How Do Neural Systems Use Probabilistic Inference That Is Context-Sensitive to Create and Preserve Organized Complexity?

William A. Phillips

University of Stirling, Scotland, and
Frankfurt Institute of Advanced Studies, Germany
w.a.phillips@stir.ac.uk

Abstract. This paper claims that biological systems will more effectively create organized complexity if they use probabilistic inference that is context-sensitive. It argues that neural systems combine local reliability with flexible, holistic, context-sensitivity, and a theory, Coherent Infomax, showing, in principle, how this can be done is outlined. Ways in which that theory needs further development are noted, and its relation to Friston's theory of free energy reduction is discussed.

Keywords: self-organization, complexity, probabilistic inference, induction, neural systems, Coherent Infomax, context-sensitivity.

1 Unsolved Problems in Theoretical Neurobiology

Biological systems create and preserve organized complexity despite the ever present forces of noise and disorder. This self-organization occurs in open, holistic, far-from-equilibrium, non-linear systems with feedback, to which the classical paradigms of physics are not well-suited. Though usually implicit, probabilistic inference can be seen as being central to self-organization, and useful inference is only possible because the laws of physics are sufficiently reliable. The endless variety of individual circumstances and the prevalence of deterministic chaos and quantal indeterminacy make many things unpredictable, however; so, to thrive, biological systems must combine reliability with flexibility.

Erwin Schrödinger (1944) played an important role in the discovery of the genetic code by correctly estimating the balance between reliability and flexibility (e.g. mutation) in the case of genetics, and by showing, contrary to the opinion of many physicists at that time, that the required balance could be achieved at the molecular level. Analogous insights are now needed to guide our search for general principles of information coding and processing in neural systems. We need to know whether it is possible to state in general abstract terms what is coded by neural activity, how it is coded, and what use is to be made of that information.

Many forms of organized complexity have arisen in nature's long journey from uniformity to maximal disorder, but it is in neural systems that the importance of probabilistic inference is most obvious. Helmholtz correctly emphasized the centrality of unconscious inference to perception, and many examples of its use for contextual disambiguation can be given (e.g. Phillips, von der Malsburg and Singer 2010). Friston (2010) has now shown formally how such unconscious inference may also be central to reinforcement learning, motor control, and many other biological processes.

These arguments suggest several issues on which we need to make progress. What is organized complexity? What are the capabilities and constraints of various forms of inductive inference, e.g. classical versus Bayesian (Jaynes 1998), conscious versus unconscious (Engel and Singer 2008)? How is reliability combined with flexibility, i.e. how is information about reliable generalities combined with information about individual particularities? How is localism combined with holism? Do various neural systems or sub-systems perform inductive inference in different ways with differing degrees of accuracy or generality? Do biological capabilities for probabilistic inference evolve towards forms of inference with greater accuracy or generality? What learning do the inferences require, and how is that learning implemented at the synaptic level in neural systems? Information theory measures such as Shannon entropy and free-energy have been applied to these issues, but how can they be tested and what do they contribute to our understanding?

Better formalisation of these issues is clearly needed, so I will outline an elementary neurocomputational perspective that uses information theory measures to shed some light on them (Phillips, Kay and Smyth 1995; Kay, Phillips and Floreano 1998; Kay and Phillips 2010). A major advantage of this perspective is that it has wide-ranging interdisciplinary roots, and is related, often in considerable detail, to much empirical data from neuroanatomy, cellular and synaptic physiology, cognitive psychology, and psychopathology. I will also argue, however, that this perspective is still greatly in need of further development and testing.

2 Evidence for Local Reliability and Holistic Flexibility

Within neurobiology the contrast between the requirements of reliability and flexibility is reflected by two frequently opposed perspectives that have arisen from the neuroscience of the last century. First, there is the classical perspective, such as that of Hubel and Wiesel. This sees sensory features and semantic attributes as being signalled by single cells, or small local populations of cells, with well-specified receptive fields about which they transmit information. Though modifiable by experience, these codes are highly reliable. They do not change from moment to moment, and do not depend upon what is going on elsewhere. From this perspective feature detection, object recognition, and other higher cognitive processes, are thought to be achieved through a fixed or slowly adapting feedforward projection through a hierarchy of cortical areas.

