Journal of Mathematical Imaging and Vision 15: 55–82, 2001

Complexity, Confusion, and Perceptual Grouping. Part I: The Curve-like Representation

BENOIT DUBUC
Espace Courbe, 642 de Courcelle, suite 303 Montreal, Canada H4C 3C5
benoit@espacecourbe.com

STEVEN W. ZUCKER
Center for Computational Vision and Control, Departments of Computer Science and Electrical Engineering, Yale University New Haven, CT 06520-8285, USA
zucker-steven@cs.yale.edu

Abstract. Intermediate-level vision is central to form perception, and we outline an approach to intermediate-level segmentation based on complexity analysis. We focus on the problem of edge detection, and how edge elements might be grouped together. This is typical because, once the local structure is established, the transition to global structure must be effected and context is critical. To illustrate, consider an edge element inferred from an unknown image. Is this local edge part of a long curve, or part of a texture? If the former, which is the next element along the curve? If the latter, is the texture like a hair pattern, in which nearby elements are oriented similarly, or like a spaghetti pattern, in which they are not? Are there other natural possibilities? Such questions raise issues of dimensionality, since curves are 1-D and textures are 2-D, and also of complexity. Working toward a measure of representational complexity for vision, in this first of a pair of papers we develop a foundation based on geometric measure theory. The main result concerns the distribution of tangents in space and in orientation, which serves as a formal basis for the concrete measure of representational complexity developed in the companion paper.

Keywords: perceptual organization, segmentation, tangent maps, image curves texture, texture flow, complexity

1. Introduction

> *How complex or simple a structure is depends critically upon the way we describe it. Most of the complex structures found in the world are enormously redundant, and we can use this redundancy to simplify their description. But to use it, to achieve the simplification, we must find the right representation.*
> Herbert A. Simon (1968)

Although edges provide the foundation upon which much of visual processing is built, it is curious that the definition of an edge remains unsettled. The usual one, which has roots back to Ernst Mach in the mid-19th century, is that edges are bright-to-dark transitions (or vice-versa), and that these transitions can be detected from measurements provided by (discrete approximations to) linear differential operators. Mach, for example, preferred the Laplacian operator, and Marr has revived modern interest in it. However, this view has turned out to be problematic, and much frustration in visual shape analysis derives from the unsatisfactory nature of early edge processing. In brief, the edges returned by standard tools do not correspond in any systematic way to the boundaries of objects; they arise in other ways as well, such as in textures. Nevertheless, it is boundaries—not bright-to-dark transitions—that are needed for visual shape analysis. Our goal in this paper is to propose a geometric foundation for early vision in an attempt to get right to the matter of what

comprises a boundary. Our proposal is based on concepts from differential geometry and geometric measure theory. In the process of achieving this goal, we shall overview our recent research in edge detection, and shall concentrate on the geometric problems associated with interpreting the role an edge might be playing in an image. The geometric approach in the end leads us to conclude that one should think not in terms of edges, which are image-defined constructs, but rather in terms of tangents, which are differential-geometric constructs. It is this abstraction that reveals the connection back to boundaries.

Our contribution is divided into two papers. The first one is, in a sense, the more theoretical, and it motivates the philosophy of our approach, while the the second is more relevant to applications, in that it contains experiments. But both papers are linked, in that they both cover material from geometric measure theory, with the first building on Hausdorff measures and the second on Minkowski constructions. These different approaches to dimension theory lead to different types of results: we use Hausdorff (and Besicovich) constructions in the first paper to prove "tangent separation theorems", or statements in principle about how many tangents can arise in different positions and in different orientations. This structure is fundamental to understanding certain aspects of the geometry of grouping, or separating boundaries from textures. These "density" results in turn motivate the algorithms developed in the second paper that are based on the *Minkowski sausage* and more practical area computations. Because the material is abstract and dimension theory is not commonly used in computational vision, both papers are written in a discursive, informal style. Examples to stimulate the reader's intuition are spread throughout, as are references to the mathematical literature. We believe this motivation for the role of complexity analysis in computational vision is perhaps the most important contribution of these papers.

1.1. Motivation

A classical problem in the design of a general purpose artificial vision system is the localization and description of image curves (edges or bars). For instance, imagine a dark cube against a white background. The task of early vision is to abstract a description of this cube sufficiently rich to enable its recognition, while segmenting it as a figure from the background. Such a description must certainly involve the bounding contour around this cube, and it is the task of boundary detection to recover this contour.[1] Complexity and dimension issues arise immediately. From an intuitive geometrical point of view, surfaces are boundaries of solids, lines are boundaries of surfaces and points are boundaries of lines, as was pointed out by the French mathematician Poincaré [40]. Because the cube may subtend a large visual angle covering an enormous number of pixels, "re-presenting" it by its edges reduces the amount of information tremendously while keeping the essence of the information about the object. Line drawings are another example of abstraction of information in which the essence of the scene is kept and reduced to its minimal expression. Thus, it made intuitive sense thirty years ago to begin to build computer vision systems by developing algorithms that would extract these edges and segment images automatically. At the same time, neurophysiology was providing important conceptual support for these techniques. The result is that now, among the stages for the processing of visual information, edge detection is one of the best understood.

1.1.1. Structure Detection in Early Vision. Edge detection implies however a basic problem in perceptual grouping: once the local structure is established, the transition to global ones must be effected. To illustrate, imagine standing on an edge element in an unknown image, as in Fig. 2(a) or (b). Is this edge element part of a curve, or perhaps part of a texture? If the former, which is the next element along the curve? If the pattern is a texture, is it a hair pattern (in which nearby elements are oriented similarly) or a spaghetti pattern (in which they are not)? These questions are in part about complexity since curves are "simpler" than textures, and in part about dimensionality, since some discontinuities are 0-D, curves are 1-D, and textures are 2-D. In this paper a complexity measure designed to address these questions will be proposed. The ultimate goals are to show, in the context of curve detection, that the choice of representation and support for the grouping process is an important issue, and to provide a means of making an appropriate decision regarding the choice of representations.

Measures, dimensionality, and complexity are coupled concepts, and the relationships between them are important practically as well as theoretically. Measures for curves might include their length, the number of components (cardinality), or the area covered. However, the situation is more subtle than this, as is

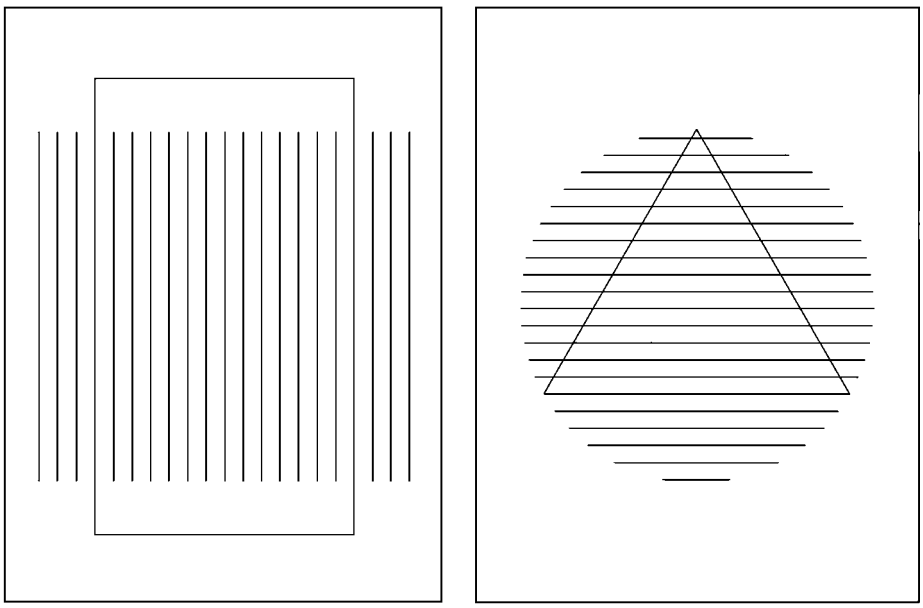

Figure 1. Two Kanizsa patterns. (a) a pattern due to Kanizsa [25] (b) circle/triangle pattern due to Galli and Zama [15]. Both these examples illustrate the need of using different representations for integration. Why does the texture "absorb" portions of the rectangle in (a) and of the triangle in (b)? In both cases the grating seems to predominate over the perception of the closed curve.

illustrated in several examples. The first example is taken from a classical demonstration by the Italian psychologist Kanizsa [25]. Figure 1(a), in which a pinstriped surface appears to be occluding a rectangle, demonstrates that curves, or sets of curves, can actually connote either the outline of objects (as in the rectangle) or surfaces (the pinstripes). Closer examination reveals that the rectangle is actually continuous through the surface, suggesting that visual inferences somehow group the pinstripes together and ignore the fact that the rectangle is the longest curve in the image. Neither the length of the curve, defining the rectangle, nor the number of components, defining the pinstripes, is dominant. Figure 1(b), a triangle, is similarly camouflaged within a horizontal grating, and once more illustrates what Kanizsa [25] called the "social conformity of a line".

Our second example is an image of a statue (Canova's sculpture "Paolina", see Fig. 2(e)). The result of an edge operator [24] at a given scale is shown in Fig. 2(f), and the problem of integrating local information raises the following observations. For the shoulder region (Fig. 2(c)), the underlying object is simple and a curve representation seems appropriate to group the edge elements. If we examine instead regions subtending part of the hair structure (Fig. 2(d)), then choosing a curve representation and walking along a hair would lead very quickly to confusion, since it will be difficult to know on which part of the curve one is. A texture representation in this case seems more appropriate.

1.2. Problem Definition

The leading question of this paper is: given the output of edge/line detectors at a given scale and for a given resolution, how can these be grouped together? This clearly involves a local-to-global transition which has been described as "collecting individual edge points together to form continuous curves" [7]. However, it assumes that edge points should be grouped into curves, a common assumption in computer vision (see also [3, 16, 28, 42, 50–52]). But local edges can arise from other image structures, such as a texture (hair, or fields of grass, for instance). We therefore question this assumption that edges should be grouped only into curves, and rather seek to determine which representation should be chosen for the grouping, together with the dimension of its support. While the problem of contours "bleeding" into textures is classical [2], even recent systems still have the problem [29].

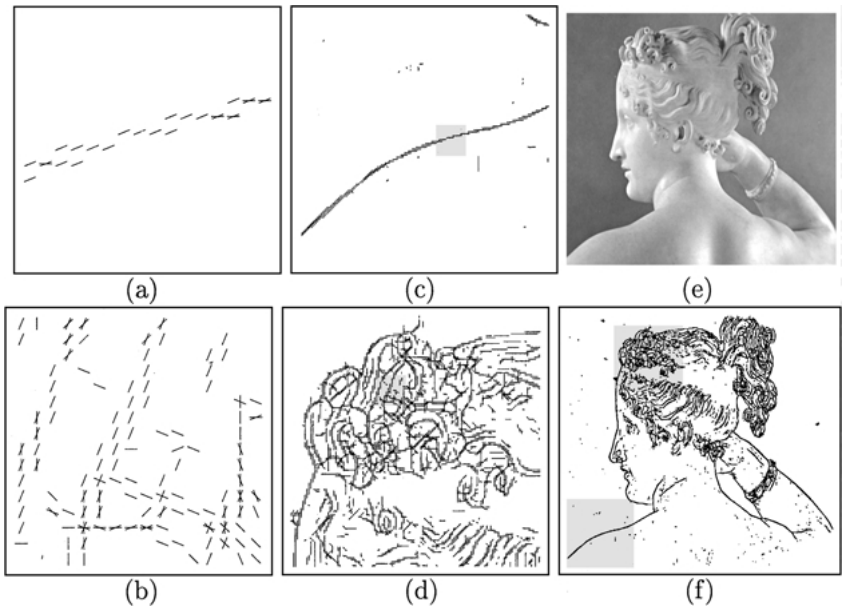

Figure 2. The subtlety of "walking through" a tangent map: (a and c) curves, (b and d) texture. Moving from right to left, the gray shaded areas are expanded to show the need for different representations to support the grouping of local edge elements. Integrating the responses of local edge detectors in the hair region is problematic. By what principles should the tangents be grouped?

A second view of grouping is that it is a noise problem [54]. Since there are spurious responses from the local detectors, a global estimation procedure is necessary to eliminate them (Kalman filtering, for instance). A third is that it is simply an image-domain phenomenon linked to scale. Since larger operators have more image support, they should be less susceptible to such local variations. However, they are also more likely to average across features belonging to different objects.

This paper questions all these assumptions and will try to shine light on the grouping and representation problems through arguments of complexity. The starting point of our investigation will be the search for edges, and for positive or negative contrast lines. The local structure will be given by the output of an edge/line detector [24] sometimes followed by a few iterations of relaxation labeling [23, 56]. The process that will be described in this paper in some cases would decide where those of Cox et al. [7], David and Zucker [8] or Mumford [36] could be used, i.e. where a curve support is appropriate and over which extent. In the case when the support indicates a surface (textures for instance), then approaches such as [53] and [26] for integral curves, or [30] for oriented texture characterization should rather be used.

1.3. Complexity in Early Vision?

Complexity has entered computer vision systems in two ways. First, there is the classical notion of computational complexity and proofs that problems are in given classes (e.g., labeling blocks world diagrams is NP-complete). Second, there are issues of efficency, and intermediate-level representations are commonly thought to be required to improve it (e.g., segmenting images into regions, to reduce search domains). However, we believe the relationship between complexity and representation is deep in another sense, and the following examples are intended to illustrate this.

Example 1 (The loop and the spaghetti). Consider two images, one of a circle, and the other of a plate of spaghetti. Two tasks can be envisioned: one, to draw the curves and the other, to follow them. To draw the circle, the algorithm can be very simple: repeat n times: take 1 pixel step, rotate $2\pi/n$. In the case of the spaghetti pattern: repeat n times: take a long step, rotate random amount. Now, suppose one would like to reproduce the two patterns. In the first case the complexity is bounded, while in the second case we need remember the various random rotations.

Another issue is the one of following the path. In the case of the circle, the task is very easy. In the case of the spaghetti pattern, the task is much more complicated and confusing. The fact that there are many branching points makes the path not unique: there are many different paths. In fact it leads to a combinatorial explosion of possibilities. The "curve" representation then fails to be efficient. Following the pattern leads to integration but the representation that needs to be used varies from one case to the next.

