THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Abstraction in time

Link:
Link to publication record in Edinburgh Research Explorer

OPEN ACCESS

Abstraction in time: Finding hierarchical linguistic structure in a model of relational processing

Leonidas A.A. Doumas (alex.doumas@ed.ac.uk)
Andrea E. Martin (andrea.martin@ed.ac.uk)
Department of Psychology
School of Philosophy, Psychology and Language Sciences
University of Edinburgh
7 George Square, Edinburgh, EH8 9JZ, United Kingdom

Abstract

Abstract mental representation is fundamental for human cognition. Forming such representations in time, especially from dynamic and noisy perceptual input, is a challenge for any processing modality, but perhaps none so acutely as for language processing. We show that LISA (Hummel & Holyoak, 1997) and DORA (Doumas, Hummel, & Sandhofer, 2008), models built to process and to learn structured (i.e., symbolic) representations of conceptual properties and relations from unstructured inputs, show oscillatory activation during processing that is highly similar to the cortical activity elicited by the linguistic stimuli from Ding et al. (2016). We argue, as Ding et al. (2016), that this activation reflects formation of hierarchical linguistic representation, and furthermore, that the kind of computational mechanisms in LISA/DORA (e.g., temporal binding by systematic asynchrony of firing) may underlie formation of abstract linguistic representations in the human brain. It may be this repurposing that allowed for the generation or emergence of hierarchical linguistic structure, and therefore, human language, from extant cognitive and neural systems. We conclude that models of thinking and reasoning and models of language processing must be integrated—not only for increased plausibility, but in order to advance both fields towards a larger integrative model of human cognition.

Keywords: computational models, sentence processing, analogy, relational reasoning, concepts, binding, temporal asynchrony, oscillations, computational neuroscience

Introduction

Abstract hierarchical representations are the hallmark of human language (Chomsky, 1957). Forming them is certainly necessary during language processing (see Martin, 2016 for a possible process model). But what are the computational origins of such an ability? One possibility is that the brain repurposed a mechanism or process already at its disposal when abstraction become an efficient solution to a problem posed by the environment, such as communicating information across time and space, or predicating novel arguments that you have never encountered before.

Ding et al. (2016) recently showed evidence of cortical tracking of abstract, hierarchical linguistic structures in oscillatory patterns in data from an electrocorticography (ECoG) and an magnetoencephalography (MEG) experiment. This tracking crucially could not be attributed to processing of acoustic information, transitional probability, or word predictability (Ding et al., 2016). Strikingly, LISA (Hummel & Holyoak, 1997, 2003), a symbolic-connectionist model of analogy making, and DORA (Doumas, Hummel, & Sandhofer, 2008), a symbolic connectionist model of relational reasoning, predict such a representational pattern. We tested the hypothesis that the representational patterns produced by LISA and DORA during their processing will give rise to hierarchical structures matching the linguistic structures observed by Ding et al. (2016) without any formal or structural changes to the model. Such an approach would be particularly compelling because it shines a light both on how the brain might parse language (i.e., the class of possible parsing mechanisms underlying cortical tracking of linguistic representations as seen in Ding et al., 2016), and about how linguistic structures might have come to be the way there are. In order to test these predictions, we simulated oscillatory unit data in LISA/DORA using the same English sentence stimuli as Ding et al. (2016). We tested whether LISA/DORA parsed the sentences correctly, and we observed the pattern of unit firing LISA/DORA exhibited while parsing the sentences.

The LISA/DORA models

LISA (Hummel & Holyoak, 1997, 2003) is a model of analogy and relational reasoning. DORA (Doumas et al., 2008) is an extension of LISA that learns structured (i.e., symbolic) representations of relations from unstructured (holistic flat feature vector) inputs. That is, DORA provides an account of how the structured relational representations LISA uses to perform relational reasoning are learned from examples. Both LISA and DORA are symbolic-connectionist models, or models based on traditional connectionist computing principles that effectively implement symbolic representations as they solve the binding problem and

can, therefore, represent compositional structures (see Doumas & Hummel, 2005, 2012).

