



Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2014 March ; 40(2): 326–347. doi:10.1037/a0034903.

Lexical interference effects in sentence processing: Evidence from the visual world paradigm and self-organizing models

Anuenue Kukona, Pyeong Whan Cho, James S. Magnuson, and Whitney Tabor

Department of Psychology, University of Connecticut Haskins Laboratories, New Haven, Connecticut

Abstract

Psycholinguistic research spanning a number of decades has produced diverging results with regard to the nature of constraint integration in online sentence processing. For example, evidence that language users anticipatorily fixate likely upcoming referents in advance of evidence in the speech signal supports rapid context integration. By contrast, evidence that language users activate representations that conflict with contextual constraints, or only indirectly satisfy them, supports *non*-integration or late integration. Here, we report on a self-organizing neural network framework that addresses one aspect of constraint integration: the integration of incoming lexical information (i.e., an incoming word) with sentence context information (i.e., from preceding words in an unfolding utterance). In two simulations, we show that the framework predicts both classic results concerned with lexical ambiguity resolution (Swinney, 1979; Tanenhaus, Leiman, & Seidenberg, 1979), which suggest late context integration, and results demonstrating anticipatory eye movements (e.g., Altmann & Kamide, 1999), which support rapid context integration. We also report two experiments using the visual world paradigm that confirm a new prediction of the framework. Listeners heard sentences like “The boy will eat the white...,” while viewing visual displays with objects like a *white cake* (i.e., a predictable direct object of “eat”), *white car* (i.e., an object not predicted by “eat,” but consistent with “white”), and distractors. Consistent with our simulation predictions, we found that while listeners fixated *white cake* most, they also fixated *white car* more than unrelated distractors in this highly constraining sentence (and visual) context.

Keywords

Anticipation; Artificial neural networks; Local coherence; Self-organization; Sentence Processing

1. Introduction

Linguistic structure at the phonological, lexical, and syntactic levels unfolds over time. Thus, a pervasive issue in psycholinguistics is the question of how information arriving at any instant (e.g., an “incoming” word) is integrated with the information that came before it (e.g., sentence, discourse, and visual context). Psycholinguistic research spanning a number

Correspondence concerning this article should be addressed to Anuenue Kukona, School of Psychology, University of Dundee, Dundee, DD1 4HN, Scotland, UK. a.b.bakerkukona@dundee.ac.uk.
Anuenue Kukona is now at the School of Psychology, University of Dundee.

of decades has produced diverging results with regard to this question. On the one hand, language users robustly anticipate upcoming linguistic structures (e.g., they make eye movements to a *cake* in a visual scene on hearing “The boy will eat the...;” Altmann & Kamide, 1999; see also Chambers & San Juan, 2008; Kamide, Altmann, & Haywood, 2003; Kamide, Scheepers, & Altmann, 2003; Knoeferle & Crocker, 2006, 2007), suggesting that they rapidly integrate information from the global context in order to direct their eye movements to objects in a visual display that satisfy contextual constraints. On the other hand, language users also seem to activate information that only *indirectly* relates to the global context, but by no means best satisfies contextual constraints (e.g., “bugs” primes “SPY” even given a context such as “spiders, roaches, and other bugs;” Swinney, 1979; see also Tanenhaus, Leiman, & Seidenberg, 1979).

These findings pose a theoretical challenge: they suggest that information from the global context places very strong constraints on sentence processing, while also revealing that contextually-inappropriate information is not always completely suppressed. Crucially, these results suggest that what is needed is a principled account of the balance between context-dependent and context-independent constraints in online language processing. In the current research, our aims were as follows: first, to show that the concept of *self-organization* provides a solution to this theoretical challenge; second, to describe an implemented self-organizing neural network framework that predicts classic findings concerned with the effects of context on sentence processing; and third, to test a new prediction of the framework in a new domain. The concept of self-organization refers to the emergence of organized, group-level structure among many, small, autonomously acting but continuously interacting elements. Self-organization assumes that structure forms from the bottom up, such that responses that are consistent with some part of the bottom-up input are gradually activated. Consequently, it predicts bottom-up interference from context-conflicting responses that satisfy some but not all of the constraints. At the same time, self-organization assumes that the higher-order structures that form in response to the bottom-up input can entail expectations about likely upcoming inputs (e.g., upcoming words and phrases). Thus, it also predicts anticipatory behaviors. Here, we implemented two self-organizing neural network models that address one aspect of constraint integration in language processing: the integration of incoming lexical information (i.e., an incoming word) with sentence context information (i.e., from the preceding words in an unfolding utterance).

The rest of this article is comprised of four parts. First, we review psycholinguistic evidence concerned with effects of context on language processing. Second, we describe a self-organizing neural network framework that addresses the integration of incoming lexical information (i.e., an incoming word) with sentence context information (i.e., from preceding words in an unfolding utterance). We show that the framework predicts classic results concerned with lexical ambiguity resolution (Swinney, 1979; Tanenhaus et al., 1979), and we extend the framework to address anticipatory effects in language processing (e.g., Altmann & Kamide, 1999), which provide strong evidence for rapid context integration. Third, we test a new prediction of the framework in two experiments in the visual world paradigm (VWP; Cooper, 1974; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995).

1.1. Rapid, immediate context integration

Anticipatory effects in language reveal that language users rapidly integrate information from the global context, and rapidly form linguistic representations that best satisfy the current contextual constraints (based on sentence, discourse, and visual constraints, among others). Strong evidence for anticipation comes from the visual world paradigm, which presents listeners with a visual context, and language about or related to that context. Altmann and Kamide (1999) found that listeners anticipatorily fixated objects in a visual scene that were predicted by the selectional restrictions of an unfolding verb. For example, listeners hearing “The boy will eat the...,” while viewing a visual scene with a *ball*, *cake*, *car*, and *train*, anticipatorily fixated the edible *cake* predicted by the selectional restrictions of “eat.”¹ By contrast, listeners hearing “The boy will move the...,” in a context in which all items satisfied the selection restrictions of “move,” fixated all items with equal probability.

Kamide, Altmann, et al. (2003) also demonstrated anticipatory effects based on the combined contextual constraints from a subject and verb. Listeners hearing “The girl will ride the...,” while viewing a visual display with a (young) *girl*, *man*, *carousel*, and *motorbike*, made more anticipatory fixations during “ride” to *carousel*, which was predicted by “girl” and “ride,” than to *motorbike*, which was predicted by “ride” but not “girl.” Thus, anticipatory eye movements at the verb were directed toward items that were predicted by the contextual constraints (e.g., *carousel*), rather than toward items that were closely related to the unfolding utterance (e.g., the rideable *motorbike*), but not predicted by the context (for further discussion, see the General Discussion).

Closely related work on syntactic ambiguity resolution provides further support for immediate integration of contextual constraints. Tanenhaus et al. (1995) presented listeners with temporary syntactic ambiguities like “Put the apple on the towel in the box.” They found that these sentences produced strong garden paths in “one-referent” visual contexts, which included objects like an *apple* (on a towel), an *empty towel* (with nothing on it), a *box*, and a *pencil*. On hearing “towel,” they found that listeners tended to fixate the *empty towel*, in preparation to move the *apple* there. However, there was no evidence for garden paths in “two-referent” visual contexts, which included a second *apple* (on a napkin). On hearing “towel,” they found that listeners rapidly fixated the *apple* on the towel, and that fixations to the *empty towel* were drastically reduced (there were roughly as many fixations to the *empty towel* as with unambiguous control sentences [e.g., “Put the apple that’s on the towel in the box”]). These findings suggest that “on the towel” was immediately interpreted as specifying *which* apple to move, rather than *where* to move an apple, when it could be used to disambiguate between referents in the visual context. Thus, eye movements were immediately constrained by the contextual constraints, such that fixations were not launched (i.e., relative to unambiguous controls) toward objects that were closely related to the unfolding utterance (e.g., *empty towel*), but not predicted by the context.

Related work on spoken word recognition also supports immediate effects of contextual constraints. Dahan and Tanenhaus (2004; see also Barr, 2008b) showed that “cohort

¹Throughout, we use “quotes” to indicate linguistic information (e.g., words, sentences, etc.), and *italics* to indicate visual display information (e.g., objects).

effects,” or the co-activation of words that share an onset (e.g., Allopenna, Magnuson, & Tanenhaus, 1998), were eliminated in constraining sentence contexts. They found that listeners hearing a neutral context (e.g., Dutch: “Nog nooit is een bok...” / “Never before has a goat...;” “een bok” [“a goat”] is the subject) showed a typical cohort effect at “bok:” listeners fixated the cohort competitor (e.g., *bot* / *bone*) of the target noun (e.g., *bok* / *goat*) more than unrelated distractors (e.g., *eiland* / *island*). However, listeners hearing a constraining context (e.g., “Nog nooit klom een bok...” / “Never before climbed a goat...;” again, “a goat” is the subject, and the cohort *bot* [*bone*] is not predicted by “climbed”) showed no difference in looks between the cohort competitor and unrelated distractors. Thus, eye movements at noun onset were immediately directed toward objects that were compatible with the contextual constraints (e.g., *bok*), rather than toward items that were closely related to the speech signal (e.g., *bot*, which was related to the onset of the noun), but not predicted by the context. Magnuson, Tanenhaus, and Aslin (2008) found similar results in a word-learning paradigm that combined the VWP with an artificial lexicon (Magnuson, Tanenhaus, Aslin, & Dahan, 2003): there was no phonological competition from noun competitors when the visual context predicted an adjective should be heard next, and no phonological competition from adjective competitors when the visual context did not predict an adjective.

In summary, the findings reviewed in this section support the rapid, immediate integration of contextual constraints. They suggest that at the earliest possible moment (e.g., word onset; Barr, 2008b; Dahan and Tanenhaus, 2004; Magnuson et al., 2008), listeners form linguistic representations that best satisfy contextual constraints. Moreover, these results suggest that listeners also do so *anticipatorily* (e.g., Altmann & Kamide, 1999). Nevertheless, results like those of Dahan and Tanenhaus (2004) are perhaps surprising in that language users do have experience with (often context-specific) uses of language that seem to violate typical expectations (e.g., a child playing with a dog bone could say, “The *bone* is climbing the dog house!”). Thus, their results suggest that verb selectional restrictions are a particularly powerful contextual constraint, which rapidly suppress representations that violate these restrictions.

1.2. Bottom-up interference

While findings concerned with anticipation and ambiguity resolution in the VWP seem to suggest that listeners immediately integrate contextual information, and immediately suppress representations that conflict with contextual constraints, closely related work supports a very different conclusion. In a number of settings, language users appear to activate representations on the basis of context-independent, “bottom-up” information from the linguistic signal, such that incoming input (e.g., an incoming word) is allowed to activate linguistic representations as if there were no context, and linguistic representations that are outside the scope of the contextual constraints, or *conflicting* with the context constraints, are also activated.

Classic work on lexical ambiguity resolution (e.g., Swinney, 1979; Tanenhaus et al., 1979) provides evidence that language users initially activate all senses of ambiguous words, irrespective of the context. In a cross-modal priming task, listeners heard sentences that

biased one interpretation of a lexically ambiguous word (e.g., “spiders, roaches, and other bugs,” biasing the insect rather than espionage sense of “bugs;” Swinney, 1979). At the offset of the ambiguity, they performed a lexical decision on a visual target word. Relative to unrelated words (e.g., SEW), Swinney found equivalent priming for targets related to both senses of the homophone (e.g., ANT and SPY). Tanenhaus et al. found similar effects with ambiguous words with senses from different syntactic categories (e.g., “rose” has both a noun “flower” sense, and a verb “stood” sense). For example, participants hearing “rose” in either a noun context (e.g., “She held the rose”) or a verb context (e.g., “They all rose”) showed cross-modal priming for “FLOWER” when it was presented 0 ms after “rose,” although “FLOWER” is inappropriate in the verb context. These results suggest that ambiguous words *initially* activate lexical representations independent of the sentence context (as we discuss later, these priming effects eventually dissipate when contextual constraints are available).

Local coherence effects provide further evidence that language users also form syntactic representations that conflict with contextual constraints (e.g., Tabor, Galantucci, & Richardson, 2004; see also Bicknell, Levy, & Demberg, 2010; Konieczny, Müller, Hachmann, Schwarzkopf, & Wolfer, 2009; Konieczny, Weldle, Wolfer, Müller, Baumann, 2010). Tabor et al. compared sentences like “The coach smiled at the player (who was) tossed/thrown the frisbee” in a word-by-word self-paced reading experiment. While the string “the player tossed the frisbee” forms a locally coherent active clause (e.g., which cannot be integrated with “The coach smiled at...”), “the player thrown the frisbee” does not. Critically, they found that reading times on “tossed” were reliably slower than on “thrown,” suggesting that the grammatically ruled-out active clause was activated, and interfering with processing. These results suggest that ambiguous phrases activate “local” syntactic representations independent of the “global” sentence context.

Cases where similarity in the bottom-up signal overwhelms a larger context have also been reported in spoken word recognition. Allopenna et al. (1998) found evidence of phonological competition between words that rhyme, despite mismatching phonological information at word onset. For example, listeners hearing “beaker,” while viewing a visual scene with a *beaker*, *beetle*, *speaker*, and *carriage*, fixated *beetle* more than distractors, but they also fixated rhyme competitors like *speaker* more than completely unrelated distractors like *carriage* (with a time course that mapped directly onto phonetic similarity over time). Thus, listeners activated rhyme competitors even though they could be ruled out by the context (e.g., the onset /b/).

Findings on the activation of lexical-semantic information also suggest that language users activate information that only indirectly relates to the contextual constraints. Yee and Sedivy (2006) found that listeners fixated items in a visual scene that were semantically related to a target word, but not implicated by the unfolding linguistic or task constraints. For example, listeners instructed to touch a “lock,” while viewing a visual scene with a *lock*, *key*, *apple*, and *deer*, fixated *lock* most, but they also fixated semantically-related competitors like *key* more than unrelated distractors. Even more remarkably, such effects can be phonologically mediated; if *lock* is replaced with *log*, people still fixate *key* more than other distractors upon hearing “log,” suggesting “log” activated “lock” which spread semantic activation to “key.”

