Constructing Meaning in Small Increments

Peter Lindes (plindes@umich.edu)
Computer Science and Engineering, 2260 Hayward Street
Ann Arbor, MI 48105 USA

Abstract

Humans comprehend natural language sentences in real time, processing the elements of each sentence incrementally with immediate interpretation, while working within the limitations of general cognitive abilities. While much research has been devoted to human sentence comprehension, a detailed computational theory of how this is done has been lacking. In this paper we explore some fundamental principles of human sentence comprehension, propose a novel computational theory of knowledge representation and incremental processing to comprehend sentences using general cognitive abilities, and discuss results of an implementation of this theory in a robotic agent using the Soar cognitive architecture. We then explore the theory's implications for future work in various areas of cognitive science.

Keywords: sentence comprehension; construction grammar; immediate interpretation; ambiguity resolution; cognitive architecture; language and the brain.

Introduction

Humans comprehend natural language sentences in real time, processing the elements of each sentence incrementally with immediate interpretation, while working within the limitations of general cognitive abilities. Much research has been devoted to human sentence comprehension.

Psycholinguists have measured many aspects of human language processing. A classic study (Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995) shows humans commit to the meanings of words and phrases before a complete sentence has been processed, doing *immediate interpretation*. A theory called *chunk-and-pass* (Christiansen & Chater, 2016) argues that this is necessary to comprehend sentences using humans' limited working memory capacity. However, theories on immediate interpretation and chunk-and-pass processing do not have computational implementations.

Linguists have built complex theories of syntax and semantics. Theories of *construction grammar* have been developed within cognitive linguistics for several decades (Hoffman & Trousdale, 2013). While they provide an approach that is more cognitively plausible than much traditional linguistic theory, only a few have computational implementations (Eppe, Trott, Raghuram, Feldman, & Janin, 2016; Steels, 2013) to allow humans to interact with robots.

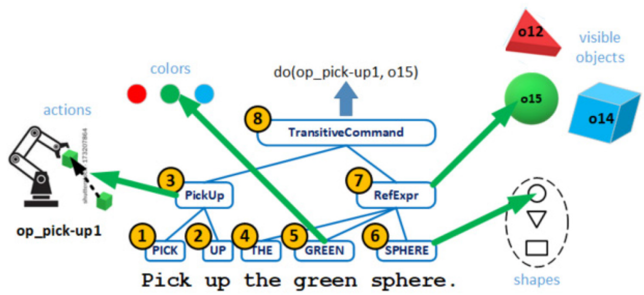

Figure 1: The meaning of a sentence

However, these do not do incremental processing or use general cognitive abilities.

In the artificial intelligence community, current work on natural language processing is concentrated on statistical techniques using deep neural networks (Young, Hazarika, Poria, & Cambria, 2018). Recently some work has been done on combining symbolic and neural network techniques (Mao, Gan, Kohli, Tenenbaum, & Wu, 2019). These techniques do not reliably produce grounded, actionable meanings for each individual sentence.

General computational models of cognition called *cognitive architectures* have been developed to model a wide variety of psychological and neural theories of cognition. A few prominent models have emerged which have been implemented and applied in a large number of research projects. Recently three of these architectures have been combined into an abstract model called the *Common Model of Cognition* (CMC; Laird, Lebiere, & Rosenbloom, 2017).[1] The CMC has also been related to the structure of the brain (Steine-Hanson, Koh, & Stocco, 2018). Previous work on modeling language comprehension with cognitive architectures (Lewis & Vasishth, 2005; Lewis, 1993) does not include full comprehension or immediate interpretation. Our implementation uses the Soar cognitive architecture (Laird, 2012).

Researchers have measured brain activity while subjects listen to readings of naturalistic language (J. Brennan, 2016; J. R. Brennan & Hale, 2019; Schwartz & Mitchell, 2019). These studies compare brain measurements to measurements on a variety of computational models of language processing.

