

# Temporary disruption in language processing reflected as multiscale temporal discoordination in a recurrent network

Benjamin Nguyen and Michael J. Spivey\*

Department of Cognitive and Information Sciences  
University of California, Merced  
Merced, CA 95343, USA

\*Correspondence: [spivey@ucmerced.edu](mailto:spivey@ucmerced.edu) (Michael Spivey)  
DOI: <https://doi.org/10.56280/1570699699>

This is an open access article under the CC BY 4.0 license (<https://creativecommons.org/licenses/by/4.0/>).  
Received 2 December 2022; Accepted 23 February 2023; Published 28 April

## Abstract

By juxtaposing time series analyses of activity measured from a fully recurrent network undergoing disrupted processing and of activity measured from a continuous meta-cognitive report of disruption in real-time language comprehension, we present an opportunity to compare the temporal statistics of the state-space trajectories inherent to both systems. Both the recurrent network and the human language comprehension process appear to exhibit long-range temporal correlations and low entropy when processing is undisrupted and coordinated. However, when processing is disrupted and discoordinated, they both exhibit more short-range temporal correlations and higher entropy. We conclude that by measuring human language comprehension in a dense-sampling manner similar to how we analyze the networks, and analyzing the resulting data stream with nonlinear time series analysis techniques, we can obtain more insight into the temporal character of these discoordination phases than by simply marking the points in time at which they peak.

**Keywords:** Psycholinguistics, Sentence Processing, Mouse-tracking, Recurrent Networks, Recurrence Quantification Analysis, Multiscale Systems, Nonlinear Dynamics, Time Series

## 1. Introduction

In recurrent neural networks, as activation is passed back and forth among the nodes, the system will typically drift toward a stable or metastable pattern of activity over time (e.g., Amit, 1989; Anderson, 1995; Grossberg, 2012). When new inputs are introduced to the network that fit into its established regime, or are already anticipated by the direction of its flow in state space, then the network will smoothly accommodate those external inputs. However, when a disruptive perturbation is introduced to that activity flow, or one that is antithetical to the direction of its flow in state space, then the smoothly coordinated metastable pattern of activity across the network will typically become uncoordinated for some period of time. Such discoordination may result in a sizeable error signal that can be used by a learning algorithm, and it may also manifest itself as an increase in high-variance activity

across the network. Although the error signals are often measured individually at a few selected points in time, a detailed understanding of the continuous temporal dynamics of this high-variance activity will typically require nonlinear time series analysis instead.

In the case of human language comprehension, the network scenario is much like that just described. The central nervous system's recurrent networks are passing activation back and forth based on previous linguistic input (and other context) at the same time that they are being forced to accommodate new external input that continuously flows in (Spivey, 2007). The result is that readily-accommodated new input maintains coordination of the networks but difficult-to-accommodate (or unanticipated) input typically generates discoordination and an increase in high-variance activity across the cortical language networks. The field of cognitive neuroscience has developed some understanding of

these indices of discoordination, particularly with the event-related potential (ERP) measures of the N400 for semantic anomaly (Kutas & Hillyard, 1980) and the P600 for structural anomaly (Osterhout & Holcomb, 1992).

In this article, we review and examine these kinds of measures of network error, network coordination, and network discoordination, in the context of sigmoidal-activation recurrent networks, spiking recurrent networks, and real-time naturalistic spoken language comprehension. We then use nonlinear time series analysis to compare measures of temporary network discoordination to a novel measure of temporary disruption in language comprehension. We conclude that by measuring human language comprehension in a manner similar to how we analyze the networks, and analyzing the resulting data stream with nonlinear time series analysis techniques, we can obtain more insight into the temporal character of these discoordination phases than simply marking the moments in time when they peak.

At one scale of analysis, language processing might be described as a “sender” delivering a message to a “receiver” (Campbell, 1982; Shannon & Weaver, 1949). However, substantial recent research is showing that at other spatial and temporal scales of analysis, naturalistic language use behaves more like two or more complex networks (e.g., people) becoming coordinated in their word choices (Brennan & Clark, 1996), syntactic choices (Branigan, 2007), perceptual choices (Fusaroli et al., 2012), facial expressions and hand gestures (Louwerse et al., 2012), eye movements (Richardson & Dale, 2005), brain activity (Kuhlen et al., 2012) and even their postural sway (Shockley et al., 2003). Rather than pretending that naturalistic language use progresses in a discrete trade-off of regimented turn-taking between uninterrupted message-delivery events, treating it as a pair of recurrent networks that have become connected by a continuously shared information flow (Dale & Kello, 2018; Falandays et al., 2020; Spivey & Richardson, 2009) may provide a new perspective that improves our understanding of the predictable and unpredictable multiscale temporal dynamics that emerge. Such accounts of cognition that rely on recurrent causal feedback loops of information flow spanning across multiple spatial scales and temporal scales support a

“contextual emergence” (Atmanspacher & Beim Graben, 2009; Bishop et al., 2022) of coordinated cognitive behavior that might look as though it was guided by centralized rules and symbols even though it was not.

## 2. Measures of network error

A key component of language understanding is the ability to detect errors in the language input we receive or in our individual parsing or processing. The importance of learning by extrapolating from and mimicking examples has long been of interest to the linguistic development research community (Bloom et al., 1974). Children and other language learners often need feedback in order to correct their mistakes (Carroll et al., 1992; Hirsch-Pasek et al., 1984). However, with time and enough exposure, linguistic expectations begin to consolidate and get applied in use, such as words should have mutually exclusive meanings or that plurality can be expressed by adding ‘s’ to the end of a word (Markman & Wachtel, 1998; Ramscar et al., 2013). Violations of these expectations become important events that brains are sensitive to and, as such, are reflected in brain activity.

For quite some time, neuroimaging techniques have been providing insight into how brains respond to anomalous sequences of linguistic stimuli. Brain waves are electrical signals with rhythmic and oscillatory properties that can be captured by methods such as electroencephalography (EEG). Substantial evidence has repeatedly validated EEG activity in certain cortical regions as reflective of activity responding to specific types of stimuli and thereby related to specific cognitive functions. Abrupt changes in the brain's electrical activity in response to semantic and syntactic anomalies, namely the N400 and P600 event-related potentials (ERPs), have been canonized according to their valence and time course.

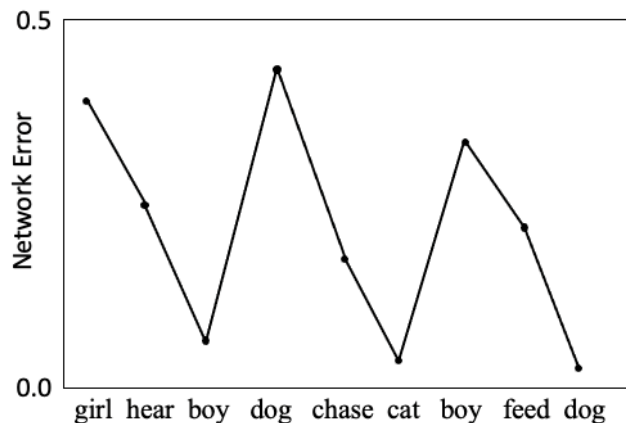
The N400 is a negative event-related potential (ERP) or hyperpolarization that starts at around 200 ms and peaks at 400 ms in response to the violation of semantic expectations and is generally spatially localized to the frontal, central and parietal channels (Kutas & Hillyard, 1980). Even among semantically valid sentences, statistical differences in the predictability of a given word

can result in subtle differences in the magnitude of the N400 elicited (Kutas & Federmeier, 2011). Moreover, the N400 elicited by an anomalous sentence can be reduced or eliminated by the presence of a preceding context that makes the sentence more understandable (Nieuwland & Van Berkum, 2006) or by a “noisy-channel” that allows a listener to mentally-revise the anomalous word (Ryskin et al., 2021). Importantly, these semantic expectations, and their violations as evidenced by the N400, are not limited to language but extend to vision, action, and nonlinguistic sounds, to name a few (Kutas & Federmeier, 2011).

Often discussed in tandem with the N400 is the P600, a positive ERP or depolarization, that peaks at 600ms to stimuli that violate syntactic or structural expectations. As with the N400, the P600’s exact function is a topic of debate, the primary tension being whether these ERPs reflect a global sentence or discourse-level reanalysis or whether they reflect difficulty with integration from expectations accrued by prior context (Kaan et al., 2000). While we do not presently weigh-in on this debate, we are inspired by both sides’ recognition of language processing as dynamic. In language research, the P600 has been specifically dissociated from the N400 using stimuli that simultaneously violated both types of expectations (e.g. “the meal was devouring.”). For these stimuli, researchers found participants most often elected a thematic parse where the error was syntactic (incorrect ending to the verb), rather than semantic (incorrect subject-object relations), which was accompanied by a P600 response and lacking an N400 response (Kim &

Osterhout, 2005). Like the N400, the P600 also extends beyond language, as it is observed during structural violations in music as well (Patel, 2003).

Simulations of language use have long been employed in the study of language processing. Often, these simulations abstract the units of language into their functional roles in the meaning-making of the linguistic unit, whether that be letters in words, parts-of-speech in sentences, or successive sentences within narratives. Despite this abstraction, these simulations have been successful and useful for deriving and theorizing about language processing. Moreover, in language simulations, the utility of error as an indicator about the semantic or syntactic parse is also well established. With a simple recurrent network architecture, using sigmoidal-activation units, the Elman (1990, 1991) is able to learn basic sentence structures by predicting upcoming words and deriving an error signal by comparing its prediction to the actual next word in the input. After training with backpropagation, the network is able to use a few words of previous context to predict what will follow. Importantly, at the joints between completed sentences, however, the network’s error shows a steep increase (Figure 1). This jump in error results from a diffuse prediction of the many possible words that may start a new sentence compared to the specific individual word that actually starts that particular next sentence. This error spike can be framed as an implicit acknowledgement of uncertainty in processing the current input, not unlike the P600 and the N400.



**Figure 1.** Schematic example of a network’s mean squared error in a simple recurrent network being exposed to sequences of words that form short sentences. As a sentence progresses, its upcoming words get easier to predict. However, when the end of a sentence is reached, and the word that begins the next sentence is not easily predicted, the Elman net produces a spike in its error measure.