Example 2 (Pick-up sticks). Suppose we are only working with line segments and we want to find out when the segment representation is no longer appropriate. Three instances are shown in Fig. 3. Now, let us make a parallel with the game called "Pick-Up Sticks".[2] Complexity will be related to the difficulty of picking up a stick without moving the other sticks. Clearly this is dictated by the tweezers: two scales are involved, one for the width and one for the length. Thus, in Fig. 3(a), the task is trivial since there is only one segment. In Fig. 3(b), it is a little more difficult, but still very easy. In Fig. 3(c), the task is hard. The task is not simply a matter of the number of line segments. Even with only two segments, the task can be hard if the sticks are parallel and close one to another, or if they cross at a small angle.

Example 3 (Frequency). An example taken from the representation of numbers states that, at first glance, it is very unlikely that somebody would be able to distinguish between the numbers 99999999999999 and 9999999999999 without serially scanning each digit. The same point can be made with grating patterns, suggesting that our organization of understanding for visual patterns is richer than the simple enumeration of the individual curves. In Fig. 4, we reconsider the Kanizsa pattern to determine whether the percept changes when a pinstripe is added or removed.[3] For instance, compare Fig. 1(a) and Fig. 4(a): which has 20 and which has 21 lines? The distinction is very hard to make, but the one between Fig. 4(a) and Fig. 4(b) is immediate. Rounding the corners would be noticeable, however, even though these events are isolated.

Taken together, Examples 1 and 2 show how algorithmic complexity relates to computational vision. The first example stressed the complexity of communication of a pattern and of a visual task. The second example introduced some of the central principles we will develop in this paper. Most importantly, it showed how the class of all patterns could be partitioned through a complexity measure into equivalence classes of equally hard tasks. There are many games of Pick-Up Sticks that are equivalent in terms of difficulty—i.e., in terms of complexity—even though the particular arrangement of the sticks may differ enormously.

How can one relate algorithmic complexity to vision? Recalling the intuitive description of algorithmic complexity presented earlier, let us build the following experiment. Take three persons: Robert, Bruno, and Veronica. Bruno has a pattern that he shows to Veronica who then describes it to Robert. From this description, Robert reproduces the pattern. He then shows his reproduction to Veronica. If Veronica cannot distinguish the reproduced pattern from the original in a bounded amount of time, then we will say that the *representational complexity* is bounded by the length of the description. Representational complexity would thus

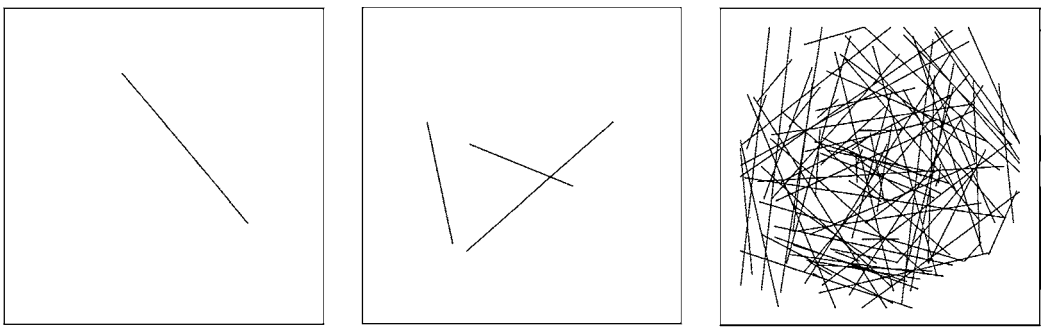

(a) one single line *(b) three lines* *(c) one hundred lines*

Figure 3. Line segments or texture? Three examples of sets composed of line segments. Notice how difficult it is to follow every path in (c).

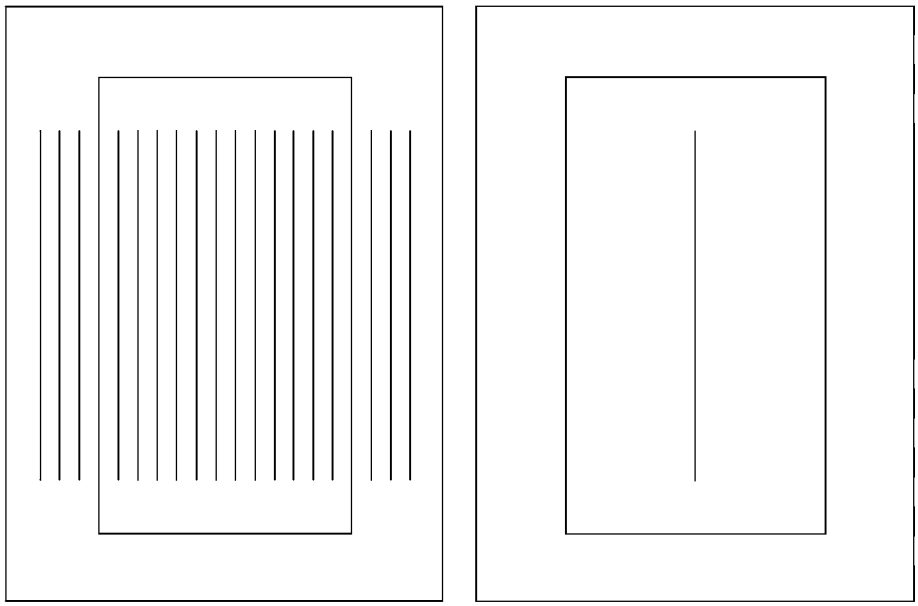

Figure 4. Kanizsa patterns again but this time with different grating frequencies. Notice how easy it is to tell these two apart but how difficult it is to differentiate between (a) and Fig. 1(a).

build equivalence classes of patterns. Relating back to our Pick-Up Sticks example, within each equivalence class, the visual tasks are equally hard. The question now is how to build a measure of complexity that will define these equivalence classes of patterns, and then how to assess it.

Returning to our examples, the last one showed that, for some patterns, approximation is sufficient. In the case of grating patterns, the exact number of lines and their exact location might not be relevant. There is an obvious difference between a pattern with one segment and one with two segments provided the lines are (i) sufficiently long and (ii) sufficiently apart from one another. In the grating part of the Kanizsa pattern, the difference between n and $n + 1$ lines is irrelevant for our percept if n is large enough and if the lines are reasonably distributed. Compare the gratings between Fig. 1(a) and Fig. 5 and try to say at a quick glance the difference between the two. The difference between 20 and 200 lines would however be noticeable (provided again that they are reasonably positioned), again reinforcing the suggestion of equivalence classes of patterns that would be indexed by complexity measures.

The take-home message from all these examples is that, within a single representation and within a general setup, integration from local to global representations is intractable. However, integration is a key step in moving from an image-centered representation to an object-centered one. It has been done successfully in controlled environments where one knows the complexity of the scene *a priori*, or assumes the complexity to be within some bounds. The *blocks world* is an example of such a constrained environment. How does one build a general theory of integration for edge detection?

The core of this paper will show that different actions in the integration stage should be taken depending on the context. First, we need to choose and define both an *intermediate representation* and a *complexity measure*. Given this representation, the local information can be integrated only if the underlying object is simple enough. Identifying such simple cases (in our case, curves will be simple compared with textures) will facilitate the transition from early to intermediate levels of vision. If the complexity of an object at a certain scale and for a particular representation exceeds some value, then two choices could be made:

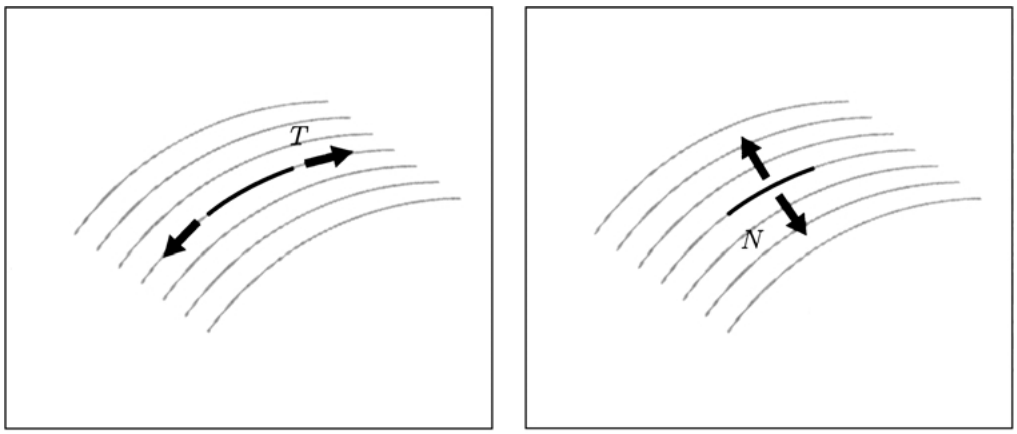

Figure 5. The main idea: examine the rate of growth of oriented dilations, in the normal direction N to test density ("space-fillingness") and in the tangential direction T to test continuity.

1. keep the same scale but adopt another representation, or
2. change scale.

Keeping status quo, i.e. keeping the sole curve representation under the current scale, is bound to failure.

1.4. Mapping Complexity and Indexing Representation

As curve detection is central to vision, what is required is a measure of the complexity of curves, and our specific goal in this paper is to propose one. We will show how it successfully handles the Kanizsa and the Paolina examples, among others. It is based on an intermediate representation—the *discrete tangent map*, or a discretized tangent field—and a consequence of our analysis is that such intermediate representations are necessary for a proper segregation of curve-like patterns that fill areas, from curve-like patterns that extend mainly along their length and also from dust patterns (discontinuities, for instance). These representational differences capture the first stages of segmentation; but via complexity analysis not pixel grouping.

The complexity measure we derive will be tailored to discrete "curve-like" sets such as those we seek in edge detection. The basic idea will be to look in two directions: in the tangential direction to assess continuity and in the normal direction to assess density of the object within a local extent (Fig. 5). This will lead to two complexity indexes, that we call the *normal* and *tangential complexity indexes*, and constitute the basis for our *complexity map*. Although the tangential complexity captures the same line of thought as previous researchers such as Ullman and Mumford, it is the normal complexity that provides some further insight into segregating textures from curve patterns. Both must be used together.

A Note on the Experiments. A series of examples will be carried over throughout this paper and the companion paper. These will serve to illustrate the points we are trying to make and clarify different concepts related to the algorithms involved. We will constantly refer to these as

1. the Kanizsa pattern: Fig. 1(a);
2. the Ullman pop-out pattern (Fig. 6(a));
3. the Ullman hidden pattern (Fig. 6(b)); and
4. the Paolina image: Fig. 2(e).

These examples are not all equivalent. Some are sets that can be described easily, others are images in which the underlying curve-like sets must be inferred. For the Kanizsa pattern the structure was known *a priori*, so the tangent map will be discretized from the continuous one precisely. The discrete tangent map of the Ullman patterns will be obtained by analyzing the corresponding images as described later. The Paolina image is our example of a full grey level image for which we will

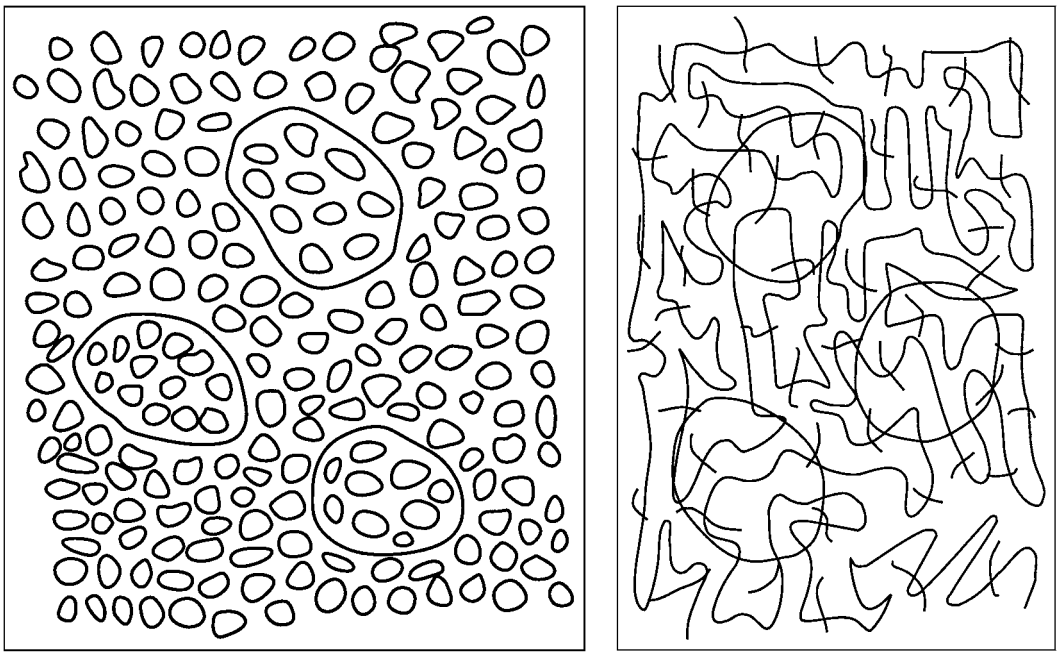

(a) Ullman pop-out (b) Ullman hidden

Figure 6. Curve complexity and context. In these two figures developed from Ullman [48], notice how in (a) the big circles pop-out as opposed to (b) where they tend to blend more into the scene.

infer edges and interpret the inference process. Note the difference between the Ullman patterns, in which we will try to detect *negative contrast lines* (i.e. dark lines on a light background), and the Paolina image, for which we are seeking *edges* (i.e. boundaries between light/dark regions).

In all cases we will project the set or the image into the unit square, so our numbers (scale for grouping and resolution) will be relative to the unit square. Two different scales will be considered: (i) the *scale of the operator* σ, expressed in pixels and (ii) the *scale for complexity analysis* δ. δ is linked with the *spatial extent* Ω over which grouping should be considered. *Resolution* is the inverse of the size of smallest element of the digitized grid. Finally, whenever a variable is "hatted", $\hat{\omega}$, for instance, it means it is expressed in image coordinates (pixels).

2. Boundaries and Their Detection

What, precisely, is a boundary? Mach's intuition was founded in photometry, and is based on the observation that since the reflectance of an object is typically different from the background, the object will project to an image region with different intensity from those regions to which the background projects. This intuition breaks down, however, when the object has internal structure.

