta))$, product $(x, y, p) \cdot (x', y', p') = (x + x', y + y' + px', p + p')$, and contact planes generated by

$$X_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y} = (1, p, 0) \quad , \quad X_2 = \frac{\partial}{\partial p} = (0, 0, 1) \quad ,$$

with Lie bracket $[X_1, X_2] = -X_3 = (0, -1, 0) = -\partial/\partial y$. Following the approach used by Agrachev and Sachkov [8], we formulate the geodesic problem as a *control problem*. If $\Gamma = v(s)$ is a smooth parametrized curve in \mathbb{V}_J , to say that it is an integral curve of the contact structure is to say that $\dot{v}(s) = u_1 X_1 + u_2 X_2$ for appropriate controls

u_1 and u_2 , or in other words that $\dot{x} = u_1$, $\dot{y} = pu_1$, $\dot{p} = u_2$, the integrability condition $\dot{y}/\dot{x} = p$ being automatically satisfied. To find the geodesics for the chosen sub-Riemannian metric SR with scalar product $\langle \cdot, \cdot \rangle_{SR}$ and norm $\|\cdot\|_{SR}$, we minimize the Lagrangian given by the kinetic energy $L = \|\dot{v}\|_{SR}^2/2$ along such curves. L is defined on the tangent bundle $T\mathbb{V}_J$. Using the Legendre transform, it can be transformed into the Hamiltonian (where $\langle \cdot, \cdot \rangle$ is the natural pairing between 1-forms and tangent vectors)

$$\begin{aligned} h(v, \varpi) &= \langle \varpi, \dot{v} \rangle - \frac{1}{2} \|\dot{v}\|_{SR}^2 \\ &= \varpi (u_1 X_1 + u_2 X_2) - \frac{1}{2} \|u_1 X_1 + u_2 X_2\|_{SR}^2, \end{aligned}$$

defined on the cotangent bundle $T^*\mathbb{V}_J$. If $\varpi = \xi^* dx + \eta^* dy + \pi^* dp = (\xi^*, \eta^*, \pi^*)$ is a 1-form on \mathbb{V}_J , then

$$h(v, \varpi) = \xi^* u_1 + \eta^* u_1 p + \pi^* u_2 - \frac{1}{2} \left(u_1^2 \|X_1\|_{SR}^2 + 2u_1 u_2 \langle X_1, X_2 \rangle_{SR} + u_2^2 \|X_2\|_{SR}^2 \right).$$

It is natural to choose a left-invariant metric, namely the sub-Riemannian metric SR_J making $\{X_1, X_2\}$ an orthonormal basis of the contact plane K_v , since $\{X_1, X_2\}$ is the left-invariant basis translating the standard Euclidean orthonormal basis of K_0 . This invariant metric is not the Euclidean metric $\langle \cdot, \cdot \rangle_E$, $\|\cdot\|_E$ since, due to non-holonomy, the Euclidean metric is not left-invariant. Incidentally, even if $\|X_2\|_E = 1$ and $\langle X_1, X_2 \rangle_E = 0$, we have $\|X_1\|_E = 1 + p^2 \neq 1$ if $p \neq 0$: it is only on the (x, y) plane $p = 0$ that the two metrics are the same. If we choose SR_J , then we have $\|X_1\|_{SR_J} = \|X_2\|_{SR_J} = 1$, $\langle X_1, X_2 \rangle_{SR_J} = 0$, and

$$\begin{aligned} h(v, \varpi) &= \xi^* u_1 + \eta^* u_1 p + \pi^* u_2 - \frac{1}{2} (u_1^2 + u_2^2) \\ &= \varpi (u_1 X_1 + u_2 X_2) - \frac{1}{2} (u_1^2 + u_2^2). \end{aligned}$$

One can then apply a fundamental result of control theory called the *Pontryagin maximum principle*. This generalizes the classical method of variational calculus using the Euler–Lagrange equations and Lagrange multipliers (see Agrachev and Gamkrelidze [9]) which we employed in our first synthesis [10]. It says that geodesics are projections on \mathbb{V}_J of the trajectories of the Hamiltonian H with *maximizing* controls $u_{j,\max}$. The maximization conditions are

$$\frac{\partial h}{\partial u_1} = \varpi (X_1) - u_1 = 0, \quad \frac{\partial h}{\partial u_2} = \varpi (X_2) - u_2 = 0,$$

and therefore,

$$\begin{aligned}
 H(v, \varpi) &= u_1 \varpi(X_1) + u_2 \varpi(X_2) - \frac{1}{2} (u_1^2 + u_2^2) = \frac{1}{2} (u_1^2 + u_2^2) \\
 &= \frac{1}{2} (\langle \varpi, X_1 \rangle^2 + \langle \varpi, X_2 \rangle^2) ,
 \end{aligned}$$

and in terms of coordinates

$$H(x, y, p, \xi^*, \eta^*, \pi^*) = \frac{1}{2} [(\xi^* + p\eta^*)^2 + \pi^{*2}] .$$

The geodesic structure implies that the sub-Riemannian *sphere* S (the ends of geodesics starting at 0 and having unit sub-Riemannian length, which are *global* minimizers) and the *wave front* W (the ends of geodesics starting at 0 and having unit sub-Riemannian length, which are not necessarily global minimizers) are rather strange. In particular, the *cut locus* of 0 (that is, the ends of geodesics when they cease to be globally minimizing) and the *conjugate locus* or *caustic* of 0 (that is the singular locus of the exponential map \mathcal{E} integrating geodesics) are rather complex.

The Hamilton equations on $T^*\mathbb{V}_J$ derived from the Hamiltonian H are

$$\left\{ \begin{aligned}
 \dot{x}(s) &= \frac{\partial H}{\partial \xi^*} = \xi^* + p\eta^* , \\
 \dot{y}(s) &= \frac{\partial H}{\partial \eta^*} = p(\xi^* + p\eta^*) = p\dot{x}(s) \text{ (i.e. } p = \dot{y}/\dot{x} = dy/dx, \text{ integrability) ,} \\
 \dot{p}(s) &= \frac{\partial H}{\partial \pi^*} = \pi^* , \\
 \dot{\xi}^*(s) &= -\frac{\partial H}{\partial x} = 0 , \\
 \dot{\eta}^*(s) &= -\frac{\partial H}{\partial y} = 0 , \\
 \dot{\pi}^*(s) &= -\frac{\partial H}{\partial p} = -\eta^*(\xi^* + p\eta^*) = -\eta^*\dot{x}(s) .
 \end{aligned} \right.$$

Given that H is independent of x and y , the derivatives $\dot{\xi}^*(s) = -\partial H/\partial x$ and $\dot{\eta}^*(s) = -\partial H/\partial y$ vanish, and the momenta ξ^* and η^* are therefore constant along any geodesic: $\xi^* = \xi_0^*$ and $\eta^* = \eta_0^*$. This fact simplifies the equations since

$$\dot{x}(s) = \xi_0^* + p\eta_0^* , \quad \dot{y}(s) = p(\xi_0^* + p\eta_0^*) , \quad \dot{\pi}^*(s) = -\eta_0^*(\xi_0^* + p\eta_0^*) .$$

We particularly draw attention to the relations $\ddot{p} = \dot{\pi}^* = -\eta^*\dot{x}$ and $\ddot{x} = \eta^*\dot{p}$, or $(\ddot{x}, \ddot{p}) = \eta^*(\dot{p}, -\dot{x})$, which mean that, in the (x, p) plane, the acceleration is orthogonal to the velocity, and geodesics are circles whose radius increases when η_0^* decreases (in the limit $\eta_0^* = 0$, the circle becomes a straight line). Incidentally, $H(x, y, p, \xi^*, \eta^*, \pi^*) = (\dot{x}^2 + \dot{p}^2)/2$, since by construction, the Hamiltonian H is the kinetic energy of the projection of the trajectories on the (x, p) plane.

The $(x, p) = z$ part of the geodesics from 0 to $(x_1 = x(\tau), y_1 = y(\tau), p_1 = p(\tau))$ is given by the formulas:

$$\begin{cases} x(s) = \frac{\sin\left(\frac{s}{2}\eta_0^*\right)}{\sin\left(\frac{\tau}{2}\eta_0^*\right)} \left[\cos\left(\frac{\tau-s}{2}\eta_0^*\right) x_1 - \sin\left(\frac{\tau-s}{2}\eta_0^*\right) p_1 \right], \\ p(s) = \frac{\sin\left(\frac{s}{2}\eta_0^*\right)}{\sin\left(\frac{\tau}{2}\eta_0^*\right)} \left[\sin\left(\frac{\tau-s}{2}\eta_0^*\right) x_1 + \cos\left(\frac{\tau-s}{2}\eta_0^*\right) p_1 \right], \end{cases}$$

which are effectively the equations of a circle

$$x^2 + p^2 - x \left[x_1 + p_1 \cot\left(\frac{\eta_0^*\tau}{2}\right) \right] - p \left[p_1 - x_1 \cot\left(\frac{\eta_0^*\tau}{2}\right) \right] = 0,$$

passing through 0 and (x_1, p_1) , with center

$$x_c = \frac{1}{2} \left[x_1 + p_1 \cot\left(\frac{\eta_0^*\tau}{2}\right) \right], \quad y_c = \frac{1}{2} \left[p_1 - x_1 \cot\left(\frac{\eta_0^*\tau}{2}\right) \right],$$

and radius

$$r^2 = \frac{1}{4} (x_1^2 + p_1^2) \left[1 + \cot\left(\frac{\eta_0^*\tau}{2}\right) \right] = \frac{1}{4 \sin^2(\eta_0^*\tau/2)} |z_1|^2.$$

We then check that the constant value of the Hamiltonian along a trajectory is

$$H_0 = \frac{\eta_0^{*2}}{8 \sin^2(\eta_0^*\tau/2)} |z_1|^2 = \frac{\eta_0^{*2}}{2} r^2.$$

For $y(s)$, the calculations are more involved. We get

$$\begin{aligned} y(s) = \frac{1}{8(\cos(\eta_0^*\tau) - 1)} & \left[-2\eta_0^*s(x_1^2 + p_1^2) - 4x_1p_1 \cos(\eta_0^*(s - \tau)) \right. \\ & + 2(x_1^2 - p_1^2) \sin(\eta_0^*(s - \tau)) \\ & + 2x_1p_1 \cos(\eta_0^*(2s - \tau)) - (x_1^2 - p_1^2) \sin(\eta_0^*(2s - \tau)) \\ & + 2x_1p_1 \cos(\eta_0^*\tau) + (x_1^2 - p_1^2) \sin(\eta_0^*\tau) \\ & \left. + 2(x_1^2 + p_1^2) \sin(\eta_0^*s) \right]. \end{aligned}$$

In terms of ξ_0^* , π_0^* , η_0^* , and τ , this becomes

$$y(s) = \xi_0^{*2} \frac{2\eta_0^*s + \sin(2\eta_0^*s)}{4\eta_0^{*2}} - \xi_0^{*2} \frac{\sin(\eta_0^*s)}{\eta_0^{*2}} + \xi_0^*\pi_0^* \frac{\sin^2(\eta_0^*s)}{\eta_0^{*2}} - \xi_0^*\pi_0^* \frac{1 - \cos(\eta_0^*s)}{\eta_0^{*2}} + \pi_0^{*2} \frac{2\eta_0^*s - \sin(2\eta_0^*s)}{4\eta_0^{*2}} .$$

The key point is that these equations explain the origin of the striking *multiplicity* of sub-Riemannian geodesics connecting two points. Indeed, if we compute $y_1 = y(\tau)$, we find

$$y_1 = \frac{1}{2}x_1p_1 + \frac{x_1^2 + p_1^2}{4} \left[\frac{\eta_0^*\tau/2}{\sin^2(\eta_0^*\tau/2)} - \frac{\cos(\eta_0^*\tau/2)}{\sin(\eta_0^*\tau/2)} \right] .$$

If we introduce the new variable $\varphi = \eta_0^*\tau/2$, we see that we must solve the equation

$$4 \left(y_1 - \frac{1}{2}x_1p_1 \right) = \mu(\varphi) |z_1|^2 ,$$

where $\mu(\varphi)$ is the function

$$\mu(\varphi) = \frac{\varphi}{\sin^2(\varphi)} - \cot(\varphi) .$$

The function $\mu(\varphi)$ is the key to the strange behaviour of sub-Riemannian geodesics. It is an odd function that diverges for $\varphi = k\pi$ ($k \neq 0$), i.e. $\eta_0^*\tau = 2k\pi$, and presents critical points when $\varphi = \tan(\varphi)$. But when $\varphi = \tan(\varphi)$, we have


$$\begin{aligned} \mu(\varphi) &= \frac{\tan(\varphi)}{\sin^2(\varphi)} - \cot(\varphi) \\ &= \frac{1 - \cos^2(\varphi)}{\cos(\varphi) \sin(\varphi)} = \tan(\varphi) = \varphi , \end{aligned}$$

and the minima of $\mu(\varphi)$ are on the diagonal. The graph of $\mu(\varphi)$ is represented in Fig. 6.1.

