Behavioral interference or facilitation does not distinguish between competitive and noncompetitive accounts of lexical selection in word production

Gary M. Oppenheim (g.m.oppenheim@bangor.ac.uk)
School of Psychology, Bangor University
Bangor, Gwynedd LL57 2AS, UK

Nazbanou Nozari (bnozari@andrew.cmu.edu)
Department of Psychology, Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213

Abstract

One of the major debates in the field of word production is whether lexical selection is competitive or not. For nearly half a century, semantic interference effects in picture naming latencies have been claimed as evidence for competitive (relative threshold) models of lexical selection, while semantic facilitation effects have been claimed as evidence for non-competitive (simple threshold) models instead. In this paper, we use a computational modeling approach to compare the consequences of competitive and noncompetitive selection algorithms for blocked cyclic picture naming latencies, combined with two approaches to representing taxonomic and thematic semantic features. We show that although our simple model can capture both semantic interference and facilitation, the presence or absence of competition in the selection mechanism is unrelated to the polarity of these semantic effects. These results question the validity of prior assumptions and offer new perspectives on the origins of interference and facilitation in language production.

Keywords: word production; computational modeling; thematic representation; taxonomic representation; semantic interference

Introduction

When choosing a word for production, do speakers consider each candidate's activation independently—selecting the first to exceed a simple threshold—or do they compare each candidate's activation to that of the alternatives? This question of competitive lexical selection has been a focus of considerable debate for the past two decades. The empirical fuel for the debate has focused on semantic interference effects in picture naming experiments, especially those using picture-word interference and blocked cyclic naming.

In the picture-word interference task, participants attempt to name pictures while ignoring visually or auditorily presented distractor words. Participants are typically slower to correctly name pictures when the accompanying distractors are semantically similar to the targets than when they are unrelated (see, e.g., Bürki, Elbuy, Madec, & Vasishth, 2020, for a review), a result that is typically claimed as evidence for a competitive lexical selection process. The most common interpretation is that perceiving a distractor directly activates its lexical representation in the production system, making it stronger competitor, and participants need to work harder to select a target name when an alternative has support from both the distractor process and the production

process itself (e.g., Levelt, Roelofs, & Meyer, 1999). However, that competitive interpretation is complicated by another robust observation, that the precise nature of the semantic relationship between the target and the distractor also matters: while distractors from the same taxonomic category as a target (e.g., dog/cow) will slow target naming, distractors that are semantically associated with a target (e.g. milk/cow) can instead produce facilitation (Alario et al., 2000; Costa et al., 2005; Mahon et al., 2007). It is not immediately obvious how a competitive system could account for such a reversal in the polarity of the effect—according to the basic competitive selection hypothesis, anything that supports an alternative response should thereby hinder target selection—so that facilitation has been interpreted as support for the noncompetitive claim.

Blocked cyclic naming elicits a similar pattern of results. This task does not involve 'distractors'; instead, participants simply name pictures one at a time. The paradigm draws its name from a sequencing manipulation: researchers present the pictures 'cycling' repeatedly through small sets of items in 'blocks' where the items are either semantically related or unrelated. Participants and are slower to correctly name items when they occur in blocks from a single taxonomic category (e.g., mammals) than when they represent different categories, a difference that increases over successive cycles. Like taxonomic picture-word interference, this taxonomic interference effect is typically described as evidence for competitive lexical selection (see Nozari & Pinet, 2020, for a review): priming taxonomically related words makes them more effective competitors, thus slowing target word selection. And although there have been some reports of interference from thematically associated blocks (e.g., Abdel Rahman & Melinger, 2007), thematic interference effects have invariably been weaker than the taxonomic interference effects in the same experiments, and sometimes altogether absent (e.g., de Zubicaray, Johnson, Howard, & McMahon, 2014). In other words, while the evidence for interference induced by taxonomic relations is robust, the consequences of thematic relations are less clear-cut, partly because many studies that test thematic relations have not fully disentangled them from taxonomic relations.

