Analysis of a stochastic dynamical model of word-by-word sentence comprehension

Garrett Smith

May 4, 2021

Abstract

An important class of computational cognitive models is based on stochastic dynamical systems. The SWIFT model of eye-movement control (Engbert, Longtin, & Kliegl, 2002) and the drift-diffusion model of perceptual decision making (Ratcliff, 1978), for example, have driven progress in part because of the mathematical tools that the models afford researchers in analyzing experimental data and deriving new predictions. Stochastic dynamical models have also been applied to word-by-word reading comprehension; however, their influence has been limited so far by relatively ad hoc, analytically opaque implementations. Here, I describe a new model of incremental sentence comprehension that models the parsing process as a series of continuous-time, discrete-state random walks among potential syntactic analyses of the sentence so far. Reading can now be framed as a first-passage problem: how long does it take arrive at a complete parse of the sentence so far, given the preceding words? This note describes in detail the derivation of various stochastic quantities of interest, including predicted reading time distribution functions, and illustrates them with example simulations. The hope is that these analytical tools can drive new progress in the theory of word-by-word reading comprehension.

1 Introduction

Dynamical systems models of cognitive processes characterize behavior in terms of how the system under consideration changes in time (Beer, 2000; van Gelder, 1998). Dynamical systems models have generated new insights about cognitive phenomena. For example, the SWIFT model of eye-movement control during reading (Engbert et al., 2002; Engbert, Nuthmann, Richter, & Kliegl, 2005; Seelig et al., 2020), which is closely related to the model discussed below, has been influential in understanding

how certain properties of written text affect fine-grained saccade planning and timing. Similarly, the drift-diffusion model has guided studies of decision making over multiple decades (Ratcliff, 1978; Ratcliff, Smith, Brown, & McKoon, 2016). These stochastic, dynamical models can attribute their theoretical and empirical success in part to their clear, tractable mathematical framing.

Stochastic dynamical models have also been used to understand and predict word-by-word sentence comprehension effects (beim Graben, Gerth, & Vasishth, 2008; Cho, Goldrick, & Smolensky, 2017; Kempen & Vosse, 1989; Smith, 2018; Smith, Franck, & Tabor, 2018, 2021; Stevenson, 1994; Tabor & Hutchins, 2004; Vosse & Kempen, 2000, 2009). The most common type of data to be explained is reading times, measured by recording eye movements or by self-paced reading (Just, Carpenter, & Woolley, 1982). Examples of previous dynamical sentence comprehension models include Tabor and Hutchins (2004) and Smith et al. (2018), who used stochastic differential equations to model changes in the relative strengths of competing parses. While the range of empirical phenomena that these and related models can explain is considerable, they have not been widely adopted. One reason for this is their relatively complicated mathematical implementations, which hamper formal analysis and parameter fitting. This makes extending or even reproducing previous results difficult.

Recently, Smith (2021) proposed a new dynamical theory of the sentence comprehension process. The model, called *mparse*, is based on the idea that sub-steps in the parsing process at each word correspond to discrete states that the model explores stochastically. This continuous-time Markov process can be described using the master equation, a set of deterministic ordinary differential equations that describe how the probability of different states changes in time (van Kampen, 2007). In contrast with most previous stochastic dynamical models in sentence comprehension, implementing *mparse*'s dynamics using the master equation provides a well-understood and readily analyzable framework for exploring the model's behavior. As described below, *mparse* makes word-by-word reading time predictions by asking how long it takes the model to reach a complete parse of the sentence so far, thus framing the question as a *first-passage time* problem. The purpose of this note is to explain and summarize the relevant results on first passage time statistics as a complement to the descriptions of one-step systems typically discussed in textbooks (Gardiner, 1985; Redner, 2001; van Kampen, 2007). While *mparse* is used to illustrate the math, the results are applicable to any continuous-time, discrete state stochastic process, with potential for broader application in cognitive science. The results are not new, but applying them to a model of sentence comprehension is new (to the best of my knowledge). The hope is that by summarizing the first-passage results here, they can be applied more readily in psycholinguistic models and cognitive science more generally.

2 The mparse model

At each word w in a sentence, mparse creates n_w discrete states which represent partial or complete parses of the string of words that has been encountered so far. For simplicity, we assume that the model “knows” a set of head-dependent relationships that are allowed in the language it is processing. The head is the word that licenses the existence of the dependent, which in turn modifies the head (Gaifman, 1965; Hays, 1964; Kübler, McDonald, & Nivre, 2009). For simplicity, dependencies are not labeled with a type (like subject or determiner), although the grammar can be extended to include typed dependencies (Kübler et al., 2009; Smith, 2021).