Let us compute geodesic lengths. Let γ be a geodesic starting at 0 and ending at time τ at $(x_1, y_1, p_1) = (z_1, p_1)$. If L is its length, we have $L = \int_0^\tau \ell ds$, with $\ell^2 = (\xi^* + p\eta^*)^2 + \pi^{*2}$ the squared norm of $\dot{\gamma}$ in the contact plane endowed with the orthonormal basis $\{X_1 = \partial_x + p\partial_y, X_2 = \partial_p\}$. But $\ell^2 = 2H = 2H_0$, since the Hamiltonian is constant along its trajectories, and we know that

$$H_0 = \frac{\eta_0^{*2}}{8} \frac{1}{\sin^2(\eta_0^*\tau/2)} |z_1|^2 .$$

Fig. 6.1 The function $\mu(\varphi)$ arising in the construction of sub-Riemannian geodesics of the polarized Heisenberg group \mathbb{V}_J . The axes have different scales



So, with $\eta_0^* \tau / 2 = \varphi$,

$$L = \sqrt{2} \left(\frac{\eta_0^* \tau}{2} \right) \frac{1}{|\sin(\eta_0^* \tau / 2)|} |z_1| = \sqrt{2} \frac{\varphi}{|\sin(\varphi)|} |z_1| .$$

In the sub-Riemannian geometry of \mathbb{V}_J , the sphere S and the wave front W (with radius $\sqrt{2}$) are given by the fundamental equation

$$|z_1| = \frac{|\sin(\varphi)|}{\varphi} .$$

We thus obtain

$$\left\{ \begin{array}{l} x_1 = \frac{|\sin(\varphi)|}{\varphi} \cos(\theta) , \\ p_1 = \frac{|\sin(\varphi)|}{\varphi} \sin(\theta) , \\ y_1 = \frac{1}{2} x_1 p_1 + \frac{\varphi - \sin(\varphi) \cos(\varphi)}{4\varphi^2} \\ = \frac{1}{2} \frac{\sin^2(\varphi)}{\varphi^2} \cos(\theta) \sin(\theta) + \frac{\varphi - \cos(\varphi) \sin(\varphi)}{4\varphi^2} \\ = \frac{\varphi + 2 \sin^2(\varphi) \cos(\theta) \sin(\theta) - \cos(\varphi) \sin(\varphi)}{4\varphi^2} . \end{array} \right.$$

Fig. 6.2 A piece of the sub-Riemannian wave front W . The external surface is the sub-Riemannian sphere S . The internal part is $W - S$. It presents smaller and smaller circles of cusp singularities which converge to 0




Figure 6.2 shows pieces of S and W . The external surface is the sub-Riemannian sphere S . It has two singularities at the intersections with the y -axis. The internal part is $W - S$. It presents smaller and smaller circles of cusp singularities which converge to 0. Such a complex behaviour is impossible in Riemannian geometry.

6.3 The \mathbb{V}_S Model

The jet-space \mathbb{V}_J model implies choosing a privileged x -axis. We saw in Sect. 4.3.8 of Chap. 4 that the group $SE(2)$ naturally acts on the model, but the asymmetry between x and y in the base space \mathbb{R}^2 was reflected in the ‘polarization’ of the Heisenberg group. As in every case in which a group G acts on a fibre bundle, here $G = SE(2)$ on $\pi_J : \mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$, it is relevant to consider the associated *principal* fibre bundle, here

$$\pi_S : G = SE(2) = \mathbb{R}^2 \rtimes SO(2) \cong \mathbb{V}_S = \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2 .$$

We developed this natural idea with Giovanna Citti and Alessandro Sarti (see their paper [11]). In this case, the contact form is

$$\omega_S = -\sin(\theta) dx + \cos(\theta) dy ,$$

that is, $\cos(\theta)(dy - pdx) = \cos(\theta)\omega_J$. The contact planes are spanned by the tangent vectors

$$\begin{cases} X_1 = \cos(\theta) \frac{\partial}{\partial x} + \sin(\theta) \frac{\partial}{\partial y}, \\ X_2 = \frac{\partial}{\partial \theta}, \end{cases}$$

with Lie bracket

$$[X_1, X_2] = \sin(\theta) \frac{\partial}{\partial x} - \cos(\theta) \frac{\partial}{\partial y} = -X_3.$$

In contrast to the polarized Heisenberg case, the X_j constitute an Euclidean orthonormal basis and are therefore more natural. The distribution \mathcal{K} of contact planes is still bracket generating (Hörmander condition) and maximally non-integrable since $d\omega_S = \cos(\theta) dx \wedge d\theta + \sin(\theta) dy \wedge d\theta$, and $\omega_S \wedge d\omega_S = -dx \wedge dy \wedge d\theta$ cannot vanish because it is a volume form. The Frobenius condition $\omega_S \wedge d\omega_S = 0$ is not satisfied, and there exists no integral surface of \mathcal{K} in \mathbb{V}_S (but there exist a lot of integral curves of \mathcal{K} , viz., all the Legendrian lifts Γ in \mathbb{V}_S of curves γ in the base plane \mathbb{R}^2). As for the characteristic vector field (or Reeb field) X_3 , it is orthogonal to K_ν for the Euclidean metric and defines a *scale* through

$$\omega_S(X_3) = [-\sin(\theta) dx + \cos(\theta) dy](X_3) = \sin^2(\theta) + \cos^2(\theta) = 1.$$

When we work with \mathbb{V}_S , the natural *sub-Riemannian* metric is the one making $\{X_1, X_2\}$ an orthonormal basis of the contact plane K_ν .

The two contact structures on $\mathbb{V}_J = \mathbb{R}^2 \times \mathbb{R}$ and $\mathbb{V}_S = \mathbb{R}^2 \times \mathbb{S}^1$ seem to be alike but are nevertheless very different. Indeed, let us look at their respective Lie algebras. For \mathbb{V}_J , we have the algebra \mathcal{V}_J generated by $\{t_1, t_2, t_3\}$, where

$$t_1 = \frac{\partial}{\partial x} + p \frac{\partial}{\partial y}, \quad t_2 = \frac{\partial}{\partial p}, \quad t_3 = \frac{\partial}{\partial y},$$

with $[t_1, t_2] = -t_3$ and $[t_1, t_3] = [t_2, t_3] = 0$ (we denote these vectors by t_i and no longer by X_i to avoid any confusion). As we have seen, \mathcal{V}_J is a nilpotent algebra because the coefficients $\{1, p, 1\}$ are polynomials whose derivatives vanish beyond a certain rank (here 2).

In contrast, for \mathbb{V}_S , we have the algebra \mathcal{V}_S generated by $\{X_1, X_2, X_3\}$ where

$$X_3 = -\sin(\theta) \frac{\partial}{\partial x} + \cos(\theta) \frac{\partial}{\partial y},$$

satisfying $[X_1, X_2] = -X_3$, $[X_1, X_3] = 0$ and $[X_2, X_3] = -X_1$, which is therefore *not nilpotent*. Nevertheless, we note that, for small θ , we have to first-order $p \sim \theta$, $\sin(\theta) \sim \theta$, and $\cos(\theta) \sim 1$, so $\omega_S = -\sin(\theta) dx + \cos(\theta) dy$ can be approximated by $\omega = -\theta dx + dy$, which is nothing else than the 1-form $\omega_J = dy - pdx$. \mathbb{V}_J is in some sense ‘tangent’ to \mathbb{V}_S . In fact, it is called the *tangent cone* of \mathbb{V}_S or its

nilpotentization (see, e.g. Mitchell [12], Rothschild and Stein [13], Margulis and Mostow [14], and Bellaïche [15]).

So we get two sub-Riemannian models \mathbb{V}_I and \mathbb{V}_S . The former is defined on a nilpotent group (a Carnot group) and the latter on a non-nilpotent group. Methods for neurogeometry, that is for the modelling of neural functional architectures of vision, thus become part of a sub-Riemannian geometry with its geodesics, Laplacians, heat kernels, harmonic analysis, and so on.

As we will see, many great geometers have studied these very rich structures. We have been personally interested in the work by Misha Gromov, Andrei Agrachev, Richard Beals, Bernard Gaveau, Peter Greiner, Luca Capogna, Vladimir Gershkovich, John Mitchell, Richard Montgomery, Robert Strichartz, Anatoly Vershik, Pierre Pansu, Jean-Michel Bismut, André Bellaïche, and Jean-Jacques Risler.