In contrast, the second perspective emphasizes flexibility. From the early 1980s onwards, there has been a rapidly growing body of evidence showing that, even in sensory systems, neural activity is influenced by an ever-changing stimulus context that reaches far beyond the classical receptive field, and by high-level cognitive state variables such as attention. This has led many to conclude that the simple classical assumption of cells with reliable receptive-fields is no longer viable, and that information is conveyed only by the rich non-linear dynamics of very large and ever-changing populations of cells.

Our perspective combines these two views. It emphasizes dynamic contextual interactions, but claims that, instead of robbing the local signals of their reliability, they increase both their reliability and their relevance. Its central hypothesis is that there are two classes of synaptic interaction: primary driving inputs that specify the information content of the output signals transmitted by the local processor, and gain-controlling inputs that coordinate those computations so as to achieve current goals in current circumstances. This theory emphasises processes of contextual disambiguation and dynamic grouping that choose between alternative interpretations of the data. They do so by amplifying activity that is relevant to the current task and stimulus context, and by suppressing activity that is irrelevant. These coordinating interactions also group activity into coherent subsets, and combat noise by context-sensitive redundancy.

They are crucial to Gestalt perception, selective attention, working memory, strategic coordination, and perhaps also to reinforcement learning and motor control. The contextual coordinating inputs can be seen as learning to predict the activity driven by the primary inputs, and as using those predictions to emphasize and organize activity relevant to current tasks.

Contextual disambiguation and dynamic grouping require many locally specific coordinating interactions between all the detailed processes that compute the cognitive contents. This implies that, in the case of the mammalian cerebral cortex, coordinating interactions must occur within and between cortical regions, because it is only they that know the detailed cognitive contents. Our working assumption is that there is a special class of ubiquitous synaptic interactions within the cortex that selectively amplify and synchronise relevant activities. They are predominantly located on long-range lateral and descending connections and influence post-synaptic activity via various mechanisms, including NMDA receptors (Phillips and Silverstein 2003) and the control of synchronised disinhibition (Tiesinga, Fellous and Sejnowski 2008). They do not themselves provide primary drive to post-synaptic cells, but modulate the effects of those that do. We call them coordinating, or gain-controlling, interactions to distinguish them from the diffuse effects of the classical neuromodulators. There is now clear evidence for a variety of mechanisms that implement such gain-control (von der Malsburg, Phillips and Singer 2010; Silver 2010). Their variety and ubiquity provide evidence of their importance to neural function. One goal for formal studies within theoretical neurobiology is therefore to clarify the capabilities and limitations of such coordinating or gain-controlling mechanisms in more general and abstract terms.

In the cerebral cortex there is anatomical evidence for hierarchical data selection and organization, with extensive context-sensitivity at each stage. Douglas and Martin (2007) review much evidence suggesting that feedforward driving signals provide only about 5% of the input to the layer 2/3 pyramidal cells that transmit information through the system, with the remaining 95% being composed of various forms of contextual input. The small percentage of inputs that are driving may be telling us something crucial about constraints on the discovery of latent structure in data, i.e. it becomes much less feasible as the number of signals in the input to be summarised rises beyond a few hundred. The remaining 95% (several thousand inputs) then suggests that information from a much larger context can usefully guide that discovery. Adams and Cox (this Volume) argue that learning in neuronal systems requires highly accurate pairing of pre- and post-synaptic spikes, and that this faces unavoidable biophysical difficulties. Thus, this may be one of the constraints limiting the number of driving feedforward connections to no more than a few hundred.

These broad claims of close relations between particular local neuronal interactions and particular aspects of cognitive function are based upon much psychophysical, neurobiological, and clinical evidence, as reviewed in Phillips and Singer (1997), Phillips and Silverstein (2003), and von der Malsburg, Phillips and Singer (2010). Though there will not be room to review this evidence here it is of great importance because it suggests that formal clarification of the role of context-sensitivity in probabilistic inference would be worthwhile.