To take advantage of this internal structure, consider (informally) a smooth surface such as a ball. The surface of the ball has a 2-dimensional tangent plane everywhere, and this tangent plane projects smoothly into the image except along the locus of positions at which the viewing vector just grazes the ball; at these positions the tangent plane becomes singular, and collapses to a single (one-dimensional) tangent. That is to say, the tangent plane folds away from the line of sight. An integral curve through these tangents defines the boundary of the ball. Thus, each tangent defines a proper edge element.

To see that edge elements can arise from other configurations, consider a bulge of cloth. Edges can now arise from within the cloth, with the tangent plane disappearing along the top of the bulge and then reappearing on the other side. The tangent plane transition across a bulge (from cloth-to-cloth) is thus different from that across the ball, because it never

re-appears. Moreover, this type of fold can end within the body of the cloth, as in the wrinkles around one's shirt sleeve or the muscles in one's shoulder.

2.1. *Whitney's Classification Theorem*

Whitney has classified the generic maps from smooth surfaces into smooth surfaces, and has shown that these two situations are the only ones that can occur generically (see [27]). Generically means that the configuration does not change with small changes in viewpoint. He referred to them as the fold and the cusp (the position where the cloth fold disappears). Furthermore, by considering a view of a mountain range it is easy to observe that generic boundaries are not globally smooth; rather, they are punctuated by discontinuities. Thus we have two observations regarding boundaries:

1. *local edge elements arise in different contexts* such as exterior boundaries or interior folds. Textures, of course, provide further examples of complex contexts for edge elements.
2. *discontinuities should be expected in bounding contours* from occlusion relationships.

The structure of this map is developed formally in Huggins and Zucker [21].

2.2. *Normal Intensity Configurations and Logical/Linear Operators*

We begin with the first observation to provide formal and computational support for the idea that local edge elements can be interpreted as tangents to contours. The first advantage is that, with the tangent direction established, one can examine the intensity profiles in the tangent and the normal directions separately. Note that most other edge operators average these together, although not necessarily uniformly. Informally we observe immediately that, in many situations and at a finite scale,

- *Normal direction*: The fold condition in Whitney's Theorem often takes on a different intensity profile for a bounding edge (which involves a dark-to-light transition) than for an interior fold (which involves a light-to-dark-to-light transition) or vice versa. This latter profile is often called a line, and must be separated from the former. Standard linear operators blur both together. Examples of this (light-dark-light) interior fold configuration can been seen in the shoulder musculature of the Poalina image. It is curious how such surface markings (and line operators) relate to edge operators. Considering a "line" as back-to-back edges suggests how surfaces can "fold" into a crack from both sides (in the normal direction to the edge). The need for a separate operator is clear numerically: back-to-back edge operators cannot be evaluated without overlap. In general, such observations stress the importance of separating analysis along tangential and normal directions, a point that will recur in these papers.
- *Tangential direction*: The differential interpretation demands that continuity conditions exist (that is, that the limit of one point approaching another must exist). This corresponds to continuity contraints on the intensity pattern.

A necessary condition for a tangent to exist is that one or the other of the above intensity and continuity conditions must be satisfied. We have developed a class of non-linear local operators, called logical/linear operators, that use Boolean conditions to test whether the above structural criteria are met. "Edge operators" are separated from "line operators", and lines can arise either in light-dark-light conditions (typical of a crack or a crease) or dark-light-dark conditions (typical of a highlight). Note that both of these latter conditions refer to surface markings, rather than to surface boundaries.

The tangent interpretation also leads to a solution for representing orientation discontinuities, but requires a more modern view of a discontinuity. While it is classically the case that no unique tangent exists at an irregular point on a curve, it is also the case that multiple tangents can be defined to exist there. Informally, taking a limit into the point of discontinuity from one side yields one tangent; a limit from the other side yields another. The two tangents span a 2-dimensional space; it is precisely this difference in dimensionality between the tangent spaces at regular and singular points that we exploit. The logical/linear operators are arranged into columns, so that multiple values of tangent orientation are possible at each point. Those points at which multiple tangents are established are the discontinuities.

Code for a system that implements the three logical/linear operators [24] in columns, one for edges, one for bright lines, and one for dark lines, is available from via anonymous ftp;[4] Researchers are invited

to experiment with this system. Results are shown in Fig. 2.

2.3. Edge Maps, Tangent Maps, and Grouping

The result of our logical/linear edge operator [24] at a given scale is shown in Fig. 2(f) on the image of the statue ("Paolina", see Fig. 2(e)) and raises the following observations. For the shoulder region (Fig. 2(c)), the underlying object is simple and a curve representation seems appropriate to group the edge elements. If we examine instead regions subtending part of the hair structure (Fig. 2(d)), then choosing a curve representation and walking along a hair would lead very quickly to confusion, since it will be difficult to know on which part of the curve one is. A texture representation in this case seems more appropriate. The remainder of the paper will deal with these issues.

Although we have been speaking of the the logical/linear operators as if they signal tangents, actually they only return a distribution of positions at which the above intensity signatures are obtained. This is a very different notion firm that of tangent. As we show in the next two sections, however, there is a geometric-measure-theoretic definition of tangent that suggests a connection between them.

3. Curve-like Sets

> *In the case of the world on a human scale you don't care much about problems involving infinities or infinitesimals, whereas you certainly care whether something is line-like or point-like.*
>
> Jan. J. Koenderink (1990)

An elementary analysis textbook [43] states that "*a continuous curve is usually thought as the path of a continuous moving point and this rather vague notion is often felt to carry with it the even vaguer attribute of* '*thinness*' *or* '*one-dimensionality*'". This definition of a curve is bound to the one of dimension. The first part of this section will informally address the issue of dimension to give a feel for what will follow. Then the Jordan definition of a curve will be presented, stating that a curve Γ is the range of a continuous map α from an interval I to Euclidean space (typically $\mathbf{R}^2$ or $\mathbf{R}^3$). It will remind us of concepts such as the one of *length*, denoted $L(\Gamma)$, and the one of best linear local approximation, namely the *tangent* to Γ at $x = \alpha(t)$, denoted $T(x)$. Within these are embedded the notions of local representation (the tangent) and global measure (the length), which are tied together through the map α. Unfortunately, however, the map α is not given for general vision problems, but must be inferred. More abstract notions of a curve are thus required. A generalization of the previous ideas through measure theory to maps that are not necessarily smooth will thus be introduced. The resulting *curve-like sets* and their associated parametrization-free tangents will constitute a much better basis for our needs and will represent the core apparatus for our work.

When inferring curve-like objects from images, one is confronted with issues such as discretization, quantization and choice of scale. One of the main characteristics of curve-like sets is that they lie in the continuous domain. How can one make a parallel between these sets and the output of a finite set of "edge detectors"? An answer to this will lead to a definition of what we will call *discrete curve-like sets*. The underlying philosophy will be somewhat different since we are trying to infer the set from images. We start by getting the discrete tangent map and then infer the curve-like set from its local properties. The local structure of the discrete curve-like sets will be obtained through edge detection. From our previous description of the edge we will be able to present the implications of our choice of operators and decision method.

3.1. Elements of Dimension Theory

The history of the various notions of dimension involves some of the greatest mathematicians of the turn of the century: Poincaré, Lebesgue, Brouwer, Cantor, Peano, Hilbert, just to name a few. That history is very closely related to the creation of space-filling curves and early fractals [39]. Hausdorff remarked that the problem of creating the right notion of dimension is very complicated. People had an intuitive idea about dimension: the dimension of a set, say E, is the number of independent parameters (coordinates), which are required for the unique description of its elements. This turned out to be incorrect, as the counterexamples of Cantor and Peano showed. In this section, we will review three different approaches to defining the concept of dimension. These will help in setting up the framework for our own intuitive requirements (presented last) which we call the *curve assumption*.

3.1.1. Poincaré's Cut Dimension. Poincaré's cut dimension is inductive by nature and starts with a point.

A point has dimension 0. Then he observes that "...if to *divide* a continuum it suffices to consider as cuts a certain number of elements all distinguishable from one another, we say that this continuum is of *one dimension*; if on the contrary, to divide a continuum it is necessary to consider as cuts a system of elements themselves forming one or several continua, we shall say that this continuum is of *several dimensions*" [40].

From this definition we get that a segment has dimension 1 since it can be split by a point (dimension 0). The same happens for a circle, since it can be disconnected using a pair of points (dimension 0). The unit square has dimension 2, since it needs a line (dimension 1) to get disconnected. Finally, the cube has dimension 3, since can be disconnected using a plane (dimension 2).

Poincaré's definition has the advantage of being very intuitive and easy to grasp. It will be used later to motivate the proofs of the *tangent separation theorems* (Section 3.4). This idea of dimension also formed the basis of the now accepted one developed by Menger and Urysohn [22]; namely that

1. the empty set has dimension -1,
2. the dimension of a space is the least integer n for which every point has arbitrarily small neighborhoods whose boundaries have dimension smaller than n.

3.1.2. Lebesgue Covering Dimension. The Lebesgue covering dimension is the most frequently used in point set topology to define the notion of dimension for a topological space [10, 37]. It consists in covering the set E with little disks (such as those used in point set topology) and then focusing on the maximal number of disks in the cover which have non-empty intersection. This is called the *order of the cover*.

An object E has covering dimension n provided any cover admits an open refinement of order $n+1$, but not of order n. Taking the line segment as an example, it is easy to see that the order of the cover cannot exceed 2, leading to a topological dimension smaller than 1 (see Fig. 7).

Another equivalent definition would be that the topological dimension of a set E is the smallest integer k such that, for all $\epsilon > 0$, there exists a covering A_i of E by closed sets of diameter $\leq \epsilon$, with the following property: the intersection of any $k+2$ sets A_i is empty.

3.1.3. Measure and Dimension. In their classical monograph, Hurewicz and Wallman [22] presented different approaches to defining dimension. The one we will adopt here will associate the concepts of measure and dimension. An object will be called one-dimensional if it has length (one-dimensional measure), 2-dimensional if it has an area (2-dimensional measure), 3-dimensional if it has a volume (3-dimensional measure), and so on. The measure that will

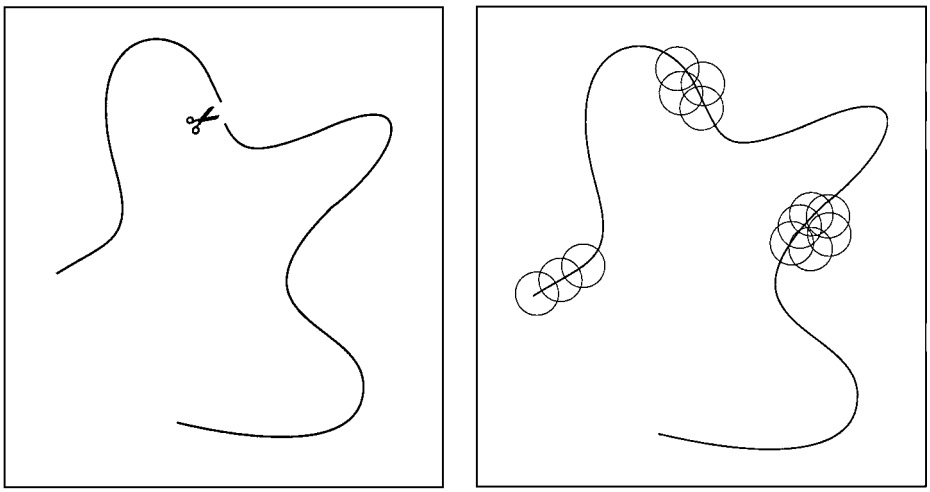

(a) toward cut dimension (b) toward covering dimension

Figure 7. Illustration of Poincaré's cut and Lebesgue's covering dimension of a line. (a) A simple line disconnected by a single point; (b) covers of orders 2, 3, and more. Since it is possible to cover the curve with a cover of order 2, its dimension is not bigger than 1 from the definition.

be used in this case is the one developed by Hausdorff. It will also allow us to study in more detail the local structure of the set.

Before introducing geometric measure theory, we will review the concepts of local linear approximation and rectifiability for Jordan simple curves in $\mathbf{R}^2$. This is most easily approached through differential geometry. We will then try to extend it to a more general class of objects that we will call *curve-like sets*. It will be shown that the concept of rectifiability allows us to derive important properties and constraints on the local structure of the sets to be studied.

3.2. *Elementary Differential Geometry: Jordan Curves*

The most common definition of a curve is the one of Jordan, namely that a curve Γ is the range of a continuous map γ from an interval I to Euclidean space (typically $\mathbf{R}^2$ or $\mathbf{R}^3$). In elementary differential geometry, this definition precedes two other basic notions, namely the *length* and the best local linear approximation or the *tangent* to Γ at $\gamma(t)$.

Definition 1 (Jordan curve [47]). A *curve* Γ in $\mathbf{R}^2$ is the range of a continuous function $\gamma(t) = (\gamma_1(t), \gamma_2(t))$ defined on an interval $[a, b]$. If γ is an injection, the curve Γ is called *simple*. Its endpoints are $\gamma(a) = A$ and $\gamma(b) = B$. The mapping γ is called a *parametrization* for the curve Γ.

Remark 1. In this subsection, when using the word "curve", we mean a simple Jordan curve. Although too restrictive a definition for the type of patterns detected in edge detection, it will be useful to study the basic concepts. We will later widen the definition to include the types of sets sought after in computer vision.

Two examples of Jordan curves are shown on Fig. 8. In (a), a simple Jordan curve, i.e. a Bézier curve with 7 control points. In (b), a Jordan curve that is not simple (since the curve cuts itself). This is a drawing due to Pablo Picasso that we adapted from Mendès-France [31], and it illustrates the fact that the definition of "curve" is indeed very large. The following definition builds an equivalence relation between different parametrizations of a curve.

Definition 2 (Fréchet equivalence [6]). A mapping $\Gamma : \gamma = \gamma(t), t \in I$ is said to be Fréchet equivalent to another mapping $\Gamma_1 : \gamma = \gamma_1(s), s \in I_1$ if for every $\epsilon > 0$ there exists a homeomorphism h_ϵ from I_1 to I such that $|\gamma(h_\epsilon(s)) - \gamma_1(s)| < \epsilon$ for all $s \in I_1$. This defines an equivalence relation between Γ and Γ_1, and then we write $\Gamma \sim \Gamma_1$.