DORA accounts for over 25 phenomena from the literature on relational learning, as well as its development (e.g., Doumas & Hummel, 2010; Doumas et al., 2008; Lim et al., 2014; Morrison et al., 2012; Sandhofer & Doumas, 2008). In addition, as DORA learns relational representations, it comes to take LISA as a special case, and can simulate the additional 30-plus phenomena in relational thinking simulated by LISA. The description of LISA/DORA that follows is a brief overview due to space constraints. For full details of the models and their operations see Doumas et al. (2008) and Hummel and Holyoak (1997, 2003).

LISAese Representations In LISA (and DORA *after* it has gone through learning) relational structures are represented by a hierarchy of distributed and localist codes (see Figure 1). At the bottom, "semantic" units represent the features of objects and roles in a distributed fashion. At the next level, these distributed representations are connected to localist units (POs) representing individual predicates (or roles) and objects. Localist role-binding units (RBs) link object and predicate units into role-filler binding pairs. At the top of the hierarchy, localist P units link RBs into whole relational propositions.

chase (goblin, gnome)

chaser+goblin   chased+gnome

chr.   goblin   chd.   gnome

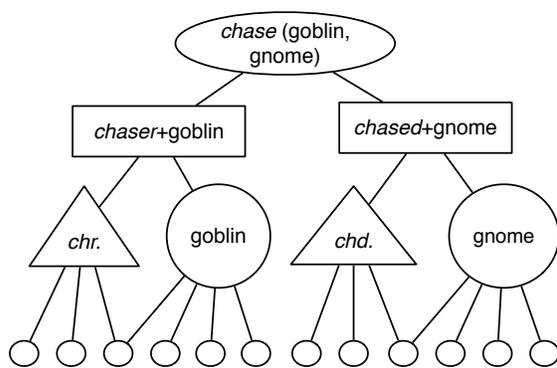P units

RB units

PO units

semantic units

Figure 1. Representation of a LISA/DORA representation of the proposition *chase* (goblin, gnome). We use different shapes to represent units in different layers (ovals for P units, rectangles for RB units, triangles and large circles for PO units, and small circles for semantic units) for the purposes of clarity. In the model these units are simply nodes in different layers of the network.

Propositions are divided into two mutually exclusive sets: a driver and one or more recipients. In LISA/DORA, the sequence of firing events is controlled by the driver. We take the driver to be the focus of attention in LISA/DORA (i.e., what LISA/DORA is attending to at a specific moment). The driver contains one (or at most three) proposition(s).

Activation flows from the driver units to their semantic units. Units in the driver and recipient share the same pool of semantic units. Thus, units in the recipient become active in response to the pattern of activation imposed on the semantic units by the active driver proposition. The flow of activation from driver to recipient through shared semantic units is important for many of LISA and DORA's processes including comparison, analogical mapping, relation learning, schema induction, and memory retrieval. We will not discuss these processes further as they are not important for the purposes of the current paper, but full details may be found in Hummel & Holyoak, 1997, 2003) and Doumas et al. (2008).

Representing binding information What is most important about LISA/DORA for the purposes of the present paper is the manner in which the models solve the binding problem. As noted above, LISA and DORA are symbolic-connectionist models. That is, they are based on traditional connectionist computing principles (i.e., layers of interconnected nodes passing activation via weighted connections that are modified via Hebbian learning), but unlike traditional connectionist systems, they can process symbolic structure.

