UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Modeling structural priming in sentence production via analogical processes

Permalink

Journal

ISSN

Authors
Taylor, Jason
Friedman, Scott
Goldwater, Micah
et al.

Publication Date
2011

Peer reviewed

Modeling structural priming in sentence production via analogical processes

Jason L. M. Taylor (jason.taylor@u.northwestern.edu)
Scott E. Friedman (friedman@northwestern.edu)
Kenneth Forbus (forbus@northwestern.edu)
Qualitative Reasoning Group, Northwestern University, 2133 Sheridan Road
Evanston, IL 60208 USA

Micah Goldwater (micahbg@gmail.com)
Dedre Gentner (gentner@northwestern.edu)
Department of Psychology, Northwestern University, 2029 Sheridan Road
Evanston, IL 60208 USA

Abstract

Recently there has been a surge of interest in using structural priming to examine sentence production. We present an analogical model of sentence production that exhibits structural priming effects. It uses analogical generalization to acquire abstract language patterns from experience. To construct utterances, it uses analogical retrieval to find semantically similar utterances and generalizations, and constructs a new sentence by analogy to them. Using the stimulus generator of Chang et al (2006), we show that this model can exhibit structural priming effects similar to those observed in humans, but with orders of magnitude less prior experience than required by a previous simulation.

Keywords: structural priming; sentence production; syntax acquisition; analogy.

Introduction

What mechanisms underlie sentence production? In particular, how do speakers choose among the multiple grammatical forms that are capable of expressing something they intend to convey? Recently, there has been a surge of interest in *structural priming* as a way to examine sentence production processes in adults and children (Bock, 1986; Bock & Griffin, 2000; Chang, Dell, & Bock, 2006; Kaschak & Borreggine, 2008; Savage et al., 2003). In structural priming, the structure of one sentence is repeated in the structure of a second sentence (Bock, 1986). Structural priming occurs without any intention to create syntactic parallelism. It does not require semantic or thematic overlap between the utterances, although the effects can be stronger when lexical items are repeated, and sentences are semantically similar (Branigan, Pickering, & Cleland, 2000; Goldwater et al., 2011; Hare & Goldberg, 1999; Pickering & Garrod, 2004; Snider, 2009).

To illustrate, consider a scene of a man giving cake to his son. It could be described either by a double object dative construction (DO), as in 1, or by a prepositional dative construction (PP), as in 2.

1. The man gave his son some cake.
2. The man gave some cake to his son.

If an experimenter describes this scene with the DO (the *prime* utterance), and then shows a picture of a girl telling her friend a story, structural priming would be shown by the increased likelihood that the participant's description of the scene (the *target* utterance) would use a DO as in 3 (rather than a PP as in 4).

3. The girl told her friend a story
4. The girl told a story to her friend.

Structural priming is seen as evidence of abstract syntax because it can operate across semantically different utterances and across intervening sentences (Bock, 1986; Chang, Bock, & Goldberg, 2003; Thothathiri & Snedeker, 2008). Thus the development of structural priming in children has been used to mark the development of syntactic abstraction (e.g., Savage et al., 2003). Indeed, Chang et al. (2006) have suggested that the mechanisms underlying structural priming are the same mechanisms involved in learning grammar. Pickering and Garrod (2004) additionally suggest that structural priming is one mechanism by which conversational fluency between interlocutors is achieved.

Two highly influential models, by Chang et al. (2006) and by Pickering & Garrod (2004), each account for many of the phenomena of structural priming. However, research by Goldwater et al. (2011) shows that some phenomena of structural priming can best be captured by using the mechanism underlying analogical reasoning—*structure-mapping* (Gentner, 1983; Gentner & Markman, 1997).

We propose that structural priming can be modeled as a species of analogy. This proposal might seem surprising, given that analogy is often considered to be a conscious phenomenon while priming is clearly implicit. However, recent results show that structure-mapping from a prior analog can occur without attention or awareness (Day & Gentner, 2007). We describe an initial computational model, based on analogical processes of matching, generalization, and retrieval. To provide a solid basis for comparison, we use the experimental design and stimulus generator developed by Chang et al (2006). We begin by summarizing the psychological experiments and the Chang et al (2006) model. We then describe our analogy-based simulation, including its structure and operation. The results of three simulation experiments are presented.