[1] The term *Common Model of Cognition* has replaced the term *Standard Model of the Mind* used in the original referenced paper.

However, these models are rather simple, not doing grounded, end-to-end comprehension.

The research reported here addresses the lack of a complete and detailed computational theory of how human sentence comprehension works.

We propose a computational theory of sentence comprehension that attempts to explain both the functional capabilities of human sentence comprehension and the process humans use, implemented using domain-general cognitive abilities. At the core of this theory is the idea that human comprehension involves actively constructing the meaning of each sentence incrementally in small steps, where each step integrates an additional element of form-meaning correspondence (Goldberg, 2013) into the developing meaning representation (Christiansen & Chater, 2016), and immediately grounds that element to the agent's knowledge of the world. Figure 1 is a diagram of comprehending a simple sentence, which we discuss in detail below.

In this paper we explore high-level characteristics of human language comprehension which serve as constraints to define our computational theory. We explain our theory of how linguistic knowledge can be represented and operated on in a manner that fits these constraints, and how it fits with existing theories of sentence comprehension, then describe a concrete implementation of the theory. Finally we discuss how this theory may contribute to further research on human language processing, both at a symbolic level of cognitive processing and in terms of how such a computational model might be implemented in the brain.

Characteristics of Human Comprehension

We focus on three levels of human sentence comprehension: its *functional* capabilities as seen from outside, characteristics of its *processing* on the inside, and the *underlying mechanisms* by which the processing is accomplished.

Functionality

From a functional perspective, the most important characteristic is what we call *end-to-end comprehension*. By this we mean that the comprehension system produces, for each individual input sentence, an internal meaning representation that the human agent can act upon. The resulting action might be responding to the speaker, adding new knowledge, performing some action in the world, or any combination of these things. In addition, the production of this actionable meaning must be performed in *real time*, meaning fast enough that the hearer can react in an appropriate way in the real world.

Another key functional aspect of human comprehension is its *generality*. This means that the linguistic knowledge that a person has can be composed in many ways to allow understanding of many new sentences that have never been heard or seen before. Often this characteristic of human language ability is called *creativity* (Chomsky, 1965) or *productivity*, but these terms imply language production. We will use the term *generality* to refer to the corresponding aspect of language comprehension.

Processing

From a processing perspective, human comprehension is characterized by *immediate interpretation*. Both psychological (Tanenhaus et al., 1995) and neuroscience (Hagoort, 2019) experiments show that humans do not wait until the end of a sentence to construct meaning. Rather, for each input word all available knowledge is used to construct as much meaning as possible, including grounding the meaning representation, when possible, to the agent's current perception, action, and world knowledge.

Immediate interpretation is necessary (Christiansen & Chater, 2016), yet it imposes significant constraints on the processing algorithm. Human language contains many local ambiguities, so when there are multiple options available, the processing system must make its best guess at each point. At times, it will turn out that a choice is not consistent with subsequent input, so that the system must make a correction or local repair (Lewis, 1993) to its ongoing representation of the meaning of the sentence. Both the initial decision about which option to commit to, and any later repair, provide *local ambiguity resolution*. This capability is essential to avoid an exponential increase in ambiguity.

Underlying Mechanisms

There has been much controversy over the years about whether the human brain has innate language ability or whether language processing is a skill using domain-general cognitive abilities. Here we assume the latter, and explore whether humanlike comprehension is possible with the tools of *general cognitive abilities*.

To accomplish this, we need a theory of what those cognitive abilities are and how they work. Cognitive psychology has developed theories of these abilities, and a substantial amount of research in cognitive neuroscience has explored how these abilities might be implemented in the brain. Often, however, that work does not directly provide computational mechanisms for cognition that can be implemented on a computer.