As a first example, consider the string *the cat*. After reading these two words, mparse creates three discrete states (see Fig. 1). The first is a no-structure state: no dependencies are established between the words. The second is a state consisting of the parse in which *the* attaches as the dependent of *cat*, the most complete parse of these two words. The dependency between *the* and *cat* in Fig. 1 is shown by the thin arcs pointing from the head to the dependent word. The final state is identical to the previous state but represents the decision to move on to the next word. This state is an absorbing state; once the model reaches it, it cannot go back and visit other states. Instead the next word is input and processing resumes again with an updated state space.

Figure 1: Network of possible transitions between the states (surrounded by boxes) available after reading *the cat*. Thick arrows indicate possible jumps between states. The no-structure and complete-parse states are transient states. Within each state, the thin arcs point from the head to the dependent. Note that the parses in the complete parse and absorbing states are identical; jumping to the latter marks the decision to move on to the next word of the sentence.

Having enumerated the states available at word w , mparse then jumps stochastically between states that differ by a single link (nearest-neighbor transitions). This continuous-time Markov process can be described using the master equation (van Kampen, 2007). For each state n , the master equation describes how the probability p_n of that state changes in time. Jumps between state n and state m (shown as thick arrows in Fig. 1) occur at the rates W_{nm} , which is the probability of jumping from m to n per unit time. The master equation for state n can thus

be written as

$$\frac{dp_n(t)}{dt} = \sum_{m \neq n} W_{nm} p_m(t) - p_n(t) \sum_{m \neq n} W_{mn}$$

The master equation describes how probability flows into state n from other states m (first term) and how it flows away from n to other states m (second term). The master equation can also be written in matrix form as Eq. 1:

$$\frac{d\mathbf{p}(t)}{dt} = \mathbf{W}\mathbf{p}(t) \quad (1)$$

The matrix $\mathbf{W}$ contains the transition rates W_{nm} with the diagonal elements $W_{nn} = \sum_{m \neq n} W_{mn}$, and $\mathbf{p}(t)$ is a column vector containing all of the $p_n(t)$. Note that for an absorbing state a , $W_{na} = 0 \forall n \neq a$. Once mparse reaches a , it cannot return to the other states (called transient states; note the lack of thick arrows leaving the absorbing state in Fig. 1). Instead, mparse reads in the next word, updates its state space with any new states made available by the presence of the new word, and resumes the random walk again.

To solve the master equation, we must provide an initial probability distribution over states $\mathbf{p}(0)$. At first word of the sentence, all probability is placed on the no-structure state ($p_n(0) = \delta_{n, \text{no struct.}}$), but for later words, $\mathbf{p}(0)$ depends on the processing at the previous word (see below). Given the initial state, the formal solution to the master equation is

$$\mathbf{p}(t) = e^{\mathbf{W}t} \mathbf{p}(0) = \left(\sum_{l=0}^{\infty} \frac{t^l}{l!} \mathbf{W}^l \right) \mathbf{p}(0) \quad (2)$$

The solution can also be written in terms of the n_w left and right eigenvectors ($\mathbf{l}$ and $\mathbf{r}$) and eigenvalues (λ) of the matrix $\mathbf{W}$ (Oppenheim, Shuler, & Weiss, 1977):

$$\mathbf{p}(t) = \sum_{j=1}^{n_w} (\mathbf{l}_j^T \mathbf{p}(0)) \mathbf{r}_j e^{\lambda_j t}$$

where $\top$ denotes the transpose of a vector.

Using the dynamics governed by this master equation, mparse explores the set of partial parses (i.e., its states) that is possible at each word w until it gets absorbed in an absorbing state. For the string *the cat*, mparse can jump between the transient no-structure and complete-parse states shown in Fig. 1 until it jumps to the absorbing state. At that point, the model inputs the next word (e.g., *sleeps*), adds new states depending on the parses available with the new word (e.g., having a dependency between *sleeps* and *cat*), and repeats the whole process until there are no more words in the sentence to input. The state it reaches at the end of the last word corresponds to the parse it builds for the whole sentence. In the

next section, I show how we can use the master equation approach of mparse to provide formulas for a number of first-passage statistics and apply them to a simple example sentence.