6.4 Geodesics of the \mathbb{V}_S Model

Andrei Agrachev has found the formulas of the geodesics for $\mathbb{V}_S = SE(2)$ endowed with the sub-Riemannian metric making $\{X_1, X_2\}$ an orthonormal basis of K_V . The formulation of the problem in terms of control theory yields the differential system $\{\dot{x} = u_1 \cos(\theta), \dot{y} = u_1 \sin(\theta), \dot{\theta} = u_2\}$. Applying the Pontryagin maximum principle, one gets the Hamiltonian on $T^*\mathbb{V}_S$ in the form

$$\begin{aligned} H(v, \varpi) &= \frac{1}{2} (u_1^2 + u_2^2) = \frac{1}{2} [(\varpi, X_1(v))^2 + (\varpi, X_2(v))^2] \\ &= \frac{1}{2} (\varpi_1^2 + \varpi_2^2) \\ &= \frac{1}{2} \left[[\xi^* \cos(\theta) + \eta^* \sin(\theta)]^2 + \vartheta^{*2} \right], \end{aligned}$$

where $\{\varpi_1, \varpi_2, \varpi_3\}$ are the components of the covector ϖ in the dual basis of $\{X_1, X_2, X_3\}$, and $\{\xi^*, \eta^*, \vartheta^*\}$ its components in the basis $\{dx, dy, d\theta\}$. Hence, the Hamilton equations are

$$\left\{ \begin{aligned} \dot{x} &= \frac{\partial H}{\partial \xi^*} = \xi^* \cos^2(\theta) + \eta^* \cos(\theta) \sin(\theta), \\ \dot{y} &= \frac{\partial H}{\partial \eta^*} = \eta^* \sin^2(\theta) + \xi^* \cos(\theta) \sin(\theta), \\ \dot{\theta} &= \frac{\partial H}{\partial \vartheta^*} = \vartheta^*, \\ \dot{\xi}^* &= -\frac{\partial H}{\partial x} = 0, \\ \dot{\eta}^* &= -\frac{\partial H}{\partial y} = 0, \\ \dot{\vartheta}^* &= -\frac{\partial H}{\partial \theta} = [\xi^* \cos(\theta) + \eta^* \sin(\theta)] [-\xi^* \sin(\theta) + \eta^* \cos(\theta)]. \end{aligned} \right.$$

The sub-Riemannian geodesics are the projections of the solutions on \mathbb{V}_S . Since $\xi^* = \xi_0^*$ and $\eta^* = \eta_0^*$ are constant, if we write them in the form $(\xi_0^*, \eta_0^*) = \rho_0 e^{i\beta_0}$, then

$$\dot{\vartheta}^* = \frac{1}{2} \rho_0^2 \sin(2(\theta - \beta_0)) ,$$

and the constant Hamiltonian

$$H = \frac{1}{2} [\rho_0^2 \cos^2(\theta - \beta_0) + \vartheta^{*2}]$$

yields the energy first integral $\rho_0^2 \cos^2(\theta - \beta_0) + \vartheta^{*2} = c$ (with $c = 1$ if $H = 1/2$) and the ODE for $\dot{\theta}$ (c, ρ_0 , and β_0 are constants) $\dot{\theta}^2 = \vartheta^{*2} = c - \rho_0^2 \cos^2(\theta - \beta_0)$. For $\beta_0 = 0$ (which is allowed by rotation invariance), the equations become:

$$\begin{cases} \dot{x} = \rho_0 \cos^2(\theta) , \\ \dot{y} = \rho_0 \cos(\theta) \sin(\theta) = \frac{1}{2} \rho_0 \sin(2\theta) , \\ \dot{\theta} = \vartheta^* , \\ \ddot{\theta} = \dot{\vartheta}^* = \frac{1}{2} \rho_0^2 \sin(2\theta) . \end{cases}$$

For $\rho_0 = 1, 2\theta = \pi - \mu$, and $\mu = 2\varphi = \pi - 2\theta$, we get a *pendulum equation*, viz., $\ddot{\mu} = -\sin(\mu)$ with first integral $\dot{\varphi}^2 + \sin^2(\varphi)$. As

$$dt = \pm \frac{1}{\sqrt{c}} \frac{d\varphi}{\sqrt{1 - \frac{1}{c} \sin^2(\varphi)}} ,$$


the system can be integrated explicitly using *elliptic functions*.

A fundamental property of these geodesics is that, under certain conditions, when the deviation from coaxiality between the boundary conditions (a_1, θ_1) and (a_2, θ_2) gets too great, they become *singular* and exhibit some cusps. At the level of the underlying pendulum equation, this corresponds to *oscillating* solutions. Yuri Sachkov and Igor Moiseev [16], then Remco Duits, studied these cusps and constructed the sphere, the wave front, and the cut locus of the sub-Riemannian geometry of $SE(2)$ explicitly. Their complexity is remarkable (see Duits [17]).

These singularities are interesting and can perhaps explain an intriguing aspect of illusory contours, namely their *bistability*. Consider for example the cross in Fig. 6.3. With the cooperation of V_2 which induces orientations orthogonal to segments at their end points, the segments induce an illusory contour which can be perceived as a square or a circle. If we look at the image for a sufficiently long time, the percepts spontaneously and periodically bifurcate from one case to the other.

In the context of a variational explanation of illusory contours, this shows how two models can compete: a geodesic model where the curvature is maximally spread out, and on the other hand, a piecewise linear model which concentrates all the

Fig. 6.3 Segments induce an illusory contour which can be perceived as a square or a circle. In general, the initial perception is of a circle, but it can bifurcate after looking at the picture for a certain time



curvature at some angular points (the curvature is null everywhere except at these points, where it is infinite). In the case of an illusory contour between (a_1, θ_1) and (a_2, θ_2) , experimental data seem to show that, if the coaxiality defect exceeds a certain threshold, the geodesic model is replaced by the piecewise linear model. We can formulate the conjecture that this *bifurcation of variational models* occurs when the geodesics become singular because of the emergence of cusps. In fact, a cusp occurs when a geodesic has a ‘vertical’ tangent, i.e. is tangent to the fibre of the bundle $\pi_S : \mathbb{V}_S = \mathbb{R}^2 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2$. But in neurophysiological terms, this means that some ‘horizontal’ excitatory connections between different hypercolumns must be identified with ‘vertical’ inhibitory connections internal to a single hypercolumn, which is not possible. Bimodality could then be caused by the fact that the period of fixation ‘stresses’ the selected model and allows its bifurcation.

6.5 Elastica Revisited

In 1992, David Mumford proposed an ‘elastica’ variational model [18] for illusory contours. This consists in minimizing an energy

$$E = \int_{\gamma} (\alpha \kappa^2 + \beta) ds ,$$

where γ is a curve in \mathbb{R}^2 with element of arc length ds . For $\alpha = \beta = 1$, its formulation as a control problem on the group $G = SE(2) = \mathbb{V}_S$ can be written $\{\dot{x} = \cos(\theta), \dot{y} = \sin(\theta), \dot{\theta} = \kappa\}$, where the derivatives are taken with respect to the arc length s and where $\kappa = d\theta/ds$ is the curvature of γ . This model is defined in the base plane \mathbb{R}^2 with Euclidean metric, and not with respect to the sub-Riemannian metric in G , because in G , ds is *not* the element of arc length. The element of arc length in G is $dt = \sqrt{1 + \kappa^2(s)} ds$, and the curvature

$$\kappa_G(s) = \frac{d\theta(t(s))}{dt} = \frac{d\theta}{ds} \frac{ds}{dt} = \frac{\kappa(s)}{\sqrt{1 + \kappa^2(s)}}$$

satisfies the pendulum equation (with respect to s), viz., $\ddot{\kappa}_G(s) = \kappa_G(s)$. Yuri Sachkov [19] investigated the elastica model curves and compared them to the sub-Riemannian geodesics. The problem is difficult because of the cusps. He later pursued this investigation with Ugo Boscaïn, Remco Duits, and Francesco Rossi in [20].

David Mumford gave a stochastic interpretation of his elastica model by supposing that the curvature $\kappa(s)$ of γ in \mathbb{R}^2 is white noise and that the angle $\theta(s)$ is therefore Brownian motion.¹ In terms of control theory, this is equivalent to considering the stochastic process $\{\dot{x} = \cos(\theta), \dot{y} = \sin(\theta), \dot{\theta} \sim N(0, \sigma^2)\}$, where $\dot{\theta}$ is now a normal random Gaussian variable of mean 0 and variance σ^2 . This process has been studied by Gonzalo Sanguinetti (in his thesis [21] supervised by G. Citti and A. Sarti) and also by Remco Duits and Markus van Almsick. It is no longer a simple diffusion mechanism, but an advection–diffusion mechanism described by a Fokker–Planck equation. The advection (the drift) occurs along the X_1 direction, and the diffusion of θ occurs along the X_2 direction. The fundamental solution of the Fokker–Planck equation being too complex in the \mathbb{V}_S model, the authors went back to the first-order approximation of $G = SE(2)$ (its tangent cone or nilpotentization), that is, to our \mathbb{V}_J model based on the polarized Heisenberg group. Let $v_0 = (x_0, y_0, \theta_0) = (a_0, \theta_0)$ be an initial point in G and let us follow a random walk starting at v_0 . Without noise, the trajectory is of course deterministic and is a straight line satisfying the principle of strict coaxiality (without any curvature): $\{\theta = \theta_0, x = x_0 + \cos(\theta_0)t, y = y_0 + \sin(\theta_0)t\}$. If $v = (x, y, \theta) = (a, \theta)$ is a generic element of G and if $P(v, t)$ is the probability of finding the random walk at v at time t , the evolution equation for P with initial condition $P_0(v) = P(v, 0)$ is

$$\frac{\partial P}{\partial t}(v, t) = - \left[\cos(\theta) \frac{\partial P}{\partial x}(v, t) + \sin(\theta) \frac{\partial P}{\partial y}(v, t) \right] + \frac{\sigma^2}{2} \frac{\partial^2 P}{\partial \theta^2}(v, t) ,$$

whence

$$\frac{\partial P}{\partial t}(v, t) = -X_1(P(v, t)) + \frac{\sigma^2}{2} (X_2)^2(P(v, t)) .$$

For the \mathbb{V}_J model where θ is small, $\theta \sim \tan(\theta) = p$, and the Fokker–Planck equation is therefore

$$\frac{\partial P}{\partial t}(v, t) = - \left[\frac{\partial P}{\partial x}(v, t) + p \frac{\partial P}{\partial y}(v, t) \right] + \frac{\sigma^2}{2} \frac{\partial^2 P}{\partial p^2}(v, t) .$$

The authors solve this equation and, to complete a contour with boundary conditions $v_0 = (a_0, \theta_0)$ and $v_1 = (a_1, \theta_1)$, consider two direction processes, a forward process starting at v_0 and a backward process starting at v_1 . They compute the probability of collision of these two random walks.

¹Named after Robert Brown.

6.6 Sub-Riemannian Diffusion, Heat Kernel, and Non-commutative Harmonic Analysis

We have stressed the importance and the difficulties of sub-Riemannian diffusion techniques, mentioning the work of some specialists and also the neurogeometrical applications made by G. Citti and A. Sarti in this area. These techniques belong to the general theory of heat kernels on Riemannian and sub-Riemannian manifolds. The specialized literature on these subjects is enormous.

Building on earlier results by Hulanicki [22], A. Agrachev, J-P. Gauthier, U. Boscain, and their Ph.D. student F. Rossi gave an ‘intrinsic’ formulation of the sub-Riemannian Laplacian and in 2009 [23] proved a general theorem for the 3D unimodular Lie groups (i.e. those whose left- and right-invariant Haar measures are identical) endowed with a left-invariant sub-Riemannian geometry. They used the non-commutative generalized Fourier transform (GFT) defined on the dual space G^* of G (the set of irreducible unitary representations in Hilbert spaces) to compute the heat kernel associated with the hypoelliptic Laplacian $\Delta_{\mathcal{X}} = X_1^2 + X_2^2$, i.e. the sum of squares of the generators $\{X_1, X_2\}$ of the distribution \mathcal{X} . The Laplacian is *hypoelliptic* due to the fact that \mathcal{X} is bracket generating, i.e. satisfies the Hörmander condition.