Processing symbolic structure requires that representational elements in a system can be composed into meaningful structures in a manner that does not violate the independence of those elements (see e.g., Markman, 1999; Russell & Norvig, 2003). For example, representing a relational proposition like *chase* (goblin, gnome) requires representing that *chasing,* a goblin, and a gnome are all present, and that goblin is bound the *chaser* role and gnome to the *chased* role. Importantly, the binding of *chaser* to goblin must not change the fundamental meaning of either what it means to be a goblin or what it means to be a *chaser*—i.e., the binding system must not violate role-filler independence.

In LISA and DORA roles and fillers are represented independently in the PO and semantic units. In order to behave symbolically, however, when a proposition in the driver becomes active, role-filler bindings must be represented dynamically on these units (i.e., POs and semantic units; see Hummel & Holyoak, 1997). Both LISA and DORA use time to carry dynamic binding information.

Binding information is represented in LISA with bound role-filler pairs firing in synchrony. To illustrate, when a proposition like *chase* (goblin, gnome) becomes active in the driver (Figure 2a), the units representing *chaser* and goblin become active and fire together (representing the binding between *chaser* and goblin; Figure 2a[i]). Subsequently, the units representing *chased* and gnome become active and fire together (representing the binding between *chased* and gnome;

Figure 2a[ii]). Importantly, bound role-filler pairs fire together, and out of synchrony with other bound role-filler pairs. The distinct firing bursts allow LISA to code bindings between roles and their fillers, and process these structures symbolically, forming the basis of LISA's capacity to solve analogical mappings, and perform relational inference (see Hummel & Holyoak, 2003)
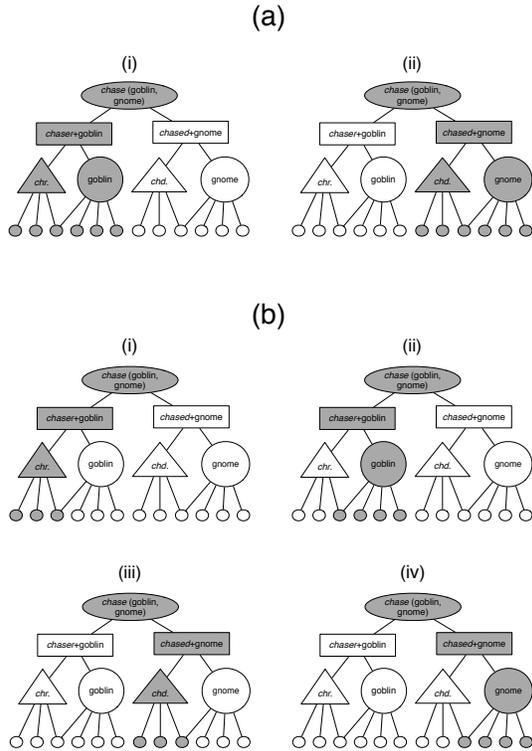
(a)

(b)

Figure 2. Dynamic binding in LISA and DORA. (a) Binding in LISA. (i) To bind the role *chaser* to goblin, units coding for the *chaser* role and goblin (as well as those coding conjunctively for *chaser*+goblin) fire. (ii) To bind *chased* to gnome, units coding for the *chased* role and gnome (as well as those coding conjunctively for *chased*+gnome) fire. (b) Binding in DORA. , (i-ii) To bind the role *chaser* to goblin, units coding for the *chaser* role (as well as those coding conjunctively for *chaser*+goblin) fire, followed by the units coding for goblin (as well as those coding conjunctively for *chaser*+goblin). (iii-iv) To bind the role *chased* to gnome, units coding for the *chased* role (as well as those coding conjunctively for *chased*+gnome) fire, followed by the units coding for gnome (as well as those coding conjunctively for *chased*+gnome).

