

Cognitive architecture of perceptual organization: from neurons to gnosons

Peter A. van der Helm

Received: 13 July 2011 / Accepted: 26 October 2011 / Published online: 16 November 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract What, if anything, is cognitive architecture and how is it implemented in neural architecture? Focusing on perceptual organization, this question is addressed by way of a pluralist approach which, supported by metatheoretical considerations, combines complementary insights from representational, connectionist, and dynamic systems approaches to cognition. This pluralist approach starts from a representationally inspired model which implements the intertwined but functionally distinguishable subprocesses of feedforward feature encoding, horizontal feature binding, and recurrent feature selection. As sustained by a review of neuroscientific evidence, these are the subprocesses that are believed to take place in the visual hierarchy in the brain. Furthermore, the model employs a special form of processing, called transparallel processing, whose neural signature is proposed to be gamma-band synchronization in transient horizontal neural assemblies. In neuroscience, such assemblies are believed to mediate binding of similar features. Their formal counterparts in the model are special input-dependent distributed representations, called hyperstrings, which allow many similar features to be processed in a transparallel fashion, that is, simultaneously as if only one feature were concerned. This form of processing does justice to both the high combinatorial capacity and the high speed of the perceptual organization process. A naturally following proposal is that those temporarily synchronized neural assemblies are “gnosons”, that is, constituents of flexible self-organizing cognitive

architecture in between the relatively rigid level of neurons and the still elusive level of consciousness.

Keywords Cognitive architecture · Cognitive processing · Distributed representations · Feature binding · Neuronal synchronization · Perceptual organization

Introduction

The term cognitive architecture refers to computational models of not only resulting behavior but also structural properties of intelligent systems. These structural properties can be physical properties as well as more abstract properties implemented in physical systems such as computers and brains. There is no consensus about what these structural properties should be, and indeed, many different cognitive-architecture models have been proposed (for extensive reviews and references, see, e.g., Langley et al. 2009; Sun 2004). These models differ, for instance, in whether they involve fixed or flexible architectures, in what forms of processing they allow (e.g., serial or parallel processing), and the extent to which they are based on a set of symbolic information-processing rules applied by one central processor or rely on emergent properties of many interacting processing units. Most models agree, however, that a cognitive architecture is a parameter-free blueprint for a system that acts like the human cognitive system as a whole.

Cognitive-architecture models differ from cognitive models and expert systems which focus on particular competences such as language, concept learning, or problem solving. Even so, many cognitive-architecture models seek compliance with higher (conscious) cognitive faculties rather than with lower (nonconscious) faculties like

P. A. van der Helm (✉)
Donders Institute for Brain, Cognition, and Behaviour,
Radboud University Nijmegen, Montessorilaan 3,
6525 HR Nijmegen, The Netherlands
e-mail: p.vanderhelm@donders.ru.nl
URL: <http://www.socsci.ru.nl/~peterh>

visual perception. In this article, I do not pretend to present a full-blown cognitive architecture, but I aim to contribute to understanding the architecture of the human cognitive system by discussing a neurally plausible algorithmic model of perceptual organization.

To give a first gist, this model implements the intertwined but functionally distinguishable subprocesses of feedforward feature encoding, horizontal feature binding, and recurrent feature selection. As I sustain by a review of neuroscientific evidence, these are the subprocesses that are believed to take place in the visual hierarchy in the brain. The model further employs a special form of processing, called transparallel processing, whose neural signature is proposed to be gamma-band synchronization in transient neural assemblies. This is argued to lead to a picture of how flexible self-organizing cognitive architecture might be implemented in the neural architecture of the brain. Next, by way of further introduction, I briefly sketch the problem of perceptual organization, the presumed role of neuronal synchronization in perceptual organization, and the pluralist approach I adopt to arrive at this picture of cognitive architecture.

Perceptual organization

Perceptual organization refers to the neuro-cognitive process that takes the light in our eyes as input and that enables us to perceive scenes as structured wholes consisting of objects arranged in space (see Fig. 1). This automatic process may seem to occur effortlessly, but by all accounts, it must be very complex and yet very flexible. To give a gist (following Gray 1999), multiple sets of features at multiple, sometimes overlapping, locations in a stimulus must be grouped simultaneously. This implies that the process must cope with a large number of possible

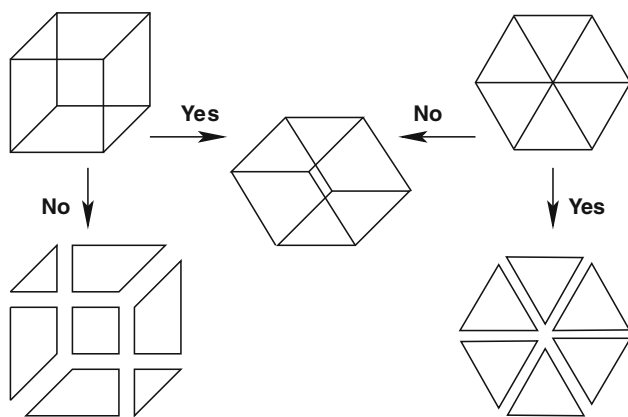


Fig. 1 Perceptual organization. Both images at the top can be interpreted as 3-D cubes and as 2-D mosaics, but as indicated by “yes” and “no”, humans preferably interpret the one at the left as a 3-D cube and the one at the right as a 2-D mosaic of triangles

combinations in parallel, which also suggests that these possible combinations are engaged in a stimulus-dependent competition between grouping criteria. This indicates that the combinatorial capacity of the perceptual organization process must be very high. This, together with its high speed (it completes in the range of 100–300 ms), reveals the truly impressive nature of the perceptual organization process.

My algorithmic model was developed to account for both the high combinatorial capacity and the high speed of the perceptual organization process. To this end, it implements the earlier-mentioned subprocesses of feedforward feature encoding, horizontal feature binding, and recurrent feature selection. Most distinguishing, it employs this special form of processing, called transparallel processing, whose neural signature is proposed to be neuronal synchronization. This issue is introduced next.

Neuronal synchronization

Neuronal synchronization is the phenomenon that neurons, in transient assemblies, temporarily synchronize their activity. Not to be confused with neuroplasticity which involves changes in connectivity, such assemblies are thought to arise when neurons shift their allegiance to different groups by altering connection strengths (Edelman 1987), which may also imply a shift in the specificity and function of neurons (Gilbert 1992). Both theoretically (Milner 1974; von der Malsburg 1981) and empirically (Eckhorn et al. 1988; Gray and Singer 1989), neuronal synchronization has been associated with cognitive processing, and 30–70 Hz gamma-band synchronization in particular has been associated with feature binding in perceptual organization.

As I discuss in section “[The visual hierarchy](#)”, physical properties of neuronal synchronization have been studied, but thus far, it lacked a computational account explaining what is being processed, and how. My algorithmic model now suggests that those transient neural assemblies can be conceived of as cognitive information processors—which I call “gnosons” (i.e., fundamental particles of cognition) and which I propose to be the constituents of flexible self-organizing cognitive architecture. The idea that cognition is a dynamic process of self-organization is not new (see, e.g., Attneave 1982; Kelso 1995; Koffka 1935; Köhler 1920; Lehar, 2003; Wertheimer 1912, 1923), and the idea that those assemblies are the building blocks of cognition is not new either (see, e.g., Buzsáki 2006; Finkel et al. 1998). What my model adds, however, is the idea that those assemblies are involved in transparallel feature processing. As I discuss in section “[A representationally inspired algorithmic account](#)”, this special form of processing is enabled by special input-dependent distributed representations, called hyperstrings,

which allow one processor (also, e.g., a single computer) to recode many similar features in one go, that is, simultaneously as if only one feature were concerned. This is key in my account of the high combinatorial capacity and speed of perceptual organization.

Transparallel processing is basically an idea about feature binding. The classical binding problem is often taken to refer to binding of different features. This is a form of binding which I rather would call integration (think of Treisman and Gelade's 1980, feature integration theory) and which, in my model, is the result of feature selection. Preceding this selection, however, there is also binding of similar features, and this what neuronal synchronization seems to mediate (see section "[The visual hierarchy](#)"). Binding of similar features may seem a limited basis to focus on, but in my model, it enables a high combinatorial capacity and speed which remain effective until selection and integration (see section "[A representationally inspired algorithmic account](#)"). Furthermore, my notion of features is broader than first-order features, like orientation, as considered usually in neuroscience. I focus on second-order features, such as symmetry and repetition, in terms of correlations between elements in a stimulus. I do not think this conflicts with existing neuroscientific evidence (cf. Tyler et al. 2005), and pre-attentive detection of such second-order features is believed to be an integral part of the automatic perceptual organization process (Simon 1972; Tyler 1996; van der Helm and Leeuwenberg 1996; Wagemans 1997).

Pluralist approaches

David Marr (1945–1980) probably would have been thrilled by the present state of cognitive (neuro)science. When he died, classical representational theory dominated the research field, in which connectionism and dynamic systems theory (DST) had not yet gained the impact they have nowadays. Even so, in his book *Vision* (Marr 1982/2010), he envisioned a theory comprising three separate but complementary levels of description of the visual system—the computational, algorithmic, and implementational levels—to which, as I argue in section "[Towards a pluralist account](#)", representational, connectionist, and DST approaches run sort of parallel. In line with Marr's complementarity idea, I argue further that insights from all these three modeling approaches must be combined to address the question of how cognitive architecture might be implemented in the neural architecture of the brain.

It is true that, at least according to some, those three modeling approaches exhibit differences in underlying philosophy (e.g., DST proponents tend to reject the existence of representations), and they certainly reflect different modeling stances. Roughly, representational theory

proposes that cognition relies on regularity extraction to get structured mental representations; connectionism proposes that it relies on activation spreading through a network connecting pieces of information; and DST proposes that it relies on dynamic changes in the brain's neural state. Not surprising therefore, during the past decades, many things have been written for and against each of these three approaches (see, e.g., Fodor and Pylyshyn 1988; Smolensky 1988; van Gelder and Port 1995).

However, instead of thinking that these approaches are mutually exclusive, I think they are complementary—precisely because they focus on different aspects. The idea that intelligent systems need a pluralist approach is already quite common in artificial intelligence research (cf. Dale 2008; Dale and Spivey 2005; Edelman 2008a; Jilk et al. 2008) and is gaining in acceptance in cognitive science (cf. Abrahamsen and Bechtel 2006; Bem and Looren de Jong 2006; Kelley 2003; Lehar, 1999, 2003; Pavloski 2011; Smith and Samuelson 2003). In this article, I aim to go farther than just promoting this idea. My algorithmic model was inspired by a representational approach, but I adopt a pluralist approach to investigate how cognitive architecture might be implemented in neural architecture. Pivotal in this investigation is the phenomenon of neuronal synchronization which, thus far, has been studied in DST, less so in connectionism, and to my knowledge not in representational theory. Also pivotal is the returning topic of distributed representations, which is argued to connect those three modeling approaches.

Organization of this article

In this article, insights from representational, connectionist, and DST approaches are combined to sustain the proposal that the cognitive architecture of perceptual organization is constituted by gnosons, that is, by transient neural subnetworks exhibiting synchronization as a manifestation of transparallel processing of similar features. To elaborate these issues, I hardly discuss details of specific models within the three above-mentioned modeling approaches to cognition. Rather, I aim to assess differences and parallels between the modeling tools they provide to understand the role of neuronal synchronization in perceptual organization. To this end, the organization of this article is as follows.

- In section "[The visual hierarchy](#)", I review neuroscientific evidence on the intertwined but functionally distinguishable subprocesses that are believed to constitute the perceptual organization process in the visual hierarchy in the brain—followed by a discussion of the dynamics and earlier-proposed meanings of neuronal synchronization.

- In section “[A representationally inspired algorithmic account](#)”, I discuss my algorithmic model of the perceptual organization process—introduced by an overview of theoretical ideas and developments within the representational approach that underlies this algorithmic model.
- In section “[Towards a pluralist account](#)”, to substantiate my pluralist approach, I discuss metatheoretical issues such as metaphors of cognition, levels of description, and forms of processing—now and again expanding on traditional views in a way that, in my view, is appropriate to relate representational, connectionist, and DST approaches to each other.
- In section “[Cognitive architecture](#)”, I discuss implications regarding cognitive architecture—grounding gnosons as constituents of flexible self-organizing cognitive architecture.

Before I proceed, a few general remarks seem in order. In this article, I present an idea about the meaning and role of neuronal synchronization. Whether neuronal synchronization indeed exhibits the specific behaviors I suggest is a question I gladly leave to future research by expert experimentators. My objective as a theorist is to provide arguments for a hopefully innovative idea that is not in conflict with existing evidence—I think that such ideas are needed to round the empirical cycle.

Furthermore, this is a multidisciplinary article, and probably the biggest challenge for such articles is the usage of different terminologies by different domains. Therefore, now and again, I state things repeatedly but in different terminologies, which may look redundant but which is needed to assess whether statements from different domains really express different things or merely look different because they are stated in different “languages”. In other words, without denying that different domains model things in different ways (I in fact cherish differences, because that is what complementarity is about), I

want to stress that different languages can also express the same things.

Finally, a multidisciplinary article unavoidably contains domain-specific parts which reflect textbook material to some readers—they may skip such parts—but which are yet necessary to serve other readers. Some readers may also feel that some parts of this article still lack some pertinent domain-specific details and related literature references. I hope, however, that readers agree that such features are inherent to attempts to find common ground for different approaches to the same problem.

The visual hierarchy

This section sets the stage for my algorithmic model. First, with a representationalist eye, I review neuroscientific evidence on the intertwined but functionally distinguishable subprocesses that are believed to take place in the visual hierarchy in the brain. Then, I discuss the phenomenon of neuronal synchronization, DST studies on its dynamics, and neuroscientific ideas about its role in perceptual organization.

To begin with standard textbook material, the top end of the visual hierarchy seems to involve a smooth transition into higher cognitive structures, while the bottom end can be said to be in the primary visual area V1 in the occipital lobe, which receives its main input from the lateral geniculate nucleus (LGN) (see Fig. 2a). In the LGN, a distinction can be made between retinal input entering the parvocellular pathway and retinal input entering the magnocellular pathway. Via V1 and higher visual areas, these pathways bifurcate into a ventral and a dorsal stream which seem to be dedicated to object perception and spatial perception, respectively (Ungerleider and Mishkin 1982; see Fig. 2b).