The Chang et al. dual-path model

In a typical structural priming experiment (e.g., Bock & Griffin, 2000) participants alternate between prime trials, on which they hear and then repeat a sentence, and target trials, on which they are given a depicted scene to describe in any way they choose. For example, in Bock and Griffin's (2000) Experiment 1 there were 48 such sequences. Any prime sentence would be in one of two syntactic alternates, e.g., the DO or PP dative, and the dependent measure is the frequency of matching the structure of the prime in the target scene description vs. using the alternate structure.

Chang et al. (2006) present a connectionist model of sentence production, the *dual-path* model, which simulates several structural priming phenomena. Their model includes one system for representing the message, i.e., the meaning of the sentence, and a second system for producing sentence structure from the message. Before simulating structural priming, the model was trained with 60,000 message-sentence pairs, each consisting of a meaning and a word sequence for that meaning. Using error-based learning, the model learned to produce grammatical word sequences when given a message. The model was then tested using conditions that mirrored structural priming experiments. In the prime trials, the model received both a message and a sentence structure. In the target trials, only the message was given to the model. On every prime sentence, the weights between nodes in the sequencing system were updated based on prediction error, just as in training.

The stimulus-producing grammar consisted of a set of message-sentence templates corresponding to the kinds of constructions used in the experiments (see Table 1 for examples). Random satisfactory fillers were chosen from a small fixed lexicon of concepts and bound to empty abstract thematic role slots in the message portion of the template, as well as to the event-semantic categories indicating the *tense* (e.g. present, past) and *aspect* (e.g. progressive) of the event represented. Their model uses a XYZ thematic role representation scheme, wherein the roles roughly correspond to agent, theme/patient, and recipient/location, respectively. The corresponding word strings for the concepts in the message were then bound to corresponding slots in the sentence template. Finally, a small set of transformations (e.g. morphemes for tense) were applied to produce grammatical sentences for the given sentence type.

Every structural priming test set took 100 prime–target message pairs. Each target message was presented twice, preceded each time by a prime with the same message but with a different syntactic alternate. There were two versions of each target message with a built-in bias towards one of the alternates, creating 4 trials per prime-target message pair, yielding 400 total prime-target trial sequences.

Chang et al. (2006)'s model was able to capture several key phenomena. They simulated priming both the dative alternation and the passive/active alternation. These constructions remained primed across multiple filler items between prime and target, as in people (Bock & Griffin, 2000). (see Table 1 for examples). The dual-path model was used to simulate other phenomena that we eventually plan to simulate as well, but these are our current focus.

The success of the dual-path model in simulating structural priming phenomena is very impressive. It has set a standard against which future models of structural priming will be measured. We use this model as a basis for calibrating our analogy based model, showing that we can capture some of the same phenomena with many fewer training examples.

Analogical learning and priming of constructions

Our model uses analogical processing in both the training phase and the priming phase. Training is modeled as analogical generalization (using SAGE, described below). During testing, when a target message is presented, analogical retrieval (via MAC/FAC) is used to efficiently retrieve a small number of utterances (or generalizations) from memory that overlap in content with the target. Then analogical mapping (SME) is used to map their sentence structure onto the role bindings in the target message.

We now review the components of our model. The major components—SME, MAC/FAC, and SAGE—were developed prior to this study, and have been shown to be useful in modeling other analogy-driven phenomena. We begin with SME, which underlies the others. Then we discuss analogical generalization via SAGE, which here models prior language learning. Finally, we turn to retrieval, which (along with mapping) is used to model priming.

Mapping: The Structure-Mapping Engine (SME) (Falkenhainer, Forbus, & Gentner, 1989) is a computational model of Gentner's (1983) structure-mapping theory of analogy. Its inputs are *base* and *target* structured representations. Its output is one or more *mappings* that describe how the two descriptions can be aligned (where alignment requires finding a like relational structure in which the relations match identically).[1] Each mapping consists of a set of *correspondences* linking elements from the base and target, a *score* based on the degree of overlap between them, and *candidate inferences* that represent hypotheses about what elements in one description could be projected to the other, based on the correspondences for that mapping. SME is used as a component in the other two analogical processes, and is also used here to generate word sequences to describe new utterances.

Retrieval: MAC/FAC (Forbus, Gentner, & Law, 1995) models similarity-based retrieval over structured representations. Its inputs are a probe case and a case library. The first stage of MAC/FAC rapidly retrieves up to three candidate matches using a crude parallel vector match, where the vectors are automatically constructed from the structured representations. The second stage uses SME in parallel to compare the probe to the structured representations for the candidates produced by the first

---

[1] Under some circumstances, nonidentical relations are re-represented to find identical subcomponents.

stage, returning the best mapping (or up to three, if very close) as the *reminding* for that probe.