The use of Fourier transform on groups to compute heat kernels and fundamental solutions of diffusion equations has a long history. In the case of the polarized Heisenberg Lie group \mathbb{V}_J , according to the Stone–von Neumann theorem, the non-trivial unitary irreducible representations (unirreps) are group morphisms π_λ from \mathbb{V}_J to the group $\mathcal{U}(\mathcal{H})$ of unitary automorphisms of the Hilbert space $\mathcal{H} = L^2(\mathbb{R}, \mathbb{C})$, morphisms parametrized by a real scalar $\lambda \neq 0$. They are of the form:

$$\begin{aligned} \pi_\lambda : \mathbb{V}_J &\longrightarrow \mathcal{U}(\mathcal{H}) \\ v &\longmapsto \pi_\lambda(v) : \mathcal{H} \longrightarrow \mathcal{H} \\ u(s) &\longmapsto e^{i\lambda(y+xs)}u(s+p). \end{aligned}$$

There exists a measure on the dual space \mathbb{V}_J^* , called the *Plancherel measure*, given by $dP(\lambda) = \lambda d\lambda$, which allows us to do integration. To compute the Fourier transform of the sub-Riemannian Laplacian $\Delta_{\mathcal{X}}$, we examine the action of the differential of the unirreps on the left-invariant vector fields X on \mathbb{V}_J , which are given by the left translation of vectors $X(0)$ of the Lie algebra \mathcal{V}_J of \mathbb{V}_J . By definition,

$$d\pi_\lambda : X \rightarrow d\pi_\lambda(X) := \left. \frac{d}{dt} \right|_{t=0} \pi_\lambda(e^{tX}),$$

and we obtain the Fourier transform $\widehat{X}_i^\lambda = d\pi_\lambda(X_i)$. Carrying out the calculations, we obtain $X_1(0) = (1, 0, 0)$, $e^{tX_1} = (t, 0, 0)$, $\pi_\lambda(e^{tX_1})u(s) = e^{i\lambda ts}u(s)$,

$$\begin{aligned} \widehat{X}_1^\lambda u(s) &= d\pi_\lambda(X_1) u(s) = \left. \frac{d}{dt} \right|_{t=0} \pi_\lambda(e^{tX_1}) u(s) \\ &= \left. \frac{d}{dt} \right|_{t=0} e^{i\lambda t s} u(s) = i\lambda s u(s) , \end{aligned}$$

and $X_2(0) = (0, 0, 1)$, $e^{tX_2} = (0, 0, t)$, $\pi_\lambda(e^{tX_2}) u(s) = u(s + t)$,

$$\begin{aligned} \widehat{X}_2^\lambda u(s) &= d\pi_\lambda(X_2) u(s) = \left. \frac{d}{dt} \right|_{t=0} \pi_\lambda(e^{tX_2}) u(s) \\ &= \left. \frac{d}{dt} \right|_{t=0} u(s + t) = \frac{du(s)}{ds} . \end{aligned}$$

The GFT of the sub-Riemannian Laplacian is therefore the Hilbert sum, i.e. the integral on λ with the Plancherel measure $dP(\lambda) = \lambda d\lambda$, of the $\widehat{\Delta_{\mathcal{H}}^\lambda}$ with

$$\widehat{\Delta_{\mathcal{H}}^\lambda} u(s) = \left[\left(\widehat{X}_1^\lambda \right)^2 + \left(\widehat{X}_2^\lambda \right)^2 \right] u(s) = \frac{d^2 u(s)}{ds^2} - \lambda^2 s^2 u(s) .$$

This equation is nothing else than the equation of the *harmonic oscillator*.

The heat kernel is then

$$P(v, t) = \int_{\mathbb{V}_j^*} \text{Tr} \left[e^{t \widehat{\Delta_{\mathcal{H}}^\lambda}} \pi_\lambda(v) \right] dP(\lambda) , \quad t \geq 0 ,$$

where Tr denotes the trace. If the $\widehat{\Delta_{\mathcal{H}}^\lambda}$ have discrete spectrum and a complete set of normalized eigenfunctions $\{u_n^\lambda\}$ with eigenvalues $\{\alpha_n^\lambda\}$, then

$$P(v, t) = \int_{\mathbb{V}_j^*} \left[\sum_n e^{\alpha_n^\lambda t} \langle u_n^\lambda, \pi_\lambda(v)(u_n^\lambda) \rangle \right] dP(\lambda) , \quad t \geq 0 .$$

This is the case that concerns us here. The eigenfunctions of the harmonic oscillator are well known and satisfy

$$\frac{d^2 u_n^\lambda(s)}{ds^2} - \lambda^2 s^2 u_n^\lambda(s) = \alpha_n^\lambda u_n^\lambda(s) ,$$

with $\alpha_n^\lambda = -(2n + 1)/\lambda$. They are essentially the *Hermite functions* scaled by λ :

$$u_n^\lambda(s) = (2^n n! \sqrt{\pi})^{-1/2} \lambda^{1/4} e^{-\lambda s^2/2} H_n(\sqrt{\lambda} s) ,$$

where H_n is the n th Hermite polynomial.

In the case of $SE(2) = \mathbb{V}_S$, the authors found explicit formulas for the heat kernel. The dual \mathbb{V}_S^* of \mathbb{V}_S is now the set of unirreps in the Hilbert space $\mathcal{H} = L^2(\mathbb{S}^1, \mathbb{C})$.

These unirreps are parametrized by a positive real λ and are of the form:

$$\begin{aligned} \mathcal{X}^\lambda : \mathbb{V}_S &\longrightarrow \mathcal{U}(\mathcal{H}) \\ v &\longmapsto \mathcal{X}^\lambda(v) : \mathcal{H} \longrightarrow \mathcal{H} \\ \psi(\theta) &\longmapsto e^{i\lambda[x \sin(\theta) + y \cos(\theta)]} \psi(\theta + \alpha) . \end{aligned}$$

The Plancherel measure on \mathbb{V}_S^* is still $dP(\lambda) = \lambda d\lambda$. As explained previously, we get $\mathcal{X}^\lambda(e^{tX_1}) \psi(\theta) = e^{i\lambda t \sin(\theta)} \psi(\theta)$ and

$$\begin{aligned} \widehat{X_1}^\lambda \psi(\theta) &= d\mathcal{X}^\lambda(X_1) \psi(\theta) = \left. \frac{d}{dt} \right|_{t=0} \mathcal{X}^\lambda(e^{tX_1}) \psi(\theta) \\ &= \left. \frac{d}{dt} \right|_{t=0} e^{i\lambda t \sin(\theta)} \psi(\theta) = i\lambda \sin(\theta) \psi(\theta) , \end{aligned}$$

and $\mathcal{X}^\lambda(e^{tX_2}) \psi(\theta) = \psi(\theta + t)$ and

$$\begin{aligned} \widehat{X_2}^\lambda \psi(\theta) &= d\mathcal{X}^\lambda(X_2) \psi(\theta) = \left. \frac{d}{dt} \right|_{t=0} \mathcal{X}^\lambda(e^{tX_2}) \psi(\theta) \\ &= \left. \frac{d}{dt} \right|_{t=0} \psi(\theta + t) = \frac{d\psi(\theta)}{d\theta} . \end{aligned}$$

The GFT of the sub-Riemannian Laplacian is therefore the Hilbert sum of the $\widehat{\Delta_{\mathcal{X}}^\lambda}$ with

$$\widehat{\Delta_{\mathcal{X}}^\lambda} \psi(\theta) = [(\widehat{X_1}^\lambda)^2 + (\widehat{X_2}^\lambda)^2] \psi(\theta) = \frac{d^2\psi(\theta)}{d\theta^2} - \lambda^2 \sin^2(\theta) \psi(\theta) ,$$

which is nothing else than the *Mathieu equation*. The heat kernel is

$$P(v, t) = \int_{\mathbb{V}_S^*} \text{Tr} \left[e^{t \widehat{\Delta_{\mathcal{X}}^\lambda}} \mathcal{X}^\lambda(v) \right] dP(\lambda) , \quad t \geq 0 .$$

The $\widehat{\Delta_{\mathcal{X}}^\lambda}$ have a discrete spectrum and a complete set of normalized eigenfunctions $\{\psi_n^\lambda\}$ with eigenvalues $\{\alpha_n^\lambda\}$, and therefore,

$$P(v, t) = \int_{\mathbb{V}_S^*} \sum_n e^{\alpha_n^\lambda t} \langle \psi_n^\lambda, \mathcal{X}^\lambda(v) \rangle \langle \psi_n^\lambda \rangle dP(\lambda) , \quad t \geq 0 .$$

The 2π -periodic eigenfunctions of the Mathieu equation satisfy

$$\frac{d^2\psi(\theta)}{d\theta^2} - \lambda^2 \sin^2(\theta) \psi(\theta) = E\psi(\theta) ,$$

and as $\sin^2(\theta) = [1 - \cos(2\theta)]/2$, this means

$$\begin{aligned} \frac{d^2\psi(\theta)}{d\theta^2} - \frac{\lambda^2}{2}\psi(\theta) - E\psi(\theta) + \frac{\lambda^2}{2}\cos(2\theta)\psi(\theta) &= 0, \\ \frac{d^2\psi(\theta)}{d\theta^2} + [a - 2q\cos(2\theta)]\psi(\theta) &= 0, \end{aligned}$$

where $a = -(E + \lambda^2/2)$ and $q = -\lambda^2/4$. The normalized 2π -periodic eigenfunctions are known: they are even or odd and denoted $\text{cen}(\theta, q)$ and $\text{sen}(\theta, q)$. The associated $a_n(q)$ and $b_n(q)$ are called characteristic values. There can be *parametric resonance* phenomena (Arnold tongues) when $a = -(E + \lambda^2/2) = n^2$.

The authors also solved the problem for $SU(2)$, $SL(2)$, and $SO(3)$.

Sub-Riemannian diffusion is highly *anisotropic* since it is restricted to an angular diffusion of θ and a spatial diffusion only along the X_1 direction. It is strongly constrained by the good continuation Gestalt law and the difference with classical (Euclidean) diffusion is spectacular.

6.7 Confluence Between \mathbb{V}_J and \mathbb{V}_S Models

We analyzed two neurogeometrical models of $V1$, \mathbb{V}_J and \mathbb{V}_S . It is interesting to note that one can easily construct an *interpolation* between the two models. Mohammed Brahim Zahaf and Dominique Manchon [24] constructed just such an interpolation, given by a family of models \mathbb{V}^α , and studied the confluence of the corresponding differential equations in the Fourier space. The model \mathbb{V}^α is summarized in the following:

$$\begin{aligned} X_1^\alpha &= \cos(\theta) \frac{\partial}{\partial x} + \frac{1}{\alpha} \sin(\alpha\theta) \frac{\partial}{\partial y}, \\ X_2^\alpha &= \frac{\partial}{\partial \theta}, \\ X_3^\alpha &= -\alpha \sin(\alpha\theta) \frac{\partial}{\partial x} + \cos(\theta) \frac{\partial}{\partial y}, \\ [X_1^\alpha, X_2^\alpha] &= -X_3^\alpha, \\ [X_2^\alpha, X_3^\alpha] &= \alpha^2 X_1^\alpha, \\ [X_1^\alpha, X_3^\alpha] &= 0, \\ \mathbb{V}^\alpha &= SE_\alpha(2), \text{ with } \mathbb{S}_\alpha^1 = \frac{\mathbb{R}}{2\pi\alpha^{-1}\mathbb{Z}}, \\ \mathcal{H} &= L^2(\mathbb{S}_\alpha^1, \mathbb{C}), \\ X_1^\alpha(\psi(\theta)) &= i\lambda\alpha^{-1} \sin(\alpha\theta) \psi(\theta), \\ X_2^\alpha(\psi(\theta)) &= \psi'(\theta), \end{aligned}$$

$$\begin{aligned} \text{operator } \hat{\Delta}^\lambda &: \psi''(\theta) - \frac{\lambda^2}{\alpha^2} \sin^2(\alpha\theta) \psi(\theta) , \\ \text{equation} &: \psi''(\theta) + \left[\mu - \frac{\lambda^2}{\alpha^2} \sin^2(\alpha\theta) \right] \psi(\theta) = 0 . \end{aligned}$$

For $\alpha = 1$, ∇^1 gives the model ∇_S , and when $\alpha \rightarrow 0$, for small θ denoted by p , ∇^0 gives the model ∇_J .