The neural network in the visual hierarchy is organized with 10–14 distinguishable hierarchical levels (with

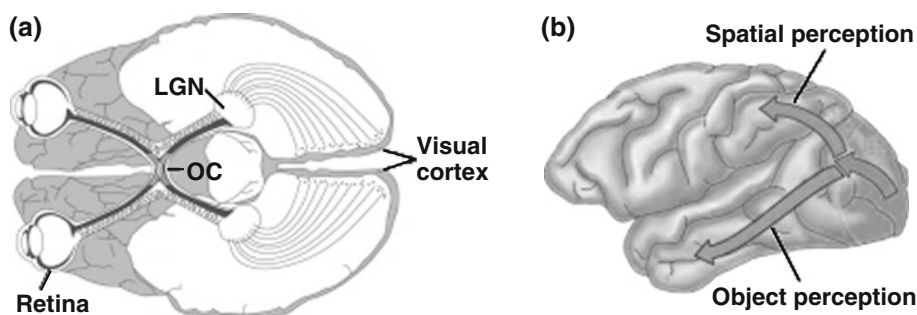


Fig. 2 Visual pathways. **a** Retinal signals go, via the optic chiasm (OC) and the lateral geniculate nucleus (LGN), to the visual cortex; the OC arranges that the left-hand visual fields of both eyes are projected onto the right-hand cortex, and vice versa; in the LGN,

retinal signals enter parvocellular and magnocellular paths, which perform a spatial frequency analysis. **b** In the visual cortex, the signals bifurcate into ventral and dorsal streams which are dedicated to object perception and spatial perception, respectively

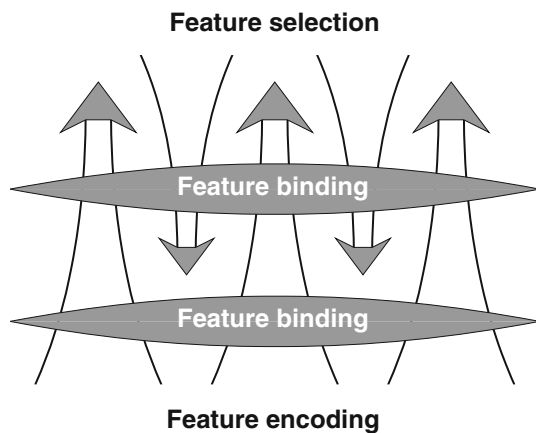


Fig. 3 The three intertwined subprocesses that are believed to take place in the visual hierarchy in the brain. Feedforward connections seem responsible for an initial feature encoding; horizontal connections seem responsible for binding similar features within visual areas; and recurrent connections seem responsible for selecting and integrating different features into percepts

multiple distinguishable areas within each level), contains many short-range and long-range connections (both within and between levels), and it can be said to perform distributed hierarchical processing (Felleman and van Essen 1991). Furthermore, as depicted in Fig. 3, the intertwined but functionally distinguishable subprocesses of feature encoding, feature binding, and feature selection seem to be mediated by feedforward (or ascending), horizontal (or lateral), and recurrent (or feedback, or reentrant, or descending) connections, respectively (see, e.g., Lamme et al. 1998; Lamme and Roelfsema 2000). The horizontal connections, in particular, have been associated with neuronal synchronization, but for a complete picture, I first discuss the others by conveying impressions I get from the available evidence.

Feedforward feature encoding

Feedforward connections seem responsible for a fast bottom-up processing of incoming stimuli. This so-called feedforward sweep takes about 100 ms to reach the top end of the visual hierarchy, and it is thought to yield an initial, autonomous, tuning to features to which the visual system is sensitive (which does not exclude top-down influences; see both this and the next subsections). It is generally thought that, during this feedforward sweep, more complex things are coded in higher visual areas. Traditional ideas about this increase in complexity lean upon the concept of the classical receptive field (cRF). The cRF corresponds to the region of the retina to which a neuron is connected by way of feedforward connections (Hubel and Wiesel 1968). This region is larger in higher visual areas, which suggests that the difference between simple and complex things

corresponds merely to the spatial difference between small (or local) and large (or global) features.

However, by way of horizontal and recurrent connections, neurons also receive input from neurons at the same and higher levels in the visual hierarchy. This suggests that a neuron is responsive to local features outside its cRF and to global features extending beyond its cRF (Gilbert 1992; Lamme et al. 1998; Salin and Bullier 1995). This suggests that the feedforward sweep is part of a more intricate process than just tuning and that, during this process, higher visual areas accommodate features which, perceptually, turn out to be more categorical (cf. Ahissar and Hochstein 2004; Hochstein and Ahissar 2002). I use the term categorical to refer to dominant or salient features which give the gist of a scene—for instance, because they reflect statistical regularities in the environment (cf. Howe and Purves 2004, 2005) or because they reflect geometrical regularities in terms of correlations between elements in a stimulus (cf. Kimchi and Palmer 1982; Leeuwenberg and van der Helm 1991; Leeuwenberg et al. 1994).

A more categorical feature may correspond to a larger feature, but not necessarily so. For instance, in visual search studies, a target usually is a local feature (e.g., one red item among many blue items; Treisman and Gelade 1980). The search for such a target is easier as the distractors are more similar to each other and more different from the target (Donderi 2006; Duncan and Humphreys 1989; Wolfe 2007). Hence, a target may pop-out but only if allowed by the distractors. This means that, for a target to become a pop-out, the distractors have to be processed first—this may well involve lateral inhibition among similar things so that the target rises above the distractors, but in any case, it seems plausible that the similarity of the distractors is processed first in lower visual areas and that the pop-out nature of the target ends up in higher visual areas.

Recurrent feature selection

Recurrent connections seem responsible for a top-down selection and integration of different features into percepts. Somewhat related to the question of whether this subprocess relies on environmental regularities or on stimulus regularities (see above), a question is whether or not this subprocess involves top-down processing starting from beyond the visual hierarchy. For instance, Hochstein and Ahissar (2002) proposed that, via recurrent connections from beyond the visual hierarchy, attention can be deployed in a top-down fashion to any level in the visual hierarchy (see also Wolfe 2007). This would imply that it first captures things coded in higher visual areas and that, if required by task and allowed by time, it may descend along recurrent connections to capture things coded in lower

areas. Given the above picture of the feedforward sweep, this suggests that a pop-out is not a pop-out because it is (nonconsciously) processed first during the bottom-up feedforward sweep, but because its pop-out nature ends up in higher visual areas so that it is among the first things (consciously) encountered by top-down attentional processes.

This picture of the role of recurrent connections in the deployment of attention agrees with Lamme et al. (1998) and Lamme and Roelfsema (2000), who also noted that it may explain the effect of backward masking. A structured stimulus and a subsequent random mask trigger successive feedforward sweeps, and the second sweep (by the mask) then may perturb the trace of the first sweep (by the stimulus) in lower visual areas, so that attention can capture only the more categorical stimulus features coded in higher visual areas. This agrees with the above idea that, in general, less-structured parts (as in a random mask) are coded in lower areas than more-structured organizations into wholes (as in a structured stimulus). It also explains Leeuwenberg et al. (1985) finding that, if a part and a whole are presented briefly and with small stimulus onset asynchrony (SOA), then not only their presentation order but also their structural relationship determines how well the part is identified afterward. It further explains van der Vloed et al. (2007) similar finding which, by way of example, I discuss next in more detail.

Van der Vloed et al. (2007) considered stimuli composed of one symmetrical (S) or random (R) part surrounding another symmetrical or random part (see Fig. 4). The parts were presented for 200 ms each, either simultaneously (SOA = 0) or not (SOA = 20–100 ms), and the task was to identify a given stimulus as being partly symmetrical (for SOA > 0, presented in the orders SR or RS) versus either completely random or completely




Fig. 4 Time course of a trial in van der Vloed et al. (2007). First, one part of the stimulus is presented (here, a symmetrical center). This part remains visible for 200 ms in total, but after an SOA of 0–100 ms, it is complemented with the remaining part (here, a random surround). After 200 ms, the first part disappears and the second part still remains visible for as long as the SOA was so that it is also visible for 200 ms in total

symmetrical (for SOA > 0, referred to by RR and SS, respectively). For SOA = 0, the partly symmetrical stimuli behaved like normal noisy symmetries, with the well-known quantitative effect that, compared to symmetry in the surround, symmetry in the center yields better discrimination from completely random stimuli and worse discrimination from completely symmetrical stimuli (Barlow and Reeves 1979). For SOA > 0, however, there was a qualitative effect of order, no matter whether symmetry was in the surround or in the center: compared to SOA = 0, SR showed no difference (just as RR and SS), but RS yielded better discrimination from RR and worse discrimination from SS.

This order effect again agrees with the idea that, in general, less-structured (e.g., random) information is coded in lower areas than more-structured (e.g., symmetry) information. That is, in SR, the code of the symmetry first settles relatively high and the code of the later-presented random information remains relatively low—just as when the parts were presented simultaneously. In RS, however, the symmetry—on its way to be coded relatively high—passes through the lower areas where the code of the preceding random information already resides; thereby, it perturbs (or masks) the encoded random information, resulting in a percept that reflects less randomness than there really is.

Notice that the foregoing suggests that structural relationships within and between stimuli presented subsequently with small SOA form a factor to be reckoned with (e.g., in experiments involving priming or masking; see also Hermens and Herzog 2007). That is, it asserts that structural factors are at least as relevant as spatio-temporal factors (probably also in, e.g., apparent motion; see Moore et al. 2007).

Also notice, however, that the examples above involve experimental paradigms in which participants respond consciously, that is, they respond on the basis of attentional scrutiny of already-encoded percepts. The question therefore still is whether the formation of these percepts is controlled by endogenous, attention-driven, recurrent processing starting from beyond the visual hierarchy (see, e.g., Lamme et al. 1998; Lamme and Roelfsema 2000) or by exogenous, stimulus-driven, recurrent processing within the visual hierarchy (see, e.g., Gray 1999; Moore et al. 2007; Pylyshyn 1999). The latter reflects my modeling stance in this article, but as I clarify next, it leaves room for the former (see also, e.g., van Leeuwen et al. 2011).

The combination of feedforward and recurrent processing in the visual hierarchy might be analogous to the cascade formed by a fountain under increasing water pressure. That is, as the feedforward sweep progresses along ascending connections, each passed level in the visual hierarchy forms the starting point of integrative recurrent

processing along descending connections. This yields a gradual buildup from partial percepts at lower levels in the hierarchy to complete percepts near its top end. This implies, on the one hand, that top-down attentional processes may intrude before a percept has completed, but on the other hand, that the perceptual organization process has already done much of its integrative work by then. To paraphrase Neisser (1967), before you can pick an apple from a tree, you first have to perceptually organize the scene to at least some degree.

Horizontal feature binding

In between the two just-discussed intertwined subprocesses, horizontal connections seem responsible for binding similar features. This seems to yield feature constellations from which, as mentioned above, recurrent processing seems to select and integrate different features into percepts. For instance, as Lamme et al. (1998) noted, a well-established property of horizontal fibers is that they interconnect cells with similar orientation preferences and that these connections are strongest when cRFs are also co-axially aligned (see, e.g., Bosking et al. 1997; Gilbert 1993, 1996; Malach et al. 1993; Schmidt et al. 1997).

Horizontal binding is a relatively underexposed topic, but to be clear, it seems to concern binding of similar features, with, at least in my model, also a very positive efficiency effect on the subsequent selection and integration of different features. Notice that, in my model, I focus on second-order features such as symmetry and repetition. In section “[Introduction](#)”, I already mentioned that I do not think this conflicts with neuroscientific evidence (cf. Tyler et al. 2005) and that pre-attentive detection of such regularities is believed to be an integral part of the perceptual organization process (Simon 1972; Tyler 1996; van der Helm and Leeuwenberg 1996; Wagemans 1997). In fact, horizontal binding may well be the neuronal counterpart of the regularity extraction operations which, in representational theory, are proposed to lead to structured mental representations.

The subprocess of horizontal feature binding seems to start in V1 and seems to be followed by feature recoding in higher visual areas (Pollen 1999; see also Eckhorn 1999; Gray 1999; Tyler et al. 2005). Furthermore, I can only imagine that it is intertwined with the already intertwined subprocesses of feedforward feature encoding and recurrent feature selection. In any case, such intertwining is key in my model (see section “[A representationally inspired algorithmic account](#)”). Finally, the horizontal feature binding seems to be mediated by transient neural assemblies which also have been implicated in the phenomenon of neuronal synchronization (see, e.g., Eckhorn 1999; Eckhorn et al. 1988; Engel et al. 1990; Gilbert 1992; Gray

et al. 1989, 1990; Gray and Singer 1989). Because my investigation into cognitive architecture revolves around a computational account of this phenomenon, I next discuss it in more detail.

Neuronal synchronization

In representational approaches, a mental representation of a scene (or a percept, or a Gestalt) is said to carry information about the perceptual structure of the scene—that is, about properties (such as shape, parts, and spatial arrangement) of the perceived objects. DST proponents tend to reject the existence of representations, but the term representation can also be said to refer to a relatively stable cognitive state which arises during the dynamic neural process (cf. Kelso 1995). Such a state constitutes the brain’s response to a scene, and it can therefore be said to represent what representationalists call the information about the perceptual structure of the scene (cf. Bem and Looren de Jong 2006).

In any case, for a specific scene, this response (or this information) must also be given (or represented), probably isomorphically, by a specific neural activation pattern (Köhler 1920; Lehar 1999, 2003; Pavloski 2011). That is, it is no surprise that, as shown in brain-imaging studies, different stimuli evoke different neural responses. The question, however, is how to explain these differences. Therefore, cracking the neural code is a central issue in neuroscience. Traditionally, the spike rate of neurons (i.e., the firing rate, or the rate of action potentials) is seen as an important component of the neural code. For instance, the spike rate of neurons may increase as the intensity of a stimulus increases (Adrian and Zotterman 1926). Nowadays, however, as I discuss next, correlations which rely on the precise timing of spikes are seen as being probably more important.

It has been argued that, in general, correlations between spike trains can only reduce, and never increase, the total amount of information in spike trains (Johnson 1980a, b). This, however, may hold if one adopts Shannon’s (1948) classical probabilistic quantification of information, but not if one adopts modern descriptive quantifications of information (see Li and Vitányi 1997; van der Helm 2000). For instance, the equality of two equal messages (e.g., spike trains) is not coded in these messages themselves, so that this equality forms a message in itself. This message may be conveyed by a code which captures the correlation between the two equal messages so that, this way, correlations increase the total amount of conveyable information (Nirenberg and Latham 2003).

Particularly interesting are temporal correlations in the form of neuronal synchronization. As said, neuronal synchronization is the phenomenon that neurons, in transient

assemblies, temporarily synchronize their activity (the aggregate of their cRFs then forms what Eckhorn 1999 called an association field). It has been related to cortical integration and, more generally, to cognitive processing (Milner 1974; von der Malsburg 1981). It is true that, as Shadlen and Movshon (1999) noted, one speaks of synchronization when neurons fire within a fairly arbitrarily chosen small time window, that is, the spikes do not have to be completely coincident in time. Empirically, however, it is a well-established phenomenon that has been associated with a broad range of cognitive processes (for reviews, see, e.g., Finkel et al. 1998; Gray 1999).

For instance, oscillatory synchronization in the theta, alpha, and beta bands (4–30 Hz) seems involved in interactions between relatively distant brain structures, while oscillatory synchronization in the gamma band (30–70 Hz) seems involved in relatively local computations (see, e.g., Kopell et al. 2000; von Stein and Sarnthein 2000). More specifically, theta, alpha, and beta synchronization have been found to be correlated with, for instance, top-down processes dealing with aspects of memory, expectancy, and task (see, e.g., Kahana 2006; van der Togt et al. 2006; von Stein et al. 2000). Furthermore, gamma synchronization has been found to be correlated particularly with visual processes—such as those dealing with change detection, interocular rivalry, feature binding, Gestalt formation, and form discrimination (see, e.g., Börgers et al. 2005; Fries et al. 1997; Keil et al. 1999; Lu et al. 2006; Singer and Gray 1995; Womelsdorf et al. 2006).