Simulations have also been made of EEG activity, some of which have focused on biological plausibility, while others have prioritized alignment with real human behavior and responses. Rabovsky & McRae's (2014) simulation had a functional focus, starting with the premise that the N400 response is consequent to difficulty processing semantic information. Informed by this premise, they constructed a network with a 30-node input layer fully connected to 2,526 semantic nodes or features. These features were determined *a priori* in a separate study where people were asked to generate tokens of semantic information for a given word (McRae et al., 1997). Rabovsky & McRae's (2014) was trained on 541 concepts or words, encoded as unique distributed patterns of 1s and 0s to be fed to the input layer. The network used backpropagation to adjust its weights and was evaluated on the whole network error, as measured by cross-entropy error, and on the whole network activation. The network was tested on several different phenomena to which the N400 has been found sensitive, namely, priming effects, semantic richness, word frequency, repeated exposure, and the number of orthographic neighbors (words that share all letters except for a single letter, preserving letter position). The network was successful in replicating the N400 patterns in these various phenomena by exhibiting increases in cross-entropy error (see also, Rabovsky et al., 2018).

Successful simulation of error detection has also been conducted for the P600 response. With a reservoir network, Hinaut & Dominey (2013) were able to simulate the P600 response to syntactically anomalous stimuli. In most cases, a reservoir network contains a fixed reservoir of nodes with recurrent connections. Inputs are connected to this reservoir, and thereby low-dimensional inputs become projected into a higher-dimensional space. This projection enables the network to be sensitive to both spatial and temporal patterns in the input stream. The reservoir is often connected to a "readout" or output layer. Unlike the connections in the input layer and reservoir itself, these connections are adjusted and serve as the site of learning. Hinaut and Dominey's network, in particular, contained 300 nodes with 10% connectivity and conceptually represented processing in the prefrontal cortex outputting to the striatum, which is purported to support the speed of syntactic error detection. The network was fed an input stream of sentences and was trained to produce the part

of speech for nouns, whether the noun was the subject or the object of the verb. To adjust reservoir-output connections, regression techniques were used. In general, the network was successful in identifying word roles and was generalizable to untrained constructions. For subject- and object-relative sentences, the network would consider multiple competing role identifications until disambiguating words were inputted. At that point, because of the lower frequency of relative sentences in the corpus, the instantaneous change in output activity spiked, akin to a P600 response.

Expanding on reservoir-type networks and responses to unexpected stimulus inputs, Falandays et al. (2021) trained a spiking reservoir-like network (100 reservoir nodes, 10% connectivity with the input layer) on a simple set of word sequences. Like a reservoir network, the network had a fixed number of nodes. However, unlike canonical reservoir networks, nodes in this network updated their connections after each timestep to maintain a homeostatic critical level of activity, producing approximately one outgoing spike for every incoming spike. With no output layer, instead, the activity of the reservoir was interpreted directly. The network was fed a series of sentences created according to a probability matrix. This matrix included the nouns "dog" and "man" and the verbs "bites" and "walks." All eight grammatically possible noun-verb-noun combinations were used as input sentences (separated by an end-of-sentence input marker), but differential probabilities were assigned according to the more typical semantic relationships. The network was able to differentiate between the words (and subject/object roles), creating regular, consistent, and unique distributed representations in the reservoir layer for each of them. Importantly, when the input layer was fed the word "man" but then nothing else for two-time steps, the reservoir continued to generate activity in the absence of input – as though it was anticipating the likely next words. The following time step involved the reservoir producing a pattern of activity that was extremely close to the distributed representation for the verb "walks." On the time step after that, the reservoir produced a pattern of activity that was moderately consistent with the distributed representation for the input "dog" (as a direct object). Anomaly-detection phenomena were also observed with this network, similar to a P600. When a

subject noun was followed immediately by an object noun instead of the regular verb, mean activity in the reservoir increased substantially higher than the typical activity level – not unlike the sharp increase in error seen with the Elman net when its predictions are violated (**Figure 1**).

This simulated network demonstrates several key points. Firstly, the homeostatic “learning” algorithm was highly local, with each node adjusting its connections simply to maintain a homeostatic balance in its input/output ratio. However, population-level collective representations emerged for not only the word itself but for the word’s role in the sentence (subject/object). Secondly, this work questions how explicit the recognition of “error” needs to be in a learning algorithm. In this network, there is no teaching signal of any kind; rather, the learning is localized to each node and based on a homeostatic process of balancing local connectivity. Finally, the anticipatory persistence of the network activity in the absence of input calls into question the need for any explicit treatment of prediction as though it were an internal representation of expected upcoming input ([Clark, 2015](#); [Hohwy, 2013](#)). In this network, the same connectivity pattern, and its concomitant flow of activity, carry out the system’s processing of incoming input and its anticipation of future input – without the need for a separate forward model that generates predicted representations. Evidently, behavior that mimics predictive processing can emerge in a network that is doing little more than becoming dynamically entrained with the sequences in its environment ([Dubois, 2003](#); [Falandays et al., 2021](#); [Falandays et al., 2023](#); [Marmelat & Delignières, 2012](#); [Stepp & Turvey, 2010](#); see also [Di Paolo et al., 2022](#); [Raja, et al., 2021](#)).

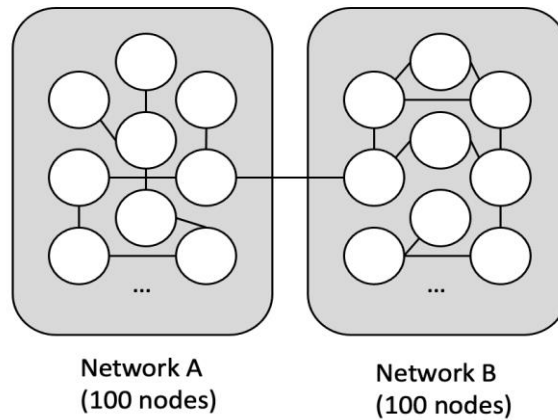
### 3. Measures of network coordination

Dynamic entrainment between the rhythms of a brain and the rhythms of the environment can be similar to the dynamic entrainment between the rhythms of two neighboring brain regions. In the cognitive neuroscience of language processing, increasing evidence has mounted that synchrony across brain regions is fundamental to what is typically described as context effects ([Hauk et al., 2017](#)). Perhaps the method by which

ambiguity and uncertainty in one subsystem are resolved by contextual information from another subsystem involves those two subsystems becoming synchronized to some degree. Some of this synchrony and interaction can be captured by detangling brain waves at the frequency level. For example, [Lewis et al. \(2017\)](#) recorded EEG data from participants while they read sets of four sentences that were either unrelated to each other or formed a coherent narrative. They found that oscillatory activity in the beta frequency range (13-21 Hz) was significantly stronger with the coherent narratives. For some time now, it has been understood that since neural networks behave as complex dynamic systems, synchrony in brain waves will be a particularly fruitful indicator of interdependencies between subsystems in the network ([Lopes da Silva, 1991](#)).

Emergent synchronization of oscillatory activity has frequently been observed in neural dynamics ([Chialvo, 2010](#)), particularly when the network is flexibly balanced on a critical point of metastability ([Kello et al., 2023](#)), somewhere in between stagnant stability and rampant chaos. For spiking neural network simulations, this critical point is achieved when each node emits approximately one outgoing spike for every incoming spike ([Kello, 2013](#)). However, with sigmoidal-activation networks, the critical point that fosters emergent synchrony may look slightly different. [Falandays et al. \(2020\)](#) found that a fully-connected 100-node network with a logistic activation function, random weights centered around a mean of zero, and a random starting activation was mostly likely to drift toward a global oscillatory activity regime (instead of flatlining), particularly when the standard deviation of its random weights was close to 1.2. When two such networks were minimally connected (**Figure 2**), [Falandays et al. \(2020\)](#) found that they would become entrained with one another, coaxing each other to slightly change their rhythms to approach one shared rhythm.


This can be seen with “single-cell recordings” from the simulated networks. In new analyses of the behavior of these networks, we report here new statistical evidence for temporal coordination between the two networks after a single bi-directional connection is introduced .



**Figure 2.** Fully-connected random networks are given a single bi-directional connection (between Network A's node 18 and Network B's node 6) that instigates coordination between the two networks.

across the two of them, between two arbitrary nodes: Network A's node 18 and Network B's node 6. By recording the activity of Network A's node 52 and Network B's node 58 (which are not directly connected to each other), we were able to track their uncoordinated behavior before the two networks were connected and then compare it to their coordinated behavior after the two networks were connected. For best comparison, the time series for Network A's node 52 and Network B's node 58 were z-scored and then overlaid in **Figure 3A**. Network A's node 52 (in black) completes a cycle every 5 time steps whereas Network B's node 58 (in red) completes a cycle every 12 time steps. The two-time series are uncorrelated no matter what lag one applies in a cross-correlation.

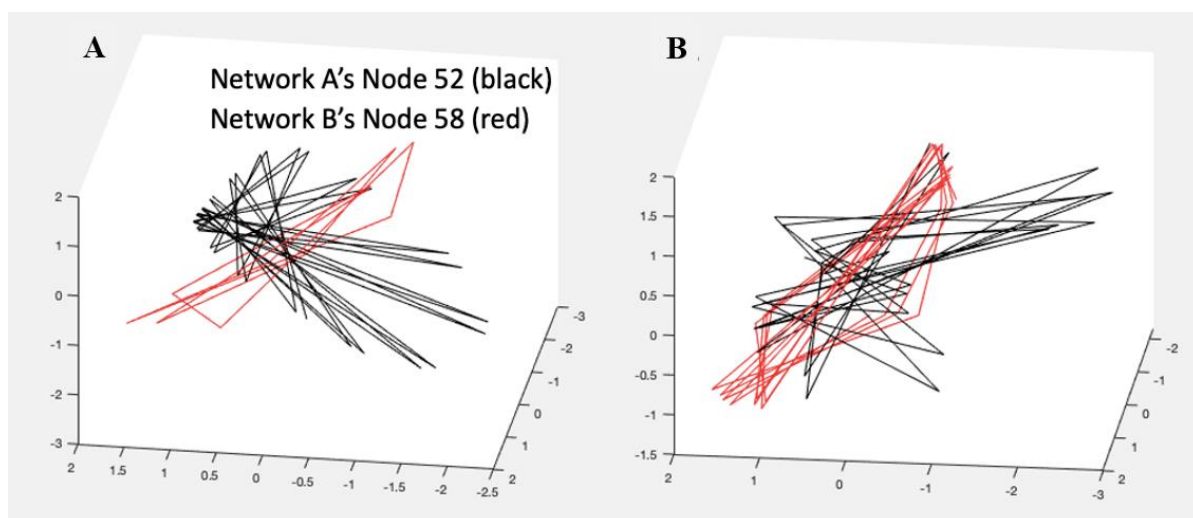
By contrast, when the two nodes are compared *after* the bi-directional connection has been introduced, one can see that the two time series have developed some similarity in their patterns (despite the fact that these two nodes are not directly connected to each other). See **Figure 3B**. Network A's node 52 (in black) now tends to have a two-peak cycle that repeats every 11 time steps and Network B's node 58 (in red) also has a cycle that repeats every 11 time steps. With a lag of 2 time steps shifting the time series for Network A's node 52 rightward, the cross-correlation between these two nodes is quite strong;  $r = .45$ . Most randomly selected pairs of nodes across the two networks exhibit a temporal coordination of some kind after the bi-directional connection has been introduced (Falandays et al, 2020).