3 First-passage time statistics

Given this formal model, we want to know how long it takes to process each word or a string of words. We can model this as a first passage time problem, where we ask how long and with what probability mparse reaches a particular absorbing state. To make the discussion more concrete, I will use the ambiguous sentence in Ex. (1):

- (1) The ranger saw the hunter with binoculars.

This sentence can either be interpreted as the ranger using the binoculars to see the hunter (the “instrument” parse) or as the ranger seeing the binoculars-carrying hunter (the “modifier” parse; see Fig. 2). The first interpretation requires that the prepositional phrase *with binoculars* attaches as the dependent of the verb *saw*, while the second requires that *with binoculars* be the dependent of the noun phrase *the hunter*. All of the results below are implemented in Python at <https://tinyurl.com/MparseDemo>.

After reading *saw the hunter with the binoculars*, mparse explores the transient states ("No structure," "Partial parse," "Instrument parse," and "Modifier parse" in Fig. 2) by jumping between states that differ only by a single dependency link. If the model is in either the instrument or modifier parse, there is some probability that it can jump to the corresponding absorbing state instead of jumping back to a different transient state. To model reading times, we ask how long it takes mparse to find an absorbing state. We can also ask what the probabilities are of ending up in the instrument or modifier absorbing states. I now show how to calculate these quantities and use them to make reading time predictions.

For concreteness, the $\mathbf{W}$ matrix for this example is listed in 3 below.

$$\mathbf{W} = \begin{bmatrix} W_{00} & W_{10} & 0 & 0 & 0 & 0 \\ W_{10} & W_{11} & W_{12} & W_{13} & 0 & 0 \\ 0 & W_{21} & W_{22} & 0 & 0 & 0 \\ 0 & W_{31} & 0 & W_{33} & 0 & 0 \\ 0 & 0 & W_{42} & 0 & 0 & 0 \\ 0 & 0 & 0 & W_{53} & 0 & 0 \end{bmatrix} \quad (3)$$

Figure 2: Network of possible transitions between the states available after reading Ex. (1). The numbers serve as labels for the states in the equations in the text.

3.1 Exit time distribution

For many sentences, there will be a single absorbing state for most words in the sentence. But in general, there can be more than one absorbing state, as shown in Fig. 2 for Ex. (1). A first question we can ask is how long it takes to find either one of the absorbing states (it does not matter which), the *exit time* from the transient states into one of the absorbing states (van Kampen, 2007). To do this, we begin by partitioning the matrix $\mathbf{W}$ into submatrices $\mathbf{T}$ and $\mathbf{A}$.

$$\mathbf{W} = \begin{bmatrix} \mathbf{T} & \mathbf{0} \\ \mathbf{A} & \mathbf{0} \end{bmatrix}$$

The matrix $\mathbf{T}$ contains the transitions within the set $\mathcal{T}$ of transient states, and the matrix $\mathbf{A}$ contains the transitions from the transient states to the absorbing states (given by the set $\mathcal{A}$). The $\mathbf{0}$ are zero matrices of appropriate dimension. The vectors $\mathbf{p}_{\mathcal{T}}$ and $\mathbf{p}_{\mathcal{A}}$ denote the vectors of probabilities of the transient and absorbing states, respectively, with $\mathbf{p}_{\mathcal{T}}(t) + \mathbf{p}_{\mathcal{A}}(t) = 1 \forall t$. The matrices and vectors for the example sentence are:

$$\mathbf{T} = \begin{bmatrix} W_{00} & W_{10} & 0 & 0 \\ W_{10} & W_{11} & W_{12} & W_{13} \\ 0 & W_{21} & W_{22} & 0 \\ 0 & W_{31} & 0 & W_{33} \end{bmatrix} \quad \mathbf{A} = \begin{bmatrix} 0 & 0 & W_{42} & 0 \\ 0 & 0 & 0 & W_{53} \end{bmatrix}$$

$$\mathbf{p}_{\mathcal{T}} = \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix} \quad \mathbf{p}_{\mathcal{A}} = \begin{bmatrix} p_4 \\ p_5 \end{bmatrix}$$

The total probability of having already reached an absorbing states at time t is $S(t) = \sum_{a \in \mathcal{A}} p_a(t)$. Thus, the probability of being absorbed into any absorbing state in the time between t and $t + dt$ is given by the time derivative of $S(t)$. This quantity is the exit time distribution $f(t)$:

$$f(t) = \frac{dS(t)}{dt} = \sum_{a \in \mathcal{A}} \frac{dp_a(t)}{dt} = \sum \frac{d\mathbf{p}_{\mathcal{A}}}{dt}$$