Similarly, listeners also fixated *trumpet* when hearing “piano” (Huettig & Altmann, 2005), *rope* when hearing “snake” (Dahan & Tanenhaus, 2005), and *typewriter* when hearing “piano” (Myung, Blumstein, & Sedivy, 2006). These results suggest that the representations that are activated during language processing are not merely those that fully satisfy the constraints imposed by the context (which predicts fixations to the target alone, in the absence of potential phonological competitors). Rather, rich lexical-semantic information that is related to various aspects of the context (e.g., *key* is related to the word “lock,” but not the task demand of touching a “lock”) is also activated.

Kukona, Fang, Aicher, Chen, and Magnuson (2011) also re-assessed evidence for rapid context integration based on anticipation. Their listeners heard sentences like “Toby arrests the...,” while viewing a scene with a character called *Toby* and items like a *crook*, *policeman*, *surfer*, and *gardener*. Listeners were told that all sentences would be about things *Toby* did to someone or something, and an image of *Toby* was always displayed in the center of the screen (thus, it was clear that *Toby* always filled the agent role). Critically, their visual displays included both a good patient (e.g., *crook*) and a good agent (e.g., *policeman*) of the verb, but only the patient was a predictable direct object of the active sentence context. Kukona et al. hypothesized that if listeners were optimally using contextual constraints, they should not anticipatorily fixate agents, since the agent role was filled by *Toby*. During “arrests,” they found that listeners anticipatorily fixated the predictable patient (e.g., *crook*) most. However, listeners also fixated the agent (e.g., *policeman*) more than unrelated distractors (in fact, fixations to the patient and agent were nearly indistinguishable at the verb), based on context-independent thematic fit between the verb and agent. Thus, items were gradually activated in proportion to the degree to which they satisfied the contextual constraints. These results suggest that in addition to anticipating highly predictable items that satisfy the contextual constraints, language users also activate items based on thematic information independent of the sentence context.

In summary, the bottom-up interference effects reviewed in this section suggest that contextually-inappropriate information is not always completely suppressed by contextual constraints. Rather, language users activate information that may only be indirectly related to, or even in conflict with, contextual constraints. Critically, these studies provide a very different insight into the representations that language users are activating: they suggest that language users’ representations include rich information that is outside of what is literally or directly being conveyed by an utterance.

Additionally, there is an important temporal component to these various effects: they tend to be highly “transient.” For example, recall that Tanenhaus et al. (1979) found equivalent priming for “FLOWER” in both the noun context “She held the rose” and the verb context “They all rose” when probes were presented 0 ms after “rose.” However, when “FLOWER” was presented 200 ms after “rose,” they found greater priming in the noun context as compared to the verb context (suggesting initial exhaustive access to items matching the bottom-up input, followed by a later stage of context-based selection). Thus, while these findings suggest that contextually-inappropriate information is *ultimately* suppressed, they nevertheless reveal an important time course component: the transient activation of context-

conflicting representations (e.g., the noun sense of “rose” in a verb context like “They all rose.”).

1.3. Simulations and experiments

Results like those of Swinney (1979) and Tanenhaus et al. (1979) were historically taken as support for encapsulated lexical and syntactic “modules” (e.g., Fodor, 1983), because they suggested that lexical processes were not immediately subject to the influences of context. However, this kind of modularity has been all but abandoned given more recent evidence for wide-ranging influences of context on language processing (see Section 1.2). Thus, these diverging results regarding the impact of contextual constraints on language processing raise an important question: what kind of theory would predict the language system to be both context-dependent, such that language users activate representations that best satisfy the contextual constraints, and context-independent, such that language users activate representations without respect to the context?

One potential solution to this question involves re-assessing what we take to be the relevant contextual time window. For example, following “Toby arrests the...” (Kukona et al., 2011), *policeman* only conflicts with the context within a narrow time window (e.g., it is unlikely as the next word). Alternatively, in a longer (e.g., discourse) time window *policeman* may be very predictable, in so much as we often talk about policemen when we talk about arresting (e.g., “Toby arrests the crook. Later, he’ll discuss the case with the on-duty policeman;” by contrast, note that Kukona et al. interpreted their data as indicating that listeners were activating information that was *not* necessarily linguistically predictable). However, this account is problematic in two senses: first, eye movements to *policeman* quickly decrease following “arrest,” although *policeman* presumably becomes more discourse-predictable further on from the verb (e.g., it is likely in a subsequent sentence). More importantly, this account also cannot explain other effects of context-independent information: for example, there is not an apparent (e.g., discourse) time window in which “FLOWER” is relevant to “They all rose” (Tanenhaus et al., 1979), “the player tossed the frisbee” is relevant to “The coach smiled at the player tossed the frisbee” (Tabor et al., 2004), or *speaker* is relevant to “beaker” (Alloppenna et al., 1998). Thus, our goal with the current research was to develop a unifying framework that explains a wide range of evidence supporting both bottom-up interference from context-independent constraints *and* rapid context integration.

Here, we argue that the concept of *self-organization* offers a key insight into this puzzle. The term self-organization has been used to describe a very wide range of phenomena. In physics, it is associated with a technical notion, self-organized criticality (Jensen, 1998), which describes systems that are poised between order and chaos. In agent-based modeling, the term is used to describe systems of interacting agents that exhibit systematic behaviors. For example, Reynolds (1987) showed that “coordinated” flocking behavior emerges (“self-organizes”) among agents following very simple rules (e.g., maintain a certain distance from neighbors and maintain a similar heading). We adopt the following definition:

Definition. Self-organization refers to situations in which many, small, autonomously acting but continuously interacting elements exhibit, via convergence under feedback, organized structure at the scale of the group.

Self-organizing systems in this sense can be specified with sets of differential equations. We use such a formulation in the simulations we describe below. Well-known examples of self-organization from biology include slime molds (e.g., Keller & Segel, 1970; Marée & Hogeweg, 1999), social insects (e.g., Gordon, 2010), and ecosystems (Sole & Bascompte, 2006). In each of these cases, the continuous bidirectional interactions between many small elements cause the system to converge on coherent responses to a range of environmental situations. We claim that language processing works in a similarly bottom-up fashion: the stimulus of the speech stream prompts a series of low-level perceptual responses, which interact with one another via continuous feedback to produce appropriate actions in response to the environment. Previously proposed bottom-up language processing models such as the Interactive activation model of letter and word recognition (McClelland & Rumelhart, 1981), TRACE (McClelland & Elman, 1986), and DISCERN (Miikkulainen, 1993) are self-organizing in this sense.

Critically, structure thus forms from the bottom up in self-organization. Initially, all possible responses that are consistent with some part of the bottom-up input (e.g. phoneme, morpheme, word, etc.) are activated. But the response that is most consistent with all the parts of the input will generally get the strongest reinforcement through feedback, and will thus come to dominate, while the other possible responses will be shut down through inhibition. Thus, signs of bottom-up interference (e.g., activation of context-conflicting responses; see Section 1.2) will be transitory. In fact, if the contextual constraints are very strong, and the bottom-up signal supporting context-conflicting responses is very weak, then there will be no detectable interference. This latter case is hypothesized to hold in cases of “rapid, immediate context integration” (see Section 1.1). Indeed, even in the classic case of lexical ambiguity resolution (e.g., Swinney, 1979; Tanenhaus et al., 1979), the early activation of context-conflicting lexical items can be eliminated with a sufficiently strong context (e.g., closed-class words like “would” do not prime open-class homophones like “wood” in strongly constraining contexts that only predict a closed-class word; Shillcock & Bard, 1993). Here, we focus on the integration of incoming lexical information (i.e., an incoming word) with sentence context information (i.e., from preceding words in an unfolding utterance), and we show that adopting a self-organizing perspective unifies a range of findings from the psycholinguistics literature.

2. The self-organizing neural network framework

Artificial neural networks are systems of simple, interacting neuron-like units. Continuous settling recurrent neural networks are systems whose activations change continuously and whose units have cyclic connectivity. These networks have the architectural prerequisites for self-organization as defined above: their many small units interact via bidirectional feedback connections to converge on a coherent response to their input. Previously, continuous settling recurrent neural networks have been used to model, and to generate predictions about, many aspects of human sentence processing (e.g., Kukona & Tabor, 2011; Spivey & Tanenhaus, 1998; Spivey, 2007; Tabor & Hutchins, 2004; Vosse & Kempen, 2000; see also MacDonald, Pearlmuter, & Seidenberg, 1994; Trueswell & Tanenhaus, 1994). Here, we describe continuous settling recurrent neural networks that address the integration of incoming lexical information (i.e., an incoming word) with sentence context information

(i.e., from preceding words in an unfolding utterance). We focus on two key findings in the literature: evidence for context-independent processing, based on studies of lexical ambiguity resolution (e.g., Swinney, 1979; Tanenhaus et al., 1979), and evidence for rapid context integration, based on studies of anticipatory effects in the VWP (e.g., Altmann & Kamide, 1999).

Prior research suggests that artificial neural networks are well suited to capturing effects of both context-dependent and context-independent information. Tabor and Hutchins (2004) implemented an artificial neural network model of sentence processing, called self-organizing parsing (SOPARSE), in which lexically-activated treelets (Fodor, 1998; Marcus, 2001) self-organize into syntactic parse structures. The interacting lexical elements in SOPARSE act *semi-autonomously*: while each element has a context-independent mode of behavior, its behavior is continuously dependent on its interactions with other elements. Because organized structure (as reflected in activation patterns across the network) emerges out of the interactions of the semi-autonomous elements in SOPARSE, the model naturally predicts context-independent processing effects: the semi-autonomous elements are free to simultaneously form various (e.g., context-conflicting) syntactic parses, even as they form a syntactic parse that best satisfies the contextual constraints. Consequently, the global coherence of syntactic parses is not enforced in SOPARSE, thus predicting effects of context-independent information, like local coherence effects (e.g., Tabor et al., 2004).

However, feedback mechanisms in SOPARSE (Tabor & Hutchins, 2004) also allow structures that best satisfy the constraints to inhibit competing, context-conflicting structures over time (i.e., from the point of encountering the ambiguous bottom-up input), thus predicting that context-independent processing effects will be transient (e.g., SOPARSE predicts priming of “FLOWER” immediately after “They all rose,” but the absence of priming after a short delay; Tanenhaus et al., 1979).

2.1. Simulation 1: Self-organizing neural network approach to lexical ambiguity resolution

In Simulation 1, we implemented a network that addressed the classic results of Swinney (1979) and Tanenhaus et al. (1979). Again, these authors showed that listeners activated both contextually-appropriate and contextually-inappropriate senses of lexically ambiguous words (homophones) immediately after hearing them in constraining sentence contexts.

In Simulation 1, our neural network learned to activate sentence-level phrasal representations while processing sentences word-by-word. The simulation included only essential details of the materials and task: the network was built of simple word input units with feedforward connections to many phrasal output units, which were fully interconnected and continuously stimulated each other. Additionally, the network acquired its representations via a general learning principle. In Simulation 1, the network was trained on a language that included a lexically ambiguous word like “rose,” which occurred as both a noun (“The rose grew”) and a verb (“The dandelion rose”). The critical question is whether we see activation of the verb sense of “rose” in noun contexts, and activation of the noun sense of “rose” in verb contexts.

Our strategy with the current set of simulations was to model the integration of lexical and sentential information given a minimally complex network that was capable of capturing the critical phenomena. While implementing models that capture a wide variety of phenomena can provide insight into cognitive phenomena, such an approach often requires very complex models that can be very difficult to understand. Consequently, we implemented simulations that were complex enough to make critical behavioral distinctions, but also relatively simple to analyze. As a result, the current simulations are considerably less assumption-laden than, for example, SOPARSE (Tabor & Hutchins, 2004). Also in contrast to SOPARSE, the current simulations used simple phonological-form detectors, like Elman (1990, 1991)'s Simple Recurrent Network (SRN), rather than linguistic treelets.²

2.1.1. Simulation 1

2.1.1.1. Architecture: The architecture of the network is depicted in Figure 1. The network consisted of a word input layer, with 12 localist word nodes (i.e., one word per node) of various syntactic classes (determiners: “the,” “your;” nouns: “dandelion,” “tulip;” verbs: “grew,” “wilted;” adjectives: “pretty,” “white;” adverbs: “slowly,” “quickly;” and a lexically ambiguous noun/verb: “rose”), and a phrasal output layer, with 27 localist noun or verb phrasal nodes (reflecting the grammatical combinations of word nodes; e.g., NP: “the dandelion,” “the pretty dandelion,” “the white dandelion,” etc.; VP: “grew,” “slowly grew,” “quickly grew,” etc.). Thus, the output phrasal layer used distributed representations to a limited degree, in that the network activated both an NP and VP node for each sentence. We included this limited degree of abstraction in order to show that the network can interact effectively with a mix of local and distributed representations (distributed representations are more plausible given complex combinatorial structures, which would alternatively require a very large number of localist representations). The network had feedforward connections from the word input layer to the phrasal output layer, and recurrent connections within the phrasal output layer. Activations in the word input layer corresponded to words the network “heard,” and activations across the phrasal output layer corresponded to the network's sentence-level representations.

2.1.1.2. Processing: Our network is closely related to a number of artificial neural network models in the literature, although it differs in certain key respects. Like the Simple Recurrent Networks used by Elman (1990, 1991) to model sentence processing, the network processed sentences word-by-word: processing of a word was implemented by setting the activation for a word node to 1 in the word input layer, and all other word nodes to 0. Unlike the typical approach with SRNs, word units were turned on for an interval of time, corresponding to the temporal durations of words. In this regard, the network resembles TRACE (McClelland & Elman, 1986), a model of spoken word recognition that takes in input across continuous time (however, in TRACE the input is gradient acoustic/phonological information, rather than categorical localist word inputs). Our localist inputs are a convenient simplification that were sufficient to characterize the phenomena we focus

²The current model is also simpler than the SRNs used by Elman (1990, 1991), in that it did not use hidden units. One shortcoming of the current model is that it cannot handle arbitrary long-distance dependencies (whereas the Elman net and its relatives plausibly can; e.g., Tabor, 2000; Tabor, Cho, & Szudlarek, 2013).

on here (however, graded, distributed representations, which can characterize similarity in phonological and semantic structure, are quite compatible with this framework).