**Figure 3.** Before and after the networks are connected: Z-scored time series for Network A's Node 52 (in black) and Network B's Node 58 (in red).

Since Network A's 100 nodes are fully interconnected, and thus their activity patterns are interdependent, extracting a time series from just one of those nodes and embedding it in multiple dimensions can provide a glimpse into the network's overall activity (Takens, 1981). This state-space reconstruction technique simply requires taking regularly-spaced samples from the single time series (e.g., time slices 1, 4, 7, 10, etc.) and treating them as the x-dimension. Then, one takes the remaining regularly-spaced samples (e.g., time slices 2, 5, 8, 11, etc. along with time slices 3, 6, 9, 12 etc.) and treats them as the y- and z-dimensions. The reconstruction can take place in a state-space with any number of dimensions, but for data visualization, three dimensions is useful. In this state-space reconstruction, the first time slice of an embedded 3-dimensional trajectory carries with it information about the original time series' 1<sup>st</sup>, 2<sup>nd</sup> and 3<sup>rd</sup> activity levels and the second time slice of this

embedded trajectory carries information about the original data's 4<sup>th</sup>, 5<sup>th</sup>, and 6<sup>th</sup> activity levels, etc. As long as the activity of the other 99 nodes is interdependent with the activity of the measured node, then embedding that one node's time series into multiple dimensions can produce a multi-dimensional trajectory that bears substantial resemblance to the overall network's own multi-dimensional trajectory (Spivey, 2018; Stephen et al., 2009). **Figure 4A** shows the absence of shared locations visited by Networks A and B before the connection is introduced. The vertices in the red trajectory (Network B's Node 58) do not overlap with any of the vertices in the black trajectory (Network A's Node 52), thus revealing the lack of coordination between the networks. By contrast, **Figure 5B** shows substantial visiting of shared locations after the cross-network connection is added, thus revealing significant coordination between the two networks.



**Figure 4.** Before connected (A) and after connected (B): State-space reconstruction of Z-scored time series for Network A's Node 52 and Network B's Node 58 embedded in 3 dimensions. Very little recurrence (i.e., overlap in locations visited by red and by black) is evident before the networks are connected (A). However, substantial recurrence is evident after the networks are connected (B).

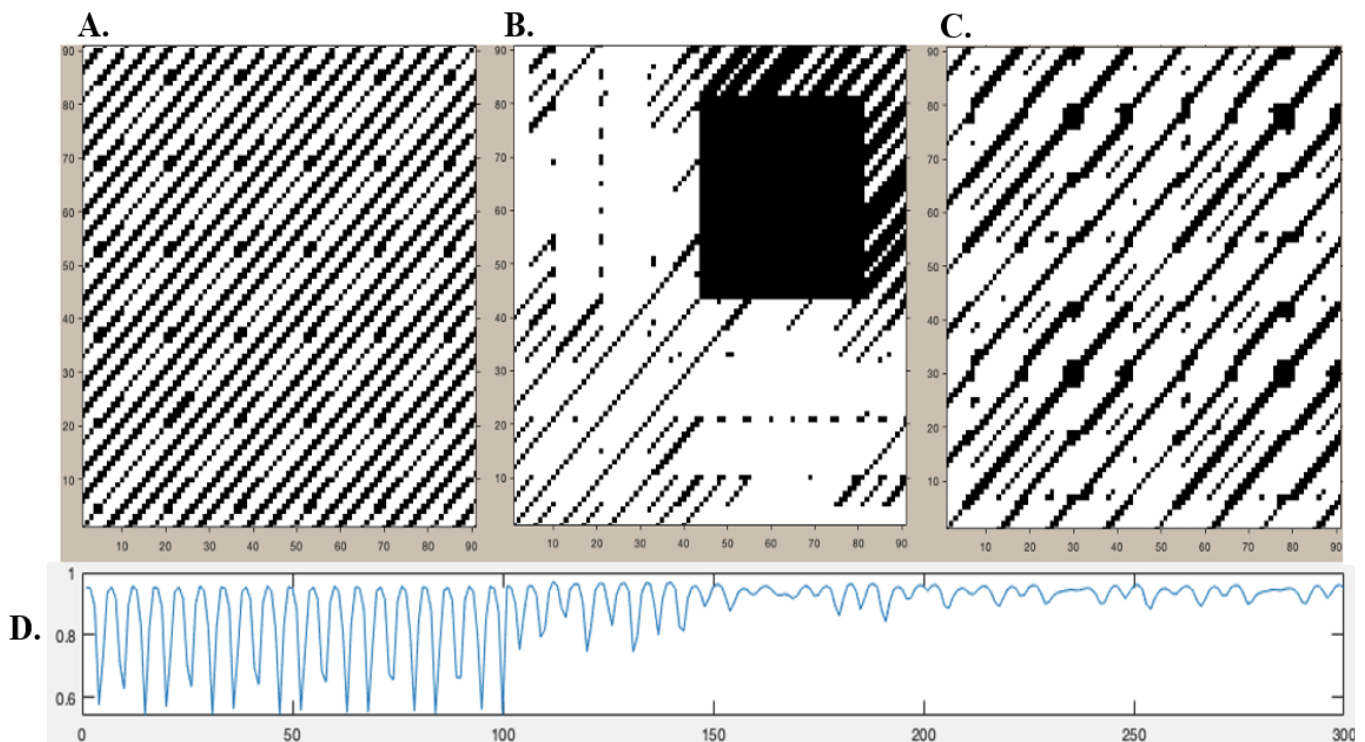
Results show that a single bi-directional connection between the networks made it so that almost any randomly selected node from one of the original independent networks became correlated (or anti-correlated) with any randomly selected node from the other network (Falandays et al., 2020). The individual and collective networks were self-propagating, and over time, stabilized to a regular, periodic pattern of activity. When the trajectories through state space of these networks are examined, the connected network's trajec-

tory becomes a meld of the two individual network paths. Given the immense level of connectivity in the brain, we expect that multiscale interactions like this are pervasive. Unpacking these interdependencies may lend unique insights beyond examination at any one particular scale. To that end, recurrence quantification analysis (RQA) is a statistical method – not relying on assumptions of normality or stationarity – which can be leveraged for that express purpose (Marwan et al., 2007).

In general, RQA describes a system's trajectory through state space by quantifying the degree of return or revisitation to specific portions of this space. When performed on a single system, the results can reveal the systematicity with which it carries out coherent steady cyclic loops in its state space, exhibiting recurrences of sequenced values at long- and short-time scales, stabilization at certain values, as well as degrees of determinism and entropy. RQA can corroborate whether the underlying generative system is stabilized in a simple cyclic pattern, born from multiscale interaction between numerous nested subprocesses, or exhibiting chaotic unpredictable behavior.

First, we use [Marwan et al.'s \(2007\)](#) RQA software to examine how the internal dynamics of Network A were changed by the introduction of this one bi-directional connection with Network B (**Figure 2**). As we will see, RQA produces a wide variety of measures, but since we have a record of all the data inside these 100-node network simulations and we know when the external

perturbation was introduced, we can examine these different measures to see which ones provide an index of these system transitions that we know are taking place in the networks. When we examine Network A's node 52 by itself, in a (auto)recurrence analysis, we line up the node-activation time series with itself on both x- and y-axes and look for recurrences in the state-space, which are plotted as black dots on the recurrence plot (**Figure 5**). This is why in a (auto)recurrence analysis, the main diagonal (or Line of Identity, where each time slice is lined up with itself) is always a solid line of dots. Since we know that the actual state space of each network's activity has 100 dimensions, for these recurrence plots, we embedded the time series in more than 3 dimensions. For ease of computability, we settled on a 10-dimensional embedding and a radius criterion of 2. Figure 5 shows the (auto)recurrence plots for Network A's node 52 *before* the cross-connection with Network B was introduced (panel A), *immediately after* it was introduced (panel B), and *100-timesteps after* it was introduced (panel C).



**Figure 5.** Recurrence plots of Network A's node 52 before (A), immediately after (B), and long after (C) the connection from Network B was introduced. Panel D plots node 52's raw time series throughout this disruption and recovery.



Among these RQA measures, Overall Recurrence (proportion of the plot filled with black dots) is always analyzed in recurrence plots. In this case, Recurrence was .34 before the connection was introduced, .30 immediately after it was introduced, and .23 100-timesteps later. Entropy in the recurrence plot measures the unpredictability of the deterministic structure in the system. Not surprisingly, Entropy starts low (2.49) before the connection is introduced (**Figure 5A**), increases to 3.41 right after the connection is added (**Figure 5B**), and then returns to low levels again (2.34) after the network has had some time to accommodate to the new dynamics (**Figure 5C**). Trapping Time measures the average length of the vertical lines in the plot and serves as an index of how long the system remains in any specific state. This measure, too, seems to document the transition that Network A goes through. Before the connection is introduced, Trapping Time is 2.03. Immediately after the connection is introduced, it increases to 7.42. Then, after 100 timesteps have elapsed, it returns to its original range: 2.58. These measures serve as indices for the pre-, during - and post-disruption phases of behavior we imposed on this simulated network.

