

On the Persistence of Structural Priming: Mechanisms of Decay and Influence of Word-Forms

Gaurav Malhotra (gaurav.malhotra@ed.ac.uk)
Martin Pickering (martin.pickering@ed.ac.uk)
Holly Branigan (holly.branigan@ed.ac.uk)
James A. Bednar (james.bednar@ed.ac.uk)
University of Edinburgh, Edinburgh, EH8 9JZ, UK

Abstract

When people speak, they have to choose a syntactic structure for their utterances. Research shows that this choice is not made independently for each utterance, and that it is influenced by properties of recent utterances (Bock, 1986). But the learning mechanisms responsible for this influence, and its duration are not completely clear. Pickering and Branigan (1998) suggest that structural repetition is due to trailing-activation in the language production system. However, this account lacks a formal implementation. Chang, Dell, and Bock (2006) have developed a computational model that tries to explain structural priming based on error-based learning. In this paper we report a formal model that relies on an unsupervised learning mechanism and use this to replicate behavioral data for both structural priming and lexical influence on this priming. We use this model to investigate the factors that lead to these priming effects and its rate of decay.

Introduction

A key goal of cognitive science is to understand the processes underlying language production. This entails understanding the mechanisms underlying lexical, structural and semantic choice during sentence formulation. Research shows that the choice of syntactic structure is not made independently for each episode of language production. Rather, recent utterances that the speaker has heard, or produced, influence this choice – i.e. language production shows structural priming (Bock, 1986; Branigan, Pickering, & Cleland, 2000).

Furthermore, Pickering and Branigan (1998) found that structural priming increases when speakers use the same verbs in prime and target utterances. Consider the following Prime and Target pairs for a sentence-completion task:

Prime1 The teacher gave the girl the book.

Prime2 The teacher gave the book to the girl.

Prime3 The teacher showed the girl the book.

Prime4 The teacher showed the book to the girl.

Target The patient gave ...

The prime sentences use either a direct-object (DO) or prepositional-object (PO) syntactic structure. The target sentence can also be completed using either a DO or a PO phrase. Pickering and Branigan (1998) show that subjects are more likely to repeat the structure of prime sentence if the prime and target use the same verb – i.e. Prime1 and Prime2 elicit more priming than Prime3 and Prime4. This phenomenon is called the *lexical-boost* effect.

There are two existing accounts of structural priming. The *trailing-activation* account attributes priming to unsupervised, associative learning which leads to traces of activation

in the system. Pickering and Branigan (1998) proposed that this trailing activation links structural options to individual lexical items, causing lexical context to influence structural priming (thereby explaining the *lexical-boost* effect).

The alternative *error-based learning* account is a supervised learning account. It proposes that listeners actively predict what they will hear and use the error between prediction and actual outcome to adjust their structural-decision rules; this error-based learning is hypothesized to underlie structural priming (Chang et al., 2006). Such an account therefore emphasises that structural repetition is a byproduct of a larger function of human cognition: language acquisition.

Both *error-based* and *trailing-activation* accounts are able to explain data from structural priming. Since the latter account is based on learning associations between lexical and structural constructs, it is more compatible with the data showing lexical boost. But the trailing-activation account lacks a formal implementation. Also, some recent studies have found no lexical boost in long-term structural priming (Kaschak & Borreggine, 2008; Hartsuiker, Bernolet, Schoonbaert, Speybroeck, & Vanderelst, 2008) leading them to conclude that listeners store structural information in long-term memory after discarding its lexical context. These data are consistent with Chang et al's model, which assumes that an independent *syntax module* abstracts structural information from utterances through error-based learning.

In contrast, we present a nonlinear dynamical system that explains such data in terms of trailing activation. This computational model does not rely on a strategic learning process. Instead, each episode of training leaves a memory trace based on the units that it activates, recorded as a *fixed* amount of adjustment to the system. The structural choices made by the system during subsequent episodes are dependent on these memory traces. In addition to this, the architecture of the system assumes a considerable degree of similarity in lexical and structural information processing. As a consequence, it allows these kinds of processes to closely inform each other.

We show that this system can successfully model short-term effects of both structural priming and its lexical-enhancement. We also use this system to model results from Kaschak and Borreggine (2008), and Hartsuiker et al. (2008) – two studies which fail to find lexical-enhancement on long-term structural priming. We use these results to investigate the time-course of structural priming and the mechanisms by which lexical representations can influence this priming.

