

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Toward a Functional Model of Human Language Processing

Permalink

<https://escholarship.org/uc/item/8kr28525>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN

1069-7977

Authors

Ball, Jerry
Freiman, Mary
Rodgers, Stuart
et al.

Publication Date

2010

Peer reviewed

Toward a Functional Model of Human Language Processing

Jerry Ball¹, Mary Freiman², Stuart Rodgers³ & Christopher Myers¹

Air Force Research Laboratory¹, L3 Communications², AGS TechNet³

Jerry.Ball@mesa.afmc.af.mil, Mary.Freiman@mesa.afmc.af.mil, Stu@agstechnet.com,

Christopher.Myers@mesa.afmc.af.mil

Abstract

This paper describes a computational cognitive model of human language processing under development in the ACT-R cognitive architecture. The paper begins with the context for the research, followed by a discussion of the primary theoretical and modeling commitments. The main theoretical commitment is to develop a language model which is at once functional and cognitively plausible. The paper continues with a description of the word recognition subcomponent of the language model which uses a perceptual span and ACT-R's spreading activation mechanism to activate and select the lexical unit that most closely matches the perceptual input. Next we present a description of the linguistic structure building component of the model which combines parallel, probabilistic processing with serial, pseudo-deterministic processing, including a non-monotonic context accommodation mechanism. A description of the mapping of linguistic representations into a situation model, follows. The paper concludes with a summary and conclusions.

Keywords: human language processing (HLP); functional; cognitively plausible; pseudo-deterministic.

Introduction

The capability to model the cognitive processes associated with language is a long sought-after goal of cognitive science. Computational cognitive process models help researchers to not only understand language processes in their own right, but to determine how they affect and interact with other cognitive processes (e.g., reasoning, decision-making, situation assessment, etc.). Scaled-up versions of these models also support the development of cognitive agents with communicative capabilities based on human linguistic processes (Ball et al., 2009; Douglass, Ball & Rodgers, 2009). In this paper we present a “snapshot” of a functional language comprehension model under development within the ACT-R architecture (Anderson, 2007). The model implements a referential and relational theory of human language processing (Ball, 2007; Ball, Heiberg & Silber, 2007) within ACT-R¹.

A key commitment of the language comprehension research is development of a model which is at once cognitively plausible and functional. We believe that adherence to well-established cognitive constraints will

facilitate the development of functional models by pushing development in directions that are more likely to be successful. There are short-term costs associated with adherence to cognitive constraints; however, we have already realized longer-term benefits. For example, the integration of a word recognition capability with ACT-R's perceptual system and higher-level linguistic processing has facilitated the recognition and processing of multi-word expressions and multi-unit words in ways that are not available to systems with separate word tokenizing and part of speech tagging processes. Using an available tokenizer and part of speech tagger would have initially facilitated development, but the cognitive implausibility of using staged tokenizing and part of speech tagging led us to reject this approach. The benefits that we have realized as a result of this decision are described below.

Theoretical & Modeling Commitments

There is extensive psycholinguistic evidence that human language processing is incremental and interactive (Gibson & Pearlmutter, 1998; Altmann, 1998; Tanenhaus et al., 1995; Altmann & Steedman, 1988). Garden-path effects, although infrequent, strongly suggest that processing is essentially serial at the level of phrasal and clausal analysis (Bever, 1970). Lower level processes of word recognition suggest parallel, activation-based processing mechanisms (McClelland & Rumelhart, 1981; Paap et al., 1982). Summarizing the psycholinguistic evidence, Altmann & Mirkovic (2009, p. 605) claim “The view we are left with is a comprehension system that is ‘maximally incremental’; it develops the fullest interpretation of a sentence fragment at each moment of the fragment’s unfolding”.

These cognitive constraints legislate against staged analysis models. All levels of analysis must at least be highly pipelined together, if not, in addition, allowing feedback from higher to lower levels. They also suggest the need for hybrid systems which incorporate a mixture of parallel and serial mechanisms, with lower levels of processing being primarily parallel, probabilistic and interactive, while higher levels of analysis are primarily serial, deterministic and incremental.