In this article, I have this “visual” gamma synchronization in mind. Next, I first briefly review DST research into the dynamics of synchronization, and then I discuss existing neuroscientific ideas about its function and meaning.

The dynamics of synchronization

Synchronization is a long-standing topic in DST (see, e.g., Pikovsky et al. 2001; Wu 2007). It probably started with Huygens (1673/1986) who observed that two pendulum clocks, coupled by suspending them from the same wooden beam, tend to synchronize their motion. From a DST point of view, this topic is intriguing because, in general, DST describes system behavior that, at first glance, seems chaotic and unpredictable—such systems seem to defy an orderly thing like synchronization (Pecora and Carroll 1990). To describe seemingly chaotic system behavior, DST uses the powerful mathematical tools called nonlinear partial differential equations (NPDEs) which, traditionally, find application mainly in physics (e.g., to make weather forecasts).

A differential equation typically describes the development of a system over time (where the “system” may be

anything one chooses it to be). It does not specify system states as such but, instead, it specifies the difference between any one state and the next (with arbitrarily small time steps). This implies that, to determine actual system states, also a starting state must be given. So-called linear differential equations can usually be solved analytically (yielding one formula which, for every starting state, specifies subsequent system states) and imply that a change in the starting state yields a proportional change in subsequent states. This does not hold for NPDEs, however. For different starting states, an NPDE may have different solutions, and a small change in the starting state may yield a dramatic change in subsequent states. Therefore, actual system states can usually only be determined numerically, that is, by way of subsequent applications of the NPDE.

To add some flavor, the state space refers to the set of all states, over all starting states, a system may arrive at according to an NPDE. A trajectory then is the sequence of states the system passes from a specific starting state, and an attractor is a state for which the system can be said to have a preference, that is, a relatively stable state reached for relatively many nearby starting states. Applied to perceptual organization, attractors can be said to correspond to cognitive states, or percepts (Eliasmith 2001)—they should not be too stable, though, because the system must be able to switch from one percept to another (Spivey 2007; van Leeuwen 2007). Furthermore, a strong point of DST is that potential behavior of a system under various imaginable settings can be investigated by varying parameters in the starting state or in the NPDE. This method is also used in DST studies on synchronization in networks, mostly in the context of vision research.

For instance, van Leeuwen et al. (1997) performed simulations with a sparsely connected network of nonlinear maps. They found that the coupling strength between the maps, in proportion to the rate of chaotic divergence, determines whether rapid transitions occur between unsynchronized and synchronized states of varying assemblies of maps (see also Buzsáki and Draguhn 2004). Furthermore, for networks of locally coupled integrate-and-fire oscillators, Campbell et al. (1999) investigated (de)synchronization parameters and found that the time to synchronize seems proportional to the logarithm of the network size, or in other words, that synchronization propagates exponentially. Moreover, gamma and beta rhythms seem to have different synchronization properties (Kopell et al. 2000), and for gamma rhythms, the time to synchronize seems to fit the gamma cycle (Harris et al. 2003).

These are in fact just a few of the many studies into the dynamics of synchronization in networks (see also, e.g., Izhikevich 2006; Li 1998; Roelfsema et al. 1996; Sporns et al. 1991; Yen and Finkel 1998; Yen et al. 1999). This

DST research does not affect the information-processing ideas in the model I discuss in section “[A representationally inspired algorithmic account](#)”, but it does provide necessary complementary insights into a question left open by this model. That is, in Marr’s (1982/2010) terms, this DST research is not about the computational goal or algorithmic method of the information process I attribute to gnosons (i.e., the transient assemblies of synchronized neurons), but it is about how the implementational means might allow gnosons to go in and out of existence.

Proposed meanings of synchronization

As said, neuronal synchronization seems to occur most notably in neural assemblies formed by horizontal connections, and these assemblies are also thought to mediate the binding of similar features. A binding function, but then referring to integration of different features, is reflected in the temporal correlation hypothesis (Milner 1974; von der Malsburg 1981; for a review, see Gray 1999). This hypothesis holds that synchronization binds those neurons that, together, represent one perceptual entity, say, an object or a Gestalt (see also Eckhorn et al. 2001; but see also Thiele and Stoner 2003). I think that synchronization is indeed related to perceptual organization, but I do not think it is a binding force, because that would beg the question of which neurons are to be bound (Shadlen and Movshon 1999). In other words, synchronization may signal what is going on, namely, perceptual organization, but it does not account for how perceptual organizations are computed.

Other ideas about neuronal synchronization are, for instance, that it underlies consciousness (Crick and Koch 1990; later, Crick and Koch 2003, rejected this idea), or that it is under the control of selective attention (Womelsdorf and Fries 2007), or that it is a marker that a steady state has been achieved (Pollen 1999), or that its strength is an index of the salience of features (Finkel et al., 1998; Salinas and Sejnowski 2001). In line with the latter idea, Fries (2005) proposed that more strongly synchronized assemblies in a visual area are locked on more easily by higher visual areas.

These ideas all sound plausible and may all contain some truth: as Sejnowski and Paulsen (2006) argued, neuronal synchronization may reflect a flexible and efficient mechanism subserving the representation of information, the regulation of the flow of information, and the storage and retrieval of information (see also Tallon-Baudry 2009). All those ideas, however, are about cognitive factors associated with synchronization rather than about the nature of the underlying cognitive process itself. Therefore, instead of saying that synchronization mediates cognitive processes, I prefer to say that it is a manifestation

of cognitive processing—just as the bubbles in boiling water are a manifestation of the boiling process (see also Bojak and Liley 2007; Shadlen and Movshon 1999).

This does not make synchronization less interesting—on the contrary, it raises the question of what form of processing it might be a manifestation. The goal of this process seems to be feature binding, but its method does not seem to be a simple form of parallel processing. In section “[Forms of processing](#)”, I go into more detail on forms of processing, but basically, parallel processing is performed by different agents who simultaneously do different things. When these agents simultaneously do the same thing, however, they seem to enter another processing mode—think of flash mobs or of groups of singers going from cacophony to harmony. Indeed, considering the complexity of perceptual organization, with its high combinatorial capacity and high speed, it must be a special form of processing that manifests itself by synchronization. In the next section, I discuss my algorithmic model of perceptual organization, incorporating not only the three intertwined subprocesses discussed above but also this special form of processing, called transparallel processing, whose neural signature is proposed to be neuronal synchronization.

A representationally inspired algorithmic account

In this section, I discuss my algorithmic model of perceptual organization. To give a proper impression of this model, it is expedient to begin by reviewing Leeuwenberg’s (1969, 1971) structural information theory (SIT), which is its underlying representational approach. SIT’s information-theoretic approach differs fundamentally from Shannon’s (1948) classical approach in that it starts from a totally different idea about how information is to be measured (for more details, see van der Helm 2000; see also Luce 2003). In the 1980s, SIT received considerable criticism, but as this section may be proof of, it has fully recovered from that criticism, and nowadays, it is probably the most elaborated representational approach to perceptual organization (Palmer 1999).

Structural information theory

For a proper appreciation of SIT, it is crucial to distinguish between the theory and the representational coding model implemented in my algorithmic model. SIT’s theory, on the one hand, is a coherent set of ideas about visual form perception (see this section “[Structural information theory](#)”)—its central idea being that the visual system selects the most simple interpretation of a given stimulus. SIT’s coding model and my implementation thereof, on the other hand, constitute a formal model that implements SIT’s

theoretical ideas, but then applied to patterned sequences of symbols (see section “[A transparallel processing model](#)”). This distinction is crucial because, as I address first, a persistent misunderstanding about SIT seems to be that it is thought to assume that the visual system converts visual stimuli into symbol strings

As I discuss more extensively in section “[Metaphors of cognition](#)”, any formal model uses and manipulates symbols. This holds for SIT’s model, just as it holds for DST and connectionist models. To design a formal model, the modeler decides what the symbols stand for, and more importantly, which principles are implemented. In DST models, these principles are reflected by NPDEs; in connectionist models, they are reflected by activation spreading through networks; and in SIT’s model, they are reflected by regularity extracting operations. Notice that, in each case, the principles are implemented to capture relationships between the things the symbols stand for, and that in this respect, SIT’s model is no exception.

It is true that, in the SIT literature, relatively much attention has been paid to how symbol strings might represent interpretations of visual stimuli, but this merely serves to illustrate how, in the empirical practice, the formal principles might be applied to visual stimuli in order to get testable quantitative predictions. That is, to be clear, SIT does not assume that the visual system converts visual stimuli into symbol strings. Furthermore, like any theory, SIT has limitations and open ends. For instance, it does not provide an algorithm that can take visual stimuli as input; hence, in the empirical practice, it is up to experimentators to choose and analyze relevant candidate interpretations in a perceptually plausible way. This may involve both 2-D and 3-D interpretations, and what matters in such analyses is that SIT’s theory assumes that the visual system employs the same information-processing principles as those which SIT’s model considers for strings.

Theoretical starting points

Representational approaches aim to gain insight into cognitive processes, and they do so by modeling systematicities in the output as a function of the input (i.e., what characterizes the nature of the output?). In the past, representational models may not have paid much attention to process mechanisms, but the idea of course was and still is that unraveling input-output systematicities is a first and necessary step towards proposing process mechanisms—after all, one has to know the goal before proposing a method to reach that goal. To this end, they focus on the informational content of mental representations which, as indicated before, can be taken to be relatively stable cognitive states arising during a dynamic neural process. Unlike DST and connectionist approaches, representational

approaches assume this process involves regularity extraction to get structured representations.

SIT takes the output to be a perceptual organization of an incoming visual stimulus. Detection of regularities such as symmetry and repetition subserves object perception and is believed to be an integral part of this perceptual organization process (Simon 1972; Tyler 1996; Wagemans 1997). Accordingly, SIT assumes that such regularities are extracted to construct candidate interpretations for a given stimulus, that is, candidate hierarchical organizations of the stimulus in terms of wholes and parts. It assumes further that the interpretation with the most simple descriptive code (i.e., the code that captures a maximum of regularity) is selected as the preferred interpretation.

SIT’s selection criterion, which is called the simplicity principle, is a descendant of Hochberg and McAlister’s (1953) minimum principle. Both are modern information-theoretical translations of the law of Prägnanz which Koffka (1935) proposed as a general principle in cognition (cf. Attneave 1954). In vision, this law has been proposed to underlie the various Gestalt laws of perceptual grouping (e.g., the laws of proximity, symmetry, similarity, and closure; Wertheimer 1923). Inspired by the minimum principle in physics, which refers to the tendency of physical systems to settle into relatively stable energy states, it states more specifically: of several geometrically possible organizations that one will actually occur which possesses the best, the most stable shape (Koffka 1935).

Hence, SIT models such a stable state as corresponding to a most simple descriptive code. As I discuss later on, connectionism models it as corresponding to a steady pattern of activation in a network, which, in DST terms, corresponds to an attractor in the network’s state space. Indeed, nowadays, all three approaches to cognition tend to find their roots in the Gestaltist motto that the whole is something else than the sum of its parts (cf. Sundqvist 2003; van der Helm 2006). Hence, they all aim to model aspects of the same thing—albeit in different terms and with noteworthy modeling differences.

For instance, to obtain good data fits, DST and connectionist modeling involves tuning of model parameters, whereas SIT’s approach is basically parameter-free (see section “[A transparallel processing model](#)”). Furthermore, unlike DST, both connectionism and SIT assume a competition between simultaneously present candidate outputs—but with a crucial difference. In connectionist models, a pre-defined network represents an output space for all possible inputs, and the process of activation spreading merely serves to select, for a given input, an output from this total output space. This contrasts with my SIT model which (a) first constructs an output space for only the input at hand and (b) then selects an output from this limited, input-dependent, output space. The selection

in (b) is performed in a way that, computationally, is comparable to connectionist activation spreading (see section “[Distributed processing](#)”). The construction in (a), however, is not standard in connectionist modeling and is probably the most distinguishing aspect of my model (see also sections “[A transparallel processing model](#)”, “[Connectionist modeling](#)”, and “[Distributed representations](#)”).

Theoretical developments

Since the 1960s, and in interaction with empirical research, SIT developed from a classical coding model of pattern classification (Leeuwenberg 1969, 1971; cf. Simon 1972) into a competitive theory of perceptual organization (Palmer 1999). To further specify the theoretical context of my algorithmic model, I next give a brief overview of these developments (see the included literature references for further details).

Nowadays, SIT includes a theoretically sound and empirically successful quantification of pattern complexity (van der Helm 1994; van der Helm et al. 1992), and an empirically successful quantitative model of amodal completion (van Lier 1999; van Lier et al. 1994). To predict preferred interpretations, this model applies a distinction and interaction between (viewpoint-independent) structural properties of candidate distal objects and (viewpoint-dependent) spatial relationships between these objects—reflecting the distinction and interaction between object perception and spatial perception, or between the ventral and dorsal streams in the brain (see Fig. 2b). Using findings from algorithmic information theory (see Li and Vitányi 1997), a Bayesian translation of this model led to the assessment that the simplicity principle is a general-purpose principle in that it promises to be fairly veridical in many different environments. This contrasts, in my view favorably, with the likelihood principle (von Helmholtz 1909/1962) which is a special-purpose principle in that it, by definition, is highly veridical in only one environment (for more details, see van der Helm 2000, 2002, 2007, 2011).

In addition, SIT nowadays includes an empirically successful quantitative model of symmetry perception (van der Helm and Leeuwenberg 1996, 1999, 2004). This model does not start from the traditionally considered transformational formalization of regularity (Garner 1974; Palmer 1983) which suits object recognition, but from a formalization that suits object perception (van der Helm and Leeuwenberg 1991). The latter defines visually relevant regularities as being holographic and hierarchically transparent. To give a gist, a stimulus regularity is holographic if all its substructures reflect the same kind of regularity; this allows its code to be built step-wise by going from

small to large substructures (think of an organism preserving its shape symmetry while growing). Furthermore, a stimulus regularity is hierarchically transparent if regularities nested in its code are stimulus regularities too (i.e., are also accessible separately from this code); this ensures that codes specify stimulus organizations with properly nested wholes and parts.

The properties of holography and hierarchical transparency pinpoint the unique formal status of the regularities called repetition, symmetry, and alternation (the latter covers, e.g., Glass patterns; Glass 1969). These regularities are generally considered to be visual regularities (i.e., regularities to which the visual system is sensitive), and in SIT, they are proposed to be extracted to construct candidate organizations of a given stimulus. As I discuss next, these regularities also have remarkable computational properties.

A transparallel processing model

SIT’s formal model of perceptual organization takes symbol strings as input. As said, this does not mean that SIT assumes that the visual system converts visual stimuli into strings—instead, the idea is that the visual system employs the same information-processing principles as those which SIT’s model considers for strings. The main principle is the simplicity principle, which implies that all candidate organizations of an input are considered and that the one with the most simple descriptive code is selected as the preferred organization. This principle is theoretically and empirically sound (see previous subsection), but it also suggests a daunting tractability problem (cf. Hatfield and Epstein 1985). Next, for strings, I first explicate this problem, and then I discuss my solution.