So as to not introduce a duration-based bias, all words processed by the network were of the same duration (40 timesteps). Explorations of the network's behaviors suggest that while it required more training for shorter word durations (i.e., to reach a similar level of performance), the behavior of the network at the end of training remained similar across a wide range of word durations.

The phrasal output layer activations changed according to Equations 1 and 2. In these equations, net_i is the net input to the i th unit, a_j and a_i are the activation of the j th and i th units, respectively, w_{ij} is the weight from the j th to i th unit, and η_a is a scaling term ($\eta_a = 0.05$). Equation 2 creates a sigmoidal response as a function of net input (net_i).³

$$net_i = \sum_{j \in \text{units}_{sending}} a_j w_{ij} \quad (1)$$

$$\Delta a_i = \eta_a net_i a_i (1 - a_i) \quad (2)$$

2.1.1.3. Training: The network was trained on a language of 162 sentences (e.g., “The pretty dandelion slowly wilted,” “Your tulip grew,” etc.), reflecting the grammatical combinations of words from the network's language. The set of training sentences is specified by the grammar in Appendix A. The network was trained to activate both the NP and VP node in the phrasal output layer that was associated with each sentence (e.g., NP: “The pretty dandelion” and VP: “slowly wilted” for “The pretty dandelion slowly wilted”) as each word in a sentence was “heard.” The network was trained on a set of 5,000 sentences, which were selected at random from the network's language. We stopped training after 5,000 sentences, at which point the network showed a high level of performance in activating target nodes (performance statistics are reported in the Simulation results). We report the average behavior of the network across 10 different simulations. Learning was implemented via the delta rule in Equation 3, with weights updated once per word presentation. In this equation, t_i is the target activation of the i th unit in the phrasal output layer, and η_w is the learning rate ($\eta_w = 0.02$).

$$\Delta w_{ij} = \eta_w (t_i - a_i) a_i \quad (3)$$

2.1.2. Simulation results—To confirm that the network learned the training language, we tested the network on one sentence of each of the four syntactic types (i.e., Det N V; Det Adj N V; Det N Adv V; Det Adj N Adv V) at the end of training. At the offset of each test sentence, the network activated the target NP node more than all non-target nodes with all test sentences (accuracy: $M = 1.00$, $SD = 0.00$), and the target VP node more than all non-

³Equation 2 combines the excitatory and inhibitory equations used in other language processing models (e.g., Interactive activation model of letter and word recognition, McClelland & Rumelhart, 1981; TRACE, McClelland & Elman, 1986) into a single equation, without decay terms.

target nodes in all but two instances across the 10 simulations (accuracy: $M = 0.95$, $SD = 0.11$). While we examined variations in a range of free parameters (e.g., word durations, number of simulations, etc.), we did not need to engage in any formal or informal parameter fitting; the behavior of the network was consistent across a wide range of parameter settings given training to a similar level of performance. Following training, we tested the network with the example sentences “The rose grew” (noun context) and “The dandelion rose” (verb context). Activations of relevant NP and VP nodes are plotted in both the noun and verb contexts in Figures 2A and 2B. Again, the critical question is whether we see greater activation in the noun context of the VP nodes associated with the verb “rose” (e.g., “VP: rose”) as compared to VP nodes not associated with the verb “rose” (e.g., “VP: grew,” “VP: wilted”), and whether we see greater activation in the verb context of the NP nodes associated with the noun “rose” (e.g., “NP: rose”) as compared to NP nodes not associated with the noun “rose” (e.g., “NP: dandelion,” “NP: tulip”).⁴

First, in both contexts we observed greatest activation of the appropriate NP and VP structures at the end of each sentence: in Figure 2A, the network activated “NP: The rose” and “VP: grew” most when hearing “The rose grew,” and in Figure 2B, the network activated “NP: The dandelion” and “VP: rose” most when hearing “The dandelion rose.” These results suggest that the network learned the structure of the language. Second, we also found evidence for bottom-up interference from context-independent constraints in both the noun and verb contexts. In Figure 2A, the network activated “VP: rose” more than both “VP: grew” and “VP: wilted” when hearing “rose” in a noun context. In Figure 2B, the network activated “NP: The rose” more than “NP: The tulip” when hearing “rose” in a verb context (note that “NP: The dandelion” is activated more than “NP: The rose” because the network just heard “The dandelion”).

2.1.3. Discussion— The network highly activated phrasal structures that were appropriate for the sentence context (e.g., “NP: The rose” and “VP: grew” when hearing “The rose grew”). However, consistent with Tanenhaus et al. (1979), the network also activated context-conflicting structures based on context-independent lexical information (e.g., “VP: rose” given “rose” in a noun context, and “NP: rose” given “rose” in a verb context). These results reveal that the network was gradiently sensitive both to context-dependent constraints, based on the sentence context, and context-independent constraints, based on context-independent lexical information, consistent with a large body of evidence supporting both rapid context integration (see Section 1.1) and bottom-up interference (see Section 1.2).

The transitory activation of both senses of the ambiguous word “rose” stems from the self-organizing nature of the trained network. The learning rule causes positive connections to form between the input unit for “rose” and all output units involving “rose.” It also causes inhibitory connections to form among all the NP output nodes (because only one NP percept occurs at a time) and among all the VP output nodes (for the corresponding reason).

⁴Note that the model was not trained on semantic category information (i.e., it was not trained that the nouns “dandelion,” “rose,” and “tulip” all belong to the semantic category of “flowers”), and thus it would not be expected to show semantic priming effects (i.e., between nouns).

Consequently, even though the VP nodes associated with “rose” (e.g., “VP: rose”) are activated when the input is “rose” in a noun context, this errant activation is shut down once the appropriate VP node (e.g., “VP: grew”) becomes sufficiently activated to inhibit competitors (e.g., “VP: rose”).⁵

2.2. Simulation 2: Self-organizing neural network approach to anticipatory effects

In Tanenhaus et al. (1979), the contextual information logically ruled out words from particular syntactic classes (e.g. verbs in a noun context, and nouns in a verb context). However, their contexts were not “completely” constraining in the sense that the word classes that were not ruled out by the contextual constraints still contained massive numbers of words (cf. Shillcock & Bard, 1993). By contrast, in studies like Altmann and Kamide (1999) and Dahan and Tanenhaus (2004), the contextual information ruled out all but a single outcome (e.g., the *cake* in the scene following “The boy will eat the...”). Critically, in this *completely* constraining context (i.e., in which there was a single outcome), Dahan and Tanenhaus (2004; see also Barr, 2008b; Magnuson et al., 2008) found that context-independent activation of cohort competitors was eliminated (e.g., the cohort competitor *bot* [*bone*], which was not predicted by the verb “kloem” [“climbed”], was not fixated more than unrelated distractors in the context of “Nog nooit kloem een *bok*...” [“Never before climbed a *goat*...”). In Simulation 2, we tested the network in a similarly constraining anticipatory VWP context. Again, the simulation included only essential details of the materials and task: the network was built of simple word input units with feedforward connections to object output units, which interacted with one another via recurrent connections.

In Simulation 2 (and the experiments that follow), we used contexts involving strong verb selectional restrictions that were closely based on the materials of Altmann and Kamide (1999). We presented the network (followed by human listeners) with sentences like “The boy will eat the white cake,” while viewing visual displays with objects like a *white cake*, *brown cake*, *white car*, and *brown car*. The direct object of the sentence (e.g., *white cake*) could be anticipated at the adjective, based on the contextual constraints from “eat” and “white” (predicting *white cake*, and not any of the other objects in the visual display). However, the visual display also included a “locally coherent” competitor (e.g., inedible *white car*), which fit the context-independent lexical constraints from the adjective (e.g., “white”), but was not predicted by the contextual constraints (e.g., the verb selection restrictions of “eat”). Critically, we tested for the anticipation of objects predicted by the contextual constraints (e.g., *white cake*), and for the activation of objects based on context-independent lexical constraints (e.g., *white car*, which was ruled out by “eat,” but consistent with “white”). In Experiments 1 and 2, which follow, we tested the predictions of Simulation 2 with human language users.

⁵For clarity, it is worth noting that there is another timescale at which the present network model is self-organizing: the timescale of the learning. In this case, the successive events of word processing that occur over many sentences during the training phase cause the weights to move into relationships in which the network coherently responds to its environment. This convergence of the weight relationships occurs via feedback among the weights themselves: adjustment of one weight often requires complementary adjustment of other weights. We do not discuss the self-organizing aspects of weight change further because this paper is focused on online processing as opposed to learning (for other recent modeling work addressing interactions between learning and processing, see e.g., Dell & Chang, in press; Fine & Jaeger, 2013; Jaeger & Snider, 2013).

2.2.1. Simulation 2

2.2.1.1. Architecture: The architecture of the network is depicted in Figure 3. The network consisted of a language input layer, with six localist word nodes, and a visual output layer, with four localist object nodes. As in Simulation 1, the network had feedforward connections from the input layer to the output layer, and recurrent connections within the output layer. Activations in the language input layer corresponded to words the network “heard,” and activations across the visual output layer corresponded to a probability distribution of “fixating” each object in the visual context. Thus, the network modeled the central tendency of fixation behaviors (e.g., average proportions of fixations) rather than individual eye movements.

2.2.1.2. Processing: Processing was identical to Simulation 1. Additionally, activations across the four object nodes in the visual output layer were interpreted as fixation probabilities, p_i , after normalization via Equation 4.

$$p_i = \frac{a_i}{\sum_{j \in \text{units}_{\text{output}}} a_j} \quad (4)$$

2.2.1.3. Training: The network was trained on a set of 4 verb-consistent sentences (e.g., “Eat [the] white/brown cake,” “Drive [the] white/brown car”). To simplify the structure of the language, we used imperative sentences that lacked an overt subject noun phrase (as well as a direct object determiner). However, these sentences were comparable to the materials in Altmann and Kamide (1999) because the subject in those sentences could not be used to anticipate the direct object (e.g., “The boy...” does not predict edible vs. inedible). The network was trained to activate the object node in the visual output layer that was the direct object of each sentence (e.g., *white cake* for “Eat [the] white cake”), as each word in a sentence was “heard” by the network (e.g., as the word nodes for “eat,” “white,” and “cake” in the language input layer were turned on word-by-word). As in Simulation 1, learning was implemented via the delta rule. Training epochs consisted of training on each sentence once in a randomized order. The network was trained for 250 epochs.

2.2.2. Simulation results—To confirm that the network learned the training language, we tested the network on each of the four training sentences at the end of training. At the offset of each test sentence, the network activated the target object node more than all non-target nodes with all test sentences (accuracy: $M = 1.00$, $SD = 0.00$). At various points during training, we also tested the network with the example sentence “Eat [the] white cake.” Simulated proportions of fixations to objects in epochs 1, 50, 100, 150, and 250 are plotted in Figures 4A–E. The total sum squared error in the network across training is also plotted in Figure 4F. Again, the critical question is whether we see greater activation following “Eat [the] white...” of *white car*, the locally coherent competitor, compared to *brown car*, an unrelated distractor.

Following epoch 1, the fixation pattern was as follows: simulated proportions of fixations to the cakes diverged from the cars at “eat,” and simulated fixations to *white cake* diverged

from *brown cake* at “white.” Critically, the network also fixated *white car* more than *brown car* at “white” (and “cake”). Further training of the network had the effect of increasing looks to the *white cake*, and decreasing looks to the remaining objects. Additionally, the magnitude of the difference between *white car* and *brown car* diminished with further training (for further discussion, see Section 2.3).⁶

2.2.3. Discussion—The network anticipated target objects based on the combined contextual constraints from the verb and adjective (e.g., *white cake* given “Eat [the] white...”). The network’s incremental sensitivity to the contextual constraints is consistent with prior demonstrations of anticipatory eye movements in the VWP (e.g., Altmann & Kamide, 1999; and specifically, anticipatory effects based on multi-word constraints; e.g., Kamide, Altmann, et al., 2003). However, the network also activated context-conflicting structures based on context-independent lexical information from the adjective (e.g., the locally coherent competitor *white car* given “Eat [the] white...”). As in Simulation 1, these results reveal that the network was gradiently sensitive both to context-dependent constraints, based on the sentence context, and context-independent constraints, based on context-independent lexical information, despite the *completely* constraining context. Again, these gradient effects are compatible with a large body of evidence supporting both rapid context integration (see Section 1.1) and bottom-up interference (see Section 1.2). In Experiments 1 and 2, we directly test the predictions of Simulation 2 with human language users.

2.3. Evaluation of the self-organizing neural network approach

Simulations 1 and 2 reveal that the network robustly predicts gradient effects of context-dependent and context-independent information in language processing, thus providing a coherent account of results showing both rapid context integration and bottom-up interference. In both simulations, the network learned to balance context-dependent and context-independent information via a general learning principle. The networks are self-organizing by the definition stated above: bottom-up stimulation causes the output units to continuously interact via bidirectional feedback connections, and to converge on a coherent response to each environmental stimulus. Unlike some other self-organizing models (e.g., Interactive activation model of letter and word recognition [McClelland & Rumelhart, 1981]; TRACE [McClelland & Elman, 1986]; DISCERN [Miikkulainen, 1993]), the converged state is not a distributed representation with many activated units. Unlike some other self-organizing phenomena, the converged state is also not an elaborate, temporally extended pattern of activity, as in the case of foraging for nectar in environmental locations where flowers are abundant (e.g., bee colony self-organization; Hölldobler, & Wilson, 2008). Instead, the converged state in the current networks (i.e., as reflected in the training outputs) is a two-on bit vector specifying the syntactic structure of a sentence (Simulation 1), or a localist vector specifying a locus of eye-fixation (Simulation 2). Despite the natural intuition that self-organization involves the formation of very complex (possibly temporally extended) structures, we are suggesting here that the crucial property that distinguishes self-

⁶The magnitude of the difference between *white car* and *brown car* initially increased in size during training: the average difference in fixations was roughly zero in epoch 1 (see Figure 4A), and grew to a peak around roughly epoch 33 (shortly before Figure 4B).

organizing systems from other kinds of systems is that they respond coherently to their environment via continuous feedback among bidirectionally interacting elements. The current networks fit this description. On the one hand, the current networks also have simple converged states; however, this is because we specified a simple encoding for their output behaviors. Crucially, we did not use these simple encodings because we think that human responses to language stimuli are simple, or because we believe that a two-layer featural network is adequate; hidden units and multiple, recurrently connected layers are likely needed. Rather, we made these assumptions to highlight crucial features of self-organization, and their relation to key data points, while using a network that can be straightforwardly analyzed.