3.2.1. A Curve in the Small: Tangent. The notion of a tangent, the best linear approximation to a curve at a point, is key in the study of curves in the small. Intuitively, it is defined as follows:

Definition 3 (Tangent to a curve [19]). If Γ is a simple (parametrized) curve and x is a point on Γ, the *tangent*

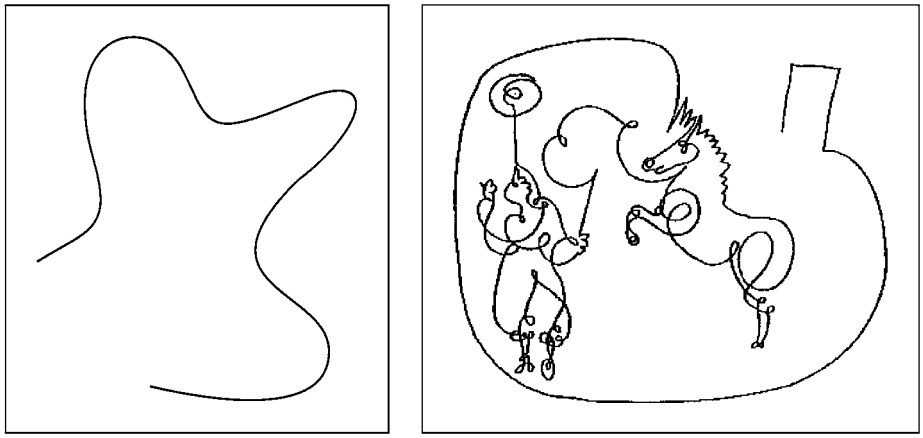

(a) a very simple curve (b) a Jordan curve that is not simple!

Figure 8. Examples of Jordan curves. In (b) we reproduced a drawing due to Picasso entitled "Le Jongleur" (adapted from Mendès-France, [31]).

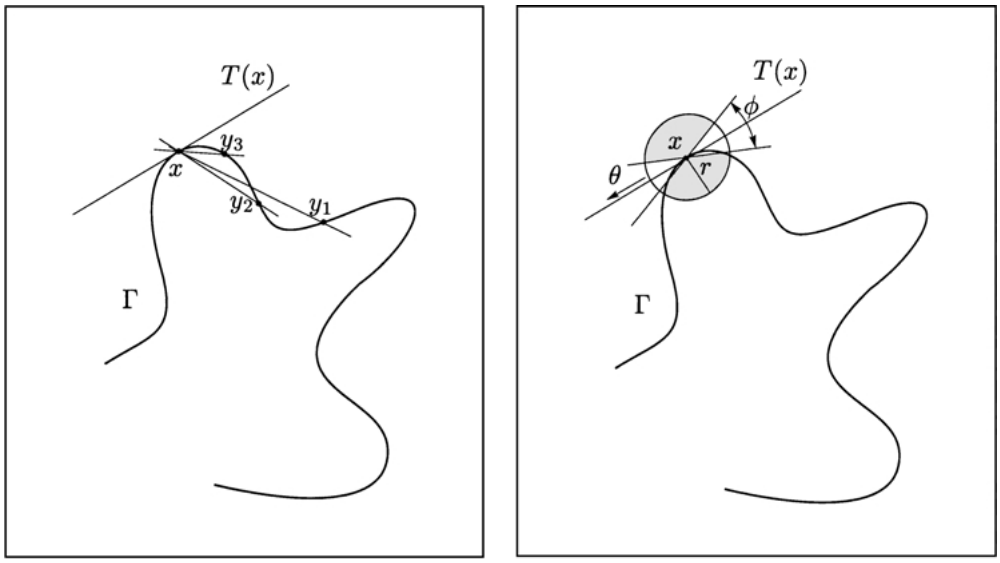

(a) the tangent $T(x)$ to Γ at x (b) Besicovitch tangent

Figure 9. This figure illustrates in (a) the intuitive definition of a tangent T to a curve Γ at a point $x = \gamma(t)$. Take a sequence $\{y_1, y_2, \ldots\}$ of points on the curve converging to x. Draw the lines passing through y_i and x. The "limit line" gives us the tangent $T(x)$ to the curve at x. In (b) we illustrate the parametrization-free definition of the tangent, looking at a cone that shrinks around the point x (presented later in the text).

$T(x)$ to the curve at x is the limit (if it exists) of the straight line passing through x and y when $y \in \Gamma$ and $y \to x$. (see Fig. 9(a).)

The reason why we called this an intuitive definition is the fact that the limit $x \to y$ is not always defined. Furthermore, in applications such as computer vision, for one, the parametrization is exactly what one is trying to infer. Thus a more general model is required. However, when the parametrization is given, the tangent is the first derivative of the map γ. We will see in the next section that this local notion is tightly linked to the global one of *length*. In this highly structured situation, there is a clear model of the local-to-global transition.

3.2.2. The Length of a Curve. We mentioned previously that a curve was a set extending along its length. How can we compute such a length? And even before that, does such a measure exist for a particular set? One formalisation of the intuitive definition of the length of a parametrized simple curve can be derived from the ancient device of inscribed polygons:

Definition 4 (Partition and its norm [47]). Let $[a, b] \subset \mathbf{R}$, then a partition $\mathcal{P}$([a, b]) is a set of points $\{t_0, t_1, \ldots, t_n\}$ such that

$$a = t_0 < t_1 < \cdots < t_n = b$$

The *norm* $|\mathcal{P}|$ of a partition is defined as being

$$|\mathcal{P}| = \max(t_{i+1} - t_i), \quad i = 0, 1, 2, \ldots, n-1$$

Given Γ, a simple curve in $\mathbf{R}^2$, let $\mathcal{P}_n$ be a sequence of partitions such that $\lim_{n\to\infty} |\mathcal{P}_n| = 0$. Define

$$L(\mathcal{P}_n, \Gamma) = \sum_{i=0}^{n-1} |\gamma(t_{i+1}) - \gamma(t_i)|$$

From the triangle inequality we see that the insertion of new points of subdivision will produce an increase in $L(\mathcal{P}_n, \Gamma)$.

Definition 5 (Jordan length [4]). If $L(\mathcal{P}, \Gamma)$ is bounded for all dissections $\mathcal{P}$ of $[a, b]$, the *length* of a curve Γ in the Jordan sense is as follows:

$$L(\Gamma) = \sup_{\mathcal{P}} L(\mathcal{P}, \Gamma).$$

It is sufficient to obtain the length from a limit of a sequence of partitions:

Theorem 1 (*Length as a limit* [47])**.** *If* $(\mathcal{P}_n)$ *is a sequence of partitions such that* $\lim_{n\to\infty} |\mathcal{P}_n| = 0$ *then*

$$L(\Gamma) = \lim_{n\to\infty} L(\mathcal{P}_n, \Gamma).$$

This last theorem implies that the length is independent of the choice of polygonal approximations as these get finer and finer.

Definition 6 (Rectifiability [4]). A curve Γ is called *rectifiable* if it has finite length in the Jordan sense.

Remark 2. It is interesting to see that the word *Rectifiable* derives from the Latin word *rectus* which means *Straight*. In French, the expression *Rectification d'une Courbe* means calculating the length of a curve as if it were a straight line segment. Not being able to unfold a curve into a straight line segment implies that the curve is not rectifiable.

A few more results need to be mentioned. These provide the invariance properties one would like for the calculation of length and link a local notion (the tangent) to a global one (the length):

1. *Length and rigid body motions*: the length of a curve is invariant under rigid body motions, i.e., translations and rotations;
2. *Length and arclength* [45]: if the mapping Γ is differentiable, then

$$L(\Gamma) = \int_a^b |\Gamma'(t)|dt$$

   i.e., the notion of length as just presented corresponds with the one of arclength in differential geometry;
3. *Length and parametrization* [6]: length $L(\Gamma)$ is independent of the parametrization (Fréchet-independent), i.e., $\Gamma \sim \Gamma_1$ implies $L(\Gamma) = L(\Gamma_1)$.

3.3. Curve-like Sets in Geometric Measure Theory

As we mentioned before, we need a wider class of objects as an underlying model for curve recovery. Simple curves are too restrictive since they do not allow multiple curves and various kinds of discontinuities that are key in our description and understanding of the visual world. Even some of the simplest patterns could not be expressed by the Jordan definition. Mathematicians have however described a wider class of objects which would be more suited for our needs and these are called 'curve-like sets' (regular sets with finite positive Hausdorff measure). Instead of considering a mapping, we will rather consider sets. The notion of length will be kept implicit: it will be one of the

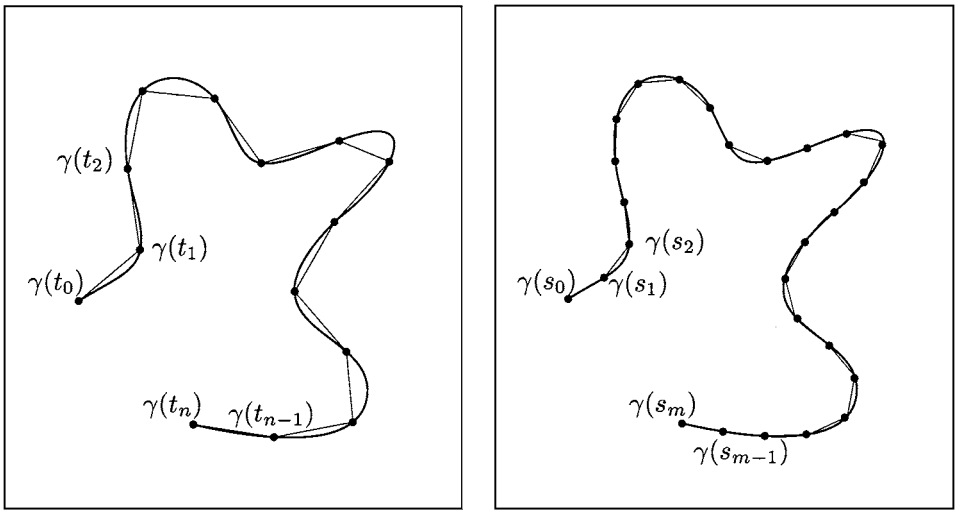

(a) coarse partition (b) fine partition

Figure 10. Two different partitions leading to two different approximations of the length of a simple Jordan curve. If the curve is rectifiable, the finer the partition gets, the more accurate the estimate of length will be.

basic requirements for these sets to exist. This section is a brief introduction to curve-like sets. It starts with the definition of the Hausdorff measure, and leads to the one of 1-sets, or those with finite length. Then we present density properties which will provide a hierarchy for one-dimensional sets. It is within this hierarchy that the type of sets to be considered throughout this paper arise: the *curve-like sets*. A fine study of the local structure of curve-like sets will lead in the next section to constraints on the distribution of tangents and discontinuities.

3.3.1. Hausdorff Measure. One way to compute the length, area or volume of an object is to use the Hausdorff s-dimensional measure $\mathcal{H}^s$, where, in the case of a smooth rectifiable curve, $s = 1$, in the case of a surface, $s = 2$ (classical references for this are [13, 41], but one can also look at [11], a more readable presentation). Consider the problem of defining the length $\mathcal{H}^1$ of a set E in the plane. Hausdorff's idea was to cover the set with small circles and to take the sum of the diameters (Fig. 11). If the balls are restricted to be smaller than some given value $\delta > 0$, and if the 'most economical' covering is chosen, we get an approximation of the length of the set at resolution δ. Allowing arbitrary covers, instead of covers by balls, gives us an outer measure, and for $\delta > 0$ we write

$$\mathcal{H}^1_\delta(E) = \inf \sum_i |U_i|$$

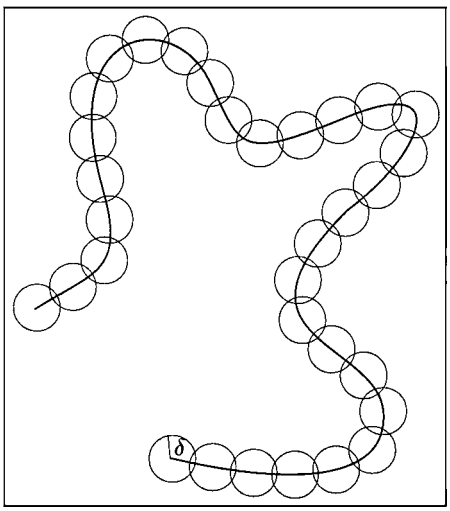

Figure 11. Parametrization-free approximation of the length of a set by a covering with δ-balls: a first step to the calculation of the Hausdorff measure.

where $|U|$ is the *diameter* of U, (i.e., $|U| = \sup\{|x - y| : x, y \in U\}$) and $\{U_i\}$ is any sequence of sets of diameter less than δ covering E. The infimum here is taken over all (countable) δ-cover $\{U_i\}$ of E. It can be shown that $\mathcal{H}_\delta(E)$ increases as δ decreases, therefore:

Definition 7 (Hausdorff measure [11]). The *one dimensional Hausdorff* measure of E is given by

$$\mathcal{H}^1(E) = \lim_{\delta \to 0} \mathcal{H}^1_\delta(E) = \sup_{\delta > 0} \mathcal{H}^1_\delta(E).$$

Since no confusion will arise in this paper, we will write $\mathcal{H}$ for $\mathcal{H}^1$.

Remark 3. One can show that the Hausdorff measure is in fact a *measure* in the measure theoretic sense [11]. It then can be used to define the notion of "almost". In this paper we will use the terms "for almost all x in E" and "almost everywhere" (sometimes denoted a.e.). This means that the property applies for all $x \in E$, except maybe on a (very small) set G with $\mathcal{H}(G) = 0$. When writing $\mathcal{H}$-almost everywhere, or $\mathcal{H}$-a.e., we want to emphasize that this is with respect to the Hausdorff measure $\mathcal{H}$, and not with respect to another measure (the Lebesgue measure, $|\cdot|_1$, for instance). The term "almost nowhere" (which is used only once in this document) means that the property holds at most on a set of measure 0.

One might wonder if the Hausdorff measure coincides with the Jordan length for simple Jordan curves, or with the Lebesgue one-dimensional measure for Lebesgue-measurable subsets of the real line.

Theorem 2 (*Hausdorff, Lebesgue and Jordan measures* [11])**.** *If* Γ *is a curve, and* E *a Lebesgue measurable subset of* $\mathbf{R}$ *then*

1. *the Jordan length* L *and the Hausdorff measure coincide, i.e.,* $\mathcal{H}(\Gamma) = L(\Gamma)$.
2. *the Lebesgue measure* $|\cdot|_1$ *and the Hausdorff measure coincide, i.e.,* $\mathcal{H}(E) = |E|_1$.