Defining the problem

To construct all candidate hierarchical organizations of a string, SIT’s formal model encodes the string by means of coding rules which extract the hierarchically transparent holographic regularities called repetition (or iteration I), symmetry (S), and alternation (A). These coding rules can be applied to any substring of the input string, and a code of the entire input string consists of a string of symbols and coded substrings, such that decoding the code returns the input string. In formal terms, SIT’s coding language is defined by:

Definition 1 A code \bar{X} of a string X is a string $t_1 t_2 \dots t_m$ of code terms t_i such that $X = D(t_1) \dots D(t_m)$, where the decoding function $D : t \rightarrow D(t)$ takes one of the following forms:

I-form:	$n * (\bar{y})$	$\rightarrow yyy \dots y$	$(n \text{ times } y; n \geq 2)$
S-form:	$S[(\bar{x}_1)(\bar{x}_2) \dots (\bar{x}_n), (\bar{p})]$	$\rightarrow x_1x_2 \dots x_n p x_n \dots x_2x_1$	$(n \geq 1)$
A-form:	$\langle (\bar{y}) \rangle / \langle (\bar{x}_1)(\bar{x}_2) \dots (\bar{x}_n) \rangle$	$\rightarrow yx_1yx_2 \dots yx_n$	$(n \geq 2)$
A-form:	$\langle (\bar{x}_1)(\bar{x}_2) \dots (\bar{x}_n) \rangle / \langle (\bar{y}) \rangle$	$\rightarrow x_1yx_2y \dots x_ny$	$(n \geq 2)$
Otherwise:	$D(t) = t$		

for strings y , p , and x_i ($i = 1, 2, \dots, n$). The code parts (\bar{y}) , (\bar{p}) , and (\bar{x}_i) are chunks; the chunk (\bar{y}) in an I-form or an A-form is a repeat; the chunk (\bar{p}) in an S-form is a pivot which, as a limit case, may be empty; the chunk string $(\bar{x}_1)(\bar{x}_2) \dots (\bar{x}_n)$ in an S-form is an S-argument consisting of S-chunks (\bar{x}_i) , and in an A-form, it is an A-argument consisting of A-chunks (\bar{x}_i) .

Hence, a code may involve not only recursive encodings of strings inside chunks, that is, from (y) into (\bar{y}) , but also hierarchically recursive encodings of S- or A-arguments $(\bar{x}_1)(\bar{x}_2) \dots (\bar{x}_n)$ into $(\overline{\bar{x}_1}(\overline{\bar{x}_2}) \dots (\overline{\bar{x}_n}))$. For instance, below, a string is encoded in two ways, and for each code, the resulting hierarchical organization of the string is given:

String:	$X = abacdacdababacdacdab$
Code 1:	$\bar{X} = ab2 * (acd) S[(a)(b), (a)] 2 * (cda) b$
Organization:	$a \ b \ (acd)(acd) \ (a)(b)(a)(b)(a) \ (cda)(cda) \ b$
Code 2:	$\bar{X} = 2 * (\langle (a) \rangle / \langle S[(\langle (b) \rangle)(\langle cd \rangle)] \rangle)$
Organization:	$(\ (\langle (a) \rangle)(\langle (a) \rangle cd) \ (\langle (a) \rangle cd) \ (\langle (a) \rangle)(\langle b \rangle)) \ (\ (\langle (a) \rangle)(\langle (a) \rangle cd) \ (\langle (a) \rangle cd) \ (\langle (a) \rangle)(\langle b \rangle))$


Code 1 does not involve recursive encodings, but Code 2 does: it is an I-form with a repeat that has been encoded into an A-form with an A-argument that, in turn, has been encoded into an S-form. These examples also illustrate the problem that a string generally has many codes—which all have to be considered to select a most simple one.

Notice that the exact definition of SIT’s complexity metric is not relevant in this article (the number of remaining symbols in a code can be taken as a good approximation) and that the problem lies in the huge number of candidate codes. This is analogous to the problem the visual system faces (see section “Introduction”). In fact, to expand this analogy, the code $2*(ab)$ of string $abab$, for instance, reflects a higher-level organization $2*(y)$ in which y refers to lower-level parts ab . This is analogous to how I imagine that wholes and parts are represented at different levels in the visual hierarchy in the brain (see section “The visual hierarchy”).

One may infer from Def. 1 that I-forms do not pose a big computational problem, but that a substring of length k can be encoded into $O(2^k)$ S-forms and $O(k2^k)$ A-forms. [The “big O” notation $O(g)$, with g some function, has a precise mathematical definition, but it means essentially “in the order of magnitude of g ”.] To pinpoint a most simple one, also most simple codes of the arguments of these S- and A-forms have to be determined, and so on—with $O(\log N)$ recursion steps because, for a substring of length k , the argument of a covering S- or A-form has maximally length $k/2$. Hence, if each S- and A-argument were to be recoded separately, then the entire process would require a super-exponential $O(2^{N \log N})$ amount of work which, to both computers and brains, could easily require more time than is available in this universe (cf. van Rooij 2008).

To solve this problem, I implemented the transparallel processing algorithm I presented earlier (see van der Helm 2004, also for its full formal and tractability details). Only later, I realized that the three intertwined subprocesses of feature encoding, feature binding, and feature selection—which this algorithm implements—correspond to the three subprocesses which, in neuroscience, are believed to take place in the visual hierarchy in the brain (see Fig. 5). To specify this correspondence, I next sketch how I modeled the three subprocesses, with a special eye for feature binding which is relevant to the synchronization issue (see section “Towards a pluralist account”) and, thereby, also

Fig. 5 **a** Copy of Fig. 3, depicting the three intertwined subprocesses that are believed to take place in the visual hierarchy in the brain. **b** The three corresponding and also intertwined methods implemented in the transparallel processing model of perceptual organization



to the cognitive architecture issue (see section “Cognitive architecture”).

Feature encoding

In the model, the subprocess of feature encoding involves an exhaustive search for hierarchically transparent holographic regularities (i.e., repetitions, symmetries, and alternations) in the input string, and hierarchically recursively, in the arguments of S- and A-forms. This subprocess corresponds to the feedforward sweep yielding an initial tuning, from lower to higher visual areas, to regularities to which the visual system is sensitive.

The search for regularities in the input string or in an S- or A-argument starts with a so-called all-substrings identification. This preprocess assigns identical numerals to identical substrings, so that the regularity search can identify identical substrings by these numerals instead of by, each time, a cumbersome symbol-by-symbol comparison. A naive method to do this preprocess would require $O(N^4)$ computing steps for a string of length N , but the model uses an $O(N^2)$ method which, in computer science, informally is called a smart method (I return to such methods in section “Distributed processing”).

Hence, this preprocess corresponds to an initial pick-up of information by which identical stimulus parts as such are encoded by identical neuronal responses. After this preprocess, it is easy to find separate regularities, but because of the hierarchically recursive nature of the search for regularities, a naive algorithm for an exhaustive search would require an unacceptable superexponential amount of work and time (see previous subsection). As I discuss next, a solution to this problem lies in feature binding by hyperstrings.

Feature binding

In the model, feature binding is implemented by gathering similar regularities in so-called hyperstrings—not as a goal in itself, but to allow for transparallel recoding of these regularities. To specify this crucial point, I begin with van der Helm’s (2004) graph-theoretical definition of hyperstrings (for details on graph theory, see Harary 1994).

Definition 2 A hyperstring is a simple semi-Hamiltonian directed acyclic graph (V, E) with a labeling of the edges in E such that, for all vertices $i, j, p, q \in V$, either $\pi(i, j) = \pi(p, q)$ or $\pi(i, j) \cap \pi(p, q) = \emptyset$, where a substring set $\pi(v_1, v_2)$ is the set of label strings represented by the paths from vertex v_1 to vertex v_2 ; the subgraph formed by the vertices and edges in these paths is a hypersubstring.

Hence, a hyperstring is a graph with, for N nodes, $O(N^2)$ links between the nodes and $O(2^N)$ paths from the first node




Fig. 6 A hyperstring. The 15 paths from vertex 1 to vertex 9 represent normal strings; for instance, the path along vertices 1, 3, 4, 5, 9 represents the string $xcfw$. Characteristic of hyperstrings is that the substring sets represented by hypersubstrings are either completely identical or completely disjoint, that is, never something in between. Here, as indicated in gray, the substring sets $\pi(1,4)$ and $\pi(5,8)$ are identical: the paths from vertex 1 to vertex 4 represent the same substrings (i.e., abc , ay , and xc) as those represented by the paths from vertex 5 to vertex 8

to the last node (see Fig. 6 for an example). Each of the links represents a string element, so that each of the paths through the graph represents a string (in which the nodes represent locations). In other words, a hyperstring on N nodes is a distributed representation of $O(2^N)$ strings, that is, it represents $O(2^N)$ strings in a distributed fashion (notice that this characteristic is usually associated with connectionist modeling). Presently most relevant is the special property of hyperstrings that substring sets represented by hypersubstrings are either identical or disjoint—never something in between. For instance, in Fig. 6, the substring sets $\pi(1,4)$ and $\pi(5,8)$ are identical, that is, they both represent the substrings abc , ay , and xc . The relevance hereof may be explicated, in two steps, by means of the following examples.

The string $ababfababgbabafbaba$ of length $N = 19$ can be encoded into $O(2^N)$ S-forms, for instance into $S[(a)(b)(a)(b)(f)(a)(b)(a)(b), (g)]$ and $S[(aba)(b)(f)(a)(bab), (g)]$. In Fig. 7a, the arguments of all these S-forms have been gathered in a distributed representation. For instance, the arguments of the two S-forms above are represented by the path along all vertices and by the path along vertices 1, 4, 5, 6, 7, and 10, respectively. In general, after the above-mentioned $O(N^2)$ all-substrings identification, the arguments of all S- and A-forms in a string can be gathered in $O(N)$ distributed representations like the one in Fig. 7a. Such a distributed representation can be constructed in $O(N^2)$ computing steps and, crucially, it consists provably of one or more independent hyperstrings (van der Helm 2004). In other words, the arguments of S- and A-forms group by nature into hyperstrings, so that, during the encoding, one does not have to check whether they do form hyperstrings—which is precisely what one would expect of an automatic binding mechanism.




Fig. 7 Hyperstrings of symmetry arguments. **a** The hyperstring representing the arguments of all S-forms into which the string *ababfababgbabafababa* can be encoded. **b** The hyperstring

representing the arguments of all S-forms into which the slightly different string *ababfababgbabafabab* can be encoded. The substring sets $\pi(1,5)$ and $\pi(6,10)$ are identical in (a) but disjoint in (b)

Furthermore, Fig. 7b shows that a small change in the input string may imply that substring sets represented by hypersubstrings turn from completely identical to completely disjoint. This illustrates that substring sets represented by hypersubstrings are either identical or disjoint, which implies that a hyperstring can be treated as if it were a single normal string. More specifically, it implies that all $O(2^N)$ S- or A-arguments in a hyperstring can be recoded simultaneously as if only one S- or A-argument were concerned, that is, in one go or, as I call it, in a transparallel fashion. For instance, the hyperstring in Fig. 7a can be seen as a string $h_1h_2 \dots h_9$ in which the substrings $h_1 \dots h_4$ and $h_6 \dots h_9$ are identical because the substrings sets $\pi(1,5)$ and $\pi(6,10)$ are identical. This implies that the string $h_1h_2 \dots h_9$ can be recoded into the S-form $S[(h_1 \dots h_4), (h_5)]$, without bothering about the different options $h_1 \dots h_4$ stands for (i.e., as if only one option were concerned).

Here, $h_1 \dots h_4$ stands for the substring set comprising $(a)(b)(a)(b)$, $(aba)(b)$, and $(a)(bab)$, so that $S[(h_1 \dots h_4), (h_5)]$ stands for the S-forms $S[((a)(b)(a)(b)), ((f))]$, $S[((aba)(b)), ((f))]$, and $S[((a)(bab)), ((f))]$. Eventually, one of these initial options may have to be selected, but also my selection method is indifferent to the number of these options (see below). The crucial point thus is that these options never have to be processed separately.

Hence, the underlying idea is that the visual system is sensitive to specific regularities (determined by identity relationships between parts), and that similar regularities automatically yield (or are bound into) hyperstring-like assemblies which allow these similar regularities to be hierarchically recoded in a transparallel fashion. Notice that this yields the combination of combinatorial capacity and speed the perceptual organization process is believed to have. Furthermore, notice that the hierarchically recursive recoding of hyperstrings yields a tree of hyperstrings, which represents all possible codes (of only the input string) in a hierarchical distributed representation. The final step then is to backtrace this hyperstring tree to select a most simple code of the input string.

Feature selection

In section “[Recurrent feature selection](#)”, I used the analogy of the cascade formed by a fountain under increasing water pressure, to illustrate what I think is the role of recurrent processing in the perceptual organization process. To recall, as the feedforward sweep progresses along ascending connections, each passed level in the visual hierarchy forms the starting point of integrative recurrent processing along descending connections. This yields a gradual buildup from partial percepts at lower levels in the visual hierarchy to complete percepts near the top end of the visual hierarchy. The model proceeds in the same way.

Already during the buildup of the hyperstring tree by the intertwined subprocesses of feature encoding and feature binding, the subprocess of feature selection starts to select most simple codes of increasingly larger (hyper)substrings, to select eventually a most simple code of the entire input string. This selection mechanism is implemented by applying, to each hyperstring, the $O(N^3)$ all-pairs version of Dijkstra’s (1959) $O(N^2)$ shortest path method (cf. Cormen et al. 1994; van der Helm and Leeuwenberg 1986). This is the method which, as I mentioned earlier and as I illustrate in section “[Distributed processing](#)”, is comparable to selection by activation spreading in connectionist models.

It is true that the encoding of a (hyper)string yields candidate subcodes of its (hyper)substrings, which in case of a hyperstring, add to the options represented initially in the hyperstring (see previous subsection). However, the intertwined selection of most simple subcodes implies that, no matter the number of these initial options, the maximum number of options in case of a hyperstring remains the same as in case of a single normal string. Hence, the transparallel treatment of those initial options also allows the selection mechanism to deal with a hyperstring as if it were a single normal string. In other words, the mechanism to select different features preserves the combination of high combinatorial capacity and high speed yielded by the transparallel recoding of similar features.

As said, full formal and tractability details can be found in van der Helm (2004), but to sum up, for a hyperstring on N nodes, the all-substrings identification requires $O(N^2)$ computing steps. Furthermore, the construction of all hyperstrings representing S- and A-arguments requires $O(N^3)$ steps, that is, $O(N^2)$ steps for each of $O(N)$ distributed representations. Finally, the all-pairs shortest path method requires $O(N^3)$ steps. Thus, for each hyperstring in the hyperstring tree, $O(N^3)$ steps are required. The depth of the hierarchical recursion is $O(\log N)$, so that the total process requires $O(N^{3+\log N})$ steps.

