Neural Oscillations as Representations

Abstract

We explore the contribution made by oscillatory, synchronous neural activity to representation in the brain. We closely examine six prominent examples of brain function in which neural oscillations play a central role, and identify two levels of involvement that these oscillations take in the emergence of representations: enabling (when oscillations help to establish a communication channel between sender and receiver, or are causally involved in triggering a representation) and properly representational (when oscillations are a constitutive part of the representation).

We show that even an idealized informational sender-receiver account of representation makes the representational status of oscillations a non-trivial matter, which depends on rather minute empirical details.

1   Introduction

A foundational hypothesis in cognitive science is that cognition progresses through the manipulation of representations: entities that are about, or stand for, other, generally extra-mental entities (Frankish and Ramsey 2012). According to this hypothesis, information about the external world flows through the brain, encoded somehow in its states, modulating behavior in increasingly sophisticated ways. Much cognitive science aims at identifying those brain states, and uncovering their content—what it is that they say about the external world. In its turn, one of the central research programs in the philosophy of cognitive science aims at formulating a metaphysics of representations: what needs to be the case for a certain vehicle to qualify as a representational state? What determines its representational content? (Cummins 1991; Ramsey 2007; Shea 2018)

1

The paragraph above is written in the abstract way characteristic of philosophical discussion on representations (e.g. Fodor 1974; Pylyshyn 1984); one that glosses over particular details of implementation. For example, philosophers often use the term 'vehicle' with the intention that it range over all possibly representational brain states (say, single neurons, populations of neurons, or neurotransmitter discharges) and all possible encodings (say, rate- or phase-based). Studying the metaphysics of representation in such non-committal terms is methodologically smart: it allows the resulting theory to apply to vehicles in general, and to capture whatever is common to all instances of representation in the brain. On the other hand, this approach tends to neglect the particular way in which representations are implemented, and the keys those particular ways might hold to the more general question of representation—apart from their intrinsic theoretical interest (Craver 2007; Boone and Piccinini 2016).

In this paper we redress this situation with respect to one very interesting, philosophically underexplored kind of vehicle. Frequently, brain activity is organized into synchronous, quasi-periodic patterns of activation which appear to contribute to many aspects of cognition, from pre-attentional grouping (Jensen, Kaiser, and Lachaux 2007; Zion Golumbic et al. 2013; Fries 2015; Pritchett et al. 2015) to the construction and modification of spatial maps in the hippocampus (Skaggs et al. 1996; Colgin et al. 2004). The question that will interest us here is whether these contributions are representational: do synchronous, periodic patterns of activity (we will call them 'neural oscillations', for brevity) in and of themselves constitute representations?

Investigation of these implementational details is not only interesting in its own right; it can also help philosophers reach conclusions about representation in general. In particular, the results of this paper can be used as a response to the 'trivialization' objection against naturalistic theories of content: many theorists have recently argued that theories that try to explain the notion of 'representation' in terms of functions, information, or causal relations (for example, teleosemantic approaches) are too liberal. According to this complain, if representational status merely depended on the presence of these features, almost any brain event would qualify as such (Ramsey 2007; Burge 2010; Schulte 2017; Gładziejewski and Miłkowski 2017; Butlin 2018; Williams and Colling 2018). This outcome would trivialize the notion of 'representation' and would put its explanatory role into question. Partly for this reason, some of these theorists suggest that attribution of representational status should be restricted to relatively sophisticated processes, such as those

involving map-like structures or constancy mechanisms.

Neural oscillations show this to be a hasty conclusion: as we will see, under an understanding of representations that would certainly count as liberal by those theorists' lights, some forms of simple oscillatory processes qualify as representational, whereas others do not. "Liberal" naturalistic theories of content can draw a meaningful, non-trivial distinction between brain events that are genuinely representational and those that are not.

Our paper is structured in six main sections. In section 2 we sketch our preferred general approach to the question of the metaphysics of representation. Our aim in this section is not to offer a fully developed metasemantics but, rather, to identify a common core to which many theorists working on this topic adhere. This common core will be enough to distinguish two roles neural oscillations can play, in the process of generating representations: we call them *enabling* and *properly representational*. We present these roles in section 3. The following two sections consider some scientifically prominent cases of brain function, with the goal of ascertaining which one, if any, of these roles neural oscillations play in them: section 4 discusses neural oscillations that probably play an enabling role; and section 5 deals with what, we argue, are the properly representational cases. This taxonomy exemplifies a second lesson about the usefulness of investigating particular details of implementation for the study of representations in general: while we have arrived at it from the consideration of rather minute such details, it is not unreasonable to think that it will prove helpful in the investigation and description of many other, unrelated representational phenomena. Section 6 wraps up and offers some conclusions.

2 The core metaphysics of content

In subsequent sections we will ask of certain kinds of brain activity whether they count or not as representations. In order to answer this question, we need a theory of what makes a certain state representational. A substantial discussion of this question is well beyond the scope of this paper; instead, we will present, without argument, two tenets that are widely (though by no means universally) thought to be part of what it is for a representation to be a representation. These tenets are most closely related with so-called *teleosemantic* naturalistic metasemantics (Millikan 1984; Papineau 1987; Neander 2017) but they also draw from the partially overlapping *signaling games*

framework (Skyrms 2010; Godfrey-Smith 2014).

The first tenet is that all representational systems share an architectural motif: in all of them, representational vehicles, or 'signals', mediate between a mechanism producing the representation (called 'sender' or 'producer') and a mechanism using the representation ('receiver' or 'consumer'). For a mechanism to qualify as a sender or a receiver, it needs to have the right sort of biological function: the sender's function is, roughly, to emit the above mediating signals when certain states of affairs obtains, whereas the receiver's functions is, roughly, to act in certain ways upon reception of the signal.

The content of the representation will be fixed, among other things, by the relation it bears to the behavior of the producer/sender and the consumer/receiver. Although the specification of these relations vary from account to account, information generated in the external world flowing from producer to the representation, and from the representation to the consumer, is usually taken to play a role: under a first approximation, contents are fixed by the information that representations carry about the world, under some designated set of optimal conditions. What these optimal conditions amount to is somehow fixed by the receiver/consumer's biological function—this is the second tenet.

Putting these two ideas together, our metasemantic sketch says that a certain mental vehicle, V, is a representation if:

1. V is an intermediate state in a sender-receiver architecture, and

2. V helps fulfil its consumer/receiver's biological function by carrying the information that the state of affairs S is the case (information which has flown to V through its producer/sender)

In such a situation, moreover, we may say that the content of V is that S is the case. In this paper, though, our main focus will be on establishing that certain neural oscillations are representations, and not so much on ascertaining their content.

This metasemantic sketch is heavily simplified—Ruth Millikan, and many other theorists after her, have been developing related ideas for the best part of four decades. Still, it is detailed enough to accommodate a number of properties that many associate with, or even take to be necessary for, the existence of representations. First, the possibility of erroneous representation, or misrepresentation: this will happen, among other things, in some cases

4

of sender *malfunction*, the possibility of which is part and parcel of its having a biological function. Second, the fact that representations ought to be action-guiding—or, at least, somehow contribute to the production of behavior: this will fall out of understanding biological functions as contributions to the subsistence of the individual (Mossio, Saborido, and Moreno 2009) or contributions to selection for the functional device in question (Millikan 2002). From section 4 onwards we will rely on these ideas in our discussion of the representational status of oscillations. Before that, we will briefly distinguish two possible degrees of involvement for oscillations to have in the generation of representations.

3   Enabling and representational roles

The taxonomy we are about to present is not a theoretical assumption, but one of the results of the present investigation: after having surveyed many of the relevant empirical cases, it is the classification that suggests itself as most conducive to understanding and organizing the involvement of neural oscillations in the generation of representations. In any event, of course, other useful taxonomies are certainly possible, and our preferred one does not aim at being exhaustive. The two roles under which we will classify the involvement of neural oscillations in representational phenomena are:

Enabling: Neural oscillations enable or trigger the activation of a representation.

Representational: Neural oscillations properly are (or are a constitutive part of) a representation.

First of all, neural oscillations play an enabling role when their main job is to help set up the communication channel between sender and receiver. To see what this means, we first note that a sender-receiver configuration is just a *point-to-point information-processing pipeline* in the Shannonian tradition (Shannon 1948; El Gamal and Kim 2011, section 3.1). Compare the entirely analogous figures 1 and 2. The sender/producer can be thought of as well as an *encoder*, and the receiver/consumer as a *decoder*. Efficient encoding/decoding is useful in order to get relevant information through inherently noisy, limited-capacity channels, and some of the time it is coding-related roles that neural oscillations will play: for example, oscillations can

5

help communicate to the encoder/decoder what is the task-relevant information they should focus on communicating (this is related to what information theorists call a *distortion measure*—see Martínez (2019) for details). This is our interpretation of the phenomena discussed in Subsection 4.1. They can also constitute the mechanism that allows the brain to move from reading representations to writing them, and back (Subsection 4.2). A third kind of enabling role is to help build a representation, without being a proper part of it (Subsection 4.3). There are probably many other purely representation-enabling roles that neural oscillation can play.
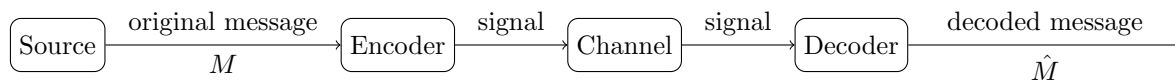
Source $\xrightarrow[\text{M}]{\text{original message}}$ Encoder $\xrightarrow{\text{signal}}$ Channel $\xrightarrow{\text{signal}}$ Decoder $\xrightarrow[\hat{M}]{\text{decoded message}}$

Figure 1: A point-to-point information-processing pipeline

World $\xrightarrow{\text{state}}$ Sender $\longrightarrow$ Messages $\longrightarrow$ Receiver $\xrightarrow{\text{act}}$

Figure 2: A sender-receiver model

The job of neural oscillations, however, is not limited to facilitating (or disrupting) communication. Sometimes they seem to play a *bona fide* representational role. In other words, they are a constitutive part of the vehicle of communication—of the signals in figure 1. Neural oscillations not only enable the flow of information, but also convey information themselves. In Subsections 5.1 and 5.2 we discuss two cases in which oscillatory phenomena not just enable but *are* representations. Finally, in Subsection 5.3 we take a step back and canvass the general role that oscillations play in the generation of an important kind of representational vehicle—so-called *neural sequences*—across the brain.

4 Enabling oscillations

In this section we review cases in which oscillations provide background conditions for representations to exist, or are causally involved in the tokening of representations, without being themselves representational.