6.8 Other Themes

All these technical points will be developed in detail in the second volume. As discussed in the *Preface*, we will also present the following themes, already addressed in [25–28]:

1. Modal and amodal illusory contours (this is an ‘interpolation’ problem, see, e.g. [29]).
2. Physical models of neural networks and their synchronization properties. Synchronization is a key feature of brain processes. Indeed, as emphasized by Jan Koenderink, their functional order is a *simultaneous* functional order (see Toet et al. [30]).
3. The work of Bard Ermentrout, Jack Cowan, Paul Bressloff, and Martin Golubitsky on Hopfield networks in which synaptic weights encodes the functional architecture of V1.
4. Our work with Alessandro Sarti and Giovanna Citti on the scale parameter and the resulting symplectic structure (see also [31]).
5. Multiscale differential geometry and the Thom–Mather theory of singularities as reworked in this framework by James Damon (see also [32]).
6. Image segmentation algorithms, including those where segmentation is achieved by applying nonlinear anisotropic diffusion equations and those arising from David Mumford and Jayant Shah’s variational models.
7. Epistemological problems ranging from the phenomenology of perception in the sense of Husserl and Merleau-Ponty (see [33, 34]) to the Kantian problem of transcendental aesthetics. Indeed, as we have seen with Koenderink [35], Dehaene [36] or O’Keefe [37], Kant is vindicated by visual neurosciences.

References

1. Brockett, R.: Control Theory and Singular Riemannian Geometry. Springer, Berlin, New York (1981)
2. Montgomery, R.: A Tour of Subriemannian Geometries, Their Geodesics and Applications. American Mathematical Society (2002)
3. Strichartz, R.S.: Sub-Riemannian geometry. J. Differ. Geom. **24**, 221–263 (1986)

4. Ge, Z.: Horizontal path spaces and Carnot-Carathéodory metrics. *Pacific J. Math.* **161**(2), 255–286 (1993)
5. Hammenstädt, U.: Some regularity in Carnot-Carathéodory metrics. *J. Differ. Geom.* **32**, 192–201 (1991)
6. Beals, R., Gaveau, B., Greiner, P.C.: Hamilton-Jacobi theory and the heat kernel on Heisenberg groups. *Journal de Mathématiques Pures et Appliquées* **79**(7), 633–689 (2000)
7. Gaveau, B.: Principe de moindre action, propagation de la chaleur et estimées sous-elliptiques sur certains groupes nilpotents. *Acta Mathematica* **139**, 96–153 (1977)
8. Agrachev, A.A., Sachkov, Y.: *Control Theory from the Geometric Viewpoint*. Springer (2004)
9. Agrachev, A.A., Gamkrelidze, R.V.: The Pontryagin maximum principle 50 years later. In: *Proceedings of the Steklov Mathematical Institute, Dynamical Systems: Modeling, Optimization, and Control*, pp. S4–S12 (2006)
10. Petitot, J., (with a contribution of Tondut, Y.): Vers une Neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. *Mathématiques, Informatique et Sciences Humaines* 145, 5–101 (1999)
11. Citti, G., Sarti, A.: A cortical based model of perceptual completion in the roto-translation space. *J. Math. Imaging Vis.* **24**(3), 307–326 (2006)
12. Mitchell, J.: On Carnot-Carathéodory metrics. *J. Differ. Geom.* **21**, 35–45 (1985)
13. Rothschild, L.P., Stein, E.M.: Hypoelliptic differential operators and nilpotent groups. *Acta Mathematica* **137**, 247–320 (1976)
14. Margulis, G.A., Mostow, G.D.: Some remarks on the definition of tangent cones in a Carnot-Carathéodory space. *Journal d'Analyse Mathématique* **80**, 299–317 (2000)
15. Bellaïche, A.: The tangent space in sub-Riemannian geometry. In: Bellaïche, A., Risler, J. (eds.) *Sub-Riemannian geometry*, pp. 4–78. Basel, Birkhäuser (1996)
16. Moiseev, I., Sachkov, Y.: Maxwell Strata in sub-Riemannian problem on the group of motion of a plane. *ESAIM: Control. Optim. Calc. Var.* **16**(2), 380–399 (2010)
17. Duits, R., Ghosh, A., Dela Haije, T., Sachkov, Y.: Cuspless sub-Riemannian geodesics within the Euclidean motion group $SE(d)$. In: Citti, G., Sarti, A. (eds.) *Neuromathematics of Vision*, pp. 173–215. Springer, Berlin (2014)
18. Mumford, D.: *Elastica and computer vision*. In: Bajaj, C. (ed.) *Algebraic Geometry and Applications*, pp. 491–506. Springer, Heidelberg (1992)
19. Sachkov, Y.L.: Conjugate points in the Euler elastica problem. *J. Dyn. Control Syst.* **14**(3), 409–439 (2008)
20. Boscain, U., Duits, R., Rossi, F., Sachkov, Y.: Curve cuspless reconstruction via sub-Riemannian geometry. [arXiv: 1203.3089v4](https://arxiv.org/abs/1203.3089v4) (2013)
21. Sanguinetti, G.: *Invariant models of vision between phenomenology, image statistics and neurosciences*. Montevideo University, Thesis (2011)
22. Hulanicki, A.: The distribution of energy in the Brownian motion in the Gaussian field and analytic hypoellipticity of certain subelliptic operators on the Heisenberg group. *Studia Mathematica* **56**(2), 165–173 (1976)
23. Agrachev, A.A., Boscain, U., Gauthier, J.-P., Rossi, F.: The intrinsic hypoelliptic Laplacian and its heat kernel on unimodular Lie groups. *J. Funct. Anal.* **256**, 2621–2655 (2009)
24. Zahaf, M.B., Manchon, D.: Confluence of singularities of differential equations: a Lie algebra contraction approach, <https://hal.archives-ouvertes.fr/hal-00292676>, 2008
25. Petitot, J.: Neurogeometry of V1 and Kanizsa contours. *Axiomathes* **13**, 347–363 (2003)
26. Petitot, J.: Neurogéométrie des architectures fonctionnelles de la vision. In: *Mathématiques et Vision, Journée annuelle de la SMF*, pp. 69–128 (2006)
27. Petitot, J.: Neurogeometry of neural functional architectures. *Chaos, Solitons & Fractals* **50**, 75–92 (2013)
28. Petitot, J.: *Neurogéométrie de la vision. Modèles mathématiques et physiques des architectures fonctionnelles*. Les Éditions de l'École Polytechnique, Distribution Ellipses, Paris (2008)
29. Kellman, P.J., Shipley, T.F.: A theory of visual interpolation in object perception. *Cogn. Psychol.* **23**, 141–221 (1991)

30. Toet, A., Blom, J., Koenderink, J.J.: The construction of a simultaneous functional order in nervous systems, I. *Biol. Cybern.* **57**, 115–125 (1987)
31. Sarti, A., Citti, G., Petitot, J.: On the symplectic structure of the primary visual cortex. *Biol. Cybern.* **98**(1), 33–48 (2008)
32. Petitot, J.: Théorie des singularités et équations de diffusion, pp. 89–107. *Gazette des Mathématiciens*, René Thom (2004)
33. Petitot, J.: Morphological eidetics for phenomenology of perception. In: Petitot, J., Varela, F.J., Roy, J.-M., Pachoud, B. (eds.) *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*, pp. 330–371. Stanford University Press, Stanford (1999)
34. Petitot, J.: Neurogéométrie et phénoménologie de la perception. In: Bouveresse, J. Rosat, J.-J. (eds.) *Philosophie de la Perception*, Collège de France–Odile Jacob, pp. 53–76. Paris (2003)
35. Koenderink, J.J.: The brain a geometry engine. *Psychol. Res.* **52**, 122–127 (1990)
36. Dehaene, S., Brannon, E.M.: Space, time, and number: a Kantian research program. *Trends Cogn. Sci. Spec.* **14**(12), 517–519 (2010)
37. O’Keefe, J.: Immanuel Kant: pioneer neuroscientist, public lecture. Royal Institution of Great Britain, London, June 2 (2014)

Author Index

A

Abel, N.H., 21
Abrahamsen, P., 196
Adelson, E., 76
Adler, R., 57
Afgoustidis, A., 196
Agazzi, E., 18
Agrachev, A., 5, 7, 17, 18, 356
Alais, D., 291
Albertazzi, L., 17
Albright, T.D., 319
Alekseevsky, D., 17
Alembert, J. le Rond d', 12
Alexander, D., 277
Allilaire, J.-F., 14
Almsick, M. van, 7, 359
Alonso, J.-M., 66
Ambrosio, L., 17
Andler, D., 15
Angelucci, A., 124
Arikawa, K., 57
Aristotle, 13, 35, 263
Arnold, V., 11, 323, 363
Atick, J., 95, 97
Atlan, H., 12, 15
Attneave, F., 95
Azencott, R., 157

B

Badea, T., 59, 331
Balasubramanian, M., 123
Barbara Chapman, Michael Stryker, and Tobias Bonhöffer, 151
Barbieri, D., 194
Barbut, M., 14
Barilari, D., 17

Barlow, H.B., 95, 319
Barry, C., 18, 59, 105
Basole, A., 147
Bassett, J.P., 107
Baxter, W., 149, 151
Baye, L., 331, 334
Bayes, T., 96, 286
Beals, R., 5, 348, 356
Beaudot, W., 84, 113, 124
Bednar, J., 203, 245
Bellaïche, A., 5, 17, 356
Bellettini, G., 17
Bellman, R., 17
Bennequin, D., 17, 231
Benoît, E., 18
Ben-Shahar, O., 132, 304
Berestycki, H., 15
Berry, M., 4, 13, 172, 173, 180, 181, 194, 195, 200
Berthier, M., 18
Berthoz, A., 12, 14, 18, 25, 102, 259, 302
Bessel, F.W., 197
Besson, M.-J., 14
Bettioli, P., 17
Bienenstock, E., 16
Bismut, J.-M., 6, 356
Bitbol, M., 18
Blake, R., 255
Blasdel, G.G., 138–140
Bloom, J., 74
Blum, H., 8, 41
Bochner, S., 196, 197
Bolzano, B., 28, 38
Bonhöffer, T., 138, 208
Boole, G., 287
Borel, É., 17
Born, R., 319

Boscain, U., 5, 7, 17, 359
 Bosking, W., 141, 220, 222, 259, 278, 281
 Bouligand, Y., 13
 Bourbaki, 157
 Bourguine, P., 15
 Bourguignon, J.-P., 12
 Bouveresse, J., 18
 Boyd, J., 250, 252
 Bradley, D., 319
 Braitenberg, C., 150
 Braitenberg, V., 147
 Brandt, P.-A., 18
 Brannon, E., 2
 Brentano, F., 12
 Bressloff, P., 8, 16, 18, 40, 224
 Broca, P., 45
 Brockett, R.W., 348
 Brodmann, K., 114
 Brown, R., 359
 Bryant, R., 303
 Buchsbaum, G., 267
 Bumsted-O'Brian, K., 331
 Burgess, N., 105
 Burnod, Y., 15
 Buser, P., 46
 Butts, D.A., 62
 Byrne, C., 18