Helpfully, consideration of the current simulations leads to an answer to the question of *why* human language users show bottom-up interference effects: the simulations suggest that they may be a residue of the learning process. In the network, context-independent effects depended on first-order associations between inputs (e.g., words) and outputs (e.g., phrases/objects), whereas context-dependent effects depended on higher-order associations across sequences of inputs. Crucially, the network must first be able to detect first-order input-output associations, before it can detect higher-order associations across sequences of inputs. Therefore, context-independent processing of inputs occurs first during training (thus, bottom-up interference from *white car* is greatest early in training; see Figure 4B), and context-dependent processing develops as a modulation of context-independent effects later in training (thus, fixations to *white car* diminish with further training). In addition, Simulation 1 provides an indication of how the framework can scale up to languages of greater complexity, and Simulation 2 reveals that the framework captures results from a number of experimental paradigms (e.g., VWP).

In connection with prior models, the framework also provides coverage of a wide variety of phenomena that support bottom-up interference from context-independent constraints. In addition to the results of Swinney (1979) and Tanenhaus et al. (1979), we highlighted two closely related findings in the Introduction: local coherence effects (e.g., activation of “the player tossed the frisbee” when reading “The coach smiled at the player tossed the frisbee;” Tabor et al., 2004) and rhyme effects (e.g., activation of *speaker* when hearing “beaker;” Allopenna et al., 1998). As we reviewed, Tabor and Hutchins (2004)’s SOPARSE model captures local coherence effects. Relatedly, Magnuson et al. (2003) used a Simple Recurrent Network (Elman, 1990, 1991) to capture rhyme effects. Their model used localist phonetic feature inputs, and localist lexical outputs. Critically, both SOPARSE and SRNs are compatible with our self-organizing neural network framework: they share the core property that coherent structures emerge via repeated feedback interactions among multiple interacting elements.

However, our simulations also raise a number of important issues. One critical simplification is the network’s feedforward connections: for example, activations in Simulation 2 flowed from the language layer to the object layer, but not back from the object layer to the language layer. The implication is that language processing influences eye movements, but that visual context does not influence language processing, counter to a number of findings (e.g., Tanenhaus et al., 1995). In fact, the self-organizing framework is

fully compatible with visual to language feedback connections. For example, Kukona and Tabor (2011) modeled effects of visual context on the processing of lexical ambiguities using an artificial neural network that included visual to language feedback connections. Critically, they observed closely related bottom-up interference effects even in the context of these visual to language feedback connections.

A second issue raised by the current simulations concerns training. In fact, the magnitude of the difference between *white car* and unrelated distractors in Simulation 2 diminished over training, such that if training were allowed to continue for an *infinite* length of time, the effect would disappear altogether (given limitations on the numerical precision of our computer-based simulations, the effect actually disappears before then). Thus, a critical question is: at which point in the network's training should its behavior be mapped onto human language users? Our predictions are drawn from the network when it is in a transient state, and there is room for error to improve (even if to a very small degree; e.g., prior to epoch 250), rather than when it is in an asymptotic state, and error has stabilized (e.g., beyond epoch 250). This strategy closely follows Magnuson, Tanenhaus, and Aslin (2000), who similarly ended training in their SRN before it learned the statistics of its training language *perfectly*. They suggest that ending training before asymptote provides a closer “analog to the human language processor,” given the much greater complexity of human language (i.e., than typical training languages) and the much greater amount of training required to reach asymptote (p. 64; see also Magnuson et al., 2003).

Moreover, the network also fixates *white car* more than unrelated distractors over essentially its *entire* training gradient (with the exception of epoch 1): when total sum squared error is anywhere between its maximum and minimum (see changes in error along the y-axis in Figure 4F), we observe the same *qualitative* pattern, which changes only in *magnitude* (e.g., see Figures 4B - 4E). Thus, our predictions do not reflect precise training criteria; rather, our predictions reflect the behavior of the network over all but a small (initial) piece of its training gradient. Additionally, the pattern of results we observed over the course of training provides predictions about human language development and individual differences, which we return to in the General Discussion.

Finally, it is important to note that the current predictions stem from the self-organizing nature of the network in *combination* with constraints on its training. Certainly, we are not simply drawing our predictions from the network when it is untrained, and exhibiting largely random behavior (e.g., epoch 1; see Figure 4A). Rather, the network simultaneously shows bottom-up interference while also showing a massive anticipatory effect (e.g., epoch 250; see Figure 4E), suggesting that it has considerable knowledge of the structure of its language even while showing the crucial interference effect (i.e., not unlike human language users). Moreover, the current predictions are not a direct and unambiguous consequence of training to less-than-asymptote. For example, an alternative approach might plausibly predict a very similar pattern to Dahan and Tanenhaus (2004; e.g., no *white car* vs. *brown car* difference, suggesting no bottom-up interference), but with fewer looks to targets, and more looks to *all* non-targets (i.e., irrespective of whether they are locally coherent or not), early in training. Crucially, the current predictions depend on the assumption that responses that are only consistent with parts of the bottom-up input also get activated.⁷

3. Experiments

In Experiments 1 and 2, we tested for bottom-up interference in contexts involving strong verb selectional restrictions (e.g., Altmann & Kamide, 1999; Dahan & Tanenhaus, 2004). Our experiments followed directly from Simulation 2: human language users heard sentences like “The boy will eat the white cake,” while viewing a visual display with objects like a *white cake*, *brown cake*, *white car*, and *brown car* (see Figure 5). The direct object (e.g., *white cake*) could be anticipated based on the combined contextual constraints from “eat” and “white.” Critically, the visual display also included a locally coherent competitor (e.g., *white car*), which fit the lexical constraints from “white,” but was not predicted by the selection restrictions of “eat.” Finally, the visual display also included a verb-consistent competitor (e.g., *brown cake*, which was predicted by “eat” but not “white”), to test for the impact of the joint verb plus adjective constraints, and an unrelated distractor, which was either similar (Experiment 1: e.g., *brown car*) or dissimilar (Experiment 2: e.g., *pink train*) to the locally coherent competitor.

In both experiments, we tested two key questions. First, do listeners anticipatorily fixate predictable target objects (e.g., *white cake*) that best satisfy the contextual constraints (e.g., as compared to all other objects)? And second, do listeners fixate locally coherent competitor objects (e.g., *white car*) based on context-independent lexical constraints (e.g., as compared to unrelated distractors)?⁸ Simulation 2 predicts most anticipatory fixations following the adjective to the predictable target. However, the simulation also predicts more fixations to the locally coherent competitor than unrelated distractors, on the basis of context-independent lexical information. By contrast, if verb selectional restrictions completely rule out representations that are inconsistent with those constraints, then we expect anticipatory fixations following the adjective to the predictable target, and no difference in fixations between the locally coherent competitor and unrelated distractors, consistent with the results of Dahan and Tanenhaus (2004).

3.1. Experiment 1

3.1.1. Methods

3.1.1.1. Participants: Twenty University of Connecticut undergraduates participated for partial course credit. All participants were native speakers of English with self-reported normal hearing and self-reported normal or corrected-to-normal vision.

⁷An anonymous reviewer asked us to clarify how the network's behavior would differ given a language with different word ordering. The structure of the language in Simulation 1 is based closely on English, in which adjectives precede nouns (e.g., “white cake”). This is not true of all languages: in French the order is reversed (“gâteau blanc”). Critically, our simulation results would thus differ for languages like French: note that in English, following “Eat the white,” *white car* is consistent with one word (e.g., “white”), and inconsistent with one word (“eat”). In French, following an utterance like “Tu manges gâteau blanc,” *white car* is consistent with one word (e.g., “blanc” / “white”), but inconsistent with two words (e.g., “manges” / “eat” and “gâteau” / “cake”). Thus, fixations to *white car* at “blanc” (“white”) would be predicted to be smaller in French, given the greater number of contextual constraints.

⁸Note that the *white car* effects we are testing for are also “anticipatory,” in that we expect them while sentences are still unfolding. Consistent with this prediction, a large body of research (e.g., Eberhard, Spivey-Knowlton, Sedivy, & Tanenhaus, 1995; Sedivy, 2003; Sedivy, Tanenhaus, Chambers, & Carlson, 1999) reveals that listeners rapidly use adjective information to anticipate objects in the VWP. For example, listeners hearing sentences like “Pick up the yellow banana” while viewing a visual display with objects like a *yellow banana*, *yellow notebook*, *green pea*, and *pink eraser*, fixated the yellow objects more than the non-yellow objects shortly after hearing “yellow.”

3.1.1.2. Materials: We modified 16 predictive sentences from Altmann and Kamide (1999) to include a post-verbal direct object adjective (e.g., “The boy will eat the white cake.”). Each sentence was associated with a verb-consistent direct object (e.g., cake), a verb-inconsistent competitor (e.g., car), and two color adjectives (e.g., white/brown). The full item list is presented in Appendix B.⁹ Each sentence was also associated with a visual display, which included four objects in the corners of the visual display, and a fixation cross at the center. The four objects reflected the crossing of verb-consistent/inconsistent objects with the color adjectives (see Figure 5).

A female native speaker of English recorded sentences using both adjectives in a sound-attenuated booth. Recordings were made in Praat with a sampling rate of 44.1 kHz and 16-bit resolution. The speaker was instructed to produce the sentences naturally. We used clip-art images of the direct objects and competitors for our visual stimuli.

3.1.1.3. Design: We used a 2 X 2 design with verb consistency (consistent and inconsistent) and adjective consistency (consistent and inconsistent) as factors, and with all four conditions represented in each trial. For participants hearing “The boy will eat the white cake,” for example, the visual display included a verb-consistent and adjective-consistent *white cake*, a verb-consistent and adjective-inconsistent *brown cake*, a verb-inconsistent and adjective-consistent *white car*, and a verb-inconsistent and adjective-inconsistent *brown car*. Thus, participants viewed 16 objects of each type across the experiment (i.e., one of each type per visual display). Half the participants heard the other associated color adjective (e.g., “The boy will eat the brown cake.”). In this case, the adjective consistency labeling was reversed (e.g., *brown car* counted as verb-inconsistent and adjective-consistent). The target adjective in sentences, the location of objects in the visual display, and the order of sentences was randomized for each participant. Because objects rotated randomly between adjective-consistent and -inconsistent conditions, differences in looks between the adjective conditions are not likely to be due to differences in saliency.

3.1.1.4. Procedure: Participants initiated each trial by clicking on a fixation cross in the center of the computer screen. Then, the visual display for the trial was presented for a 500 ms preview, and it continued to be displayed as the associated sentence was played over headphones. Participants were told that they would be hearing sentences about people interacting with various things, and they were instructed to use the mouse to click on the object that was involved in each interaction (i.e., the direct object of the sentence). A trial ended when participants clicked on an object in the visual display, at which point the scene disappeared and was replaced by a blank screen with a fixation cross at the center. The experiment began with four practice trials with feedback, and there were no additional filler trials. We used an ASL R6 remote optics eye tracker sampling at 60 Hz with a head-tracking

⁹Although some competitors began with the same phonological segment (e.g., “cake” and “car”), this is generally not considered sufficient phonological overlap to result in robust competition; the standard definition of “cohort” competitors is overlap in the first two segments or in the onset and first vowel (e.g., Magnuson, Dixon, Tanenhaus, & Aslin, 2007). Moreover, the analyses we report below focus on fixations prior to bottom-up information from the noun. Thus, onset competition effects from the noun are neither expected, nor capable of contaminating our analyses of primary interest. For the sake of completely ruling out any such possibility, we also report analyses that exclude all items with first-segment overlap (i.e., items 1, 10, 11, and 15; see Appendices A and B, and Footnotes 10 and 11).

device (Applied Scientific Laboratories, MA, USA) and E-Prime software (Version 1.0, Psychology Software Tools, Inc., Pittsburgh, PA). The experiment was approximately 10 minutes in length.

3.1.1.5. Mixed-effects models: Eye movements from Experiments 1 and 2 were analyzed using linear mixed-effects models (lme4 in R). We focused our analyses on the onset of the direct object noun (e.g., “cake”), where we expected listeners’ eye movements to reflect information from the verb and adjective, but not direct object noun. Eye movements were submitted to a weighted empirical-logit regression (Barr, 2008a). First, we computed the number of trials with a fixation to each object for each condition for each participant or item (in order to allow for participants and items analyses) at the onset of the direct object noun. Second, we submitted the aggregated data (i.e., across participants or items) to an empirical-logit transformation. Third, the transformed data was submitted to mixed-effects models with fixed effects of condition, and random intercepts for participants/items (e.g., participants models used the following syntax: “outcome ~ condition + (1| participant)”). We did not include random slopes because there was only one data point (i.e., in the aggregated data) for each condition for each participant/item (e.g., see Barr, Levy, Scheepers, & Tily, 2013). We computed *p*-values for our fixed effects by using a model comparison approach: we tested for a reliable improvement in model fit (-2 log likelihood) against an otherwise identical model that did not include a fixed effect of condition.

3.1.2. Results—The average proportions of fixations to each object type in the visual display are plotted in Figure 6A. Fixations were time locked to the onsets of the verb, determiner, adjective, and noun. The full plotted window extends from the onset of the verb to 200 ms past the offset of the direct object noun. Our critical comparison involves looks to the verb inconsistent objects (e.g., cars) following the adjective modifier, which are plotted across a zoomed-in window in Figure 6B, to magnify potential differences between conditions.

We submitted eye movements to separate mixed-effects models to test (A) for anticipation, by comparing looks to the verb-consistent objects (e.g., cakes), and (B) for context-independent processing, by comparing looks to the verb-inconsistent objects (e.g., cars). We did not directly compare verb-consistent and inconsistent objects because objects did not rotate between verb conditions (e.g., *car* was never the target), and thus differences could be due to saliency. Both models included a fixed effect of adjective consistency (consistent vs. inconsistent). Below, we report means and standard errors by participants for both the proportions of fixations data and (empirical-logit) transformed data.