To disentangle these factors, McDonagh, Fisher, & Nozari (2020) recently used a simplified version of the blocked cyclic naming paradigm, including just two pictures per

block, controlled for quantitative measures of both taxonomic similarity (Resnik scores) and thematic similarity (association scores). They paired the same target pictures once with taxonomically related pictures (high Resnik and low association scores), once with thematically related pictures (high association and low Resnik scores), and once with pictures that were unrelated to the target in either dimension (low scores on both measures). This double dissociation, along with controlling for segmental overlap, which can itself induce interference, allowed for disentangling the two effects: compared to the baseline, taxonomic similarity induced clear interference, while thematic similarity produced marginal facilitation that also significantly differed from the taxonomic interference.

To summarize, the behavioral evidence is clear: taxonomic relations reliably induce interference, while thematic relations often do not, and may even facilitate production. The question is: Can these seemingly opposing effects be captured using a single mechanism, and if so, is that mechanism competitive or non-competitive selection?

In this paper, we use a computational modeling approach to critically evaluate two core assumptions underlying the use of both paradigms to assess the presence or absence of lexical competition: (1) behavioral interference is evidence for competitive lexical selection, and (2) behavioral facilitation is evidence against competitive lexical selection. The first of these assumptions has already received some scrutiny. Regarding taxonomic picture-word interference, researchers have proposed alternative explanations that do not require competition at the level of lexical selection. The Response Exclusion Hypothesis (e.g., Mahon et al., 2007), for instance, posits that semantic picture-word interference results instead reflect late, post-selection monitoring processes that participants engage in response to the peculiarities of the picture-word interference task. A participant might initially choose a response quite easily but covertly re-evaluate it before actually speaking it aloud. Such alternatives abound because the complexity of the task hinders associating even robust empirical results with specific properties of the production system (Oppenheim & Balatsou, 2019).

Regarding interference in the blocked cyclic naming paradigm, Oppenheim, Dell, & Schwartz, (2010) have previously shown that such taxonomic interference naturally emerges from domain-general incremental learning processes, regardless of whether they are paired with competitive or noncompetitive algorithms for lexical selection. While the learning account is compatible with competitive selection models, it renders any possible competition in the selection algorithm superfluous.

The question that we now turn to is whether thematic facilitation in such paradigms implies the absence of competition. Under the competitive selection account, any manipulation that increases the activation of an alternative response is generally expected to delay target selection and

thereby increase naming latencies, so the account struggles to explain any systematic exception to that pattern (e.g., Levelt et al., 1999). Associative facilitation in picture-word interference might therefore require ad hoc assumptions that semantic associates do not compete with targets for selection, or be explained away as caused by other, non-selection mechanisms. But because thematic interference in blocked cyclic naming has repeatedly been claimed as evidence for competitive lexical selection (e.g., Abdel Rahman & Melinger, 2007), its absence in better controlled studies (e.g., McDonagh et al., 2020) is more problematic for the account.

In the following simulations, we offer a novel account of the range of observed thematic effects in blocked cyclic naming as emerging from variation in the context-dependence of thematic representations. Simulation 1 shows that, when we represent thematic features in precisely the same way as taxonomic features, contributing to lexical retrieval regardless of context, they produce cumulative semantic interference. Simulation 2 shows that, when we instead assume that thematic features contribute more to lexical activation in supporting contexts—for instance, when they are shared by successive items—thematic relationships can instead yield facilitation. In each simulation, we simulate response times using both competitive and noncompetitive algorithms for lexical selection, and in each simulation the contrasting algorithms yield equivalent results.