This contrasts very favorably with the superexponential $O(2^{N \log N})$ amount of work a naive algorithm would require. Due to the factor $\log N$, the model should probably be qualified as weakly exponential or near-tractable, but the $O(N^{3+\log N})$ is a generous worst-case upperbound, and in the average case, this factor $\log N$ hardly seems a problem. One could also restrict the hierarchical depth to the number of hierarchical levels in the visual hierarchy in the brain (see section “The visual hierarchy”), which would yield a fully tractable model.

Towards a pluralist account

Above, starting from a representational approach, I discussed an algorithmic model which is neurally plausible in that it incorporates the intertwined but functionally distinguishable subprocesses of feature encoding, feature binding, and feature selection. A pivotal point now is that this model has additional value in that it suggests that transparallel processing by hyperstrings provides a computational account of synchronization in transient neural assemblies—which complements DST research into this phenomenon. Even if details of this proposal turn out to be controversial, I think its pluralist nature indicates a promising direction for research in cognitive (neuro)science. To substantiate this, I next give a pragmatic line-up of meta-theoretical considerations which now and again expand on traditional views in a way that, in my view, is appropriate to relate representational, connectionist, and DST approaches to each other. First, I discuss philosophical metaphors of cognition; then, I discuss Marr’s (1982/2010) paradigmatic levels of description; finally, I discuss generic forms of processing to position the ones in my model.

Metaphors of cognition

Reality is something we experience subjectively. People may agree something is an objective reality, but this agreement is based on shared subjective experiences. Like traditional story-telling and religion, science is basically an endeavor to understand or control what many people

experience as reality, using metaphors whether or not expressed in concrete theories and models. The idea that science is about useful metaphors instead of objective truths may be uncomfortable, but to vision scientists in particular, it is evident that reality is in the eye of the beholder (cf. Lyons 1977; Socrates, 469–399 BC).

The currently dominant but often challenged metaphor in cognitive science is the computer metaphor. It is related to Putnam’s (1961/1980) computational theory of mind which, in the tradition of functionalism, promotes the idea that the workings of the mind can be understood in terms of information processing defined as computation, that is, as the conversion of an input by a set of rules into an output (see also, e.g., Edelman 2008b; Fodor 1981, 1997 2001; Haugeland 1982; Newell and Simon 1972; Pylyshyn 1984).

Opponents of this idea usually argue that the brain is a dynamic physical system and that the mind should be described accordingly (e.g., Smolensky 1988; van Gelder and Port 1995). However, having been trained in both, I see differences but no opposition. Some dynamicists, and perhaps even some computationalists, may interpret computationalism as assuming that the brain really manipulates discrete symbols, but as I argue next, this interpretation mistakes modeling tools for the things being modeled.

First, to be clear, the usage of symbols is inherent to all formal modeling, also within dynamic systems approaches. The very idea of formalization is that things, at a certain semantic level, are labeled by symbols—not for the sake of it, but to capture potentially relevant relationships between these things. For instance, in physics, formulas like Newton’s $F = ma$ are not assumed to be real things in nature but are merely tools to describe allegedly relevant relationships between allegedly relevant things in nature. Furthermore, even within the same research domain, formal models may differ in modeling tools, but this is often merely because some tools are more convenient than others to investigate potentially relevant relationships between things at the chosen semantic level.

Second, in my view, computationalism does not assume that the brain manipulates discrete symbols (which, to me, would be as odd as assuming that nature applies formulas like Newton’s $F = ma$). It merely uses conversion rules as formal tools to model the semantic structure of relatively stable cognitive states—independently of how the brain goes physically from one state to the next. These physical transitions, in turn, are modeled in dynamicism using other formal tools, namely, differential equations. Hence, whereas computationalism focuses on semantic structure, dynamicism focuses on physical change. This is analogous to the difference between the semantic structure of a computer algorithm, on the one hand, and the electrical currents in a computer, on the other hand.

Indeed, already before the dynamics versus computation debate began, Neisser (1967) characterized cognition as a dynamic information-processing system whose mental operations might be described in computational terms. In other words, instead of either dynamics or computation, it is both, and theories about either aspect may contribute equally to a more comprehensive understanding of cognition as a whole, precisely because they address different aspects. One might object that they use different tools and metaphors, but this is precisely one of the challenges which I, also in this article on perceptual organization, aim to overcome to understand cognition as a whole (see also Mitchell 1998).

For instance, thanks to Gestalt psychology (Koffka 1935; Köhler, 1920; Wertheimer 1912, 1923), it is nowadays commonly accepted that a percept is a relatively stable cognitive state which arises during a dynamic neural process. Initially, representational theory focused on the informational content of such stable cognitive states, and later, DST focused on the dynamics of the neural transitions from any one state to the next—of course, insight in both aspects is needed for a full understanding of perceptual organization. Connectionism is, in many respects, in between representational theory and DST, and as mentioned, all three approaches nowadays tend to find their roots in the Gestaltist motto that the whole is something else than the sum of its parts. That is, all three approaches aim to account for nonlinear behavior, meaning that a small change in the input may yield a dramatic change in the output. This is often presented as a trade-mark of DST, but it also holds for many connectionist and representational models (including SIT's model).

To return to the computer metaphor, it is of course just a metaphor, and by its metaphorical nature, it is about general processing principles rather than about specific process instantiations. Yet, related to the latter, I would like to make the following distinction between a narrow version (as the metaphor sometimes is interpreted by opponents) and a broad version (as the metaphor usually is interpreted by proponents):

Narrow computer metaphor: The digital computer is a model of the neural brain.

Broad computer metaphor (a.k.a. information-processing metaphor): Information processing by computers is a model of cognitive processing by the brain.

The narrow computer metaphor, on the one hand, follows the tradition of comparing the brain to the most sophisticated machine known at the time. In the past, machines such as the clock and the steam-engine had served as model of the brain, and in the twentieth century, it was the computer's turn to serve as model. A concrete model within this tradition aims to capture the serial

development over time of a system that, as a whole, goes from one state to the next. Such a system may, for instance, be a single neuron, or a group of neurons, or the brain as a whole. DST proponents may tend to reject the computer metaphor (e.g., van Gelder and Port 1995), but DST models do fit in this tradition: as I discussed in section “[The dynamics of synchronization](#)”, DST employs differential equations, which describe the strictly serial process by which a system goes from one state to the next.

The broad computer metaphor, on the other hand, suggests that cognitive processing can be modeled usefully in terms of information close to the everyday meaning of the word; these are also the terms in which computers can be programmed to process things. Hence, in contrast to previous metaphors, the broad computer metaphor does not refer to the hardware principle that the brain is a physical system, but it refers to software principles implemented in the brain to allow for cognition (see also Neisser 1967).

Such software principles are, in representational models, modeled by regularity extracting operations to get structured representations, and in connectionist models, by activation spreading through a network. Such a network typically is a distributed representation which, via combinations of connected pieces of information, represents many wholes. This concept stems from graph theory (see Harary 1994), and it is powerful in that the metaphor of interacting pieces can be used to efficiently evaluate many wholes (for more details, see section “[Distributed processing](#)”). Notice, however, that also my representationally inspired algorithmic model employs distributed representations (see section “[A transparallel processing model](#)”).

The latter suggests that the concept of distributed representations may bridge the gap between representational theory and connectionism. Furthermore, as I discussed in section “[The dynamics of synchronization](#)”, synchronization in networks is a topic in DST. It is true that DST models the states of such a network as a whole rather than individual interpretations represented by those states, but implicitly, such a network can also be seen as a distributed representation. This suggests that the concept of distributed representations may bridge the gap between connectionism and DST as well (see also, e.g., Spencer et al. 2009). Indeed, I think that, regarding cognitive architecture, distributed representations constitute the proverbial coin, with DST highlighting its neuronal side and representational theory highlighting its cognitive side. This may leave less room for connectionism as a theory, but it asserts connectionist modeling as a most powerful tool to implement realistic simulations of ideas within DST and representational theory (see also section “[Connectionist modeling](#)”).

Levels of description

Proponents of representational theory, connectionism, or DST may have criticized the others for not telling the whole story, but I actually think that none of these approaches alone tells the whole story. However, I also think that, together, they might tell a more complete story. For instance, as indicated above, connectionist modeling has both a representational side and a dynamic systems side, which suggests that the three approaches form a continuum (cf. Bem and Looren de Jong 2006). In other words, I think that the three approaches are complementary rather than mutually exclusive.

This agrees with Marr's (1982/2010) distinction between three separate but complementary levels of description of information processing systems:

1. *The computational level*—at which the goal of a system is specified in terms of systematicities in the system's output as a function of its input. Applied to the visual system, this level concerns the question of what logic defines the nature of resulting mental representations of incoming stimuli.
2. *The algorithmic level*—at which the method of a system is specified in terms of the mechanisms that transform the system's input into its output. Applied to the visual system, this level concerns the question of how its input and output are represented and how one is transformed in the other.
3. *The implementational level*—at which the means of a system is specified in terms of the hardware of the system. Applied to the visual system, this level concerns the question of how those representations and transformations are neurally realized.

To avoid misunderstandings, notice that Marr's distinction is a general distinction which can be applied recursively to any part of any system (or to any part of any model thereof) and that, just as Marr did, I apply to the visual system.

The labels Marr assigned to these levels were inspired by the rise of computers: computer programmers are well aware of the problem to compute something (the goal) by way of an algorithm (the method) implemented in certain hardware (the means). Others assigned different labels to basically the same levels. For instance, Dennett (1978) labeled them similarly by the intentional stance, the design stance, and the physical stance; Glass et al. (1979) labeled them similarly by the levels of content, form, and medium; and Pylyshyn (1984) labeled them similarly by the semantic level, the syntactic level, and the physical level. In fact, the relevance of the distinction between goal, method, and means was already emphasized by Aristotle (384–322 BC), and indeed, whatever the labels are, the distinction is relevant in many domains. For instance,

cooks are well aware of the problem to prepare a dish (the goal) by way of a recipe (the method) using certain ingredients (the means). Furthermore, in evolution theory, Darwin (1859) specified the goal (i.e., survival), Mendel (1866/1965) specified the method (i.e., heredity rules), and Watson and Crick (1953) specified the means (i.e., DNA).

The foregoing illustrates that the computational, algorithmic, and implementational levels yield descriptions of different aspects, and that they are complementary in that, together, they may explain how the goal is reached by a method that is allowed by the means. Cognitive (neuro)science still has a long way to go before it may arrive at a comprehensive theory which, even then, might well accommodate explanations at different levels of description. For instance, neuroscientists may argue that near-death and love experiences are the result of biochemical processes in the brain—and they may be right—but this does not yet do justice to people's conscious experiences which call for another story. In other words, I am open to what is called a metaphysical (or ontological) reading of pluralism (which assumes that a “grand unifying theory” is possible), but for the moment, I adopt an explanatory (or epistemological) reading of pluralism—which, more pragmatically, focuses on differences and parallels between existing explanations at different levels of description to see whether and how they might be combined (see also, e.g., Jilk et al. 2008).

Of course, it remains perfectly legitimate to focus on only the one or two levels of description that are most relevant to a research question at hand. Yet, also then, it is fruitful to have an eye for ideas that are compatible with all three levels—as I experienced in research on symmetry perception (see Csathó, van der Vloed and van der Helm 2003; Treder and van der Helm 2007; van der Helm and Leeuwenberg 1999, 2004). Furthermore, there are no strict borders between the three levels, but the distinction is useful not only to position ideas in the total field of cognitive science but also to assess whether ideas formulated at different levels, and thereby perhaps seemingly opposed, might yet be compatible.

Representational theory, connectionism, and DST are not confined to one level of description each, but their operating bases can be said to be the computational level, the algorithmic level, and the implementational level, respectively. That is, all three approaches are (at least verbally) concerned with all three levels, but as a rule, representational models start from ideas about the nature of mental representations, connectionist models from ideas about the transformations from input to output, and DST models from ideas about the neural realizations. This suggests that, like Marr's levels, also these three approaches are complementary rather than mutually exclusive. As mentioned in section “[Introduction](#)”, I aim to go farther

than just promoting this idea which can also be framed as follows.

Notice that a distinction can be made between representations and processes. The brain does not make this distinction, as DST proponents surely emphasize, but it is a crucial scientific distinction because it stresses that there are two basic questions: (a) the “what” question, which is the mostly computational and partly algorithmic question I addressed in section “[A representationally inspired algorithmic account](#)”, and (b) the “how” question, which is the partly algorithmic and mostly implementational question I addressed in section “[The visual hierarchy](#)”. This distinction reverberates the distinction which, according to Koffka (1935), Wertheimer made between the molar (or behavioral, or cognitive) and molecular (or physiological, or neural) levels.

As Marr noted, answering the what and how questions may be totally different endeavors, but answers to both questions are needed for a complete understanding. For instance, one might argue that gamma synchronization has already been explained in some sense by the empirically supported association with perceptual organization (see section “[Proposed meanings of synchronization](#)”). Side-stepping my feeling that this association is not an explanation but rather an observation to be explained, it could indeed be said to explain synchronization in some sense, namely, in the sense that it provides sort of an answer to the question of what synchronization is involved in—however, it does not answer the question of how it is involved.

Traditionally, representational models focus on the what question, whereas DST models focus on the how question (with, again, connectionist models somewhere in between). Thus far, DST approaches have addressed the phenomenon of synchronization (see section “[The dynamics of synchronization](#)”), but to my knowledge, representational approaches have not (in section “[Distributed representations](#)”, I discuss the few connectionist models that addressed it). The additional value of my algorithmic model now is that it implements a representational specification of this association with perceptual organization, employing a special form of processing that might be the form of cognitive processing that manifests itself by neuronal synchronization.

Forms of processing

Apart from the foregoing philosophical and paradigmatic issues, there is the metatheoretical issue of the forms of processing a theory or model might employ in its proposed process from input to output. Therefore, here, I discuss generic forms of processing to position the ones employed in my algorithmic model of perceptual organization.

To be clear, I do not aim to present a detailed taxonomy. For instance, Flynn (1972) distinguished classes of computer processes involving single or multiple instruction streams executed serially or in parallel on single or multiple data streams. Furthermore, Townsend (Townsend and Nozawa 1995) distinguished elementary cognitive processes, classifying them in terms of architecture, capacity, and stopping rule. Such taxonomies are helpful but also known to be nonexhaustive, and due to the novelty of transparallel processing, my model does not seem to fit neatly in existing taxonomies. Closest seems to be its qualification, in Townsend’s terms, as an exhaustive process using a coactive architecture yielding supercapacity—where coactive means that input from separate parallel channels is consolidated in a resultant common processor. This is not only close to what hyperstrings do, but it is also what Townsend feels is needed to account for perceptual organization.

What both taxonomies do indicate is that, apart from the number of processors involved, one also has to reckon with the structure of the data operated on. I therefore begin with the notion of distributed processing which sounds like referring to a specific form of processing, but which rather refers to a specific organization of data to be processed.

Distributed processing

The term distributed processing is often used to refer to a process that, instead of being executed by one processor, is divided over a number of processors. The latter does not yield a reduction in the work to be done, but it may yield a proportional reduction in the time needed—at least, if those processors operate in parallel. For instance, in the search for extraterrestrial intelligence project (SETI), a central computer divides the sky into parts, and it assigns each part to a different computer which analyzes this part and which returns its findings to the central computer. Thus, each of the computers does only part of the total job, and the total job is done by the computer network as a whole, which therefore is said to perform distributed processing. Saving time this way is of course relevant in practice, but theoretically, most interesting is the division of the sky into parts, which implies that the central computer maintains a distributed representation of the sky.