4.1 Attention and Communication Through Coherence

Attention is a kind of modulation of sensory processing, whereby incoming stimuli are preferentially processed, or ignored, based on their current saliency, or behavioral relevance to the perceiver (Treue 2001; Gilbert and Li 2013). Attentional influence helps target finite and expensive computational resources in higher cognition to the treatment of those aspects of the incoming sensory stream that most directly affect the perceiver. "As a consequence the perceptual quality of a visual stimulus located within an attended region is improved ... at the expense of the perceptual quality of stimuli located elsewhere ..." (Chalk et al. 2010, p. 114).

Attentional influence on sensory processing has been demonstrated to exist as early on as V1 (Posner and Gilbert 1999) but probably not earlier than that (Alilović et al. 2019); this influence increases as we move on to higher visual areas (Maunsell and Cook 2002) although "the magnitude of attentional effects depends on the nature of the task and the configuration of the stimulus" (Gilbert and Li 2013, p. 352).

Spatial attention is often directed to concrete spatial areas in the visual field. The once ubiquitous metaphor was that attention is like a 'spotlight': it illuminates the attended area and singles it out for further processing (Brefczynski and DeYoe 1999; Hurlbert and Poggio 1985) More recent research has uncovered some limitations of this spotlight metaphor. First, the spatial resolution of neurons decreases as we go up in the visual processing hierarchy (that is, the receptive field size of these neurons increases, Dumoulin and Wandell 2008): the higher the visual area, the bigger chunks of the visual fields particular neurons are affected to. This makes intuitive sense: these neurons are often, though not always, attuned to global properties of the visual scene—to its gist; and not so much to fine-grained details, say, of texture or color. This means that attention cannot be just a matter of upstream neurons specializing on (directing a spotlight to) smaller, spotlight-sized, visual-field regions. Furthermore, the existence of attentional mechanisms that target objects and features, and not spatial regions, has been amply demonstrated

(Maunsell and Treue 2006; O'Craven, Downing, and Kanwisher 1999; Treue and Trujillo 1999). Attentional mechanisms, therefore, somehow are able to single out, from the whole input to a neuron's receptive field, a relevant subset of activity for further treatment.

One of the most widely accepted hypothesis as to how attention accomplishes this, Pascal Fries's *communication through coherence* (also CTC henceforth, Fries 2005, 2015), relies on synchronized[1] oscillations in the gamma band (or *gamma oscillation*, for short, Pritchett et al. 2015; Jensen, Kaiser, and Lachaux 2007; Fell et al. 2003; Zion Golumbic et al. 2013; Fries 2009). In this section we first discuss the mechanism by which synchronized oscillatory behavior is accomplished; then how this solves the problem that attention faces; and, finally, whether attention-related gamma oscillation could be plausibly regarded as representational.

Gamma oscillation depends crucially on the behavior of fast-spiking inhibitory interneurons (Pritchett et al. 2015; Fries 2009). Simplifying somewhat current computational models of the emergence of oscillatory behavior (see, e.g., Börgers, Epstein, and Kopell 2008, for a fuller picture), the main idea is that a stimulus presentation excites a group of interneurons, which in turn inhibit a larger group of excitatory neurons in their vicinity. When the inhibition wears out (after approximately 15ms, hence the gamma frequency of the resulting rhythm, Pritchett et al. 2015, p. 254) there is a window of opportunity for the activity of excitatory neurons, which in turn generate the next cycle of inhibition by interneurons.

Attention is hypothesized to work roughly as follows: the stimulus attended to by a population of neurons downstream, D, oscillating in the gamma band, is the one encoded by the population upstream, U, also oscillating and synchronized with D. The way in which this helps fix attention to the stimulus encoded by U is by enhancing the *effective connectivity* (Friston 2011) between D and U (and impeding the connectivity between D and other possible neuronal populations upstream). First, focusing on the downstream neuron population, "input consistently arriving at high-gain phases benefits from enhanced effective connectivity" (Ni et al. 2016, p. 240). Second, focusing on the upstream population, "[presynaptic synchronization] ...ensures that a presynaptic activation pattern arrives at postsynaptic neurons in a temporally coordinated manner" (Fries 2015), which in turn results in much enhanced impact in postsynaptic neurons because of feedforward coincidence

---

1. Here, by "synchronized" we mean "in (delayed) coherence", see Fries (2015).

detection (Fries 2009). This is, in a nutshell, the mechanism that Fries calls CTC.

Our main question is: are attention-related gamma oscillations representational? To be precise, the putative representational vehicle is

Vehicle: Populations U and D being synchronized in the gamma band.

Under the simplified analysis of representations we are working with (as presented in section 2) a representation is, at least, a vehicle that mediates between a sender (or producer) and a receiver (or consumer). Now, on the one hand, it *is* possible to find a sender, or set of senders, for *Vehicle*. First, for saliency-based attention, the most salient stimulus is the one that will drive excitatory neurons most vigorously. This will kick off the process described above, that results in gamma oscillation for these neurons, which in turns "manages to entrain postsynaptic neurons and thereby achieves the increase in input gain at the postsynaptic neurons." (Fries 2015, p. 226). The sender of the attentional vehicle, according to this picture, would be neural (e.g., retinal) activity vigorous enough to entrain gamma oscillation. For top-down attention, the CTC picture is somewhat less clear. Roughly, whatever volitional processes that result in a mandate to attend to a certain spatial region, feature or object would directly communicate with U (recall, this is the population of upstream, presynaptic neurons), causing them to synchronize, and thereby entraining D. The sender, according to this picture, would be something like the neural correlate of an intention to focus one's attention on the feature encoded by D.

But, on the other hand, it is very unlikely that *Vehicle* have a receiver: that D attends to U is not something that needs to be communicated to another area of the brain for further treatment. There is no further component that is sensitive to the synchronization and employs it to gather information about some state of the world. Here neural oscillations enable a better communication between two brain areas, but they are not supposed to represent or carry information about any particular aspect of the world. Attentional modulation of sensory processing is a very efficient way of optimizing the brain's limited computational budget, but the information that this optimization might carry is of no use to other brain areas.[2]

---

2. At least in usual cases, leaving aside comparatively uncommon situations of self-monitoring of attention (Harris et al. 2005).

Note that our claim here is just that *Vehicle* (that is, the actual neural-oscillatory implementation of attention) is not representational. We are not discussing other phenomena in the vicinity, such as, for example, the intention to focus our attention on some aspect of the visual field, which, as suggested above, could plausibly trigger an attentional process in some cases. The communication of this intention to neural population U *could*, for all we know, constitute an imperative representation with a content along the lines of "Attend to stimulus S!". We will not discuss here this possible imperative representation[3] among other things because the philosophical treatment of imperative representations is less mature than that of indicative ones. See Butlin (2018) and Artiga (2013) for related discussions.[4]

What attention does falls squarely under *channel management*: given the available channel capacity from early vision to the extrastriate cortex, information relevant to current behavioral goals (top-down attention) or information that might potentially inform new behavioral goals (saliency-based attention) needs to be prioritized. Attention-related gamma oscillations play a purely enabling, non-representational role.

4.2   Retrieval and encoding of cognitive maps

Our second example concerns place cells. Research on these neurons began in the 1970s, when O'Keefe and Dostrovsy (1971) discovered that activity in some cells of the rat's hippocampus (specially in areas CA1 and CA3) were not well correlated with significant events (such as food finding or lever pressing), but were instead specially sensitive to location. Subsequent research has confirmed that these cells carry information about particular places and many have taken this to suggest that rats possess a cognitive map of the environment (O'Keefe and Nadel 1978; Muller et al. 1994; McNaughton et al. 2006; Hartley et al. 2017).

Cognitive maps must be stable enough that they can be re-used whenever the rat needs to navigate the same space. At the same time, in certain circumstances (e.g. in new environments, or when familiar environments change in significant ways) a new map needs to be created, a process called 'remapping'

---

3. A representation which, by the way, would also be possibly mediated by gamma synchrony, this time between the frontal eye field and V4 (Gregoriou et al. 2009). See Baluch and Itti (2011) for a review of top-down influences in attention.

4. We would like to thank an anonymous reviewer for drawing our attention to this point.

(Muller and Kubie 1987). The exact circumstances that trigger remapping are not well understood, but it has been extensively shown that this process involves cells changing their place fields and establishing new connections between them. In case of *global* or *complete* remapping, the cell's new place field does not bear any specific relationship with the previous one (Knierim 2003; Latuske et al. 2018).

The existence of remapping, however, raises a puzzle. When the rat needs to orient itself in a familiar environment, action needs to be driven by the information stored in the hippocampus and, when it finds itself in a new environment or when the old environment changes in significant ways, the new information needs to be stored. The problem, however, is that some of the very same brain regions containing place cells (e.g. C3, and CA1) are involved in both processes of encoding and retrieval (Leutgeb et al. 2005). How can a system deal with these two different goals? How can the hippocampus *encode* new information without interference from old memories, and *retrieve* an old map without interference from incoming input? One possible mechanism involves the use of neural oscillations.

The entorhinal cortex (EC) is one of the main extrahippocampal relays delivering new information to the hippocampus (Buhl and Whittington 2007). In encoding a new map, activity in the EC and CA1 is synchronized, meaning that cells in both regions tend to fire in-phase, whereas synaptic transmission between CA3 and CA1 is weak, which prevents interference. In contrast, in retrieval there is strong input from CA3 to CA1 and weak input from EC (Hasselmo, Bodelón, and Wyble 2002; Montgomery and Buzsáki 2007). In a nutshell, the idea is that, in encoding, oscillations are used in order to promote the transmission of information between extrahippocampal areas and the hippocampus and inhibit interference from CA3, while, in retrieval, it is information between CA3 and CA1 that is privileged (Hasselmo and Stern 2014; Colgin 2016). This seems, again, to be a clear example of communication through coherence.

Now, if this hypothesis is on the right track, what role are neural oscillations playing? As in the case of attention reviewed earlier, here synchronization seems to to be the mechanism used for privileging the transmission of certain kinds of information: it facilitates communication between two neuronal assemblies and, at the same time, obstructs possible interferences from other brain areas. Neural oscillations do not seem to provide any new content; they are just the mechanism that opens or closes the channels from two areas that store preexisting information or relay new information.

11

It is important to note that the fact that synchronized oscillation is established between CA1 and EC (CA3) *does* carry information about the system's task currently being one of encoding (retrieval). It is sometimes suggested that the kind of metasemantic account sketched in Section 2 has the unwelcome consequence that most any informational connection will come out representational. Encoding/retrieval in the hippocampus shows this to be false: as no other part of the brain is receiving (or consuming, or decoding) this information, the relevant vehicle does not qualify as a representation. Again here, neural oscillations seem to play a merely enabling role.

4.3   Theta precession in remapping

Neural oscillations can also play a different sort of enabling role: sometimes they are instrumental in building a certain representation, but they are not themselves part of the actual representational vehicle. We offer theta precession as an example of this kind of enabling.