Model Description

We assess the consequences of competitive lexical selection and approaches to taxonomic and thematic representation within the framework of Oppenheim, Dell, & Schwartz's (2010) Dark Side model of incremental learning in semantically driven word production, implemented in R/C++. This simple model maps shared semantic features onto words via a single layer of connections, adjusting the connections after each retrieval to reduce activation error. It was originally proposed as an explanation of taxonomic cumulative semantic interference in tasks including blocked cyclic picture naming, allowing us to straightforwardly compare selection rules under two approaches to representing thematic associations.

Architecture

As illustrated in Figure 1, the model consists of an array of semantic nodes, described[1] as semantic features, directly mapped to an array of word nodes via a single layer of learnable connections, initialized at zero.

Representations. Following Oppenheim et al. (2010), this model represents each word as the unique intersection of two semantic features. For instance, within its small vocabulary, the output word 'dog' is the intersection of [mammal] and [terrestrial] input features, while the output word 'car' is the

---
[1] Though the term 'feature' can imply a localist approach to semantic representation, the model is similarly compatible with distributed approaches.
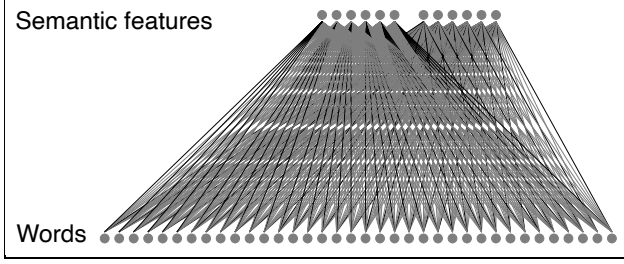
Figure 1: Architecture of the base model as implemented in Simulation 1. All connection weights are learned from experience: excitatory in black, inhibitory in grey.

intersection of the [vehicular] and [terrestrial] input features. A simple extension of this approach could represent thematic features like [birthday party] in exactly the same way.

Algorithms

Following Oppenheim et al. (2010), each trial consists of three phases: activation, selection, and learning.

Activation. Each trial begins with the activation of two semantic features, setting their activations to 1 and the activations of all other nodes to zero. From these, we calculate the initial activation of each output word as the sigmoidal transformation of its net input, including a small amount of normally distributed noise. Thus, the net input for each word output node i is calculated in Eq. 1 as the sum of the activation of each input node, a_j, times the weight of its connection to node i, w_{ij}:

$$\text{net}_i = \sum_j w_{ij} a_j \qquad \text{Eq. 1}$$

And the activation for each word output node i is given in Eq. 2 as a sigmoidal function of the net input, plus a small amount of normally distributed noise (with a mean of 0 and a variance of v):

$$a_i = \frac{1}{1 + e^{-(net_i + N(0, v))}} \qquad \text{Eq. 2}$$

Selection. An active selection process plays out over time, linking each word's activation to its likelihood and timing of selection. We assume that this selection emerges from an evidence accumulation process, which 'boosts' task-relevant activations toward a selection threshold, τ, at a rate, β. The functional nature of this selection process remains an active topic of debate. While researchers universally assume that the process imposes a winner-take-all function that links target activation to the time of target selection, accounts differ on the possible role of non-target activations in determining target selection times. To address the range of theoretical perspectives, we report selection times according to two formulae.

First, a non-competitive formula (Eq. 3, following Eq. 11 from Oppenheim et al, 2010) assumes that the time of word

selection is a simple threshold function of the activation of the most strongly activated word.

$$t_{selection} = \log_\beta \left(\frac{\tau}{a_{i\,t_1}} \right) \qquad \text{Eq. 3}$$

Second, a competitive formula (Eq. 4, following Eq. 12 from Oppenheim et al., 2010; cf. Krajbich & Rangel, 2011; Nozari & Hepner, 2019) assumes that the time of word selection is a *relative* threshold function of the activation of the most strongly activated word compared to the next most strongly activated word.

$$t_{selection} = \log_\beta \left(\frac{\tau}{a_{i\,t_1} - a_{strongest\,competitor\,t_1}} \right) \qquad \text{Eq. 4}$$

We have selected these algorithms for the sake of expository clarity, with the assumption that other plausible algorithms will yield similar predictions when considered at the same grain.