Thus, even when the only thing that is measured is a single one-dimensional times series from one element of the complex system that is undergoing a disruptive perturbation, a great deal of information can be extracted from that signal by embedding it in a state space, building recurrence plots, and applying RQA to those plots. As we will see in Section 4, with real-world systems whose connectivity we do not know in advance, these measures can also provide insight into their degree of coordination or discoordination.

Analyzing these network simulations with state-space reconstruction and recurrence quantification analysis provides us with tools to track the coordination (and lack thereof) that may emerge in systems whose subcomponents may be interdependent and whose computations may involve multiscale nested processes. When a system is functioning smoothly and traversing layered rhythmic loops in its state space, the coordination of its subcomponents will be evident in the state-space reconstruction, in the recurrence plots, and in the recurrence quantification

analysis measures. Moreover, when a system is *not* functioning smoothly, the discoordination should also be visible with these methods.

#### 4. Measures of linguistic discoordination

So far, we have seen how state-space reconstruction and RQA can provide quantitative insight into the global character of a simulated interconnected high-dimensional system, even when looking at just a one-dimensional time series emitted by that system. We can now transfer that understanding to a human language process that similarly relies on an interconnected central nervous system from which we might extract a one-dimensional time series. The one-dimensional time series we focus on here is one generated by human participants listening to a naturalistic speech that contains temporary syntactic ambiguities. As opposed to the full-blown syntactic anomalies (e.g., ungrammatical sentences) that elicit P600 effects in EEG signals (Kim & Osterhout, 2005), these temporary syntactic ambiguities are typically resolved as non-anomalous within the next few words (Bever, 1970). As a result, the discoordination in comprehension is often brief and mild, though still detectable in eye-movement patterns (Spivey & Tanenhaus, 1998), reaction times (Spivey-Knowlton et al., 1993), and graded grammaticality judgments (Bard et al., 1996).

In the present study, participants listened to spoken narratives and used a computer mouse to continuously stream a running tally of their ease of comprehension. In contrast to previous methods, such as measuring saccadic eye movements or reaction times, this continuous meta-cognitive measure provides a time series of output throughout the course of the spoken narrative that can be analyzed with nonlinear time series analysis techniques such as RQA. Previous research has used dense-sampling measurements (e.g., mouse-tracking, eye-tracking) to examine theories of sentence processing in target selection and reading tasks (Farmer et al., 2007; Spivey & Tanenhaus, 1998). These approaches often convert the complex movement data into a short sequence of discrete eye fixation durations or reduce the mouse movement task to a two-alternative forced choice (but cf. Dotov et al.,

2010; Meyer et al., 2023). In such experiments, the data of interest are the mouse and eye trajectories. However, the quality and character of these movements (e.g., pauses, fixations, angular velocity, etc.) are rarely examined with time series analysis techniques.

By contrast, we introduce an adaptation of the mouse-tracking paradigm to explore a time series measure of comprehension felicity in the form of continuous self-reporting of ease of comprehension. By providing a direct and informationally dense data stream, this novel method may reveal new insights into real-time meta-cognitive sentence processing because it allows one to apply nonlinear time series analyses to the data. In this experiment, we recorded continuous ratings of ease of comprehension while participants listened to reduced relative clause “garden path” sentences that have temporary syntactic ambiguities and also to unambiguous paired control sentences.

These sentences are referred to as “garden path” sentences because they can give the reader/listener the impression that they have been led down the garden path to a dead end if their parsing of the syntax “took a wrong turn” and built the wrong syntactic structure early on. Take, for example, Bever's (1970) famous sentence, “*The horse raced past the barn fell.*” At first read, it may seem like an ungrammatical sentence, but this is simply due to the fact that the reader probably parsed “The horse” as the Subject of the verb “raced,” when in fact, it should be the Object of that verb. The phrase “raced past the barn” is not the main verb phrase of the sentence but instead is a relative clause describing some details about what was previously done to the horse by an unspecified agent. The main verb of the sentence, which describes what the horse is doing, is “fell.” An identical syntactic structure is seen in the following sentence that does not lead the reader down the garden path, “The landmine buried in the sand exploded.” The semantic content of these words, despite identical syntax, is what makes the difference in their processing (Spivey & Tanenhaus, 1998). The semantics of “horse” and “race” naturally lead the reader to assume the horse is the Subject of the racing event, which leads the reader into building a syntactic structure with no place for the final word

to fit in. However, the semantics of “landmine” and “buried” tend to lead the reader into correctly treating the landmine as the Object of the burying event, which involves building a syntactic structure with room for the final word to be included. Between the extremes of “garden path” and “no garden path” effects are sentences that may induce only partial confusion or intermediate discoordination in comprehension. For example, “The disc spun in the drive broke.” and “The child hurried through the doorway tripped.”

Previous work with meta-cognitive sentence acceptability (or grammaticality) judgments has suggested that gradations in the acceptability of a sentence can indeed be captured with careful methods. Bard et al. (1996) adapted the psychophysical method of magnitude estimation to demonstrate that a consistent hierarchy of ratings can be found among participants for sentences that are not fully acceptable and also not fully unacceptable. Results like this have been interpreted as evidence that theories of syntax require some mechanism for gradience in acceptability to be included (Bard et al., 1996), perhaps in a fashion similar to how the field of phonology has allowed for gradations in acceptability via optimality theory (Prince & Smolensky, 2004).

By extracting a continuous stream of graded acceptability measurements – instead of obtaining it only at the end of a sentence – our novel method may provide additional insights into the temporal dynamics of this gradience in syntactic processing. Moreover, since it is a densely-sampled data stream that uses nonlinear time series analyses, we can reveal hidden structures in the multi-dimensional temporal dynamics of these meta-cognitive garden path reports. The same indices of coordination and discoordination observed with our network simulations above can now be measured from human participants’ continuously streaming mouse-movement data.

## 4.1 Methods

### 4.1.1 Participants

Forty-six participants were recruited from the undergraduate population at the University of California, Merced. The mean age was 20 years old,

and all participants reported normal hearing ability. Each participant contributed 5 ratings while listening to garden path sentences and 5 ratings while listening to control sentences. One participant contributed slightly less than 5 reports for each condition due to voluntary noncompliance.

#### 4.1.2 Materials

Sentence stimuli were generated using the open-source text-to-speech software Balboa v.2.15.0.803. Embedded in a natural continuous narrative, we created 10 garden path sentences, where the unambiguous control version contained the phrase “who was” between the noun and the verb, whereas the ambiguous version had “who was” edited out with Audacity v.2.4. Eight of these garden path sentences were stereotypical reduced relative clauses, similar to those used in previous sentence processing experiments (e.g., [McRae et al., 1998](#)). The other two garden path sentences were less stereotypical natural sentence forms that still introduced relative clause ambiguities that can result in garden path effects (indicated by asterisks in sentence numbers 1 and 7 below). These reduced-relative clauses introduce temporary syntactic ambiguity as to whether the noun phrase is the Subject or Object of the verb. Then, when the *subsequent* verb is heard, the syntactic ambiguity is resolved, forcing the initial noun phrase to be understood as the Object of the initial verb. The narrative was ~17 minutes or 173 sentences in total length and constant across participants, except that half heard a version of the narrative where 5 relative clause sentences were in their syntactically ambiguous reduced form and the other 5 in their unambiguous control form.

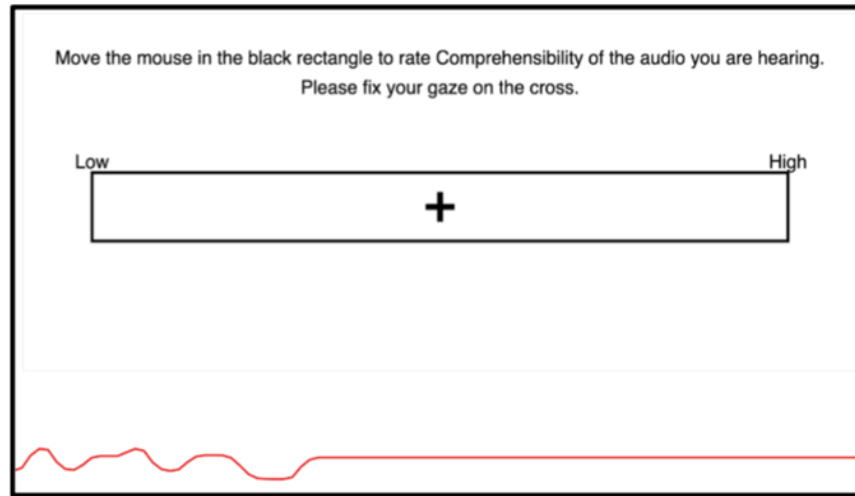
In contrast, the other half of the participants heard the inverse version of that arrangement. Using these two slightly different stimulus lists allowed us to obtain data for both versions of each garden path sentence without having any one participant hear both versions of it. The semantic content of the rest of the narrative

did not contextually resolve the ambiguities introduced inside the critical garden path sentences.

1. Ana recalls one child (who was) bought an ice cream cone smiling at the driver. \*
2. The chef (who was) gifted the knife greeted the audience.
3. The babysitter (who was) purchased a gift card, thanked the parents and went home.
4. The manager (who was) prepared the report praised the employees.
5. The doctor (who was) delivered the donor blood assisted several surgeons.
6. The child (who was) read a story hugged the nanny.
7. Ana nodded to Mr. Koga, and the man (who was) prayed for last Sunday. \*
8. One child (who was) saved a front-row seat surprised the clown with a mud pie.
9. The dispatcher (who was) handed a pink slip, worried about the driver.
10. The flight attendant (who was) called for arrived to address his concerns.

#### 4.1.3 Procedure

Participants sat at computers in a private room and personally configured over-ear-headphones to comfortable hearing levels. We asked participants to move their mouse cursor freely within a rating box (**Figure 6**) as they listened to our stimulus sentences. To increase compliance, participants were allowed to take self-paced breaks at specified intervals. The box’s sides were labeled as ‘low’ and ‘high’ levels of comprehensibility. Participants’ mouse position in the box indicated their self-assessed current level of comprehension. Each participant (N=46) continuously listened to and rated the entire narrative. However, we are interested in the ratings of the garden path sentences and their unambiguous control versions. Each participant was randomly assigned to listen to one version (garden path or control) of these 10 sentences of interest.