Methods

Network architecture

The model consists of three layers, each of which can be thought of as independent cognitive module in the brain (see 1). Syntactic and lexical processing are performed by the two layers of representation *Layer1* and *Layer2* in the figure. Each of these layers consists of a group of nodes representing a basic construct of the layer: verbs for lexical layer and grammatical constructions for the syntax layer. In addition to this, the system has a layer consisting of binding nodes that bind representations in lexical and syntactic layers. This layer can be thought of as a cognitive module providing an activation-based short-term memory (STM) for associations between the lexical and syntactic modules. Such an activation-based memory is coherent with the notion of cortico-cortical associations being maintained by persistent activation of neurons in the prefrontal cortex (Funahashi, Bruce, & Goldman-Rakic, 1989).

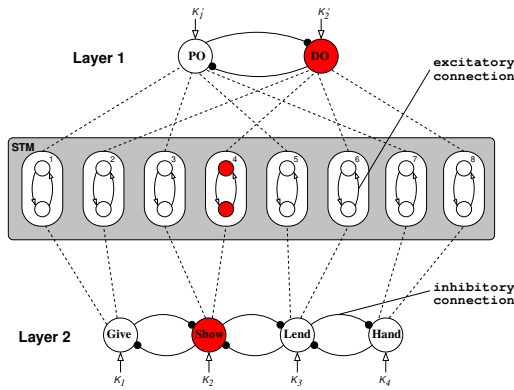


Figure 1: The model consists of two winner-take-all layers, connected via binding nodes. Activation flows across layers (dotted lines) according to the logic discussed in Experimental Design.

Each node in layers 1 and 2 receives an external input, K_n . During comprehension, this input is assumed to be from feed-forward connections coming from either lower levels of processing or afferent connections from the sensory system. During production, nodes in the lexical layer receive input from the meaning system, while the syntax nodes receive no external input, except the activation that flows from lexical layer via the binding nodes (discussed later).

Dynamics

Each node sums its input and produces an output governed by the Naka-Rushton transformation function (Naka & Rushton, 1966):

$$S(x) = \begin{cases} \frac{Mx^N}{\sigma^N + x^N} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases} \quad (1)$$

where x is the input to the function, M is the maximum activation, σ is semi-saturation constant (the point at which $S(x)$ reaches half its value), and N determines the slope of the function.

Each node is connected to all other nodes in the layer through mutually-inhibitory connections. This ensures that each layer has **winner-take-all** (WTA) dynamics: the nodes of a particular layer compete with each other for achieving maximum activation; the winning node then suppresses the other nodes completely. The strength of mutual inhibition is predefined and fixed. The dynamic equation for the rate of change of activation of each node in *Layer1* and *Layer2* is:

$$\frac{dE_i}{dt} = \frac{1}{\tau}(-E_i + S(K_i - 3 \sum_{j \neq i} E_j)) \quad (2)$$

where E_i is the activation of the node, K_i is the external input, and τ is the time-constant for rate of change of activation of each node (Wilson, 1999). Therefore, $(K_i - 3 \sum_j E_j)$ is the net input to each node – the difference between external input and sum of inputs from inhibitory connections. Each external input tries to pull the activation of the node towards its own value and each node tries to suppress all other nodes. Since the system always starts from the *rest-state* ($E_i = 0$) at the beginning of a prime trial, the winning-node at the end of this trial is completely dependent on the external input.

The dynamical equation of the binding nodes is similar, except there are only two binding nodes in each group and they are mutually excitatory:

$$\frac{dE_i}{dt} = \frac{1}{\tau_{stm}}(-E_i + S(3E_j)) \quad (3)$$

Once activated, a node of such a WTA or STM network shows reluctance to move from that stable state – i.e. it shows *hysteresis* (Wilson, 1999). This hysteresis serves as a short-term memory in the network as the activation of the node (during the target trial) is governed not only by its current input, but also by its history of activation.

Learning. There are two mechanisms for recording episodic memory traces in the system: (i) hysteresis in the nodes – which serves as a short-term memory for lexical, syntactic and binding nodes – and (ii) an incremental adjustment to the inputs of the winning node. Whenever a node wins a competition, its input connection, K_{int} , receives a fixed amount of boost, ρ . This boost means that the system favours the activation of a recently activated (or frequently activated) node. Thus every time a node gets activated, it becomes easier to activate it the next time. This mechanism of incremental adjustment is Hebbian learning in its most primitive form: when an input frequently contributes to the firing of a particular neuron, then synapses from the input to the neuron should be strengthened. It should be noted that, as against an error-based learning account, learning in this account is fixed and unsupervised.