Substituting in the relevant terms from Eq. 1 and Eq. 2 and noting that the absorbing states only receive probability from the transient states, we arrive at the explicit form of the exit time distribution:

$$\begin{aligned} f(t) &= \sum \mathbf{A} \mathbf{p}_{\mathcal{T}}(t) \\ &= \sum \mathbf{A} e^{\mathbf{T}t} \mathbf{p}_{\mathcal{T}}(0) \end{aligned} \tag{4}$$

The matrix exponential can be evaluated numerically (as done in the demonstration online at <https://tinyurl.com/MparseDemo>) or by substituting in the eigenvalue decomposition shown above.

The m -th non-central moment μ_m of $f(t)$ is given by solving $\mu_m = \int_0^\infty t^m f(t) dt$ (Oppenheim et al., 1977):

$$\mu_m = (-1)^m m! \mathbf{1}^\top \mathbf{T}^{-m} \mathbf{p}_\mathcal{T}(0)$$

The vector $\mathbf{1}$ is a vector of ones of length $|\mathcal{T}|$. Thus, the mean exit time is $\mu_1 = -\mathbf{1}^\top \mathbf{T}^{-1} \mathbf{p}_\mathcal{T}(0)$, and the variance is $\mu_2 - \mu_1^2 = 2\mathbf{1}^\top \mathbf{T}^{-2} \mathbf{p}_\mathcal{T}(0) - \mu_1^2$.

In mparse, the exit time distribution at each word or multi-word region (Eq. 4), along with its mean and variance, can be compared to human reading time distributions. Eq. 4 can also be used as the basis for a likelihood function for fitting parameters in the matrix $\mathbf{W}$ to human experimental data.

Without loss of generality, we can set all of the $W_{nm} = 1.0$ to illustrate. The exit time distribution (in arbitrary time units) for the example sentence is shown in Fig. 3. The mean of the distribution is 4.0 and the variance is 10.0. Note that the mean exit time is considerably higher than the mode of the exit time distribution (approximately 1.5); this is due to positive skew in the exit time distribution. Long exit times are possible because the model will sometimes jump around among the transient states for a long time before finally jumping to an absorbing state.

Figure 3: Exit time distribution for Ex. (1).

3.2 Multiple absorbing states

The example sentence in Ex. (1) has more than one absorbing state, each corresponding to a possible analysis of the sentence. In cases such as this, we can go beyond the exit time distribution and calculate the probability of finding different absorbing states (called the splitting probabilities) and the probability distribution of times to reach one absorbing state before any of the others (the conditional first passage times; van Kampen, 2007). As noted above, probability flows to the absorbing states exclusively from the transient states, so the solution to the master equation for just the absorbing states is Eq. 5:

$$\mathbf{p}_A(t) = \mathbf{A}e^{\mathbf{T}t}\mathbf{p}_T(0) \quad (5)$$

To figure out the splitting probability π_a for the a -th absorbing state, we integrate Eq. 5 over time (Valleriani, Li, & Kolomeisky, 2014; van Kampen, 2007):

$$\pi_a = \int_0^\infty p_a(t)dt$$

Plugging in Eq. 5 and solving the integral, we arrive at

$$\begin{aligned} \pi_a &= \int_0^\infty [\mathbf{A}e^{\mathbf{T}t}\mathbf{p}_T(0)]_a dt \\ &= -[\mathbf{A}\mathbf{T}^{-1}\mathbf{p}_T(0)]_a \end{aligned} \quad (6)$$

where $[\cdot]_a$ denotes the a th element of the vector in brackets. Eq. 6 holds for all absorbing states a .

For the example sentence, if we assume mparses starts in the no-structure state ($\mathbf{p}_T(0) = [1, 0, 0, 0]^\top$), the splitting probabilities are $[0.5, 0.5]^\top$. If, on the other hand, we assume that speakers of English prefer the modifier parse (i.e., the hunter has the binoculars; Frazier, 1978) and increase the transition rates W_{31} and W_{53} to 2.0 instead of 1.0, we get splitting probabilities of $[0.27, 0.73]^\top$, a preference for attaching *with the binoculars* as the modifier of *hunter*.