The analysis of verb-consistent objects (e.g., *white* vs. *brown cake*, when hearing “The boy will eat the white...”) revealed a reliable effect of adjective consistency in both the participants analysis, $Estimate = 2.88, SE = 0.15, \chi^2 = 64.48, p < .001$, and items analysis, $Estimate = 2.90, SE = 0.17, \chi^2 = 45.82, p < .001$, with reliably more fixations to adjective-consistent objects (proportions of fixations: $M = 0.71, SE = 0.03$; transformed: $M = 0.91, SE = 0.16$) as compared to adjective-inconsistent objects (proportions of fixations: $M = 0.06, SE = 0.01$; transformed: $M = -2.61, SE = 0.18$). The analysis of verb-inconsistent objects (e.g., *white* vs. *brown car*, when hearing “The boy will eat the white...”) revealed a reliable

effect of adjective consistency in both the participants analysis, $Estimate = 0.72$, $SE = 0.19$, $\chi^2 = 11.27$, $p < .001$, and items analysis, $Estimate = 0.78$, $SE = 0.23$, $\chi^2 = 8.85$, $p < .01$, with reliably more fixations to adjective-consistent objects (proportions of fixations: $M = 0.06$, $SE = 0.01$; transformed: $M = -2.52$, $SE = 0.16$) as compared to adjective-inconsistent objects (proportions of fixations: $M = 0.02$, $SE = 0.01$; transformed: $M = -3.23$, $SE = 0.12$).¹⁰

3.1.3. Discussion—Listeners anticipated target objects based on the joint contextual constraints from the verb and adjective (e.g., *white cake* given “The boy will eat the white...”). This incremental impact of contextual constraints on anticipatory eye movements is consistent with Simulation 2 and prior findings from the VWP (e.g., Kamide, Altmann, et al., 2003). However, as predicted by Simulation 2, listeners also fixated locally coherent competitors more than unrelated distractors, based on context-independent lexical information from the adjective (e.g., *white car* vs. *brown car* given “The boy will eat the white...”). These results reveal that listeners were gradiently sensitive both to context-dependent constraints, based on the sentence context, and context-independent constraints, based on context-independent lexical information, despite the *completely* constraining context.

Closer examination of the materials from Experiment 1, however, raises a potential issue: in our visual displays, the locally coherent competitor (e.g., *white car*) was always in a “contrast set” with a similar object of a different color (e.g., *brown car*). Sedivy (2003) found that eye movement behaviors in the VWP were modulated by the presence of a contrast set: for example, listeners hearing “Pick up the yellow banana” in the context of a *yellow banana* and *yellow notebook* were equally likely to look to both items during the adjective. However, when another item of the same type as the target was also included (e.g., *brown banana*), which created a contrast set (among the *bananas*), fixations favored the *yellow banana* over the *yellow notebook* in that same adjective window. Sedivy (2003)'s results suggest that listeners rapidly map adjectives (from the unfolding language) onto specific items in the visual display that are in a contrast set (presumably, to disambiguate between them). Consequently, listeners in Experiment 1 may have fixated the locally coherent competitor simply because it happened to always be in a contrast set. To address this concern, locally coherent competitors were not in a contrast set in Experiment 2.

3.2. Experiment 2

We tested for bottom-up interference in contexts similar to Experiment 1, but in which the verb-inconsistent objects were not in a “contrast set.” Our predictive contexts involving strong verb selectional restrictions were unchanged, except that we replaced the verb- and adjective-inconsistent objects (e.g., *brown car*) with a completely unrelated distractor (e.g., *pink train*). Again, Simulation 2 predicts more looks to the locally coherent competitor (e.g., *white car*) with adjective-consistent sentences (“The boy will eat the white...”) as compared

¹⁰Exclusion of cohort competitors revealed a similar pattern among verb-inconsistent objects: a reliable effect of adjective consistency in both the participants analysis, $Estimate = 0.73$, $SE = 0.19$, $\chi^2 = 14.16$, $p < .001$, and items analysis, $Estimate = 0.86$, $SE = 0.30$, $\chi^2 = 6.31$, $p < .05$.

to adjective-inconsistent sentences (“The boy will eat the brown...”), even in the absence of a contrast set.

3.2.1. Methods

3.2.1.1. Participants: Thirty University of Connecticut undergraduates participated for course credit. All participants were native speakers of English with self-reported normal hearing and self-reported normal or corrected-to-normal vision who had not participated in Experiment 1.

3.2.1.2. Materials: We used identical materials to those used in Experiment 1, except that we replaced one of the verb-inconsistent objects (e.g., *brown car*) in each visual display with a completely unrelated distractor (e.g., *pink train*), which was unrelated to the verb and the two adjectives associated with each sentence (see Figure 7). Sentences were recorded as in Experiment 1, except that they were recorded by a male native speaker of English. The full item list is presented in Appendix C.

3.2.1.3. Design: We used a 1-factor design of adjective consistency (consistent and inconsistent). In the adjective-consistent condition, the locally coherent competitor (e.g., *white car*) in the visual display was consistent with the adjective modifier of the direct object in the sentence (e.g., “The boy will eat the white...”). In the adjective-inconsistent condition, this same competitor was inconsistent with the adjective modifier of the direct object (e.g., “The boy will eat the brown...”). In both conditions, visual displays also included two verb-consistent objects of the same type, each consistent with one of the associated adjectives (e.g., *white cake* and *brown cake*), and a completely unrelated distractor (e.g., *pink train*), which was unrelated to both the verb and the associated adjectives. In our analyses, we compared looks to the same object (e.g., *white car*) in two language contexts, so differences in looks between the contexts cannot be attributed to differences in saliency (this differs from the design in Experiment 1, where saliency was controlled by rotating a given object across e.g., locally coherent competitor and unrelated distractor conditions). We created two lists to rotate objects across the adjective-consistent and -inconsistent conditions, with half of each type on each list. Participants were randomly assigned to a list, and the location of objects in the visual display and the order of sentences was randomized for each participant.

We created 16 additional filler items to counterbalance biases related to the direct object target, which was always paired with a like object (e.g., *white cake* and *brown cake*) in critical trials. In filler trials, the visual display contained two objects of the same type and two unpaired objects, and the target of the trial was always one of the unpaired objects (the target was also predictable, as the only object related to the verb in the sentence, mirroring the critical conditions).

3.2.1.4. Procedure: Procedures were identical to Experiment 1, except that the experiment was slightly longer (approximately 15 minutes instead of 10 minutes, due to the additional filler items).

3.2.2. Results—The average proportions of fixations to each object type in the visual display are plotted in Figure 8A. Our critical comparison involves looks to the locally coherent competitor (e.g., *white car*) with adjective consistent (e.g., “white”) versus inconsistent (e.g., “brown”) sentences, which is plotted across a zoomed in window in Figure 8B. We did not directly compare the verb-inconsistent objects (e.g., *white car* vs. *pink train*), as in Experiment 1, because objects did not rotate across these conditions, and differences could thus be due to saliency. Rather, to control for effects of saliency, we compared looks to the same object (e.g., *white car*) in the two language contexts (e.g., “white” vs. “brown”).

As in Experiment 1, we focused our analyses on the onset of the direct object noun (e.g., “cake”), and we submitted eye movements to a weighted empirical-logit regression (Barr, 2008a). Again, our linear mixed-effects models included fixed effects of condition, and random intercepts for participants/items. The analysis of verb-consistent objects (e.g., *white* vs. *brown cake*, when hearing “The boy will eat the white...”) revealed a reliable effect of adjective consistency in both the participants analysis, $Estimate = 2.66, SE = 0.17, \chi^2 = 84.37, p < .001$, and items analysis, $Estimate = 2.85, SE = 0.22, \chi^2 = 41.62, p < .001$, with reliably more fixations to adjective-consistent objects (proportions of fixations: $M = 0.72, SE = 0.03$; transformed: $M = 0.93, SE = 0.15$) as compared to adjective-inconsistent objects (proportions of fixations: $M = 0.05, SE = 0.02$; transformed: $M = -2.36, SE = 0.13$). The analysis of the locally coherent competitor (e.g., *white car*) with adjective consistent (e.g., “white”) versus inconsistent (e.g., “brown”) sentences revealed a reliable effect of adjective consistency in both the participants analysis, $Estimate = 0.42, SE = 0.17, \chi^2 = 9.60, p < .01$, and items analysis, $Estimate = 0.67, SE = 0.24, \chi^2 = 8.53, p < .01$, with reliably more fixations to the locally coherent competitor with adjective consistent sentences (proportions of fixations: $M = 0.04, SE = 0.02$; transformed: $M = -2.53, SE = 0.13$) as compared adjective inconsistent sentences (proportions of fixations: $M = 0.01, SE = 0.01$; transformed: $M = -2.75, SE = 0.06$).¹¹

3.2.3. Discussion—Listeners again anticipated target objects based on the joint contextual constraints from the verb and adjective (e.g., *white cake* given “The boy will eat the white...”). This incremental impact of contextual constraints on anticipatory eye movements is consistent with Experiment 1, Simulation 2, and prior findings from the VWP (e.g., Kamide, Altmann, et al., 2003). However, as predicted by Simulation 2, listeners again also fixated locally coherent competitors more with consistent rather than inconsistent adjectives, based on context-independent lexical information from the adjective (e.g., *white car* given “The boy will eat the white...” vs. “The boy will eat the brown...”). These results reveal that listeners were gradiently sensitive both to context-dependent constraints, based on the sentence context, and context-independent constraints, based on context-independent lexical information, despite the *completely* constraining context, and the lack of a contrast set.

¹¹Exclusion of cohort competitors revealed a similar pattern among verb-inconsistent objects: a reliable effect of adjective consistency in both the participants analysis, $Estimate = 0.32, SE = 0.15, \chi^2 = 8.32, p < .01$, and items analysis, $Estimate = 0.64, SE = 0.30, \chi^2 = 5.19, p < .05$.

One potential alternative explanation of the *white car* effect is that listeners were using low-level visual cues about color from their periphery. Consequently, they may have erroneously launched eye movements to locally coherent competitors because they had sufficient visual acuity to know they were white, but they did not have sufficient visual acuity to know whether they were edible. In part, the rapid, massive reduction in fixations to cars during “eat” does suggest that listeners had good visual acuity in general (i.e., to know that the cars were cars, and thus inedible). Moreover, the current results are also consistent with a number of closely related findings (e.g., Borovsky, Elman, & Fernald, 2012; Kamide et al., 2003; Kukona et al., 2011), which do not depend on low-level visual cues to color. For example, Kukona et al. (2011) found that listeners fixated *policeman* (a locally coherent competitor like *white car*) more than the distractors *gardener* or *surfer* when hearing “Toby will arrest the,” although *all* of these objects are quite visually distinct from *crook*, the predictable direct object.

4. General Discussion

A growing body of research reveals that language users rapidly form linguistic representations that best satisfy the contextual constraints provided by the sentence, discourse, and visual context, *and* that they also activate representations that are outside the scope of these constraints. While anticipatory effects (e.g., Altmann & Kamide, 1999; Chambers & San Juan, 2008; Kamide, Altmann, et al., 2003; Kamide, Scheepers, et al., 2003; Knoeferle & Crocker, 2006, 2007) reveal that the language processing system rapidly uses the contextual constraints to predict upcoming linguistic input, effects of context-independent information reveal that the system simultaneously activates representations that are not consistent with the context (e.g., Allopenna et al., 1998; Dahan & Tanenhaus, 2005; Huettig & Altmann, 2005; Kukona et al., 2011; Myung et al., 2006; Swinney, 1979; Tabor et al., 2004; Tanenhaus et al., 1979; Yee & Sedivy, 2006). In two simulations, we showed that the self-organizing neural network framework predicts classic results showing both bottom-up interference from context-independent constraints (e.g., Swinney, 1979; Tanenhaus et al., 1979) and rapid context integration (e.g., Altmann & Kamide, 1999). In two experiments, we also found evidence for bottom-up interference even in highly constraining anticipatory VWP contexts involving strong verb selection restrictions, as predicted by the simulations. For example, listeners hearing “The boy will eat the white...” predominately looked to targets (e.g., *white cake*) predicted by the joint contextual constraints from the verb and adjective (consistent with, e.g., Kamide, Altmann, et al., 2003). However, they also made looks following the adjective to locally coherent competitors (e.g., *white car* vs. unrelated distractors). These results suggest that listeners were sensitive to the fit between “white” and the *white car*, although this object was not predicted by the verb selectional restrictions of the sentence context.

Note that our self-organizing neural networks do not have a grammar in the form of a system of rules that operate on symbolic objects. Rather, the networks’ interacting units encode “grammatical constraints,” which the network learns through the interactions of its units with the language. Critically, these grammatical constraints, which shape how the units interact, do not enforce total grammatical coherence, thus predicting the activation of the verb sense of “rose” in a noun context in Simulation 1, and fixations to *white car* following

“Eat the white...” in Simulation 2. By contrast, our experimental findings provide evidence against approaches that *do* use a grammar to enforce global coherence by only allowing for the activation of structures that fully satisfy the global constraints (we describe Bayesian models that make this assumption below; e.g., Hale, 2001; Levy, 2008b).

Moreover, the pattern of results we observed over the course of training in the self-organizing neural network (see Figure 4) also provides rich predictions about both human language development and individual differences in language skill. In Simulation 2, we found that activation of the locally coherent competitor (e.g., *white car*) was largest early in training, and diminished over the course of training. One prediction that follows from this pattern, given the assumption that the amount of training in the network maps onto *age*, is that bottom-up interference will be largest early in development (e.g., in children), and will diminish over development (e.g., in adults). Another prediction, given the assumption that the amount of error in the network maps (inversely) onto language *skill*, is that bottom-up interference will be larger in less skilled language users (i.e., with higher error) as compared to more skilled language users (i.e., with lower error).