Learning. At the end of the trial, the weight of each connection is adjusted according to the Widrow-Hoff (1960) Delta Rule (Eq. 5), implementing a supervised corrective learning algorithm:

$$\Delta w_{ij} = \eta(a_i(1 - a_i)(d_i - a_i))a_j \qquad \text{Eq. 5}$$

Parameters. Model parameters are given in Table 1. These are identical to the parameters used in Oppenheim et al. (2010), with one exception: because we are only interested in response times in these simulations, we reduced the activation noise parameter from 0.5 to 0.1.

Simulation 1

In Simulation 1, we show that the model can capture the basic taxonomic cumulative semantic interference effect in McDonagh et al.'s (2020) modified blocked cyclic naming paradigm, and test whether competitive selection is necessary for demonstrating such interference. This demonstration and test naturally extend to thematic interference if we assume that speakers similarly use thematic 'features' to support word retrieval.

Methods

The model was trained as in Oppenheim et al.'s (2010) Simulation 2, with identical parameters as described above. The only major difference is that, following McDonagh et al.'s (2020) method, the testing phase implemented 6 two-item cycles instead of four 6-item cycles.

Table 1: Model parameters.

Parameter	Value
Learning rate (η)	0.75
Activation noise (ν)	0.1
Boosting rate (β)	1.01
Threshold (τ)	1
Deadline (Ω)	100
Scaling constant for thematic input (λ; Sim 2 only)	0.25

Vocabulary. As in Oppenheim et al.'s (2010) Simulation 2, the model was trained on a vocabulary that mapped 12 semantic features onto 36 words. Each feature was shared by exactly six words, and each word represented the unique intersection of two features.

Pre-training. All connection weights were initialized at 0. Prior to simulating the experiment, we trained the model with 100 randomly ordered passes through its entire vocabulary, updating connection weights after each trial.

Experiment simulation. To approximate McDonagh et al.'s (2020) task, we ran the model separately on 1024 blocks in each of two conditions: (1) a *Baseline* condition, where the targets shared neither taxonomic nor thematic features, and (2) a *Semantically related* condition, where the targets shared one of their two input features; because this simulation assumes equivalency between taxonomic and thematic features, the shared feature can be assumed to represent either. As in the empirical study, each block consisted of 6 randomly ordered 'cycles' through the two items.

Results

The simulation produced the expected cumulative semantic interference effect: slower lexical selections in the semantically related condition than in the semantically unrelated condition, corresponding to the longer RTs that are typically observed in human data. Figure 2 depicts the mean per-cycle simulated selection times for 1024 blocks in each condition, using both the competitive (Figure 2a) and noncompetitive (Figure 2b) selection rules. Notably, the model produces the same cumulative semantic interference effect, increasing across cycles, when employing the competitive selection rule as when employing the noncompetitive one.

Discussion

Simulation 1 demonstrates that our model naturally accounts for cumulative semantic interference in this simplified blocked cyclic naming paradigm. As in Oppenheim et al. (2010), the crucial mechanism underlying this interference is incremental learning: strengthening and weakening connections to reduce error leads the model to emphasize its most recent experience even if doing so requires partially forgetting previous knowledge (cf. McCloskey & Cohen, 1989).

