CrossMark

Special issue: Review

A predictive coding framework for rapid neural dynamics during sentence-level language comprehension

Ashley G. Lewis [a,b] and Marcel Bastiaansen [a,c,*]

[a] Neurobiology of Language Department, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands
[b] Radboud University, Donders Institute for Brain, Cognition and Behaviour, Center for Cognitive Neuroimaging, Nijmegen, The Netherlands
[c] Academy for Leisure, NHTV University of Applied Sciences, Breda, The Netherlands

ARTICLE INFO

ABSTRACT

There is a growing literature investigating the relationship between oscillatory neural dynamics measured using electroencephalography (EEG) and/or magnetoencephalography (MEG), and sentence-level language comprehension. Recent proposals have suggested a strong link between predictive coding accounts of the hierarchical flow of information in the brain, and oscillatory neural dynamics in the beta and gamma frequency ranges. We propose that findings relating beta and gamma oscillations to sentence-level language comprehension might be unified under such a predictive coding account. Our suggestion is that oscillatory activity in the beta frequency range may reflect both the active maintenance of the current network configuration responsible for representing the sentence-level meaning under construction, and the top-down propagation of predictions to hierarchically lower processing levels based on that representation. In addition, we suggest that oscillatory activity in the low and middle gamma range reflect the matching of top-down predictions with bottom-up linguistic input, while evoked high gamma might reflect the propagation of bottom-up prediction errors to higher levels of the processing hierarchy. We also discuss some of the implications of this predictive coding framework, and we outline ideas for how these might be tested experimentally.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Reading, or listening to someone speaking, are the simple kinds of tasks that most people engage in every day of their lives without much difficulty. Yet if one considers that the average reader can easily manage between 250 and 300 words per minute (e.g., Rayner, Pollatsek, Ashby, & Clifton, 2012), it becomes clear that the processing carried out by the language comprehension system must be extremely fast and dynamic.

* Corresponding author. Academy for Leisure, NHTV Breda University of Applied Sciences, Archimedesstraat 17, 4816 BA, Breda, The Netherlands.
E-mail address: bastiaansen4.m@nhtv.nl (M. Bastiaansen).

One possible explanation for this speed (to be sure, one among many) is that the system may make predictions about upcoming linguistic input. From such a perspective it is surprising that models of language comprehension based on the passive building up of semantic and syntactic structures (from the lexical building blocks activated upon perception of linguistic input) dominated the psycholinguistics literature for so long (e.g., Forster, 1981; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Zwitserlood, 1989). Arguments that prediction was not likely to be involved in language comprehension were generally made based on the observation that at any point while reading or listening there are a large number of possible continuations. Processing costs involved in making incorrect predictions, along with the presumed low percentage of benefits accrued (predictions would not often be correct) made predictive processing accounts unappealing (see van Petten & Luka, 2012 for discussion).

On the other hand, a large number of studies began to show that the processing of a word in a sentence can be facilitated by the constraining sentence context (sometimes even before the word can be uniquely identified; e.g., Altmann & Kamide, 1999, 2007; Balota, Pollatsek, & Rayner, 1985; van den Brink, Brown, & Hagoort, 2001; Ehrlich & Rayner, 1981; Federmeier & Kutas, 1999; Kamide, 2008; Kamide, Altmann, & Haywood, 2003; Kamide, Scheepers, & Altmann, 2003; Knoeferle, Crocker, Scheepers, & Pickering, 2005; MacDonald, Pearlmutter, & Seidenberg, 1994; McRae, Hare, Elman, & Ferretti, 2005; van Petten, Coulson, Rubin, Plante, & Parks, 1999; Schwanenflugel & Lacount, 1988; Sussman & Sedivy, 2003). The idea that predictive processing could, at least in some circumstances, be beneficial for language comprehension has slowly grown in popularity. By now the notion that (at least some of the time) prediction plays an important role in rapid, dynamic, real-time language comprehension is a widely accepted view (Pickering & Garrod, 2007).

However, within this emerging view there are many outstanding questions. For instance, what are the details about exactly when prediction plays a role (is the system always making predictions or only under certain circumstances when this may be a useful strategy?). How do predictions interact with real-time comprehension? What kinds of information might lead to (strong) predictions? And, crucially, how does the brain implement predictive processing? While we briefly discuss each of these questions we acknowledge that it is not possible to do justice to them all in a single review. The main focus of this review is to outline some ways in which we think that the study of electrophysiology, and in particular oscillatory neural dynamics measured using electroencephalography (EEG) and magnetoencephalography (MEG) can contribute to our understanding of predictive processing during language comprehension beyond the level of individual words.

1.1. Event-related potential (ERP) studies and prediction during sentence comprehension

In the last ten to fifteen years a number of ERP studies have investigated the potential role of prediction during sentence-level language comprehension (see e.g., van Berkum, Brown, Zwitserlood, Kooijman, & Hagoort, 2005; DeLong, Urbach, & Kutas, 2005; Otten, Nieuwland, & Van Berkum, 2007; Szewczyk & Schriefers, 2013; Wicha, Moreno, & Kutas, 2004). The common ingredient used in all these studies to investigate predictive processing was agreement relations between a particular noun and some element preceding the noun. If the constraining sentence context allows readers/listeners to make predictions about the following noun, then lexical information associated with that noun should be available to the comprehension system before the noun is presented and should have an influence on the processing of agreement relations between the noun and the element preceding it.

An effect of prediction on ERP responses has been shown in the context of both gender-marked determiners (Wicha et al., 2004), and adjectives (van Berkum et al., 2005; Otten et al., 2007) preceding some highly expected noun in strongly constraining sentence contexts. These congruity (congruous or incongruous gender agreement) effects prior to the presentation of the word eliciting them are not the result of simple word-priming (Otten et al., 2007) and can occur more than a single word in advance of the target noun (van Berkum et al., 2005). Along similar lines, the effects of prediction on ERP responses have been shown to be graded in nature (DeLong et al., 2005), dependent on the target noun's cloze probability (a normative measure that in most circumstances can be taken as a proxy for how predicted a particular word is in a given sentence context; cf., Kutas & Federmeier, 2011). In addition to grammatical (van Berkum et al., 2005; Otten et al., 2007; Wicha et al., 2004) and phonological (DeLong et al., 2005) information, it has recently been shown that semantic information (in this case the semantic class of animacy) about an upcoming noun may also be predicted, and has an effect on ERP responses before the target noun (Szewczyk & Schriefers, 2013).

Taken together these studies make a strong case for (graded) predictions during sentence-level language comprehension, and not simply predictions about particular words but also about (at least some) semantic categories of words. They also show that electrophysiological brain responses (in this case ERPs) are sensitive to (at least some of) the processing consequences of these predictions.

In addition to syntactic features associated with specific lexical items (e.g., gender or number marking), other non-local syntactic dependencies may also lead to predictive processing, and the prediction of particular syntactic structures. For example, one prominent account of the P600 ERP component is as a reflection of processes of reanalysis and repair (e.g., Friederici, 2002; Friederici & Mecklinger, 1996). A P600 effect has been reported in the case of syntactic garden path sentences (e.g., Osterhout, Holcomb, & Swinney, 1994), for syntactic ambiguity resolution with object-compared to subject-relative clauses (Mecklinger, Schriefers, Steinhauer, & Friederici, 1995), and for syntactic violations (Hagoort, Brown, & Groothusen, 1993). All these cases have in common that they involve a preferred syntactic structure that is constructed and needs to be revised or repaired at a point where the input indicates that it is not correct (Friederici & Mecklinger, 1996). Although they have not traditionally be interpreted in this way, it is possible to argue that all these cases involve a prediction (by the language comprehension system) that a particular syntactic construction will

accurately account for the linguistic input. At some point during the sentence, the input provides evidence that this prediction was incorrect, and the P600 may be thought of as the brain's response to such a failed prediction. Reframing some of the classical P600 findings in this way may provide a hint that readers/listeners make predictions about likely syntactic structures or structural dependencies while reading/listening.

1.2. Prediction and the brain's language comprehension network

Before moving on to discuss oscillatory neural dynamics during sentence-level language comprehension, we first outline a framework within which we can describe the relationship between predictive processing, the language comprehension system, and their associated functional brain network dynamics. There are a number of models relating the cognitive architecture of sentence-level language comprehension to its underlying neural infrastructure (e.g., Friederici, 2002; Jung-Beeman, 2005; Lau, Phillips, & Poeppel, 2008), but we adopt the framework and terminology used by Hagoort and co-workers (Baggio & Hagoort, 2011; Hagoort, 2005, 2013; Hagoort, Baggio, & Willems, 2009). A memory component, implemented by left temporal cortical areas, is responsible for the retrieval of lexical building blocks containing phonological, syntactic, and semantic properties of individual words. A unification component on the other hand is responsible for combining these building blocks to form a meaningful interpretation of the linguistic input. The unification component is implemented by left inferior frontal cortical regions, and their dynamic, coordinated interaction with left temporal and left inferior parietal cortex (Hagoort, 2014). A third component, the control component, completes the memory, unification, and control (MUC) framework (Hagoort, 2005, 2013) but for our purposes we will focus mainly on the memory and unification components.