Generalization: SAGE (Kuehne et al., 2000) models analogical generalization. It begins by storing the first input example (here, a message-sentence pair). When the next example arrives, SAGE compares it to the first one, using SME. If there is sufficient overlap (that is, if SME's score is above a pre-set threshold) the common structure is stored as a generalization. SAGE uses MAC/FAC to retrieve generalizations and/or examples similar to new inputs. New examples are assimilated into existing generalizations if sufficiently similar, and the generalization is updated based on their common structure. Otherwise, if the new example is very similar to a retrieved example, a new generalization is formed from their common structure. Finally, if the new example is not sufficiently similar to anything retrieved, it is stored separately, and may serve as a seed for another generalization later.

In essence, this process of progressive alignment leads to the gradual wearing away of the non-overlapping aspects of the examples. SAGE's generalizations are structured representations. They may also include some specific features, though generally many fewer than in the input representations. No variables are introduced. Further, the assimilation process produces probabilities attached to each statement in the description, indicating its frequency within the generalization. For each concept to be learned, the set of generalizations and exemplars learned so far constitutes its *generalization context*.

An Analogical Model of Structural Priming

In our model, target utterances are produced by retrieving utterances (or generalizations) from memory whose meaning is similar to that of the given target meaning and mapping their sentence structure onto the target.

The system's memory has a short-term as well as a long-term component, in order to simulate the greater availability of more recently encountered utterances. A buffer of messages, each linked to its sentence representation, is stored in the system's Short Term Memory (STM); these serve as priming utterances, as well as "filler" or distractor utterances. Given the message of a target utterance as input, the system uses analogical retrieval with the STM as the case library to find similar messages. Failing to find a semantically similar utterance in STM, the system uses MAC/FAC with the system's LTM as the case library. The LTM consists of the SAGE generalization context, that is, the generalizations and ungeneralized exemplars produced during the training phase (described below). SME is then used to infer a sequence of words that situates the actors and objects of the target utterance's meaning in their corresponding roles.

Returning to the prior example, the intended behavior of the model is as follows: The system is given a representation of an event in which a girl is telling her friend a story. In the structural priming condition, the STM contains meaning-sentence pairs. The presence of a prior utterance expressing a transfer of cake from father to son in the double-object dative (DO) form "The man is giving his son some cake", should lead to an increased likelihood for the system to produce the utterance "the girl is telling her friend a story", rather than "the girl is telling a story to her friend". Absent a priming example of this kind, the system should still be able to produce an utterance that conveys the target meaning by retrieving a generalization or exemplar from its LTM with a similar meaning.

To populate the model, a set of sentences paired with their meaning was generated using an input environment grammar and simple lexicon based on those used in Chang et al. (2006). We used the same grammar and lexicon as their generator, and compared the results of our generator to theirs to ensure that the sets of meaning-sentence pairs we produced were essentially the same. Some of these meaning-sentence pairs were set aside as stimuli to use in the priming experiments, with a distinct set used to train the model, as described below. Next we describe how these pairs were encoded by our simulation, and the training process it underwent.

Table 1: Sentence types included in the input environment grammar

Sentence type	Example Sentence
Animate intransitive	"a man jump –ed"
Animate with intransitive	"the girl walk –s with a dog"
Inanimate intransitive	"the ball bounce –s"
Locative transitive	"a father is go –ing around a car"
Theme-experiencer	"a uncle scare –s a cat"
Cause-motion	"the grandfather carry –ed a cup to the store"
Transfer dative	"a woman give –s a girl a apple"
Benefactive dative	"a man bake –ed a cake for the mother"
Benefactive transitive	"the boy push –ed a chair for the man"
State-change	"a cat plug –s a sink with a ball"
Locative alternation	"the father spray –ed paint on the wall"
	"a uncle loaded –ed a plate with pie"