Forgetting. Any capacity-limited memory needs to undergo gradual forgetting in order to avoid losing all stored information (Sandberg, Lansner, Petersson, & Ekeberg, 2000). Activation-based memory has a natural mechanism of forgetting: adaptation in the firing rate of neurons, due to fatigue.

Following Wilson (1999), this adaptation is introduced by gradually changing parameter σ in equation 1. Our model consists of two kinds of connections – mutually inhibitory (WTA), and mutually excitatory (STM). Each of these dynamical systems (given by Equations 2 and 3) is extended by replacing σ_i with $(\sigma_i + A_i)$, and adding the equation:

$$\frac{dA_i}{dt} = \frac{1}{\tau_a}(-A_i + \alpha E_i) \quad (4)$$

where A_i is the amount of adaptation (or adjustments to σ_i); τ_a is the time constant for adaptation; and α is the saturation constant for A_i – i.e., it governs what value will A_i increase to, as a fraction of E_i . Effectively, this adaptation mechanism moves the equilibrium point of the dynamical system such that the firing rate decreases, and the stable point moves towards a saddle state. Beyond a particular time, the equilibrium point meets the saddle point and vanishes, thus simulating complete forgetting.

The weight-based long-term memory also needs to exhibit forgetting for the reasons stated above. This forgetting is implemented as an exponential decay in the input, K_{im} , at the end of each trial. We can control the rate of decay by varying the time-constant of this exponential decay.

Experimental design

Two different sets of experiments were carried out on the model. The first set tests the short-term effects of structural repetition (Experiments 1 & 2). In these cases the target immediately follows the prime. These experiments mirror the behavioural experiments conducted by Bock (1986), and Pickering and Branigan (1998). The second set of experiments investigate the long-term effects of structural priming and the influence of lexical repetition on structural priming when prime and target are separated over several trials (Experiments 3 & 4). These experiments mirror Kaschak and Borreggine (2008) and Hartsuiker et al. (2008).

Short-term memory. The simulation is divided into two trials: the *priming trial* and the *target trial*. At the beginning of the priming trial the network was set to the *rest state*, i.e., both nodes in the syntax layer were turned to the OFF state. In equation 2 this corresponds to the initial conditions $E_i = 0$. The external input to each node, K_i is chosen based on the prime sentence. After this, the equilibrium state is calculated by simulating the dynamical equations of both layers. Once the network has settled, the bindings are stored in the STM. This is done by simulating the dynamical equations of the mutually excitatory networks and a constant external input. This external input is positive (and above threshold) if both the lexical and structural nodes connected to the binding node are active (ON), and zero otherwise.

In order to simulate the target trial of a sentence completion task, the lexical layer is provided with an external input (biased in favour of one node) during the target trial. This is equivalent to providing an incomplete sentence with a verb in the behavioural experiment. The syntactic nodes, on the

other hand, do not receive any external input. This is because we want to simulate syntactic choice and the decision has to be based on the internal states of the system. The initial state of these nodes is set to the final state at the end of training trial. Also each node in the syntactic layer receives input from STM based on the following expression:

$$\begin{cases} K_{stm} & \text{if } E_{stm} \times E_{lex} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where K_{stm} is a constant amount of input that a node in layer 1 gets from the STM, provided the condition on the right-hand-side is met. The condition $E_{stm} \times E_{lex} > 0$ ensures that a syntactic node receives input from only those binding nodes that are themselves active *and* are connected to active lexical nodes. Thus, during the target trial, there is a causal flow of activation from lexical to syntax layer via the binding nodes.

Long-term memory. The second set of tests conducted on the model tested the persistence of structural priming and long-term lexical enhancement of this priming. In order to make direct comparisons with behavioural data, we tested the model under same experimental conditions as Kaschak and Borreggine (2008). Each of their experiments is divided into two phases: a training phase, and a testing phase¹ (Figure 2). Subjects are trained for a particular grammatical construction during the training phase by being coerced to produce it. These are shown as the ‘prime’ trials in Figure 2. We simulated this by providing high input to the coerced nodes for both lexical and syntactic constructions and simulating the model till it reaches equilibrium. After 10 such priming trials, subjects receive 6 prime-target pairs during the testing phase. Target trials are simulated in exactly the same way as for short-term memory experiments above. The model is allowed to run freely and settle into a state of equilibrium. The winning node in the syntax layers is taken as the output of the production phase.