The ERP studies discussed in Section 1.2 suggest that the system is likely engaged in ongoing predictive processing whenever possible during sentence-level language comprehension. Typically, while one reads or listens to a sentence, the predictability of upcoming linguistic information increases from beginning to end (because there are more plausible possibilities for continuing the sentence at the beginning than near the end of a sentence). In terms of the MUC framework pre-activation of specific lexical items or semantic categories, as well as biases towards particular syntactic structures due to prediction are the result of the dynamic interaction between the unification component and the memory component. The unification component sends feedback to the memory component during each word processing cycle (Hagoort, 2013), and in the case of highly constraining contexts (or in other cases where the system might use predictive processing) this information may prompt the memory component to pre-activate highly predicted lexical items (or at least some of the information associated with those items, e.g., their semantic category). Similarly, predictions about particular syntactic structures could bias the weighting of connections between nodes in the syntactic representation being built. The unification component would be responsible

for this weighting, while the memory component would activate the syntactic treelets containing relevant syntactic nodes (Hagoort, 2005).

1.3. Fast oscillatory neural dynamics during language comprehension

A large amount of evidence has accumulated over the last two or more decades suggesting that the coupling and uncoupling of functional networks in the brain is related to patterns of neural synchronization and desynchronization (Bastiaansen & Hagoort, 2006; Bastiaansen, Mazaheri, & Jensen, 2012; Pfurtscheller & Lopes da Silva, 1999b; Singer, 1993, 2011; Varela, Lachaux, Rodriguez, & Martinerie, 2001; Womelsdorf et al., 2007). One instance of this occurs when areas that are part of the same functional network are linked by synchronous oscillatory firing in the same frequency range. Conceptually, synchronous repetitive firing of neurons increases the probability that they entrain one another and thereby activates participating functional networks at particular frequencies (König & Schillen, 1991). In this way the brain achieves frequency-specific segregation of information being processed by different functional networks. On the other hand, frequency-specific oscillatory neural synchrony also binds together information represented in different elements or subcomponents of the same functional network (Gray, König, Engel, & Singer, 1989). Another instance that has recently received a large amount of interest (and which is beyond the scope of this review) is cross-frequency coupling, where the phase of low frequency oscillations modulate the amplitude of oscillations in a higher frequency range (e.g., Lisman & Jensen, 2013). Such oscillatory neural phenomena typically have similar functions across multiple spatial and temporal scales. Modulations of frequency-specific power are often associated with synchrony within local neural populations, while modulations of frequency-specific phase coupling measures (e.g. coherence or phase-locking value) are most often associated with synchrony between more distant neural populations (inter-areal synchrony). There is however no clear distinction between local and inter-areal synchrony, and hence no guarantee that power always measures local synchrony and coherence always measures inter-areal communication (Varela et al., 2001).

We should note that the alpha frequency band (around 8–12 Hz) and perhaps, to some extent also the beta frequency band (around 13–30 Hz), do not straightforwardly fit in this framework. It has been observed that desynchronization in the alpha frequency range in a specific brain area sometimes entails the engagement/activation of that brain area, especially when related to motor (Bastiaansen & Brunia, 2001; Pfurtscheller & Lopes da Silva, 1999a; Pfurtscheller, Neuper, Andrew, & Edlinger, 1997) and sensory (Bastiaansen & Brunia, 2001) processing. Similarly, beta-band desynchronization has often been related to motor cortex activation (see e.g., Parkes, Bastiaansen, & Norris, 2006; Pfurtscheller, Stancák, & Neuper, 1996; Pfurtscheller, Zalaudek, & Neuper, 1998). The exact relationship between alpha/beta desynchronization and the process of functional network recruitment (in the sense of König & Schillen, 1991) and binding (in the sense of Gray et al., 1989) is yet to be established.

A growing body of literature has accumulated relating sentence-level language comprehension to event-related changes in EEG and MEG oscillations (e.g., Bastiaansen & Hagoort, 2010; Bastiaansen, Magyari, & Hagoort, 2010; Peña & Melloni, 2012; for review see Bastiaansen et al., 2012; Lewis, Wang, & Bastiaansen, 2015). Such studies typically investigate patterns of fast temporal dynamics associated with the coupling and uncoupling of nodes in the brain's language network. Effects have been found in all the classical frequency ranges, with for example theta oscillations being linked to lexical retrieval operations and semantic working memory, and alpha being linked to task-specific working memory load (Bastiaansen & Hagoort, 2006; Bastiaansen et al., 2012; Meyer, Obleser, & Friederici, 2013; Weiss et al., 2005). Bastiaansen and Hagoort (2010) proposed the 'frequency segregation of unification' hypothesis, which integrates a substantial body of empirical data into a framework in which oscillatory activity in the beta and gamma frequency bands reflect syntactic and semantic unification operations respectively. We have recently reviewed the evidence both for and against that hypothesis, and suggested that beta and gamma oscillations during sentence-level language comprehension may be at least equally well, or perhaps even better explained in relation to predictive processing, and maintenance/change of the current cognitive set respectively (Lewis et al., 2015).

1.4. Predictive coding and hierarchical Bayesian inference

In Section 2.1 we will outline how we think (at least some aspects of) oscillatory neural dynamics during sentence-level language comprehension might be explained in a predictive coding framework. Here we provide a brief outline of the particular flavor of predictive coding we will adopt (Friston, 2005). In this framework, brain systems are considered to be hierarchically structured, with higher hierarchical levels creating probabilistic (or forward) models designed to explain cortical (or sub-cortical) activity at lower levels (Clark, 2013; Friston, 2005).

For any two hierarchical levels, predictions are made at the higher level and propagated to the lower level in a top-down fashion. Here, such predictions are matched with bottom-up information (the actual activity in that lower level based on its inputs) to compute a prediction error (the difference between the prediction and the actual activity). Bottom-up information takes the form of the prediction errors themselves, so that this information contains only the amount of surprisal (an information theoretic measure of how far off the predictions actually were) based on the mismatch between the prediction and the actual activity in the lower hierarchical level. Those prediction errors are then used by the higher hierarchical level to update the predictions being made about the activity at the level below (Clark, 2013; Friston, 2005).

In the framework of Friston (2005), each hierarchical level contains both representational units and error units. Representational units represent activity at the current hierarchical level, provide top-down predictions to lower hierarchical levels, and receive inputs from error units at lower hierarchical levels with which they update predictions (or forward models) at the current level. Error units on the other hand receive input from representational units at the current and at higher hierarchical levels. They compute prediction errors based on the mismatch between top-down predictions and bottom-up information (activity at the current level) and send those prediction errors to higher hierarchical levels. Error units at the same hierarchical level also interact in order to decorrelate and laterally inhibit one another. The system attempts to minimize prediction error, and in this way achieve an optimal model of the events or causes of activity at different hierarchical levels (for more details see Friston, 2005).

The actual organization of the brain is of course far more complicated than the simple picture outlined above, with multiple hierarchical layers (often embedded within sub-layers) and fast, dynamic interactions between layers resulting in the constant updating and refinement of myriad forward-models at different hierarchical levels. Nonetheless, the static view outlined above provides a useful descriptive tool for probing the relationship between predictive coding, oscillatory neural dynamics, and various cognitive phenomena (for a formal description of a more dynamic implementation see e.g., Friston, 2005). This framework has been highly successful in accounting for a wide range of phenomena with a relatively simple mechanistic explanation of how information flows between and within cortical (and subcortical) hierarchies (e.g., spike-time dependent plasticity (STDP) during learning, classical and extra-classical receptive field properties in vision, repetition suppression, priming effects, ERP responses to learned sequences; see Friston, 2005 for details). In the next section we will see whether applying some of these principles to sentence-level language comprehension might prove useful in better understanding its neural implementation.

2. Beta and gamma oscillatory dynamics during language comprehension

Thus far we have briefly discussed prediction during sentence-level language comprehension, oscillatory neural dynamics in relation to the formation of functional brain networks, and a predictive coding framework for understanding the flow of information between hierarchical levels in the brain. In Section 2.1 we attempt to unify these three areas of investigation, and in Sections 2.2 and 2.3 we review the available data from this new perspective, in order to evaluate its explanatory power.

2.1. Language comprehension, neural oscillations, and predictive coding

It has been shown experimentally in monkey visual cortex and in rat somatosensory cortex (although the hope is that the general principles will also apply in humans) that gamma oscillations are most prominently expressed in supragranular cortical layers (L2/3), while beta oscillations are more prominent in infragranular (and granular) layers (L4/5; Maier, Adams, Aura, & Leopold, 2010; Roopun et al., 2006, 2008). At the same time, feedforward connections predominantly originate in superficial layers (L2/3) and terminate in L4, while

feedback connections originate from deeper layers (L4/5) and terminate outside of L4 (Bastos et al., 2012). This has led to the proposal (Wang, 2010) that within cortical hierarchies, feedforward signaling may be mediated by high frequency oscillations (in the gamma range for instance) compared to feedback signaling, which may be mediated by oscillations at lower frequencies (in the beta or alpha range). Bastos et al. (2012) have suggested that this principle might constitute a canonical form of hierarchical functional organization in the brain. The proposal is that within a cortical processing hierarchy gamma oscillations might predominate for bottom-up interactions, while beta oscillations might predominate in the top-down direction. The levels of such processing hierarchies can be restricted to local cortical regions (e.g., occipital cortex for most of the visual system), but can also span non-local cortical regions (e.g., left inferior frontal cortex and left middle temporal cortex for two important parts of the core language processing hierarchy).