The CMC and the architectures it is based on include several major mechanisms based on psychological theory, including: working memory, procedural memory, long-term declarative memory, interfaces to perception and action, and a temporal processing cycle (Laird et al., 2017). A challenge for our theory is to see whether and how these domain-general mechanisms can be used to comprehend language in a humanlike way.

One consequence of using general mechanisms based on human cognition is that humans have limitations on memory and processing capacities, leading to *bounded rationality* (Simon, 1996), which places limits on the kinds of algorithms that can be performed, especially in real time. A successful computational theory of human comprehension must work within these constraints.

A Computational Theory of Comprehension

A computational theory of human comprehension must satisfy the constraints outlined above. For each sentence it must produce a meaning representation which can be acted upon. This meaning representation must be made up of small elements that can be composed in many ways. These elements must be integrated and grounded immediately, and any resulting local ambiguities must be resolved. All this must be accomplished using the abilities of, and within the limitations of, a model of human cognition. The theory can be summarized in what we call *the incremental meaning construction hypothesis*:

> *The meaning of a sentence is constructed in small increments in a repeating cycle in which each new pattern of form evokes a new element of meaning, and these elements are grounded and composed into larger and larger elements until the full sentence meaning has been constructed.*

To clarify this idea, consider the example in Figure 1. At the bottom we see a simple sentence that might be addressed to a robot. At the top we see a message to the operational part of the robot telling it to perform its "pick up" action on an indicated object it sees in the environment. In between are eight meaning elements that have been evoked by the input word sequence and grounded to the agent's long-term knowledge and current perception. These elements are constructed in the order shown by the numbers.

Functionality

Our theory proposes that an active process constructs, word by word and piece by piece, the meaning of each sentence. Each piece is an instance of an abstract pairing of form and meaning called a *construction* (Goldberg, 2013). A construction relates a particular linguistic form, which may be a word or a larger syntactic structure, to the meaning it evokes. Constructions fall into two major categories: those that represent lexical items and those that represent compositions of one or more simpler constructions into larger units.

At each step a single construction is selected from a large inventory of possibilities, instantiated and elaborated, integrated into a data structure representing the current *comprehension state*, and then grounded as possible to the agent's short- and long-term knowledge of the world. Occasionally new input shows that a selected construction was not the correct one, and a *local repair* to the comprehension state is performed (Lewis, 1993). When the complete meaning of a sentence has been constructed, a final process called *message formatting* extracts the essential elements of grounded meaning and assembles an *actionable message* that the surrounding agent then acts upon. This process provides *end-to-end comprehension*, and the composition of meaning from small pieces provides *generality*.

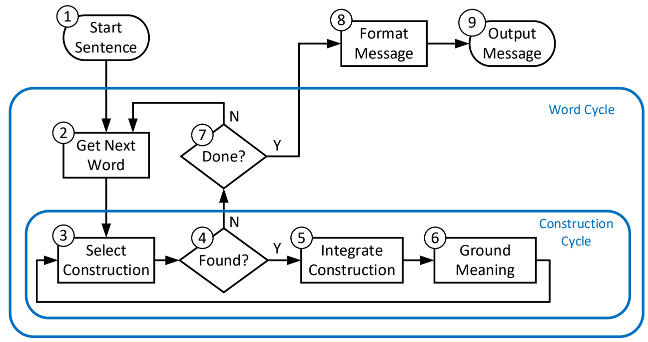

Figure 2: Sequence of operations

Processing

How is the construction process implemented? Figure 2 gives a sketch of the sequence of operations needed to do the incremental construction of the meaning of a sentence. Processing begins in (1) and proceeds until the final message is put out in (9). Blocks (2) through (7) implement two nested loops that build up the meaning representation incrementally. The *word cycle* does word-by-word processing, attending to a new word in (2) and repeating the *construction cycle* in (3) through (6) as often as necessary. Each word cycle will select one lexical construction and possibly one or more composite constructions. When (4) finds that no more constructions apply, this word cycle is complete. If (7) finds that the end of the sentence has been reached, it proceeds to (8) to complete the processing. Otherwise, the next word is attended to, and the cycles continue.