Both of these predictions were supported by recent experimental findings by Borovsky et al. (2012). They found that both adults and children hearing “The pirate hides the...,” while viewing a visual display with a *treasure chest*, *ship*, *dog bone*, and *cat*, anticipatorily fixated the *treasure chest*, which was predicted by the combined contextual constraints from the subject and verb (compatible with Kamide, Altmann, et al., 2003). However, their visual display also included a locally coherent competitor: *dog bone* was predicted by the verb “hides,” but not the subject “pirate.” They observed a numerically larger interference effect (as reflected in the length of the time window in which there were reliably more fixations to *dog bone* vs. distractors like *cat*) in children than adults, and in language users with lower vocabulary skill as compared to higher vocabulary skill (although they did not directly compare these groups; bottom-up interference was not a central focus of their study). These findings further suggest that our self-organizing predictions should not be limited to the behavior of the network at asymptote (i.e., where it predicts no *white car* vs. *brown car* difference).

4.1. Relation to previous findings

Our experimental findings are consistent with the results of Kukona et al. (2011), who found that listeners hearing a sentence like “Toby arrests the...” anticipatorily fixated a predictable (i.e., based on the contextual constraints) patient of the verb (e.g., *crook*). However, listeners were approximately as likely to make “anticipatory” fixations to a good agent of the verb (e.g., *policeman*), based on the thematic fit between the verb and agent, despite the fact that the agent role was already filled. In a second experiment with passive sentences (e.g., “Toby was arrested by the...”), in which more time elapsed between the verb and direct object noun, and more syntactic cues were available (e.g., “was...”), Kukona et al. found even clearer evidence for graded influences, with more anticipatory fixations to the *policeman* (e.g., the predictable direct object of the passive sentence) than the *crook*, but also more fixations to the *crook* than the unrelated distractors. These gradient effects with *policeman* and *crook* closely parallel our results with *white cake* and *white car*, suggesting that there is

a balance between top-down context effects and bottom-up interference effects that varies with the strength of the two sources of information.

In addition, we have argued throughout that the activation of representations that are at odds with contextual constraints (i.e., outcomes with zero probability in a sentence context) represents a kind of interference that *impedes* processing (e.g., by shifting listeners' attention away from target objects, and toward non-target objects). However, there may also be *facilitory* side effects of this activation, as emphasized in Kukona et al. (2011). For example, the activation of *policeman* when hearing "Toby will arrest the" may facilitate discourse processing because situations that involve "arrest" often also involve policemen (although, as we have suggested, this activation is highly transient). Indeed, activating representations that are not explicitly referred to linguistically, but which are associated or related to the linguistic input, may also facilitate integration with prior knowledge and subsequent memory. On the other hand, the current results suggest that not all bottom-up interference is necessarily facilitative (i.e., it is not clear that eye movements to *white car* similarly facilitate discourse processing).

Our results are also consistent with evidence for "verb effects" in Kamide, Altmann, et al. (2003; see also Borovsky et al., 2012). Again, they found that listeners hearing "The girl will ride the..." anticipatorily fixated a *carousel* more than a *motorbike*, based on the contextual constraints. However, they also found that listeners anticipatorily fixated the subject-conflicting *motorbike* more during "the" when hearing "The girl will ride the..." as compared to "The girl will taste the..." These results suggest that listeners *were* sensitive to the thematic fit between "ride" and *motorbike*, and that they were activating this object even though it conflicted with the context, consistent with both Kukona et al. (2011) and our current results.

Critically, however, our findings extend the results of Borovsky et al. (2012), Kamide, Altmann, et al. (2003), and Kukona et al. (2011), as well as other evidence for context-independent processing in language, in a number of important ways. The current results reveal that these effects compete against strongly constraining verb selectional restrictions, which have previously been shown to drive robust anticipatory effects (e.g., Altmann & Kamide, 1999). For example, our sentence contexts involved a constraining verb (e.g., "The boy will eat the..."), while the preceding sentence contexts in Borovsky et al., Kamide, Altmann, et al., and Kukona et al. involved only a weakly constraining noun phrase (e.g., "The girl..." or "Toby..."). As well, the effects in the current study were based on adjectival information (e.g., linguistic-visual match that runs counter to the sentence context), rather than thematic association.

Finally, additional findings on ambiguity resolution in the VWP also provide a means for integrating the results of Tanenhaus et al. (1995) and Dahan and Tanenhaus (2004), which we described in the Introduction. Recall that Tanenhaus et al. (1995) demonstrated that contextual constraints from the visual context (e.g., two-referents) continuously and immediately constrained the interpretation of syntactic ambiguities (e.g., "Put the apple on the towel in the box"). However, Snedeker and Trueswell (2004; see also Novick, Thompson-Schill, & Trueswell, 2008) showed that these contextual effects can be

modulated by lexical biases in adult listeners. They compared ambiguous sentences like “Choose/Tickle the cow with the stick:” whereas the verb “Choose” favors a modifier interpretation (e.g., choose the cow that's holding the stick), the verb “Tickle” favors an instrument interpretation (e.g., tickle the cow by using the stick). Consistent with Tanenhaus et al. (1995), they found a reliable effect of visual context: two-referent contexts (e.g., a visual display with a *cow holding a stick* and a *cow holding nothing*) favored modifier interpretations relative to one-referent contexts. However, they also found a reliable effect of verb bias: instrument-biased verbs like “Tickle” favored instrument interpretations relative to modifier-biased verbs like “Choose.”

Similarly, recall that Dahan and Tanenhaus (2004) demonstrated that contextual constraints from a sentence context (e.g., “Nog nooit klom een bok...” / “Never before climbed a goat...”) continuously and immediately constrained spoken word recognition (e.g., eliminating cohort effects from competitors like “bot” / “bone”). The self-organization approach, in keeping with constraint-based frameworks in general, predicts that such effects should show evidence of being graded. Indeed, in a second experiment in which they cross-spliced the onset and vowel of a cohort competitor (e.g., “bo” from “bot”) with the offset of the target noun (e.g., “k” from “bok”), they found more cohort fixations. While these results suggest that contextual constraints can *overwhelm* bottom-up constraints, they also reveal that context-independent processing effects can also be made more pronounced by shifting the relative balance of the constraints (e.g., by using verbs that bias context-inconsistent structures, Snedeker & Trueswell, 2004; or by providing more bottom-up support for context-inconsistent competitors via cross-splicing, Dahan & Tanenhaus, 2004).

4.2. Bayesian belief update

Bayesian models of sentence processing (e.g., Bicknell & Levy, 2010; Levy, 2008a; Narayanan, & Jurafsky, 2002; see also closely related probabilistic models: e.g., Hale, 2001; Jurafsky, 1996; Levy, 2008b) have also been widely used to model human sentence processing. Recent Bayesian implementations have assumed that language users have a probabilistic model of the way that words are sequenced in their language, such that during online processing, language users distribute probabilities over grammatical structures that are consistent with the observed input. Crucially, this probability distribution also allows language users to assign probabilities to future events (i.e., structures that are likely to follow from the current input). This approach is called “Bayesian belief update” because Bayes’ Rule specifies the way that each new piece of information revises current expectations.

A number of recent proposals (e.g., Hale, 2001; Levy, 2008b) have assumed that contextual information places *very* strong constraints on processing: language users are assumed to be “optimal” in the sense that they only distribute probabilities over structures that are *fully* consistent with the context. These proposals very naturally predict findings showing “rapid, immediate context integration” (see Section 1.1). However, they also predict that language users should *not* show bottom-up interference effects (see Section 1.2). For example, listeners in the current experiments would be predicted to only distribute probabilities over cakes, which satisfy the selectional restrictions of “eat,” and not cars, thus predicting no

white car vs. *brown car* difference. Thus, the current findings provide strong evidence against this class of models.

However, more recent proposals (e.g., Bicknell & Levy, 2010; Gibson, Bergen, & Piantadosi, 2013; Levy, 2008a; for a closely related proposal concerned with spoken word recognition, see Norris & McQueen, 2008) have assumed that language users also have “uncertainty” about contextual information. Consequently, these proposals allow for probabilities to be distributed over alternative grammatical structures that are not fully consistent with the context (i.e., as specified by the bottom-up input). For example, Levy, Bicknell, Slattery, and Rayner (2009) argue that the local coherence effects observed by Tabor et al. (2004) are a consequence of language users’ uncertainty about the sentence context “The coach smiled at...” Accordingly, when readers encounter “tossed” in “The coach smiled at the player tossed the frisbee,” they entertain alternative interpretations of the prior words in the sentence (e.g., instead of “at,” they read “and,” which is orthographically similar) due to uncertainty about this input (crucially, “and” also allows for the more predictable sentence, “The coach smiled *and* the player tossed the frisbee”). Indeed, such an approach offers a very clear alternative to self-organization: rather than assuming that language users activate structures that conflict with the context (which is implicitly assumed to have no uncertainty) due to the self-organizing nature of structure formation, such an approach assumes that context-conflicting structures are activated because contextual information itself is uncertain. Nevertheless, Levy et al.’s proposal relies on the fact that the sentence context in Tabor et al. is (orthographically) confusable with other sentence contexts that *do* allow for the locally coherent parse. By contrast, it is not apparent that there are also phonologically similar parses in the case of Swinney (1979) and Tanenhaus et al. (1979; i.e., Simulation 1), or Experiments 1 and 2 (i.e., Simulation 2), that would likewise lead language users to expect the respective “locally coherent” structures in each of those cases.

4.3. Relation to other theories of language processing

Our experimental results also bear on a number of other theories of sentence processing that are closely related to self-organization. Ferreira and colleagues (e.g., Ferreira & Patson, 2007) proposed the Good-Enough Theory of sentence processing that, like self-organization, assumes that contextual coherence of linguistic representations is not fully maintained. According to this theory, linguistic representations or structures may be neither “accurate” nor “detailed;” rather, representations are often cobbled together, using quick and dirty heuristics, so as to be “good enough” for the task at hand. Consistent with Good-Enough Theory, they have shown that readers misinterpret garden path sentences like “While Mary bathed the baby played in the crib,” to mean both that Mary bathed the baby, as well as that Mary bathed herself (although only the latter is grammatically accurate; Christianson, Hollingworth, Halliwell, & Ferreira, 2001). Our results are generally consistent with Good-Enough Theory’s claims that structures can be cobbled together despite their contextual incoherence. However, our results suggest that such heuristics may even be used for straightforward, routine sentences, which presumably place very little burden on the language processing system (e.g., relative to NP-Z garden path sentences), consistent with evidence for closely related misinterpretation effects in passive sentences (e.g., Ferreira, 2003). For example, our results suggest that listeners may have a heuristic that leads them to

rely on information within the noun phrase (e.g., adjective information), rather than outside it (e.g., verb information), in anticipating nouns. However, on our view these effects are not due to the system literally employing *heuristics*, but instead are emergent effects of a self-organizing system in which structures that are not fully coherent or compatible with the global context can become strongly activated.

Our results are also consistent with the feature-based approach of Altmann and Kamide (2007). They describe a linking hypothesis for language-mediated eye movements that accounts for findings showing that listeners fixate items in a visual context that are variously related to an unfolding utterance (e.g., Allopenna et al., 1998; Altmann & Kamide, 1999; Dahan & Tanenhaus, 2005; Huettig & Altmann, 2005; Myung et al., 2006; Yee & Sedivy, 2006). According to their approach, listeners activate representations in a “mental world” that correspond to input from both language (e.g., words) and the visual world (e.g., objects within a visual display to which the discourse refers). Critically, these representations are featural in nature, and they include information like phonological form, verb selectional restrictions, physical form (visual, tactile, etc.), category membership, function, and so forth. Visual representations receive an activation boost when they share features with the linguistic input, which increases the likelihood of a fixation to the relevant item. Very much in this line, the behavior of the self-organizing neural network in Simulation 2 appears to similarly reflect the influence of verb- and adjective-based featural-type (e.g., edible/rideable; white/brown) information.

In summary, we have presented new evidence for bottom-up interference in highly constraining anticipatory contexts in the VWP. These results, in conjunction with prior findings, suggest that the language systems acts on, and maintains, transient information structures at many levels of language structure that are outside the scope of, or in conflict with, the contextual constraints. Critically, these phenomena are hallmarks of the self-organizing neural network framework. Moreover, our simulations, which address the integration of incoming lexical information (i.e., an incoming word) with sentence context information (i.e., from preceding words in an unfolding utterance), explicitly show that the framework predicts a range of findings that involve gradient effects of both context-dependent and context-independent information.

Acknowledgments

Preliminary results from this study were presented at the 16th Annual Conference on Architectures and Mechanisms for Language Processing in York, UK (September, 2010). We thank Klinton Bicknell, Florian Jaeger, and Paola Dussias for their helpful feedback on earlier versions of this manuscript. We also thank Karen Aicher, Joshua Coppola, Reed Helms, Matthew Madruga, and Elyse Noccioli for their help with this project. We gratefully acknowledge support from NICHD grant HD40353 to Haskins Laboratories, Predoctoral NRSA HD060414 to AK, NSF grant 1059662 to WT, and NSF grant BCS CAREER 0748684 to JSM.

Appendix

Appendix A

Grammar used to generate the training sentences in Simulation 1. Within each syntactic class (e.g., Adj, Adv, Det, N, V), lexical items were used with equal probabilities.

1.0 S → NP VP
 0.5 NP → Det N
 0.5 NP → Det Adj N
 0.5 VP → V
 0.5 VP → Adv V

Adj: pretty, white

Adv: slowly, quickly

Det: the, your

N: dandelion, **rose**, tulip

V: grew, **rose**, wilted

Appendix B

Materials used in Experiment 1, which were closely based on materials from Altmann and Kamide (1999). Sentences always had a predictable direct object target, and one of two possible color adjectives. Visual displays included four objects, reflecting the crossing of the target (e.g., cake) and competitor (e.g., car) objects with the color adjectives (e.g., brown/white).