Bastos et al. (2012) have also proposed that this canonical hierarchical organizing principle might provide physiological correlates of the implementation of predictive coding within cortical hierarchies. From a predictive coding perspective, top-down information (conveyed by feedback connections) provides context for lower-level processing over slower time scales (beta oscillations), while bottom-up information (conveyed by feedforward connections) works on faster time scales (gamma oscillations) propagating prediction errors up the hierarchy in order to rapidly adapt predictions at these higher levels. This implies that oscillatory activity in the beta frequency range might be a proxy for top-down predictions about activity at lower hierarchical levels within a cortical hierarchy, while gamma oscillations might be an indication of the forward propagation of prediction errors to higher cortical levels in order to update predictions (Bastos et al., 2012; Friston, Bastos, Pinotsis, & Litvak, 2015).

Bressler and Richter (2015) propose that cortical areas recruited under task-specific conditions may be linked by inter-areal beta synchrony to form NeuroCognitive Networks (NCNs – self-organizing, large-scale distributed cortical networks) at the highest hierarchical levels. They propose that beta activity may serve the dual purpose of maintenance of such networks, and carrying top-down signals to lower levels of the cortical hierarchy. We propose that the construction of a sentence-level meaning representation during unification entails the formation of such a NCN, encompassing areas in left inferior frontal cortex, left temporal cortex, and left inferior parietal cortex (Hagoort, 2013, 2014), along with other relevant areas outside the core language network depending on the particular context in which sentence-level meaning construction is taking place (e.g., recruitment of the theory of mind network for taking another person's perspective). By 'sentence-level meaning' we are referring not just to the semantics associated with the individual words comprising a sentence, but also to the semantics derived from the syntactic structure governing the hierarchical relations between those words. The above ideas about beta fit well within the framework of Engel and Fries (2010), who suggest that oscillatory activity in the beta range reflects the active maintenance or change of the current cognitive set. In this case an NCN constitutes the neural implementation of the current cognitive

set (Bressler & Richter, 2015).Beta increases indicate that the current NCN configuration is being actively maintained, while beta decreases indicate that the current NCN configuration is under revision/change.

Our previous proposal about the relationship between beta oscillations and maintenance/change of the current cognitive set (Lewis et al., 2015) is thus subsumed under the current proposal of an NCN for representation and construction of the current sentence-level meaning. As pointed out by Bressler and Richter (2015), beta oscillations could be simultaneously involved in both the maintenance of the current NCN, as well as the propagation of top-down predictions (perhaps based on the information represented or processed in the NCN) to lower levels in the cortical processing hierarchy (Fig. 1). We would like to emphasize that these two related roles of beta oscillations are entirely compatible with one another. The representation of information in a distributed NCN and the use of that information to make predictions in a top-down fashion in our opinion constitute closely related and heavily interdependent forms of neural processing. In the case of sentence-level language comprehension, beta synchrony may therefore be responsible for the active maintenance of the current NCN supporting sentence-level meaning construction, as well as the top-down transfer of predictions that the sentence-level meaning might convey to lower levels (e.g., the memory component responsible for lexical retrieval) of the cortical processing hierarchy. Such predictions can be about individual words, but also about other units of linguistic information (e.g., particular syntactic constructions; cf. Levy, 2011).

We have suggested a role for oscillatory activity in the gamma band during sentence-level language comprehension in matching incoming linguistic input with pre-activated lexical representations (Lewis et al., 2015). This was proposed based on the 'match-and-utilization' model of Herrmann, Munk, and Engel (2004), as well as on earlier proposals for language processing from Peña and Melloni (2012), and from Wang, Zhu, and Bastiaansen (2012). The idea is that gamma activity reflects a match between strong top-down predictions and bottom-up linguistic input. We now tentatively suggest that the gamma activity in this case may reflect synchrony between neural populations (containing representational and error units), related to both the pre-activated lexical representation that matches the input (synchrony results in resonance within the local neural population for that representation) and the suppression of neural populations (and hence lateral inhibition of competing error units at the same hierarchical level) related to competing lexical representations (synchrony results in increased inhibition of connected neural populations that are not part of the matching representation). Reciprocal lateral connections (which tend to be inhibitory in nature) within the same hierarchical processing level would perform the function of lateral inhibition of competing representations, and such reciprocal lateral connections are indeed present in the cortex (Bastos et al., 2012). Functionally, once a clear match is made between a pre-activated lexical item (strong prediction) and the input, competing lexical representations should be strongly suppressed. In the case of input that is incongruent with the preceding sentence context no gamma synchrony occurs because the linguistic input does not match any top-down

Simplified illustration of hierarchical information flow during language comprehension
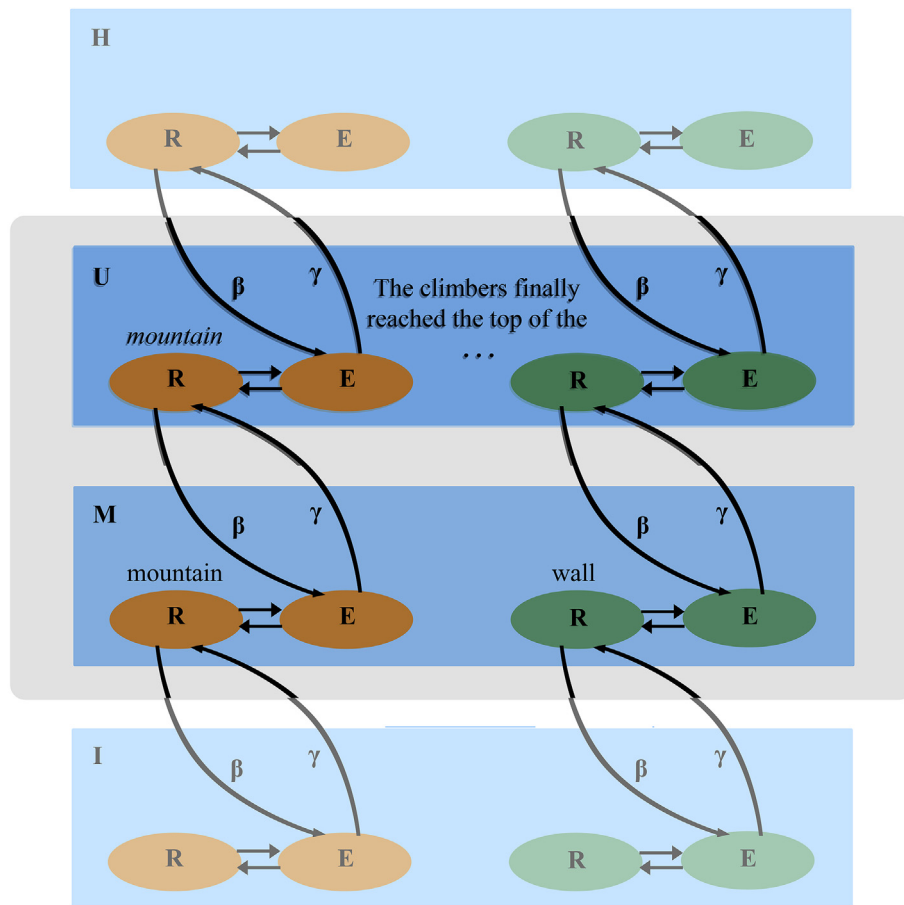

Fig. 1 – **Simplified illustration of the proposed hierarchical flow of information during language comprehension. Blue boxes refer to different levels of the processing hierarchy. We focus (highlighted portion of figure) on levels corresponding to the memory component (M) and the unification component (U) from the MUC framework (Hagoort, 2013). Also pictured (but not highlighted) are an input level (I) for auditory, visual or other types of input to the language comprehension system, and a higher level (H) for things like cognitive control or other forms of higher-level processing. Within each hierarchical layer there are multiple representation (R) and error (E) units. Those in orange are relevant for our example, while those in green indicate (more than one pair of) potential competing representation and error units. Not pictured in the figure (to avoid the figure becoming too cluttered) are inhibitory lateral connections between these (orange and green) representation and error units at the same hierarchical level, responsible for lateral inhibition of competing representations. In the example from the figure, a reader has read the input 'The climbers finally reached the top of the' and makes a strong prediction that the next word will be 'mountain'. We have suggested that nodes in the NCN responsible for constructing and maintaining a representation of the sentence-level meaning are linked via oscillatory activity in the beta frequency range. This is not depicted in the figure but the network comprises (amongst others) nodes from left inferior frontal gyrus, left inferior parietal and left temporal cortex. Based on this sentence-level meaning, the system generates a prediction at U. This prediction is sent down the processing hierarchy (via oscillatory activity in the beta band) from an R unit at U, to an E unit at M, where the prediction is matched with incoming linguistic information (sent from the R unit at M to the E unit at that same level) to compute a prediction error. That prediction error is sent back up the processing hierarchy (via oscillatory activity in the high gamma range) from an E unit at M to an R unit at U (but also to an R unit at M to update representation units at the same level) so that the generative model (and hence the prediction) at this higher level can be updated if necessary. When the input matches a strong prediction (e.g., in our example the input is 'mountain') this results in an increase in low-mid gamma power reflecting the match, as well as strong lateral inhibition (not depicted in the figure) of competing representation and error units at the same hierarchical level.**

predictions. Similarly, when the target word is less predictable (e.g., for less constraining sentence contexts) but still congruent with the preceding sentence context, there are no (or perhaps only weak) top-down predictions with which the

incoming bottom-up linguistic input can match, and hence no gamma synchrony is observed.