In DORA, binding information can be carried either by synchrony (as in LISA) or by systematic asynchrony of firing, with bound role-filler pairs firing in direct sequence.[1] During asynchronous binding, when a proposition like *chase* (goblin, gnome) becomes active in the driver (Figure 2b), the units representing *chaser* fire (along with units conjunctively coding for *chaser*+goblin and for the *chase* (goblin, gnome) proposition; Figure 2b[i]), followed directly by the units representing goblin (along with units conjunctively coding for *chaser*+goblin and for the *chase* (goblin, gnome) proposition; Figure 2b[ii]), representing the binding of *chaser* to goblin. Then, the units representing *chased* fire (along with units conjunctively coding for *chased*+gnome and for the *chase* (goblin, gnome) proposition; Figure2b[iii]), followed directly by the units representing gnome (along with units conjunctively coding for *chased*+gnome and for the *chase* (goblin, gnome) proposition; Figure 2b[iv]), representing the binding of *chased* to gnome. In short, bound role-filler pairs fire in direct sequence, and out of synchrony with other bound role-filler pairs. These patterns of sequential oscillation dynamically code role-filler bindings in DORA, and underlie DORA's capacity to use the representations that it learns to support relational reasoning (e.g., analogical mapping, schema induction, and relational induction; see Doumas et al., 2008) and to learn structured relational representations from unstructured object representations.

Crucially, sequential firing of related constituent elements is a necessary property of binding via synchrony and systematic asynchrony. When LISA/DORA perform any structured processing, a pattern will invariably emerge wherein bound elements within a larger compositional proposition will fire in direct sequence and at a different time-scale than units coding for conjunctions of independently bound elements and full propositional compounds. In the following section we show that the pattern produced by LISA/DORA as it processes compositional structures matches very closely the temporal pattern of spike activity observed in Ding et al.'s (2016) when people process compositional propositions.

Simulation

Ding et al. (2016) presented auditory strings of synthesized speech in Mandarin Chinese in an MEG experiment, and strings of synthesized speech in American English in an ECoG experiment. They manipulated the structural relationship between the units in the auditory string, i.e., the syllables. In one condition, there was no meaningful relationship

---

[1] Asynchrony-based binding allows roles and fillers to be coded by the same pool of semantic units, which allows DORA to learn representations of relations from representations of objects (Doumas et al., 2008).

between the strings of syllables, in the second condition, phrases were formed from adjacent syllables, and in the third condition, sentences emerged from the string of syllables. Using this design, they observed peaks in the MEG-based oscillatory response on the timescale of syllabic rate (4Hz), phrasal rate (2Hz), and sentential rate (1Hz). Importantly, for trials with Mandarin sentences only speakers of Mandarin (compared to English speakers in a control group) showed tracking of phrasal and sentential representations in the form of peaks at 2Hz and 1Hz, respectively, although both English and Mandarin speakers showed tracking of speech/acoustic-syllabic stimuli regardless of language comprehension. The ECoG data from English speakers showed a similar pattern to the Mandarin MEG data, but without direct one-to-one acoustic-syllabic-to-phrase correspondence. Importantly, Ding et al. also observed cortical activity coding for sentence structure when English speakers tracked sentences of varying syllabic durations. For example, when English speakers tracked sentences with a noun phrase followed by a verb phrase wherein the initial noun phrase was three or four syllables (e.g., "the gold lamp", or "mahogany desk"), cortical activity tracked the entire phrasal structure, with a burst firing for the duration of the phrasal unit. Ding et al. controlled for effects of predictability in a string by showing that tracking of phrasal and sentence forms is not confounded by transitional probability.

Ding et al.'s results suggest definite structural form emerging during sentence processing. Specifically, beyond processing information at the level of syllables (or the basic features of a sentence tracked even by non-speakers of a language), speakers of a language process information that appears to track phrase structure and sentence structure. Moreover, when processing simple 2 argument verb structures, the structural pattern that emerges is two significant cortical response peaks (seemingly capturing phrasal information) firing within (at twice the rate) of a single cortical response peak (seemingly capturing sentence information).