**Figure 6.** Rating box within which participants moved their mouse cursor to rate their ease of comprehension. A redline streamgraph at the bottom depicts their ratings updated in real-time. This graph served to aid in understanding the task and to further yoke changes in mouse position to changes in concurrent sentence processing.

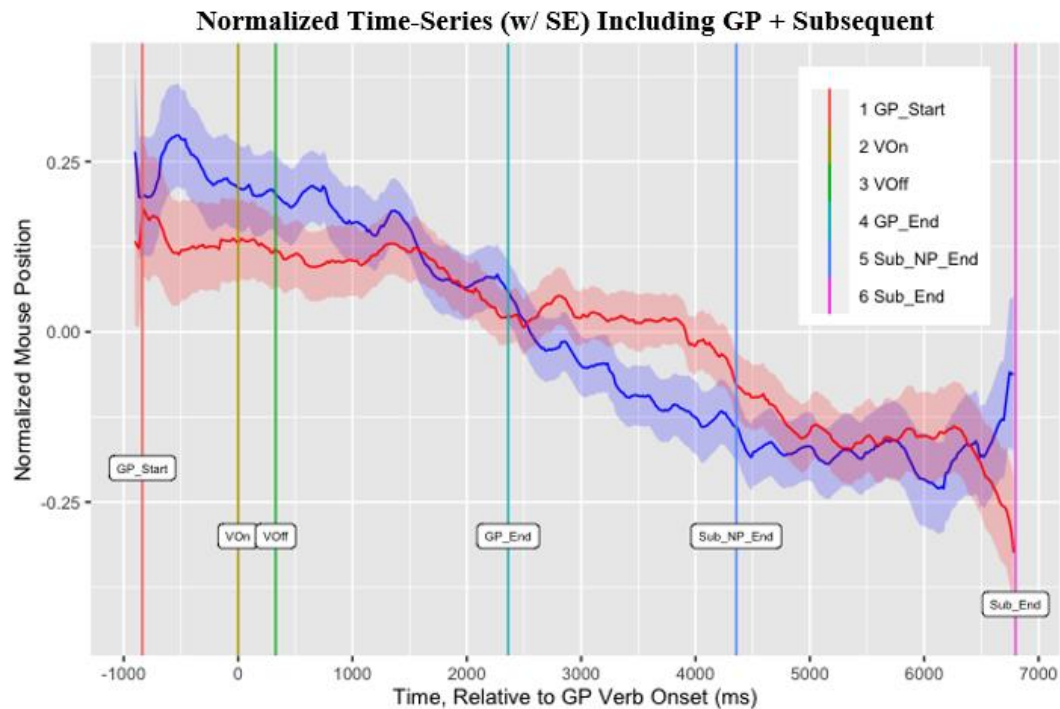
## 5. Results

We compared the averaged normalized ratings of ease-of-comprehension between the garden path and control conditions of each critical sentence (**Figure 7**). At the start of the sentence, there is a sizeable but non-statistically-significant difference between garden path (blue) and control (red) versions of the sentence, where the garden path version appears to show greater ease of comprehension than the control version. This could be due to the increased complexity of processing the “[noun] who was [verbed]” portion of the control sentence, which forces the reader to non-canonically treat the head noun of the sentence as the object of that verb. The fact that the garden path version shows *greater* ease-of-comprehension during the start of the sentence may indicate that the listener is blithely treating the head noun of the sentence as the canonical subject of the verb, which is exactly what later leads to the garden path effect in these sentences. Over the course of the garden path sentence and shortly after it (blue line in **Figure 7**), a steep decline in ease-of-comprehension is apparent – much steeper than what is observed with the control version of the sentence (red line in **Figure 7**). Indeed, later in the time series, the difference in ease-of-comprehension between control and garden path versions reverses dramatically during the two-second period beginning around the end of the garden path sentence (GP End in **Figure 7**) and persisting until the

end of the noun-phrase in the subsequent sentence (Sub\_NP End in **Figure 7**). During that later temporal window, the garden path version of the sentence elicits substantially lower ease-of-comprehension values than the control version. Much of that period of time is the silence between sentences, suggesting that the decrease in reported ease-of-comprehension is a delayed response to the end of the garden path sentence – slightly similar to the delay seen with self-paced word-by-word button-press reading experiments compared to eye-tracking reading experiments (McRae et al., 1998).

We averaged across this later temporal window for each participant (middle of **Figure 7**), and using a one-sided paired t-test compared the mean of the garden path condition ( $M = -0.071$ ) to the mean of the control condition ( $M = 0.011$ );  $t(45) = 1.60$ ,  $p = 0.058$ . This result affirms that this method can adequately capture participants’ fleeting awareness of encountering a brief garden path effect.

Similar to the network error measures discussed above (recall **Figure 1**), the difference between ease-of-comprehension for the garden path condition and for the control condition can pose as a kind of measure of error, or network discoordination, in the language comprehension process. **Figure 8** plots the garden path condition’s ease-of-comprehension data stream minus



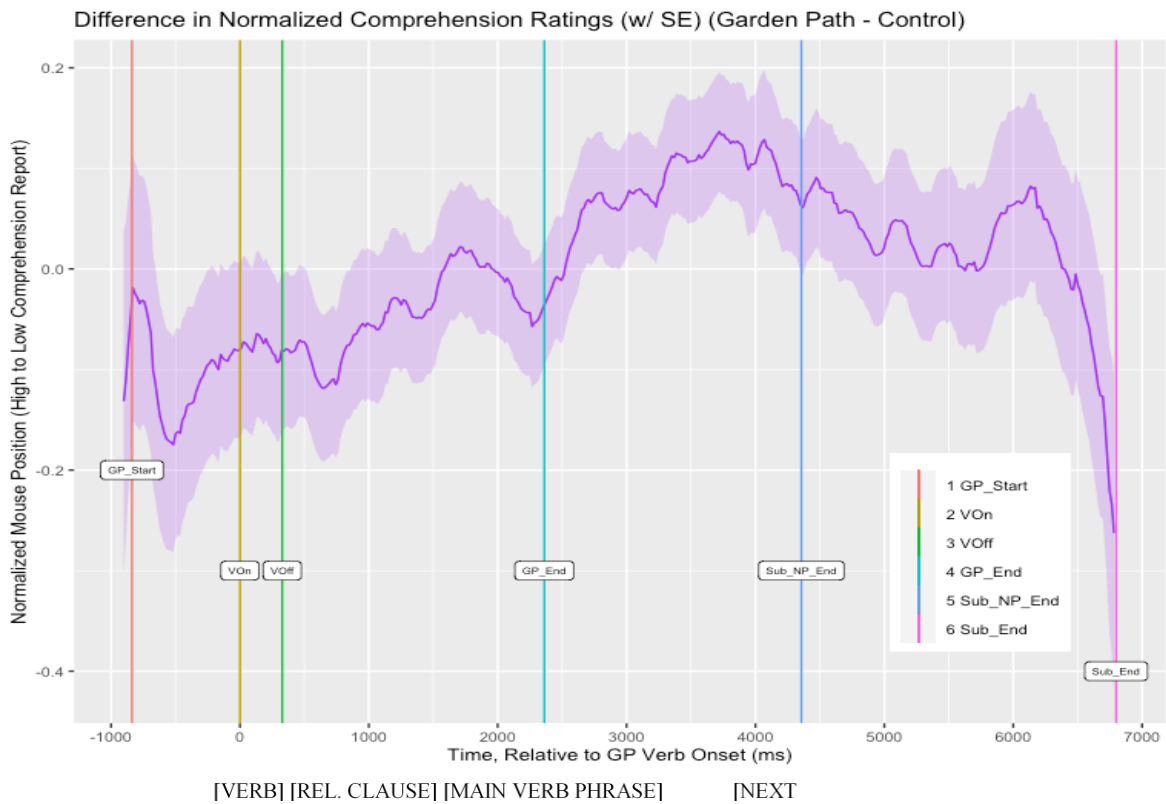
**Figure 7.** Averaged, normalized mouse positions (with standard error) for the time-window including the garden path sentence and subsequent sentence. As ease-of-comprehension drifts downward for the complicated critical sentence, it drops lower for the garden path version of the sentence (in blue) compared to the control version (in red).

that of the control condition, inverted so that greater difficulty with the garden path goes upward. Like **Figure 1**, the result shows a gradual rise in awareness of comprehension difficulty (or perhaps a form of error) as the garden path effect unfolds and spills partially into the beginning of the next sentence. However, as the new sentence continues, this “error measure” returns down to baseline as the sentence reveals itself to be rather simple syntactically and thus predictable.

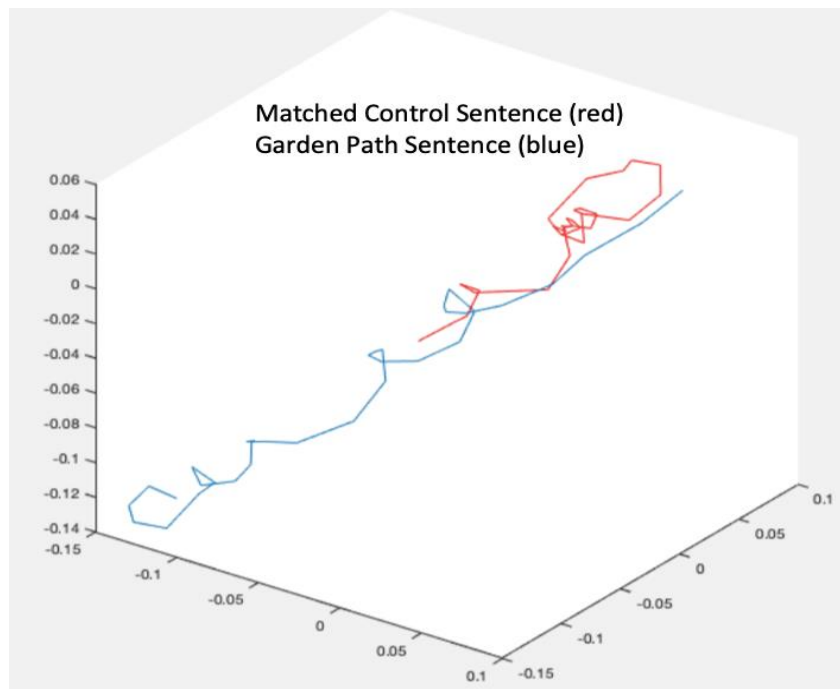
To further characterize the structure of these comprehension rating time series, RQA was applied to the data. Based on the RQA treatments of the simulated networks above, examining the higher-dimensional dynamics of this time series may provide insight into the processes involved during a garden path effect. For example, recent work shows that RQA can distinguish the portions of an orchestra’s performance where the instruments are coordinated and those where the instruments are uncoordinated – even just from a one-dimensional time series from a single audio source (Proksch et al., 2022). If the cognitive mechanisms that

process the syntax and semantics of a sentence are well coordinated during smooth comprehension and less coordinated during a garden path effect, perhaps RQA can reveal the characteristic dynamics of that reduced coordination.