Hausdorff measure permits a classification of sets. One of its most popular uses is as the basis for *Hausdorff dimension*, an important construction which was used as an abstract formulation for the concept of dimension. Among other results, it led to the definition of a class of sets called 'fractals' [31]: sets with non-integer Hausdorff dimension. In this section, we

are not considering non-rectifiable curves and rather concentrate on those with positive but finite measure:

Definition 8 (1-set [11]). An $\mathcal{H}$-measurable set E with $0 < \mathcal{H}(E) < \infty$, will be called a 1-*set* (originally called 'linearly measurable set' by Besicovitch [1]).

3.3.2. Basic Density Properties. The notion of densities for sets will be used in the definition of the local approximation of a set (tangent). Intuitively, densities indicate the local measure of a set compared with the expected measure [11]. The definition is as follows:

Definition 9 (Density [11]). Let $B_r(x)$ denote the closed ball of centre x and radius r. The *upper* and *lower densities* of a 1-set E at a point $x \in \mathbf{R}^2$ are defined as

$$D_u(E,x) = \limsup_{r\to 0} \frac{\mathcal{H}(E \cap B_r(x))}{2r}$$

and

$$D_l(E,x) = \liminf_{r\to 0} \frac{\mathcal{H}(E \cap B_r(x))}{2r}$$

respectively. If $D_u(E,x) = D_l(E,x)$, we say that the *density* of E at x exists and write $D(E,x)$ for the common value.

Example 4. To better understand the previous definition, let us consider the following subset of $\mathbf{R}^2$ (Fig. 12):

$$E = \{(u,v) \in B_r((0,0)) : u^2 = v^2\}$$

has density

$$D(E,x) = \begin{cases} 1 & \text{for } x \in E \setminus \{(0,0), \text{end points}\}, \\ 0 & \text{for } x \notin E, \\ \frac{1}{2} & \text{for } x \in \{\text{end points}\}, \\ 2 & \text{for } x = (0,0) \end{cases}$$

Note that in the last example the density is zero when outside the set and non-zero otherwise. In fact, one of the most interesting results about densities is that the density is zero almost everywhere outside the set:

Proposition 1 [11]. *If E is a 1-set in $\mathbf{R}^2$, then*

1. *$D(E,x) = 0$ at $\mathcal{H}$-almost all x outside E, and*
2. *$2^{-1} \le D_u(E,x) \le 1$ at almost all $x \in E$.*

The last proposition is used mainly in the structure study of one-dimensional sets. Requiring $D_u(E,x)$ to be greater than zero insures that we are almost surely on the set E.

Definition 10 (Regular and irregular sets [11]). A point $x \in E$ at which $D_u(E,x) = D_l(E,x) = 1$ is called a *regular* point of E; otherwise it is called an *irregular* point. A 1-set is said to be *regular* if $\mathcal{H}$-almost all of its points are regular, and *irregular* if $\mathcal{H}$-almost all of its points are irregular.

Remark 4. Examples of irregular 1-sets in $\mathbf{R}^2$ include constructions similar to the one of the Cantor set [11]. An example taken from Morgan [35] defines $E \subset \mathbf{R}^2$ by starting with an equilateral triangle and removing triangles at different scales. Start with E_0, a closed equilateral triangular region of side 1 (Fig. 13(a)). Let E_1 be the three equilateral triangular regions of side 1/3 in the corners of E_0 (Fig. 13(b)). In general let E_{j+1} be the triangular regions, a third the size, in the corners of the triangles of E_j. Finally, let $E = \bigcap E_j$ (an approximation is shown in (Fig. 13(e)).

E is a 1-set since the projection of each E_j onto the x-axis is the unit segment, therefore the projection of $E = \bigcap E_j$ is also the unit segment which gives that $\mathcal{H}(E) \ge 1$. As for the other inequality we have that E_j is covered by 3^j equilateral triangles of side $\left(\frac{1}{3}\right)^j$, therefore $\mathcal{H}(E) \le 1$, confirming that E is a 1-set. The

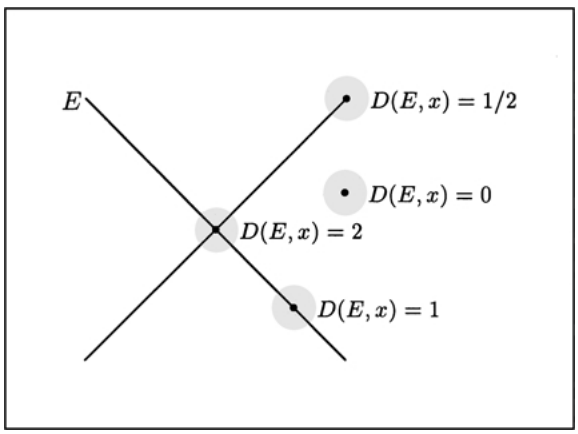

Figure 12. Example of an set where the density is not always 1: the truncated cone in $\mathbf{R}^2$. The cross is the set studied. The grey regions are places where we wanted to focus attention for the density.

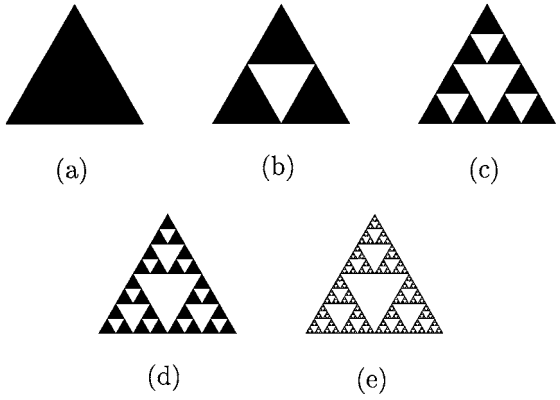

Figure 13. Example of an irregular 1-set in $\mathbf{R}^2$: the Sierpinski triangle.

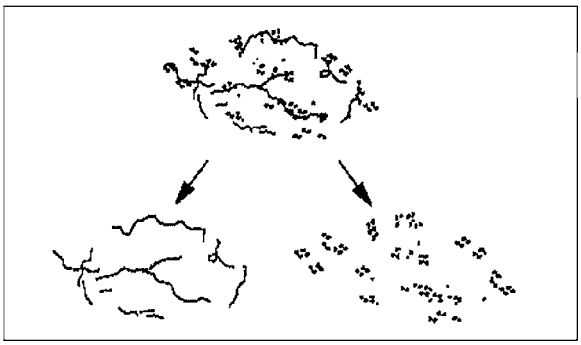

Figure 14. Decomposition of a 1-set. This figure, after Falconer (1990), illustrates the concept of decomposing a rectifiable one-dimensional set into a regular "curve-like" part and an irregular "curve-free" part.

proof that E is irregular, i.e., that its density is different than 1 $\mathcal{H}$-a.e. on E, can be found in Tricot [46].

1-sets and regularity lead to the cornerstone definition for our work. They constitute the basic types of objects to be studied within the rest of this paper. Originally, they were called *Y-sets* by Besicovitch [1] (and this nomenclature still persists in some books [11, 47]), but we decided to adopt the nomenclature used by Falconer [12], since it fits more closely to intuition:

Definition 11 (Curve-like set; [12]). A 1-set contained in a countable union of rectifiable curves will be called a *curve-like set*.

This definition is more general than the one of Jordan. It allows for multiple curves and these can intersect. It does not require a parametrization, since it is rather based on the notion of a set. Moreover

Theorem 3 (*Regularity* [12])**.** *A curve-like set is a regular* 1*-set.*

This therefore assures us that curve-like sets are free of the potential curve-free structures that the more general class of 1-sets could contain. Curve-like sets will suit our needs for edge detection, where we know that some kind of curve structure is present, since the set to be inferred will be provided by the output of oriented line/edge operators.

Hierarchy for one-dimensional sets. In the continuous domain, measures and densities have allowed mathematicians to partition the space of one-dimensional sets (not only the 1-sets), and to build a hierarchy for them. The first distinction, a rather crude one, is between those that have finite measure, the 1-sets, and those that have infinite length. Among the 1-sets, a finer subdivision provides regular (curve-like) and irregular (curve-free) sets. One nice result, called the *decomposition theorem* [12], enables a split of 1-sets into a regular and an irregular part, as shown in Fig. 14. It can be shown that each part from the set can be analyzed separately and then recombined without affecting density properties. The spirit of this decomposition is similar to what we will do with the discrete tangent map obtained though edge/curve detection. While our decomposition scheme will be different than the one presented here, the underlying idea is very similar. For vision applications, of course, further types of structures will be important. This will be discussed in detail in Section 3 of the companion paper [9]. Until then, the reader should keep in mind what was shown in Fig. 14.

3.3.3. Local Structure of Curve-like Sets. Before discussing the existence of tangents for curve-like sets, we will present an alternate definition of a tangent that does not rely on a parametrization of the set. This definition is due to Besicovitch [1]: (c.f. Fig. 9):

Definition 12 (Tangent [11]). A curve-like set E has a *tangent* $T_B(x)$ at x in the direction $\pm\theta$ if

1. $D_u(E, x) > 0$ and
2. for every angle $\phi > 0$,

$$\lim_{r \to 0} \frac{\mathcal{H}(E \cap (B_r(x) \backslash S_r(x, \theta, \phi) \backslash S_r(x, -\theta, \phi)))}{r} = 0 \quad (1)$$

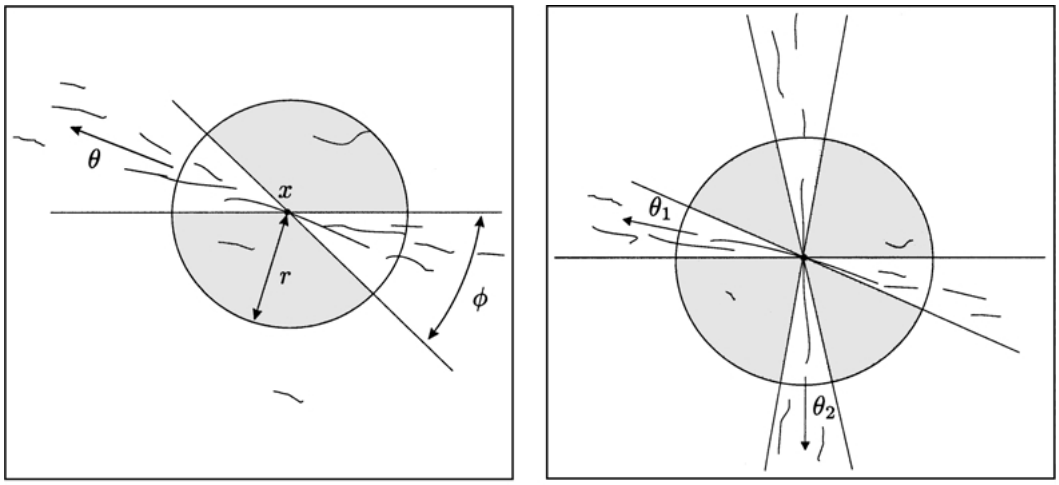

(a) Besicovitch tangent $T(x)$ (b) Tangent set $\Theta(x)$

Figure 15. Illustrating the parametrization-free definition of tangents (Besicovitch tangent) at a point. Such definitions require that a significant part of E lies near x, of which a negligible amount lies outside the wedges. In (a) we illustrate the Besicovitch tangent at one scale, adapted from Falconer [12], and in (b) the tangent set (multiple tangents) at one point.

where $B_r(x)$ is the ball of radius r centered at x, $S_r(x, \theta, \phi)$ is the sector of radius r at angle θ with opening ϕ, and $S_r(x, -\theta, \phi)$ is the sector in the opposite direction (it could have been written $S_r(x, \theta + \pi, \phi)$).

Suppose $x \in E$, then this definition means that at x the set E is locally concentrated around the line $T_B(x)$ with orientation θ passing through x. The first condition in this definition ensures that x is indeed on the set. The second condition ensures that the concentration is around the tangent line only. Figure 15(a) illustrates the definition, namely that the second condition consists in looking at the rate of growth of what is found outside the local angular sector centered at x. If this rate of growth is much faster than r, the curve is ensured (from the first condition) to be concentrated around the line with orientation θ at x.

How does this definition of tangent relate to the usual definition of a tangent to a parametrized curve? In his seminal work, Besicovitch [1] showed that this parametrization-free definition was indeed equivalent to the classical definition we presented in Section 3.2.

Theorem 4 (*Besicovitch and classical tangent* [1])**.** *Let E be a parametrized simple curve. If $x \in E$, and if both $T(x)$ and $T_B(x)$ exist, then the Besicovitch and the usual definition of the tangent at x correspond, i.e., $T(x) = T_B(x)$.*

Proof: The original proof can be found in Besicovitch [1], but a modern presentation can be found in Tricot [47]. □

Remark 5. Since we know now that the Besicovitch tangent and the usual tangent to a parametrized curve coincide, we will denote the tangent to a set E at x by $T(x)$ and always imply the Besicovitch construction.

One can now wonder if the Besicovitch tangent is solving some of the problems encountered with the classical definition for representing data obtained from edge detection. We will focus here on line endings and intersections. For this, let us recall Example 4 in which one of the lines was at an angle $\theta_1 = \pi/4$, while the other was at $\theta_2 = 3\pi/4$. Suppose we are at one of the end points of the line with orientation θ_1. The density $D(E, x) = 1/2 > 0$, and the rate of growth outside a sector is zero since

$$E \cap (B_r(x) \backslash S_r(x, \theta_1, \phi) \backslash S_r(x, -\theta_1, \phi)) = \emptyset$$

for all $0 < \phi < \pi/8$ and $r > 0$. A tangent at the end points is thus defined with the same orientation as for the rest of the line. At the intersection, i.e., for $x = (0, 0)$, the density $D(E, x) = 2 > 0$, but

$$\frac{\mathcal{H}(E \cap (B_r(x) \backslash S_r(x, \theta_1, \phi) \backslash S_r(x, -\theta_1, \phi)))}{r} = 1,$$

for all $r < 1$ and $0 < \phi < \pi/8$, therefore the Besicovitch tangent at $x = (0, 0)$ fails to exist. Since these events (curve intersections) are of paramount importance for the description of an edge map, we will present in the next section a wider representation for the local structure of a set than the Besicovitch tangent.