Structure, Structural Priming & Sentence Production

Analogical processing assumes that people use structured, relational representations. Our input encodings, automatically produced from message-sentence pairs, reflect a reasonable approximation to what people would encode in similar situations. A complete example of the message-sentence pair representation used by our model can be seen in Figure 1. Sentence structure is represented by a series of

word slot entities (e.g. `w1`), each corresponding to a word in the sequence (e.g. `(isa w1 (WordFn "grandmother")))`. Sequentiality in the sequence is represented by a set of relationships between word slots. Semantic structure is represented by entities representing abstract thematic role fillers (e.g. `x0`), whose interrelationships are described via binary relations (e.g. `(roleX a0 x0)`). The referential structure ties the thematic roles to their corresponding word slot in the sentence structure, e.g. `(sameObject x0 (WordReferentFn w1))`. The use of words in the semantic structure (e.g. `(isa x0 (WordFn "grandmother")))` follows the Chang et al model, which used words rather than internal concepts as fillers in their meaning representation. Consequently, we did the same, in order to keep the simulations as comparable as possible.
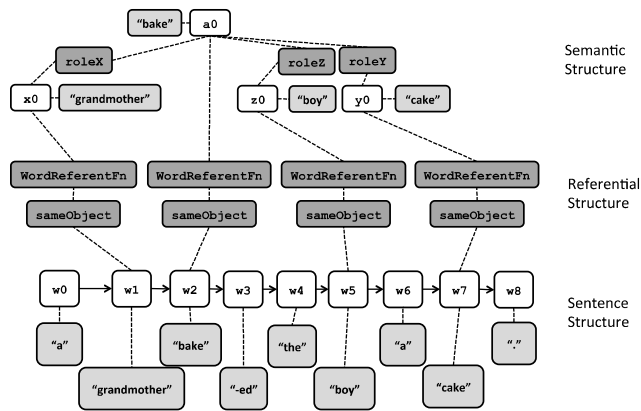

Figure 1: Interrelations between semantic and sentence structure through the referential structure layer.

Seeding LTM via analogical generalization

Prior to the priming tests, a training procedure was run to seed the system's Long Term Memory (corresponding to Cheng et al.'s Training phase). Our training utilized analogical generalization via SAGE. Five examples of each of the 24 variants of the 11 construction types in the input environment grammar were produced: 120 message-sentence pairs in total. These stimuli were incrementally generalized by SAGE, using a similarity threshold of 0.9. This resulted in 45 separate generalizations of message-sentence pairs and 15 concrete, ungeneralized message-sentence exemplars. SAGE required just one pass through the 120 examples, which is two orders of magnitude less exposures than the dual-path model required.

Sentence production

Given a new semantic message m_i, a prime p_i, and a set of filler message-sentence pairs the prime and fillers are stored in STM. Then, the system uses MAC/FAC to find the most similar semantic message to m_i from among the messages present in STM, and if that fails, MAC/FAC is used on the LTM. In either case, once a sufficiently similar message is

retrieved, SME's alignment of that message-sentence pair with the input message produces candidate inferences representing hypotheses about the structure of the target sentence. These candidate inferences are used to produce sentence structure for the target, by projecting word information and order relationships from the retrieved utterance (or generalization) to the description of the target message.

Priming Experiments

We next evaluate the model's ability to produce sentences from messages without primes (Experiment 1) and with two kinds of priming alternations (Experiments 2 and 3). All three studies used the LTM generated by SAGE as described above.

Experiment 1

In Experiment 1, we tested the model's production in two LTM-only conditions: a dative production condition and a transitive production condition. This examines the model's ability to select a proper grammatical form for messages in the absence of specific prime sentences in STM. In each condition, the model was given a sequence of 50 examples of messages corresponding to the given construction type and required to produce a sentence for each. As noted above, this means that the model will use MAC/FAC to retrieve generalizations and exemplars from the LTM contents produced via SAGE to do the generation.

We applied a twofold evaluation to the output of our model, similar to that used by Chang et al. (2006). Each sentence produced by the model is evaluated in terms of its *grammaticality* and its *message accuracy*. *Grammaticality* measures the degree to which the output sentence matches the prototype defined in the input environment grammar. *Message accuracy* measures the degree to which the semantic message retrieved from memory maps to the target message given as input. The results are summarized in Figure 3. For both kinds of constructions the model's message accuracy and grammaticality is quite high. Even with an extremely limited training set, our analogy-based model produces sentences conforming to the input grammar.