The mammal hippocampus presents a very strong pattern of oscillation in the theta frequency band, around 3–10 Hz (Buzsáki 2002). There appears to be a systematic timing relationship between the activities of single place cells in CA1 or CA3 and this background theta rhythm: O'Keefe and Recce (1993) found that as rats move through the place field corresponding to a certain place cell, the phase of spike trains of this cell tend to change their phase with respect to the background theta period: when the rat enters its place field, a place cell starts firing at the end of the first theta cycle, and subsequent spike trains progressively move forward, as the rat progresses through the field. By the time the rats leaves the place field, bursts might have advanced a whole cycle, i.e. almost 360°, but never more, and most of the time the phase precession spans at most about 180° (O'Keefe and Recce 1993; Maurer et al. 2006; Schmidt et al. 2009). This process is called 'theta precession' (Figure 3).

Among the different proposals on offer about the role of theta precession (Jensen and Lisman 2000; Hasselmo 2005; Huxter, Senior, and Allen 2008; Jeewajee et al. 2014), another one of which we will discuss in Subsection 5.1, we will here focus on the relation between precession and remapping (Skaggs et al. 1996). Spatial representation in the hippocampus is not topological, in the sense that two cells that are close together are not more likely to represent adjacent areas than more distant cells. How can a stable map be formed in such a structure? How can distant neurons come to steadily represent
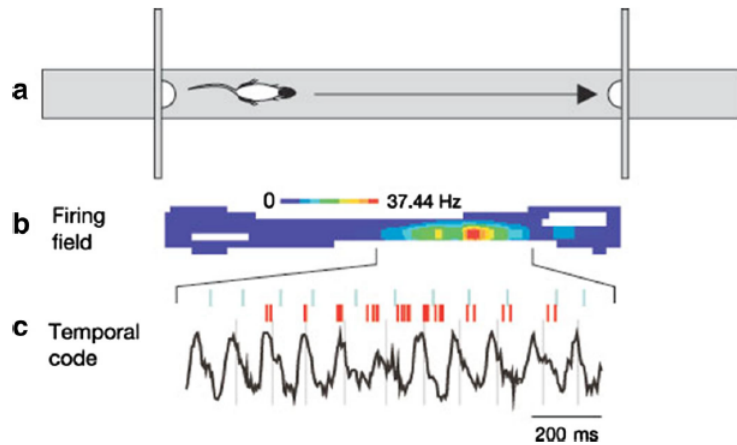
Figure 3: The colored area in b represents the place field of a particular hippocampal place cell when the rat crosses the field represented in a. c shows that place cell spikes (shown in red) precess against the theta oscillations: firings begin near the peak and progressively move earlier in subsequent cycles (from Huxter, Burgess, and O'Keefe 2003).

adjacent locations and how are their connections established? William Skaggs and colleagues suggest that precession plays a key role.

In short, the suggestion runs as follows. When the rat initiates a process of remapping, place cells lose they preferred place fields and gain new fields that bear no predictable relationship to the old ones. Suppose that a rat engages in remapping and in the new environment an A-cell fires when the rat is in location A.[5] Here precession is to be expected: the first spike train will take place near the peak of the last gamma cycle within the first theta cycle,[6] but in subsequent theta cycles the firing pattern will take place at earlier phases as the rat traverses the field. At some point the rat will enter a new place field B and a B-cell (which might be located far away from any A-cells) will become active. Now, due to precession these two cells will fire

_____

5. In this paper we follow the convention of calling the cell that represents, e.g., location A an "A-cell"; *mutatis mutandis* for other cells and their place fields.

6. For more on the relation between theta and gamma cycles see Lisman and Jensen (2013) and Section 5.3 below.

13

in the order AB and, since place fields partly overlap, this pattern will be repeated several times as the rat moves (see figure 4).
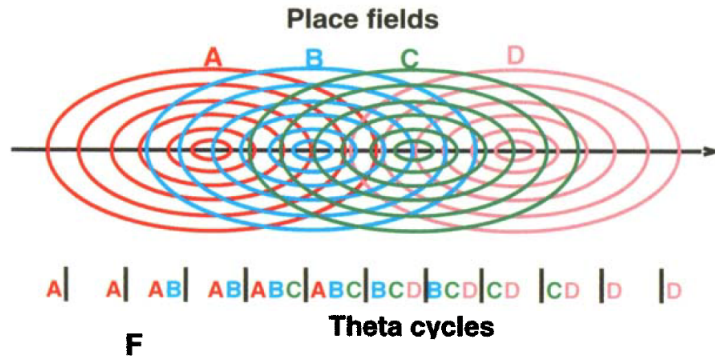

Figure 4: From Skaggs et al. 1996: 169. Explanation in the main text.

More generally, as the rat traverses a sequence of places ABCDE, the following pattern will emerge (vertical lines represent the beginning of a theta cycle): $\dots |ABC|ABC|BCD|BCD|CDE|CDE\dots$. Note that spike trains in the A-cell will systematically take place a bit earlier than spike trains in the B-cell (or vice versa, if the rat is traveling in the contrary direction). This short time span between the spikes of two cells that represent adjacent locations A and B is crucial because it will help strengthen the connections between A- and B-cells through long-term potentiation (LTP).[7] It has been suggested that LTP reinforces better the synaptic connections from A-cell to B-cell when the A-cell fires slightly earlier in time than the B-cell (Larson and Lynch 1989). If this is true, precession might facilitate LTP between neurons that carry information about adjacent places and this might explain how maps can be formed in which cells located at a (relatively) long distance from each other represent nearby places (Skaggs et al. 1996; Bechtel 2016).

Let us suppose that this explanation is on the right track. Do neural oscillations play a representational role here? We lean towards a negative answer: neural oscillations should be construed as a mechanism that enables the construction of a map-like representation, but are not part of it. Certainly, the stable time lag between spike bursts of A- and B-cells, afforded by their occupying different phases in the gamma cycle, is sensitive to (carries information

---

7. Long-term potentiation is a persistent strengthening of synapses caused by co-activation patterns (Cooke and Bliss 2006).

14

about) the adjacency of the related place fields, but we suggest that this is not enough to credit the underlying oscillations with a representational role. There are two main reasons for this.

First, LTP occurs whenever there is a particular temporal delay between the activity of two cells, independently of how this delay is produced. In LTP, timing, not phase, is essential: firings of A- and B-cells need to be sufficiently close *in time* and, for example, keeping phase-delay constant, LTP will happen if the oscillation is sufficiently fast, but not if it is too slow. Compare this with the cases of communication through coherence in attention reviewed above: there, persistently rhythmic, coincident activity (and not merely a certain time lag) seemed to play an essential role in upstream neurons entraining downstream activity. The second reason is that, in this mechanism, there does not seem to be any receiver sensitive to oscillations as such. There is no internal downstream mechanism sensing this phase delay and using it go gain information that could be used in some computational process.

5 Representational oscillations

So far we have discussed cases in which neural oscillations play a role in bringing about representational phenomena without being representations themselves, but rather partly constituting the communicative scaffolding necessary for representations to emerge. In this section we present cases in which neural oscillations do seem to play a properly representational role.

5.1 Theta precession in prediction

Apart from its contribution to remapping, theta precession plays a second role that might qualify as genuinely representational. In their seminal paper, O'Keefe and Recce (1993) found that place cells tend to fire more vigorously in earlier phases of the background theta wave, as the rat approaches the center of their place field. In fact, they observed that the phase of the background theta wave at which place-cell activity is maximal appears to correlate much better with the exact location of the rat within the place field than with the time it has spent in it, or its velocity. Furthermore, in normal conditions this precession (i.e., the phenomenon by which the phase of predominant place-cell activity moves towards the beginning of the theta wave)

15

vanishes when the rat is in the running wheel (Czurko et al. 1999; Hirase et al. 1999), which lends support to the idea that the relevant connection is not with its speed or the time it has spent in a place field, but with its location. As a result, some have suggested that theta phase is used to provide a more fine-grained representation of location: whereas activity in a given place cell indicates the broad area in which the rats finds itself (this would be *rate coding*), the degree of precession carries information about the rat's location within the place field (and this would be *phase coding*). Jensen and Lisman (2000) provided further support for this hypothesis, when they used the activity recorded in 38 pyramidal cells to pinpoint the rats' place within a linear track, and showed that by including phase information they could improve the accuracy by 43% (p. 2607). They were able to predict the rat's position within a 2-meter track with a precision of 3 cm.

Other, more recent work interprets the same results in a slightly different way: instead of taking them to support the idea that phase codes for fine-grained location information within one and the same place field, John Lisman and colleagues have claimed that phase coding is actually used to make *predictions* about which place field the rat is moving towards (Lisman and Redish 2009a; Lisman and Jensen 2013). This is, in a sense, a different gloss on the same main idea: you can think of place fields as being comparatively big, and therefore interpret theta phase as coding for location within that field; or you can think of place fields as comparatively smaller, and then think of theta phase as predicting which place field the rate will be visiting next. Indeed, an important feature of the more recent Lisman and colleagues model is that "the 'true place field' . . . is taken to be approximately one-seventh of the apparent place field (the entire field where rate is elevated)" (Lisman and Redish 2009b, p. 1194). Below, though, we will review empirical data that seems to support the prediction interpretation as more than a mere gloss. It is also possible that there be *both* downstream consumers for this phase-coded information that use it as an aid to prediction or as fine-grained information about location (Colgin 2016, p. 245; Maurer and McNaughton 2007, p. 325f). In the remainder of this section we will stick to prediction.

The predictive role relies on the fact that neural oscillations at different frequencies can be nested, with faster oscillations locked to concrete phases of the slower ones. In particular, within a single theta cycle (called a 'sweep'), there can be between 5 and 14 gamma cycles (Lisman and Redish 2009a, p.
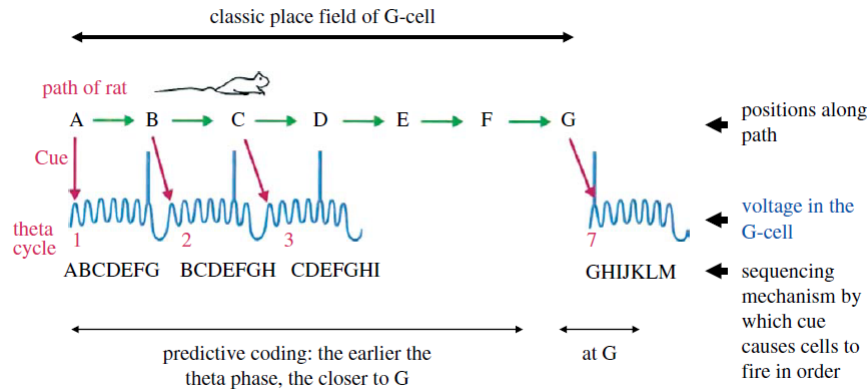
Figure 5: From Lisman and Redish (2009a, p. 1196)

1194).[8] In the current context, this means that a theta sweep accommodates a sequence of 5 to 14 place-cell gamma bursts. Lisman and colleagues' idea is that this sequence is used for prediction: the order in which place cells fire within a single theta cycle corresponds to the order in which the rat expects to visit their place fields.