The study of the distribution of tangents for curve-like sets will be based on this critical result about the existence on tangents:

Theorem 5 (*Tangent a.e* [11])**.** *A curve-like set E has a tangent at almost all its points.*

Sketch of proof: The proof has several steps and is the subject of ([11], Chapter 3). First, one proves that a rectifiable curve Γ has a tangent at almost all its points. This can be done using the following

Lemma 1 ([11])**.** *If $\phi > 0$ and E is the set of points on a rectifiable curve Γ that belong to pairs of arbitrarily small subarcs of Γ subtending chords that make an angle of more than 2ϕ with each other, then $\mathcal{H}(E) = 0$.*

which characterizes the distribution of tangents for a rectifiable curve. It says that on a rectifiable curve, the chords defined by triples of points that are sufficiently close should almost never make a large angle between them. The existence of a tangent almost everywhere for a single rectifiable curve then follows from the continuity of the mapping. Once one knows that a rectifiable curve has a tangent almost everywhere, properties of densities, together with Theorem 3 provide us with the final result. □

Although for this paper the central theorem will be Theorem 5, we end our review of the classical results from geometric measure theory by the *structure theorem*, which constitutes a very deep result about the structure of arbitrary subsets of $\mathbf{R}^n$ [35]. We will cite here its one-dimensional version and a corollary that partly follows from the previous theorem. Recalling that a *continuum* is a compact connected set, we have

Theorem 6 (*Structure theorem* [11])**.** *If E is a continuum with $\mathcal{H}(E) < \infty$, then it consists of a countable union of rectifiable curves together with a set of $\mathcal{H}$-measure zero.*

Corollary 1. *If E is a continuum with $\mathcal{H}(E) < \infty$, then E is regular and has a tangent at almost all its points.*

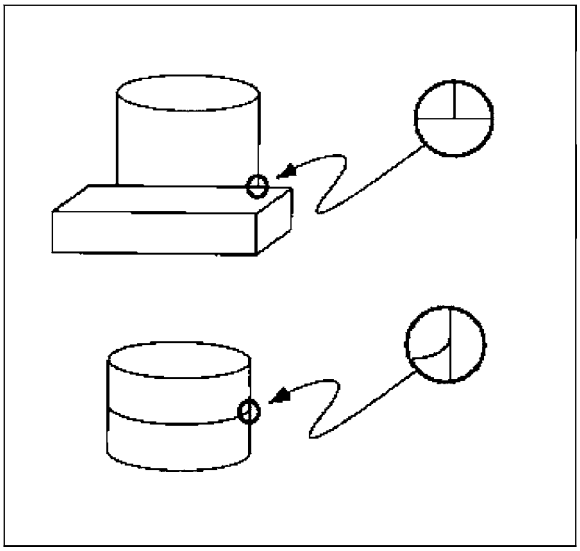

Figure 16. T-junctions. When objects occlude one another, they induce discontinuities in the edge map. T-junctions such as those shown in this figure are one type of these occlusion discontinuities. Reproduced from [Nitzberg, Mumford and Shiota [38].

The material we presented in this section so far is standard in elementary differential geometry and in geometric measure theory. We have briefly established links with relevant issues in computer vision, and have presented the mathematical apparatus needed for the development of our intermediate representation. The rest of this section presents our original contributions.

3.3.4. Multiple Tangents. The Besicovitch tangent must be extended for applications in computer vision. This is due to the fact that when objects occlude, they create discontinuities in bounding contours, leading to T-junctions [17, 18, 38, 49] such as those presented in Fig. 16. At these points of discontinuity in orientation, it is natural to represent "multiple tangents" [55]. Intuitively the rationale is as follows: in the limit, as the discontinuity is approached from one side, one tangent is obtained, while from the other side, the second tangent is obtained. There is also an important motivation from neurobiology, with multiple orientations represented in the same orientation hypercolumn.

The following is an extension of the Besicovitch tangent to allow the representation of multiple tangents at a point:

Definition 13 (Multiple tangents). A curve-like set E has a *tangent set* $\Theta(x)$ at x if $D_u(E, x) > 0$ and for

every angle $\phi > 0$,

$$\lim_{r\to 0} \frac{\mathcal{H}(E\cap(B_r(x)\setminus(\bigcup_{\theta\in\Theta(x)}[S_r(x,\theta,\phi)\cup S_r(x,-\theta,\phi)])))}{r} = 0 \tag{2}$$

but also, for each $\theta \in \Theta(x)$, $\exists\ r_0$ and ϕ_0 such that $\forall 0 < \phi < \phi_0$ and $0 < r < r_0$,

$$\limsup_{r\to 0} \frac{\mathcal{H}(E\cap(B_r(x)\cap(S_r(x,\theta,\phi)\cup S_r(x,-\theta,\phi))))}{2r} > 0 \tag{3}$$

As in the definition of the Besicovitch tangent, the density condition makes (almost) sure we are on the set. The condition given by Eq. (2) prevents things from being too crumpled around the point, while the third (Eq. (3)) ensures that indeed there is something going on in the directions contained in the tangent set. Equation (3) can be interpreted as the requirement that the conical density around each tangent direction be positive. In the case of the usual Besicovitch tangent (Definition 12), we know that this is true (see for instance [34]), therefore both definitions agree.

Remark 6. The value chosen for ϕ_0 will be linked to the orientation resolution of the operators when doing curve detection, while r_0 will be linked to their tangential extent (partly defining the scale of the operator).

Existence of multiple tangents. It is easy to build a set with multiple tangents. The set E from Example 4, for instance, has multiple tangents at $x = (0, 0)$. In this case $\Theta(x) = \{\pi/4, 3\pi/4\}$. For both $\theta \in \Theta(x)$, we have for all $r > 0$ and $0 < \phi < \pi/8$

$$\frac{\mathcal{H}(E\cap(B_r(x)\cap(S_r(x,\theta,\phi)\cup S_r(x,-\theta,\phi))))}{2r} = 1 > 0,$$

while

$$E\cap\left(B_r(x)\setminus\left(\bigcup_{\theta\in\Theta(x)}[S_r(x,\theta,\phi)\cup S_r(x,-\theta,\phi)]\right)\right) = \emptyset$$

which ends the verification.

Until now we did not put any constraint on the cardinality of the tangent set $\Theta(x)$ at a point. This first result addresses part of the issue:

Corollary 2 (*Unique tangent a.e.*)**.** *If E is a curve-like set and $x \in E$, then the set of tangents $\Theta(x)$ at x is composed of a unique tangent for almost all x in E.*

Proof: The proof follows from Theorem 5.

3.3.5. The Tangent Map. The definition of multiple tangents allows us to define a new structure essential for the development of our complexity measure. This structure will give the orientation of the set at each of its points:

Definition 14 (Tangent map). Given a curve-like set E, the *tangent map* τ is

$$\tau = \bigcup_{x\in E}(x, \Theta(x)).$$

The tangent map will provide the mechanism for relating geometric structure to visual structure. The links will be provided by showing that the geometric structure is directly analogous to that obtained in computer vision. Thus we must define the tangent map in the discrete domain. This is how the structure we just defined will be linked to the output of edge detectors. The resulting intermediate representation will be the one used to characterize the complexity of the tangent (edge) map, and to provide a decision scheme for the representation underlying the grouping process. Before describing the discrete counterpart of the tangent map, we will investigate the structure of the underlying tangent map for continuous curve-like sets.

3.4. Tangent Separation Theorems

Rectifiability constrains the global and local distribution of tangents. Basically, for a one-dimensional set to be rectifiable, it cannot be too crumpled and cannot cut itself too often (Fig. 17). The following theorems will try to capture this last statement, and will provide constraints on the underlying local approximation to a curve-like set. Recall that a set is *totally disconnected* if no two of its points lie in the same connected component. Thus, given any pair of points in the set, there is a decomposition into two disjoint closed subsets, each containing one of the points [11]. Now, looking in the neighborhood of a tangent, in the normal direction for instance, we get the following:

Theorem 7 (*Separation of parallel tangents*)**.** *Let E be a curve-like set. Consider a point w on E with tangent $T(w)$. If $N(w)$ is a line passing through w with*

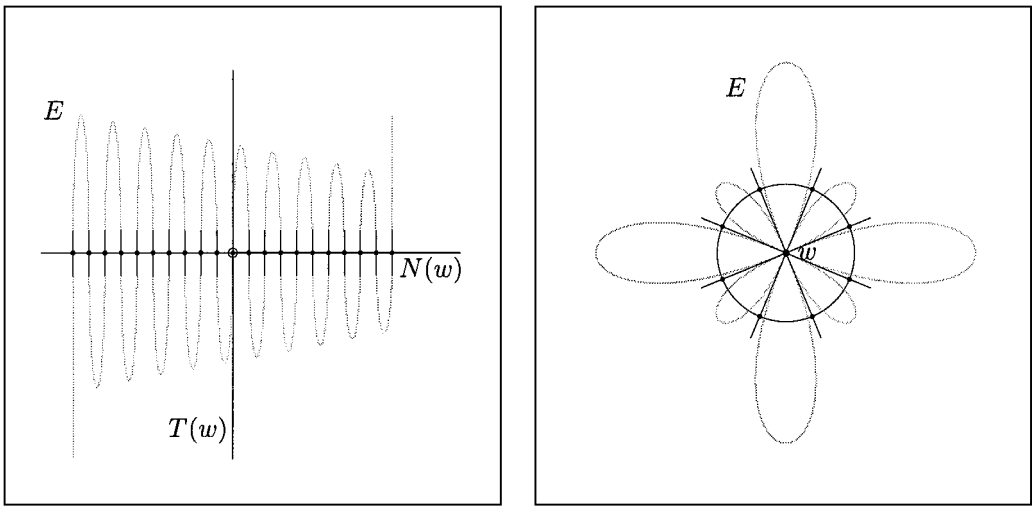

(a) parallel tangents separation (b) multiple tangents separation

Figure 17. Illustration of the tangents separation theorems. The key idea is that for a curve to be rectifiable, the parallel or the multiple tangents cannot form a continuum. (a) illustrates the parallel tangents separation, where $w \in E$ is the point with the circle around it. The tangent $T(w)$ and the normal $N(w)$ are drawn and for this particular example the set $\mathcal{D}(w, E)$ is composed of all the other dots. (b) multiple tangents separation.

orientation different than $T(w)$, then $\mathcal{D}(w, E) = \{y \in E \bigcap N(w) : T(w) = T(y)\}$ is a totally disconnected set.

This result says that a curve can't be squeezed into itself too much, or a point may be reached where the tangents become dense. At this point it will be difficult to "walk" along the curve-like set, in the same sense as it was difficult to walk along Paolina's hair in Section 1.1. Sets with tangents dense in the normal direction require different representations for local to global transitions.

Proof: We will proceed by contradiction. Suppose there exists a connected component $C \subset \mathcal{D}(w, E)$. Since $C \subset N(w)$ we have

$$L(C) = |C|_1 = \mathcal{H}(C) = \text{diam}(C) > 0$$

where $|\cdot|_1$ is the usual one-dimensional Lebesgue measure and $L(\cdot)$, the Jordan length. Take now any point z inside C. For small enough ρ we get

$$\mathcal{H}(C \cap B_\rho(z)) = 2\rho$$

Therefore, if θ is the angle for $T(w)$, then we have for all ϕ sufficiently small

$$\lim_{r \to 0} \frac{\mathcal{H}(C \cap (B_r(z) \backslash S_r(z, \theta, \phi) \backslash S_r(z, -\theta, \phi)))}{r} \geq 1$$

which shows that there is no tangent at that point. This being true for all interior points of C, and since $\text{diam}(C) > 0$, we have found a set of $\mathcal{H}$-measure > 0 for which there does not exist a tangent. However E being a rectifiable set, this contradicts the fact that it should have a tangent almost everywhere. □

Now consider the distribution of multiple tangents. More importantly, however, Theorem 8 will be the equivalent of Theorem 7 but in orientation space rather than in the spatial domain.

Theorem 8 (*Separation of multiple tangents*)**.** *If w is a point on a curve-like set E, then $\Theta(w)$, the set of multiple tangents at w, is a totally disconnected set.*

Sketch of proof: Suppose there exists a connected component $C \subset \Theta(w)$. We can then find a circular arc in the neighborhood of w for which each point has a multiple tangent. That means we would be able to build a set of measure greater than zero with more than one tangent everywhere, which contradicts Theorem 2 since E is rectifiable.

The above proofs of Theorems 7 and 8 show the spirit of Poincaré's cut dimension idea [40]. In both proofs (using contradiction) we show that in order to disconnect the set, we need a continuum; a straight

line in the case of Theorem 7, an arc of circle in Theorem 8. This leads to a contradiction because then the object couldn't be a rectifiable curve.

4. From Curve-like Sets to Edge Detection

One of the central features of the theory of curve-like sets was the Besicovitch tangent formulated in a parametrization-free manner. This is interesting for computer vision, because it is analogous to parametrization-free methods for estimating tangents, namely edge detection. We shall place tremendous emphasis on this analogy. In particular, just as Besicovitch sought a dense collection of points within a cone, "edge" operators seek a dense collection of pixels at a certain contrast [5, 20, 32]. There are two important differences, however, and these differences will motivate the rest of this paper. First is the notion of *scale*. The Besicovitch tangent was defined in the infinitesimal limit, and is related to a classification of curves as being either finite length (rectifiable) or infinite (non-rectifiable). Those with finite length were called 1-sets. The second difference is *resolution*. Orientation for the Besicovitch tangent is a real variable, as is spatial location; for any computation on a computer, these will be quantized numbers.

To develop the analogy between Besicovitch tangent sets and curve detection, we must face several subtleties. Finite scale and finite resolution have deep consequences, which we shall now attempt to illustrate. The result will both help us to frame the measure-theoretic problems that are appropriate for vision, and will lead to a statement of what we seek formally: discrete curve-like sets. To avoid the impression that all the mathematical questions are resolved, we also switch to a more informal style of presentation.

4.1. From Edge Detection to Discrete Curve-like Sets

Finite scale and finite resolution, as they arise in edge detection, make the definition of *discrete* curve-like sets a very delicate one. Since edge operators are all band limited, the inferred objects will all be finite, and the subtlety lies in their organization. Returning to Fig. 2(c), we see that the tradeoff between finite scale and finite resolution results in a distribution of tangents that is more elongated than thick, but it is still a distribution. This certainly captures the spirit of the Besicovitch construction, provided the limiting process is stopped at a finite cone. However inside the hair texture the tangent distribution is more turbulent (Fig. 2(d)). Even though the length of the detected hair is finite, many orientations are seen in a dense neighborhood. This is a situation that required a modification of the pure Besicovitch construction which we called the *tangent set*. The *multiple tangents separation* theorems (Theorem 7 and Theorem 8) lead to the observation that, while more than one tangent can exist at a point, not all tangents can exist at all positions, even for a complicated object like a bowl of spaghetti. Rather, in the "spaghetti" case, the structural criterion will be not how well the tangents continue "along" the curve, but how they fall orthogonally to it. As we show in Section 3 of the companion paper [9], the classification of discrete curve-like sets at finite scale and resolution is much richer for vision than abstract rectifiability.