Experiment 2

Next we tested the model's performance when presented with a dative prime from one of two alternates: the prepositional form and the dative form. We also wished to test whether the model would capture the finding that structural priming can persist across intervening sentences (Bock & Griffin, 2000). Therefore we varied whether there were intervening intransitive filler sentences in STM. This led to a 2X2 design: Alternative constructions (prepositional dative vs. double-object dative) crossed with Filler conditions (no fillers vs. intransitive fillers). In each condition the model was given a sequence of 100 prime-target pairs with dative messages. The prime message-sentence stimulus was stored in STM and the system was required to produce an appropriate sentence for the target

message. In the no-filler condition, no additional message-sentence pairs were entered into STM. In the intransitive filler condition, 10 intransitive stimuli were entered into STM in addition to the prime stimulus. For both kinds of constructions, and with fillers and no fillers in STM, the model matched the sentence structure of the prime in every trial. That is, the model was able to find a proper match in STM on every trial and to map its structure to the target without using the LTM store of sentences.
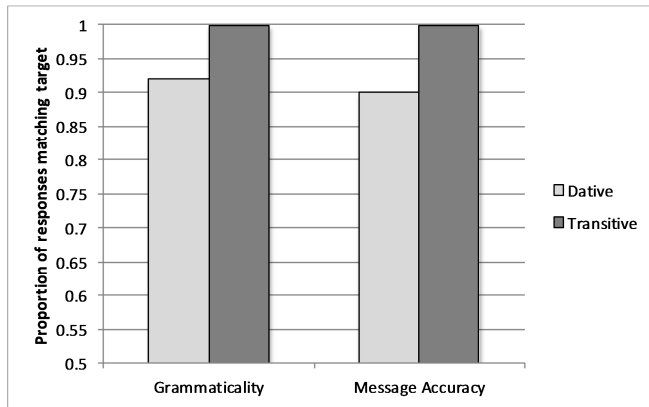

Figure 3: Sentence production performance of model in LTM-only retrieval condition

Experiment 3

Next we tested the model's performance when presented with transitive primes that were either active or passive. The experiment used the same basic 2x2 design and procedure as in Experiment 2, and the same number of prime-target message pairs. As usual, the system first attempted to retrieve a structural match from its STM before retrieving from its LTM.

For both active and passive priming conditions without fillers, the model matched the target structure to the prime for 100 out of 100 trials using the STM store. When there were fillers, the model was able to do so for 98 of the passive trials, and 99 of the active trials. That is, LTM was used as a basis for target sentence structure a total of 3 times across 400 trials. The model produced grammatically appropriate sentences with both STM and LTM retrievals.

Discussion

These experiments show that our analogy-based model is (1) capable of forming generalizations over meaning-sentence pairs; (2) able to use its learned memory of generalizations and exemplars to produce sentences conforming to the input grammar when given a meaning (Experiment 1); (3) able to match the structure of prime sentences for either the dative alternation (Experiment 2) or the active/passive transitive alternation (Experiment 3). As per human data, the presence of intransitive fillers had minimal effect on the effects of a prime. The model can simulate structural priming when there is no lexical overlap

between prime and target utterances across structurally dissimilar fillers, matching human findings.

These findings provide evidence for the viability of analogical mechanisms in learning constructions and in applying them to form utterances. That analogical processes readily accommodate both learning and priming phenomena is in accord with the idea that the two phenomena are intimately related, as suggested by Chang et al (2006). We now discuss these two aspects in more detail, including both implications and limitations of the current model. We begin with structural priming and then turn to grammar learning.

Structural priming

While the strong priming effects our model shows is encouraging support for analogical mapping as a mechanism of structural priming, in some sense the model's performance is too good. Across Experiments 2 and 3, over 90% of the targets conform to the structure of the prime. Priming effects are typically much smaller in humans; in general, roughly 60% of targets conform to the prime. We believe there are two reasons for this. The first is that we only consider structural priming, and not other types of constraints, such as distributional and semantic preferences connected with individual words and phrases, pragmatic constraints, and discourse constraints that enter into construction selection in natural language use (e.g., Bresnan et al., 2007). Chang et al. (2006) dealt with this issue by building in a bias into every message towards a particular construction; these bias effects can act as a competing (or facilitating) force on priming. We are exploring ways to capture these effects. The other reason may be the overly strong reliance on an STM buffer in the current model. Recall that analogical retrieval is used on LTM only when retrieval on STM fails. This happened only three times across Experiments 2 and 3. We suspect that reducing the bias towards STM retrieval, or even eliminating the STM-LTM distinction entirely, might more closely match human data. Such a model would take into account both recency (thereby favoring STM) and strength of generalization (favoring LTM).

Learning grammatical patterns

An intriguing result is the effectiveness of analogical generalization, as modeled by SAGE, in learning grammatical patterns. SAGE was given only one pass through 120 example message-sentence pairs, yet it produced a set of generalizations (along with some isolated examples) that was sufficient to support the construction of grammatically and semantically accurate sentences over 90% of the time. In contrast, the dual-path model required 8,000 examples, each trained an average of 7.5 times-- around 60,000 trials.