Let us consider a particular example. Suppose A, B, C, D, E, F and G are the different sections of a path that leads to a certain goal. When the rat is in A, the A-, B-, C-, D-, E-, F- and G-cells will fire within a single theta cycle, in this order. The A-cell firing early in the theta cycle represents that the rat is *in* A. In contrast, the G-cell firing at the end of the cycle represents that the rat is *moving towards* G. In other words, the order in which cells fire within a sweep represents the location of their place field with respect to each other. A place cell firing early in the theta cycle represents the actual location, whereas firings in later phases of the cycle predict future positions (Jensen and Lisman 1996).

Different kinds of evidence have been presented in support of this hypothesis. First of all, the postulated predictive role coheres very well with accounts of hippocampal memory (Jensen and Lisman 1996; Lenck-Santini, Fenton, and Muller 2008): the idea being that prediction relies on the hippocampus operating in "recall mode" (Jensen and Lisman 1996). There is also more direct evidence for a predictive role: as the rat familiarizes itself with a certain environment, it should be able to predict its future location

8. "Seven to nine", according to Buzsáki (2010, p. 370).

17

earlier; and this is indeed what is observed: as rats becomes more familiar with an environment, more gamma cycles are nested within every theta cycle—indicating that prediction starts earlier (Jensen and Lisman 1996). There is also evidence that rats make predictions about future locations before choosing a path, by relying on this phase-coded information. Johnson and Redish (2007) showed that, at bifurcation points in a T-maze, theta sweeps go, successively, through the sequence of place fields corresponding to both arms, which suggests that the rat evaluates available possibilities in advance of deciding.

Let us suppose that this phase-coding-as-prediction hypothesis is correct: a G-cell firing at the beginning of the theta cycle represents that the rat is in G, whereas the very same pattern of activity at late stages of the theta cycle would instead represent that the rat is heading towards G. In this case, it seems that oscillations-involving states such as, e.g., *Vehicle* below satisfy all the requirements for qualifying as representational states:

Vehicle: G-cells firing early in the theta cycle.

On the one hand, it is very plausible that *Vehicle* has a sender—that is, an internal mechanism that generates it in response to location-related worldly states of affairs.[9] There are two main current hypotheses about this mechanism: that it relies on two different oscillators, and that it emerges from asymmetric connections among place cells (Maurer and McNaughton 2007). Hybrid models have also been formulated. But our understanding of this mechanism is still in flux and, as Maurer and McNaughton note in their review, it might "turn out to be extraordinary" (p. 332).[10]

Vehicle also has a receiver, as it makes a clear functional difference downstream (again, assuming that the prediction hypothesis is correct): distinguishing a representation of the actual location from a representation of a future location the rat will be in. This is made most vivid in the role they

---

9. Or perhaps we should think of this case as involving two senders: one in charge of the actual place cell that gets activated, and another in charge of the phase in which this activation happens. As far as we know, our current understanding of the relevant mechanistic details does not allow us to adjudicate this question.

10. The asymmetric-connectivity model is related to the facilitation of long-term potentiation discussed in Subsection 4.3. As far as we are aware, the particular details about how the look-ahead role can be made compatible with the remapping role are still unknown. In any event, the provisional consensus appears to be that both roles *are* compatible (*cf.* Colgin 2016).

seem to play in navigation-related decision making (see above.) We conclude that here neural oscillations appear to play a representation-constitutive role, rather than just an enabling role. What is their representational content? The most plausible reading, as we have argued, is that a late theta phase modifies a place-cell representation: G-cell firing in early theta-phases represents that the rat *already* is in G, whereas G-cell firing in late theta-phases represents that the rat is moving toward G.

It could be objected that, in fact, it is just *G-cell firings* that are properly representational, with the background theta oscillation merely providing a syntactic scaffolding. After all, it is only place cell firings, not oscillations, that carry the relevant information about present and future locations. But theta phase does not have a merely syntactical role. Early (late) phases of the theta wave have something akin to predicative content:

$$\lambda x\{\text{I am currently at (moving towards) } x\}$$

Another possible deflationary understanding of the role theta phase plays is that it is merely contextual: its contribution to the meaning of *Vehicle* would be analogous to the contribution time makes to an utterance such as 'Whiskers is at G *now*'. While the time at which the utterance is made supplies the reference for "now", this is compatible with refraining from claiming that time is literally part of the representation. But this is not a good model for the role theta phases play: time contributes itself to the meaning of "now", but theta phases are not themselves part of the content—they *stand in for* times, like representations do.

Finally, phase differences are as information-carrying as firing rates: if we want to infer where the rat currently is from hippocampal activity, simply focusing on the rate of activation of place cells will not do. We need to take into account the phase of activation as well.[11]

5.2   Feature Binding

One of the first modern discussions of the role that synchrony plays in brain function is von der Malsburg (1981). Here, von der Malsburg suggests that the "correlation between two cellular signals" should be defined "in terms of synchrony and asynchrony between spike trains" (we are citing from the

---

11. We would like to thank an anonymous reviewer for pressing us on these points.

19

4th edition of this paper, von der Malsburg 1994, p. 110). Von der Malsburg presents his synchronicity-based "correlation theory" of brain function as a way of solving, among others, the problem of detecting specific perceptual patterns—among the very many possible combinations of colors, shapes, movement, etc. that could make up a visual scene. Postulating the existence of a set of specialized units that differentially respond to each such pattern (the variously called *cardinal*, *gnostic*, *pontifical* or *grandmother* cells, see Quiroga 2013), he claims, will not do, as "the number of [such units] required would ... be forbidding" (von der Malsburg 1994, p. 101). This is known as the problem of *combinatorial explosion*: if we are to keep track of n simultaneous features (color, shape, etc.), each of which could take m values (blue, yellow, square, etc.), we would need m^n grandmother cells. Synchronization between cells, on the other hand, can result in dynamically created "synaptic networks", in which individual cells respond to rather simple stimuli, but which collectively function as "complex composite feature detectors" (p. 112).

Another important early theory of brain function which predicts a role for synchrony in feature binding along these lines is Grossberg's *Adaptive Resonance Theory* [ART]. In a recent review of ART, Grossberg states that

> Coherent binding of the attended features to the category give them a meaning as a context-sensitive "event" rather than as just isolated pixels. Such coherent states between distributed features and symbolic categories are often expressed dynamically as synchronously oscillating activations across the bound cells ... (Grossberg 2013, p. 9)

This "coherence between distributed features and symbolic categories" is what we would now call feature binding (of the former features into an object conceptualized under the latter symbolic category).

Many experimental results support these ideas: Gray et al. (1989) famously demonstrated that oscillatory responses were evoked by stimuli which showed coherent motion, but not by stimuli which moved in opposite directions (Gray et al. 1989; Gray et al. 1990, p. 335). That is: synchronous activity marks the presence of coherent motion, which (as Gestalt theorists suggested) is in its turn evidence that the two stimuli in question are not in fact separate entities, but belong to one and the same object. Kreiter and Singer (1996) showed that two cells with overlapping receptive fields, but such that each of them is tuned to a different direction of movement, can be driven to fire

in synchrony by the movement of a single bar, yet not by two bars each one moving in the preferred direction of one cell—again here, synchronization goes with co-boundedness, even trumping the tuning features of individual neurons. Synchronization appears to be responsive to other Gestalt-grouping principles as well, besides coherent motion (Singer and Gray 1995; Engel et al. 1992; Gray 1999). Tallon-Baudry and Bertrand (1999) review much other relevant evidence. Modeling work also supports the link between oscillations and feature binding (Eckhorn et al. 1990) and suggests that synchronization can be sufficiently fast in long-range interactions (Yazdanbakhsh and Grossberg 2004).

The most ambitious contemporary version of this *binding by synchrony* [BBS] hypothesis claims that synchrony is the main code for feature binding. This seems to be the position taken by proponents of the *temporal binding model* (Engel and Singer 2001; Engel, Fries, and Singer 2001, and papers cited therein), who claim that "neural synchrony ... is crucial for object representation" (Engel, Fries, and Singer 2001, p. 706). There are good reasons (both theoretical (Shadlen and Movshon 1999; Di Lollo 2012)[12] and empirical (Palanca and DeAngelis 2005)) to think that synchrony is not the only feature-binding code. Still, the results discussed above and many others (including ones by Palanca and DeAngelis in the critical paper just cited) suggest that it does make a contribution to representing the fact that different features are co-bound to the same object (*cf.* Hommel 2004, Box 1).

Let us assume that this putative contribution takes the following form: "the mechanism that evaluates temporal relations among responses for perceptual grouping interprets synchronous responses as related and segregates them from responses that are temporally offset." (Singer 1999, p. 51) where, as we have just discussed, "interprets" should not be taken to mean that synchronicity *determines* the status of responses as related or unrelated, but rather that it provides evidence for it, possibly to be combined with other mechanisms such as, e.g., task-dependent alterations in neural tuning (Gilbert and Li 2013, p. 5). Should we interpret this less ambitious version of BBS as vindicating a representational role for synchrony?

In a nutshell, the hypothesis is that synchronous activity between two populations that code for two different perceptual features would represent

---

12. Di Lollo is sceptical about the very existence of a feature-binding problem. He defends that something like gnostic cells are actually available in the brain, but doesn't explicitly discuss how combinatorial explosion is therefore dealt with.

that these features are co-bound. Again, we unpack this idea by trying to find occupants for the different roles in our metasemantic sketch. First, there is a plausible occupant for the role of representational vehicle:

Vehicle: A population of neurons P_1 (that encodes a perceptual feature F_1) and another population P_2, (that encodes a perceptual feature F_2) firing in synchrony.[13]

Now, is there a plausible sender/encoder for *Vehicle*? What this encoder has to do is to subsume populations P_1 and P_2 into a larger coherent assembly, as a reaction to the fact that onset and offset of activity in P_1 and P_2 is more or less simultaneous (which in turn is explained by features F_1 and F_2 being actually co-bound in the world, and therefore appearing and disappearing more or less simultaneously).

There are at least two goals that this encoder must meet. First, note that the fact that F_1 and F_2 are reliably co-instantiated is not the same as, nor sufficient for, synchrony. Co-instantiation is indeed sufficient for *sameness of onset and end of firing*, but not for the fact that, during the duration of the firing episode, this firing is synchronized, which it is, "over periods ranging from tens to thousands of milliseconds" (Gray 1999, p. 38), moreover showing sophisticated oscillatory structure (Singer and Gray 1995, p. 1094). Second, this synchronic firing must start quickly and reliably after stimulus onset if it is to be an ecologically viable way of signaling feature boundedness. This job description is far from computationally trivial, and it is met by a rather specific pattern of lateral interconnections in the relevant neuronal population, with the right mix of excitation and inhibition (Fries 2015; Yazdanbakhsh and Grossberg 2004, see Section 4.1 above). This kind of network architecture is a good candidate for our sender/encoder.