As with the network analyses above, RQA begins with embedding the time series in a higher-dimensional state space (recall **Figure 4**). A recurrence analysis on a single time series is thus aimed at observing revisitations of locations in the embedded state space. For best data visualization, in this first treatment, we embedded the later temporal window time series (where the garden path effect is statistically detected) in three dimensions. **Figure 9** shows (in red) that the control sentence’s embedded trajectory loops back onto itself many times on both short and long-time scales, revealing numerous smooth revisitations of locations – as though generally settled into a comfortable metastable attractor. By contrast, the garden path version of the sentence produces a trajectory that loops onto itself only a few times locally and mostly extends out into state space – as though searching in vain for a metastable attractor.



**Figure 8.** The garden path condition's ease-of-comprehension data stream, minus that of the control condition, is inverted so that greater difficulty goes upward during the garden path (similar to the network error measure in **Figure 1**).



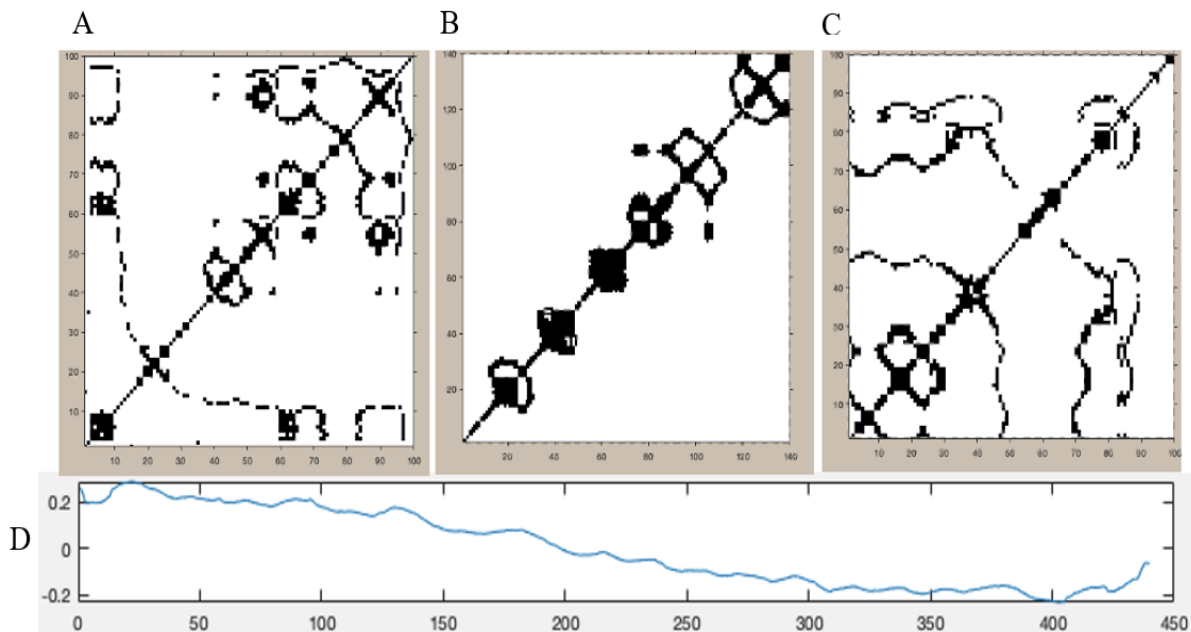
**Figure 9.** State-space reconstruction of the later temporal window for the garden path sentence (in blue) and the control version of the sentence (in red).

Following our analysis of Network A's node 52 (**Figure 5**) undergoing its difficult transition of accommodating disruptive input from Network B, here we apply the same analyses to the normalized ease-of-comprehension measure extracted from these participants while they undergo the experience of a garden path effect during naturalistic spoken language comprehension. See **Figure 10**. The long-distance recurrences in the state-space trajectory are immediately detectable in panels A and C (before and after the garden path) as black dots far from the main diagonal. By contrast, the garden path phase itself, panel B, has no long-distance recurrences in those upper-left and lower-right corners (similar that observed in **Figure 5B**). Much like with our RQA of Network A's node 52, it turns out that the measures of Entropy and Trapping Time again serve to distinguish the disruption phase of the garden path effect here (**Figure 10B**) from the rest of the time series before and after (**Figure 10A and C**).

In **Figure 10**, overall Recurrence goes from .087 (panel A) down to .076 (panel B) and back up to .085 (panel C). Correspondingly, Entropy starts out at 1.28 (panel A), climbs to 2.03 (panel B) and then returns to 1.25 – much

like it did in the data for **Figure 5**. Trapping Time (which indexes how long the system remains in any specific state) also starts low at 2.98 (panel A), climbs to 4.78 (panel B) and then returns to 2.91 (panel C). These RQA metrics for temporary network discoordination (**Figure 5**) and for temporary garden path confusion (**Figure 10**) are remarkably similar to one another. This observation offers support for treating the disrupted cognition that results from a garden path effect during sentence processing in terms of network dynamics (Gerth & Beim Graben, 2009; Tabor et al., 1997), and not solely in terms of abstract computational modules that are attempting to resolve conflicting commands for syntactic parsing (Frazier & Clifton, 1995; Staub, 2015).

Now we turn to comparing RQA of the garden path condition to the control condition. In this circumstance, it turns out that the set of RQA metrics that happen to be diagnostic of the difference between garden path and control is larger. The RQA metrics on which we focus here include *recurrence rate*, *entropy*, and *trapping time*, as before, but added to that set is *determinism* (based on the number of diagonals on the recurrence plot, quantifying the number of specific sequences of short-term local continuous recurrence) and *laminarity* (based



**Figure 10.** Recurrence plots of ease-of-comprehension before (A), during-and-after (B), and long after (C) the garden path was introduced. Panel D plots the normalized time series throughout this garden path disruption and recovery. (Compare to **Figure 5**.)

on the number of the vertical lines, which quantifies recurring visitations to particular values). When run on the full temporal window defined by the start of the garden path sentence to the end of the subsequent sentence (**Figure 11A, B**), RQA actually yields rather similar recurrence patterns across garden path and control conditions. In both conditions, recurrence is highly local, with temporal clustering, which appears as square-like shapes along the diagonal line of incidence. This similarity in structure is also reflected in the accompanying quantitative descriptions of the recurrence plot. Overall recurrence was .099 for the garden path condition and .115 for the control condition. Entropy was 2.97 for the garden path condition and 2.98 for the control condition. Trapping Time was 10.60 for the garden path condition and 10.82 for the control condition. Determinism was .98 for the garden path condition and .98 for the control condition. Laminarity was .99 for the garden path condition and .99 for the control condition. Thus, RQA of the full temporal window spanning from the beginning of the critical sentence to the end of the subsequent sentence covers such a wide swath of sentence processing dynamics that no substantial differences are observed between the garden path condition and the control condition.

However, a more focused RQA was applied on the temporal window between the end of the garden path sentence and the end of the next noun phrase. This is where the meta-cognitive awareness of the garden path was detected in **Figure 7** and where one sees a distinctive recurrence plot in **Figure 10B**. These analyses yielded robustly different recurrence patterns across garden path and control conditions. These differences were reflected qualitatively, by comparing the two recurrence plots in **Figure 11C & D**, and in the quantitative measures as well. For this more focused temporal window, overall recurrence was .082 for the garden path condition and .125 for the control condition. Entropy was 1.99 for the garden path condition and 1.40 for the control condition. Trapping Time was 4.52 for the garden path condition and 3.47 for the control condition. Determinism was .93 for the garden path condition and .77 for the control condition. Laminarity was .91 for the garden path condition and .83 for the control condition. Thus, recurrence patterns during the garden path effect were relatively more sporadic (lower recurrence rate, higher entropy) with a greater emphasis on short-timescale

recurrence (higher determinism and laminarity) than in the control condition during that same temporal window.

The dense-sampling nature of this continuous meta-cognitive ease-of-comprehension measure allows us to apply nonlinear time series analysis to the data, such as RQA. In particular, the discussions of **Figures 9, 10** and **11** suggest some insights into what exactly a subtly reportable syntactic garden path effect might be made of – not at an abstract computational level of description where rules and symbols are encountering conflict but at a concrete implementational level of analysis where the physical system of a set of cortical language networks is exhibiting quantitatively verifiable discoordination. When comprehension is going relatively smoothly, as in the control sentence condition, that trajectory appears to have a substantial amount of predictability (low entropy) and numerous loops or temporary limit cycles (e.g., high recurrence) that revisit certain regions even much later in the trajectory. See **Figure 9** (red trajectory) and **Figure 11D**. However, when comprehension is disrupted in the garden path condition, some of those smooth predictable loops (and later revisitations) disappear as the trajectory undergoes an increase in entropy. See **Figure 9** (blue trajectory), **Figure 10B**, and **Figure 11C**.

## **6. Measuring coordination and discoordination in language and in networks**

The viability and richness of mouse-tracking to measure fine-grain real-time meta-cognitive awareness during sentence processing is supported by the method's ability to capture the substantial difference in ease-of-comprehension ratings following the end of the garden path sentence, as compared to the matched control sentence. Once the method was proven to detect garden path effects, this dense-sampling measure opened up opportunities to use nonlinear time series analysis to investigate the temporal structure of the cognitive processes involved in resolving a temporary syntactic ambiguity – and potentially other effects in real-time language comprehension as well.