Why is our analogical model of construction generalization so effective? In an important sense, we believe this finding is real: Structural alignment and abstraction is a highly effective way of extracting common relational structure. For example, Kuehne et al. (2000) used

SEQL (a predecessor of SAGE) to simulate the Marcus et al. (1999) studies, in which 7-month old infants abstracted a grammar-like rule from exemplars. The model required only the amount of exposure given to the infants—16 strings repeated 3 times each, a total of 48 strings.

However, the obvious challenge to our results is that children do not master grammar in 120 utterances, nor even after many thousands. We suggest that a major source of the disparity lies in the nature of the input. We can characterize learning environments on a continuum from high-alignable to low-alignable. In a high-alignable environment, the learner encounters juxtaposed alignable pairs, as in the Marcus et al. studies. Lab studies show dramatic learning under these conditions (Gentner, 2010). On the other hand, children's language learning takes place in a low-alignable environment; they only occasionally receive perfectly alignable juxtapositions (Cameron-Faulkner et al., 2003)

A unified approach to language

Despite the differences in specific mechanisms between our models, we share an important commitment with Chang et al.: namely, that the mechanisms of structural priming can also be applied to grammar learning in children. Goldwater et al. (2011) found a developmental sequence towards less reliance on high semantic similarity in structural priming— an effect specifically predicted by a structure-mapping account of grammar learning. There is also evidence that analogical processes enter into learning word meanings, particularly for relational terms such as verbs (Childers, 2008). If further studies bear out the hypothesis that analogical processes are involved in grammar learning, this will implicate analogy as a major force in language learning.

Acknowledgments

References

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive Psychology, 18*(3), 355-387.

Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General, 129*(2), 177-192.

Branigan, H., Pickering, M., & Cleland, A. (2000). Syntactic co-ordination in dialogue. *Cognition, 75*(2), 13-25.

Cameron-Faulkner, T., Lieven, E., & Tomasello, M. (2003). A construction-based analysis of child-directed speech. *Cognitive Science*, 27, 843-874.

Chang, F., Bock, K., & Goldberg, A. (2003). Can thematic roles leave traces of their places? *Cognition, 90*(1), 29-49.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review, 113*(2), 234-272.

Childers, J. B. (2008). The structural alignment and comparison of events in verb acquisition. *Proceedings of the 30th Annual Cognitive Science Society*. Austin ,TX

Day, S. B., & Gentner, D. (2007). Nonintentional analogical inference in text comprehension. *Memory & Cognition, 35*(1), 39-49.

Falkenhainer, B. F., Forbus, K. D., & Gentner, D. (1989). The Structure Mapping Engine: Algorithm and examples. *Artificial Intelligence, 41*, 1-63.

Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of Similarity-based Retrieval. *Cognitive Science, 19*(2), 141-205.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-170.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52*, 45-56.

Gentner, D. & Markman, A. B. (2006). Defining structural similarity. *The Journal of Cognitive Science*, 6, 1-20.

Goldwater, M. B., Tomlinson, M. T., Echols, C. H., & Love, B. C. (2011). Structural Priming as Structure-Mapping: Children Use Analogies From Previous Utterances to Guide Sentence Production. *Cognitive Science, 35*(1), 156-170.

Hare, M., & Goldberg, A. (1999). Structural priming: Purely syntactic? *Proceedings of the 21st Annual Conference of the Cognitive Science Society* (pp. 208-211). Vancouver.

Kaschak, M. P., & Borreggine, K. L. (2008). Is long-term structural priming affected by patterns of experience with individual verbs? *Journal of Memory and Language, 58*(3), 862-878.

Kuehne, S. E., Forbus, K. D., Gentner, D., & Quinn, B. (2000). SEQL: Category learning as progressive abstraction using structure mapping. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*: Philadelphia, PA.

Pickering, M., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*(02), 169-190.

Savage, C., Lieven, E., Theakston, A., & Tomasello, M. (2003). Testing the abstractness of children's linguistic representations: lexical and structural priming of syntactic constructions in young children. *Developmental Science, 6*(5), 557-567.

Snider, N. (2009). Similarity and structural priming. *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 815–820). Amsterdam.

Thothathiri, M., & Snedeker, J. (2008). Syntactic priming during language comprehension in three-and four-year-old children. *Journal of Memory and Language, 58*(2), 188-213.