There is also a plausible receiver for *Vehicle*: the whole point of binding by synchrony is that dynamic assemblies formed by synchronized neurons are treated as a unit. One often proposed mechanism in this connection is that coincidence-sensitive neurons (Engel and Singer 2001, p. 18; Abeles 1982; König, Engel, and Singer 1996) would be specially driven by neurons firing synchronously (see Fries 2015, on *effective connectivity*). We will discuss

---

13. Shouldn't the vehicle be just the *coinstantiation* of P_1 and P_2 firings, rather than full-blown synchrony? Not according to the defenders of the BBS hypothesis: synchrony (i.e., rhythmic, congruent activity), and not mere coinstantiation, is necessary to entrain postsynaptic activity in a sufficiently vigorous way (Fries 2015; Engel, Fries, and Singer 2001, p. 705).

"reader-classifiers" of these neuronal assemblies in more detail in Subsection 5.3.

It seems, thus, that in the case of BBS all of the links in our metasemantic sketch have a plausible occupant. We therefore tentatively conclude that *Vehicle* is a representation. Note that here synchrony is not merely an enabling mechanism: synchronized oscillations are a constitutive part of the representational vehicle. Moreover, and more importantly, these same oscillations are directly causally involved in the decoding behavior. That is, it is by virtue of its oscillatory profile that *Vehicle* means what it means.

5.3 Sequences

An impressive body of work associated mainly (but not only) with György Buzsáki's lab (Buzsáki 2006; 2010; Buzsáki and Watson 2012, among many others) has aimed at uncovering "syntactical" units in brain processing: how does the brain go about providing vehicles over which computations can be performed? As we will see, neural oscillations play an essential role in the construction of these various syntactical units. It will turn out, though, that there is some, perhaps ineliminable, indeterminacy between taking some cases of neural activity as constituting one such unit (e.g., a "neural word") or as a process of computation of one unit from another. In our framework, this will translate to indeterminacy in their status as properly representational.

Constructing these vehicles is a difficult task, if only because it involves arbitrating between two desiderata that pull in different directions. First, the repertoire of available vehicles needs to be *sensitive* enough: it should be possible to token, reliably, different vehicles in the presence of different, but similar, external conditions. Brains need to keep track of fine-grained differences in external events (say, in the velocity of looming or receding objects, Maier et al. 2004) and this requires provisioning a sufficient number of different vehicles to stand in for each member in a sufficiently fine-grained partition of such events (say, for a sufficiently fine-grained range of different looming velocities). Note that this sensitivity requirement is not met by simply generating states that count as different under some third-personal criterion: it must be possible for these different states to make a difference to the system itself—they must be possibly *treated as different* downstream.

Second, this repertoire needs to be *robust* enough: the vehicles in question need to be somehow resistant to the very noisy environment in which they are tokened (see Faisal, Selen, and Wolpert 2008, for a review of noise in the

brain). Sensitivity and robustness are in tension: sensitive processing involves tokening noticeably different vehicles in response to very similar world states; yet, robust processing involves preventing ambient noise from conflating the processing trajectories of the vehicles that are triggered by those world states. As we are about to see, it has been suggested that neural oscillations play an important role in solving this conundrum.

Buzsáki and colleagues have claimed that the fundamental "syntactical" units in brain processing (that is to say, the fundamental unit in the construction of vehicles) is not the single neuron but the *cell assembly*: "a collection of neurons that come together . . . to produce a composite downstream effect that cannot be produced by single neurons alone" (Buzsáki 2010, p. 364, recall that von der Malsburg uses a very similar expression). Because the existence of a cell assembly depends on the existence of these "composite downstream effects", there must be a downstream "reader-classifier" that treats the assembly as a functional unit (*ibid.*). According to Buzsáki, the most basic assembly reader-classifier in the brain is the *integration of presynaptic activity*: the process whereby a certain neuron treats presynaptic events (say, action potentials coming from different presynaptic neurons at slightly different times) as a unit. How far apart from each other these events can be and still be treated as a single whole depends on the so-called *membrane time constant* τ (*ibid.*) but, in general, the closer in time those events are, the most likely it will be that they will be treated as a unit. Obviously, a reader-classifier such as this one, that mainly detect coincidences in presynaptic activity, will greatly benefit from synchronized activity upstream. This is the first place where neural oscillations play a role in the generation of vehicles: locking presynaptic firings to concrete phases of an oscillation cycle leads to postsynaptic neurons treating each such phase-locked collection of firings as a functional unit.

These cell assemblies should be thought of as the phonemes (or perhaps letters) of the neural syntax. One step up in the Buzsákian hierarchy of vehicles we find *neural words*, made up of "sequences" of these assemblies, quickly following one another through the duration of a (typically gamma) cycle (*ibid.*, p. 365). E.g., we can think of a neural word as assemblies A, B, C, D succeeding one another through a gamma cycle (each, say, occupying 90 degrees of the full 360 degrees in a cycle). The sequence consisting of assemblies A, C, D, B would constitute a different word, possibly discharging an entirely different processing role. Again, neural words will only be real insofar as some entity downstream treats them as a unit: for example, mechanisms

24

with a longer integration window than the membrane time constant, such as NMDA or GABAe receptors (*ibid.*, p. 366). A more important, and more general, class of reader-classifiers with different integration windows is again provided by neural oscillations (*ibid.*). From the point of view of electronics, neural oscillations are *relaxation oscillators* (Wang 1999): this kind of circuit (which is used, for example, in the blinking turn signals of cars) works by periodically charging and discharging a capacitor (a storage of electrical energy); when the capacitor reaches a certain threshold, the oscillator is "reset". The concept of neural word relies on two features of relaxation oscillators: first, the phase during which the capacitor is being charged naturally corresponds to an integrator window (Buzsáki 2010, p. 366)—recall that this is just a period during which incoming neural activity is treated as an undiferentiated whole. Second the oscillator "reset" acts as a natural gap between different neural words.

There is ample empirical evidence of the existence of neural words, understood along those lines. For example, Jones et al. (2007) show how neurons in the gustatory cortex of rats go through four different sequences, each comprised of four different states, whenever the rat is exposed to sweet, bitter, sour or salty flavors, respectively. The four states that are part of the sequences are different for different flavors, but always the same within each flavor (see figure 6). Laurent (2002, p. 886) similarly reports population-level representations in the antennal lobe of insects and the olfactory bulb of mammals that are "dynamic, carried by an assembly of neurons . . . that evolves in a stimulus-specific manner over time". In sections 4.3 and 5.1 we have reviewed in detail theta sweeps in the hippocampus, which provides yet another important example of assembly sequence.

Further syntactic structure is provided by nested oscillatory rhythms: for example, the number of assemblies that fit in a theta cycle has consequences for the "memory 'buffer' of the gamma-nested theta-cycle" (*ibid.*). More generally, "[theta-gamma] oscillations form a code for representing multiple items in an ordered way." (Lisman and Jensen 2013, p. 1002)

Now, how do neural vehicles constructed out of these syntactic building blocks trade off sensitivity and robustness? First of all, reader-classifiers will help with robustness by *failing* to distinguish between different sequences of events (say, different sets of arrival times of presynactic activities) as long as they fall within the same activation window—e.g., *neuron 1 firing now and neurons 2 and 3 firing in the next 5ms* will be indistinguishable from *neuron*
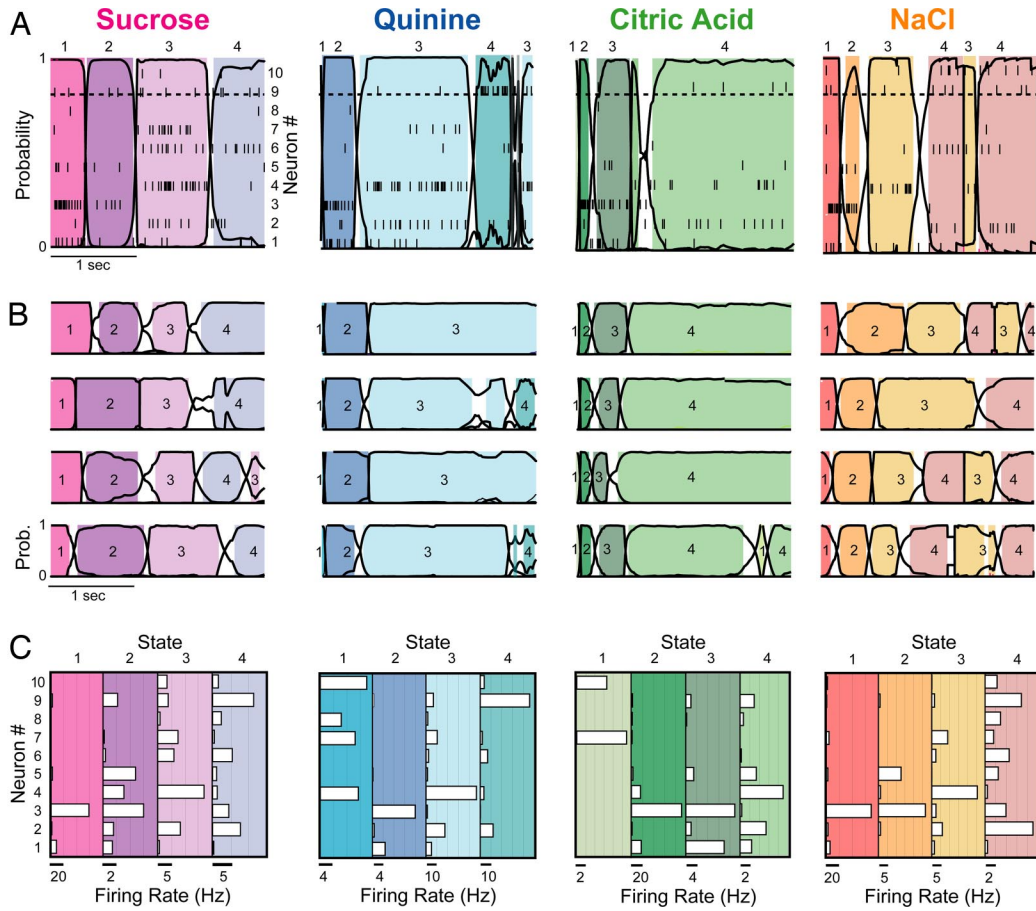
Figure 6: Neural words in the gustatory cortex. From Jones et al. (2007)

3 firing now and neurons 1 and 2 firing in the next 5ms.[14] As for sensitivity, the production of sequences might be such that initially minute differences in input are taken, as more evidence comes in, into rapidly divergent trajectories corresponding to different sequences that are much easier to distinguish: those initially minute differences, for example, might result in two different sequences with an overlapping initial part, as in the ABCD / ACBD example above (Harvey, Coen, and Tank 2012).