Once we have the splitting probabilities π_a for the absorbing states a , we can now calculate the conditional first-passage time distributions for each a . The conditional first-passage time distribution for a is the distribution of first passages to a conditional on not having been absorbed in another absorbing state a' . This contrasts with the exit time, which is the time to reach any absorbing state at all instead of a particular one. Instead of summing across all of the absorbing states in Eq. 4, we take each element of the vector $\mathbf{A} \exp(\mathbf{T}t)\mathbf{p}_T(0)$ and normalize by the corresponding splitting probability:

$$f_a(t) = \frac{[\mathbf{A}e^{\mathbf{T}t}\mathbf{p}_T(0)]_a}{\pi_a} \quad (7)$$

Eq. 7 thus gives the probability distribution of getting absorbed in state a given the initial condition $\mathbf{p}(0)$. This can lead directly to differing predictions about reading times conditional which structure is built when processing a word. The conditional first-passage time distributions for the two absorbing states are shown in Fig. 4.

Figure 4: Conditional first-passage time distributions for Ex. (1), assuming a bias toward the modifier parse.

The mean conditional first passage time is given in Eq. 8 (Polizzi, Therien, & Beratan, 2016):

$$\begin{aligned}\tau_a &= \frac{1}{\pi_a} \int_0^\infty t f_a(t) dt \\ &= \frac{[\mathbf{T}^{-2} \mathbf{p} \tau(0)]_a}{\pi_a}\end{aligned}\tag{8}$$

For the example sentence (including the bias for the modifier parse), the mean first-passage time to the instrument parse is 2.85 and 2.68 to the modifier parse. These are considerably shorter than the mean exit time. This is because the mean first-passage times are conditioned on *not* being absorbed by a different absorbing state. Intuitively, if mparse is to be absorbed in a particular absorbing state, then the probability of that happening decreases each time it makes a jump within the transient states. Getting absorbed in a particular state is more probable if mparse takes a faster, more direct path, rather than jumping around in the transient states for a while. Indeed, Valleriani et al. (2014) show via an analysis of the Taylor series expansion of the first-passage time density that the dominant terms correspond to the shortest paths from the initial state to the absorbing state.

In most of these equations, it is necessary to specify the initial probability distribution at each word $\mathbf{p}(0)$. As noted above, at the first word, all of the probability is placed on the no-structure state. This initial state is also used here for illustration purposes. For each subsequent word, though, the splitting probabilities from the previous word should be used as the initial conditions. In this way, the results of parsing at word $w - 1$ affect the parsing at word w . In the example, though, $\mathbf{p}(0)$ is set to have all probability on the no-structure state for simplicity.

4 Discussion and conclusion

Using the example sentence in (1), *the ranger saw the hunter with the binoculars*, I have summarized a set of first-passage time statistics for the continuous-time, discrete-state Markov model of sentence comprehension, mparse (Smith, 2021). These statistics can be used to generate predictions from mparse and compare them to human reading time data. The exit time and conditional first-passage time distributions at a word can be used as the predicted distribution of reading times at that word in the human experiment. These distributions can also form the basis of a likelihood function for fitting free parameters in the $\mathbf{W}$ matrix to human data. In Smith (2021), the entries of $\mathbf{W}$ are scaled by a free parameter τ that controls the time scale of the processing time predictions so that they are directly comparable to human reading times. Work is underway to fit individual-level τ parameters hierarchically to understand how individual differences in parsing speed relate to effect sizes in experimental sentences. Such model-based data analysis has led to great progress in other fields of cognitive science (e.g., in understanding forced-choice decision making; Ratcliff et al., 2016). The hope is that by making a process model a part of the data analysis pipeline, we can discover areas where the process model provides a good fit to the data (and therefore possibly a plausible explanation for them) and where the fit is poorer (indicating a need to revise the model; see also, Guest & Martin, 2021).

Having a tractable likelihood function also facilitates quantitative model comparison (Schütt et al., 2017), either using likelihood ratio tests or Bayes factors (which utilize the marginal likelihood, taking into account uncertainty in the parameters). To my knowledge, mparse is the first dynamical model of sentence process for which a likelihood function is available, which opens the door to much more informative model comparison efforts than have been possible with previous models (e.g., Cho et al., 2017; Kempen & Vosse, 1989; Smith et al., 2018, 2021; Stevenson, 1994; Tabor & Hutchins, 2004; Vosse & Kempen, 2000, 2009). These newly possible model comparison techniques will finally make testing the viability of the dynamical approach to sentence comprehension feasible.