The Bastos et al. (2012) framework explicitly suggests that prediction errors should be sent up the cortical hierarchy via

gamma frequencies, while thus far we have only discussed gamma within the same hierarchical processing level (resonance when top-down predictions match bottom-up input and concomitant lateral inhibition of competing representations). The gamma effects reported in Lewis et al. (2015) are primarily effects in the low and middle gamma range (approximately 35–75 Hz). We suggest that prediction errors might be propagated to higher levels of the processing hierarchy by evoked gamma at higher frequencies (approximately 80–130 Hz; see e.g., Fontolan, Morillon, Liegeois-Chauvel, & Giraud, 2014; Giraud & Poeppel, 2012), while gamma in the lower and middle frequency ranges are a reflection of lateral connections within the same hierarchical processing level (Fig. 1). Our proposal is admittedly somewhat speculative at this stage, but we feel justified in speculating since in our opinion further empirical investigation could yield great benefits for our understanding of the relationship between language comprehension and hierarchical neural information processing.

The proposal then is that during sentence-level language comprehension, beta activity reflects the formation and active maintenance/change of an NCN responsible for the construction of a sentence-level meaning representation. Beta may also reflect the propagation of predictions (based on the information represented in that NCN) to lower levels of the processing hierarchy, sometimes leading to lexical preactivation. Low and middle gamma activity reflects a match between highly predictable (and so pre-activated) lexical representations and the incoming linguistic input. This gamma synchrony may be an index of both resonance between representational (and error) units related to the activated lexical item, and lateral inhibition of competing representations when a clear match comes about. We have also suggested that prediction errors may be propagated up the processing hierarchy by evoked high gamma oscillations. In Sections 2.2 and 2.3 we review the available evidence to see whether it supports our claims about beta and gamma respectively.

As an important aside, we have suggested that it might be viable to understand neural oscillations from the perspective of the underlying network dynamics they reflect and the flow of information between and within cortical processing hierarchies. On this view different cognitive processes and/or functions differentially recruit these networks in various ways in order to carry out computations, or to represent relevant information (for related ideas see e.g., Siegel, Donner, & Engel, 2012). The frequency-specific oscillatory dynamics reflect the underlying cortical network dynamics rather than being a direct reflection of some cognitive process or function (i.e., nothing in the network dynamics itself determines what kind of cognitive function is being implemented, but if we know what cognitive function is being investigated we can certainly link that function in that context to the underlying pattern of network dynamics). This proposal implies that a description of some cognitive phenomenon at the level of oscillatory neural dynamics (and from a predictive coding perspective) provides just one level of description (linking cognitive function to the dynamics in underlying neural systems), and in order to fully understand that phenomenon an appropriate description is also necessary at the cognitive-functional level (which is how the system is defined in the first place). While it may sometimes be the case that there is a one-to-one mapping between a neural, systems level explanation and a cognitive-functional level explanation, we would suggest that this is rarely the case for higher-order cognitive functions, and so both levels of explanation are necessary.

2.2. Beta oscillations and NeuroCognitive Networks

All the studies discussed in this and the next section were specifically designed to investigate the relationship between neural oscillations and some particular aspect of sentence-level language comprehension. Next we discuss these studies in relation to our suggested predictive coding framework, however in so doing we do not wish to detract from any insights they may have provided about language processing.

A number of studies have compared syntactically legal sentences (e.g., 'Janneke got the *blessing* at the river') to sentences containing a syntactic violation (e.g., 'Janneke got the *to bless* at the river'; Bastiaansen et al., 2010; Davidson & Indefrey, 2007; Kielar, Meltzer, Moreno, Alain, & Bialystok, 2014; Pérez, Molinaro, Mancini, Barraza, & Carreiras, 2012). They have all reported that power in the beta frequency range is higher at the target word for syntactically legal sentences compared to sentences containing a syntactic violation. Bastiaansen et al. (2010) have also shown that power in the beta band is higher at the target word for syntactically legal sentences compared to the same words in random order (e.g., 'The the Janneke *blessing* got river at'). Extending these findings, Bastiaansen et al. (2010) showed that beta power increases linearly over the course of syntactically legal sentences, remains consistently low over the course of random word lists (see also Bastiaansen & Hagoort, 2010 for a replication of these findings), and for sentences containing syntactic violations shows a linear increase up to the point of the violating word before rapidly returning to baseline levels. These findings have been taken as support for the idea that oscillatory activity in the beta frequency range might be related to syntactic unification operations during sentence-level language comprehension (e.g., Bastiaansen & Hagoort, 2010; Bastiaansen et al., 2010; Lewis et al., 2015).

Further support for this idea comes from studies showing that beta power is higher for sentences which are more demanding in terms of syntactic unification load than for less demanding sentences. Bastiaansen and Hagoort (2006) reported that beta power was higher for syntactically more demanding center-embedded (e.g., 'The mouse that the cat chased ran away') compared to right-branching relative clauses (e.g., 'The cat that chased the mouse ran away'). Similarly, syntactically more demanding object-relative clauses showed higher beta coherence just after the relative clause than their simpler subject-relative counterparts (Weiss et al., 2005). Meyer et al. (2013) showed that beta power was higher for long-compared to short-distance subject-verb agreement dependencies at the point in a sentence where the agreement relation between a subject and subsequent verb had to be computed. Since syntactic working memory load is higher for long-compared to short-distance dependencies, this leads to higher load on the system responsible for syntactic

unification, and this result can thus be interpreted as support for a link between beta band oscillations and syntactic unification.

So far, all the evidence seems to point strongly to a link between oscillatory activity in the beta band and syntactic unification operations during sentence-level language comprehension. Not all the data are consistent with this interpretation however. For one thing, semantic violations (e.g., 'The climbers finally reached the top of the *tulip*') elicit decreases in beta power relative to semantically (and syntactically) legal sentences (e.g., 'The climbers finally reached the top of the *mountain*'; Kielar et al., 2014; Luo, Zhang, Feng, & Zhou, 2010; Wang, Jensen, et al., 2012). Luo et al. (2010) also showed a beta power decrease for rhythmically abnormal target nouns (in verb−noun pairs in Chinese) compared to their rhythmically normal counterparts. Furthermore, Pérez et al. (2012) showed that for Spanish 'Unagreement' (where there is a mismatch between the grammatical person feature marking on the subject and the verb of a sentence, but where that sentence still remains perfectly grammatical; see Pérez et al., 2012 for more details) there is a decrease in beta power (similar to the beta power decrease reported above for the genuine agreement violation condition in that study) compared to syntactically legal sentences. Since 'Unagreement' does not strictly speaking represent a case of syntactic violation it is not clear why syntactic unification should be disrupted in this case, and so the beta power decrease is unlikely to reflect syntactic unification difficulties.

We therefore proposed (Lewis et al., 2015) that the more domain-general framework of Engel and Fries (2010) might provide a better explanation for the beta findings reported above. On that account beta power increases reflect active maintenance of the current cognitive set (which we defined as the current sentence-level meaning representation under construction), while decreases in beta power reflect a change in the current cognitive set (and an associated change in the underlying functional network configuration). If we extend this idea to the framework proposed in Section 2.1, beta power reflects the maintenance of the NCN responsible for the construction and representation of the current sentence-level meaning during unification. A decrease in beta power would signal a change in the current NCN as a result of some cue in the linguistic signal indicating to the system that the current sentence-level meaning representation needs to be revised. In addition to maintenance of the current NCN, beta synchrony reflects the top-down flow of predictions from higher to lower levels of the processing hierarchy. This fits well with the idea that when unification is proceeding normally, the NCN is maintained and predictions can be confidently propagated down the processing hierarchy (and this is reflected in increased beta activity). When there are cues in the linguistic input indicating that the current NCN needs to change, the context provided by the current sentence-level meaning no longer provides reliable predictions to the lower levels of the processing hierarchy, and so it makes sense that the system would suppress the flow of such top-down information (resulting in lower beta activity in those cases).