Each form-meaning construction selected in (3) is integrated into the comprehension state in (5) and grounded in (6). The first selection in (3) for each word cycle will be for a lexical construction matching the current input word. After that the selection in (3) looks for the best composite construction to match the current state. When no matching composition is found, the process goes on to (7). After the last word has been processed, the message formatting operation in (8) summarizes the whole comprehension that has been built as an internal message the agent can act on.

Grounding is a key step in this process. When a sentence refers to an object that is currently within the agent's world model, the comprehension process identifies this object, and the message formatting need only pass along a unique internal identifier of the object to the rest of the agent, as Figure 1 illustrates. When an object is not currently in the agent's model of the world, than a description of the object's properties must be passed along to be grounded to perception at some later time.

As the meaning of each construction is integrated into the comprehension state and grounded, we have *immediate interpretation*, and selection and local repairs provide *local ambiguity resolution*. The end result is to construct and ground the meaning representation incrementally. Typically incremental comprehension is thought of as processing the input one word at a time (Lewis & Vasishth, 2005; Lewis, 1993), as implemented by the word cycle. Our theory provides the finer-grained concept of a construction cycle, as shown in (3) through (6) of Figure 2. With all this working in a model of comprehension, we have satisfied our constraints to model humanlike comprehension.

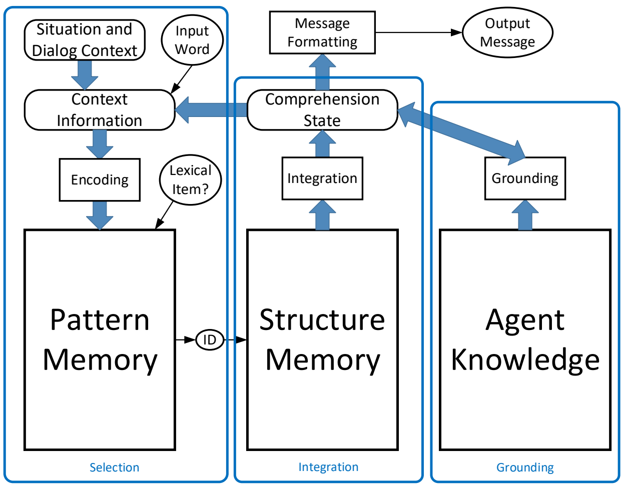

Figure 3: Phases of a construction cycle

Underlying Mechanisms

Figure 3 gives a sketch of mechanisms used for the three major phases of each construction cycle. The *Selection* phase encodes all the contextual information the agent has, including the dialog context, the current input word, and the current comprehension state, as input to a Pattern Memory, which chooses which known construction best fits the current state. This memory operates in two modes: one to match a lexical pattern and one to match a composite pattern in the current comprehension state. It retrieves a unique identifier for the construction that best matches the current context.

The *Integration* phase uses the retrieved identifier as a cue to access the Structure Memory, which retrieves the complex data structures needed to integrate the construction into the comprehension state and evoke its meaning representation. The *Grounding* phase then connects the new state to the agent's long-term knowledge and current perception. Thus data flows around the loop of these memories once for each construction cycle, gradually building up the complete meaning of the sentence. When the test in (7) determines that the end of the sentence has been reached, the Format Message operation assembles the message that the agent will act on.

An Implementation of the Theory

To test whether the theory we have outlined is computationally feasible, and to explore the implications of the theory, we have written a computer program we call *Lucia* (Lindes, 2018; Lindes & Laird, 2016) using the Soar cognitive architecture (Laird, 2012), which implements the theory in detail.