The most straightforward way to think of these neural words is as representational vehicles: they are hypothesized to solve a problem (the sensitivity / robustness tradeoff) that is precisely the kind of problem a representational system in a noisy environment would face. Furthermore, the empirical evidence for neural words links them in every case to concrete representational roles (the representation of flavors in Jones et al. (2007), odorants in Laurent (2002), or paths to be taken in space in Jensen and Lisman (1996)).

Still, closer examination of the available evidence, and the attitude of researchers toward that evidence, reveals that this representational reading is not without problems. First, note that the mechanisms through which, we suggested, sensitivity and robustness are accommodated by neural words are somewhat at odds with one another: robustness depended on a reader-classifier with a large enough activation window that it may be able to take the full neural word in as a unified whole; while sensitivity depended on neural words being interpreted as diverging *trajectories*—crucially, their diachronic unfolding, and their responsiveness to incoming evidence that coheres with the trajectory chosen, is part of what makes them robust. Indeed, many researchers are skeptical that sequences have readers of their own. So, for example, according to Lauren Jones and colleagues, "[c]oherent state sequences . . . *probably do not represent pure 'sensory codes' to be interpreted by downstream 'grandmother neurons'*" Jones et al. (2007, p. 18776, emphasis added). According to these researchers, the sequences we see in sensory cortical ensembles should be interpreted, not as pure codes, but as computations of motor codes from purely sensory ones.

Laurent (2002), on the other hand, claims that sequences of assemblies do play a role in generating a "*large coding space* in which to spread representation clusters" (*ibid.*, p. 885f, emphasis added). While this is one of the main points of his paper, highlighted in abstract and conclusions, Laurent

---

14. This is just a straightforward example of what information theorists call *channel coding*, or error correction (MacKay 2003).

also warns that, conceivably, "slow temporal patterns, although crucial for the separation of representations, *are never actually decoded as such.* More generally, the creation of spatiotemporal representations by circuit dynamics might be a transient phase in signal processing, used simply to spread out those representations in a larger coding space and to facilitate decoding" (*ibid.*, p. 891, emphasis added). That is, it is not clear that the olfaction-related neural sequences in insects and mammals that Laurent is interested in have receivers—i.e., are pure sensory codes. Rather, they would amount to *encoding stages* along which representational vehicles are progressively optimized.

This hybrid (part code, part encoding process) nature of spatio-temporal sequences is fleshed out most clearly by Harvey, Coen, and Tank (2012). They show that an ideal observer is able to predict behavioral choices in mice from spatio-temporal neural sequences but not from the synchronic behavior of neuronal populations (see also Yuste 2015, p. 492). The model proposed by Harvey and colleagues is one in which incoming sensory information initiates a (firstly decision-agnostic) sequence which gets progressively decision-specific as incoming information modifies it. When the sequence overlaps sufficiently with a decision-specific trajectory, this corresponds to the personal-level state of the mouse having decided. The subsequent, endogenously generated unfolding of the decision-specific sequence corresponds to the personal-level state of keeping the decision in mind.

The upshot of this discussion is that, while very prominent models of brain function accord an important role to neural oscillations in the generation of representational vehicles, there is some vacillation in the literature regarding what counts as code (of incoming sensory information); what counts as computation (of decisions from sensory information); and what counts as encoding processes (whereby sensory codes gets optimized into decision codes). This indeterminacy between merely enabling and properly representational status for neural-sequence-related oscillations will perhaps be remedied as our knowledge of brain processing improves; perhaps it is ineliminable.

6 Conclusion

The main goal of this paper was to examine whether neural oscillations in the brain actually are representations. As we have seen, whether they do often depends on rather intricate facts about the relevant mechanisms where

28

they participate, and the role they play in these mechanisms. We offer this to philosophers as a cautionary tale: representational status, like the devil, is often in the detail.

But we can also glean a few big-picture lessons from the foregoing discussion. First, our results shows that the core metaphysical theory of content put forward in Section 2 is a useful tool for understanding and modelling neural representations. It suggests important questions, and provides theoretical instruments with which to answer them.

Secondly, the discussion of case studies revealed that neural oscillations can play at least two different roles, enabling and properly representational, in the implementation of representations in the brain. This distinction is a result of the careful consideration of specific cases, rather than a pre-conceived analysis. We hypothesize that this distinction will also apply to the mechanisms subserving other representational phenomena; and that it can potentially be extended by considering other roles that cognitive processes can play that are not representational, but are defined in relation to properly representational states.

Thirdly, following a recent call for multi-level analysis (Craver 2007; Boone and Piccinini 2016), our results show that paying attention to implementational details is relevant for understanding higher-order levels. Even though 'representation' is a functional category, considering which specific brain structure actually play a representational role can help us vindicate a particular analysis of the nature of representation and distinguishes different non-representational roles brain states can play, among others.

Finally, against recent suggestions to the contrary, we argued that low-level processes can qualify as representational. Furthermore, this result has not been obtained by defining representation so cheaply that just anything can count as such, since we have identified some processes employing oscillations in which they fail to play a representational role (e.g. attention, retrieval and encoding). The results of this paper support the idea that representations can be found all the way down without trivializing this notion. This provides some vindication for naturalistic theories of representation.[15]

---

15. For discussion of the objection that naturalistic theories of content are too liberal, see Artiga (2016, 2020).

References

Abeles, M. 1982. "Role of the cortical neuron: integrator or coincidence detector?" *Israel J. Med. Sci.* 18:83—92.

Alilović, Josipa, Bart Timmermans, Leon C. Reteig, Simon van Gaal, and Heleen A. Slagter. 2019. "No Evidence That Predictions and Attention Modulate the First Feedforward Sweep of Cortical Information Processing" [in en]. *Cerebral Cortex,* https://doi.org/10.1093/cercor/bhz038.

Artiga, Marc. 2016. "Liberal Representationalism: A Deflationist Defense." *dialectica* 70 (3): 407–430.

———. 2020. "Signals Are Minimal Causes" [in en]. *Synthese* (February). ISSN: 1573-0964. https://doi.org/10.1007/s11229-020-02589-0.

———. 2013. "Teleosemantics and Pushmi-Pullyu Representations." *Erkenntnis,* 1–22.

Baluch, Farhan, and Laurent Itti. 2011. "Mechanisms of Top-down Attention" [in en]. *Trends in Neurosciences* 34, no. 4 (April): 210–224. ISSN: 0166-2236. https://doi.org/10.1016/j.tins.2011.02.003.

Bechtel, W. 2016. "Investigating Neural Representations: The Tale of Place Cells." *Synthese.* 193 (5): 1287–1321.

Boone, W., and G. Piccinini. 2016. "The cognitive neuroscience revolution." *Synthese.* 193 (5): 1509–1534.

Börgers, Christoph, Steven Epstein, and Nancy J. Kopell. 2008. "Gamma Oscillations Mediate Stimulus Competition and Attentional Selection in a Cortical Network Model" [in en]. *Proceedings of the National Academy of Sciences* 105, no. 46 (November): 18023–18028. ISSN: 0027-8424, 1091-6490. https://doi.org/10.1073/pnas.0809511105.

Brefczynski, Julie A., and Edgar A. DeYoe. 1999. "A Physiological Correlate of the'spotlight'of Visual Attention." *Nature neuroscience* 2 (4): 370.

Buhl, E., and M. Whittington. 2007. "Local Circuits." In *The Hippocampus Book,* edited by P. Andersen, R. Morris, D. Amaral, T. Bliss, and J. O'Keefe, 297–320. Oxford University Press.

Burge, Tyler. 2010. *Origins of Objectivity.* Oxford University Press.

Butlin, Patrick. 2018. "Representation and the Active Consumer" [in en]. *Synthese* (September). ISSN: 1573-0964. https://doi.org/10.1007/s11229-018-01941-9.

Buzsáki, György. 2010. "Neural Syntax: Cell Assemblies, Synapsembles, and Readers." *Neuron* 68, no. 3 (November): 362–385. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2010.09.023.

————. 2006. *Rhythms of the Brain.* Oxford University Press.

————. 2002. "Theta Oscillations in the Hippocampus" [in en]. *Neuron* 33, no. 3 (January): 325–340. ISSN: 0896-6273. https://doi.org/10.1016/S0896-6273(02)00586-X.

Buzsáki, György, and Brendon O. Watson. 2012. "Brain Rhythms and Neural Syntax: Implications for Efficient Coding of Cognitive Content and Neuropsychiatric Disease." *Dialogues in Clinical Neuroscience* 14, no. 4 (December): 345–367. ISSN: 1294-8322.

Chalk, Matthew, Jose L. Herrero, Mark A. Gieselmann, Louise S. Delicato, Sascha Gotthardt, and Alexander Thiele. 2010. "Attention Reduces Stimulus-Driven Gamma Frequency Oscillations and Spike Field Coherence in V1." *Neuron* 66, no. 1 (April): 114–125. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2010.03.013.

Colgin, L. L. 2016. "Rhythms of the hippocampal network." *Hippocampus* 17 (4): 239–249.

Colgin, L.L., D. Kubota, Y. Jia, C.S. Rex, and G. Lynch. 2004. "Long-term potentiation is impaired in rat hippocampal slices that produce spontaneous sharp waves." *J Physiol.* 558 (03): 953–61.

Cooke, S.F., and T. P. Bliss. 2006. "Plasticity in the human central nervous system." *Brain* 129:1659—1673.

Craver, C. 2007. *Explaining the Brain.* OUP.

Cummins, Robert. 1991. *Meaning and Mental Representation.* The MIT Press.

Czurko, A., H. Hirase, J. Csicsvari, and G. Buzsaki. 1999. "Sustained activation of hippocampal pyramidal cells by 'space clamping' in a running wheel." *Eur. J. Neurosci.* 11:344—352.

Di Lollo, Vincent. 2012. "The Feature-Binding Problem Is an Ill-Posed Problem." *Trends in Cognitive Sciences* 16, no. 6 (June): 317–321. ISSN: 1364-6613. https://doi.org/10.1016/j.tics.2012.04.007.

Dumoulin, Serge O., and Brian A. Wandell. 2008. "Population Receptive Field Estimates in Human Visual Cortex." *Neuroimage* 39 (2): 647–660.

Eckhorn, Reinhard, Herbert J. Reitboeck, Mt Arndt, and P. Dicke. 1990. "Feature Linking via Synchronization among Distributed Assemblies: Simulations of Results from Cat Visual Cortex." *Neural computation* 2 (3): 293–307.