3 The Theory of Coherent Infomax

Our contribution to this effort has produced the theory of Coherent Infomax (Phillips Smyth and Kay 1995; Kay, Floreano and Phillips 1998; Kay and Phillips 2010). Only a brief outline is given here. For full formal presentations see the original publications. The theory of Coherent Infomax uses three-way mutual information and conditional mutual information to show how it is possible in principle for contextual inputs to have large effects on the transmission of information about the primary driving inputs, while transmitting little or no information about themselves, thus influencing the transmission of cognitive content, but without becoming confounded with it. Guided by neuroanatomy, the gross system architecture assumed is that of at most a few tens of hierarchical layers of processing, with very many specialized but interactive local processors at each stage. Feedforward connections between layers are driving, whereas lateral and feedback connections provide coordinating gain-control. Minimally, the function of local processors is to select and compress that information in their primary input that is relevant to the current task and situation, as indicated by the contextual input that modulates primary information transmission. This is formalized in information theoretic terms as an objective function describing the signal processing work to be done. In short, the goal is to maximise the information transmitted about the primary inputs subject to the constraints of substantial data reduction while emphasizing the three-way mutual information between output and primary and secondary inputs and minimizing the information transmitted specifically about the secondary inputs. To show how that objective could be met in neural systems, a biologically plausible activation function for idealized local neural processors was formulated to include the required gain-control, and a learning rule for modifying the synaptic strengths of the connections between these local processors was derived analytically from the objective function by a statistician, Jim Kay. What most impressed us about the consequent learning rule is that, although it was deduced formally from the objective function, assuming none of the physiological evidence concerning the dependence of synaptic plasticity on current and prior activity, it is broadly in agreement with that evidence. The theory of Coherent Infomax thus shows how it is possible for neural systems to perform probabilistic inference in a way that combines reliability with flexibility, and localism with holism.

4 Unsolved Problems in the Theory of Coherent Infomax

Despite its extensive roots in the relevant empirical sciences, however, there are still many ways in which this perspective requires improved conceptual development and empirical testing. One unresolved set of issues arises from what I think of as the impossibility of perfection. In the limit, coherence and information are opposed in that the total information within a system is reduced by correlations between its elements. Therefore, simultaneous maximisation of both considered separately is not possible. We assume that the objective is to increase the total information on which there is agreement, not the attainment of some final and complete optimum. Furthermore, applying the idealized Coherent Infomax objective to systems with many inputs and outputs is not computationally feasible unless simplifying approximations are used. Possible approximations were formally specified by Kay and Phillips (2010), but no attempt was made to explore either their applicability to realistic tasks or their biological plausibility.

A second unresolved problem concerns relations between Coherent Infomax and concepts of complexity. Proposed measures of organized or structured complexity try to combine order (organization/coherence) with disorder (entropy/information) as does Coherent Infomax, and they often do so using mutual information (Sporns 2007). These measures are designed to ascribe high complexity to systems of many elements that interact in such a way as to achieve effective integration but without imposing such uniformity that their joint entropy is low. The contextual interactions of Coherent Infomax seem well-designed to contribute to this because they coordinate activities while not becoming confounded with the information that those activities variously transmit. Furthermore, Coherent Infomax is highly compatible with the small-world network architectures conducive to high complexity on these measures. Therefore, it may be possible to formulate or modify Coherent Infomax so as to relate it explicitly to these measures of complexity, but nothing of that sort has yet been done.

The final unresolved issue concerns the relationship between Coherent Infomax and predictive coding. The current growth of interest in inference and prediction as possible keys to a fundamental understanding of neuronal systems is exemplified by the many groups that work on the ‘Bayesian brain’ and ‘predictive coding’. I will briefly discuss some relations of Coherent Infomax to the work of one of them, i.e. that proposing a unifying brain theory based on ideas from statistical physics and machine learning (Friston 2011) That theory interprets many aspects of neural structure and function as having evolved to reduce Helmholtz free-energy using a form of predictive coding in which ascending activities predicted by feedback descending from higher levels in the hierarchy are suppressed. In contrast to this, Coherent Infomax proposes that activities predicted by contextual input are amplified. Thus, prima facie, predictive coding theories and Coherent Infomax propose opposing effects of context. There are at least three grounds for thinking that these theories are not fundamentally opposed, however. First, both Friston's theory and ours imply that reduction of the difference between predicted and observed probability distributions is a major objective of neuronal dynamics. Furthermore, the reduction of free energy is central to Friston's theory, and in or about 1993 unpublished work by John Hertz (a statistical physicist) gave a proof that Coherent Infomax implies the reduction of free energy. That proof is now lost, but if valid, it shows a deep unity between the two theories. Second, Coherent Infomax emphasizes lateral connections between streams of processing dealing with distinct data-sets, whereas predictive coding is concerned exclusively with feedback connections from higher levels in the hierarchy. Coherent Infomax is most obviously relevant to the use of co-occurrence constraints between distinct streams of processing to select between alternative interpretations of ambiguous inputs, whereas predictive coding theories are concerned with the coding of information for transmission through a hierarchy. Thus, the two theories may be complementary, rather than opposed. Third, Spratling (2008) argued that predictive coding theories can be made formally equivalent to theories based on evidence for *amplifying* effects of top-down attentional inputs. He did so by reorganising the computations required and assuming that suppressive effects of prediction operate on intra-regional signals, not on inter-regional signals. His work therefore suggests that some form of predictive coding may also be formally equivalent to Coherent Infomax, but it is not yet known whether this is so or not.