We simulated the Ding et al. (2016) studies using the same English sentences used in their experiments 5 and 6 (with native English speakers). All of these sentences took the form modifier-noun-verb-noun, forming sentences like, "new plans give hope", and "dry fur rubs skin". LISA/DORA can represent sentences of this type in two ways. Most simply, such sentences could be represented with the modified noun represented as a single object containing both the semantics of the object and the modifier (see Figure 3a). Alternately, LISA/DORA can represent hierarchical propositions by representing propositional structures as arguments of other propositional structures. For example, to represent "dry fur rubs skin" the modified noun phrase "dry fur" can be represented

explicitly by the propositional structure *dry*(fur), which can then serve as the argument of the agent role of the *rubs* relation (see Figure 3b; details of higher-order structure representation in LISA and DORA can be found in Hummel & Holyoak, 1997 and Doumas et al., 2008). We have previously hypothesised that LISA/DORA can alternate between these types of representation depending on the properties of the current task (e.g., Doumas et al., 2008; Rabagliati et al., submitted). Specifically, when modifier information must be considered explicitly, the later type of representation (as in Figure 3b) might be employed. Alternately, when the modifier information can be considered implicitly, the former type of representation (as in Figure 3a) can be employed. For the purposes of the current simulations both types of representations would work, however we used the hierarchical representations (i.e., Figure 3b) to code the sentences following our assumption that participants coded the modifier-noun structure explicitly.
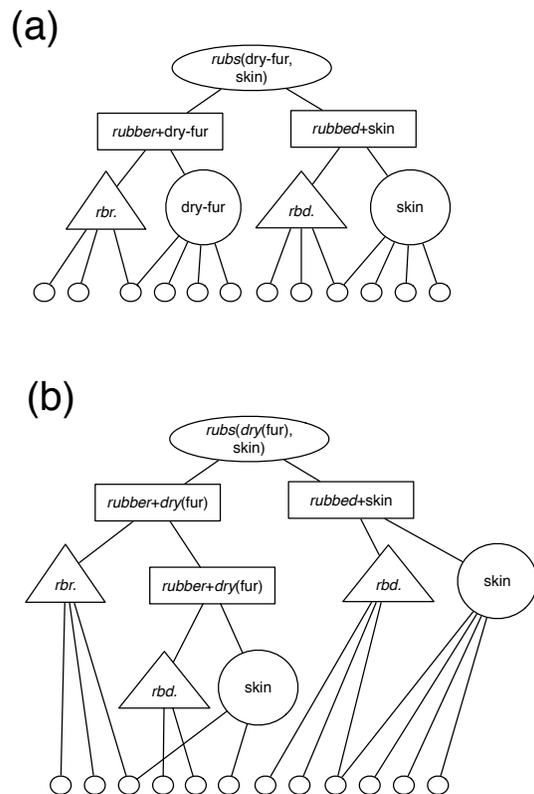

Figure 3. Representations of the sentence "dry fur rubs skin" in LISA/DORA. (a) A representation where the dry-skin modified noun is represented as a single unit connected to the semantics of both dry and skin. (b) A higher-order representation of the sentence where the modified noun is represented as a predicate structure, *dry*(skin) taken as an argument of the agent role of the *rubs* relation.

It is also important to note that the DORA model can learn all of the representations used in the current simulation from experience. As demonstrated previously (e.g., Doumas & Hummel, 2010; Doumas et al., 2008; Hamer & Doumas, 2013; Lim et al., 2014; Sandhofer & Doumas, 2008), DORA can learn explicit structured (i.e., symbolic) representations of verb structures like *give*, *rubs*, or *chases*, and of single-place modifiers like *dry*, *new*, or *golden* from experience with objects in the world involved in those relations or with those feature.