C

Cai, A., 342
 Cai, Daqing, 70
 Cai, David, 297
 Cantor, G., 10
 Capogna, L., 6
 Carandini, M., 66
 Carathéodory, C., 5
 Carnot, S., 356
 Carreiro-Perpiñán, M., 244
 Carsetti, A., 17
 Cartan, E., 21, 22, 25, 28, 102, 299
 Cartier, P., 16
 Cauchy, A.L., 177, 195
 Changeux, J.-P., 15
 Chaperon, M., 12
 Charnay, P., 15
 Chedotal, A., 331
 Chen, C., 120
 Chenciner, A., 12
 Chiarcos, O., 19
 Chitour, Y., 17
 Chklovskii, D., 156
 Choe, Y., 245

Chossat, P., 17
 Chow, W.-L., 348
 Citti, G., 5–8, 16, 173, 354, 360
 Connes, A., 12, 22, 28
 Conte, A., 18
 Coombes, S., 18
 Corradi Dell'Acqua, F., 19
 Cowan, J., 8, 16, 17, 33, 40, 224
 Crair, M., 141, 242, 244
 Cramer, H., 196
 Crick, F., 279
 Curry, H., 29

D

Daido, H., 8
 Damon, J., 8
 D'Arcy Wentworth Thompson, 11
 Das, A., 220
 Daubechies, I., 15, 87
 Daugman, J., 73
 David Alexander, Paul Bourke, Otto Konstandatos, Phil Sheridan, and James Wright, 277
 David, G., 5, 16, 95, 99
 DeAngelis, G.C., 63, 67, 69, 71, 223
 De Giorgi, E., 16
 Dehaene, S., 2
 Deligne, P., 219
 Dennis, M.R., 179, 194, 197, 200
 Descartes, R., 35, 45
 Descola, P., 18
 Deshouillers, J.-M., 12
 Desolneux, A., 286
 Destexhe, A., 16
 Diderot, D., 12
 Doeller, C.F., 104–106
 Dostrovsky, J., 103
 Dow, B., 149, 151
 Dreher, B., 65
 Dreyfus, G., 15
 Droulez, J., 16
 Dubner, R., 319
 Duits, R., 7, 17, 357, 359
 Dumoulin, S.O., 119
 Dupoux, É., 15
 Dupuy, J.-P., 12, 15

E

Eco, U., 18
 Engel, F., 23
 Erasistratus, 45
 Ermentrout, B., 8

Euclid, 10
 Euler, L., 25, 32, 303, 349
 Evans, L., 14
 Eysel, U., 143, 276

F

Farge, M., 90
 Faugeras, O., 16–18
 Feldheim, D.A., 46, 337
 Finkel, L., 267
 FitzHugh, R., 52
 Fitzpatrick, D., 142, 295
 Flanagan, J., 331, 340, 342
 Florack, L., 13, 80, 95
 Fokker, A., 7, 359
 Folland, G., 6
 Foster, D., 266
 Fourier, J., 85
 Fraassen, B. van, 18
 Frankowska, H., 17
 Frege, G., 28, 38
 Frégnac, Y., 14, 16, 17, 62, 95, 143, 290, 296
 Friedman, M., 18
 Fritsch, G.T., 45
 Frobenius, F.G., 355
 Fyhn, M., 104

G

Gabor, D., 66, 68, 73, 86, 282, 290, 291, 293
 Gage, M., 14
 Galileo, G., 12
 Gamkrelidze, R.V., 349
 Gauss, C.F., 22
 Gauthier, J.-P., 5, 7, 17
 Gaveau, B., 5, 348, 356
 Geisel, T., 4, 171, 173, 195, 201, 203, 204
 Geoffroy Saint-Hilaire, É., 12
 Georges, S., 14, 290
 Gerbino, W., 17
 Gershkovich, V., 6, 356
 Gevrey, M., 219
 Gibson, J., 96, 317
 Gilbert, C.D., 65, 293
 Giraud, J., 10
 Glotin, H., 18
 Gödecke, I., 208
 Goethe, J.H. von, 264, 265
 Goles, E., 18
 Golgi, C., 45, 50
 Golubitsky, M., 8, 11, 40
 Goodhill, G., 150, 244

Goodwin, B., 13, 338
 Gray, C., 125, 129, 208
 Grayson, M., 14
 Gregor, K., 99
 Greiner, P., 5, 348, 356
 Griffiths, P., 303
 Grinvald, A., 138, 139, 320, 327
 Gromov, M., 5, 12, 26, 28, 356
 Grossberg, S., 263, 266, 295
 Grossmann, A., 88
 Grothendieck, A., 10
 Gruner-Schlumberger, A., 16
 Guilbaud, G.-T., 14

H

Haar, A., 360
 Hafting, T., 104
 Haken, H., 12
 Hale, D., 82
 Hamdouni, A., 18
 Hamilton, R., 14
 Hamilton, W. R., 12, 135
 Hammenstädt, U., 348
 Hartline, H., 51
 Hasboun, D., 115
 Hawken, M.J., 265, 266, 316
 Hayes, A., 5, 282, 286–288, 290, 292, 295
 Heaviside, O., 256
 Hebb, D.O., 98
 Hebbian, 98, 203, 204
 Heine, E., 168
 Heisenberg, W.K., 5, 6, 74, 86, 300, 347, 348, 353–355, 359, 360
 Helmholtz, H. von, 10, 22, 23, 60, 144, 179, 180, 187, 197, 264, 347
 Helvétius, C.-A., 45
 Henry, G., 65
 Hering, E., 61, 265
 Hermite, C., 7, 75
 Herophilos, 45
 Hess, R., 5, 39, 282, 285–288, 290, 295
 Higgs, P., 25
 Hikosaka, O., 296
 Hilbert, D., 10, 85, 86, 89, 360–362
 Hippocrates, 45
 Hiripi, E., 18
 Hironaka, H., 10, 12
 Hitzig, E., 45
 Hodgkin, A.L., 7, 31
 Hoffman, W., 13
 Hooks, B.M., 47, 120
 Hopfield, J., 5, 8, 364

Hörmander, L.V., 6, 348, 355, 360
 Horwitz, L.G., 2
 Howard, A., 29
 Hoyer, P., 101, 158
 Hubel, D.H., 65, 67, 123, 128, 129, 137, 147, 201, 259
 Hübener, M., 223, 227, 242, 243
 Hughes, B.D., 342
 Hume, D., 35
 Husserl, E., 8, 9, 12, 17, 18, 22, 38, 42, 263, 317
 Huxley, A.F., 7, 31, 52
 Hyvärinen, A., 101, 158

I

Imbert, M., 14, 45, 46
 Issa, N., 223–226, 228, 242

J

Jackson, H., 45
 Jacobi, C.G.J., 17, 176, 177, 195, 197, 199
 Jancke, D., 42, 296
 Jean, F., 14
 Jirsa, V., 18
 Johnson, E., 266, 268
 Jones, J., 73
 Jong, T.M. de, 256
 Jonnal, R., 57
 Julesz, B., 285, 287

K

Kac, M., 197, 199
 Kaluza, T., 219
 Kanizsa, G., 6, 16, 26
 Kant, I., 1, 2, 25, 35
 Kapadia, M.K., 282
 Kara, P., 250, 252
 Karol Gregor, Arthur Szlam, and Yann Le-Cun, 158
 Kashube, M., 203
 Kayser, A., Priebe, Miller, 66
 Kelso, S., 12
 Kepler, J., 30, 179
 Kerszberg, P., 18
 Killing, W., 199
 Kim, Y.-J., 255
 Klein, F., 10
 Klüver, H., 22, 40
 Koch, C., 65
 Koechlin, E., 15

Koenderink, J., 13, 17, 25, 75, 80, 124, 133, 134, 302, 304, 364
 Kohonen, T., 247
 Kolb, H., 58
 Kolodkin, A., 331, 332, 338
 Kopell, N., 8
 Koulakov, A., 156, 157
 Kourtzi, Z., 289, 290
 Kovács, I., 284, 285, 287
 Kronecker, L., 206
 Kuffler, S., 62, 123
 Kuramoto, Y., 8

L

Lagrange, J.-L., 32, 303, 348, 349
 La Métrie, J.O. de, 45
 Lamme, V.A.F., 62
 Landman, K.A., 342
 Laplace, P.-S. de, 7
 Le Bihan, D., 15
 Leadbetter, M.R., 196
 Lebesgue, H.-L., 197
 LeCun, Y., 99
 Lê Dũng Tráng, 12
 Lee, H.Y., 26, 113, 204
 Lee, T.S., 26, 41
 Legendre, A.-D., 349
 Leibniz, G.W., 12, 25, 35
 Leopold, D., 257
 Le Penec, E., 309
 Levi-Civita, T., 199
 Lie, S., 23, 25, 30, 74
 Link, B., 331, 334
 Lipschitz, R., 199
 Liu, G.B., 239, 241
 Locke, J., 35
 Logothetis, N., 257
 Longo, G., 16
 Lorenceau, J., 14, 16, 17, 290, 291
 Lotka, A.J., 31
 Lund, J.S., 124, 142, 279
 Luo, L., 331, 340–342
 Lyle, S., 19

M

Mézard, M., 17
 MacPherson, R., 11, 199
 Maffei, L., 17, 62
 Malach, R., 278
 Maldonado, P., 206, 208, 214, 217
 Malgrange, B., 11, 219, 323
 Mallat, S., 15

Malsburg, C. von der, 16
 Mamassian, P., 15
 Manchon, D., 363
 Margulis, G.A., 356
 Mariño, J., 209, 211
 Markram, H., 127
 Marr, D., 26, 64, 82, 307
 Marre, O., 33
 Martinet, J., 219
 Martinez, L.M., 66
 Martins, R., 331
 Mather, J., 11
 Mathieu, É., 7
 McCulloch, W., 36
 McLaughlin, 66
 McLaughlin, D., 208, 297
 McLaughlin, T., 331, 338
 McLoughlin, N., 144
 McNaughton, B.L., 102
 Merker, J., 23
 Merleau-Ponty, M., 8
 Meunier, C., 15
 Meunier, J.-G., 18
 Meyer, Y., 15, 85
 Michor, P., 17
 Miikkulainen, R., 203, 245
 Milleret, C., 16, 261
 Milnor, J., 11
 Mingolla, E., 295
 Mishkin, M., 49
 Mitchell, J., 6, 356
 Mitchison, G., 279
 Moisan, L., 286
 Moiseev, I., 7, 17, 357
 Molden, S., 104
 Montgomery, R., 5, 348
 Morel, J.-M., 16, 17, 286
 Morgan, A., 74
 Morlet, J., 88
 Morse, M., 10
 Moser, E., 2, 103, 104
 Moser, M.-B., 103
 Mostow, G.D., 356
 Mountcastle, V., 123
 Mullen, K., 84, 113, 124
 Muller, Bostock, Taube, 103
 Muller, R., 103, 108
 Mulligan, K., 18
 Mumford, D., 6–8, 11, 12, 15, 16, 26, 41,
 359