I therefore prefer to define distributed processing more generally (i.e., independently of the number of processors involved) as referring to a process that operates on a distributed representation of the data to be processed. Defined this way, distributed processing can yield a reduction in work (and, thereby, also in time): as I discuss in a moment, there are distributed representations which a process may exploit effectively to substantially reduce work. This is not the case in the SETI project, but it is part and parcel of my

algorithmic model and also of connectionist models. In these models, the work reduction depends on the nature of the distributed representations employed and not on the number of processors involved. For instance, connectionist models usually postulate networks of processors operating in parallel. Such a network is therefore said to perform parallel distributed processing. One might object that this usually is sustained only by a simulation on a single serially processing computer but, though the simulation takes extra time, this does not affect the proposed work-reducing principles. The only difference is that, in the simulation, the computer can be said to perform serial distributed processing.

In general, a distributed representation is a data structure that can be visualized by a set of interconnected nodes, in which pieces of information are represented by the nodes, or by the links, or by both. An example is the Internet, which connects pieces of information stored at different places. In the 1980s, distributed representations became popular in cognitive science due to connectionism, but already since the 1950s, properties and applications of distributed representations have been studied extensively in graph theory, which is a subdomain of both mathematics and computer science (cf. Harary 1994).

Work-reducing distributed representations are typically like road maps in which roads are represented by links between nodes representing places, so that routes are represented in a distributed fashion by successive links. Different wholes (i.e., routes) thus share parts (i.e., roads), and this is key to achieve a reduction in work. That is, for N nodes, such a distributed representation typically represents $O(2^N)$ wholes by way of only $O(N^2)$ parts. A process that has to search or select a specific whole, for instance, may exploit this and may confine itself to evaluating the $O(N^2)$ parts instead of the $O(2^N)$ wholes. This principle is part and parcel of what, in computer science, informally is called smart processing—because it typically reduces an exponential $O(2^N)$ amount of work to a polynomial $O(N^2)$ amount of work. For instance, suffix trees (cf. Gusfield 1997) and the data structure used in deterministic finite automata (Hopcroft and Ullman 1979) are, in computer science, well-known distributed representations used in smart search algorithms.

These smart methods can all be said to rely on interactions between parts in order to arrive at wholes—which, noteworthy, is also a central Gestalt principle. In fact, my model implements the subprocess of feature encoding using a smart method that implicitly uses suffix trees. Furthermore, it implements the subprocess of feature selection using Dijkstra's (1959) shortest path method, which falls in the same category of smart selection algorithms as the selection by activation spreading in connectionist models (see Fig. 8 for an informal connectionist

translation of Dijkstra's method). Its implementation of the subprocess of feature binding, however, takes the foregoing to a new level by using hyperstrings, which enables a reduction of exponential $O(2^N)$ amounts of work to constant $O(1)$ amounts of work. To position this form of processing further, I next go into some more detail on the role of distributed representations in connectionist modeling.

Connectionist modeling

Inspired by the brain's neural network, connectionism entertains the idea that cognitive behavior arises from activation spreading in a network that represents pieces of information in its nodes, or in its links, or in both (Churchland 1986, 2002; Churchland and Sejnowsky 1990, 1992; Smolensky 1988). The nodes are taken to be parallel processors, each typically doing little more than (a) sum its incoming activation, (b) change its state according to some function of this sum, and (c) modulate the activation it transmits as a function of some weight (cf. Fodor and Pylyshyn 1988). Hence, each node performs only part of the total job, and the network is therefore said to perform parallel distributed processing.

A seminal example is McClelland and Rumelhart's (1981) model of word recognition. Roughly, their network consists of (a) an input layer of nodes responding to letter strokes in pictures of words, (b) an output layer of nodes representing words, and (c) an intermediate layer of nodes which regulate the flow of activation between the input and output layers (in this model, these nodes represent letters, but in other models, this layer is also called a layer of hidden nodes). When fed with a picture of a word, activation spreads through the network until it settles in a relatively stable state—then, the most highly activated output node is taken to represent the word in the picture.

Nowadays, connectionist models come in many flavors (cf. Bechtel and Abrahamsen 2002). For instance, the represented pieces of information may or may not be at different levels of aggregation—if they are, as in the example above, the network is said to be hierarchical (cf. Miikkulainen and Dyer 1991). Furthermore, so-called feedforward networks do not allow activation to flow in circles, whereas so-called recurrent networks do. Moreover, in so-called localist networks, the output is given by a node, whereas in so-called distributed networks, it is given by a trace of successive links (or by the entire pattern of activation). The latter distinction corresponds to Smolensky's (1988) symbolic-subsymbolic distinction and is formally merely a matter of decomposition (Fodor and Pylyshyn 1988; Bechtel 1994). In contrast to localist networks, however, distributed networks allow for a flexible




Fig. 8 Parallel distributed processing implementation of Dijkstra's (1959) shortest path method to select an optimal flow path in a hilly tube system with six distribution nodes (nodes 0,1,...,5). The fluid used is such that it hardens within one time unit once it stops flowing. A link between two nodes i and j is a soft tube that expands as the fluid runs through it and consists of at most $j - i$ straight segments having slopes such that the fluid takes one time unit to cross a segment. Every node has a separate outlet for each outgoing tube, but only one inlet for all incoming tubes. An inlet has the same cross section as one fluid-filled tube, so, when the fluid reaches the inlet

through one or more tubes, the remaining tubes are automatically sealed off. At time $T = 0$, the fluid starts to be poured into node 0 and reaches node 2 at time $T = 1$, sealing off the tube between nodes 1 and 2. At time $T = 2$, the fluid has filled this dead-end tube, and the then nonflowing fluid therein has hardened at time $T = 3$. By then, the fluid has also already reached node 5. After that, there is still some filling of dead-end tubes and hardening of the fluid therein, but at times $T \geq 5$, the only remaining flow path consists of a minimal number of segments

clustering of represented “subsymbolic” parts into aggregates representing “symbolic” wholes.

In applications, a network typically is first fed with many inputs to tune its activation-spreading parameters such that the desired outputs tend to result; this training technique is called backpropagation. Subsequently, it is tested by feeding it novel inputs—then, a network is said to be robust if its performance is insensitive to small variations in the parameter setting, and if it also performs well, it is proposed to capture a relevant systematicity in the input domain. This systematicity may or may not be specified explicitly, but it seems in line with the philosophy of connectionism to say that it is an emergent property

which arises “automagically” from the process of activation spreading.

The foregoing shows that connectionism uses powerful modeling tools which seem suited to simulate cognition. However, backpropagation is basically just a form of data fitting, which suggests that connectionism may not be sufficient to explain cognition. For instance, I concur with Fodor and Pylyshyn (1988) who argued that connectionism may provide, at best, an account of the neural structures in which representational cognitive architecture is implemented (see also Bechtel 1994; Fodor and McLaughlin 1990).

Furthermore, standard connectionism rejects the representational idea that the brain performs regularity

extraction to get structured representations of incoming stimuli. This connectionist stance implies, as mentioned earlier, that activation spreading is merely a mechanism to select outputs from a pre-defined output space for all possible inputs. Considering all three subprocesses that are believed to take place in the visual hierarchy (see section “[The visual hierarchy](#)”), however, I think it is more plausible that, preceding such a selection, feedforward encoding and horizontal binding create an output space for only the input at hand. This is what my model does, and as I discuss in section “[Distributed representations](#)”, this does not exclude connectionist modeling, but it does call for a more flexible version thereof.

Finally, neuronal synchronization occurs in a neural network that can be said to perform parallel distributed processing. DST research focuses on how synchronization might arise in such a network (see section “[The dynamics of synchronization](#)”), and this is also the natural way in which connectionism might look at it. This, however, ignores that synchronization reflects a processing mode which, at least in representational terms, seems to yield a combinatorial capacity and speed that surpass the capacity and speed of standard parallel distributed processing (see section “[Proposed meanings of synchronization](#)”). This issue touches upon the question—discussed next—of how a process may operate on data whether or not organized in a distributed fashion.

From subserial to transparallel processing

Many everyday processes are hybrid in that they involve a combination of serial and parallel processing (see also Wolfe 2003). For instance, in a relay race, the teams run in parallel (i.e., simultaneously), but the members of each team run serially (i.e., one after the other). Likewise, at the checkout in a supermarket, the cashiers work in parallel, but each cashier processes customer carts serially. As I discuss here, however, there is more to processing than this traditional dichotomy between serial and parallel processing.

I begin with the observation that, at the checkout in a supermarket, an additional form of processing can be distinguished. That is, not only are the cashiers working in parallel, each cashier processing customer carts serially, but the different carts are also presented serially by different customers. This example indicates that, under appropriate specifications of “items” and “processors”, not just two but at least three forms of processing can be distinguished (see also Fig. 9):

1. *subserial processing*, in which items are processed one after the other by different processors;
2. *serial processing*, in which items are processed one after the other by one processor;

	One item at a time	Many items at a time
Many processors	<i>Subserial processing</i>	<i>Parallel processing</i>
One processor	<i>Serial processing</i>	<i>Transparallel processing</i>

Fig. 9 Forms of processing defined by numbers of processors and items processed at a time

3. *parallel processing*, in which items are processed simultaneously by different processors.

The supermarket example illustrates that these are three natural forms of processing—which probably occur also in the brain (where a processor may be defined by a neuron or by a group of neurons). Furthermore, the line-up of these three forms of processing in Fig. 9 suggests the existence of the form of processing I defined by:

4. *transparallel processing*, in which items are processed simultaneously by one processor.

Transparallel processing may look like science-fiction, but as I argued in section “[A transparallel processing model](#)”, it is mathematically sound and has already been implemented in my model of perceptual organization. In fact, as I illustrate next, it is also a natural form of processing.

Imagine that, for some odd reason, the longest pencil among a number of pencils is to be selected (see Fig. 10a). Then, one or many persons could measure the lengths of the pencils in a (sub)serial or parallel fashion—after which the longest pencil can be selected by comparing the outcomes of the measurements (see Fig. 10b). A much smarter method, however, would be if one person gathers all pencils in one bundle and places the bundle upright on a table—after which the longest pencil can be selected in a glance (see Fig. 10c). The smart part of this (of course also hybrid) method is that, once gathered, the pencils are not treated as separate items by one or many processors (here, persons) in a (sub)serial or parallel fashion, but that they are treated in a transparallel fashion, that is, simultaneously by one processor as if they constitute one item (i.e., a bundle).

To be clear, this example should not be confused with Dewdney’s (1984) spaghetti metaphor which illustrates a sorting algorithm. My example illustrates that, in some cases, items can be gathered in one bin after which they can be treated simultaneously as if only one item were




Fig. 10 Transparallel pencil selection. **a** Suppose the longest pencil is to be selected from among a number of pencils. **b** Then, one or many persons could measure the lengths of the pencils in a subserial, serial, or parallel fashion—after which the longest pencil can be

selected by comparing the outcomes of the measurements. **c** A smarter, transparallel, method would be if one person gathers all pencils in one bundle and places the bundle upright on a table—after which the longest pencil can be selected in a glance

concerned. In my model of perceptual organization, such transparallel processing has a positive efficiency effect on feature selection and integration, but it is employed primarily to efficiently recode similar features. To this end, as I discussed in section “[A representationally inspired algorithmic account](#)”, those similar features are gathered in distributed representations called hyperstrings, which allows those features to be recoded in one go, that is, in a transparallel fashion. Hence, the binding role of the bundle in the pencil example is analogous to the binding role of hyperstrings in my model, but hyperstrings serve a more sophisticated purpose, namely, transparallel recoding of similar features. This transparallel recoding by way of hyperstrings can be seen as a special form of distributed processing, and as I argue in the next section, it leads to a concrete pluralist picture of cognitive architecture.

Cognitive architecture

Going from brain to model, my model of perceptual organization is neurally plausible in that it incorporates the intertwined but functionally distinguishable subprocesses of feature encoding, feature binding, and feature selection—which, in neuroscience, are believed to take place in the visual hierarchy (see Fig. 5). To recall, the subprocess of feature encoding reflects an initial feedforward tuning of visual areas to features to which the visual system is sensitive; the intertwined subprocess of feature selection reflects a recurrent integration of different features into percepts; and, in between, the subprocess of feature binding reflects a horizontal binding of similar features. The latter subprocess may be a relatively underexposed topic in neuroscience, but it can be seen as the neuronal counterpart of the regularity extraction which, in representational theory, is proposed to lead to structured mental representations. Furthermore, at least in my model, it is key to allow for transparallel processing by hyperstrings—which, to my knowledge, is the first representationally inspired

mechanism proposed to do justice to both the high combinatorial capacity and the high speed of the perceptual organization process.

Inversely, going from model to brain, this transparallel mechanism may fill a gap in the understanding of neuronal synchronization. The model suggests that hyperstrings can be seen as formal counterparts of the transient horizontal assemblies of synchronized neurons which, in neuroscience, are thought to be responsible for binding similar features. Thereby, it also suggests that the synchronization in these assemblies can be seen as a manifestation of transparallel processing. In this sense, transparallel processing by hyperstrings provides a computational explanation of the dynamic phenomenon of synchronization in transient neural assemblies. This proposal of course needs further investigation (see also below), but as said, for one thing, it does justice to both the high combinatorial capacity and the high speed of the perceptual organization process.

Although my model was developed starting from a representational approach, it reflects a truly pluralist account in the spirit of Marr (1982/2010). First, it transcends traditional definitions of representational and connectionist approaches, in that it puts the representational idea that cognition relies on regularity extraction to get structured representations in a more dynamic perspective together with a more flexible version of the connectionist idea that cognition relies on activation spreading through a network. Second, its transparallel mechanism relates plausibly to neuronal synchronization, so that it also honors the DST idea that cognition relies on dynamic changes in the brain’s neural state. To summarize this like I did in section “[Metaphors of cognition](#)”, I think that, regarding cognitive architecture, distributed representations (as highlighted in connectionism) constitute the proverbial coin, with DST highlighting its neuronal side and representational theory highlighting its cognitive side. To discuss this further, I first revisit distributed representations.

Distributed representations

In connectionist terms, the hyperstrings in my model are distributed networks in which nodes represent locations in a localist fashion, while links represent spatial features (i.e., visual regularities) in a distributed fashion. Furthermore, they are the constituents of hyperstring trees which, in connectionist terms, are hierarchical networks. In such a hyperstring tree, a hyperstring is constituted by horizontal links representing featural information at some level of aggregation, and it is anchored vertically by the spatial information in the nodes. Moreover, backtracing a hyperstring tree to select a most simple code is a recurrent process. Hence, my model shares various characteristics with standard connectionist modeling, and in fact, a hyperstring tree corresponds to a recurrent hierarchical distributed network yielding a most highly activated trace of links as output.

Though beyond the scope of this article, it would be interesting to implement a formal connectionist version of this model. Inherent to the idea of complementarity, such a connectionist version does not have to be a literal translation. For instance, the strength of outcomes usually is a discrete variable in representational models and a continuous variable in connectionist models. This difference, however, seems without much consequence because, in the end, the ranking of outcomes is what matters most.