To simulate Ding et al.'s experimental procedure we allowed LISA/DORA to process Ding et al.'s English sentences one at a time. Representations of the sentence structures entered the driver (i.e., were attended to). LISA/DORA processed the sentences as it normally would (i.e., the units fired to represent and encode binding information; see above). We tracked firing rate of all the nodes in the driver as LISA/DORA processed the sentences. Because of the controlled length and structure of the sentences, DORA, like the participants in the Ding et al. experiments, took the same amount of time to process each sentence. The results of the simulation and the comparison to the patterns observed by Ding et al. are presented in Figure 4. Interestingly, the pattern of firing of the nodes in the various layers of LISA/DORA very closely mirror the patterns observed by Ding et al. Specifically, just like the human participants, DORA showed an activation burst in it's P units that lasted throughout the processing of the sentence (i.e., firing at the 1Hz range), activation bursts at twice the rate of the whole sentence burst (i.e., the RB unit firing in the 2Hz range), and activation bursts at 4 times the rate of the whole sentence burst (i.e., the PO units firing in the 4 Hz range).
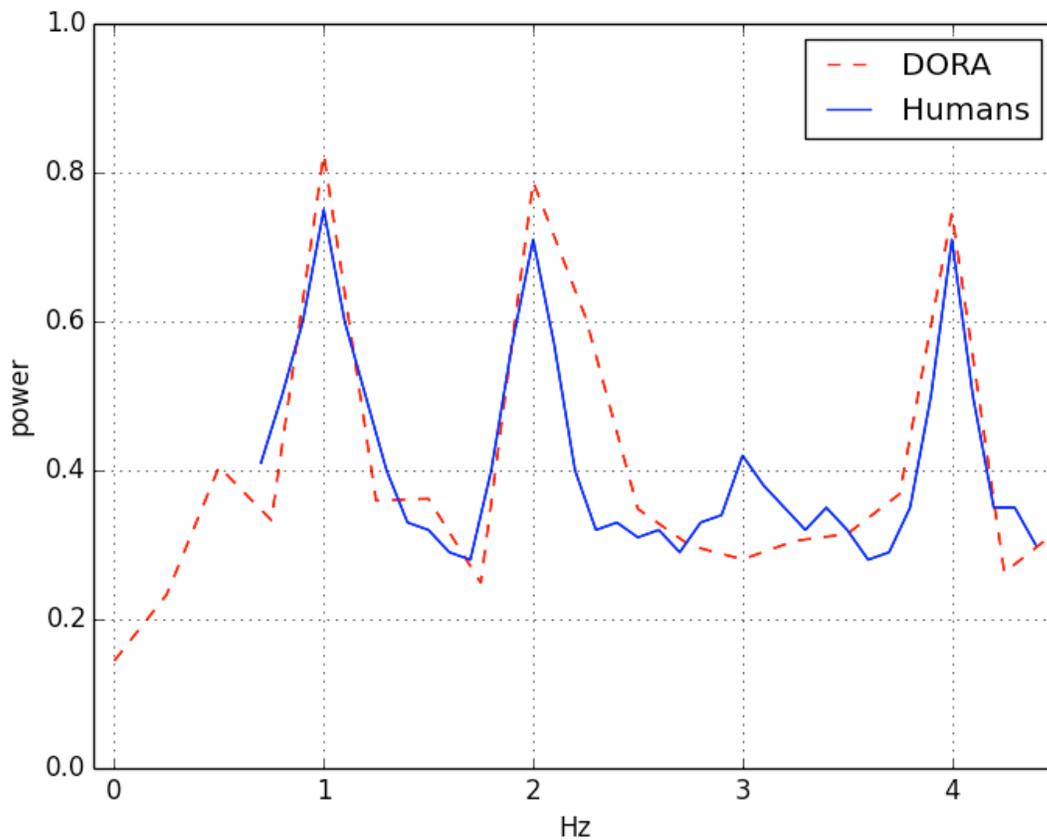

Figure 4. The solid line represents cortical power of participants listening to 4 word sentences played for 1 second in Ding et al. (2016). High cortical firing is evident at the 1Hz (the duration of the sentence), 2Hz, and 4 Hz range. The dashed line depicts firing in LISA/DORA while processing the same sentences used in Ding et al. There is evidence of units firing for the duration of the sentence, at intervals of half the length of the sentence, and at intervals lasting 1 quarter of the length of the sentence.