To adhere to and take advantage of these cognitive constraints, we have developed a *pseudo-deterministic* human language processing model—i.e. a model that presents the appearance and efficiency of serial, deterministic processing, but uses a non-monotonic context

¹ At the time of publication the model contained 6,395 declarative memory elements and 548 production rules which cover a broad range of grammatical constructions.

accommodation mechanism and relies on lower level parallel mechanisms to deal with the ambiguity that makes true deterministic processing impossible. This model makes use of the architectural mechanisms in ACT-R that are most compatible with incremental and interactive processing. For example, parallel, probabilistic processing taps into ACT-R's declarative memory (DM) and parallel spreading activation mechanism, with ACT-R's DM retrieval mechanism supporting probabilistic selection—without inhibition between competing alternatives as is typical of connectionist models (cf. Vosse & Kempen, 2000). Serial, incremental processing is based on ACT-R's procedural memory which is instantiated as a production system. ACT-R at once constrains the computational implementation and provides the basic mechanisms on which the model relies. Other than adding a collection of buffers to support language processing by retaining the partial products of retrieval and structure building, and improving the perceptual processing in ACT-R, the computational implementation does not add any language-specific mechanisms. In the following sections we discuss important subcomponents of the model, such as how the model recognizes words, builds linguistic representations, and maps linguistic representations to a situation representation.

Reading & Word Recognition

A functional language model must deal with the linguistic input as is. In an experiment involving human subjects communicating via text chat (cf. Ball, et al., 2009), we collected a text chat corpus that is riddled with variability in word forms—e.g., misspellings like “altitde”, abbreviations like “alt.”, and concatenations like “speedrestriction” and “speed=200-500”. For competent readers, misspelled words activate the intended lexical items because they contain many of the same letters and trigrams (Perea & Lupker, 2003). Further, all the letters of a word can be transposed, yet still prime the intended word (Guerrera 2004). Key requirements of a functional language model are the ability to handle variability and misspellings in input forms, the ability to separate perceptually conjoined units (e.g. separating punctuation from words as in “He went.”, but not “etc.”); separating concatenated words, and the ability to recognize multi-word expressions (e.g. “speed up”) and multi-unit words (e.g. “ACT-R”, “a priori”).

To satisfy these requirements, the model includes a word recognition subcomponent that uses ACT-R's spreading activation mechanism combined with a multi-word perceptual span to influence lexical item retrieval. It is assumed that word recognition involves mapping orthographic input directly into DM representations without recourse to phonetic processing (although a phonetic mapping is not precluded). The model does not treat each word as a sum of its parts, ignoring the complete form altogether. Rather, if the text input as a whole does not match, and thereby activate an item in the lexicon, the closest match can be retrieved based on the cues that do match, such as letters, word-length, and trigrams.

In the model's DM, word chunks have slots for letters, word-length, and trigrams. Multi-unit words and multi-word expressions have this information for all of the constituent units. Text input is distilled into this information by the model and put into buffers to spread activation to words in DM containing matching information. The activation mechanism allows the model to retrieve words from DM that are not an exact match to the input. Letters and trigrams in the text input increase the activation of word chunks containing those letters and trigrams in the mental lexicon. The most highly activated word chunk, which need not be an exact match to the input, is retrieved. These processes and encodings are based on the Interactive Activation model of word recognition (McClelland and Rumelhart 1981), with the addition of trigrams based on “letter triples” (Seidenberg and McClelland, 1989).

Besides breaking words into letters and trigrams, we modified the ACT-R architecture to better interpret multi-unit words and multi-word expressions. By default, ACT-R splits input text into perceptual units based on spaces and punctuation—even word internal punctuation, where “ACT-R” becomes “ACT” “.” “R”—and processes each perceptual unit separately. We replaced this behavior with a perceptual span that is based on human reading span data and a multi-level splitting of the input within the perceptual span into larger and smaller perceptual units which spread activation in parallel. We also added multi-word expression chunks and multi-unit lexical chunks to DM. The overall effect is a significant reduction in the number of DM retrievals per space and punctuation delimited input. Words with internal punctuation and multi-word expressions can now be retrieved as a single perceptual unit despite their internal structure (Freiman & Ball, submitted).

The new perceptual span is considerably larger than ACT-R's punctuation and space delimited span. There is a great deal of evidence that the perceptual span of adult readers is about 14-15 letters to the right of fixation (McConkie & Rayner, 1975; Rayner, 1986). We implemented a span of up to twelve letters, with the greatest amount of activation spreading from the first few letters of the span and decreasing toward the end of the span. Just as for adult readers, information to the right of fixation is obtained when the next word is predictable from the preceding text (see Rayner 1975; and Binder, Pollatsek, & Rayner, 1999).