A more delicate point concerns neuronal synchronization which, to my knowledge, is a topic addressed by only few connectionist models (e.g., Hummel and Biederman 1992; Hummel and Holyoak 2003, 2005; Shastri and Ajjanagadde 1993). These models do not associate synchronization with binding of similar features, but with integration of different features. The neuroscientific evidence is admittedly still too scanty to decide, but it may well be associated with both. For instance, different sets of similar features might be represented in different assemblies of synchronized neurons, and the integration of different features might be reflected by simultaneous synchronization of these assemblies. Anyway, notice that my model does associate it with both. It suggests that synchronization already starts pre-selection with the binding of similar features (reflecting a regularity extraction that is absent in standard connectionist modeling) into hyperstring-like assemblies of synchronized neurons, whose combinatorial capacity is primarily exploited to efficiently recode similar features but, subsequently, also to efficiently select and integrate different features.

Furthermore, a major difference with standard connectionist modeling is that the hierarchical distributed network in my model does not refer to a relatively rigid neural network but to a cognitive network that shapes itself flexibly to the input at hand (which implies an efficient usage

of storage resources without increasing the order of magnitude of work to be done; see the end of section “[A representationally inspired algorithmic account](#)”). Just as I implemented my model in a computer, this flexible cognitive network is assumed to be implemented in the brain. As I discuss next, precisely this triggers a concrete picture of cognitive architecture.

From neurons to gnosons

As I mentioned in section “[Introduction](#)”, the idea that cognition is a dynamic process of self-organization is not new, and the idea that transient assemblies of synchronized neurons are the building blocks of cognition is not new either. That is, nowadays, it is widely accepted that neuronal synchronization is a cognitively relevant phenomenon, and gamma synchronization in particular has been associated strongly with perceptual organization (see section “[Proposed meanings of synchronization](#)”). Thus far, however, this idea lacked a computational explanation. My transparallel processing model now opens a concrete pluralist perspective on the cognitive architecture of perceptual organization. That is, it suggests the following picture.

Perceptual organization is mediated by a self-organizing, hierarchical, cognitive network which arises in the neural network of the brain. This network shapes itself to the input at hand and consists of hyperstring-like neural assemblies which signal their presence by synchronization of the neurons involved. These assemblies, or gnosons as I call them, are formed automatically by the extraction of regularities to which the visual system is sensitive. They represent similar regularities in a distributed fashion, supplying high combinatorial capacity and high speed by allowing many similar regularities to be hierarchically recoded in one go, that is, as if only one feature were concerned. These assemblies, with the high combinatorial capacity and high speed they supply, remain effective during the selection and integration of different features into percepts.

Of course, my model does not cover everything, and I cordially invite other researchers to provide additional input on how gnoson-forming regularity extraction might take place in the neural network of the brain, for instance. My present point, however, is that my model gives rise to a picture of flexible cognitive architecture constituted by self-organizing gnoson hierarchies arising in the relatively rigid neural architecture of the brain.

To conclude, the concept of gnosons may be grounded further as follows. Pascal (1658/1950) observed that a particular description of things usually reflects just one of an indefinite number of semantically related nominalistic levels in a hierarchy of possible descriptions. That is, concepts used at some level build on (or can be

decomposed into) lower-level concepts and form the building blocks for (or can be combined into) higher-level concepts. Both upward and downward in such a hierarchy of descriptions, there always seems to be room for additional levels, each with its own new concepts. For instance, particle physics currently takes quarks as the concepts at the lowest description level in physics, but superstring theory is an attempt to model them, at a still lower level, as vibrations of tiny supersymmetric strings (see Greene 2003).

Going upward, from quarks to consciousness, there are various levels of description, among which are the levels of atoms, molecules, and neurons. These concepts are taken to stand for the functional entities, or “processors”, at their respective levels. In between neurons and consciousness, there is cognition, and it seems fair to assume that, size-wise, cognitive processors must lie between individual neurons and the brain as a whole. For instance, in the past, the perceptron (a small single-layered network; Rosenblatt 1958) and the cognitron (a small multi-layered network; Fukushima 1975) have been proposed as formal counterparts of cognitive processing units. This line of thinking is continued by my proposal to conceive of input-dependent hyperstrings as formal counterparts of gnosons and to conceive of gnosons as constituents of flexible cognitive architecture.

Conclusion

Cognitive (neuro)science still has a long way to go before it may arrive at a comprehensive theory of perceptual organization, let alone of cognition as a whole. As I argued in this article, however, such a comprehensive theory might be obtained by combining complementary insights from representational theory, connectionism, and DST. Inherent to the idea of complementarity, insights from these different approaches do not have to be literal translations of each other. Rather, they might concern the different, but complementary, questions of (a) what is the nature of the outcomes of a process; (b) how does the process proceed; and (c) how are the process and its outcomes neurally realized.

In search for answers, I started from a representationally inspired algorithmic model which (a) is neurally plausible in that it implements intertwined but functionally distinguishable subprocesses which, in neuroscience, are believed to take place in the visual hierarchy in the brain; and (b) suggests that synchronization in transient neural assemblies in the visual hierarchy is a manifestation of transparallel processing. In the model, this special form of processing relies on hyperstrings, that is, special distributed representations which allow many similar features to be

recoded simultaneously as if only one feature were concerned. A naturally following suggestion is that those temporarily synchronized neural assemblies, or gnosons as I call them, are constituents of flexible cognitive architecture implemented in the relatively rigid neural architecture of the brain.

This proposal qualifies rather than challenges existing ideas about neuronal synchronization in the visual hierarchy, but its specifics of course need further investigation. Furthermore, I feel it is open to modulating effects of attention, but also this needs further investigation. For one thing, however, this proposal sketches a concrete pluralist picture of a neurally plausible cognitive architecture which accounts for the high combinatorial capacity and high speed of the human perceptual organization process.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Abrahamsen A, Bechtel W (2006) Phenomena and mechanisms: putting the symbolic, connectionist, and dynamical systems debate in broader perspective. In: Stainton R (ed) *Contemporary debates in cognitive science*, Basil Blackwell, Oxford, pp 159–185
- Adrian ED, Zotterman Y (1926) The impulses produced by sensory nerve endings: part II: the response of a single end organ. *J Physiol Lond* 61:151–171
- Ahissar M, Hochstein S (2004) The reverse hierarchy theory of visual perceptual learning. *Trends Cogn Sci* 8:457–464
- Attneave F (1954) Some informational aspects of visual perception. *Psychol Rev* 61:183–193
- Attneave F Prägnanz and soap-bubble systems: a theoretical exploration. In: Beck J (ed) *Organization and representation in perception*. pp 11–29 Erlbaum, Hillsdale, NJ, (1982)
- Barlow HB, Reeves BC (1979) The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vis Res* 19:783–793
- Bechtel W (1994) Levels of description and explanation in cognitive science. *Mind Mach* 4:1–25
- Bechtel W, Abrahamsen A (2002) *Connectionism and the mind*, 2nd edn. Blackwell, Oxford
- Bem S, Looren de Jong H (2006) *Theoretical issues in psychology*. SAGE Publications Ltd, London
- Bojak I, Liley DTJ (2007) Self-organized 40 Hz synchronization in a physiological theory of EEG. *Neurocomputing* 70:2085–2090
- Börgers C, Epstein S, Kopell NJ (2005) Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc Natl Acad Sci USA* 102:7002–7007
- Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *J Neurosci* 17:2112–2127
- Buzsáki G (2006) *Rhythms of the brain*. Oxford University Press, New York
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929
- Campbell SR, Wang DL, Jayaprakash C (1999) Synchrony and desynchrony in integrate-and-fire oscillators. *Neural Comput* 11:1595–1619

- Churchland PS (1986) *Neurophilosophy*. MIT Press, Cambridge
- Churchland PS (2002) *Brain-wise: studies in neurophilosophy*. MIT Press, Cambridge
- Churchland PS, Sejnowsky TJ (1990) Neural representation and neural computation. In: Lycan WG (ed) *Mind and cognition: a reader*, Blackwell, Oxford, pp 224–252
- Churchland PS, Sejnowsky TJ (1992) *The computational brain*. MIT Press, Cambridge
- Cormen TH, Leiserson CE, Rivest RL (1994) *Introduction to algorithms*. MIT Press, Cambridge
- Crick F, Koch C (1990) Towards a neurobiological theory of consciousness. *Semin Neurosci* 2:263–275
- Crick F, Koch C (2003) A framework for consciousness. *Nat Neurosci* 6:119–126
- Csathó Á, van der Vloed G, van der Helm PA (2003) Blobs strengthen repetition but weaken symmetry. *Vis Res* 43:993–1007
- Dale R (2008) The possibility of a pluralist cognitive science. *J Exp Theor Artif Intell* 20:155–179
- Dale R, Spivey M (2005) From apples and oranges to symbolic dynamics: a framework for conciliating notions of cognitive representation. *J Exp Theor Artif Intell* 17:317–342
- Darwin C (1859) *On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life*. John Murray, London
- Dennett DC (1978) *Brainstorms: philosophical essays on mind and psychology*. Harvester, Brighton
- Dewdney AK (1984) On the spaghetti computer and other analog gadgets for problem solving. *Sci Am* 250:19–26
- Dijkstra EW (1959) A note on two problems in connexion with graphs. *Numer Math* 1:269–271
- Donderi DC (2006) Visual complexity: a review. *Psychol Bull* 132:73–97
- Duncan J, Humphreys GW (1989) Visual search and stimulus similarity. *Psychol Rev* 96:433–458
- Eckhorn R (1999) Neural mechanisms of visual feature binding investigated with microelectrodes and models. *Visual Cogn* 3, 4:231–265
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanisms of feature linking in the visual cortex? *Biol Cybern* 60:121–130
- Eckhorn R, Bruns A, Saam M, Gail A, Gabriel A, Brinksmeier HJ (2001) Flexible cortical gamma-band correlations suggest neural principles of visual processing. *Visual Cogn* 8:519–530
- Edelman GM (1987) *Neural darwinism: the theory of neuronal group selection*. Basic Books, New York
- Edelman S (2008a) On the nature of minds, or: truth and consequences. *J Exp Theor Artif Intell* 20:181–196
- Edelman S (2008b) *Computing the mind: how the mind really works*. Oxford University Press, Oxford
- Eliasmith C (2001) Attractive and in-discrete. *Mind Mach* 11:417–426
- Engel AK, König P, Gray CM, Singer W (1990) Stimulus-dependent neuronal oscillations in cat visual cortex: intercolumnar interaction as determined by cross-correlation analysis. *Eur J Neurosci* 2:588–606
- Felleman DJ, van Essen DC (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1:1–47
- Finkel LH, Yen S-C, Menschik ED (1998) Synchronization: the computational currency of cognition. In: Niklasson L, Boden M, Ziemke T (eds) *ICANN 98, proceedings of the 8th international conference on artificial neural networks*. (Skövde, Sweden, 2–4 September, 1998). Springer, New York
- Flynn MJ (1972) Some computer organizations and their effectiveness. *IEEE Trans Comput C-21*:948–960
- Fodor JA (1981) *Representations: philosophical essays on the foundations of cognitive science*. Harvester, Hassocks
- Fodor JA (1997) Special sciences: still autonomous after all these years. *Philos Perspect* 11:149–163
- Fodor JA (2001) Language, thought, and compositionality. *Mind Lang* 16:1–15
- Fodor JA, McLaughlin B (1990) Connectionism and the problem of systematicity: why Smolensky’s solution does not work. *Cognition* 35:183–204
- Fodor JA, Pylyshyn ZW (1988) Connectionism and cognitive architecture, a critical analysis. *Cognition* 28:3–71
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci* 9:474–480
- Fries P, Roelfsema PR, Engel AK, König P, Singer W (1997) Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc Natl Acad Sci USA* 94:12699–12704
- Fukushima K (1975) *Cognitron: a self-organizing multilayered neural network*. *Biol Cybern* 20:121–136
- Garner WR (1974) *The processing of information and structure*. Erlbaum, Potomac
- Gilbert CD (1992) Horizontal integration and cortical dynamics. *Neuron* 9:1–13
- Gilbert CD (1993) Circuitry, architecture and functional dynamics of visual cortex. *Cereb Cortex* 3:373–386
- Gilbert CD (1996) Plasticity in visual perception and physiology. *Curr Opin Neurobiol* 6:269–274
- Glass L (1969) Moiré effect from random dots. *Nature* 223:578–580
- Glass A, Holyoak K, Santa J (1979) *Cognition*. Addison-Wesley, Reading
- Gray CM (1999) The temporal correlation hypothesis of visual feature integration: still alive and well. *Neuron* 24:31–47
- Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci USA* 86:1698–1702
- Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–337
- Gray CM, Engel AK, König P, Singer W (1990) Stimulus-dependent neuronal oscillations in cat visual cortex: receptive field properties and feature dependence. *Eur J Neurosci* 2:607–619
- Greene B (2003) *The elegant universe: superstrings, hidden dimensions, and the quest for the ultimate theory*. W W Norton & Company, New York
- Gusfield D (1997) *Algorithms on strings, trees, and sequences*. University Press, Cambridge
- Harary F (1994) *Graph theory*. Addison-Wesley, Reading
- Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–556
- Hatfield GC, Epstein W (1985) The status of the minimum principle in the theoretical analysis of visual perception. *Psychol Bull* 97:155–186
- Haugeland, J (ed) (1982) *Mind design: philosophy, psychology, artificial intelligence*. Bradford Books, MIT Press, Cambridge
- Hermens F, Herzog MH (2007) The effects of the global structure of the mask in visual backward masking. *Vis Res* 47:1790–1797
- Hochberg JE, McAlister E (1953) A quantitative approach to figural “goodness”. *J Exp Psychol* 46:361–364
- Hochstein S, Ahissar M (2002) View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron* 36:791–804
- Hopcroft JE, Ullman JD (1979) *Introduction to automata theory, languages, and computation*. Addison-Wesley, Reading
- Howe CQ, Purves D (2004) Size contrast and assimilation explained by the statistics of natural scene geometry. *J Cogn Neurosci* 16:90–102
- Howe CQ, Purves D (2005) Natural-scene geometry predicts the perception of angles and line orientation. *Proc Natl Acad Sci USA* 102:1228–1233

- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol Lond* 195:215–243
- Hummel JE, Biederman I (1992) Dynamic binding in a neural network for shape recognition. *Psychol Rev* 99:480–517
- Hummel JE, Holyoak KJ (2003) A symbolic-connectionist theory of relational inference and generalization. *Psychol Rev* 110:220–264
- Hummel JE, Holyoak KJ (2005) Relational reasoning in a neurally-plausible cognitive architecture: an overview of the LISA project. *Curr Dir Cogn Sci* 14:153–157
- Huygens C (1986) *The pendulum clock* (A F Muguet, Trans.). Ames: Iowa State University Press. (Original work published 1673)
- Izhikevich EM (2006) Polychronization: computation with spikes. *Neural Comput* 18:245–282
- Jilk DJ, Lebiere C, O'Reilly C, Anderson JR (2008) SAL: an explicitly pluralistic cognitive architecture. *J Exp Theor Artif Intell* 20:197–218
- Johnson KO (1980a) Sensory discrimination: decision process. *J Neurophysiol* 43:1771–1792
- Johnson KO (1980b) Sensory discrimination: neural processes preceding discrimination decision. *J Neurophysiol* 43:1793–1816
- Kahana MJ (2006) The cognitive correlates of human brain oscillations. *J Neurosci* 26:1669–1672
- Keil A, Muller EM, Ray WJ, Gruber T, Elbert T (1999) Human gamma band activity and perception of a Gestalt. *J Neurosci* 19:7152–7161
- Kelley TD (2003) Symbolic and sub-symbolic representations in computational models of human cognition. *Theory Psychol* 13:847–860
- Kelso JAS (1995) *Dynamic patterns: the self-organization of brain and behavior*. MIT press, Cambridge
- Kimchi R, Palmer SE (1982) Form and texture in hierarchically constructed patterns. *J Exp Psychol Hum Percept Perform* 8:521–535
- Koffka K (1935) *Principles of gestalt psychology*. Routledge and Kegan Paul, London
- Köhler W (1920) *Die physischen gestalten in Ruhe und im stationären Zustand* [Static and stationary physical shapes]. Vieweg, Braunschweig, Germany
- Kopell N, Ermentrout GB, Whittington MA, Traub RD (2000) Gamma rhythms and beta rhythms have different synchronization properties. *Proc Natl Acad Sci USA* 97:1867–1872
- Lamme VAF, Roelfsema PR (2000) The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci* 23:571–579
- Lamme VAF, Supèr H, Spekreijse H (1998) Feedforward, horizontal, and feedback processing in the visual cortex. *Curr Opin Neurobiol* 8:529–535
- Langley P, Laird JE, Rogers S (2009) Cognitive architectures: research issues and challenges. *Cogn Syst Res* 10:141–160
- Leeuwenberg ELJ (1969) Quantitative specification of information in sequential patterns. *Psychol Rev* 76:216–220
- Leeuwenberg ELJ (1971) A perceptual coding language for visual and auditory patterns. *Am J Psychol* 84:307–349
- Leeuwenberg ELJ, van der Helm PA (1991) Unity and variety in visual form. *Perception* 20:595–622
- Leeuwenberg ELJ, Mens L, Calis G (1985) Knowledge within perception: masking caused by incompatible interpretation. *Acta Psychol* 59:91–102
- Leeuwenberg ELJ, van der Helm PA, van Lier RJ (1994) From geons to structure: a note on object classification. *Perception* 23:505–515
- Lehar S (1999) Gestalt isomorphism and the quantification of spatial perception. *Gestalt Theory* 21:122–139
- Lehar S (2003) Gestalt isomorphism and the primacy of the subjective conscious experience: a gestalt bubble model. *Behav Brain Sci* 26:375–444
- Li Z (1998) A neural model of contour integration in the primary visual cortex. *Neural Comput* 10:903–940
- Li M, Vitányi P (1997) *An introduction to Kolmogorov complexity and its applications*, 2nd edn. Springer, New York
- Lu HJ, Morrison RG, Hummel JE, Holyoak KJ (2006) Role of gamma-band synchronization in priming of form discrimination for multiobject displays. *J Exp Psychol Hum Percept Perform* 32:610–617
- Luce RD (2003) Whatever happened to information theory in psychology? *Rev Gen Psychol* 7:183–188
- Lyons J (1977) *Semantics*. University Press, Cambridge
- Malach R, Amir Y, Harel M, Grinvald A (1993) Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc Natl Acad of Sci USA* 90:10469–10473
- Marr D (2010) *Vision*. MIT Press, Cambridge (Original work published 1982 by Freeman)
- McClelland JL, Rumelhart DE (1981) An interactive activation model of context effects in letter perception: part 1. An account of basic findings. *Psychol Rev* 88:375–407
- Mendel G (1965) Experiments in plant hybridisation. In: Bennett JH (ed) *Oliver and Boyd, London* (Original work published 1866)
- Miikkulainen R, Dyer MG (1991) Natural language processing with modular PDP networks and distributed lexicon. *Cogn Sci* 15:343–399
- Milner P (1974) A model for visual shape recognition. *Psychol Rev* 81:521–535
- Mitchell M (1998) A complex-systems perspective on the “computation vs. dynamics” debate in cognitive science. In: Gernsbacher MA, Derry SJ (eds) *Proceedings of the 20th annual conference of the cognitive science society*. Lawrence Erlbaum, Hillsdale, pp 710–715
- Moore CM, Mordkoff JT, Enns JT (2007) The path of least persistence: evidence of object-mediated visual updating. *Vis Res* 47:1624–1630
- Neisser U (1967) *Cognitive psychology*. Appleton-Century-Crofts, New York
- Newell A, Simon HA (1972) *Human problem solving*. Prentice Hall, Englewood Cliffs
- Nirenberg S, Latham PE (2003) Decoding neuronal spike trains: how important are correlations? *Proc Natl Acad Sci USA* 100:7348–7353
- Palmer SE (1983) The psychology of perceptual organization: a transformational approach. In: Beck J, Hope B, Rosenfeld A (eds) *Human and machine vision*, Academic Press, New York, pp 269–339
- Palmer SE (1999) *Vision science: photons to phenomenology*. MIT Press, Cambridge
- Pascal B (1950) *Pascal's pensées* (HF Stewart, Trans.). Routledge and Kegan Paul, London (Original work published 1658)
- Pavloski R (2011) Learning how to get from properties of perception to those of the neural substrate and back: an ongoing task of Gestalt psychology. *Humana Mente J Philos Stud* 17:69–94
- Pecora LM, Carroll TL (1990) Synchronization in chaotic systems. *Phys Rev Lett* 64:821–824
- Pikovsky A, Rosenblum M, Kurths J (2001) *Synchronization: a universal concept in nonlinear sciences*. Cambridge University Press, Cambridge
- Pollen DA (1999) On the neural correlates of visual perception. *Cereb Cortex* 9:4–19
- Putnam H (1961) *Brains and Behavior*. Presented at a conference of the American association for the advancement of science.

- Reprinted in Block N (ed) (1980), *Readings in philosophy of psychology*, vol 1. Harvard University Press, Cambridge, pp 24–36
- Pylyshyn ZW (1984) *Computation and cognition: towards a foundation for cognitive science*. MIT Press, Cambridge
- Pylyshyn Z (1999) Is vision continuous with cognition? The case of impenetrability of visual perception. *Behav Brain Sci* 22:341–423
- Roelfsema PR, Engel AK, König P, Singer W (1996) The role of neuronal synchronization in response selection: a biologically plausible theory of structured representations in the visual cortex. *J Cogn Neurosci* 8:603–625
- Rosenblatt F (1958) The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev* 65:386–408
- Salin PA, Bullier J (1995) Corticocortical connections in the visual system: structure and function. *Physiol Rev* 75:107–154
- Salinas E, Sejnowski TJ (2001) Correlated neuronal activity and the flow of neural information. *Nat Rev Neurosci* 2:539–550
- Schmidt KE, Goebel R, Lowel S, Singer W (1997) The perceptual grouping criterion of collinearity is reflected by anisotropies of connections in the primary visual cortex. *Eur J Neurosci* 9:1083–1089
- Sejnowski TJ, Paulsen O (2006) Network oscillations: emerging computational principles. *J Neurosci* 26:1673–1676
- Shadlen MN, Movshon JA (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 24:67–77
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:379–423, 623–656
- Shastri L, Ajanagadde V (1993) From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behav Brain Sci* 16:417–494
- Simon HA (1972) Complexity and the representation of patterned sequences of symbols. *Psychol Rev* 79:369–382
- Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586
- Smith LB, Samuelson LK (2003) Different is good: connectionism and dynamic systems theory are complementary emergentist approaches to development. *Dev Sci* 6:434–439
- Smolensky P (1988) On the proper treatment of connectionism. *Behav Brain Sci* 11:1–23
- Spencer, JP, Thomas, MSC, McClelland, JL (eds) (2009) *Toward a unified theory of development: connectionism and dynamic systems theory re-considered*. Oxford University Press, Oxford
- Spivey M (2007) *The continuity of mind*. Oxford University Press, Oxford
- Sporns O, Tononi G, Edelman GM (1991) Modeling perceptual grouping and figure ground segregation by means of active reentrant connections. *Proc Natl Acad Sci USA* 88:129–133
- Sun R (2004) Desiderata for cognitive architectures. *Philosophical Psychol* 3:341–373
- Sundqvist F (2003) *Perceptual dynamics: theoretical foundations and philosophical implications of Gestalt psychology*. Gothenburg University, Göteborg, Sweden
- Tallon-Baudry C (2009) The roles of gamma-band oscillatory synchrony in human visual cognition. *Front Biosci* 14:321–332
- Thiele A, Stoner G (2003) Neural synchrony does not correlate with motion coherence in cortical area MT. *Nature* 421:366–370
- Townsend JT, Nozawa G (1995) Spatio-temporal properties of elementary perception: an investigation of parallel, serial, and coactive theories. *J Math Psychol* 39:321–359
- Treder MS, van der Helm PA (2007) Symmetry versus repetition in cyclopean vision: a microgenetic analysis. *Vis Res* 47:2956–2967
- Treisman A, Gelade G (1980) A feature integration theory of attention. *Cogn Psychol* 12:97–136
- Tyler CW (1996) Human symmetry perception. In: Tyler CW (ed) *Human symmetry perception and its computational analysis*, VSP, Zeist, pp 3–22
- Tyler CW, Baseler HA, Kontsevich LL, Likova LT, Wade AR, Wandell BA (2005) Predominantly extra-retinotopic cortical response to pattern symmetry. *NeuroImage* 24:306–314
- Ungerleider LG, Mishkin M (1982) Two cortical visual systems. In: Ingle DJ, Goodale MA, Mansfield RJW (eds) *Analysis of visual behavior*, MIT Press, Cambridge, pp 549–586
- van der Helm PA (1994) The dynamics of Prägnanz. *Psychol Res* 56:224–236
- van der Helm PA (2000) Simplicity versus likelihood in visual perception: from surprisals to precisals. *Psychol Bull* 126:770–800
- van der Helm PA (2002) Natural selection of visual symmetries. *Behav Brain Sci* 25:422–423, 432–438
- van der Helm PA (2004) Transparallel processing by hyperstrings. *Proc Natl Acad Sci USA* 101:10862–10867
- van der Helm PA (2006) Review of the book *perceptual dynamics: theoretical foundations and philosophical implications of Gestalt psychology* by F Sundqvist (2003). *Philos Psychol* 19:267–279
- van der Helm PA (2007) The resurrection of simplicity in vision. In: Peterson MA, Gillam B, Sedgwick HA (eds) *In the mind's eye: Julian Hochberg on the perception of pictures, film, and the world*. Oxford University Press, Oxford, pp 518–524
- van der Helm PA (2011) Bayesian confusions surrounding simplicity and likelihood in perceptual organization. *Acta Psychol* 138:337–346
- van der Helm PA, Leeuwenberg ELJ (1986) Avoiding explosive search in automatic selection of simplest pattern codes. *Pattern Recogn* 19:181–191
- van der Helm PA, Leeuwenberg ELJ (1991) Accessibility, a criterion for regularity and hierarchy in visual pattern codes. *J Math Psychol* 35:151–213
- van der Helm PA, Leeuwenberg ELJ (1996) Goodness of visual regularities: a nontransformational approach. *Psychol Rev* 103:429–456
- van der Helm PA, Leeuwenberg ELJ (1999) A better approach to goodness: reply to Wagemans (1999). *Psychol Rev* 106:622–630
- van der Helm PA, Leeuwenberg ELJ (2004) Holographic goodness is not that bad: reply to Olivers, Chater, and Watson. *Psychol Rev* 111:261–273
- van der Helm PA, van Lier RJ, Leeuwenberg ELJ (1992) Serial pattern complexity: irregularity and hierarchy. *Perception* 21:517–544
- van der Togt C, Kalitzin S, Spekrijse H, Lamme VAF, Supèr H (2006) Synchrony dynamics in monkey V1 predict success in visual detection. *Cereb Cortex* 16:136–148
- van der Vloed G, Csathó Á, van der Helm PA (2007) Effects of asynchrony on symmetry perception. *Psychol Res* 71:170–177
- van Gelder T, Port RF (1995) It's about time: an overview of the dynamical approach to cognition. In: Port RF, Gelder T (eds) *Mind as motion: explorations in the dynamics of cognition*, MIT Press, Cambridge, pp 1–44
- van Leeuwen C (2007) What needs to emerge to make you conscious? *J Conscious Stud* 14:39449, 115–136
- van Leeuwen C, Steyvers M, Nooter M (1997) Stability and intermittency in large-scale coupled oscillator models for perceptual segmentation. *J Math Psychol* 41:319–344
- van Leeuwen C, Alexander D, Nakatani C, Nikolaev AR, Plomp G, Raffone A (2011) Gestalt has no notion of attention. But does it need one? *Humana Mentis J Philos Stud* 17:35–68
- van Lier RJ (1999) Investigating global effects in visual occlusion: from a partly occluded square to a tree-trunk's rear. *Acta Psychol* 102:203–220
- van Lier RJ, van der Helm PA, Leeuwenberg ELJ (1994) Integrating global and local aspects of visual occlusion. *Perception* 23:883–903

- van Rooij I (2008) The tractable cognition thesis. *Cogn Sci* 32:939–984
- von der Malsburg C (1981) The correlation theory of brain function. Internal Report 81-2, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany
- von Helmholtz HLF (1962) Treatise on physiological optics (J P C Southall, Trans.). Dover. (Original work published 1909), New York
- von Stein A, Sarnthein J (2000) Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *Int J Psychophysiol* 38:301–313
- von Stein A, Chiang C, König P (2000) Top-down processing mediated by interareal synchronization. *Proc Natl Acad Sci USA* 97:14748–14753
- Wagemans J (1997) Characteristics and models of human symmetry detection. *Trend Cogn Sci* 1:346–352
- Watson JD, Crick FHC (1953) Molecular structure of nucleic acids: a structure for deoxynucleic acids. *Nature* 171:737–738
- Wertheimer M (1912) Experimentelle studien über das Sehen von Bewegung. *Zeitschrift für Psychologie* 12:161–265
- Wertheimer M (1923) Untersuchungen zur Lehre von der Gestalt [On Gestalt theory]. *Psychol Forschung* 4:301–350
- Wolfe JM (2003) Moving towards solutions to some enduring controversies in visual search. *Trend Cogn Sci* 7:70–76
- Wolfe JM (2007) Guided search 4.0: current progress with a model of visual search. In: Gray W (ed) *Integrated models of cognitive systems*, Oxford University Press, Oxford, pp 99–119
- Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurobiol* 17:154–160
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736
- Wu CW (2007) *Synchronization in complex networks of nonlinear dynamical systems*. World Scientific Publishing, Singapore
- Yen SC, Finkel LH (1998) Extraction of perceptually salient contours by striate cortical networks. *Vis Res* 38:719–741
- Yen SC, Menschik ED, Finkel LH (1999) Perceptual grouping in striate cortical networks mediated by synchronization and desynchronization. *Neurocomputing* 26:609–616