Although the semantic features in our base model may most obviously approximate taxonomic features from taxonomically related items (e.g., where 'cow' and 'bear' share a 'mammal' feature), note that this simulation also addresses the simplest case for thematic associations, where thematic features are represented and used in a way that is indistinguishable from taxonomic features (e.g., 'cow' and 'milk' share a 'dairy farm context' feature). That is, if there is no real distinction between thematic features and taxonomic features in general, then extending the scope of a taxonomic feature-based model to include thematic features is as simple as re-labeling the semantic features as 'thematic' and leads to the same interference for both taxonomic and thematic relations. This demonstration is important because thematic associations have sometimes been claimed to elicit interference in the blocked cyclic naming paradigm (e.g., Abdel Rahman & Melinger, 2007). Crucially, however, a model with a non-competitive selection rule is just as successful at capturing such behavioral interference as one with a competitive selection rule. In short, regardless of whether it is attributed to taxonomic or thematic relations, behavioral interference cannot distinguish between competitive and non-competitive lexical selection.

Simulation 2

Simulation 2 extends the findings of Simulation 1, by testing whether a non-competitive selection mechanism is necessary for explaining behavioral facilitation. In this simulation, we consider the possibility that thematic features may crucially
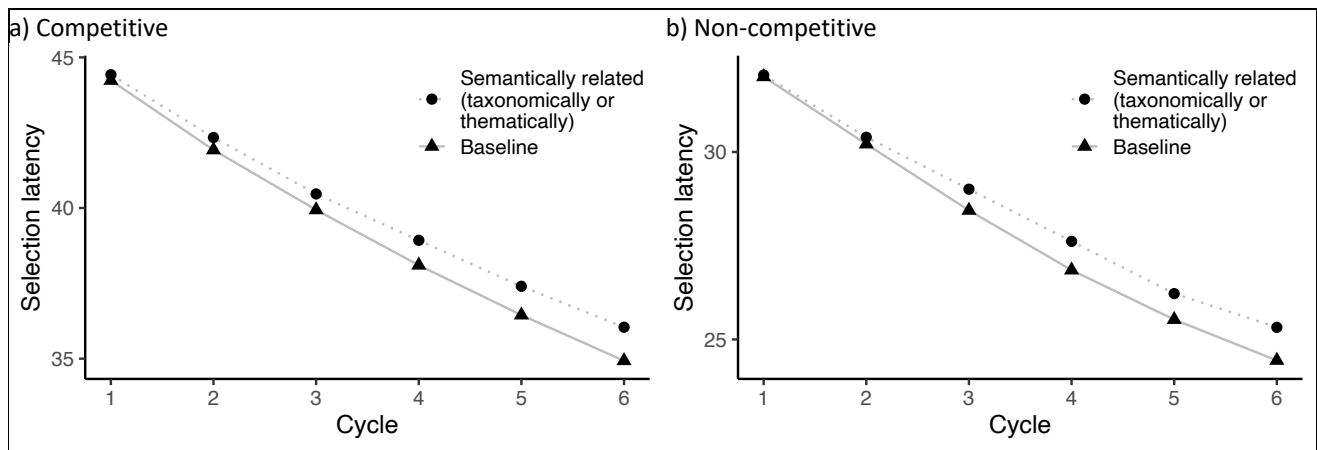

Figure 2: Results of Simulation 1. Selection latencies in Panel (a) are calculated according to a strongly competitive selection rule, while latencies in Panel (b) are calculated according to a non-competitive selection rule

differ from taxonomic features in their context-dependence. While taxonomic features tend to be central to the representation of a word/concept, reliably activated regardless of the context, we assume that thematic features are only weakly associated, making their wider contextual relevance more critical to their activation (e.g., Barsalou, 1982). As one way to implement this distinction, in Simulation 2 we represent and train the contribution of thematic associations as weak auxiliary semantic features that require supporting contexts: each word is sufficiently specified by its taxonomic features, as in Simulation 1, but we assume that a 'theme' node provides an additional retrieval cue when it has already been activated by the preceding word. For instance, naming 'cow' and 'tractor' together may activate a 'farm' theme that then provides an additional retrieval cue, but because that theme is not part of the core representation of either concept, it provides no support when encountering 'cow' in the context of farm-irrelevant words like 'bear' or 'stapler'.

Methods

Methods followed Simulation 1, except as specified below.