Under this proposal, we can easily account for the cases of syntactic and semantic violations (and similarly for rhythmic 'violations' and the case of 'Unagreement' in Spanish, where

although not ungrammatical the agreement mismatch would constitute an unexpected event for the system) by realizing that they might act as cues to the language comprehension system indicating that the current sentence-level meaning under construction is incorrect in some way and needs to be changed. This would result in a change in the underlying NCN and hence the beta power decrease observed relative to syntactically and semantically legal sentences. This decrease in beta power may also reflect diminished confidence in the predictions that can be sent to lower levels of the processing hierarchy based on the current sentence-level meaning. The cases where increased syntactic unification load results in increased beta power can be dealt with if we accept that an increase in syntactic unification load may act as a cue to the language comprehension system indicating that the current NCN needs to be actively maintained. According to Engel and Fries (2010) this would result in an increase in beta power relative to the conditions with lower syntactic unification load, and this is exactly what was observed. The increased beta may also reflect the need for higher weighting of top-down predictions (in order to prioritize information related to the current NCN) in the case of syntactically more demanding sentences. A related alternative explanation of the Meyer et al. (2013) findings might be that the long distance agreement dependencies allow for stronger predictions about when the dependency resolving verb is likely to appear (cf., Levy, 2008). If that is the case, greater reliance on such top-down predictions would explain the higher beta power at the verb for the long-compared to the short-distance dependencies.

Strong support for our proposal comes from a recent turn-taking experiment where participants listened to recordings of natural speech and had to press a button when they predicted that their interlocutor would finish their turn (Magyari, Bastiaansen, de Ruiter, & Levinson, 2014). Stimuli were constructed so that in one condition the turn-end was highly predictable, while in the other it was unpredictable. A large decrease in beta power was present just before the key-press in the highly predictable condition, while in the unpredictable condition there was an increase in beta power before the key-press. Under our proposal, in the highly predictable condition the language comprehension system predicts that the current NCN will soon need to change (in preparation for constructing a new sentence-level meaning representation) and that results in the decrease in beta power. In the unpredictable condition on the other hand, the language comprehension system in engaged in ongoing sentence-level meaning construction and has no reason to expect it to change yet, so the current NCN should be maintained and this results in the observed increase in beta power.

2.3.   *Gamma oscillations and predictive processing*

By now there are a number of studies that have compared sentences containing a semantically highly predictable target word (e.g., 'The Dutch trains are *yellow* and blue'; normally measured using a cloze probability test) to sentences containing a semantic violation (e.g., 'The Dutch trains are *sour* and blue'; Hald, Bastiaansen, & Hagoort, 2006; Penolazzi, Angrilli, & Job, 2009; Rommers, Dijkstra, & Bastiaansen, 2013;

Wang, Zhu, et al., 2012; Weiss & Mueller, 2003). They have all reported increased power (or in one case coherence; Weiss & Mueller, 2003) in the gamma frequency range for highly predictable target words compared to semantic violations. Extending these findings, Bastiaansen and Hagoort (2010) showed larger gamma power throughout the sentence for semantically legal sentences compared to sentences containing syntactic prose (syntactically acceptable sentences devoid of any sentence-level meaning, e.g., 'The dusty prison engraves the gender for the innocent throat'). In another study, higher gamma power was found for referentially correct target words compared to referentially ambiguous or referentially failing critical words (van Berkum, Zwitserlood, Bastiaansen, Brown, & Hagoort, 2004). Using a slightly different approach Peňa and Melloni (2012) found increased gamma power when monolingual participants listened to sentences in their native language, which was absent when listening to sentences in either a phonologically related or unrelated language.

The evidence reviewed so far all support the suggestion of a strong link between oscillatory activity in the gamma frequency range, and semantic unification (Bastiaansen & Hagoort, 2010; Lewis et al., 2015). There are however data that are not consistent with this proposal. For instance, an increase in gamma power was observed following target words that violated participants' world-knowledge (while still being semantically congruent within their sentence context), but not for semantically legal sentences or sentences containing semantic violations (Hagoort, Hald, Bastiaansen, & Petersson, 2004). The increased gamma power in the world-knowledge violation condition may be interpreted as reflecting increased demands on semantic unification, whereas semantic unification is not possible in the semantic violation condition. This however does not explain the absent gamma power increase (which would be predicted if gamma power is an index of semantic unification) for the semantically legal condition. Perhaps even more problematic for the gamma-semantic unification link are studies showing no gamma power increase for semantically congruent but less highly predicted words. Wang, Jensen, et al. (2012) and Wang, Zhu, et al. (2012) had a condition where target words were semantically congruent, but had relatively low cloze probability scores (and hence were less predictable). They showed a gamma power increase only for highly predictable words, but not for the less predictable words, or the semantic violations. This seems to indicate that gamma power reflects the predictability of the upcoming word in its sentence context rather than semantic unification. Further evidence for this comes from a study that also compared high to low cloze probability target words in a sentence context (Molinaro, Barraza, & Carreiras, 2013). They used phase-locking values to show higher transient phase-coupling between frontal and posterior electrodes in the gamma frequency range for the high cloze condition compared to the low cloze condition. Finally, using mixed-effects modeling, Monsalve, Pérez, and Molinaro (2014) showed that cloze probability was a significant positive predictor of gamma power, so that higher cloze probability meant higher gamma power.

We therefore proposed (Lewis et al., 2015; see Peňa & Melloni, 2012; Wang, Jensen, et al., 2012; Wang, Zhu, et al.,

2012 for related earlier proposals) that the findings relating oscillatory neural dynamics in the gamma frequency range to sentence-level language comprehension may be better captured under the more domain general 'match-and-utilization' framework proposed by Herrmann et al. (2004). Under this account lexical items are pre-activated due to strong top-down predictions based on the context provided by the sentence-level meaning under construction during unification. When the bottom-up linguistic input matches a pre-activated lexical item (or its features) this gives rise to an increase in gamma power. When no pre-activation takes place, or when the linguistic input does not match a prediction, there should be no increase in gamma power (and perhaps even a decrease). Extending these ideas to the framework proposed in Section 2.1, gamma activity in this case reflects resonance between populations of neurons acting as representational and error units that are related to the pre-activated lexical item. This resonance binds information related to the lexical representation while suppressing competing lexical representations (by laterally inhibiting neural populations acting as error units for competing lexical items or features at the same hierarchical processing level). When the input does not match a strong prediction no resonance takes place as a different lexical item needs to be retrieved. Similarly, when no strong prediction (and hence no lexical pre-activation) is present there should also be no resonance (or perhaps delayed resonance) until an appropriate lexical item can be selected and the representational and error units associated with that item can be activated.

This proposal easily accounts for the cases where gamma power increases for highly predicted target words compared to semantic violations (Hald et al., 2006; Penolazzi et al., 2009; Rommers et al., 2013; Wang, Zhu, et al., 2012; Weiss & Mueller, 2003). In the case of highly predicted target words, these lexical items are pre-activated based on the sentence-level meaning. When they match the linguistic input, resonance occurs in the low and/or middle gamma frequency range, binding together representational and error units associated with that lexical representation, while at the same time laterally inhibiting competing error units on the same hierarchical level to suppress competing lexical representations. For semantic violations the linguistic input does not match the pre-activated lexical representation and thus no gamma power (or coherence) increase is present. For the comparison of semantically legal sentences with sentences consisting of syntactic prose (Bastiaansen & Hagoort, 2010), each upcoming word would be more predictable in the case of semantically legal sentences (where construction of a sentence-level meaning representation is possible) than in the case of syntactic prose (where no sentence-level meaning representation from which to make predictions about the next word would be present), and this would result in the observed higher gamma power throughout the sentence for the semantically legal condition compared to the syntactic prose condition. The reported higher gamma power for referentially successful target words (van Berkum et al., 2004) could reflect a match between a prediction about the upcoming referent and the actual linguistic input. For referentially ambiguous and referentially failing target words the input does not match the prediction and so no gamma power increase is observed. People are more

likely to make strong predictions about upcoming words while listening to sentences in their native language than when listening to sentences in a language they do not speak/understand, and this explains the increased gamma power observed when listening to one's native language compared to listening to phonologically related and unrelated languages (Peña & Melloni, 2012).

The cases comparing sentences containing a highly predictable target word with sentences containing a less predictable (but still semantically congruent) target word (Molinaro et al., 2013; Monsalve et al., 2014; Wang, Jensen, et al., 2012; Wang, Zhu, et al., 2012) are also easily explained under this proposal. In the case of the less predictable target words, no (or perhaps less) pre-activation occurs and so there is no pre-activated lexical item (or features) with which the incoming linguistic input could match. The absence of a gamma power (or transient phase-locking) increase thus indicates that representational and error units at that level (and hence all potential lexical representations) are still competing for selection at the time that the linguistic input reaches that level of the processing hierarchy.