Functionality

Linguistic knowledge in Lucia is defined in a grammar built on principles of construction grammar and implemented using a specific theory called Embodied Construction Grammar (ECG; Bergen & Chang, 2013). This approach is grounded in decades of research in cognitive linguistics, and provides a formalism for representing meanings and the increments of form-meaning pairing that are needed to compose sentence meanings incrementally while giving much generality. Figure 1 is a small example of how meanings are represented.

The Lucia implementation is embodied in a larger robotic agent capable of acting in the physical world (Lindes, Mininger, Kirk, & Laird, 2017). The Agent Knowledge shown in Figure 3 is shared between the comprehension system and the operational part of the agent that knows how to act on the messages Lucia produces. The entire agent is capable of using interactive instruction with naturalistic language to learn new tasks and perform them in the world. This embodiment gives us a way to confirm that the theory actually provides comprehension accurate enough for an agent to act on correctly. Lucia correctly comprehends several hundred naturalistic sentences for teaching games, puzzles, and robot navigation tasks, showing that it achieves *end-to-end comprehension*.

A key feature of human comprehension is its *generality*: its ability to correctly comprehend many sentences that have never been seen or heard before. Lucia will never match the generality of human comprehension, but it can correctly comprehend several orders of magnitude more sentences than were used to develop it.

Processing

Lucia uses a novel incremental comprehension algorithm along the lines shown in Figures 2 and 3. It implements the construction cycle by building one construction at a time, integrating each with the comprehension state, and grounding it to the agent's knowledge. It includes techniques for resolving ambiguities, both at the points where new constructions are selected and using local repairs after the fact when needed.

Inspection of internal operation of the Lucia algorithm shows that it does indeed do *immediate interpretation* and *local ambiguity resolution* (Lindes & Laird, 2017). It provides additional psycholinguistic plausibility by corresponding to the *chunk-and-pass* theory (Christiansen & Chater, 2016), which suggests how language comprehension can work within the limits imposed by human working memory. It does not yet implement explicit working memory limits, but implicitly its selection of which construction to instantiate next uses only the top few levels of the comprehension state.

Underlying Mechanisms

Lucia is implemented entirely with the *general cognitive abilities* represented by the memories and processing of the CMC and the Soar cognitive architecture. Key to implementing an algorithm of this sort in a cognitive architecture is how the declarative and processing knowledge shown conceptually in Figures 2 and 3 is represented in the actual memories of the architecture.

Lucia encodes the linguistic knowledge for the Pattern Memory and Structure Memory in Figure 3 as production rules in Soar's procedural memory. These rules are generated automatically by a program that translates them from the ECG representation (Bryant, 2008) of the agent's grammar. The Comprehension State and Context Information in Figure 3 are both stored in Soar's working memory. The short-term part of the Agent Knowledge is also stored in working memory, with the long-term part in Soar's semantic memory. The Encoder, Integration, Grounding, and Message Formatting modules are built from additional hand-coded production rules.

Using production rules for linguistic knowledge provides a mechanism that allows for rapid processing to simulate real-time comprehension as a developed skill. Each construction has a production rule that recognizes its form pattern, and these collectively implement the Pattern Memory. Then additional construction-specific rules fire to provide the retrieval of information from the Structure Memory. Measurements of its speed of processing in simulated *real time*, based on a cognitive cycle taking about 50ms (Laird, 2012), gives a comprehension rate of 138 words/minute, which is comparable to human speaking rates.

Future Work and Implications

Lucia currently has several weaknesses with respect to accurately modeling human comprehension. Like the CMC and Soar, it includes no explicit model of the limitations of human working memory. Neither does it provide a non-deliberative mechanism, similar to human priming effects, to use dialog and situation context to select among multiple senses of a word. We are exploring representing constructions in Soar's semantic memory and using activation schemes, including recency, frequency, spread, fan, and attention, to support context-biased retrieval.

Psycholinguistic and neuroscience evidence suggests that the human brain does some form of prediction of what will come next during language comprehension (Kuperberg & Jaeger, 2016). Prediction, in both the brain and our model, can speed up processing. We have not yet implemented this aspect in Lucia, but we have a plan to do so as we further develop the system.