As non-rectifiable curves at a given detection scale and resolution are uninteresting for vision, we assume that there exists a (rectifiable) curve-like set projecting into a quantized image, in a way such that each pixel through which the curve passes is a discrete trace point. The union of these trace points defines the full trace $\mathcal{T}$. Further, suppose a tangent set is given at each of these trace points by a curve detection scheme. In our examples we used logical/linear operators [24] followed by relaxation labeling [23, 55], but the abstract requirement is that multiple orientations at a point can be represented. We then have

Definition 15 (Discrete tangent map). Let $\mathcal{I}$ denote an image, let $\mathcal{T}$ be a subset of its domain, and let $\hat{x}_k \in \mathcal{T}$. Then the *discrete tangent map* $\hat{\tau}$ will be:

$$\hat{\tau} = \{(\hat{x}_k, \hat{\theta}) \mid \hat{\theta} \in \hat{\Theta}(\hat{x}_k)\}$$

$k = 1, 2, \ldots$, where $\hat{\theta}$ is a quantized orientation at the discrete image coordinate $\hat{x}_k$.

Any entry in the map in fact corresponds to an equivalence class of curves passing through the cell $\hat{x}_k = (x_i, y_i)$ with orientation $\hat{\theta}$. To illustrate the previous definition, we will show the discrete tangent maps for our test patterns. For the Kanizsa pattern, the discrete tangent map was approximated. This could be done because the tangent set was known at every point, therefore we only needed to discretize position and orientation. This approach is justified since the structure of the underlying set is relatively simple: the set was aligned with the grid, curve intersections were sparse,

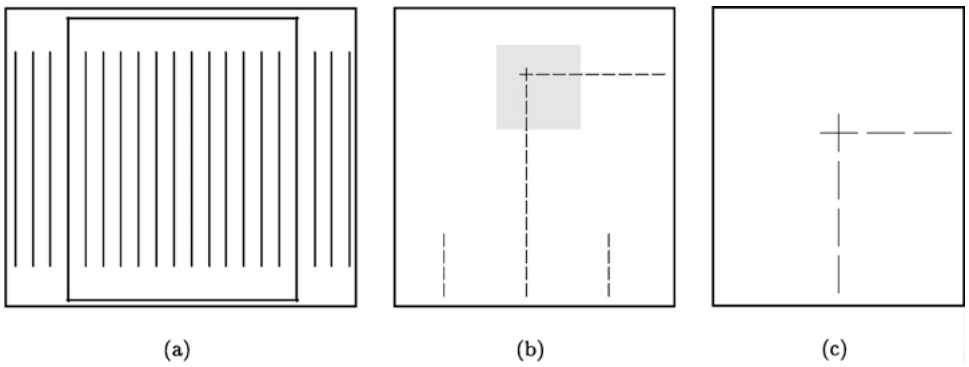

Figure 18. Discrete tangent map for the Kanizsa pattern (Fig. 4(a)). (a) tangent map on the whole pattern, (b) and (c) are close-ups at the upper left corner of the previous pattern (the blow-up in (c) is shown shaded in (b)). Notice the place where there are multiple tangents at the corner as shown in (c).

and the spacing between the lines was known *a priori*, therefore the resolution could be chosen to be high enough to allow the discrimination of the individual curves in the grating.

Example 5 (The Kanizsa pattern DTM). The discrete tangent map (DTM) for the Kanizsa pattern was obtained by first embedding the pattern in the unit square and then discretizing orientation and position. For this example, the chosen resolution was 100×100, with eight discrete orientations. The corresponding discrete tangent map is shown in Fig. 18. Most points on the set have a unique tangent, therefore, for these, the tangent set is composed of a unique element. Interesting places are the four corners of the rectangle, where the tangent set has two elements (tangents) at perpendicular orientations (see Fig. 18(c)), and the ends of lines.

For the other examples, the discrete tangent maps will be inferred through the output of edge/line detectors. For the Ullman patterns, we used *negative contrast line* detectors (seeking dark lines on a light background) based on logical/linear (L/L) operators. The kernels for the convolutions supported $N_\theta = 8$ orientations classes, and the difference between adjacent directions was 22.5°. The operators had a size of about 10×10 pixels (depending on the orientation), had a tangential support of $\sigma_T = 2.0$, and normal (lateral) support of $\sigma_N = \sqrt{2}/2$ pixels. This ensures that all curves are localized to connected regions with width ≤ 2 pixels [24]. Although these operators could have also provided curvature estimates, these were not considered for this study. A complete description of the detection algorithm can be found in either [23] or [24].

Example 6 (The Ullman pattern DTM). The discrete tangent map as computed by edge/line detection for the Ullman pattern is shown in Fig 19. Once more, zooms of regions on the global maps are presented in the bottom blowups. The lower two blowups are of significant interest since they reproduce this effect of broadening of tangents. Notice the orientation ambiguity in the DTM clearly shown in the two figures. These are due to either high curvature, line crossings, or the incapacity of representing a given orientation due to the discretization process.

The discretization process has a dramatic effect on the discrete tangent map. In this paper we used a square lattice, which leads to the question of how to represent a point on the digitization grid: it might fall at the intersection of pixels. The same line of reasoning applies for tangents: they might fall in between two discrete orientations. How should these cases be represented? For the Kanizsa pattern, the lines were intentionally aligned with the lattice, making the generation of the discrete tangent map very straightforward. For the Ullman patterns, when the curves do not align with the grid (which is almost always the case) we saw that they induce both spatial and orientation smearing of the tangents.

For the Paolina image, this time we used the *edge* operators (seeking places where there is a significant change between light and dark regions). The original image was processed with L/L operators with 8 orientations and 5 curvatures. The output was then processed with three iterations of relaxation labeling [23, 56] and only the compatibilities greater than 5% of the maximum were kept. A complete description on how the tangent map was generated for this image can be found in [23].

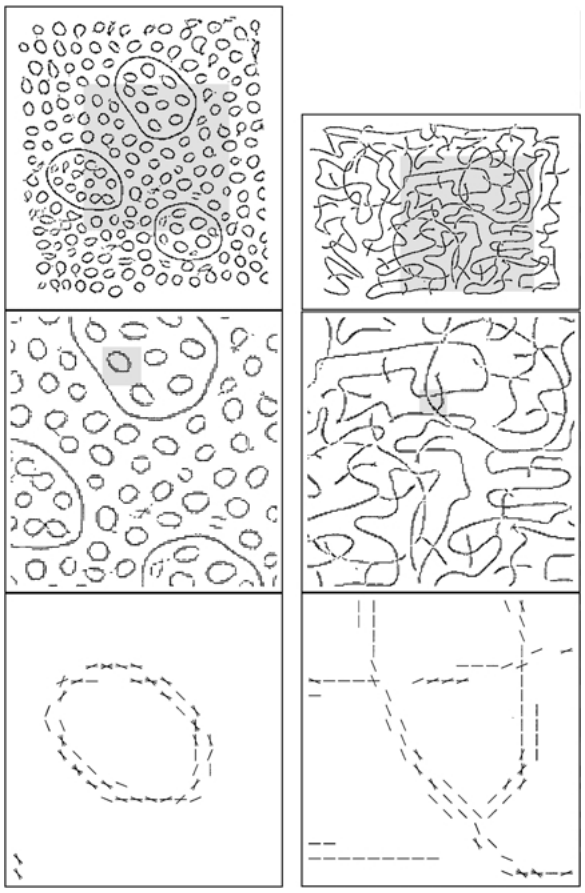

Figure 19. Discrete tangent maps for the Ullman figures (Fig. 6). Top left: discrete tangent map for the pattern in Fig. 6(a); top right: discrete tangent map for Fig. 6(b). The grey regions in upper panels indicate the extent of the close-ups shown underneath. As opposed to the previous two examples, this time the local structure needed to be inferred and was obtained from edge/line detection.

Example 7 (The Paolina image DTM). The tangent map for the Paolina image was given in Fig. 2(f) and was used as a motivation for our introductory statements for this paper. One of the main observations pointed out is that not all orientations can be "on" at the same time. Moreover, the distribution of tangents varies throughout the image: at some places it is very dense, as in the hair region, at others it is very sparse, as on the back.

The discrete tangent map is thus a reasonable interpretation of the output for certain early visual processes. Theoretically, there is a compelling analogy between the first stages of curve detection, which lead to local representations, and a discretized, extended version of the Besicovitch tangent. As the above computational experiments show, the analogy is useful practically as well. But perhaps most importantly, it provides a formal suggestion of how to proceed from a central local-to-global integration in early vision: curve detection is based upon *discrete curve-like sets*; i.e., those sets that are obtained from the integration of the discrete tangent map:

Definition 16 (Discrete curve-like sets). Recalling the property of curve-like sets that a tangent exists $\mathcal{H}$-almost everywhere, the discrete analog—*discrete curve-like sets*—are those sets of discrete points at which a discrete tangent set is defined. From the above definition this is identified as the set of discrete trace points.

For computer vision applications, it is this definition that we take to be primitive. The Besicovitch structure and the multiple tangents extension provide the mathematical rationale for it. Its usefulness is demonstrated by the framing of the following question.

4.1.1. Transversality or Quantization? In the continuous domain, tangent maps were defined in a way that allowed multiple tangents at the same position, to represent those discontinuities that can arise at, e.g. points where one object occludes another. A key concern for discrete tangent maps is how to distinguish this situation from the multiple tangents that can arise due to finite scale and quantization. This is the problem of differentiating transversality from quantization. If at a particular position there are several tangents, how can we decide if these correspond to a quantization artifact or to the fact that two lines are actually crossing at that position (as was shown in the lower panels of Fig. 19)? The standard solution to the problem of spatial broadening is to apply thinning before grouping, this however is bound to destroy the information at orientation discontinuities (as an example, one can look at Fig. 16 in [7], at the intersection between the circle and the vertical line). The standard approach to discontinuity detection is completely different. Typically, a decision process is defined, based on a "corner" model with noise, and a statistical assessment attempted [14]. The problem of corner detection has not previously been considered to be related to quantization artifacts.

Our solution to this issue is two-fold. First, we chose a representation to allow the representation of multiple, spatially coincident edges. The L/L operators not only

respond stably in the neighborhood of multiple coincident curves [24], but are also able to adequately represent them within the discrete tangent map. Second, the distinction between transversality and quantization at places where multiple tangents occur will be detected from the normal complexity of the discrete curve-like set as will be shown in Section 3 of the companion paper [9].

4.2. *Discrete Equivalent to the Tangent Separation Theorems*

Theorem 7 has its discrete equivalent: if there is a discrete tangent at position $\hat{x}$, then there is a limit on the number of tangents that can be locally parallel to it (i.e., with the same orientation). The same applies at a single image location, where all tangents cannot exist at a point, giving a discrete equivalent to Theorem 8. The detection of edges or lines therefore constrains the discrete tangent map: not all tangents can be on, even if the image is of a bowl of spaghetti.

In the continuous domain, this constraint was defined as the fact that, given an orientation θ, one cannot find an *interval* in a *direction different from* θ for which one would have tangents with the same orientation θ. Two key changes need to be considered for the discrete equivalent:

- *different orientation*: Suppose the discrete orientations considered are $\theta_1, \theta_2, \ldots, \theta_{N_\theta}$, and that a discrete tangent θ_i is on at a given position $\hat{x}$. An orientation θ_j will be said to be "different" if $d(\theta_j, \theta_i) \geq 2$ where

$$d(\theta_i, \theta_j) = \frac{|(\theta_i - \theta_j) \bmod_c r|}{\pi / N_\theta}$$

  and where $r = \pi$ for lines, and $r = 2\pi$ for edges. Here the $x \bmod_c y$ operation is a centered modulus, with the output values restricted to the interval $(-y/2, y/2]$. Details about this distance between orientation cells can be found in [23];
- *interval*: The discretization induces a lateral spreading of tangents as we saw earlier e.g., in the Ullman DTM's. An interval here will therefore be defined as a set of M adjacent (8-connected) pixels within a given orientation. It is the value that M will take that will define the minimum size of contiguous pixels to be considered as an interval. For instance, if the resolution of the image is N, then a discrete interval is definitely smaller than N. To refine this assertion, we will need to use the scale of the operator $\sigma = (\sigma_N, \sigma_T)$.

Both these notions then lead to a conjecture about the size of the largest neighborhood over which information can spread laterally:

Conjecture 1 (Separation of discrete parallel tangents). Let $\mathcal{I}$ be an image with spatial resolution N. Let $\hat{\tau}$ be its discrete tangent map obtained through a bank of L/L operators for N_θ orientations with normal scale extent $\sigma_N < N$. If the response at $\hat{x}$ for orientation θ_i is positive, then there exists a constant k such that the set of parallel tangents in direction θ_j cannot be wider than $\lfloor k\sigma_N \rfloor$ contiguous pixels whenever $d(\theta_i, \theta_j) \geq 2$.

A proof for this conjecture must involve several different ideas. First, the continuous version must be factored into discrete equivalence classes, formalizing the notions of discrete positions (traces) and orientations. Moreover, this factorization must remain consistent with geometric measure theory, so discrete concentrations of Hausdorff measure and distribution can be obtained. Finally, properties of the curve detection system must be taken into account, which defines a type of scale $\sigma = (\sigma_N, \sigma_T)$. Letting $M = \lfloor k\sigma_N \rfloor$, one needs to show that if there were more than M lines in density, then the response would not be positive. But if there were only $M - 1$ parallel tangents, then there could exist a pattern for which the response would be significant. To illustrate this for the operators we used, we know that $\lfloor k\sigma_N \rfloor$ is larger than 3 pixels. This is illustrated in Fig. 20 which is showing the discrete tangent map for a line segment at 45°. Notice there the lateral spreading of tangents over 3 pixels wide connected linear neighborhoods. In this paper, we will take the limit M for discrete parallel tangent separation to be 5 pixels (i.e. we can have 4 parallel tangents, but not 5). More generally, one would expect that such a conjecture will hold for any edge/tangent inference scheme with a bound on normal direction focus.