The only result we are aware of linking oscillatory activity in the gamma frequency range to sentence-level language comprehension that can't be captured under our predictive coding framework, is the gamma power increase for world-knowledge violations, and the absence of a gamma power increase for semantically legal sentences (Hagoort et al., 2004). We have suggested (based on arguments in Monsalve et al., 2014) that this anomaly may be due to different strategies for the allocation of attention between the studies, as a result of the different composition of the experimental lists (Lewis et al., 2015). Furthermore, a recent study (Metzner, von der Malsburg, Vasishth, & Rösler, 2014) has tried unsuccessfully to replicate the gamma findings from Hagoort et al. (2004). There are many reasons why such a replication might fail (a different analysis method for instance might be less sensitive to the detection of effects in the gamma frequency range), only one of which is that the original result was an incidental finding. We refer to this failed replication mainly to point out that the Hagoort et al. (2004) results are an interesting exception, and merit further exploration before a definitive decision can be made about whether or not they speak against the predictive coding hypothesis we are suggesting in this paper.

We have also suggested that oscillatory activity in the high gamma frequency range might reflect the bottom-up propagation of prediction errors to higher levels of the cortical processing hierarchy during sentence-level language comprehension. As far as we are aware there are no studies directly investigating this issue at the processing levels we have been discussing. There is however a large body of excellent work investigating the hierarchical flow of information within auditory cortical regions while processing speech (e.g., Arnal & Giraud, 2012; Arnal, Wyart, & Giraud, 2011; Giraud & Poeppel, 2012). These studies have shown that bottom-up information conveying prediction errors to higher processing levels within the auditory system is transmitted via evoked oscillatory activity in the high gamma range (e.g., Fontolan et al., 2014). It has also been suggested that these principles may extend to even higher levels of the cortical processing hierarchy (beyond the auditory system) during speech comprehension (Giraud & Poeppel, 2012). It is not yet clear whether this is indeed the case, nor whether a similar argument might extend to language comprehension with visual input, but the general notion fits well within the framework we are proposing here. Further evidence and more experiments investigating high gamma oscillations during sentence-level language comprehension (especially for visual input where the least is known) are clearly necessary.

3. Some suggestions for future research

To recap briefly, we have suggested a role for oscillatory activity in the beta frequency range during sentence-level language comprehension in the active maintenance or change of an NCN responsible for representing the current sentence-level meaning under construction. We have also implicated beta in providing top-down predictions to lower hierarchical levels on the basis of that sentence-level meaning. Our proposal links oscillatory activity in the low and middle gamma frequency range to a match between pre-activated lexical items and incoming linguistic input. We have tentatively suggested that high gamma reflects the bottom-up propagation of prediction errors to higher hierarchical processing levels. What we know least about so far is the role of high gamma in sentence-level language comprehension (beyond the auditory processing system), and whether or not low and middle gamma actively inhibit competing representations after a match is made between top-down predictions and bottom-up linguistic input, or mainly reflects resonance due to that match. We have already proposed a few ideas for further testing our prediction/maintenance hypothesis (Lewis et al., 2015) and below we will suggest a few more, along with some possibilities for testing some of the additional aspects (just described) introduced here under a predictive coding framework.

One consequence of the proposal is that any time a strong prediction about some upcoming lexical item matches the linguistic input there should be a gamma power increase at the associated hierarchical level(s). It might be interesting to manipulate the degree of predictability that is necessary before resonance occurs in the gamma range. We could construct stimuli with different levels of predictability and see what the effect is on gamma oscillations. In this regard it would be useful to have a reliable method for quantifying predictability. Smith and Levy (2013) have shown that the predictability of a word in a particular sentence context is logarithmically related to reading times over six orders of magnitude. This suggests that cloze probability scores are not a good measure for the effects of word predictability in a sentence context on processing difficulty (especially when attempting to experimentally match the levels of predictability of unexpected words). Instead we might try using surprisal measures (in this setting, a measure of a word's negative log-transformed conditional probability based on the preceding sentence context; e.g., Levy, 2008; however, surprisal measures can be applied to other linguistic units one finds relevant to investigate, e.g., parts of speech, and can even be modified to account for predictions about

particular syntactic structures; cf. Levy, 2011) estimated based on large data corpora using computational models (e.g., Frank, Otten, Galli, & Vigliocco, 2013; Smith & Levy, 2013). It was recently shown that word surprisal is a significant predictor for the amplitude of the N400 (Frank et al., 2013), with more surprising words giving rise to larger N400 amplitudes. This might constitute a good measure by which to manipulate the predictability of an upcoming word and hence to investigate just how predictable a word needs to be in some sentence context in order for pre-activation of its lexical representation to occur and for the matching linguistic input to give rise to gamma resonance between representational and error units at the same hierarchical processing level. In fact it may not even be necessary to experimentally manipulate predictability, as surprisal values will generally be quite variable over the course of natural texts and it may be possible to show a relationship between these values and gamma power using correlation-based methods.

An aspect of oscillatory dynamics that has not really been addressed during sentence-level language comprehension (although there is some work on this for the auditory system; e.g., Fontolan et al., 2014) is directed interactions between different levels of the cortical processing hierarchy, using measures like Granger Causality (e.g., Fontolan et al., 2014) or Dynamic Causal Modeling (e.g., Friston et al., 2015). Our proposal makes clear predictions about the direction of information flow during language comprehension and this could be an interesting avenue to explore using tasks that are expected to elicit different levels of top-down and bottom-up information flow. Semantic or syntactic violations represent such a case, because after a violation the language comprehension system should shift the emphasis to a more bottom-up mode of processing and rely less on (apparently unreliable) top-down predictions based on the sentence-level meaning. This implies that we could use the kind of well-established paradigms already discussed in Sections 2.2 and 2.3, and just focus our analyses more on the direction of information flow between hierarchical processing levels rather than the power or coherence of neural oscillations at different frequencies.

If NCNs are maintained via beta oscillations in the support and construction of the current sentence-level meaning, a question arises about exactly how far they might extend in terms of the information that is processed/represented within such networks. Do they represent only the current sentence-level meaning (in which case other information that impinges on the construction of such a meaning representation, like discourse information or other types of visual or auditory information that might be relevant in that context, would need to be represented elsewhere and only have an effect on the NCN) or do they also incorporate representations of previous discourse information and relevant auditory and visual contextual information for instance? Our account suggests that we might be able to probe these questions by manipulating various sources of information and checking how this affects oscillatory activity in the beta frequency range. We could for instance compare sentences embedded within very rich discourse contexts with the same sentences embedded within discourse contexts that do not

provide much additional information than the sentences themselves, and see how beta is modulated. It has been shown that discourse-level information can have an effect on sentence-level processing (see e.g., Nieuwland & van Berkum, 2006), so it may be the case that this discourse contextual information results in additional nodes being incorporated into the NCN responsible for processing the current sentence-level meaning. If that is the case we might expect to see modulations of beta power within this NCN, or potentially more widely spatially distributed beta synchrony. Alternatively we could vary visual or auditory contextual information (by presenting pictures or sounds) while participants read/listen to sentences and see how beta is modulated. In that case we would of course need appropriate controls for the addition of visual or auditory properties of the pictures/sounds.

Another consequence of the idea that decreases in beta power may reflect a change in the current NCN is that beta power might be negatively correlated with the amplitude of the P600 ERP component. One interpretation of the P600 has been as an index of processes of repair or reanalysis (e.g., Friederici, 2002; Kaan, 2002) and in such cases the current NCN would presumably need to change resulting in decreased beta power. Davidson and Indefrey (2007) did not find any relationship between beta power and the amplitude of the P600 in their study containing syntactic violations, but they were not able to clearly separate the alpha and beta frequency ranges in their statistical analysis and there was a clear negative relationship between alpha power and P600 amplitude. We therefore think that further investigation of this relationship between beta power and P600 amplitude is warranted.

One way to probe the idea that lexical competitors are actively suppressed by resonance at low and middle gamma frequencies once a match is made between a pre-activated lexical item and the bottom-up linguistic input, is to use target words that are highly predicted but differ systematically in the number of potential competitors they have at that position in the sentence. For more competitors there should be more lateral inhibition of the error units for competing lexical representations and this should result in higher gamma power within that processing level compared to words with fewer competitors.

Since we have suggested that high gamma oscillations might be an index of the bottom-up propagation of prediction errors, we should test whether high gamma is present in the first place during sentence-level language comprehension, as well as under what conditions it is present and at (or between) which levels of the processing hierarchy. If high gamma is really an index of prediction error being sent up the processing hierarchy we might expect that it should increase for less predictable input (larger prediction error), and decrease for more predictable input (smaller prediction error). Another factor that could result in higher gamma power in this context is a shift in processing strategy to focus more on bottom-up information during comprehension (after for example a semantic or a syntactic violation). It may also be the case that prediction errors are 'turned off' (error units at that level do not send information to higher processing levels and so there should be little or no high

gamma) when there is a match between a strong prediction and the input. All of these avenues of investigation related to high gamma are interesting and deserve further experimental exploration.

We would also like to point out that we are aware of only a handful of studies (Magyari et al., 2014; Wang, Jensen, et al., 2012) investigating neural oscillations during sentence-level language comprehension that have employed source reconstruction techniques in order to investigate the cortical sources of the effects found with greater spatial precision. We think that the use of such methods could be highly beneficial, and a move in this direction will be extremely important if we are to properly investigate some of the predictive coding ideas we have been advocating here.

