Models of retrieval in sentence comprehension: A computational evaluation using Bayesian hierarchical modeling

CrossMark

Bruno Nicenboim *, Shravan Vasishth

Department of Linguistics, University of Potsdam, Potsdam, Germany

ABSTRACT

Research on similarity-based interference has provided extensive evidence that the formation of dependencies between non-adjacent words relies on a cue-based retrieval mechanism. There are two different models that can account for one of the main predictions of interference, i.e., a slowdown at a retrieval site, when several items share a feature associated with a retrieval cue: Lewis and Vasishth's (2005) activation-based model and McElree's (2000) direct-access model. Even though these two models have been used almost interchangeably, they are based on different assumptions and predict differences in the relationship between reading times and response accuracy. The activation-based model follows the assumptions of the ACT-R framework, and its retrieval process behaves as a lognormal race between accumulators of evidence with a single variance. Under this model, accuracy of the retrieval is determined by the winner of the race and retrieval time by its rate of accumulation. In contrast, the direct-access model assumes a model of memory where only the probability of retrieval can be affected, while the retrieval time is drawn from the same distribution; in this model, differences in latencies are a by-product of the possibility of backtracking and repairing incorrect retrievals. We implemented both models in a Bayesian hierarchical framework in order to evaluate them and compare them. The data show that correct retrievals take longer than incorrect ones, and this pattern is better fit under the direct-access model than under the activation-based model. This finding does not rule out the possibility that retrieval may be behaving as a race model with assumptions that follow less closely the ones from the ACT-R framework. By introducing a modification of the activation model, i.e., by assuming that the accumulation of evidence for retrieval of incorrect items is not only slower but noisier (i.e., different variances for the correct and incorrect items), the model can provide a fit as good as the one of the direct-access model. This first ever computational evaluation of alternative accounts of retrieval processes in sentence processing opens the way for a broader investigation of theories of dependency completion.

© 2017 Elsevier Inc. All rights reserved.

There is strong evidence that the formation of syntactic dependencies between non-adjacent words relies on the memory system. An example is the so-called locality effect: increasing the distance between co-dependents (such as subjects and verbs) tends to lead to greater processing difficulty (Gibson, 2000; Grodner & Gibson, 2005). Research on interference makes a similar point: the speed and/or accuracy of dependency completion can be adversely affected by the presence of items in memory that are similar to the retrieval target (among others: Gordon, Hendrick, & Levine, 2002; Jäger, Engelmann, & Vasishth, 2015; Nicenboim, Vasishth, Engelmann, & Suckow, 2017; Van Dyke, 2007; Van Dyke & Lewis, 2003; Vasishth, Brüssow, Lewis, & Drenhaus, 2008). Such a central role for memory in sentence comprehension

is well-motivated: it is implausible that the parser could keep track of a large and in principle unbounded inventory of the dependencies that can be found in a sentence, since they easily exceed the amount of information that can be held in the focus of attention (Cowan, 1995; Marcus, 2013; McElree, 2006; McElree & Dosher, 1989; Oberauer, 2013). The evidence from studies investigating similarity-based interference (see the meta-analysis of published reading studies in Jäger, Engelmann, & Vasishth, 2017) suggests that dependency completion relies on a *content-addressable cue-based retrieval mechanism* that is subject to interference (Lewis, Vasishth, & Van Dyke, 2006; McElree, 2000; Van Dyke & Lewis, 2003). Similarity-based interference is a phenomenon that is not unique to language, and occurs when several items share a feature associated with a retrieval cue. A major implication is that the retrieval mechanism employed for the creation of linguistic dependencies is similar to the one utilized in non-language domains.

* Corresponding author.
   E-mail address: bruno.nicenboim@uni-potsdam.de (B. Nicenboim).

There are multiple implementations compatible with such a content-addressable cue-based retrieval mechanism in sentence processing. As a verbally stated model, this type of mechanism would entail that when retrieval cues fully match the target of retrieval, similarity-based interference would cause an inhibitory effect, that is, an increase of processing difficulty at the retrieval of a dependent. This processing difficulty would be reflected in longer reading times and lower accuracy. However, in some cases, shorter reading times have been observed when increased processing difficulty was clearly expected (Nicenboim, Vasishth, Gattei, Sigman, & Kliegl, 2015; Nicenboim, Logačev, Gattei, & Vasishth, 2016; Van Dyke & McElree, 2006). In these cases, it is usually assumed that the fast reading times are a consequence of a shallow parse (due to, for example, good-enough processing, Ferreira, Bailey, & Ferraro, 2002) caused by cognitive overload. There can be good reasons to assume that shorter reading times are associated with increased difficulty, for example, when shorter reading times co-occur with lower comprehension accuracy (Van Dyke & McElree, 2006) or lower working memory capacity (Nicenboim et al., 2015; Nicenboim et al., 2016). However, the trade-off between reading times and comprehension accuracy is usually left underspecified.