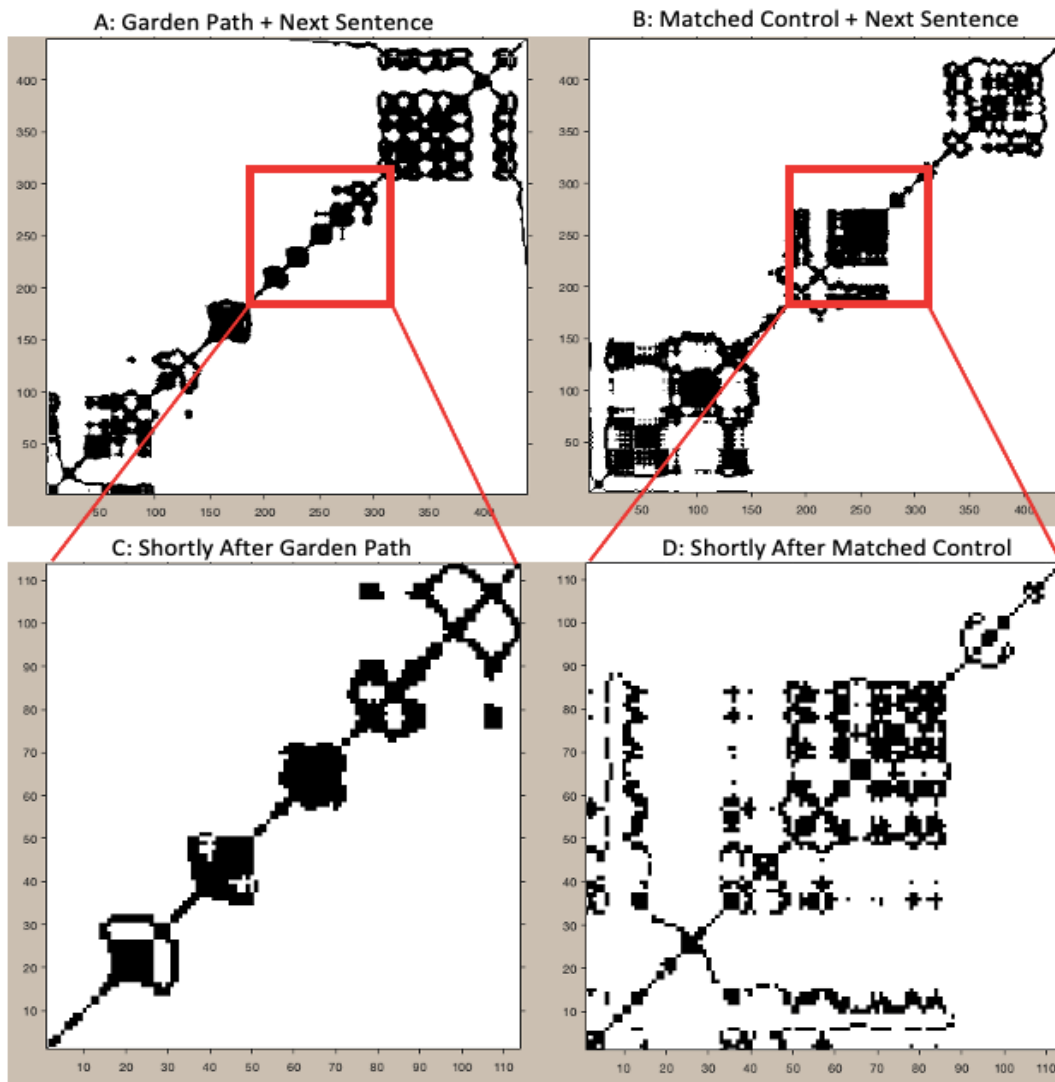
For several decades now, the account of how language processing gets disrupted when incorrect syntactic structures are pursued during comprehension (e.g., garden path effects) has been based on an abstract computational account of rules and symbols being used by a syntax



module inside the human mind (Frazier & Clifton, 1995; Frazier, & Fodor, 1978; Staub, 2015). The EEG measures of the N400 and the P600 have provided some snapshots into what the brain is doing at the moment of these disruptions in language comprehension (Kim & Osterhout, 2005; Kutas & Federmeier, 2011; Kutas & Hillyard, 1980) but have not provided insight into the continuous temporal dynamics of this process. By contrast, artificial neural network accounts (Gerth & Beim Graben, 2009; Tabor, et al., 1997) and dynamical systems theory treatments (Elman, 2009; Onnis & Spivey,

2012) have pointed to a description of language comprehension that emphasizes not rules and symbols but instead a continuous trajectory through neuronal state space – from which an extracted time series could be very informative.

By juxtaposing such time series analyses of activity measured from a fully recurrent network undergoing disruption and of activity measured from a continuous meta-cognitive report task focused on disruption in language comprehension, we have attempted here to pro-



**Figure 11.** (A) Recurrence plot from the full temporal window spanning from the beginning of the garden path sentence to the end of the subsequent sentence (the full range of **Figure 7**). (B) Recurrence plot from the full temporal window spanning from the beginning of the matched control sentence to the end of the subsequent sentence. Framed in red is the later temporal window where participant reports show sensitivity to the garden path effect. (C) Recurrence plot of that later temporal window following a garden path sentence, revealing low recurrence and greater entropy. (D) Recurrence plot of that later temporal window following a matched control sentence.

vide more insight into the statistical character of those syntactically disrupted (and non-disrupted) trajectories through state space. In the reconstructed state-space trajectory that RQA relies on, disrupted or discoordinated language comprehension appears to exhibit reduced long-time-scale auto-correlations (in the form of the overall Recurrence measure), increased short-time-scale auto-correlations (in the form of the Trapping Time, Determinism and Laminarity measures), and increased unpredictability (in the form of the Entropy measure). By contrast, in the state-space trajectory of undisrupted or coordinated language comprehension, there appear to be greater long-time-scale auto-correlations, less short-time-scale auto-correlations, and less unpredictability.

Although we explore only the garden path effect in the present work, this method is highly flexible. This continuous meta-cognitive ease-of-comprehension measure does not depend on a visual context, such as the visual world paradigm (Tanenhaus et al., 1995). Therefore, this method can use sentences of any type, including those with highly abstract or figurative semantic content. In general, the time series data from this method affords a broader range of analyses for investigating linguistic processing. In addition to these subtle garden path disruptions, other perturbations can be tested, such as overt semantic and syntactic anomalies, linguistic negation, low-frequency words and figurative language use, to name just a few. By obtaining data that can avail of nonlinear time series analysis for a wide variety of language tasks (as well as other continuous temporal signals), this novel method may open up new avenues for developing a deeper understanding of language, cognition, perception, and action.

Abstract computational treatments of cognition in general (Griffiths & Tenenbaum, 2006) and of language in particular (Ryskin et al., 2021; Staub, 2015) have made important advances in our understanding of how a mind processes complex input. However, those frameworks sometimes leave unexamined the specifics of exactly how a reconfigurable network of cortical language

subsystems (Chai et al., 2016) achieves its smooth coordination during felicitous comprehension and how it undergoes some discoordination during briefly infelicitous comprehension. By comparing the statistical character of a data stream that is extracted from a neural network simulation (where everything can be known about what is going on inside) to that of a data stream extracted from a person processing language input (where one must infer what is going on inside), we suggest that some progress can be made in understanding aspects and parameters of the simulation that might correspond to certain aspects and parameters in the person (e.g., Spivey, 2018). For example, there may turn out to be certain statistical characteristics in a time series of cognitive performance, such as multiscale temporal structure (Van Orden et al., 2003), that are naturally achievable with certain implementation-level models (Kello, 2013) but not naturally achievable with certain abstract computational models (Wagenmakers et al., 2005). In this way, implementation-level analyses (such as those from neuroscience and neural network simulations) may provide constraints on which abstract computation-level models should be pursued in the cognitive sciences.

## References

- Amit, D. (1989) *Modeling brain function: The world of attractor neural networks*. New York, New York: Cambridge University Press.
- Anderson, J. A. (1995) *An introduction to neural networks*. Cambridge, New York: MIT Press.
- Atmanspacher, H. & Beim Graben, P. (2009) Contextual emergence. *Scholarpedia* **4**, 7997.
- Bard, E. G., Robertson, D. & Sorace, A. (1996) Magnitude estimation of linguistics acceptability. *Language* **72**, 32-68.
- Bever, T. G. (1970) The Cognitive basis for linguistic structures. In, J. Hayes (Ed.), *Cognition and the Development of Language*. New York, New York: Wiley.