An interesting future application of the conditional first-passage times and splitting probabilities relates to comprehension questions. In reading experiments, it is common to include comprehension questions that query participants’ understanding of the sentences and to try to understand how they interpret them. Obviously, the link between word-by-word processing and comprehension question performance is indirect; answering a question after reading a sentence relies on, among other things, participants’ fallible memory of the sentence. However, performance on comprehension questions is commonly understood to at least partially reflect what happened when reading the sentence (e.g., Christianson, Hollingworth, Halliwell, & Ferreira, 2001; Smith et al., 2021; Swets, Desmet, Clifton, & Ferreira, 2008; Villata, Tabor, & Franck, 2018). If, in a particular experimental item, building the correct parse was less likely than in another condition, i.e., the splitting probabilities differ, we might expect participants’ performance on comprehension questions to differ to a similar degree. Thus, the splitting probabilities at a critical region in a sentence could be used as predictions of comprehension question accuracy.

The conditional first-passage time distributions could be used in a similar way. Mparse predicts that, when there is more than one absorbing state, there can be different conditional first-passage times, as we saw with the example sentence with a bias toward the modifier parse. If we assume that splitting probabilities for a critical region correlate with comprehension question accuracy, then it is also reasonable to assume that reading times in that region should differ based on comprehension question accuracy. Thus, mparse can be used to generate new, testable predictions that directly relate comprehension question performance to online reading time behavior.

Finally, we anticipate that the mparse framework will scale up to broad-coverage parsing, that is, it should be possible to implement a version of mparse that can process arbitrary input sentences, as advocated for by Crocker and Brants (2000), for example. A common critique of dynamical models of sentence processing is that they are too complicated to handle more than a couple of sentence types with a single set of parameters (e.g., Bicknell & Levy, 2009). Work is underway to extract a large set of dependency rules from gold-standard parsed corpora (e.g., Nivre et al., 2016). Once the set of possible dependencies between words is set, it is possible to set up mparse to process any sentence that uses those rules (Smith, 2021), making large-scale tests of the robustness of mparse in comparison with other broad-coverage parsing models (e.g., Roark, Bachrach, Cardenas, & Pallier, 2009) possible.

Overall, the first-passage time techniques summarized here should provide a convenient reference to spur further model development and new experiments in incremental sentence comprehension. The mparse model was used to illustrate the techniques; however, they apply to any finite-state system governed by the master

equation, including the one-step processes typically handled in detail in textbooks (Gardiner, 1985; Redner, 2001; van Kampen, 2007). By adding new mathematical tools to the toolbox of computational cognitive modeling of language comprehension, we hope to make new and unexpected discoveries in a more principled and less ad hoc way.

Acknowledgments This research was funded by the University of Potsdam. Many thanks to Prof. Dr. Shravan Vasishth for comments on an earlier version of this paper.

References

- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in cognitive sciences*, 4(3), 91–99.
- beim Graben, P., Gerth, S., & Vasishth, S. (2008). Towards dynamical system models of language-related brain potentials. *Cognitive Neurodynamics*, 2, 229–255. doi: 10.1007/s11571-008-9041-5
- Bicknell, K., & Levy, R. (2009). A model of local coherence effects in human sentence processing as consequences of updates from bottom-up prior to posterior beliefs. In *Proceedings of the 10th Annual Meeting of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT) Conference* (pp. 665–673). Boulder, Colorado, USA.
- Cho, P. W., Goldrick, M., & Smolensky, P. (2017). Incremental parsing in a continuous dynamical system: Sentence processing in Gradient Symbolic Computation. *Linguistics Vanguard*, 3(1).
- Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42, 368–407.
- Crocker, M. W., & Brants, T. (2000). Wide-coverage probabilistic sentence processing. *Journal of Psycholinguistic Research*, 29(6), 647–669.
- Engbert, R., Longtin, A., & Kliegl, R. (2002). A dynamical model of saccade generation in reading based on spatially distributed lexical processing. *Vision Research*, 42, 621–636.
- Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: A dynamical model of saccade generation during reading. *Psychological Review*, 112(4), 777–813.
- Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies* (Unpublished doctoral dissertation). University of Connecticut.