Figure 2: Experimental design for testing long-term priming and lexical influence.

¹The testing phase does not imply that the model has stopped learning – it performs learning during both phases, but is tested on target sentences only during testing phase.

Results and Analysis

Experiment 1. The first experiment was a short-term memory experiment – i.e. target trials were presented immediately after prime trials. The simulation was run for 50 different subjects, *independently*. That is the dynamical system was simulated beginning from the rest state for 50 times. Randomness is provided in the system by having noise in the strength of external inputs.

In order to quantify the amount of priming shown by the system, the system was simulated under two different conditions: (i) Priming condition: when the initial state of the target trial was influenced by the final state of the prime trial, and (ii) No-priming condition: when the initial state of the target trial was reset to the rest state. A statistic was derived by comparing (and normalising) the number of subjects who show repetition under the two conditions:

$$Priming = \frac{N_{np} - N_p}{N_{Sub}} * 100 \quad (6)$$

N_{np} = Number of subjects that do *not* show repetition in no-priming condition.

N_p = Number of subjects that do not show repetition in priming condition.

N_{Sub} = Total number of subjects.

In the absence of input from binding layer, $K_{stm} = 0$, the system shows *Priming* of 16%, and *Priming_{rep}* (which measures *Priming* when the verb was repeated between prime and target) of 10% – i.e. no enhanced priming for subjects receiving lexical boost. When K_{stm} is set to a positive value (here, 5), the *Priming* rises to 30%, and *Priming_{rep}* rises to 40.8%. Thus, under the influence of input from the binding nodes, the system starts showing lexical boost.


These results demonstrate that our system shows both priming and lexical boost (Bock, 1986; Pickering & Branigan, 1998). These results are not surprising since the system was designed to do just that, but it confirms that the given architecture is capable of showing these phenomena. The hysteresis in WTA layers ensures priming, and that in STM ensures lexical boost.

Experiment 2. This experiment introduced an *adaptation phase* between the prime and target trials. During this phase, no external input was given to the STM nodes, and a constant (and equal) external stimulus was given to the WTA nodes. The goal was to test the effect of forgetting (simulated through adaptation-time) on structural priming and lexical boost. Table 1 shows the results of varying the adaptation time on *Priming* and on *Priming_{rep}*.


Adapt time (ms)	0	1000	3000	4000	8000
<i>Priming</i>	46%	30%	28%	16%	6%
<i>Priming_{rep}</i>	57%	35%	36.8%	18%	3.8%

Table 1: Effect of forgetting on structural priming and lexical boost. (Fixed input from binding layer $K_{stm} = 10$)

We can observe from this table that the amount of priming decreases gradually as the adaptation time is increased, while the lexical boost (*Priming_{rep}* – *Priming*) shows a sudden decay around 4000ms. The reason for this behaviour can be seen by plotting the activations of STM and WTA nodes against adaptation time (Figure 3). The STM network with mutually excitatory nodes shows a catastrophic decay of memory after a certain time (Fig. 3(a)). On the other hand, the activations of the two nodes in the WTA network approach a constant difference asymptotically (Fig. 3(b)).



(a) STM adaptation



(b) WTA adaptation

Figure 3: Contrasting adaptation in WTA and STM networks

We can compare these results with the findings of Hartsuiker et al. (2008). In their study, Hartsuiker et al. tried to establish the duration of lexical enhancement on structural priming. They varied the lag between prime and target utterances by inserting filler utterances between the two. They found that structural priming persists over large lags (up to six filler items), while lexical enhancement of this priming decays quickly. The results from the current simulation (Table 1) mirror this effect. Furthermore, the model provides a mechanistic explanation for these effects and suggests that this difference in the temporal properties of priming and lexical boost is due to the difference in dynamics of hysteresis in competitive (WTA), and associative (STM) layers.

Experiment 3. The third experiment turned on long-term learning and tested the cumulative effect of several prime utterances (presented during the training phase) on a target utterance (Figure 2). Following Kaschak and Borreggine (2008), this experiment checks whether the influence of structural factors on priming changes if we change the set of verbs between training and testing phase (*Different Verb* condition) as compared to when we use the same set of verbs (*Same Verb* condition). If the pattern of priming changes during the *Dif-*

ferent Verb condition, then priming is clearly verb-specific – i.e. there is long-term lexical influence. Otherwise lexical influence is constrained to prime-target trials and does not have a cumulative effect.