Within the context of a functional language model—i.e. one that must interpret and act on the linguistic input, we are also attempting to model adult human reading rates (Freiman & Ball, submitted). Adult humans read at a phenomenal rate of 200-300 (space delimited) words per minute (Carver, 1973a; 1973b). The ACT-R architecture supports the timing of cognitive processes down to the msec level. The real-time it takes for a model to run can also be measured. Although we have not yet succeeded in achieving adult reading rates, we have improved the reading rate of the model significantly in both cognitive and real-time: 143 words per minute in ACT-R cognitive time (important for

cognitive plausibility); and 249 words per minute in real-time on a single-core, 2.1 GHz Windows Vista machine with 2 gigabytes of RAM (important for a functional model). Ultimately, we believe that achieving adult reading rates hinges on minimizing the amount of structure building and maximizing the average size of linguistic units which are retrieved. We are pursuing mechanisms and representations that will make this possible.

Building Linguistic Representations

The word recognition subcomponent typically delivers a lexical item categorized for part of speech to the higher level component that builds linguistic representations of referential and relational meaning. For example, consider the processing of “the pilot”. The processing of “the” leads to its identification as a determiner via retrieval from DM. Selection of this lexical item is based on the probabilistic, context-sensitive mechanism discussed in the previous section. The subsequent processing of the determiner “the” leads to the projection or construction of a nominal construction. The processing of the word “pilot” in the context of the preceding word “the” and the projected nominal leads to retrieval of a DM chunk identifying “pilot” as a noun. The noun “pilot” is then integrated as the head of the nominal projected during the processing of “the”.

Similar parallel, probabilistic mechanisms operate at the phrasal and clausal level, selecting between competing phrasal and clausal alternatives, and potentially interacting with lower level probabilistic mechanisms. As an example of this potential interaction, consider the processing of personal pronouns like “he” and “it”. At the lexical level, these words are categorized as pronouns, but they are also closely associated with the nominal phrasal category since they typically function as the head of a complete nominal. Processing personal pronouns may involve their recognition as pronouns followed by projection of a nominal phrase from the pronoun, but it may also be that the perceptual form can directly lead to retrieval of a nominal phrase, without the intermediate step of identifying the word as a pronoun. The word recognition component, which prefers larger and higher level units, may deliver a pre-compiled nominal unit corresponding to the pronoun, rather than a lexical unit to the higher level construction process, blurring the distinction between lexical and phrasal units. The determiner “the” may behave similarly, resulting in direct retrieval of a nominal with an empty head, without the intermediate step of identifying “the” as a determiner.

The parallel, probabilistic mechanism which is capable of retrieving existing phrasal and clausal representations as well as lexical units, competes with a mechanism which builds novel representations. DM retrieval has priority over this alternative construction mechanism. However, lexical units are more likely to be available for retrieval than phrasal and clausal representations. Further, the parallel, probabilistic mechanism is not capable of building any structure—building structure is the function of the serial construction mechanism.

There are two basic ways of building structure: 1) integration of the current linguistic unit into an existing representation which contains an expectation for the linguistic unit (i.e. substitution), and 2) projection or construction of a novel representation coupled with integration of the current linguistic unit into the novel representation. For example, the processing of the word “pilots” recognized as a plural noun by the word recognition component can lead to projection of a nominal and integration of “pilots” as the head of the nominal. On the other hand, if “the” has already projected a nominal and set up the expectation for a head to occur, the processing of “pilots” can lead to its integration as the head of the nominal projected by “the”.

The structure building mechanism is incremental in that it executes a sequence of productions that determine how to integrate the current linguistic unit into an existing representation and/or which kind of higher level linguistic unit to project. These productions execute one at a time within the ACT-R architecture which incorporates a serial bottleneck for production execution. Although supported by extensive empirical evidence, the serial production execution bottleneck is a characteristic of ACT-R that distinguishes it from other production system architectures which support parallel production execution.

The structure building mechanism uses all available information in deciding how to integrate the current linguistic input into the evolving representation. Although the parallel, probabilistic mechanism considers multiple alternatives in parallel, the output of this parallel mechanism is a single linguistic unit and the result of structure building is also a single representation. The structure building mechanism operates in a *pseudo-deterministic* manner. It is deterministic in that it builds a single representation which is assumed to be correct, but it relies on the parallel, probabilistic mechanism to provide the inputs to this structure building mechanism. In addition, structure building is subject to a mechanism of context accommodation capable of making modest adjustments to the evolving representation (Ball, 2010a). Although context accommodation does not involve backtracking or reanalysis, it is not, strictly speaking, deterministic, since it can modify an existing representation and is therefore non-monotonic. For example, in the processing of the expression “the altitude restriction”, when the word “altitude” is processed, it can be integrated as the head of the nominal projected by “the”. But when “restriction” is subsequently processed, the context accommodation mechanism can adjust the representation, shifting “altitude” into a modifying function so that “restriction” can function as the head. This context accommodation capability can apply iteratively as in the processing of “the pressure valve adjustment screw” where “screw” is the ultimate head of the nominal, but “pressure”, “valve” and “adjustment” are all incrementally integrated as the head prior to the processing of “screw”. Note that at the end of processing it appears that “pressure”, “valve” and “adjustment” were treated as

modifiers all along, giving the appearance that these alternatives were carried along in parallel with their treatment as heads.