Vocabulary As illustrated in Figure 3, we expanded the 12-feature x 36-word vocabulary from Simulation 1 to include an orthogonal set of six new 'theme' input features, each associated equally often with each of the 12 taxonomic features.

Thematic input feature activation. As discussed above, we assume that features crucially differ from taxonomic features in that they contribute more to word retrieval in supporting contexts. We implement this assumption for Simulation 2 by applying Eq. 6 in both pre-training and testing to calculate the contribution of each thematic input feature to the net input for the lexical layer. Each thematic input feature, a_j, contributes a scaled activation value, a'_j, if it is active in the current trial
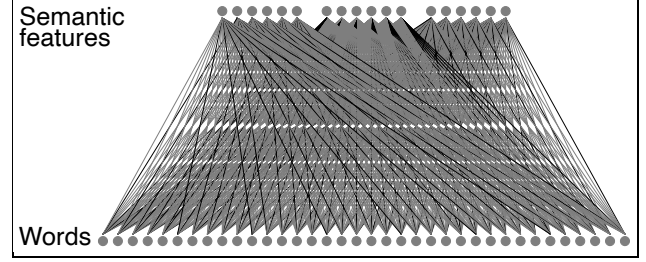

Figure 3: Architecture of the extended model as implemented in Simulation 2.

and was active in the previous trial, else its contribution reduces to 0.

$$a'_j = \lambda a_{j,current} a_{j,previous} \qquad \text{Eq. 6}$$

Although this implementation is admittedly ad hoc, Eq. 6 assumes that contextual support integrates activation patterns across multiple trials, and crucially that thematic features contribute more to word retrieval in supporting contexts.

Experiment simulation. The simulation followed the methods of Simulation 1, except that it now included 1024 blocks of each of three conditions: (1) a *Baseline* condition, where the targets shared neither taxonomic nor thematic features, (2) a *Taxonomically related* condition, where the targets shared one taxonomic feature but no thematic features, and (3) a *Thematically related* condition, where the targets shared one thematic feature but no taxonomic features.

Results

Mean selection times for Simulation 2 are depicted in Figure 4. If we assume that thematic features only contribute to word retrieval in supporting contexts, this model again generates the expected cumulative semantic interference for taxonomic features, but it now predicts net facilitation instead of
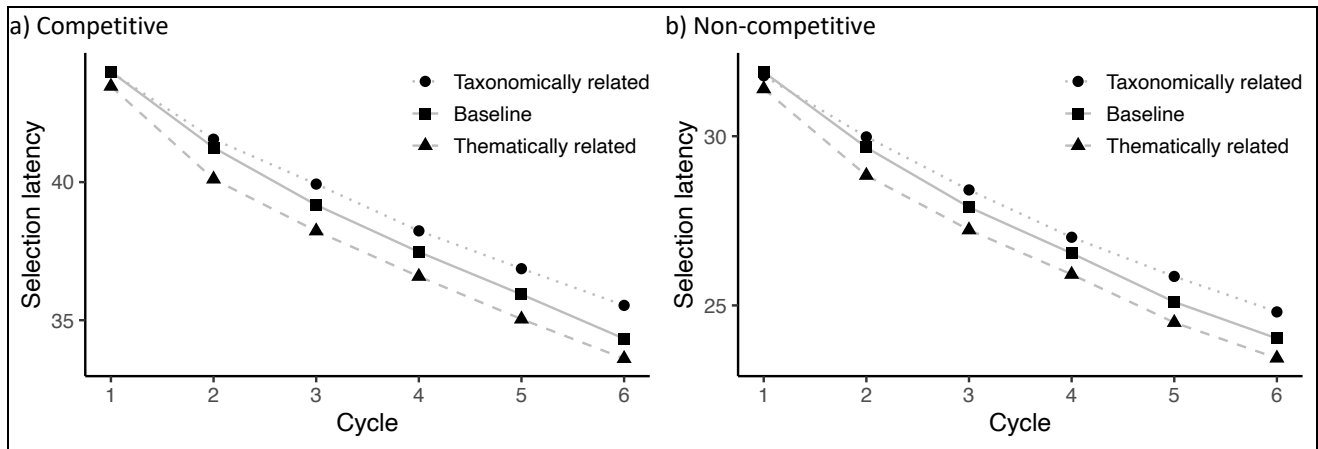