To quantify the pattern of priming, Kaschak and Borreggine (2008) compared the amount of priming when subjects saw an equal number of each structure during the training phase (*Equal* condition) to when the subjects saw only one kind of (biasing) structure during the training phase (*Unequal* condition). The difference between the amount of priming shown for *Equal* and *Unequal* conditions provides a measure the cumulative priming from training to testing phase.

We simulated this experiment on eighty independent subjects, with half of the subjects assigned to the same verb condition and the other half to the different verb condition. The results of the simulation are shown in Figure 4. Mirroring the original study, the computational model shows a difference between priming in the equal and unequal cases. But crucially, this difference is similar under both same and different verb conditions – the same result obtained by Kaschak and Borreggine (2008) and one which they use to justify the lack of influence of lexical factors on long-term structural priming.

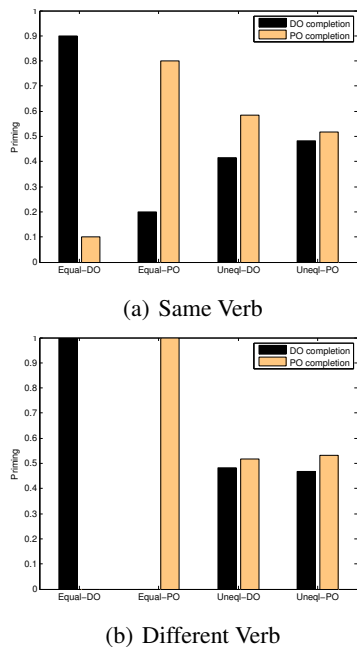


Figure 4: Results for Experiment 3. X-axis shows type of prime.

Experiment 4. Lastly, we turn our attention to the second experiment in Kaschak and Borreggine (2008). This experiment measures the amount of long-term structural priming under the *Balanced* (verb in training phase associated equally with both grammatical constructions) and *Skewed* (verb in training phase associated with only one kind of construction) conditions. The premise here is that if grammatical choice is sensitive to lexical factors, then biasing a word-form towards one (grammatical) construction during training should hinder the selection of the other construction during testing, thus giv-

ing larger priming in *Balanced* condition as compared to the *Skewed* condition.

In order to simulate this experiment, another input data set was constructed where half of the subjects were put into the *Balanced* condition and the other half into the *Skewed* condition. Again a set of eighty subjects were simulated with the same parameters as above. The results of the simulations are shown in Figure 5.




Figure 5: Simulation results for Experiment 4

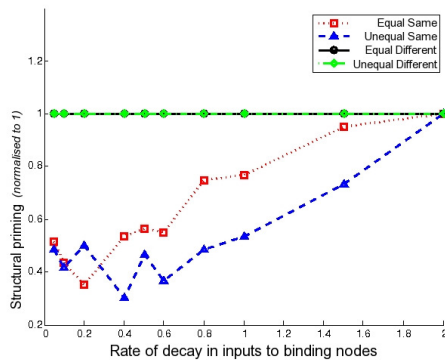
In this experiment too, we observe that the *Balanced* and *Skewed* training conditions show similar amount of structural priming. This mirrors the results of Kaschak and Borreggine (2008). We observe that even though our model is based on principles of trailing-activation, it can replicate lexical influence on long-term structural priming observed with human subjects. This clearly shows that an error-based account is not necessary for showing rapid decay in lexical boost when prime and target are separated over several trials.

Discussion

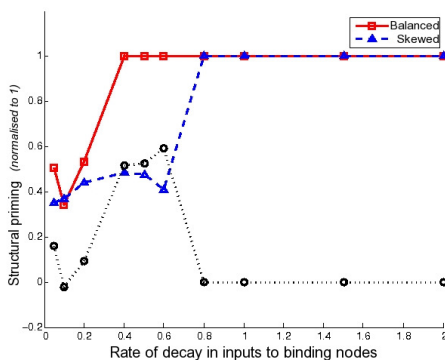
We saw in the previous section that our computational model successfully replicates the results from both Pickering and Branigan (1998) and Kaschak and Borreggine (2008). The architecture of the model is close to the one laid out in Pickering and Branigan (1998), with connections between lemma nodes and combinatorial nodes being replaced by binding nodes between lexical and syntactic layers. These binding nodes show hysteresis from priming episode to target episode and serve as a short-term memory of lexical-syntactic associations. The architecture also ensures that the model is homogeneous, with similar mechanisms responsible for learning and decision making, in lexical, syntactic and connecting nodes.