El Gamal, Abbas, and Young-Han Kim. 2011. *Network Information Theory.* Cambridge university press.

Engel, Andreas K., Pascal Fries, and Wolf Singer. 2001. "Dynamic Predictions: Oscillations and Synchrony in Top–down Processing" [in En]. *Nature Reviews Neuroscience* 2, no. 10 (October): 704. ISSN: 1471-0048. https://doi.org/10.1038/35094565.

Engel, Andreas K., Peter König, Andreas K. Kreiter, Thomas B. Schillen, and Wolf Singer. 1992. "Temporal Coding in the Visual Cortex: New Vistas on Integration in the Nervous System." *Trends in Neurosciences* 15, no. 6 (June): 218–226. ISSN: 0166-2236. https://doi.org/10.1016/0166-2236(92)90039-B.

Engel, Andreas K., and Wolf Singer. 2001. "Temporal Binding and the Neural Correlates of Sensory Awareness." *Trends in Cognitive Sciences* 5, no. 1 (January): 16–25. ISSN: 1364-6613. https://doi.org/10.1016/S1364-6613(00)01568-0.

Faisal, A. Aldo, Luc P. J. Selen, and Daniel M. Wolpert. 2008. "Noise in the Nervous System" [in en]. *Nature Reviews Neuroscience* 9, no. 4 (April): 292–303. ISSN: 1471-0048. https://doi.org/10.1038/nrn2258.

Fell, Juergen, Guillén Fernández, Peter Klaver, Christian E. Elger, and Pascal Fries. 2003. "Is Synchronized Neuronal Gamma Activity Relevant for Selective Attention?" *Brain Research Reviews* 42, no. 3 (June): 265–272. ISSN: 0165-0173. https://doi.org/10.1016/S0165-0173(03)00178-4.

Fodor, J. 1974. "Special sciences." *Synthese* 28:77–115.

Frankish, K., and W. Ramsey. 2012. *The Ccambridge Handbook of Cognitive Science.* Cambridge University Press.

Fries, Pascal. 2005. "A Mechanism for Cognitive Dynamics: Neuronal Communication through Neuronal Coherence." *Trends in Cognitive Sciences* 9, no. 10 (October): 474–480. ISSN: 1364-6613. https://doi.org/10.1016/j.tics.2005.08.011.

———. 2009. "Neuronal Gamma-Band Synchronization as a Fundamental Process in Cortical Computation." *Annual Review of Neuroscience* 32 (1): 209–224. https://doi.org/10.1146/annurev.neuro.051508.135603.

———. 2015. "Rhythms for Cognition: Communication through Coherence." *Neuron* 88, no. 1 (October): 220–235. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2015.09.034.

Friston, Karl J. 2011. "Functional and Effective Connectivity: A Review." *Brain Connectivity* 1, no. 1 (January): 13–36. ISSN: 2158-0014. https://doi.org/10.1089/brain.2011.0008.

Gilbert, Charles D., and Wu Li. 2013. "Top-down Influences on Visual Processing" [in en]. *Nature Reviews Neuroscience* 14, no. 5 (May): 350–363. ISSN: 1471-0048. https://doi.org/10.1038/nrn3476.

Godfrey-Smith, Peter. 2014. "Sender-Receiver Systems Within and Between Organisms." *Philosophy of Science* 81:866–878.

Gray, Charles M. 1999. "The Temporal Correlation Hypothesis of Visual Feature Integration: Still Alive and Well" [in English]. *Neuron* 24, no. 1 (September): 31–47. ISSN: 0896-6273. https://doi.org/10.1016/S0896-6273(00)80820-X.

Gray, Charles M., Andreas K. Engel, Peter König, and Wolf Singer. 1990. "Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Receptive Field Properties and Feature Dependence." *European Journal of Neuroscience* 2 (7): 607–619.

Gray, Charles M., Peter König, Andreas K. Engel, and Wolf Singer. 1989. "Oscillatory Responses in Cat Visual Cortex Exhibit Inter-Columnar Synchronization Which Reflects Global Stimulus Properties." *Nature* 338 (6213): 334.

Gregoriou, Georgia G., Stephen J. Gotts, Huihui Zhou, and Robert Desimone. 2009. "High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention" [in en]. *Science* 324, no. 5931 (May): 1207–1210. ISSN: 0036-8075, 1095-9203. https://doi.org/10.1126/science.1171402.

Grossberg, Stephen. 2013. "Adaptive Resonance Theory: How a Brain Learns to Consciously Attend, Learn, and Recognize a Changing World." *Neural Networks,* Twenty-Fifth Anniversay Commemorative Issue, 37 (January): 1–47. ISSN: 0893-6080. https://doi.org/10.1016/j.neunet.2012.09.017.

Głądziejewski, P., and M. Miłkowski. 2017. "Structural Representations: Causally Relevant and Different From Detectors." *Biology and Philosophy* 32 (3): 337–355.

Harris, Karen R., Barbara Danoff Friedlander, Bruce Saddler, Remedios Frizzelle, and Steve Graham. 2005. "Self-Monitoring of Attention Versus Self-Monitoring of Academic Performance: Effects Among Students with ADHD in the General Education Classroom" [in en]. *The Journal of Special Education* 39, no. 3 (November): 145–157. ISSN: 0022-4669. https://doi.org/10.1177/00224669050390030201.

Hartley, T., C. Lever, N. Burgess, and J. O'Keefe. 2017. "Space in the brain: how the hippocampal formation supports spatial cognition." *Phil. Trans. R. Soc. B* 369:663—678.

Harvey, Christopher D., Philip Coen, and David W. Tank. 2012. "Choice-Specific Sequences in Parietal Cortex during a Virtual-Navigation Decision Task" [in en]. *Nature* 484, no. 7392 (April): 62–68. ISSN: 1476-4687. https://doi.org/10.1038/nature10918.

Hasselmo, M. E. 2005. "What is the Function of Hippocampal Theta Rhythm?—Linking Behavioral Data to Phasic Properties of FieldPotential and Unit Recording Data ." *Hippocampus* 15:936—949.

Hasselmo, M. E., and Ch. E. Stern. 2014. "Theta rhythm and the encoding and retrieval of space and time." *Nat Neurosci.* 82 (02): 656—666.

Hasselmo, M., C. Bodelón, and B. P. Wyble. 2002. "A Proposed Function for Hippocampal Theta Rhythm: Separate Phases of Encoding and Retrieval Enhance Reversal of Prior Learning." *Neural Computation* 14 (4): 793–817.

Hirase, H., A. Czurko, J. Csicsvari, and G. Buzsaki. 1999. "Firing rate and theta-phase coding by hippocampal pyramidal neurons during 'space clamping." *Eur. J. Neurosci.* 11:4373—4380.

Hommel, Bernhard. 2004. "Event Files: Feature Binding in and across Perception and Action." *Trends in Cognitive Sciences* 8, no. 11 (November): 494–500. ISSN: 1364-6613. https://doi.org/10.1016/j.tics.2004.08.007.

Hurlbert, A., and T. Poggio. 1985. "Spotlight on Attention." *Trends in Neurosciences* 8 (January): 309–311. ISSN: 0166-2236. https://doi.org/10.1016/0166-2236(85)90110-9.

Huxter, J., N. Burgess, and J. O'Keefe. 2003. "Independent rate and temporal coding in hippocampal pyramidal cells." *Nature* 425 (6960): 828–32.

Huxter, J.R., T.J. Senior, and J. Allen K. Csicsvari. 2008. "Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus." *Nat Neurosci.* 11 (5): 587–94.

Jeewajee, A., C. Barry, V. Douchamps, D. Manson, C. Lever, and N. Burgess. 2014. "Theta phase precession of grid and place cell firing in open environments." *Philos Trans R Soc Lond B Biol Sci.* 396 (1635): 20120532. https://doi.org/doi:10.1098/rstb.2012.0532.

Jensen, O., and J. E. Lisman. 1996. "Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells." *Learn. Mem* 3:270–287.

————. 2000. "Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding." *Journal of Neurphysiology* 83 (5): 2602–9.

Jensen, Ole, Jochen Kaiser, and Jean-Philippe Lachaux. 2007. "Human Gamma-Frequency Oscillations Associated with Attention and Memory." *Trends in Neurosciences,* July INMED/TINS Special Issue—Physiogenic and Pathogenic Oscillations: The Beauty and the Beast, 30, no. 7 (July): 317–324. ISSN: 0166-2236. https://doi.org/10.1016/j.tins.2007.05.001.

Johnson, A., and A. D. Redish. 2007. "Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point." *J Neurosci.* 27 (45): 12176–89.

Jones, L. M., A. Fontanini, B. F. Sadacca, P. Miller, and D. B. Katz. 2007. "Natural Stimuli Evoke Dynamic Sequences of States in Sensory Cortical Ensembles" [in en]. *Proceedings of the National Academy of Sciences* 104, no. 47 (November): 18772–18777. ISSN: 0027-8424, 1091-6490. https://doi.org/10.1073/pnas.0705546104.

Knierim, J. 2003. "Hippocampal remapping: implications for spatial learning and navigation." In *The Neurobiology of Spatial Behaviour,* edited by K. Jeffery, 226–240. Oxford University Press.

Kreiter, A. K., and W. Singer. 1996. "Stimulus-Dependent Synchronization of Neuronal Responses in the Visual Cortex of the Awake Macaque Monkey" [in en]. *Journal of Neuroscience* 16, no. 7 (April): 2381–2396. ISSN: 0270-6474, 1529-2401. https://doi.org/10.1523/JNEUROSCI.16-07-02381.1996.

König, P., A.K. Engel, and W. Singer. 1996. "Integrator or coincidence detector? The role of the cortical neuron revisited." *Trends Neurosci.* 19:130–137.

Larson, J., and G. Lynch. 1989. "Theta pattern stimulation and the induction of LTP: the sequence in which synapses are stimulated determines the degree to which they potentiate." *Brain Res.* 489 (1): 49–58.

Latuske, P., O. Kornienko, L. Kohler, and K. Allen. 2018. "Hippocampal Remapping and Its Entorhinal Origin." *Frontiers in Behavioural Neuroscience.*

Laurent, Gilles. 2002. "Olfactory Network Dynamics and the Coding of Multidimensional Signals" [in en]. *Nature Reviews Neuroscience* 3, no. 11 (November): 884–895. ISSN: 1471-003X, 1471-0048. https://doi.org/10.1038/nrn964.

Lenck-Santini, P., A. A. Fenton, and R. Muller. 2008. "Discharge Properties of Hippocampal Neurons during Performance of a Jump Avoidance Task." *J Neurosci.* 28 (27): 6773—6786.