Brain measurements show one key attribute that is challenging to our model. There are properties of EEG signals during comprehension called the N400 and P600 (Delogu, Brouwer, & Crocker, 2019), which occur roughly 400ms and 600ms, respectively, after the onset of a word. However, new words are coming in every 200-300ms or so. Thus there appears to be processing related to a word that goes on after another word or two have appeared in the input. This is hard to justify based on our model, where each word, and each construction, is processed to completion before going on to the next.

One way to resolve this issue would be to develop a mechanism in the architecture to allow overlap between the processing of multiple words. This would be a major departure from most current cognitive architecture theory. Another option is to implement prediction: a prediction initiated by one word would later be confirmed or disconfirmed by a later word, producing processing correlated to the original word but happening after other words have been at least partially processed. This fits at least one theory of what causes these signals in humans (Bornkessel-Schlesewsky & Schlesewsky, 2019).

An essential feature of human language processing is the acquisition of linguistic knowledge. At the present time the grammar in Lucia is built by hand by adding to the grammar files written in the ECG formalism (Bryant, 2008). We do, however, have a nascent theory of how humanlike acquisition could be added to our model. Our grammar has been developed in small increments, each increment deriving from a piece of knowledge that was missing to process a particular sentence. These knowledge increments then have the generality to apply to many more sentences. Our theory is that human acquisition also happens in similar small increments with generalization, corresponding to research on language acquisition (Tomasello, 2003).

Our model specifies details of the use of different memory systems and the time course of accessing them. This has the potential of suggesting new ways to study language processing in the brain. Neuroscientists have gathered detailed data on the spatial and temporal patterns of neural activation in the brain during language processing (Kemmerer, 2015). One theory of the neurobiology of language (Hagoort, 2019) emphasizes the "immediacy principle" and discusses elementary linguistic units (ELUs), which correspond to our constructions, and elementary linguistic operations (ELOs), which correspond to the operations in our comprehension algorithm. Another paper (Scott, 2019) argues that lexical and phrasal processing take place in different areas of the brain, corresponding to the distinction between these two types of processing in our Pattern Memory. It should be possible to search for evidence of the construction cycle (Figure 2) and its phases and use of memories (Figure 3) in brain data.

The theory presented here and its implementation are still in their infancy. Nevertheless, they provide a working prototype that can be used to explore a wide variety of questions in human language processing, supporting Pylkkänen's (2019) argument that "More effort should be directed toward developing computational models of incremental semantic composition."

Acknowledgments

The work described here was supported by the Air Force Office of Scientific Research under Grant Number FA9550-15-1-0157. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the AFOSR or the U.S. Government.

References

Bergen, B. K., & Chang, N. C. (2013). Embodied Construction Grammar. In T. Hoffmann & G. Trousdale (Eds.), *The Oxford Handbook of Construction Grammar* (pp. 168–190). Oxford: The Oxford University Press.

Bornkessel-Schlesewsky, I., & Schlesewsky, M. (2019). Toward a Neurobiologically Plausible Model of Language-Related, Negative Event-Related Potentials. *Frontiers in Psychology*, *10*(FEB).

Brennan, J. (2016). Naturalistic Sentence Comprehension in the Brain. *Language and Linguistics Compass*, *10*(7), 299–313.

Brennan, J. R., & Hale, J. T. (2019). Hierarchical structure guides rapid linguistic predictions during naturalistic listening. *PLoS ONE*, *14*(1).

Bryant, J. E. (2008). *Best-fit Constructional Analysis*. PhD dissertation, University of California at Berkeley.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: The MIT Press.

Christiansen, M. H., & Chater, N. (2016). The Now-or-Never bottleneck: A fundamental constraint on language. *Behavioral and Brain Sciences*, 1–72.