While the model successfully replicates these behavioural experiments, the more interesting question is why the model shows no lexical influence on long-term structural priming. Kaschak and Borreggine (2008) use this very lack of lexical influence on long-term structural priming to support error-based learning account, where the syntax-module abstracts away structural information and stores it in long-term memory. In order to understand the influence of lexical memory on structural priming, we varied the rate of forgetting in the binding nodes (the part of the network which stores associations between lexical and structural representations) and noted the

amount of long-term structural priming under the conditions of Experiments 3 and 4 (Figure 6). The reader would recall that long-term learning in the model is performed by adjustment to inputs to nodes. Thus in order to vary the rate of forgetting in binding nodes, we varied the rate of decay in these input weights, plotted against the x-axis in Figure 6. The values on the left show priming for a small order or decay – i.e. very slow rate of forgetting – and the values on the right show priming under a fast rate of forgetting.



(a) Variation of priming with rate of forgetting in binding nodes – *Same* and *Diff Verb* conditions



(b) Variation of priming with rate of forgetting in binding nodes – *Balanced* and *Skewed* conditions

Figure 6: Simulation results for Experiment 3 and 4 under different durations of lexical influence.

From fig 6(a) we can observe that long-term lexical influence does affect the Same-Verb conditions more; but this affect is similar for the Equal and Unequal cases (the blue/triangle and red/square curves). This means that the test that Kaschak and Borreggine (2008) adopt for checking lexical influence – namely, checking the difference of effect across Same-Verb and Different-Verb conditions – is not useful, since the effect size ($priming_{equal} - priming_{unequal}$) is same (close to zero) for both slow and fast forgetting.

The model also shows that long-term lexical influence does not necessarily imply a difference in Balanced and Skewed cases. This difference varies as a *U-shaped* function of increase in lexical influence (Fig 6(b), black/circle curve), such that the difference is close to zero for both long and short duration of influence. Although the model does not require lex-

ical influence to decay slowly, these results show that a long-term lexical influence does not necessarily manifest itself as a difference in balanced and skewed conditions, or Same and Different-Verb conditions. Thus, these results feed back into behavioural studies and suggest further testing to establish the exact duration of lexical influence on structural priming.

The computational model that we have presented in this paper implements an unsupervised memory system. Each episode of comprehension leaves a trace in the system – just as in a trailing-activation account. This system does not try to abstract away the rules of language production from each episode, and is therefore in contrast with an error-based account like that of Chang et al. (2006). It assumes an adult subject, who has already acquired a working knowledge of language. The model hypothesises that such an adult subject shows structural repetition because of memory traces of past episodes of language comprehension. It justifies this account by replicating the effects of priming and lexical enhancement of this priming over different time intervals. Finally, it also provides a mechanistic explanation for how priming might decay, and why this decay might be independent of the decay in lexical influence on this priming.

References

- Bock, K. (1986). Syntactic persistence in language production. *Cognitive Psychology*, 18, 355–387.
- Branigan, H. P., Pickering, M. J., & Cleland, A. A. (2000). Syntactic co-ordination in dialogue. *Cognition*, 75, B13–B25.
- Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, 113, 234–272.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorso-lateral prefrontal cortex. *Journal of Neurophysiology*, 61, 331–349.
- Hartsuiker, R. J., Bernolet, S., Schoonbaert, S., Speybroeck, S., & Vanderelst, D. (2008). Syntactic priming persists while the lexical boost decays: Evidence from written and spoken dialogue. *Journal of Memory and Language*, 58, 214–238.
- Kaschak, M. P., & Borreggine, K. L. (2008). Is long-term structural priming affected by patterns of experience with individual verbs? *Journal of Memory and Language*, 58, 862–878.
- Naka, K. I., & Rushton, W. A. (1966). S-potentials from colour units in the retina of fish. *J. Physiol.*, 185, 584–599.
- Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, 39, 633–651.
- Sandberg, A., Lansner, A., Petersson, K.-M., & Ekeberg rjan. (2000). A palimpsest memory based on an incremental bayesian learning rule. *Neurocomputing*, 32-33, 987–994.
- Wilson, H. R. (1999). *Spikes, decisions and actions*. Oxford University Press.