One of the characteristics of linear/logical operators is to allow the coexistence of tangents at a given point. This was needed to allow the representation of T-junctions, and other types of discontinuities. The following is the discrete counterpart of the original continuous result:

Conjecture 2 (Separation of discrete multiple tangents). Let $\mathcal{I}$ be an image with spatial resolution N. Let $\hat{\tau}$ be its discrete tangent map obtained through a bank of

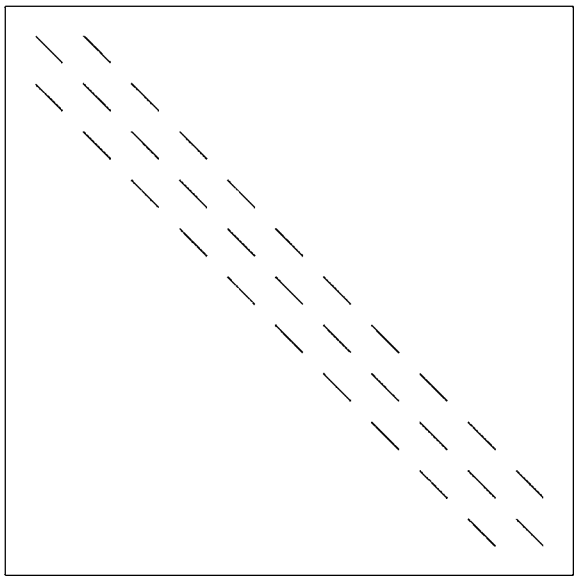

Figure 20. Discrete tangent map for a line segment at an orientation $\theta = 45^\circ$ obtained from L/L operators. Notice the spreading of tangents over a 3 pixels connected linear neighborhood: i.e. drawing a line at 0° orientation can hit 3 contiguous tangents.

L/L operators for N_θ orientations at scale (σ_T, σ_N) with $\sigma_N < \sigma_T < N$, then there exists a constant k such that the set of discrete tangents at a point $\hat{x}$ cannot have more than $\lfloor k\sigma_T/\sigma_N \rfloor$ contiguous elements.

As above, the proof of this conjecture will be involved. Let us set $M = \lfloor k\sigma_T/\sigma_N \rfloor$. We believe much of the proof will be common with the previous one: we would need to show that it is not possible to have M contiguous tangents at one point with positive response, but there could exist a pattern where the operator would provide the coexistence of $M - 1$ contiguous tangents. What needs to be understood out of this conjecture is that at a given pixel, not all tangents (orientations) should have significant response at the same time.

This type of analysis of image operator's behavior is not completely new. Canny's [5] original analysis presented something along these lines when trying to characterize x_{zc}, the mean distance between zero-crossing of f' and $x_{\max}$, the distance between adjacent maxima in the noise response to the filter f. This was used as a constraint to limit the number of peaks in the response.

5. Conclusions

Grouping is a process that inherently involves a local-to-global transition, and in this paper we studied that transition for edge elements into bounding contours. Differential geometry provides the formal framework for this problem, and we reviewed the classical constructs of (local) tangents and (global) curves to identify the need for segmentation of curves from textures. To support the tangent interpretation, we introduced classical structures from geometric measure theory, and showed how the Besicovitch tangent is analagous—at an infinitesimal scale—to edge detection (at a finite scale). Hausdorff's dimension construction then enabled us to prove an extension of the Besicovitch structure theorem that limits the density of tangents in position and orientation space ($\Re^2 \times S^1$); in effect this means that, for curve-like sets, not all tangents can be present at all positions in all orientations. This is the explicit way in which complexity enters the situation: if the tangent distribution is too complex—i.e., too dense—it cannot be arising from a curve-like set. Some other representation is required.

In the companion paper [9] we show how to compute estimates of this complexity for actual edge responses.

Notes

1. Throughout this work, when referring to curve, edge or boundary detection, we imply the process just described here.
2. Actually in this example the length of the sticks is random.
3. Note that the placement of the figures is a deliberate strategy we adopted to make the task harder: having to flip from one page to another, helps to make our point unambiguous.
4. ftp://ftp.cim.mcgill.ca/pub/people/leei/loglin.tar.gz

References

1. A.S. Besicovitch, "On the fundamental geometrical properties of linearly measurable plane sets of points," *Mathematische Annalen*, Vol. 98, pp. 422–464, 1928.
2. T. Binford, "Inferring surfaces from images," *Artificial Intelligence*, Vol. 17, pp. 205–244, 1981.
3. K. Boyer and S. Sarkar, *Perceptual Organization for Artificial Vision Systems*, Kluwer Academic Publishers, 2000.
4. J.C. Burkill and H. Burkill, *A Second Course in Mathematical Analysis*, Cambridge University Press, 1970.
5. J.F. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Machine Intell*, Vol. 8, pp. 679–698, 1986.
6. L. Cesari, *Surface Area*, Princeton University Press: Princeton, NJ, 1965.
7. I.J., Cox, J.M. Rehg, and S. Hingorani, "A Bayesian multiple-hypothesis approach to edge grouping and contour segmentation," *International Journal of Computer Vision*, Vol. 11, No. 1, pp. 5–24, 1993.
8. C. David and S.W. Zucker, "Potentials, valleys, and dynamic global coverings," *International Journal of Computer Vision*, Vol. 5, No. 3, pp. 219–238, 1990.

9. B. Dubuc and S.W. Zucker, "Complexity, confusion, and perceptual grouping. Part II: Mapping complexity," *Journal of Mathematical Imaging and Vision*, Vol. 12, No. 1, pp. 57–84, 2001.
10. G.A. Edgar, *Measure, Topology, and Fractal Geometry*, Undergraduate Texts in Mathematics, Springler-Verlag: New York, 1990.
11. K.J. Falconer, *The Geometry of Fractal Sets*, Vol. 85 of *Cambridge Tracts in Mathematics*, Cambridge University Press: Cambridge, England, 1987.
12. K.J. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, John Wiley & Sons Chichester, England, 1990.
13. H. Federer, *Geometric Measure Theory*, Springler-Verlag: Berlin, 1985.
14. W.T. Freeman and E.H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Machine Intell*, Vol. 13, No. 9, pp. 891–906, 1991.
15. A. Galli and A. Zama, "Untersuchungen über die wahrnehmung ebener geometrischen figuren die ganz oder teilweise von anderen geometrischen figuren verdeckt sind," *Zeitschrift für Psychologie*, Vol. 123, pp. 308–348, 1931.
16. G. Guy and G. Medioni, "Inferring global perceptual contours from local features," *International Journal of Computer Vision*, Vol. 20, No. 1, pp. 113–133, 1993.
17. A. Guzman, "Decomposition of a visual scene into three-dimensional bodies," *AFIPS Conf. Proc.*, Vol. 33, pp. 291–304, 1968.
18. F. Heitger and R. von der Heydt, "A computational model of neural contour processing: figure-ground segregation and illusory contours," in *Proc. of the Fourth ICCV*, 1993, pp. 32–40.
19. D. Hilbert and S. Cohn-Vossen, *Geometry and the Imagination*, Chelsea Publishing Company: New York, 1990.
20. D.H. Hubel and T.N. Wiesel, "Receptive fields and functional architecture of monkey striate cortex," *J. Physiology (London)*, Vol. 195, pp. 215–243, 1962.
21. P. Huggins and S. Zucker, "How folds cut a scene," Center For Computational Vision and Control, Yale University, New Haven, Technical Report, 2000.
22. W. Hurewicz and H. Wallman, *Dimension Theory*, Princeton University Press: Princeton, NJ, 1948.
23. L.A. Iverson, "Toward discrete geometric models for early vision," Ph.D. thesis, Dept. of Electrical Engineering, McGill University, Montréal, 1993.
24. L.A. Iverson and S.W. Zucker, "Logical/linear operators for image curves," *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 17, No. 10, pp. 982–996, 1995.
25. G. Kanizsa, *Organisation in Vision*, Praeger: New York, 1979.
26. M. Kass and A. Witkin, "Analysing oriented patterns," *Computer Vision, Graphics and Image Processing*, Vol. 37, pp. 362–385, 1987.
27. J. Koenderink, *Solid Shape*, MIT Press: Cambridge, MA, 1990.
28. T. Leung and J. Malik, J. "Contour continuity in region based image segmentation," in *Proc. 5th European Conf. on Computer Vision*, 1998, pp. 544–559.
29. J. Malik, S. Belongie, T. Leung, and J. Shi, "Contour and texture analysis for image segmentation," in *Perceptual Organization for Artificial Vision Systems*, K. Boyer and S. Sarkar (Eds.), Kluwer Academic Publishers, 2000.
30. J. Malik and P. Perona, "Preattentive texture discrimination with early vision mechanisms," *J. of the Opt. Soc. of America A*, Vol. 7, pp. 923–932, 1990.
31. B.B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman: San Francisco, 1982.
32. D. Marr and E. Hildreth, "Theory of edge detection," *Proc. Royal Society Lond.*, Vol. 207, pp. 187–217, 1980.
33. M. Mendès-France, "The Planck constant of a curve," in *Fractal Geometry and Analysis*, S. Dubuc and J. Bélair (Eds.), Kluwer Academic Publishers, 1991, pp. 325–366.
34. J.-M. Morel and S. Solimini, *Variational Methods in Image Segmentation*, Vol. 14 of *Progress in Nonlinear Differential Equations and Their Applications*, Birkhaüser: Boston, 1995.
35. F. Morgan, *Geometric Measure Theory*. Academic Press, 1987.
36. D. Mumford, "Elastica and computer vision," in *Algebraic Geometry and Applications*, Springler-Verlag: Heidelberg, 1992.
37. J.R. Munkres, *Topology: A First Course*. Prentice-Hall, 1975.
38. M. Nitzberg, D. Mumford, and T. Shiota, *Filtering, Segmentation and Depth*, Vol. 662 of *Lecture Notes in Computer Science*, Springler-Verlag, Berlin, 1993.
39. H.-O., Peitgen, H. Jürgens, and D. Saupe, *Chaos and Fractals: New Frontiers of Science*, Springer-Verlag: New York, 1992.
40. H. Poincaré, "Pourquoi l'espace a trois dimensions," in *Dernières Pensées*, Flammarion, 1926.
41. C.A. Rogers, *Hausdorff Measures*, Cambridge University Press: London, 1970.
42. E. Saund, "Identifying salient circular arcs on curves," *Computer Vision, Graphics and Image Processing*, Vol. 58, No. 3, pp. 327–337, 1993.
43. G.F. Simmons, *Introduction to Topology and Modern Analysis*, McGraw-Hill: New York, 1963.
44. H.A. Simon, *The Sciences of the Artificial*, MIT Press, 1968.
45. K.T. Smith, *Primer of Modern Analysis*, Bogden and Quigdley, 1971.
46. C. Tricot, "Rectifiable and fractal sets," in *Fractal Geometry and Analysis*, S. Dubuc and J.Bélair (Eds.), Kluwer Academic Publishers, 1991, pp. 367–403.
47. C. Tricot, *Curves and Fractal Dimension*, Springler-Verlag: New York, 1995.
48. S. Ullman, "Three-dimensional object recognition," in *Cold Spring Harbor Symposia on Quantitative Biology*, Vol. LV, Cold Spring Harbor Laboratory Press, Cold Spring Harbor, NY, 1990, pp. 889–898.
49. D. Waltz, "Understanding line drawings of scenes with shadows," in *The Psychology of Computer Vision*, P. Winston (Ed.), McGraw-Hill: New York, 1975, pp. 19–91.
50. L. Williams and A. Hanson, "Perceptual completion of occluded surfaces," *Computer Vision and Image Understanding*, Vol. 64, pp. 1–20, 1996.
51. L. Williams and D. Jacobs, "Stochastic completion fields: A neural model of illusory contour shape and salience," *Neural Computation*, Vol. 9, pp. 837–858, 1997.
52. A. Yuille and J. Coughlin, "Convergence rates of algorithms for visual search: Detecting visual contours," *Advances in Neural Information Processing*, 1998, pp. 641–647.
53. S.W. Zucker, "Early orientation selection: Tangent fields and the dimensionality of their support," *Computer Vision, Graphics and Image Processing*, Vol. 32, pp. 74–103, 1985.
54. S.W. Zucker, "Endowing AI with Vision: A biological and computational perspective," *Current Science*, Vol. 64, No. 6, pp. 407–418, 1993.

55. S.W. Zucker, A. Dobbins, and L.A. Iverson, "Two stages of curve detection suggest two styles of visual computation," *Neural Computation*, Vol. 1, pp. 68–81, 1989.
56. S.W. Zucker, R.A. Hummel, and A. Rosenfeld, "An application of relaxation labelling to line and curve enhancement," *IEEE Transactions on Computers*, Vol. 26, pp. 394–403, 1977.

Benoit Dubuc has a Ph.D. in Electrical Engineering from McGill University specializing in computer vision and was a postdoctoral fellow in the Laboratory of Theoretical and Quantum Computing at University of Montreal. Dr. Dubuc authored and coauthored about a dozen papers on the complexity of curves and surfaces and its application to roughness analysis and the integration of low-level information in computer vision. He is an active member of the IEEE and was awarded the Governor General's Gold medal in 1988.

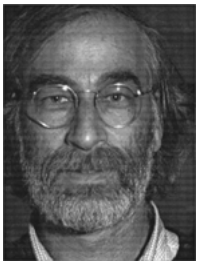

Steven W. Zucker is the David and Lucile Packard Professor of Computer Science and Electrical Engineering at Yale University. Before moving to Yale in 1996, he was Professor of Electrical Engineering at McGill University, Director of the Program in Artificial Intelligence and Robotics of the Canadian Institute for Advanced Research, and the Co-Director of the Computer Vision and Robotics Laboratory in the McGill Research Center for Intelligent Machines. He was elected a Fellow of the Canadian Institute for Advanced Research, a Fellow of the IEEE, and (by)Fellow of Churchill Colege, Cambridge.

Dr. Zucker obtained his education at Carnegie-Mellon University in Pittsburgh and at Drexel University in Philadelphia, and was a postdoctoral Research Fellow in Computer Science at the University of Maryland, College Park. He was Professor Inviteé at Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, France, in 1989, a Visiting Professor of Computer Science at Tel Aviv University in January, 1993, and an SERC Fellow of the Isaac Newton Institute for Mathematical Sciences, University of Cambridge.

Prof. Zucker has authored or co-authored more than 130 papers on computational vision, biological perception, artificial intelligence, and robotics, and serves on the editorial boards of 8 journals.