Item	Sentence	Competitor
1.	The boy will eat the brown/white cake.	car
2.	The woman will drink the white/yellow beverage.	cheese
3.	The officer will arrest the green/purple man.	house
4.	The woman will bathe the red/yellow puppy.	stool
5.	The boy will bounce the blue/red ball.	plane
6.	The man will climb the brown/gray ladder.	bunny
7.	The housewife will fry the brown/gray mushroom.	knife
8.	The doctor will inject the black/orange cat.	microscope
9.	The woman will play the black/brown piano.	table
10.	The woman will read the green/purple book.	bag
11.	The man will repair the gray/green washer.	trash
12.	The girl will ring the purple/red bell.	bricks
13.	The man will sail the red/yellow boat.	bird
14.	The man will smoke the blue/orange pipe.	glasses
15.	The boy will walk the brown/white dog.	duck
16.	The businessman will wear the black/green hat.	chair

Appendix C

Materials used in Experiment 2. Sentences were identical to Experiment 1 (see Appendix B). The visual display included four objects, including two objects reflecting the crossing of the target (e.g., cake) with the color adjectives (e.g., brown/white). However, the visual display also included an object that was consistent with one of the color adjectives (e.g., white car) as well as an unrelated distractor (e.g., pink train).

Item	Competitor, unrelated distractor	Item	Competitor, unrelated distractor
1.	white car, pink train	9.	black table, orange phone
2.	yellow cheese, brown plant	10.	green bag, blue cup
3.	purple house, black broom	11.	gary trash, pink mirror
4.	yellow stool, blue rocker	12.	purple bricks, black drum
5.	red plane, orange bike	13.	red bird, white sun
6.	gray bunny, yellow cactus	14.	blue glasses, green hen
7.	brown knife, pink scale	15.	brown duck, gray clock
8.	orange microscope, gray TV	16.	green chair, yellow magnifier

References

- Alloppenna P, Magnuson JS, Tanenhaus MK. Tracking the time course of spoken word recognition using eye-movements: Evidence for continuous mapping models. *Journal of Memory and Language*. 1998; 38:419–439.
- Altmann GTM, Kamide Y. 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*. 2007; 57:502–518.
- Altmann GTM, Kamide Y. Incremental interpretation at verbs: restricting the domain of subsequent reference. *Cognition*. 1999; 73:247–264. [PubMed: 10585516]
- Barr DJ. Analyzing ‘visual world’ eyetracking data using multilevel logistic regression. *Journal of Memory and Language*. 2008a; 59:457–474.
- Barr DJ. Pragmatic expectations and linguistic evidence: Listeners anticipate but do not integrate common ground. *Cognition*. 2008b; 109:18–40. [PubMed: 18760407]
- Barr DJ, Levy R, Scheepers C, Tily HJ. Random-effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*. 2013; 68:255–278.
- Bicknell, K.; Levy, R. Proceedings of the 48th annual meeting of the Association for Computational Linguistics (ACL). Association for Computational Linguistics; Uppsala, Sweden: 2010. A rational model of eye movement control in reading.; p. 1168-1178.
- Bicknell, K.; Levy, R.; Demberg, V. Proceedings of the 35th Annual Meeting of the Berkeley Linguistics Society. Berkeley Linguistics Society; Berkeley, CA: 2010. Correcting the incorrect: Local coherence effects modeled with prior belief update.; p. 13-24.
- Borovsky A, Elman J, Fernald A. Knowing a lot for one's age: Vocabulary skill and not age is associated with the timecourse of incremental sentence interpretation in children and adults. *Journal of Experimental Child Psychology*. 2012; 112:417–36. [PubMed: 22632758]
- Chambers CG, San Juan V. Perception and presupposition in real-time language comprehension: Insights from anticipatory processing. *Cognition*. 2008; 108:26–50. [PubMed: 18262509]
- Christianson K, Hollingworth A, Halliwell J, Ferreira F. Thematic roles assigned along the garden path linger. *Cognitive Psychology*. 2001; 42:368–407. [PubMed: 11368528]
- Cooper RM. The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive Psychology*. 1974; 6:84–107.

- Dahan D, Tanenhaus MK. Looking at the rope when looking for the snake: Conceptually mediated eye movements during spoken-word recognition. *Psychonomic Bulletin & Review*. 2005; 12:453–459. [PubMed: 16235628]
- Dahan D, Tanenhaus MK. Continuous mapping from sound to meaning in spoken-language comprehension: Evidence from immediate effects of verb-based constraints. *Journal of Experimental Psychology: Learning, Memory & Cognition*. 2004; 30:498–513.
- Eberhard KM, Spivey-Knowlton MJ, Sedivy JC, Tanenhaus MK. Eye-movements as a window into spoken language comprehension in natural contexts. *Journal of Psycholinguistic Research*. 1995; 24:409–436. [PubMed: 8531168]
- Elman JL. Finding structure in time. *Cognitive Science*. 1990; 14:179–211.
- Elman JL. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*. 1991; 7:195–224.
- Ferreira F. The misinterpretation of noncanonical sentences. *Cognitive Psychology*. 2003; 47:164–203. [PubMed: 12948517]
- Ferreira F, Patson N. The good enough approach to language comprehension. *Language and Linguistics Compass*. 2007; 1:71–83.
- Fine AB, Jaeger TF. Evidence for implicit learning in syntactic comprehension. *Cognitive Science*. 2013; 37:578–591. [PubMed: 23363004]
- Fodor JA. *Modularity of mind: An essay on faculty psychology*. MIT Press; Cambridge, MA: 1983.
- Fodor JD. Unambiguous triggers. *Linguistic Inquiry*. 1998; 29:1–36.
- Gibson E, Bergen L, Piantadosi ST. Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*. 2013; 110:8051–8056.
- Gordon, DM. *Ant encounters: Interaction networks and colony behavior (primers in complex systems)*. Princeton University Press; Princeton, NJ: 2010.
- Hale, J. *Proceedings of NAACL Volume 2*. Pittsburg, PA: 2001. A probabilistic earley parser as a psycholinguistic model.; p. 159-166.
- Hölldobler, B.; Wilson, EO. *The Superorganism: The Beauty, Elegance, and Strangeness of Insect Societies*. W. W. Norton & Company; 2008.
- Huetig F, Altmann GTM. Word meaning and the control of eye fixation: semantic competitor effects and the visual world paradigm. *Cognition*. 2005; 96:23–32.
- Jaeger TF, Snider NE. Alignment as a consequence of expectation adaptation: Syntactic priming is affected by the prime's prediction error given both prior and recent experience. *Cognition*. 2013; 127:57–83. [PubMed: 23354056]
- Jensen, HJ. *Self-organized criticality: Emergent complex behavior in physical and biological systems*. Cambridge university press; New York: 1998.
- Jurafsky D. A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science*. 1996; 20:137–194.
- Kamide Y, Altmann GTM, Haywood SL. Prediction and thematic information in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*. 2003; 49:133–156.
- Kamide Y, Scheepers C, Altmann GTM. Integration of syntactic and semantic information in predictive processing: Cross-linguistic evidence from German and English. *Journal of Psycholinguistic Research*. 2003; 32:37–55. [PubMed: 12647562]
- Keller EF, Segel LA. Conflict between positive and negative feedback as an explanation for the initiation of aggregation in slime mould amoebae. *Nature*. 1970; 227:1365–1366. [PubMed: 5465943]
- Knoeferle P, Crocker MW. The coordinated interplay of scene, utterance, and world knowledge: Evidence from eye tracking. *Cognitive Science*. 2006; 30:481–529. [PubMed: 21702823]
- Knoeferle P, Crocker MW. The influence of recent scene events on spoken comprehension: Evidence from eye movements. *Journal of Memory and Language*. 2007; 57:519–543.

- Konieczny, L.; Weldle, H.; Wolfer, S.; Müller, D.; Baumann, P. Anaphora and local coherences.. In: Ohlsson, S.; Catrambone, R., editors. Proceedings of the 32nd Annual Conference of the Cognitive Science Society. Cognitive Science Society; Austin, TX: 2010. p. 1204-1209.
- Konieczny, L.; Müller, D.; Hachmann, W.; Schwarzkopf, S.; Wolfer, S. Local syntactic coherence interpretation. Evidence from a visual world study.. In: Taatgen, NA.; van Rijn, H., editors. Proceedings of the 31th Annual Conference of the Cognitive Science Society. Cognitive Science Society; Austin, TX: 2009. p. 1133-1138.
- Kukona A, Fang S, Aicher KA, Chen H, Magnuson JS. The time course of anticipatory constraint integration. *Cognition*. 2011; 119:23–42. [PubMed: 21237450]
- Kukona A, Tabor W. Impulse processing: A dynamical systems model of incremental eye movements in the visual world paradigm. *Cognitive Science*. 2011; 35:1009–1051. [PubMed: 21609355]
- Levy R. A noisy-channel model of rational human sentence comprehension under uncertain input. Proceedings of the 13th conference on empirical methods in natural language processing. 2008a: 234–243.
- Levy R. Expectation-based syntactic comprehension. *Cognition*. 2008b; 106:1126–1177. [PubMed: 17662975]
- Levy R, Bicknell K, Slattery T, Rayner K. Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences*. 2009; 106:21086–21090.
- MacDonald MC, Pearlmutter NJ, Seidenberg MS. The lexical nature of syntactic ambiguity resolution. *Psychological Review*. 1994; 101:676–703. [PubMed: 7984711]
- Magnuson JS, Dixon J, Tanenhaus MK, Aslin RN. The dynamics of lexical competition during spoken word recognition. *Cognitive Science*. 2007; 31:133–156. [PubMed: 21635290]
- Magnuson JS, Tanenhaus MK, Aslin RN. Immediate effects of form-class constraints on spoken word recognition. *Cognition*. 2008; 108:866–873. [PubMed: 18675408]
- Magnuson JS, Tanenhaus MK, Aslin RN. Simple recurrent networks and competition effects in spoken word recognition. *University of Rochester Working Papers in the Language Science*. 2000; 1:56–71.
- Magnuson JS, Tanenhaus MK, Aslin RN, Dahan D. The time course of spoken word recognition and learning: Studies with artificial lexicons. *Journal of Experimental Psychology: General*. 2003; 132:202–227. [PubMed: 12825637]
- Marcus, GF. *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. MIT Press; Cambridge, MA: 2001.
- Marée AFM, Hogeweg P. How amoeboids self-organize into a fruiting body: Multicellular coordination in *Dictyostelium discoideum*. *Proceedings of the National Academy of Sciences*. 1999; 98:3879–3883.
- McClelland JL, Elman JL. The TRACE Model of Speech Perception. *Cognitive Psychology*. 1986; 18:1–86. [PubMed: 3753912]
- McClelland JL, Rumelhart DE. An interactive activation model of context effects in letter perception: Part I. An account of basic findings. *Psychological Review*. 1981; 88:375–407.
- Miikkulainen, R. *Subsymbolic natural language processing: An integrated model of scripts, lexicon and memory*. MIT Press; Cambridge, MA: 1993.
- Myung J, Blumstein S, Sedivy J. Playing on the typewriter, typing on the piano: Manipulation knowledge of objects. *Cognition*. 2006; 98:223–243. [PubMed: 16399263]
- Narayanan, S.; Jurafsky, D. A bayesian model predicts human parse preference and reading times in sentence processing.. In: Dietterich, TG.; Becker, S.; Ghahramani, Z., editors. *Advances in neural information processing systems*. Vol. 14. MIT Press; Cambridge, MA: 2002. p. 59-65.
- Norris D, McQueen JM. Shortlist B: A Bayesian model of continuous speech recognition. *Psychological Review*. 2008; 115:357–395. [PubMed: 18426294]
- Novick JM, Thompson-Schill S, Trueswell JC. Putting lexical constraints in context into the visual-world paradigm. *Cognition*. 2008; 107:850–903. [PubMed: 18279848]
- Reynolds CW. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*. 1987; 21:25–34.

- Sedivy J. Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research*. 2003; 32:3–23. [PubMed: 12647560]
- Sedivy J, Tanenhaus M, Chambers C, Carlson G. Achieving incremental semantic interpretation through contextual representation. *Cognition*. 1999; 71:109–147. [PubMed: 10444906]
- Shillcock, R.; Bard, EG. Modularity and the processing of closed-class words. In: Altmann, GTM.; Shillcock, R., editors. *Cognitive models of speech processing: The second Sperlonga meeting*. Lawrence Erlbaum Associates; Mahwah, NJ: 1993. p. 163-185.
- Snedeker J, Trueswell JC. The developing constraints on parsing decisions: The role of lexical-biases and referential scenes in child and adult sentence processing. *Cognitive Psychology*. 2004; 49:238–299. [PubMed: 15342261]
- Sole, RV.; Bascompte, J. *Self-Organization in Complex Ecosystems*. Princeton University Press; Princeton, NJ: 2006.
- Spivey, MJ. *The continuity of mind*. Oxford University Press; New York: 2007.
- Spivey MJ, Tanenhaus MK. Syntactic ambiguity resolution in discourse: Modeling the effects of referential context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1998; 24:1521–1543.
- Swinney D. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*. 1979; 18:645–659.
- Tabor W. Fractal encoding of context-free grammars in connectionist networks. *Expert Systems: The International Journal of Knowledge Engineering and Neural Networks*. 2000; 17:41–56.
- Tabor W, Cho PW, Szudlarek E. Fractal analysis illuminates the form of connectionist structural gradualness. *Topics in Cognitive Science*. 2013; 5:634–667. [PubMed: 23798028]
- Tabor W, Hutchins S. Evidence for Self-Organized Sentence Processing: Digging In Effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2004; 30:431–450.
- Tabor W, Galantucci B, Richardson D. Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language*. 2004; 50:355–370.
- Tanenhaus MK, Leiman J, Seidenberg M. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*. 1979; 18:427–440.
- Tanenhaus MK, Spivey-Knowlton M, Eberhard K, Sedivy J. Integration of visual and linguistic information in spoken language comprehension. *Science*. 1995; 268:1632–1634. [PubMed: 7777863]
- Trueswell, JC.; Tanenhaus, MK. Toward a lexicalist framework for constraint-based syntactic ambiguity resolution.. In: Clifton, C.; Frazier, L.; Rayner, K., editors. *Perspectives in Sentence Processing*. Lawrence Erlbaum Assoc.; Hillsdale, NJ: 1994. p. 155-179.
- Vosse T, Kempen G. Syntactic structure assembly in human parsing: A computational model based on competitive inhibition and a lexicalist grammar. *Cognition*. 2000; 75:105–143. [PubMed: 10771275]
- Yee E, Sedivy J. Eye movements reveal transient semantic activation during spoken word recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 2006; 32:1–14.