N

Nadal, J.-P., 15, 95, 97

Nadel, L., 24, 103
 Nagumo, J., 52
 Nathans, J., 59
 Necker, L.-A., 251
 Neitz, J., 59
 Neitz, M., 60
 Neumann, J. von, 36
 Newsome, W., 319
 Newton, I., 12, 30, 31, 36, 61
 Nicolis, G., 12
 Niebur, E., 145, 146
 Ninio, J., 6
 Nirenberg, S., 97
 Nobel, A., 2, 103, 123

O

O'Reagan, K., 15
 Obermayer, K., 243
 Ogawa, S., 71
 Ohki, K., 139, 162, 164, 212, 213
 O'Keefe, J., 1, 2, 103
 Olavarria, J., 261
 O'Leary, D.D.M., 331, 337, 342
 Olshausen, B., 99
 Orban, G.A., 319

P

Pachoud, B., 17
 Pakdaman, K., 15
 Palacios, A., 317
 Palmer, L., 73
 Pansu, P., 5, 356
 Parent, P., 295
 Pauls, S., 17
 Pearson, R.A., 335
 Pennek, X., 17
 Perry, V.H., 49
 Peterhans, E., 311
 Petit, J.-L., 18
 Petitot, J., 219
 Pettigrew, J.D., 239
 Peyré, G., 91, 309
 Pham, F., 12
 Pitts, W., 36
 Plancherel, M., 7
 Planck, M., 7, 359
 Poincaré, H., 10, 199
 Polat, U., 293
 Poli, R., 17
 Polimeni, J., 123, 206, 210
 Pontryagin, L.S., 356

Prigogine, I., 12, 15
 Prochiantz, A., 14, 18, 331

R

Ramón y Cajal, S., 45
 Ramis, J.-P., 219
 Ranck, J., 107
 Rangan, A., 297
 Raymond, P., 57
 Reeb, G., 355
 Reed, T., 74
 Reese, B., 331, 334, 335, 341
 Reid, C., 139
 Reimann, M., 17
 Renault, B., 15
 Ribot, J., 165, 229, 231
 Rice, S.O., 197, 199
 Richards, L., 331
 Riemann, B., 10, 22, 23, 174
 Rifford, L., 17
 Ringach, D.L., 124
 Risler, J.-J., 356
 Rochefort, N., 261
 Rockland, K.S., 279
 Roe, A., 257, 310, 313
 Romagnoni, A., 165
 Rosenstiehl, P., 15
 Rositi, F., 18
 Rossi, F., 359, 360
 Rothschild, L.P., 356
 Roy, J.-M., 17, 18
 Ruelle, D., 11, 12
 Russell, B.A.W., 38
 Ryckman, T., 18, 22

S

Sachkov, Y., 5, 7, 357
 Sagi, D., 293
 Sajda, P., 285
 Salam, A., 12
 Salama, G., 139, 140
 Sanguinetti, G., 359
 Sarti, A., 5, 6, 8, 16, 17, 173, 194, 359, 360
 Sayag, D., 52
 Schmidt, J.-P., 18
 Schnabel, M., 203, 206
 Schulte, S., 331, 337
 Schummers, J., 144, 210, 211, 214
 Schwartz, E.L., 12
 Schwartz, L., 80
 Scicolone, G., 331, 338

Seriès, P., 62, 290
 Sethian, J., 16, 17
 Shah, J., 8, 364
 Shapley, R., 77, 266, 316
 Shelley, M., 208
 Sherrington, C., 45
 Sheth, B.R., 315
 Shinbata, H., 320
 Shmuel, A., 158, 320
 Shulz, D., 62
 Sigalotti, M., 17
 Sigman, M., 306
 Simoncelli, E., 99
 Singer, W., 17
 Sirosh, J., 245
 Sirovich, L., 224, 225
 Sit, Y.F., 310
 Smale, S., 11, 12
 Smith, B., 18
 Snippe, H.P., 124
 Snodderly, D.M., 319
 Somers, D.C., 65
 Souriau, J.-M., 22
 Sperry, R., 338
 Spruck, J., 14
 Stein, E.M., 6
 Stengers, I., 12
 Stone, M., 360
 Strichartz, R., 5, 356
 Stryker, M., 223
 Stumpf, C., 22
 Sutton Weeks, M., 19
 Swindale, N., 6, 203, 246, 320

T

Tailor, D., 267
 Tanaka, S., 320
 Tani, T., 239
 Tanigawa, H., 316
 Taube, J., 107
 Taylor, B., 133
 Taylor, J., 199
 Teissier, B., 12, 16, 18
 Ter Haar Romeny, B.M., 74
 Tessier-Lavigne, M., 331, 338
 Thompson, E., 12, 317
 Thom, R., 8, 10, 11, 13, 25, 28, 249, 263
 Thorpe, S., 62
 Toet, A., 364
 Tondut, Y., 292, 293
 Tootle, R., 119, 224
 Toretti, R., 22

Touboul, J., 165
Toulouse, G., 16
Trepel, C., 223
Tritsch, D., 14
Ts'o, D.Y., 278, 310
Turing, A., 10, 28, 338

U

Uglesich, R., 224
Ullman, S., 303
Ungerleider, L., 49, 117

V

Van Essen, D.C., 319
Van Hooser, S., 125, 221
Varela, F., 12, 14, 317
Veca, S., 18
Veltz, R., 17
Vershik, A., 6, 356
Violi, P., 18
Volterra, V., 31
Von der Heydt, R., 311, 314
Von Neumann, J., 360

W

Waddington, C.H., 11, 338
Wall, C.T.C., 12, 103
Wallet, G., 18
Wandell, B.A., 71
Wardlaw, C.W., 342
Weierstrass, K.T.W., 168
Weliky, M., 137

Westheimer, G., 293
Weyl, H., 22, 28
Whitney, H., 10, 199
Wickerhauser, M.V., 94
Wiener, N., 36
Wiesel, T., 65, 67, 123
Wilson, H., 8, 255
Winaver, J., 71
Winterer, C., 19
Witkin, A., 13
Wittgenstein, L., 317
Wolf, F., 4, 17, 173, 196, 203
Wolpert, L., 338
Wörgötter, F., 65
Wu, S., 137

X

Xu, X., 147, 239

Y

Young, R.A., 75
Young, T. (Young-Helmholtz), 60, 61, 264
Yu, H., 147, 149, 248, 259

Z

Zahaf, M.B., 363
Zariski, O., 12
Zeeman, C., 11
Zeki, S., 313, 316
Zhang, J., 137
Zhu, W., 226
Zucker, S., 17, 304

Subject Index

A

Abduction, 148–150, 286
Absorption peak, 59
Action potential, 7, 31, 51, 64, 194
Adhesion factor, 332
Age-related Macular Degeneration (AMD), 58
Alanine, 60
Amacrine, 51, 57, 58, 334, 335
Amino acid, 54, 60, 333
Angular momentum, 179
Angular sector, 146
Anterior fissure, 114
Aperture problem, 147
Apical, 334, 335
Apparent contour, 317
Association field, 288, 290–293, 295, 306, 347
Axon, 6, 46, 51, 55, 59, 203, 259, 261–263, 276, 331–333, 337, 338, 340, 341
Axon guidance, 342, 347

B

Basal, 334, 335
Ben-Yishai, R., 66
Bias factor, 325
Binocular disparity, 114, 252, 310, 319
Binocular rivalry, 253, 255, 323
Biocytin, 278, 279
Bipolar, 51, 56, 58, 335
Blow-up, 347
Brodman areas, 115
Bundle, 4, 6, 132, 137, 165, 218, 219, 223, 259, 277, 310, 316, 321, 347

C

Cadherin, 332
Calcarine, 45, 115
Calcium, 51, 212, 213
Callosal, 119, 259–261
Carnot group, 5, 356
Cartan connection, 21
Cautic, 6
Checkerboard, 71, 94, 95
Chemotaxis, 343
Chemotrophic, 332
Coaxiality, 5, 357–359
Collagen, 333
Colour processing, 59, 264, 267
Cone, 53, 54, 58–60, 77, 332, 333, 337, 339, 342
Connection, 5, 7, 9, 22, 27, 30, 31, 66, 77, 101, 113, 114, 119, 133, 157, 158, 164, 198, 205, 209, 239, 260, 262, 275, 277–281, 289–291, 293, 295, 297, 331
Connection flagellum, 57
Contact element, 33, 131, 133, 280, 289, 293
Contact geometry, 5, 12, 304
Contact plane, 5, 299, 302, 303, 348, 355
Contact structure, 5, 6, 30, 193, 281, 292, 298, 299, 301, 302, 347
Contralateral, 46, 48
Convolution, 8, 76, 82
Corpus callosum, 259, 264, 338
Cortico-cortical, 78, 114, 298
Curl, 171, 177
Current, 35, 38, 180, 186, 192, 193, 253, 333
Curvature, 6, 17, 94, 102, 199, 292, 294, 303–305, 357
Cut locus, 6, 41

Cyclic guanosine monophosphate (cGMP),
56
Cytochrome oxidase, 257, 258
Cytoskeleton, 332, 338

D

Dehydrogenase, 54
Dendritic tree, 45, 59, 332, 338
2-deoxyglucose, 139
(de)oxyhemoglobin, 139, 141
Dextro rotatory, 158, 164
Differentiable, 11, 129, 134, 135, 171, 249,
309, 323
Differential form, 30
Dipole, 4, 226, 228, 229, 231, 238, 295
Direction, 2, 4, 9, 16, 27, 65, 72, 103, 105–
108, 114, 126, 128, 129, 137, 221,
228, 240, 246, 250, 280, 281, 287,
289, 291, 309, 311, 319–321, 327,
328, 347
DNA, 58, 332
Double-opponent cell, 265, 266

E

Early vision, 331
Efficient coding, 95
Eigenvalue, 193, 205
Eigenvector, 86, 205
Elastica, 7, 358, 359
Engel structure, 6, 305
Engrafted, 128, 137, 221, 259
Entoptic, 40
Entorhinal, 102, 105
Ephrin, 331, 333
Equilibrium potential, 51
Equipotential, 150, 154, 226, 227, 244
Euclidean, 2, 5, 21–23, 26, 98
Excitatory, 62, 107, 209, 210, 275, 358
Expanding ring, 71, 72
External space, 22, 24, 253, 265, 325

F

Feedback, 26, 66, 113, 293
Feedforward, 66
Fibre bundle, 6, 129–131, 133, 134, 138,
215, 217, 275, 277, 321
Fibronectin, 333
Field line, 150
Figure (Gestalttheory), 40
Firing, 63, 97, 104, 105, 256, 293
Flavoprotein, 226

Fluorescein, 261
fMRI, 71–73, 105, 119, 254, 289
Foliation, 173
1-form, 219, 347
Format, 9, 101, 133
Fovea, 114, 117, 128, 142
Frame, 102
Functional architecture, 1, 3–9, 13, 27, 31,
33, 42, 74, 83, 101, 113, 114, 125,
128, 130, 132–134, 210, 250, 259,
275, 281, 295, 298, 303–305, 310,
311, 319, 321, 331, 347, 356

G

Gabor patch, 74, 282, 283, 290, 291, 293
Ganglion cell, 3, 49, 51, 57, 58, 61, 82, 83,
97, 98, 128, 307, 334, 337–341
Gauge, 9, 25, 200
Gaussian derivative, 72, 74
Gaussian field, 195, 196
Generalized symmetry axis, 41
Generic, 158, 232, 249, 286, 287
Geodesic, 6, 7, 16, 24, 303, 347–350, 356–
358
Gestalt, 6, 9, 12, 17, 26, 39, 42, 113, 284,
285, 289, 293, 295, 306, 309, 363
Glial cells, 45, 50
Good continuation, 38, 39, 284, 285, 306,
363
Granular, 120
Grid cell, 2, 102–106
Ground (Gestalttheory), 8
Ground state, 8
Growth factor, 331
Guidance factor, 332, 333