4. Conclusions

In this paper we have suggested how the extant findings relating oscillatory neural dynamics in the beta and gamma frequency ranges to sentence-level language comprehension may be given a unified explanation under a predictive coding framework. We have proposed that beta activity reflects both the active maintenance of the current NCN responsible for the construction and representation of a sentence-level meaning, and the top-down propagation of predictions based on that meaning to lower levels of the processing hierarchy. Low and middle gamma synchrony is proposed to be related to a match between a highly predicted target word and bottom-up linguistic input, and to reflect resonance between representational and error units (related to the activated lexical item) at the current hierarchical level, as well as lateral inhibition of other error units (suppression of activity related to competing lexical representations) at that level. High-gamma on the other hand, we suggest might reflect the bottom-up propagation of prediction errors to higher hierarchical levels.

Our new predictive coding framework for sentence-level language comprehension is largely supported by the available empirical evidence, and offers clear and experimentally testable predictions. We have suggested a few concrete ideas for experiments that might test some of these predictions. In sum, we think that our proposal provides a useful framework within which to understand how the cognitive system responsible for sentence-level language comprehension dynamically recruits various neural networks responsible for the necessary representations and computations underlying sentence-level meaning construction, and how this is reflected in the measured oscillatory neural dynamics.

Acknowledgments

REFERENCES

Altmann, G. T., & Kamide, Y. (1999). Incremental interpretation at verbs: restricting the domain of subsequent reference. *Cognition, 73*(3), 247–264.

Altmann, G. T. M., & Kamide, Y. (2007). The real-time mediation of visual attention by language and world knowledge: linking anticipatory (and other) eye movements to linguistic processing. *Journal of Memory and Language, 57*(4), 502–518.

Arnal, L. H., & Giraud, A. L. (2012). Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences, 16*(7), 390–398.

Arnal, L. H., Wyart, V., & Giraud, A. L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience, 14*(6), 797–801.

Baggio, G., & Hagoort, P. (2011). The balance between memory and unification in semantics: a dynamic account of the N400. *Language and Cognitive Processes, 26*(9), 1338–1367.

Balota, D. A., Pollatsek, A., & Rayner, K. (1985). The interaction of contextual constraints and parafoveal visual information in reading. *Cognitive Psychology, 17*(3), 364–390.

Bastiaansen, M. C. M., & Brunia, C. H. (2001). Anticipatory attention: an event-related desynchronization approach. *International Journal of Psychophysiology, 43*(1), 91–107.

Bastiaansen, M. C. M., & Hagoort, P. (2006). Oscillatory neuronal dynamics during language comprehension. *Progress in Brain Research, 159*(06), 179–196.

Bastiaansen, M., & Hagoort, P. (2010). Frequency-based segregation of syntactic and semantic unification. In *16th Annual Meeting of the Organization for Human Brain Mapping, Barcelona, Spain*.

Bastiaansen, M. C. M., Magyari, L., & Hagoort, P. (2010). Syntactic unification operations are reflected in oscillatory dynamics during on-line sentence comprehension. *Journal of Cognitive Neuroscience, 22*(7), 1333–1347.

Bastiaansen, M. C. M., Mazaheri, A., & Jensen, O. (2012). Beyond ERPs: oscillatory neuronal dynamics. In S. Luck, & E. Kappenman (Eds.), *Oxford handbook of event-related potential components* (pp. 31–49). New York: Oxford University Press.

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron, 76*(4), 695–711.

Bressler, S. L., & Richter, C. G. (2015). Interareal oscillatory synchronization in top-down neocortical processing. *Current Opinion in Neurobiology, 31C*, 62–66.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences, 36*(3), 181–204.

Davidson, D. J., & Indefrey, P. (2007). An inverse relation between event-related and time-frequency violation responses in sentence processing. *Brain Research, 1158*, 81–92.

DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience, 8*(8), 1117–1121.

Ehrlich, S., & Rayner, K. (1981). Contextual effects on word perception and eye movements during reading. *Journal of Verbal Learning and Verbal Behavior, 20*, 641–655.

Engel, A. K., & Fries, P. (2010). Beta-band oscillations—signalling the status quo? *Current Opinion in Neurobiology, 20*(2), 156–165.

Federmeier, K., & Kutas, M. (1999). A rose by any other name: long-term memory structure and sentence processing. *Journal of Memory and Language, 495*, 469–495.

Fontolan, L., Morillon, B., Liegeois-Chauvel, C., & Giraud, A. L. (2014). The contribution of frequency-specific activity to hierarchical information processing in the human auditory cortex. *Nature Communications, 5*(4694).

Forster, K. I. (1981). Priming and the effects of sentence and lexical contexts on naming time: evidence for autonomous

lexical processing. *The Quarterly Journal of Experimental Psychology Section A, 33*(4), 465–495.

Frank, S., Otten, L., Galli, G., & Vigliocco, G. (2013). Word surprisal predicts N400 amplitude during reading. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics* (pp. 878–883).

Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences, 6*(2), 78–84.

Friederici, A. D., & Mecklinger, A. (1996). Syntactic parsing as revealed by brain responses: first-pass and second-pass parsing. *Journal of Psycholinguistic Research, 25*(1), 157–176.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 360*(1456), 815–836.

Friston, K. J., Bastos, A. M., Pinotsis, D., & Litvak, V. (2015). LFP and oscillations-what do they tell us? *Current Opinion in Neurobiology, 31C*, 1–6.

Giraud, A., & Poeppel, D. (2012). Cortical oscillations and speech processing: emerging computational principles and operations. *Nature Neuroscience, 15*(4), 511–517.

Gray, C., König, P., Engel, A., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Letters to Nature, 338*, 334–337.

Hagoort, P. (2005). On Broca, brain, and binding: a new framework. *Trends in Cognitive Sciences, 9*(9), 416–423.

Hagoort, P. (2013). MUC (Memory, Unification, Control) and beyond. *Frontiers in Psychology, 4*(416).

Hagoort, P. (2014). Nodes and networks in the neural architecture for language: Broca's region and beyond. *Current Opinion in Neurobiology, 28*, 136–141.

Hagoort, P., Baggio, G., & Willems, R. (2009). Semantic unification. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 819–836). Cambridge, MA: MIT Press.

Hagoort, P., Brown, C., & Groothusen, J. (1993). The syntactic positive shift as an ERP-measure of syntactic processing. *Language and Cognitive Processes, 8*, 439–483.

Hagoort, P., Hald, L. A., Bastiaansen, M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science, 304*(5669), 438–441.

Hald, L. A., Bastiaansen, M. C. M., & Hagoort, P. (2006). EEG theta and gamma responses to semantic violations in online sentence processing. *Brain and Language, 96*(1), 90–105.

Herrmann, C. S., Munk, M. H. J., & Engel, A. K. (2004). Cognitive functions of gamma-band activity: memory match and utilization. *Trends in Cognitive Sciences, 8*(8), 347–355.

Jung-Beeman, M. (2005). Bilateral brain processes for comprehending natural language. *Trends in Cognitive Sciences, 9*(11), 512–518.

Kaan, E. (2002). Investigating the effects of distance and number interference in processing subject-verb dependencies: an ERP study. *Journal of Psycholinguistic Research, 31*(2), 165–193.

Kamide, Y. (2008). Anticipatory processes in sentence processing. *Language and Linguistics Compass, 4*, 647–670.

Kamide, Y., Altmann, G. T. M., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: evidence from anticipatory eye movements. *Journal of Memory and Language, 49*(1), 133–156.

Kamide, Y., Scheepers, C., & Altmann, G. T. M. (2003). Integration of syntactic and semantic information in predictive processing: cross-linguistic evidence from German and English. *Journal of Psycholinguistic Research, 32*(1), 37–55.

Kielar, A., Meltzer, J., Moreno, S., Alain, C., & Bialystok, E. (2014). Oscillatory responses to semantic and syntactic violations. *Journal of Cognitive Neuroscience*, 1–23.

Knoeferle, P., Crocker, M. W., Scheepers, C., & Pickering, M. J. (2005). The influence of the immediate visual context on incremental thematic role-assignment: evidence from eye-movements in depicted events. *Cognition, 95*(1), 95–127.

König, P., & Schillen, T. (1991). Stimulus-dependent assembly formation of oscillatory responses: I. Synchronization. *Neural Computation, 3*, 155–166.

Kutas, M., & Federmeier, K. D. (2011). Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP). *Annual Review of Psychology, 62*, 621–647.

Lau, E. F., Phillips, C., & Poeppel, D. (2008). A cortical network for semantics: (de)constructing the N400. *Nature Reviews Neuroscience, 9*(12), 920–933.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition, 106*(3), 1126–1177.

Levy, R. (2011). Integrating surprisal and uncertain-input models in online sentence comprehension: formal techniques and empirical results. In *49th Annual Meeting of the Association for Computational Linguistics, Portland Oregon, USA*.

Lewis, A. G., Wang, L., & Bastiaansen, M. C. M. (2015). Fast oscillatory dynamics during language comprehension: unification or maintenance/prediction? *Brain and Language*. http://dx.doi.org/10.1016/j.bandl.2015.01.003 (advance online publication).