- Gaifman, H. (1965). Dependency systems and phrase-structure systems. *Information and Control*, 8, 304–337.
- Gardiner, C. W. (1985). *Handbook of stochastic methods for physics chemistry and the natural sciences*. Springer.
- Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*. doi: 10.1177/1745691620970585
- Hays, D. G. (1964). Dependency theory: A formalism and some observations. *Language*, 40(4), 511–525.
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 111(2), 228–238.
- Kempen, G., & Vosse, T. (1989). Incremental syntactic tree formation in human sentence processing: A cognitive architecture based on activation decay and simulated annealing. *Connection Science*, 1(3), 273–290.
- Kübler, S., McDonald, R., & Nivre, J. (2009). *Dependency parsing*. Morgan & Claypool Publishers.
- Nivre, J., de Marneffe, M.-C., Ginter, F., Goldberg, Y., Hajič, J., Manning, C. D., ... Zeman, D. (2016). Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the tenth international conference on language resources and evaluation (lrec 2016)*.
- Oppenheim, I., Shuler, K. E., & Weiss, G. (1977). *Stochastic processes in chemical physics: The master equation*. MIT Press.
- Polizzi, N. F., Therien, M. J., & Beratan, D. N. (2016). Mean first-passage times in biology. *Israel Journal of Chemistry*, 56, 816–824. doi: 10.1002/ijch.201600040
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Topics in Cognitive Science*, 20(4), 260–281.
- Redner, S. (2001). *A guide to first-passage processes*. New York: Cambridge University Press.
- Roark, B., Bachrach, A., Cardenas, & Pallier, C. (2009). Deriving lexical and syntactic expectation-based measures for psycholinguistic modeling via incremental top-down parsing. In *Proceedings of the 2009 conference on empirical methods in natural language processing* (pp. 324–333).
- Schütt, H. H., Rothkegel, L. O. M., Truckenbrod, H. A., Reich, S., Wichmann, F. A., & Engbert, R. (2017). Likelihood-based parameter estimation and comparison of dynamical cognitive models. *Psychological Review*, 124(4),

505–524.

- Seelig, S. A., Rabe, M. M., Malem-Shinitzki, N., Risse, S., Reich, S., & Engbert, R. (2020). Bayesian parameter estimation for the SWIFT model of eye-movement control during reading. *Journal of Mathematical Psychology*, 95(102313).
- Smith, G. (2018). *A theory of timing effects in a self-organizing model of sentence processing* (Unpublished doctoral dissertation). University of Connecticut.
- Smith, G. (2021). *Mparse: A new framework for self-organized, incremental sentence comprehension*. Retrieved from <https://osf.io/k6rnx/> (Under review) doi: 10.5281/zenodo.4674276
- Smith, G., Franck, J., & Tabor, W. (2018). A self-organizing approach to subject-verb number agreement. *Cognitive Science*, 42(S4), 1043–1074. doi: 10.1111/cogs.12591
- Smith, G., Franck, J., & Tabor, W. (2021). Encoding interference effects support self-organized sentence processing. *Cognitive Psychology*, 124(101356).
- Stevenson, S. (1994). Competition and recency in a hybrid network model of syntactic disambiguation. *Journal of Psycholinguistic Research*, 23(4), 295–322.
- Swets, B., Desmet, T., Clifton, C. J., & Ferreira, F. (2008). Underspecification of syntactic ambiguities: Evidence from self-paced reading. *Memory & Cognition*, 36(1), 201–216. doi: 10.3758/MC.36.1.201
- Tabor, W., & Hutchins, S. (2004). Evidence for self-organized sentence processing: digging-in effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 431–450.
- Valleriani, A., Li, X., & Kolomeisky, A. B. (2014). Unveiling the hidden structure of complex stochastic biochemical networks. *The Journal of Chemical Physics*, 140(064101). doi: 10.1063/1.4863997
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615–665.
- van Kampen, N. G. (2007). *Stochastic processes in physics and chemistry*. Elsevier.
- Villata, S., Tabor, W., & Franck, J. (2018). Encoding and retrieval interference in sentence comprehension: Evidence from agreement. *Frontiers in Psychology*, 9(2), 1–16.
- Vosse, T., & Kempen, G. (2000). Syntactic structure assembly in human parsing: a computational model based on competitive inhibition and a lexicalist grammar. *Cognition*, 75, 105–143.
- Vosse, T., & Kempen, G. (2009). The Unification Space implemented as a localist neural net: predictions and error-tolerance in a constraint-based parser. *Cognitive Neurodynamics*, 3(4), 331–346.