Predictions may be amplifying in some cases and suppressive in others, so a more inclusive perspective that combines Coherent Infomax with Friston's theory of free-energy reduction may be possible. As the two theories emphasize many of the same details of neuroanatomy, neurophysiology, and psychopathology (Kay and Phillips 2010; Engel et al 2010), it may not be too difficult to create such a perspective. As Friston's theory emphasizes infomax, or redundancy reduction, as one component of his theory, it is important to note that Coherent Infomax is not simply a sub-type of infomax theory; it is infomax *plus* selective amplification of that information predicting activities elsewhere. It is thus our formalisation of another central component of Friston's theory, inference. Overall, therefore, though there are important differences between the two theories, they are in broad agreement. They can be seen as examples of a perspective on biological computation that might be greatly improved by expertise such as that in the INBIOSA initiative.

References

Douglas, R.J., Martin, K.A.C.: Mapping the matrix: The ways of neocortex. Neuron 56, 226–238 (2007)

Engel, A.K., et al.: Coordination in behavior and cognition. In: Dynamic Coordination in the Brain: from Neurons to Mind. Strüngmann Forum Report, vol. 5, pp. 1–24. MIT Press, Cambridge (2010)

Engel, C., Singer, W.: Better than Conscious? Strüngmann forum report, vol. 1. MIT Press, Cambridge (2008)

Friston, K.J.: The free-energy principle: a unified brain theory? Nat. Rev. Neurosci. 11, 127–138 (2010)

Jaynes, E.T.: Probability Theory: The Logic of Science. Oxford University Press (1998)

Kay, J., Floreano, D., Phillips, W.A.: Contextually guided unsupervised learning using local multivariate binary processors. Neural Networks 11, 117–140 (1998)

Kay, J., Phillips, W.A.: Activation functions, computational goals and learning rules for local processors with contextual guidance. Neural Computation 9, 895–910 (1997)

Kay, J., Phillips, W.A.: Coherent infomax as a computational goal for neural systems. Bull. Math. Biology (2010), doi:10.1007/s11583-010-9564-x

Körding, K.P., Wolpert, D.M.: Bayesian integration in sensorimotor learning. Nature 427, 244–247 (2004)

Phillips, W.A., Kay, J., Smyth, D.: The discovery of structure by multi-stream networks of local processors with contextual guidance. Network: Computation in Neural Systems 6, 225–246 (1995)

Phillips, W.A., Silverstein, S.M.: Convergence of biological and psychological perspectives on cognitive coordination in schizophrenia. Behav. Brain Sci. 26, 65–138 (2003)

Phillips, W.A., Singer, W.: In search of common foundations for cortical computation. Behav. Brain Sci. 20, 657–722 (1997)

Phillips, W.A., von der Malsburg, C., Singer, W.: Dynamic coordination in the brain: from neurons to mind. Strüngmann forum report, vol. 5, pp. 1–24. MIT Press, Cambridge (2010)

Schrödinger, E.: What is Life? Cambridge University Press (1944)

Silver, R.A.: Neuronal arithmetic. Nat. Rev. Neurosci. 11, 474–489 (2010)

Sporns, O.: Complexity. Scholarpedia 2(10), 1623 (2007)

Spratling, M.W.: Predictive-coding as a model of biased competition in visual attention. Vis. Res. 48, 1391–1408 (2008)

Tiesinga, P.J., Fellous, J.M., Sejnowski, T.J.: Regulation of spike timing in visual cortical circuits. Nat. Rev. Neurosci. 9, 97–109 (2008)

von der Malsburg, C., Phillips, W.A., Singer, W. (eds.): Dynamic coordination in the brain: from neurons to mind. Strüngmann forum report, vol. 5. MIT Press, Cambridge (2010)