Leutgeb, J., S. Leutgeb, A. Treves, R. Meyer, C. Barnes, B. McNaughton, M. Moser, and E. Moser. 2005. "Progressive Transformation of Hippocampal Neuronal Representations in "Morphed" Environments." *Neuron* 48:345–358.

Lisman, J. E., and A. D. Redish. 2009a. "Prediction, sequences and the hippocampus." *Phil. Trans. R. Soc. B* 364:1193—1201.

Lisman, John E., and Ole Jensen. 2013. "The Theta-Gamma Neural Code." *Neuron* 77, no. 6 (March): 1002–1016. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2013.03.007.

Lisman, John, and A.D. Redish. 2009b. "Prediction, Sequences and the Hippocampus." *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, no. 1521 (May): 1193–1201. ISSN: 0962-8436. https://doi.org/10.1098/rstb.2008.0316.

MacKay, David JC. 2003. *Information Theory, Inference and Learning Algorithms.* Cambridge university press.

Maier, Joost X, John G Neuhoff, Nikos K Logothetis, and Asif A Ghazanfar. 2004. "Multisensory Integration of Looming Signals by Rhesus Monkeys" [in en]. *Neuron* 43, no. 2 (July): 177–181. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2004.06.027.

Martínez, Manolo. 2019. "Representations Are Rate-Distortion Sweet Spots." *Philosophy of Science* 86, no. 5 (December): 1214–1226. ISSN: 0031-8248. https://doi.org/10.1086/705493.

Maunsell, John H. R., and Erik P. Cook. 2002. "The Role of Attention in Visual Processing" [in en], edited by Andrew Parker, Andrew Derrington, and Colin Blakemore. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 357, no. 1424 (August): 1063–1072. ISSN: 1471-2970. https://doi.org/10.1098/rstb.2002.1107.

Maunsell, John H. R., and Stefan Treue. 2006. "Feature-Based Attention in Visual Cortex." *Trends in Neurosciences,* Neural Substrates of Cognition, 29, no. 6 (June): 317–322. ISSN: 0166-2236. https://doi.org/10.1016/j.tins.2006.04.001.

Maurer, Andrew P., and Bruce L. McNaughton. 2007. "Network and Intrinsic Cellular Mechanisms Underlying Theta Phase Precession of Hippocampal Neurons" [in en]. *Trends in Neurosciences,* July INMED/TINS Special Issue—Physiogenic and Pathogenic Oscillations: The Beauty and the Beast, 30, no. 7 (July): 325–333. ISSN: 0166-2236. https://doi.org/10.1016/j.tins.2007.05.002.

Maurer, A.P., S.L. Cowen, S.N. Burke, C.A. Barnes, and B.L. McNaughton. 2006. "Organization of hippocampal cell assemblies based on theta phase precession." *Hippocampus* 16 (9): 785–94.

McNaughton, Bruce L., Francesco P. Battaglia, Ole Jensen, Edvard I. Moser, and May-Britt Moser. 2006. "Path integration and the neural basis of the 'cognitive map'." *Nature Reviews Neuroscience* 7:663—678.

Millikan, R. 2002. "Biofunctions: Two Paradigms." In *Functions: New Essays in the Philosophy of Psychology and Biology,* edited by A Ariew, R Cummins, and M Perlman, 113–143. Oxford University Press.

Millikan, Ruth Garrett. 1984. *Language, Thought and Other Biological Categories.* The MIT Press.

Montgomery, S.M., and G. Buzsáki. 2007. "Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task performance." *Neural Computation* 104 (36): 14495–500.

Mossio, Matteo, Cristian Saborido, and Alvaro Moreno. 2009. "An Organizational Account of Biological Functions." *British Journal for the Philosophy of Science* 60 (4): 813–841.

Muller, Robert, Elizabeth Bostock, Jeffrey S. Taube, and John L. Kubie. 1994. "On the directional firing properties of hippocampal place cells." *The Journal of Neuroscience* 14 (12): 7235–7251.

Muller, R.U., and J. L. Kubie. 1987. "The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells." *Journal of Neuroscience* 7 (7): 1951–68.

Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics.* MIT Press.

Ni, Jianguang, Thomas Wunderle, Christopher Murphy Lewis, Robert Desimone, Ilka Diester, and Pascal Fries. 2016. "Gamma-Rhythmic Gain Modulation." *Neuron* 92, no. 1 (October): 240–251. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2016.09.003.

O'Craven, Kathleen M., Paul E. Downing, and Nancy Kanwisher. 1999. "fMRI Evidence for Objects as the Units of Attentional Selection" [in En]. *Nature* 401, no. 6753 (October): 584. ISSN: 1476-4687. https://doi.org/10.1038/44134.

O'Keefe, J., and M. L. Recce. 1993. "Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus." *Hippocampus* 3:317–330.

O'Keefe, John, and Jonathan Dostrovsy. 1971. "The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat." *Brain Research* 31 (4): 171–175.

O'Keefe, John, and Lynn Nadel. 1978. *The Hippocampus as a Cognitive Map.* [in English]. New York: Oxford University Press.

Palanca, Ben J. A., and Gregory C. DeAngelis. 2005. "Does Neuronal Synchrony Underlie Visual Feature Grouping?" *Neuron* 46, no. 2 (April): 333–346. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2005.03.002.

Papineau, David. 1987. *Reality and Representation.* Basil Blackwell.

Posner, Michael I., and Charles D. Gilbert. 1999. "Attention and Primary Visual Cortex." *Proceedings of the National Academy of Sciences* 96 (6): 2585–2587.

Pritchett, Dominique L, Joshua H Siegle, Christopher A Deister, and Christopher I Moore. 2015. "For Things Needing Your Attention: The Role of Neocortical Gamma in Sensory Perception." *Current Opinion in Neurobiology,* SI: Brain Rhythms and Dynamic Coordination, 31 (April): 254–263. ISSN: 0959-4388. https://doi.org/10.1016/j.conb.2015.02.004.

Pylyshyn, Z. W. 1984. *Computation and Cognition.* MIT Press.

Quiroga, Rodrigo Quian. 2013. "Gnostic Cells in the 21st Century." *Acta Neurobiol Exp (Wars)* 73:463–71.

Ramsey, William M. 2007. *Representation Reconsidered* [in en]. Cambridge University Press, June. ISBN: 978-0-521-85987-5.

Schmidt, R., K. Diba, Ch. Leibold, D. Schmitz, G. Buzsáki, and R. Kempter. 2009. "Single-Trial Phase Precession in the Hippocampus." *Hippocampus* 29 (42): 13232–13241.

Schulte, Peter. 2017. "Perceiving the World Outside: How to Solve the Distality Problem for Informational Teleosemantics." *The Philosophical Quarterly.*

Shadlen, MN, and JA Movshon. 1999. "Synchrony Unbound: A Critical Evaluation of the Temporal Binding Hypothesis." *Neuron* 24, no. 1 (September). ISSN: 0896-6273.

Shannon, Claude. 1948. "A Mathematical Theory of Communication." *The Bell System Mathematical Journal* 27:379–423,623–656.

Shea, Nicholas. 2018. *Representation in Cognitive Science* [in English]. OXFORD: Oxford University Press, November. ISBN: 978-0-19-881288-3.

Singer, W, and C M Gray. 1995. "Visual Feature Integration and the Temporal Correlation Hypothesis." *Annual Review of Neuroscience* 18 (1): 555–586. https://doi.org/10.1146/annurev.ne.18.030195.003011.

Singer, Wolf. 1999. "Neuronal Synchrony: A Versatile Code for the Definition of Relations?" [In English]. *Neuron* 24, no. 1 (September): 49–65. ISSN: 0896-6273. https://doi.org/10.1016/S0896-6273(00)80821-1.

Skaggs, W., B. NcNaughton, M. A. Wilson, and C. A. Barnes. 1996. "Theta Phase Precession in Hippocampal Neuronal Populations and the Compression of Temporal Sequences." *Hippocampus.* 6:149–172.

Skyrms, Brian. 2010. *Signals: Evolution, Learning & Information.* New York: Oxford University Press.

Tallon-Baudry, Catherine, and Olivier Bertrand. 1999. "Oscillatory Gamma Activity in Humans and Its Role in Object Representation." *Trends in Cognitive Sciences* 3, no. 4 (April): 151–162. ISSN: 1364-6613. https://doi.org/10.1016/S1364-6613(99)01299-1.

Treue, Stefan. 2001. "Neural Correlates of Attention in Primate Visual Cortex." *Trends in Neurosciences* 24, no. 5 (May): 295–300. ISSN: 0166-2236. https://doi.org/10.1016/S0166-2236(00)01814-2.

Treue, Stefan, and Julio C. Martinez Trujillo. 1999. "Feature-Based Attention Influences Motion Processing Gain in Macaque Visual Cortex." *Nature* 399 (6736): 575.

von der Malsburg, Christoph. 1981. "The Correlation Theory of Brain Function (Internal Report 81-2)." *Goettingen: Department of Neurobiology, Max Planck Intitute for Biophysical Chemistry.*

————. 1994. "The Correlation Theory of Brain Function." In *Models of Neural Networks,* edited by E. Domany, J. L. van Hemmen, K. Schulten, Eytan Domany, J. Leo van Hemmen, and Klaus Schulten, 95–119. New York, NY: Springer New York. ISBN: 978-1-4612-8736-0 978-1-4612-4320-5. https://doi.org/10.1007/978-1-4612-4320-5_2.

Wang, DeLiang. 1999. "Relaxation Oscillators and Networks" [in en]. In *Wiley Encyclopedia of Electrical and Electronics Engineering.* _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/047134608X.W2282. American Cancer Society. ISBN: 978-0-471-34608-1. https://doi.org/10.1002/047134608X.W2282.

Williams, D., and L. Colling. 2018. "From Symbols to Icons: The Return of Resemblance in the Cognitive Neuroscience Revolution." *Synthese* 195 (5): 1941–1967.

Yazdanbakhsh, Arash, and Stephen Grossberg. 2004. "Fast Synchronization of Perceptual Grouping in Laminar Visual Cortical Circuits." *Neural Networks,* Vision and Brain, 17, no. 5 (June): 707–718. ISSN: 0893-6080. https://doi.org/10.1016/j.neunet.2004.06.005.

Yuste, Rafael. 2015. "From the Neuron Doctrine to Neural Networks" [in en]. *Nature Reviews Neuroscience* 16, no. 8 (August): 487–497. ISSN: 1471-0048. https://doi.org/10.1038/nrn3962.

Zion Golumbic, Elana M., Nai Ding, Stephan Bickel, Peter Lakatos, Catherine A. Schevon, Guy M. McKhann, Robert R. Goodman, et al. 2013. "Mechanisms Underlying Selective Neuronal Tracking of Attended Speech at a "Cocktail Party"." *Neuron* 77, no. 5 (March): 980–991. ISSN: 0896-6273. https://doi.org/10.1016/j.neuron.2012.12.037.