H

Heat equation, 13, 14, 90
Heat kernel, 7, 82, 356, 360, 361
Heisenberg group, 5, 6, 74, 300, 348, 354,
359
Helmholtz equation, 4, 144, 180, 187, 195,
347
Hemifield, 118, 260, 264
Hemiretina, 46, 48, 259
Hemitarget, 122
Hermitian, 7, 74
Hessian, 95, 193
Heterogeneous, 263, 264
Hippocampus, 24, 102, 103
Homogeneous, 103, 145, 204, 263, 264, 278,
316, 317

Horizontal cell, 51, 52, 58, 59
 Hypercolumn, 123, 126–128, 130, 165, 229

I

Illusory contour, 6, 16, 39, 40, 114, 303, 304, 311, 348, 357, 358
 Imaging, 25, 35, 41, 138–140, 143, 147, 151, 210, 212, 228, 229, 257, 291, 296
 Immunoglobulin, 331
 Impulse response, 79
 Inhibitory, 107, 158, 203, 209, 276
 Inner segment, 57
 Integrability, 135, 289, 292, 298, 299, 303, 349
 Integral curve, 160, 173, 299, 303, 348, 355
 Integrin, 333
 Interkinetic Nuclear Migration (INM), 334
 Internal space, 253, 318, 323
 Internal state, 7, 253, 318, 319
 Inverse problem, 266, 286
 Ion channels, 26
 Ipsilateral, 46, 48, 49, 107, 242, 261, 338
 Isoleucine, 60

J

Jacobian, 176, 177, 195, 197
 Jet, 4–6, 11, 13, 17, 133–135, 137, 170, 303–306

K

Kernel, 5, 7, 81, 211, 216, 299, 301, 306, 347
 Koniocellular, 46, 265

L

Laminar, 125, 334, 336
 Laminin, 332, 333
 Lateral Geniculate Nucleus (LGN), 3, 15, 46, 48, 49, 62–64, 83, 121, 295, 337
 Lattice, 4, 103–105, 138, 144, 147, 150, 157, 158, 162, 164, 165, 170, 180, 201, 215, 217, 218, 220, 222, 239, 244, 347
 Layer, 4, 15, 27, 48–51, 61, 65, 101, 120–122, 124, 126, 128, 141, 217, 242, 246, 250, 257, 262, 276, 278, 310, 319, 336, 340
 Left translation, 360
 Legendrian, 4, 135, 292, 298, 299, 355
 Levo rotatory, 158, 164
 Linear momentum, 31, 33, 59

Locally trivial, 147
 Long range, 39, 275
 Loop, 102, 120, 234
 Low level vision, 26

M

Magnocellular, 46, 48, 49, 59, 120, 224, 265
 Mammal, 46, 59, 124, 259, 335
 Map, 2, 4, 11, 71, 104, 108, 122, 128, 129, 139, 144, 156, 171, 176, 194, 203, 211, 212, 220, 223, 232, 240, 247, 259, 261, 264, 277, 301, 315, 320, 322, 330, 338, 340, 341, 347
 Medial fissure, 119, 322, 337
 Mesh, 4, 103, 104, 144, 180, 218, 239, 347
 Minimal Discharge Field (MDF), 62, 295, 296
 Mitochondria, 226
 Momentum, 180, 193, 194
 Morphogen, 10, 332, 334, 342, 343
 Motion, 6, 24, 33, 49, 63, 65, 72, 102, 106–108, 128, 137, 147, 250, 291, 296, 311, 319, 334
 mRNA, 58, 332
 Multiscale, 8, 13, 80–82, 84, 87, 92, 95, 97, 128

N

Netrin, 331, 333
 Neural coding, 50, 62, 97, 306
 Neural network, 5, 26, 31, 35–37, 203, 338
 Neuroepithelial, 334
 Neurogenesis, 6, 319, 334, 335, 347
 Neurogeometric, 5, 9, 28, 34, 40, 305, 310, 347
 Neurogeometry, 1, 3, 5, 8, 9, 14, 16, 25–28, 32, 33, 38, 42, 170, 173, 303, 356
 Neuropilin, 333
 Neurotransmitter, 45
 Noise, 62, 64, 82, 94, 95, 97, 141, 255, 359
 Noradrenalin, 335
 Normal form, 4, 173, 205
 Nucleus, 51, 55

O

Occipital, 45, 115, 119
 Ocular dominance, 4, 48, 123, 137, 152, 204, 240, 242, 245, 246, 250, 261, 310
 Oculomotor field, 115, 117
 Opponent cell, 60, 266
 Opsin, 54, 57, 59, 335

Optic chiasm, 46, 337, 338
 Optic tract, 338
 Orientation (hyper)column, 3–5, 33, 41, 63,
 65, 66, 68, 72, 79, 84, 102, 123, 124,
 128, 134
 Orientation selectivity, 4, 65, 69, 99, 125,
 171, 201, 202, 208, 209, 261
 Outer segment, 53, 56, 57

P

Pale stripe, 310, 311, 316
 Parallel transport, 5, 278, 279
 Partial Differential Equation (PDE), 14, 16,
 205, 206
 Parvocellular, 46, 48, 49, 79, 120, 224, 265
 Pathway, 6, 45, 46, 48, 49, 62, 114, 117, 275,
 319, 334, 337, 338
 Peripherin, 57
 Periphery, 54, 114, 142
 Phase field, 4, 150, 170, 171, 173, 180, 181,
 183, 193, 194, 201, 217, 347
 Phenylalanine, 60
 Phosphodiesterase, 54
 Photon flux, 52, 77
 Photoreceptor, 3, 27, 51–54, 56, 59, 77, 95,
 101, 336
 Pinwheel, 3, 4, 138, 145, 149–151, 156–159,
 162, 164, 171, 173, 174, 180, 193–
 195, 200, 201, 203, 204, 206, 207,
 209, 210, 212, 215, 217, 218, 220,
 222–226, 229, 239, 242, 257, 285,
 311, 322, 329
 Place cell, 2, 102, 103, 105, 125
 Plane wave, 86, 144, 180, 200
 Plexiform, 51, 335
 Plexin, 333
 Point processor, 133, 134, 283, 302, 304
 Pontryagin's maximum principle, 349, 356
 Population coding, 33, 85, 103, 123, 138,
 143, 201
 Pulvinar, 46, 310, 319

Q

Qualitative discontinuity, 95, 266, 318

R

Radius, 144, 145, 230, 234, 307, 353
 Receptive field, 3, 13, 27, 50, 51, 54, 59, 62,
 73, 75, 79, 81, 120, 123, 147, 149,
 247, 294, 313

Receptive profile, 3, 17, 42, 57, 62, 64, 66,
 68, 70, 71, 95, 96, 266, 306, 319
 Recording, 104, 119, 139, 143, 206, 208,
 211, 303
 Reeb, G., 301
 Reproducing kernel, 89
 Retina, 3, 15, 45–47, 49–51, 61, 62, 83, 98,
 121, 264, 336, 341
 Retinal, 3, 35, 50, 54, 60, 82, 97, 123, 131,
 247, 280, 296, 317, 318, 334, 335,
 337, 340, 342
 Retinaldehyde, 54, 337
 Retinal Pigment Epithelium (RPE), 334, 335
 Retinotopy, 119, 120, 125, 204, 279
 Rhodamine, 261
 Rhodopsin, 53, 54, 58
 Ridge, 249
 RNA, 332
 RNA polymerase, 332
 Rod, 31

S

Saddle point, 147, 162, 171, 180, 193, 257,
 277
 Scale, 8, 13, 33, 74, 81, 87, 90, 94, 105, 123,
 210, 220, 266, 292
 Selection, 59, 67, 253
 Semaphorin, 331, 333, 338
 Serine, 60
 Serotonin, 335
 Short range, 66, 333
 Signal decorrelation, 95
 Simplicity, 302
 Simply connected, 188
 Single-opponent cell, 266
 Singularity, 8, 10, 12, 94, 158, 161, 162, 166,
 174, 193, 200, 232, 234, 235, 237,
 239, 324, 326, 329
 Skeleton, 8, 41, 194
 Skew curve, 298, 303
 Smooth, 14, 23, 129, 133, 135, 165, 166, 171,
 231, 298, 348
 Sparse, 97–99, 101, 158
 Spatial frequency, 4, 62, 63, 74, 99, 119, 123,
 125, 223, 231, 236, 240, 242, 246–
 248, 282
 Spike, 51, 57, 62, 63, 95, 208–210, 295, 297
 Spike train, 62, 63, 68, 78, 223, 292
 Splitting factor, 325
 Stereopsis, 114, 254, 310
 Stochastic resonance, 255
 Stratified, 11, 199

Striate cortex, 45
 Stripe, 248, 257, 310, 311
 Subcortical, 46, 107
 Subjective contour, 311, 315
 Sub-Riemannian geometry, 5–7, 16, 17, 26, 30, 304, 353, 356, 357
 Subthreshold, 62, 210, 214, 296, 297
 Superior colliculus, 46, 319, 337, 339
 Superior temporal sulcus, 319
 Symmetry breaking, 10, 205, 263
 Synaptic bouton, 262, 279, 281
 Synaptic gap, 45
 Synaptic Integration Field (SIF), 295, 296
 Synaptic weight, 7, 8, 203

T

Target, 29, 58, 59, 118, 140, 173, 293, 294, 317, 323, 332, 337–341
 Temporal frequency, 77
 Thalamic, 46, 62, 66, 319
 Thick stripe, 257, 310, 311, 316, 319
 Thin stripe, 311, 316
 Threonine, 60
 Threshold, 26, 51, 62, 143, 210, 253–255, 287
 Tonic, 47
 Topological index, 149, 158, 174
 Topological universality, 165, 166, 170, 231, 234
 Transducin, 54
 Transfert function, 62, 64
 Transition zone, 259, 261, 262, 264
 Transversality, 4, 11, 223, 239, 242, 244, 246, 248–250, 310, 347
 Trichromat, 59–61, 264, 266
 Triple point, 149, 161, 162, 166, 167, 171

Tropism factor, 332
 Tuning, 67, 85, 124, 125, 171, 209, 242, 324, 327, 328
 Tyrosine, 60, 340

U

Unfolding, 173, 323–325, 327
 Universal unfolding, 11, 254, 323, 324, 326, 327

V

Vertical meridian, 119, 140, 264, 280, 281
 Vesicle, 336, 337
 Visual area, 5, 27, 33, 46, 107, 113–116, 119, 296
 Visual system, 1, 6, 39, 46, 50, 62, 77, 80, 82, 83, 113, 133, 151, 165, 317, 331
 Visuotopy, 223
 Vitronectin, 333
 Vorticity, 174, 176–179, 192, 193

W

Wavefront, 354
 Wavelength, 59, 180, 195, 204, 205, 282, 293, 316
 Wavelet, 3, 7, 15, 85, 87–89, 92, 94, 307, 309, 310
 Wave number, 4, 77, 78, 144, 180, 181, 194, 205
 Wave vector, 180, 181, 194, 196, 198, 205–207
 Weight, 158, 203, 295
 Wiring, 4, 133, 244, 248, 332