Context accommodation uses the full available context to make modest adjustments to the evolving representation or to construe the current input in a way that allows for its integration into the representation. As an example of construal, the verb “kick” is construed as an object and functions as the head of a nominal when it occurs in the context of “the”, as in “the kick”. Function overriding and function shifting are two additional mechanisms of context accommodation. We have already seen an example of function shifting (e.g. “the altitude restriction”). In the processing of “no altitude or airspeed restrictions”, the conjoined head “altitude or airspeed” can override the initial treatment of “altitude” as the head of the nominal, with the subsequent shifting of “altitude and airspeed” into a modifying function during the processing of “restrictions”. At a lower level, there are accommodation mechanisms for handling conflicts in the grammatical features associated with various lexical items. For example, the grammatical feature *definite* is associated with “the” and the grammatical feature *indefinite* is associated with “pilots”. In “the pilots”, the *definite* feature of “the” blocks the *indefinite* feature of “pilots” from projecting to the nominal. See Ball (2010b) for more details.

Context accommodation need not be computationally expensive—a single production may effect the accommodation, just as a single production may effect integration without accommodation. In this respect, context accommodation is not a reanalysis mechanism that disrupts normal processing—it is part and parcel of normal processing. Reanalysis mechanisms need only kick in when context accommodation fails and larger adjustment is needed. The mechanism of context accommodation is most closely related to the limited repair parsing of Lewis (1998). Context accommodation may be viewed as a very modest form of repair. According to Lewis (1998, p. 262) “The putative theoretical advantage of repair parsers depends in large part on finding simple candidate repair operations”. The mechanism of context accommodation provides evidence for this theoretical advantage.

Overall, the highly interactive, parallel, probabilistic mechanism for selecting between competing alternatives combines with the incremental, serial construction and context accommodation mechanisms to provide an efficient, pseudo-deterministic language processing capability.

Mapping into the Situation Model

Although we borrow the term (cf. Zwann & Radvansky, 1998), we define *situation model* as a domain-specific mental representation of a set of objects, actions, events, and relationships related to a task, sufficient for reasoning about a set of actions within that task. The situation model is separate from the model’s world knowledge but is related to and affected by world knowledge.

The situation model is implemented in three main subcomponents: the ACT-R module definition, a set of domain general production rules, and a set of domain specific production rules. The module is instantiated like other ACT-R modules (Anderson, 2007), and includes the module buffers and handlers for module requests and queries.

The main situation buffers are: sm-subject-context, sm-related-object-context, sm-sit-context, sm-action-context, sm-event-context, and sm-prior-attention. They are named and designed to reflect the semantics of the represented situations. The buffers will contain chunks representing the objects, actions, events, and relationships discussed or encountered in the task environment. The top level chunk types were based upon the Suggested Upper Merged Ontology (SUMO) (Niles and Pease, 2001) and are: Action, Attribute, Concept, Event, Object, Relation, and Situation. All entities represented in the situation model will be sub-typed from one of these top level chunk types. Because the situations being represented in our model may span multiple sentences, the contents of the sm-subject-context buffer will frequently not equate to the subject of an individually processed sentence. Rather, the contents of the sm-subject-context buffer should be thought of as the central topic or theme of the discourse at an individual moment. The situation chunk-type and its sub-types can be thought of as instances of *schemata* or structures for mental models of stereotypical situations (Alba, 1983). In our implementation, the situation chunk contains the relevant gist of the situation, where the “gist” can be thought of as an index to a specific category of situation.

It is the responsibility of the modeler to define any needed specific chunk subtypes. Because ACT-R’s chunk inheritance mechanism does not permit inheritance from multiple supertypes, it is expected that there will be some redundancy in the definitions of the chunk subtype hierarchy. While this redundancy will create some inefficiency in the type hierarchy design, it should not preclude the modeling of necessary elements.