- Bishop, R., Silberstein, M. & Pexon, M. (2022) *Emergence in context*. New York, New York: Oxford University Press.
- Bloom, L., Hood, L. & Lightbown, P. (1974) Imitation in language development: If, when, and why. *Cognitive Psychology* **6**, 380-420.
- Branigan, H. (2007) Syntactic priming. *Language and Linguistics Compass* **1**, 1-16.
- Brennan, S. E. & Clark, H. H. (1996) Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **22**, 1482-1493.
- Campbell, J. (1982) *Grammatical Man: Information, Entropy, Language, and Life*. New York: Simon and Schuster.
- Carroll, S., Swain, M. & Roberge, Y. (1992) The role of feedback in adult second language acquisition: Error correction and morphological generalizations. *Applied Psycholinguistics* **13**, 173-198.
- Chai, L. R., Mattar, M. G., Blank, I. A., Fedorenko, E. & Bassett, D. S. (2016) Functional network dynamics of the language system. *Cerebral Cortex* **26**, 4148-4159.
- Chialvo, D. R. (2010) Emergent complex neural dynamics. *Nature Physics* **6**, 744-750.
- Clark, A. (2015) *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York, New York: Oxford University Press.
- Dale, R. & Kello, C. T. (2018) "How do humans make sense?" multiscale dynamics and emergent meaning. *New Ideas in Psychology* **50**, 61-72.
- Di Paolo, E. A., Thompson, E. & Beer, R. D. (2022) Laying down a forking path: Tensions between enaction and the free energy principle. *Philosophy and the Mind Sciences* **3**, 2.
- Dotov, D. G., Nie, L. & Chemero, A. (2010) A demonstration of the transition from ready-to-hand to unready-to-hand. *PLoS One* **5**, e9433.
- Dubois, D.M. (2003) Mathematical foundations of discrete and functional systems with strong and weak anticipations. In, M.V.Butz, O. Sigaud & P. Gerard (Eds.) *Anticipatory Behavior in Adaptive Learning Systems*. Berlin: Springer.
- Elman, J. L. (1990) Finding structure in time. *Cognitive Science* **14**, 179-211.
- Elman, J. L. (1991) Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, **7**, 195-225.
- Elman, J. L. (2009) On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive Science* **33**, 547-582.
- Falandays, J. B., Batzloff, B. J., Spevack, S. C. & Spivey, M. J. (2020) Interactionism in language: from neural networks to bodies to dyads. *Language, Cognition and Neuroscience* **35**, 543-558.
- Falandays, J. B., Nguyen, B. & Spivey, M. J. (2021) Is prediction nothing more than multi-scale pattern completion of the future? *Brain Research* **1768**, 147578.
- Falandays, J. B., Yoshimi, J., Warren, W., & Spivey, M. J. (2023) A potential mechanism for Gibsonian resonance: Behavioral entrainment emerges from local homeostasis in an unsupervised reservoir network (*submitted*).
- Farmer, T. A., Anderson, S. E., & Spivey, M. J. (2007) Gradiency and visual context in syntactic garden-paths. *Journal of Memory and Language* **57**, 570-595.
- Frazier, L. & Clifton, C. (1995) *Construal*. Cambridge, MA: MIT Press.
- Fusaroli, R., Bahrami, B., Olsen, K., Roepstorff, A., Rees, G., Frith, C. & Tylén, K. (2012) Coming to terms: Quantifying the benefits of linguistic coordination. *Psychological Science* **23**, 931-939.
- Gerth, S., & beim Graben, P. (2009) Unifying syntactic theory and sentence processing difficulty through a connectionist minimalist parser. *Cognitive Neurodynamics* **3**, 297-316.

- Griffiths, T. L. & Tenenbaum, J. B. (2006) Optimal predictions in everyday cognition. *Psychological Science* **17**, 767-773.
- Grossberg, S. (2012) Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks* **37**, 1-47.
- Hauk, O., Giraud, A. L., & Clarke, A. (2017) Brain oscillations in language comprehension. *Language, Cognition and Neuroscience* **32**, 533-535.
- Hinaut, X. & Dominey, P. F. (2013) Real-time parallel processing of grammatical structure in the fronto-striatal system: A recurrent network simulation study using reservoir computing. *PLoS One* **8**, e52946.
- Hirsh-Pasek, K., Treiman, R. & Schneiderman, M. (1984) Brown & Hanlon revisited: Mothers' sensitivity to ungrammatical forms. *Journal of Child Language* **11**, 81-88.
- Hohwy, J. (2013) *The Predictive Mind*. New York, New York: Oxford University Press.
- Kaan, E., Harris, A., Gibson, E. & Holcomb, P. (2000) The P600 as an index of syntactic integration difficulty. *Language and Cognitive Processes* **15**, 159-201.
- Kello, C. T. (2013) Critical branching neural networks. *Psychological Review* **120**, 230-254.
- Kello, C., Bhat, H., Turner, M. & Alviar, C. (2023) Hierarchical process timescales in behavioral and neural activity (*submitted*).
- Kim, A. & Osterhout, L. (2005) The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language* **52**, 205-225.
- Kuhlen, A. K., Allefeld, C. & Haynes, J. D. (2012) Content-specific coordination of listeners' to speakers' EEG during communication. *Frontiers in Human Neuroscience* **6**, 266.
- Kutas, M. & Federmeier, K. D. (2011) Thirty years and counting: finding meaning in the N400 component of the event-related brain potential (ERP). *Annual Review of Psychology* **62**, 621-647.
- Kutas, M. & Hillyard, S. (1980) Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science* **207**, 203-205.
- Lewis, A. G., Schoffelen, J. M., Hoffmann, C., Bastiaansen, M. & Schriefers, H. (2017) Discourse-level semantic coherence influences beta oscillatory dynamics and the N400 during sentence comprehension. *Language, Cognition and Neuroscience* **32**, 601-617.
- Lopes da Silva, F. L. (1991) Neural mechanisms underlying brain waves: from neural membranes to networks. *Electroencephalography and Clinical Neurophysiology* **79**, 81-93.
- Louwerse, M. M., Dale, R., Bard, E. G. & Jeuniaux, P. (2012) Behavior matching in multimodal communication is synchronized. *Cognitive Science*, **36**, 1404-1426.
- Markman, E. M. & Wachtel, G. F. (1988) Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology* **20**, 121-157.
- Marmelat, V. & Delignières, D. (2012) Strong anticipation: complexity matching in interpersonal coordination. *Experimental Brain Research*, **222**, 137-148.
- Marwan, N., Romano, M. C., Thiel, M. & Kurths, J. (2007) Recurrence plots for the analysis of complex systems, *Physics Reports* **438**, 237-329.
- McRae, K., De Sa, V. R. & Seidenberg, M. S. (1997) On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General* **126**, 99-130.
- McRae, K., Spivey-Knowlton, M. J. & Tanenhaus, M. K. (1998) Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language* **38**, 283-312.

- Meyer, T., Kim, A. D., Spivey, M. J. & Yoshimi, J. (2023) Mouse-tracking performance: A new approach to analyzing continuous mouse tracking data (*submitted*).
- Nieuwland, M. S. & Van Berkum, J. J. (2006) When peanuts fall in love: N400 evidence for the power of discourse. *Journal of Cognitive Neuroscience* **18**, 1098-1111.
- Onnis, L. & Spivey, M. J. (2012) Toward a new scientific visualization for the language sciences. *Information* **3**, 124-150.
- Osterhout, L. & Holcomb, P. (1992) Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language* **31**, 785-806.
- Patel, A. D. (2003) Language, music, syntax and the brain. *Nature Neuroscience* **6**, 674-681.
- Prince, A. & Smolensky, P. (2004) *Optimality Theory: Constraint Interaction in Generative Grammar*. New York, New York: Wiley.
- Proksch, S., Reeves, M., Spivey, M. & Balasubramaniam, R. (2022) Coordination dynamics of multi-agent interaction in a musical ensemble. *Scientific Reports* **12**, 421.
- Rabovsky, M. & McRae, K. (2014) Simulating the N400 ERP component as semantic network error: Insights from a feature-based connectionist attractor model of word meaning. *Cognition* **132**, 68-89.
- Rabovsky, M., Hansen, S. S. & McClelland, J. L. (2018) Modelling the N400 brain potential as change in a probabilistic representation of meaning. *Nature Human Behaviour* **2**, 693-705.
- Raja, V., Valluri, D., Baggs, E., Chemero, A. & Anderson, M. L. (2021) The Markov blanket trick: On the scope of the free energy principle and active inference. *Physics of Life Reviews* **39**, 49-72.
- Ramscar, M., Dye, M. & McCauley, S. M. (2013) Error and expectation in language learning: The curious absence of "mouses" in adult speech. *Language* **89**, 760-793.
- Richardson, D. & Dale, R. (2005) Looking to understand: The coupling between speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive Science* **29**, 1045-1060.
- Ryskin, R., Stearns, L., Bergen, L., Eddy, M., Fedorenko, E. & Gibson, E. (2021) An ERP index of real-time error correction within a noisy-channel framework of human communication. *Neuropsychologia* **158**, 107855.
- Shannon, C.E. & Weaver, W. (1949) *The Mathematical Theory of Communication*. Urbana, IL: The University of Illinois Press.
- Shockley, K., Santana, M. V. & Fowler, C. A. (2003) Mutual interpersonal postural constraints are involved in cooperative conversation. *Journal of Experimental Psychology: Human Perception and Performance* **29**, 326-332.
- Spivey, M. & Richardson, D. (2009) Language processing embodied and embedded. In P. Robbins and M. Aydede (Eds.), *The Cambridge Handbook of Situated Cognition*. Cambridge, UK: Cambridge University Press.
- Spivey, M. J. (2007) *The Continuity of Mind*. New York, New York: Oxford University Press.
- Spivey, M. J. (2018) Discovery in complex adaptive systems. *Cognitive Systems Research*, **51**, 40-55.
- Spivey, M. J. & Tanenhaus, M. K. (1998) Syntactic ambiguity resolution in discourse: modeling the effects of referential context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **24**, 1521-1543.
- Spivey-Knowlton, M. J., Trueswell, J. C. & Tanenhaus, M. K. (1993) Context effects in syntactic ambiguity resolution: Discourse and semantic influences in parsing reduced relative clauses. *Canadian Journal of Experimental Psychology* **47**, 276-309.

Staub, A. (2015) Reading sentences: Syntactic parsing and semantic interpretation. In A. Pollatsek & R. Treiman (Eds.), *The Oxford Handbook of Reading*. Oxford, UK: Oxford University Press.

Stephen, D. G., Boncoddio, R. A., Magnuson, J. S. & Dixon, J. A. (2009) The dynamics of insight: Mathematical discovery as a phase transition. *Memory & Cognition* **37**, 1132-1149.

Stepp, N. & Turvey, M.T. (2010) On strong anticipation. *Cognitive Systems Research* **11**, 148-164.

Tabor, W., Juliano, C. & Tanenhaus, M. K. (1997) Parsing in a dynamical system: An attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes* **12**, 211-271.

Takens, F. (1981) Detecting strange attractors in turbulence. In *Lecture Notes in Mathematics* (pp. 366-381). New York, New York: Springer-Verlag.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M. & Sedivy, J. C. (1995) Integration of visual and linguistic information in spoken language comprehension. *Science* **268**, 1632-1643.

Van Orden, G. C., Holden, J. G. & Turvey, M. T. (2003) Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, **132**, 331-350.

Wagenmakers, E. J., Farrell, S. & Ratcliff, R. (2005) Human cognition and a pile of sand: a discussion on serial correlations and self-organized criticality. *Journal of Experimental Psychology: General* **134**, 108-116.