Conclusion

We have shown that abstract, hierarchical linguistic representation can be acquired, represented, and processed by LISA/DORA, models that were built for completely different purposes (analogy making in the case of LISA and relation learning in the case of DORA). Specifically, we have shown that the oscillatory activation patterns in LISA/DORA that arise as a natural consequence of the models performing dynamic binding appear to very closely fit data from human cortical tracking of hierarchical linguistic units (Ding et al., 2016).

It is interesting that models built for completely different purposes so successfully perform another task without modification. It is notable that extant models of sentence processing would likely not generalise so seamlessly to tasks such as analogical reasoning. Both probabilistic grammar (e.g., Levy, 2008) and connectionist approaches (see Joanisse & McClelland, 2015 for a review) must either be given a set of explicit grammatical phrase structure rules, or must learn the statistical specifics of a particular syntactic structure or parsing problem, rendering them unable to generalize to problems that routinely violate statistical and featural regularity like analogical reasoning (see, e.g., Holyoak, 2012).

We take our results to provide computational support the general claim—see, e.g., Penn, Holyoak, and Povinelli (2008)—that the ability to form and represent relational roles may underlie a number of our uniquely human cognitive capacities such as language. It is, perhaps, telling that the very same mechanisms that are necessary for processing relational structure and performing relational cognition seem to so closely simulate language processing.

Given our results, we suggest that models of relational reasoning and language processing might fruitfully be integrated. Such an integrative approach offers the possibility of producing powerful, neurophysiologically and cognitively plausible models that can perform well on multiple problems. We aim to further articulate the model by testing DORA on natural speech input, varied syntactic structures, in rich discourse contexts, on multilingual input, and with different assumptions about existing knowledge representations.

References

Ding, N., Melloni, L., Zhang, H., Tian, X., & Poeppel, D. (2016). Cortical tracking of hierarchical linguistic structures in connected speech. *Nature neuroscience*, 19(1), 158-164.

Doumas, L. A. A. & Hummel, J. E. (2010). A computational account of the development of the generalization of shape information. *Cognitive Science*, 34, 698-712.

Doumas, L.A.A., Hummel, J.E., & Sandhofer, C.M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review,115,* 1-43.

Hamer, A. J. & Doumas, L. A. A. (2013). Discovering quantification and number in a role-filler model. *Proceedings of the Thirty-Fifth Annual Conference of the Cognitive Science Society*.

Holyoak, K. J. (2012). Analogy and relational reasoning. In K. J. Holyoak & R. G. Morrison (Eds.), *The Oxford handbook of thinking and reasoning* (pp. 234-259). New York: Oxford University Press.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427-466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110,* 220-264.

Joanisse, M. F., & McClelland, J. L. (2015). Connectionist perspectives on language learning, representation and processing. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(3), 235-247.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126-1177.

Lim, A., Doumas, L. A. A., & Sinnett, S. (2014). Supramodal Representations in Melodic Perception. *Proceedings of the Thirty-Sixth Annual Conference of the Cognitive Science Society*.

Markman, A. E. (1999). *Knowledge representation*. Taylor & Francis: New York, NY.

Martin, A. E. (2016). Language Processing as Cue Integration: Grounding the Psychology of Language in Perception and Neurophysiology. In press at *Frontiers in Psychology: Language Sciences*.

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Brain and Behavioral Sciences*, 31, 109-178.

Russell, S. &Norvig, P. (2003). *Artificial Intelligence: A modern approach*. Pearson.

Sandhofer, C. M. & Doumas, L. A. A. (2008). Order and presentation effects in learning categories. *Journal of Cognition and Development*, 9, 194-221.