Lisman, J. E., & Jensen, O. (2013). The theta-gamma neural code. *Neuron, 77*, 1002–1016.

Luo, Y., Zhang, Y., Feng, X., & Zhou, X. (2010). Electroencephalogram oscillations differentiate semantic and prosodic processes during sentence reading. *Neuroscience, 169*(2), 654–664.

MacDonald, M., Pearlmutter, N., & Seidenberg, M. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*(4), 676–703.

Magyari, L., Bastiaansen, M. C. M., de Ruiter, J. P., & Levinson, S. C. (2014). Early anticipation lies behind the speed of response in conversation. *Journal of Cognitive Neuroscience, 26*(11), 2530–2539.

Maier, A., Adams, G. K., Aura, C., & Leopold, D. A. (2010). Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation. *Frontiers in Systems Neuroscience, 4*(31).

McRae, K., Hare, M., Elman, J., & Ferretti, T. (2005). A basis for generating expectancies for verbs from nouns. *Memory & Cognition, 33*(7), 1174–1184.

Mecklinger, A., Schriefers, H., Steinhauer, K., & Friederici, A. D. (1995). The processing of relative clauses varying on syntactic and semantic dimensions: an analysis with event-related potentials. *Memory & Cognition, 23*, 477–494.

Metzner, P., von der Malsburg, T., Vasishth, S., & Rösler, F. (2014). Brain responses to world knowledge violations: a comparison of stimulus-and fixation-triggered event-related potentials and neural oscillations. *Journal of Cognitive Neuroscience*, 1–12.

Meyer, L., Obleser, J., & Friederici, A. D. (2013). Left parietal alpha enhancement during working memory-intensive sentence processing. *Cortex, 49*(3), 711–721.

Molinaro, N., Barraza, P., & Carreiras, M. (2013). Long-range neural synchronization supports fast and efficient reading: EEG correlates of processing expected words in sentences. *NeuroImage, 72*, 120–132.

Monsalve, I. F., Pérez, A., & Molinaro, N. (2014). Item parameters dissociate between expectation formats: a regression analysis of time-frequency decomposed EEG data. *Frontiers in Psychology, 5*(847).

Nieuwland, M. S., & van Berkum, J. J. A. (2006). When peanuts fall in love: N400 evidence for the power of discourse. *Journal of Cognitive Neuroscience, 18*(7), 1098–1111.

Osterhout, L., Holcomb, P. J., & Swinney, D. A. (1994). Brain potentials elicited by garden-path sentences: evidence of the application of verb information during parsing. *Journal of Experimental Psychology: Learning, Memory and Cognition, 20*, 786–803.

Otten, M., Nieuwland, M. S., & Van Berkum, J. J. A. (2007). Great expectations: specific lexical anticipation influences the processing of spoken language. *BMC Neuroscience, 8*(89).

Parkes, L. M., Bastiaansen, M. C. M., & Norris, D. G. (2006). Combining EEG and fMRI to investigate the post-movement beta rebound. *NeuroImage, 29*(3), 685–696.

Peña, M., & Melloni, L. (2012). Brain oscillations during spoken sentence processing. *Journal of Cognitive Neuroscience, 24*(5), 1149–1164.

Penolazzi, B., Angrilli, A., & Job, R. (2009). Gamma EEG activity induced by semantic violation during sentence reading. *Neuroscience Letters, 465*(1), 74–78.

Pérez, A., Molinaro, N., Mancini, S., Barraza, P., & Carreiras, M. (2012). Oscillatory dynamics related to the unagreement pattern in Spanish. *Neuropsychologia, 50*(11), 2584–2597.

Pfurtscheller, G., & Lopes da Silva, F. (1999a). Event-related desynchronization. In *Handbook of electroencephalography and clinical neurophysiology, revised series* (Vol. 6). Amsterdam: Elsevier.

Pfurtscheller, G., & Lopes da Silva, F. H. (1999b). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical Neurophysiology, 110*(11), 1842–1857.

Pfurtscheller, G., Neuper, C., Andrew, C., & Edlinger, G. (1997). Foot and hand area mu rhythms. *International Journal of Psychophysiology, 26*(1–3), 121–135.

Pfurtscheller, G., Stancák, A., & Neuper, C. (1996). Post-movement beta synchronization. A correlate of an idling motor area? *Electroencephalography and Clinical Neurophysiology, 98*(4), 281–293.

Pfurtscheller, G., Zalaudek, K., & Neuper, C. (1998). Event-related beta synchronization after wrist, finger and thumb movement. *Electroencephalography and Clinical Neurophysiology, 109*(2), 154–160.

Pickering, M. J., & Garrod, S. (2007). Do people use language production to make predictions during comprehension? *Trends in Cognitive Sciences, 11*(3), 105–110.

Rayner, K., Pollatsek, A., Ashby, J., & Clifton, C., Jr. (2012). *Psychology of reading* (2nd ed.). New York: Psychology Press.

Rommers, J., Dijkstra, T., & Bastiaansen, M. C. M. (2013). Context-dependent semantic processing in the human brain: evidence from idiom comprehension. *Journal of Cognitive Neuroscience*, 762–776.

Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., et al. (2008). Period concatenation underlies interactions between gamma and beta rhythms in neocortex. *Frontiers in Cellular Neuroscience, 2*(1).

Roopun, A. K., Middleton, S. J., Cunningham, M. O., LeBeau, F. E. N., Bibbig, A., Whittington, M. A., et al. (2006). A beta2-frequency (20–30 Hz) oscillation in nonsynaptic networks of somatosensory cortex. *Proceedings of the National Academy of Sciences of the United States of America, 103*(42), 15646–15650.

Schwanenflugel, P. J., & Lacount, K. L. (1988). Semantic relatedness and the scope of facilitation for upcoming words in sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(2), 344–354.

Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: some limitations of knowledge-based processing. *Cognitive Psychology, 14*(4), 489–537.

Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience, 13*(2), 121–134.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing. *Annual Review of Physiology, 55*, 349–374.

Singer, W. (2011). Dynamic formation of functional networks by synchronization. *Neuron, 69*(2), 191–193.

Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. *Cognition, 128*(3), 302–319.

Sussman, R. S., & Sedivy, J. (2003). The time-course of processing syntactic dependencies: evidence from eye movements. *Language and Cognitive Processes, 18*(2), 143–163.

Szewczyk, J. M., & Schriefers, H. (2013). Prediction in language comprehension beyond specific words: an ERP study on sentence comprehension in Polish. *Journal of Memory and Language, 68*(4), 297–314.

van Berkum, J. J. A., Brown, C. M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating upcoming words in discourse: evidence from ERPs and reading times. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 31*(3), 443–467.

van Berkum, J. J. A., Zwitserlood, P., Bastiaansen, M., Brown, C., & Hagoort, P. (2004). *So who's "he" anyway? Differential ERP and ERSP effects of referential success, ambiguity and failure during spoken language comprehension.* Paper presented at the Annual Meeting of the Cognitive Neuroscience Society, San Francisco.

van den Brink, D., Brown, C. M., & Hagoort, P. (2001). Electrophysiological evidence for early contextual influences during spoken-word recognition: N200 versus N400 effects. *Journal of Cognitive Neuroscience, 13*(7), 967–985.

van Petten, C., Coulson, S., Rubin, S., Plante, E., & Parks, M. (1999). Time course of word identification and semantic integration in spoken language. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 25*(2), 394–417.

van Petten, C., & Luka, B. J. (2012). Prediction during language comprehension: benefits, costs, and ERP components. *International Journal of Psychophysiology, 83*(2), 176–190.

Varela, F., Lachaux, J., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience, 2*(4), 229–239.

Wang, X. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews, 90*, 1195–1268.

Wang, L., Jensen, O., van den Brink, D., Weder, N., Schoffelen, J., Magyari, L., et al. (2012). Beta oscillations relate to the N400m during language comprehension. *Human Brain Mapping, 33*(12), 2898–2912.

Wang, L., Zhu, Z., & Bastiaansen, M. C. M. (2012). Integration or predictability? A further specification of the functional role of gamma oscillations in language comprehension. *Frontiers in Psychology, 3*(187).

Weiss, S., & Mueller, H. M. (2003). The contribution of EEG coherence to the investigation of language. *Brain and Language, 85*(2), 325–343.

Weiss, S., Mueller, H. M., Schack, B., King, J. W., Kutas, M., & Rappelsberger, P. (2005). Increased neuronal communication accompanying sentence comprehension. *International Journal of Psychophysiology, 57*(2), 129–141.

Wicha, N. Y. Y., Moreno, E. M., & Kutas, M. (2004). Anticipating words and their gender: an event-related brain potential study of semantic integration, gender expectancy, and gender agreement in Spanish sentence reading. *Journal of Cognitive Neuroscience, 16*(7), 1272–1288.

Womelsdorf, T., Schoffelen, J., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., et al. (2007). Modulation of neuronal interactions through neuronal synchrony. *Science, 316*(5831), 1609–1612.

Zwitserlood, P. (1989). The locus of the effects of sentential-semantic context in spoken-word processing. *Cognition, 32*, 25–64.