Figure 4: In Simulation 2, semantic blocking with a thematic association creates facilitation rather than interference. Selection latencies in Panel (a) are calculated according to a strongly competitive selection rule, while latencies in Panel (b) are calculated according to a non-competitive selection rule.

interference for thematically related blocks. Importantly, the model yields the same pattern of results whether selection is competitive (Fig. 4a) or non-competitive (Fig. 4b).

Discussion

Simulation 2 showed that taxonomically induced interference and thematically induced facilitation can be simultaneously captured within the same model. Is there something special about thematic associations that turns interference into facilitation? Not quite. The model applies the same incremental learning processes in thematically related blocks that it applies in the taxonomically related blocks, so both should have the same basis for *cumulative* semantic interference. The core reason for the thematic facilitation in Simulation 2 is the assumption that the contribution from thematic features is context-dependent: activating an extra semantic feature in supporting contexts means that the target word has more input in thematically related blocks because it receives activation from three features instead of two. It therefore predicts *net* thematic facilitation as a kind of fan effect (from the memory literature) or a semantic richness effect (e.g., Rabovsky, Schad, & Abdel Rahman, 2016).

Although behavioral studies tend not to report such robust thematic facilitation in blocked cyclic naming, this simulation captures two important features of their data. First, thematic associations *do* reliably produce less semantic interference than taxonomic associations, often resulting in null effects. Second, thematic associations *can* elicit faster responses than unrelated baseline conditions. While weaker interference might simply be attributed to weaker associations, the possibility of facilitation demands an opposing force, and the context-dependent activation of thematic features in Simulation 2 provides it. Moreover, by capturing both interference and facilitation, the model demonstrates a basis for interpolating between its states to address intermediate outcomes such as weak or null effects.

Importantly, although the polarity of the thematic effect has reversed from Simulation 1, the competitive and noncompetitive selection rules again yield equivalent results. In short, neither interference, nor facilitation, could distinguish between competitive and non-competitive mechanisms of lexical selection.

General Discussion

Our simulations have examined the importance of (non)competitive lexical selection in extending taxonomic feature-based word production models to incorporate thematic associations. Two major results emerged. The first is that simple feature-based representations of thematic associations can capture the full gamut of theme-based effects in single-word production. Representing thematic features as equivalent to taxonomic features can produce the kind of theme-based interference that has sometimes been reported in blocked cyclic naming experiments (Abdel Rahman & Melinger, 2007). Representing thematic features as context-dependent auxiliary cues can instead produce theme-based facilitation, as we have seen in our own blocked cyclic naming experiments and has often been shown in experiments utilizing picture-word interference. And by extension, the approach can account for intermediate results by assuming stochastic variation in the way that individuals represent and access specific features.

The second important result is that, within each simulation, the competitive and noncompetitive selection rules produced equivalent results. This finding clearly demonstrates that the assumption of behavioral interference and facilitation mapping directly on to competitive and non-competitive mechanisms of lexical selection is false. The presence or absence of competition in the selection mechanism is neither necessary nor sufficient to explain the existence or polarity of these effects.