Delogu, F., Brouwer, H., & Crocker, M. W. (2019). Event-related potentials index lexical retrieval (N400) and integration (P600) during language comprehension. *Brain and Cognition*, *135*.

Eppe, M., Trott, S., Raghuram, V., Feldman, J., & Janin, A. (2016). Application-Independent and Integration-Friendly Natural Language Understanding. *GCAI 2016. 2nd Global Conference on Artificial Intelligence Application-Independent*, *41*, 340–352.

Goldberg, A. E. (2013). Constructionist Approaches. In T. Hoffmann & G. Trousdale (Eds.), *The Oxford Handbook of Construction Grammar* (pp. 15–31). Oxford: The Oxford University Press.

Hagoort, P. (2019). The neurobiology of language beyond single-word processing. *Science*, *366*(6461), 55–58.

Hoffman, T., & Trousdale, G. (Eds.). (2013). *The Oxford Handbook of Construction Grammar*. Oxford: The Oxford University Press.

Kemmerer, D. (2015). *Cognitive Neuroscience of Language*. New York, NY: Psychology Press.

Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, *31*(1), 32–59.

Laird, J. E. (2012). *The Soar Cognitive Architecture*. Cambridge, MA: The MIT Press.

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*, *38*(4), 13.

Lewis, Richard L, & Vasishth, S. (2005). An Activation-Based Model of Sentence Processing as Skilled Memory Retrieval. *Cognitive Science*, *29*, 375–419.

Lewis, Richard Lawrence. (1993). *An Architectually-based Theory of Human Sentence Comprehension*. PhD dissertation, Carnegie Mellon University.

Lindes, P. (2018). The Common Model of Cognition and humanlike language comprehension. *Procedia Computer Science*, *145*, 765–772.

Lindes, P., & Laird, J. E. (2016). Toward Integrating Cognitive Linguistics and Cognitive Language Processing. *Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016)*. Retrieved from http://acs.ist.psu.edu/iccm2016/proceedings/lindes2016iccm.pdf

Lindes, P., & Laird, J. E. (2017). Ambiguity Resolution in a Cognitive Model of Language Comprehension. *Proceedings of the 15th International Conference on Cognitive Modeling (ICCM 2017)*.

Lindes, P., Mininger, A., Kirk, J. R., & Laird, J. E. (2017). Grounding Language for Interactive Task Learning. *Proceedings of the First Workshop on Language Grounding for Robotics*, 1–9.

Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., & Wu, J. (2019). The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences From Natural Supervision. *ICLR 2019*, 1–28. Retrieved from http://arxiv.org/abs/1904.12584

Pylkkänen, L. (2019). The neural basis of combinatory syntax and semantics. *Science*, *366*(6461), 62–66.

Schwartz, D., & Mitchell, T. (2019). *Understanding language-elicited EEG data by predicting it from a fine-tuned language model*. Retrieved from http://arxiv.org/abs/1904.01548

Scott, S. K. (2019). From speech and talkers to the social world: The neural processing of human spoken language. *Science*, *366*(6461), 58–62.

Simon, H. A. (1996). *The Sciences of the Artificial* (Third Edit). Cambridge, MA: The MIT Press.

Steels, L. (2013). Fluid Construction Grammar. In T. Hoffmann & G. Trousdale (Eds.), *The Oxford Handbook of Construction Grammar* (pp. 153–167).

Steine-Hanson, Z., Koh, N., & Stocco, A. (2018). Refining the Common Model of Cognition Through Large Neuroscience Data. *Procedia Computer Science*, *145*, 813–820.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K., & Sedivy, J. C. (1995). Integration of Visual and Linguistic Information in Spoken Language Comprehension. *Science*, *268*(5217), 1632.

Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA: Harvard University Press.

Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent Trends in Deep Learning Based Natural Language Processing [Review Article]. *IEEE Computational Intelligence Magazine*, *13*(August), 55–75.