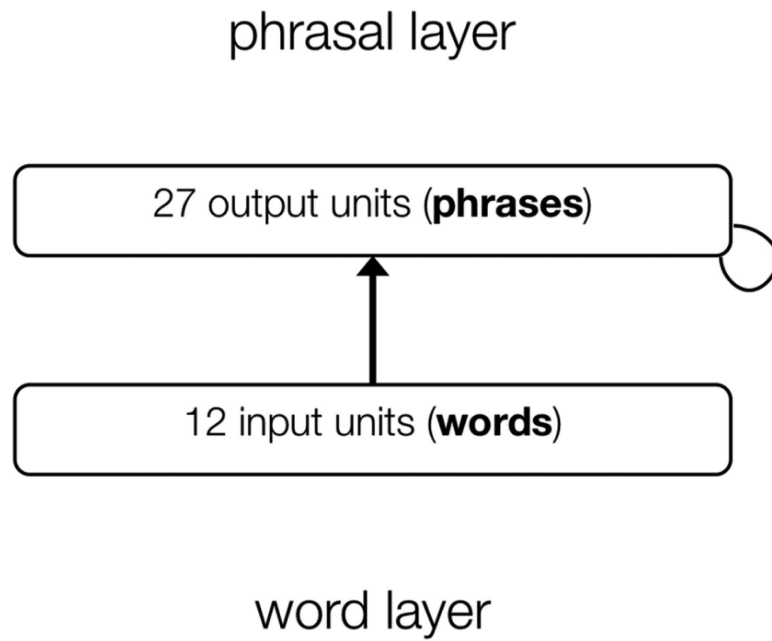


Figure 1. Simulation 1: Architecture of the self-organizing neural network. The network had feedforward connections from the word input layer to the phrasal output layer, and recurrent connections within the phrasal output layer.

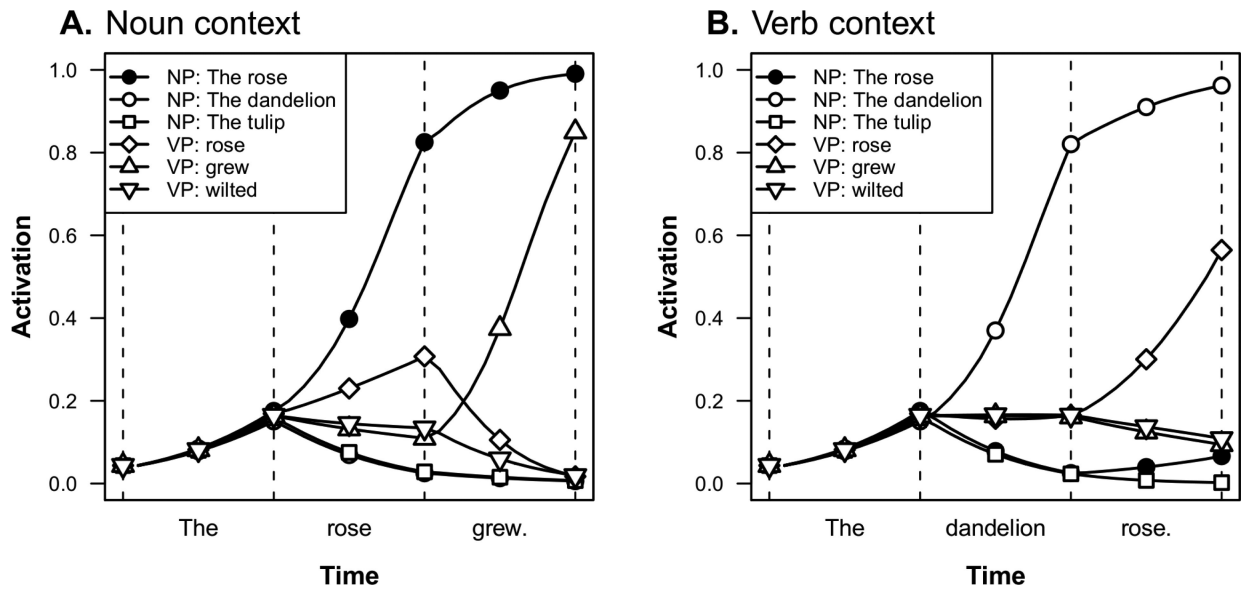


Figure 2. Simulation 1: Activations of selected phrasal output nodes in the noun context “The rose grew” (A) and the verb context “The dandelion rose” (B).




Figure 3. Simulation 2: Architecture of the self-organizing neural network. The network had feedforward connections from the language input layer to the visual output layer, and recurrent connections within the visual output layer.

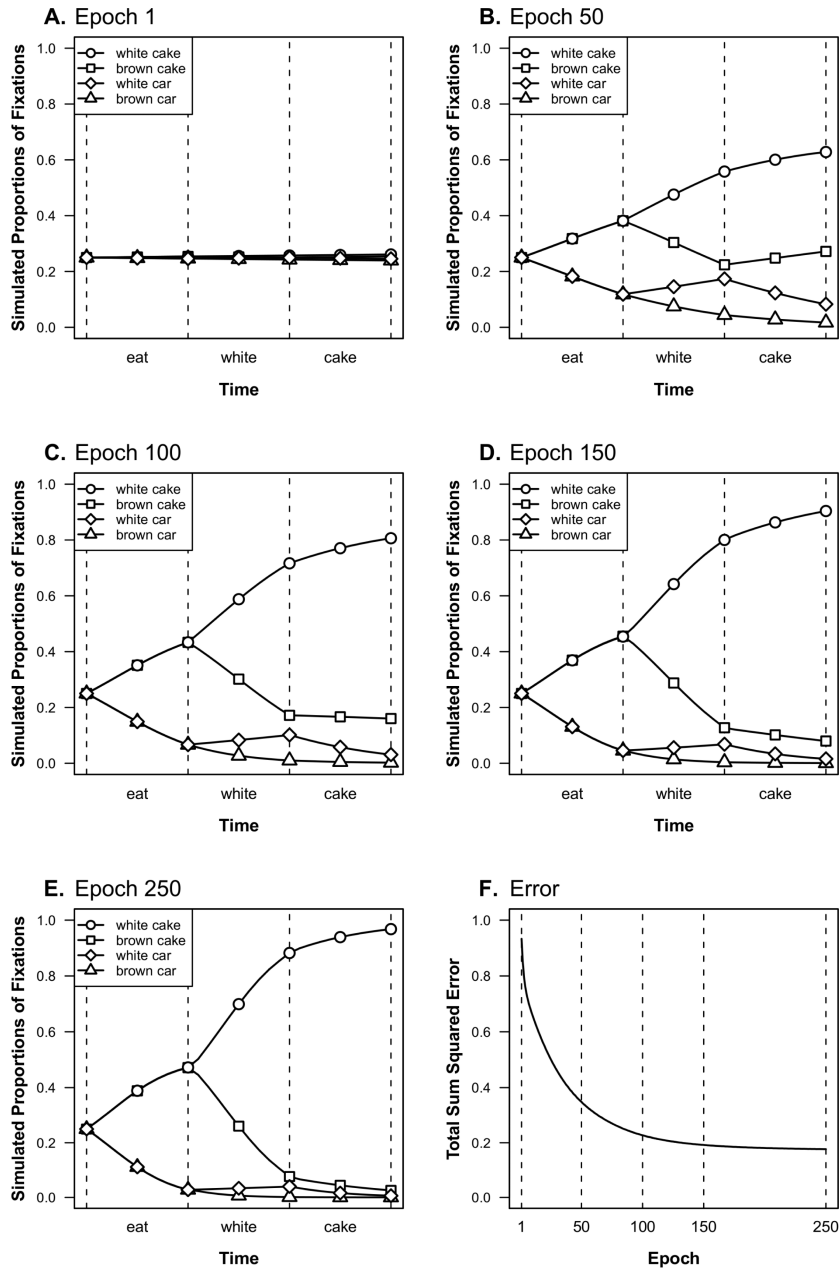


Figure 4. Simulation 2: Simulated proportions of fixations to objects in the visual layer at epoch 1 (A), 50 (B), 100 (C), 150 (D), and 250 (E), for the example sentence “Eat [the] white cake,” and the total sum squared error in the network across training (F).




Figure 5.

Example visual display from Experiment 1. Participants heard predictive sentences like “The boy will eat the white cake.” Visual displays included a target (e.g., *white cake*), which was predicted by the joint contextual constraints from the verb and adjective, and a locally coherent competitor (e.g., *white car*), which fit the lexical constraints of the adjective, but was not predicted by the verb selectional restrictions of the sentence context.

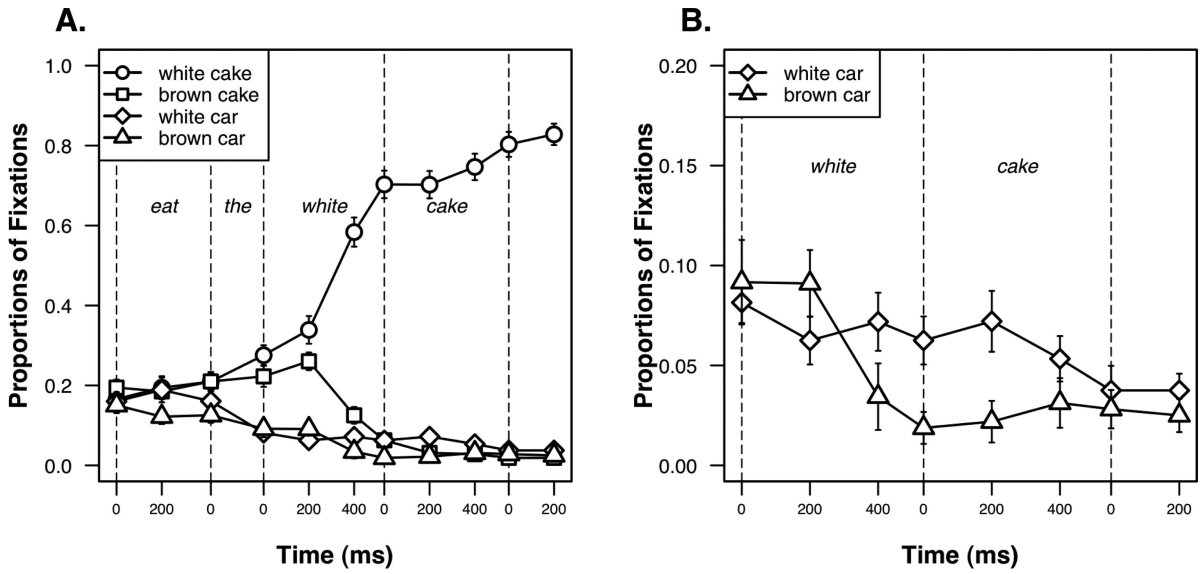


Figure 6. Experiment 1: Average (SE) proportions of fixations to the objects in the visual display between mean verb onset and mean direct object noun offset (A), and to the verb-inconsistent objects in a zoomed in window (proportions plotted to .20) between adjective onset and mean direct object noun offset (B) for the example sentence “The boy will eat the white cake.” Fixations were resynchronized at the onset of each word, and extend to the average offset of each word.


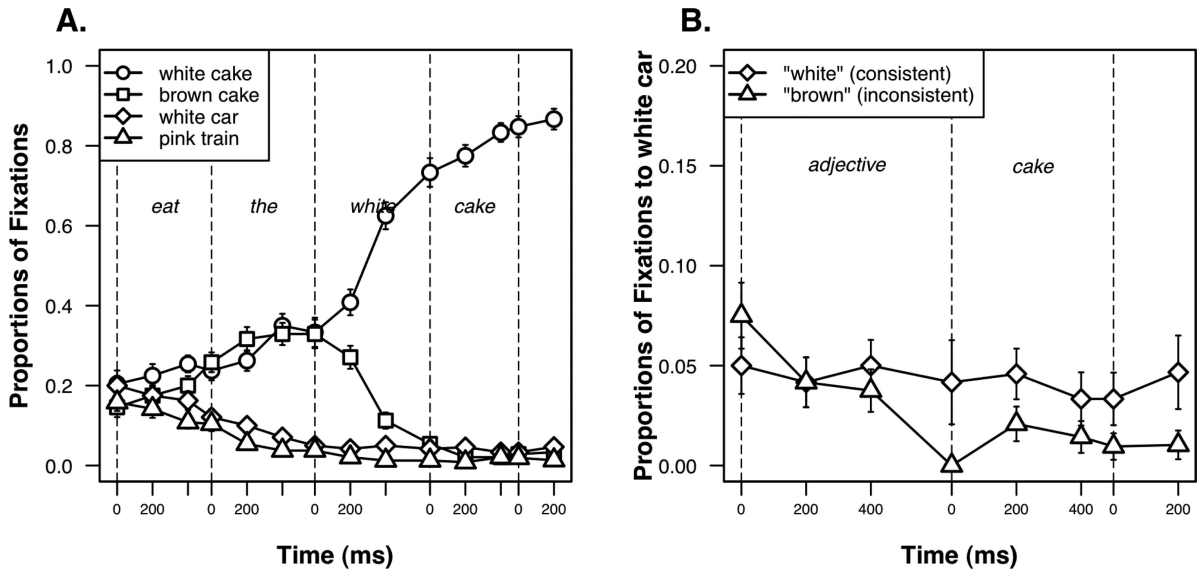


Figure 7.

Example visual display from Experiment 2. Participants heard predictive sentences like “The boy will eat the white cake.” Unlike Experiment 1, the unrelated distractor (e.g., *pink train*) was not in a competitor set with the locally coherent competitor (e.g., *white car*).

**Figure 8.**

Experiment 2: Average (SE) proportions of fixations to the objects in the visual display between mean verb onset and mean direct object noun offset with adjective-consistent sentences (A), and to the locally coherent adjective competitor object in the adjective-consistent vs. adjective-inconsistent condition in a zoomed in window (proportions plotted to .20) between adjective onset and mean direct object noun offset (B). In the adjective-consistent condition, the locally coherent competitor (e.g., *white car*) was consistent with the adjective in the sentence (e.g., “The boy will eat the white cake”), whereas in the adjective-inconsistent condition, this competitor was inconsistent with the adjective in the sentence (e.g., “The boy will eat the brown cake”). Fixations were resynchronized at the onset of each word, and extend to the average offset of each word.