There are two models that make explicit the relationship between reading times and retrieval accuracy, and even though they are sometimes not differentiated, they constitute two different implementations of the content-addressable cue-based retrieval mechanism. These model are the Lewis and Vasishth's (2005) *activation-based* model, and McElree's (2000) *direct-access* model. These models have different implications for retrieval processes in sentence comprehension. The activation-based model assumes a process that resembles a race model (Audley & Pike, 1965; Vickers, 1970), where evidence for each retrieval candidate is accumulated with different rates. This race determines both the latencies and the retrieval accuracy. By contrast, the direct-access model assumes that retrieval candidates have different levels of *availability*, which is the probability that a memory representation is retained. Availability determines only the accuracy of the retrieval and not the latency. In this model, a difference in latency between two conditions is a by-product of a mixture of directly accessed items, and retrievals that were initially incorrect, but they are reanalyzed leading to a correct retrieval.

The goal of this paper is to unpack the quantitative predictions of the activation-based and direct-access models by implementing them in a Bayesian hierarchical framework. This will allow us to compare their relative fit to a representative dataset and to assess their validity as models of retrieval that can account for similarity-based interference. We used a subset of the data from Nicenboim et al. (2017), which investigated similarity-based interference from the number feature using two relatively large-sample self-paced reading experiments. The data in this study include two dependent measures: (i) reading times for the critical region where retrieval from memory is assumed to occur, and (ii) accuracies in a comprehension task that targets specific dependency relations through a multiple choice task. This dataset is especially suitable for our modeling purposes because, apart from Van Dyke (2007), who also evaluated some of the dependencies, this is the only dataset that we are aware of that uses comprehension questions to directly assess the resolution of the dependencies. As explained in detail later, these two dependent measures (reading times and accuracy) are necessary for evaluating the models. We begin by describing the Nicenboim et al. (2017) study.

The Nicenboim et al. (2017) study

Nicenboim et al. (2017) used stimuli like (1). There were two conditions, high vs. low interference, which were assumed to affect the dependency between the subject (i.e., *Der Wohltäter* "The philanthropist") and the verb (i.e., *begrüßt hatte* "had greeted"). In the high interference condition, two nouns intervened between these two co-dependents that had the same number marking as the target noun, the subject of the sentence, namely, singular marking. In the low interference case, the two intervening nouns had plural marking while the target noun remained singular. In German, the verb (i.e., *begrüßt hatte*) agrees in number with its subject; in the high interference condition, the retrieval cue set at the verb to seek out a singular noun would match three nouns. By contrast, in the low interference condition, only one noun matches this retrieval cue. Thus, reading time at the critical region, the verb *begrüßt hatte*, provides an estimate of any interference effect.

Each target sentence was followed by a question that queried either the subject of the matrix verb (e.g., "sat"), the subject of the embedded verb (e.g.,"had greeted"), or the object of the embedded verb. The possible answers were provided in multiple-choice format in pseudo-randomized order. For all the questions, participants had the option to answer "I don't know", when they did not remember or could not answer.

(1)	a.	High Interference						
		Der **The.sg.nom**	**Wohltäter,** **philanthropist,**	der who.sg.nom	den the.sg.acc	Assistenten assistant	(of)	
		des the.sg.gen Spendenausschuss. donations committee.	Direktors director	**begrüßt** **greeted,**	**hatte,** **had.sg**	saß sat.sg	später later	im in the

'The philanthropist, who had greeted the assistant of the director, sat later in the donations committee.'

	b.	Low Interference						
		Der **The.sg.nom**	**Wohltäter,** **philanthropist,**	der who.sg.nom	die the.pl.acc	Assistenten assistants	(of)	
		der the.pl.gen Spendenausschuss. donations committee.	Direktoren directors	**begrüßt** **greeted**	**hatte,** **had.sg,**	saß sat.sg	später later	im in the

'The philanthropist, who had greeted the assistants of the directors, sat later in the donations committee.'