Collectively, these results show that behavioral interference and facilitation as a function of semantic similarity have more to do with how semantic knowledge is mapped on to the lexical items, as opposed to the details of lexical selection. Our second simulation is a step in formalizing this view: while classical accounts of knowledge representations assumed objects and concepts to have stable representations, contemporary behavioral and neuroimaging evidence has shifted this view in favor of context-dependent representations (see Yee & Thompson-Schill, 2016, for a review). Naturally, the features most dependent on context are thematic features, i.e., those that determine the association between an item and other items under specific themes. In contrast, taxonomic features are relatively stable across contexts. Simulation 2 in this paper implements this logic and demonstrates that it is sufficient to explain the patterns of behavior observed when speakers name pictures with different kinds of semantic relations. This general approach can easily be extended to test how emphasizing certain features or probabilistically manipulating a set of features can shift the balance between interference and facilitation, which we recommend as a future direction.

As for whether lexical selection is competitive or not, we have provided unambiguous evidence that the question cannot be addressed by appealing to behavioral interference or facilitation. There is of course a wide range of possible algorithms for competitive selection, including the possibility that the system could shift between competitive and noncompetitive states in response to task goals (Nozari & Hepner, 2019). We have aimed to implement and assess *representative* forms of the algorithms, because our goal here has not been to identify the precise form of the selection process. Instead, our main point is that the empirical evidence does not warrant a strong commitment to a selection process that is definitely competitive or non-competitive, so although the process of computational modelling may require a researcher to tentatively assume *some* specific form, language production theories that strongly depend on such specific properties of a selection algorithm are built on shaky ground.

Acknowledgments

References

Abdel Rahman, R., & Melinger, A. (2007). When bees hamper the production of honey: lexical interference from associates in speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(3), 604–614.

Alario, F.-X., Segui, J., & Ferrand, L. (2000). Semantic and associative priming in picture naming. *The Quarterly Journal of Experimental Psychology*, *53A*(3), 741–764.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, *10*(1), 82–93.

Bürki, A., Elbuy, S., Madec, S., & Vasishth, S. (2020). What did we learn from forty years of research on semantic interference? A Bayesian meta-analysis. *Journal of Memory and Language*, *114*(April), 104125.

Costa, A., Alario, F.-X., & Caramazza, A. (2005). On the categorical nature of the semantic interference effect in the picture-word interference paradigm. *Psychonomic Bulletin & Review*, *12*(1), 125–131.

de Zubicaray, G. I., Johnson, K., Howard, D., & McMahon, K. (2014). A perfusion fMRI investigation of thematic and categorical context effects in the spoken production of object names. *Cortex*, *54*(1), 135–149.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(33), 13852–13857.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*(01), 1–38.

Mahon, B. Z., Costa, A., Peterson, R., Vargas, K. A. A., & Caramazza, A. (2007). Lexical selection is not by competition: A reinterpretation of semantic interference and facilitation effects in the picture-word interference paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(3), 503–535.

McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, *24*(C), 109–165.

McDonagh, D. C., Fisher, A. V, & Nozari, N. (2020). Do Taxonomic and Associative Relations Affect Word Production in the Same Way? *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, May.

Nozari, N., & Hepner, C. R. (2019). To select or to wait? The importance of criterion setting in debates of competitive lexical selection. *Cognitive Neuropsychology*, *36*(5–6), 193–207.

Nozari, N., & Pinet, S. (2020). A critical review of the behavioral, neuroimaging, and electrophysiological studies of co-activation of representations during word production. *Journal of Neurolinguistics*, *53*(April 2019), 100875.

Oppenheim, G. M., & Balatsou, E. (2019). Lexical competition on demand. *Cognitive Neuropsychology*, *3294*(5–6), 216–219.

Oppenheim, G. M., Dell, G. S., & Schwartz, M. F. (2010). The dark side of incremental learning: a model of cumulative semantic interference during lexical access in speech production. *Cognition*, *114*(2), 227–252.

Rabovsky, M., Schad, D. J., & Abdel Rahman, R. (2016). Language production is facilitated by semantic richness but inhibited by semantic density: Evidence from picture naming. *Cognition*, *146*, 240–244.

Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits*.

Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin and Review*, *23*(4), 1015–1027.