The domain general productions manage the relationships between elements within each individual situation. For instance, in a situation involving an uninhabited air vehicle altitude restriction for a reconnaissance waypoint, a situation chunk would contain a subject slot and a related object slot. The subject slot value would refer to the reconnaissance waypoint and the related object slot value would refer to the waypoint’s altitude restriction. The domain general productions provide the mechanisms that manage the references between the situation elements.

The domain specific productions primarily consist of task knowledge and responses to the situations, events, actions, and objects that are learned from interacting with a specific task environment. It is the modeler’s responsibility to define the needed domain specific productions. A central goal of current research is to discover regularities and useful abstractions within the domain specific production rules that can be generalized.

The situation model represents the domain specific objects and situations to which the linguistic representations refer. The linguistic comprehension system interfaces to the non-linguistic situation model via the identification of referring expressions in the linguistic input. For example, recognition of a nominal, or object referring expression, results in the mapping to a corresponding object in the situation model. There are two basic cases: 1) recognition of a definite object referring expression typically results in identification of an existing object in the situation model or surrounding context, and 2) recognition of an indefinite object referring expression typically results in the introduction of a new object into the situation model. Extensions to these basic cases are considered in Ball (2010c) which expands the ontology of referential types to include types, collections, exemplars, prototypes and even negative instances. The extended ontology has the important benefit of simplifying the mapping from referring expressions to situation model entities.

An object referring expression from the comprehension system is mapped to the situation model when the head of the object referring expression is identified. For example, if the input is “the altitude”, then recognition of “altitude” as the head triggers the mapping to the situation model. Note that if the input is actually “the altitude restriction”, an altitude object will still be mapped to at the processing of “altitude”. At the processing of “restriction” an “altitude restriction” object will be mapped. Further, if a post-head modifier occurs as in “for Waypoint-A” in “the altitude restriction for Waypoint-A”, the mapping may need to be modified following processing of the post-head modifier. The model does not currently attempt to map to an object on the basis of pre-head modifiers as in “the red...” although there is evidence that humans may do so in Visual World Paradigm tasks (Tanenhaus et al., 1995). It should be noted that object referring expressions contain ambiguous words, not word senses or abstract concepts. It is the mapping to objects in the situation model which disambiguates the words in the linguistic representation.

Other challenges include anaphora and co-reference resolution. We currently use grammatical features to constrain the possible co-referents of a pronoun (e.g. “it” is *inanimate* and *singular*). We plan to adhere to the constraints of binding theory with respect to binding pronouns and anaphors (Chomsky, 1981) and to adopt mechanisms of Centering Theory (Grosz, Joshi & Weinstein, 1995) in a more complete implementation. We are not proposing a general solution in our research program; however, we expect to implement an initial capability for co-reference resolution by relying on ACT-R's chunk merging feature. So long as the specific context for a chunk is the same for newly introduced references to previously referenced knowledge elements, some amount of the new references automatically merge with previously constructed chunks in DM. For a more general solution, existing approaches to co-reference resolution are being investigated for inclusion in our design.

Summary and Conclusions

This paper describes a model of human language processing which is intended to be both functional and cognitively plausible. It includes a linguistic structure building mechanism which combines a serial, deterministic processing mechanism with a non-monotonic mechanism of context accommodation, and a lower level parallel, probabilistic mechanism for selecting between competing alternatives. Overall, the model is pseudo-deterministic—it presents the appearance and efficiency of deterministic processing, and can handle much of the more mundane ambiguity evident in human language via the parallel, probabilistic and non-monotonic context accommodation mechanisms. The model adheres to well-established cognitive constraints on human language processing including incremental and interactive processing. This commitment led to the integration of a cognitively plausible word recognition subcomponent, rather than adopting an off-the-shelf tokenizer and part of speech tagger that lacked cognitive plausibility.

A key attribute of the language comprehension model is the capability to handle variability and mismatch at all levels of analysis from word recognition, through the generation of linguistic representations and the mapping into the situation model, to the determination of the conversational implicatures not literally described in the linguistic input (although the capability to handle conversational implicatures is not yet implemented). There is no level of analysis at which variability and mismatch can be ignored.

The language comprehension model is a key component of a larger synthetic teammate model which is capable of functioning as the pilot in a three-person simulation of an uninhabited air vehicle reconnaissance mission task (Ball, et. al, 2009). The main objective of the synthetic teammate project is to develop cognitive agents capable of being integrated into team training simulations while maintaining training efficacy. To achieve this goal, synthetic teammates must be capable of closely matching human behavior. To this end, we have developed and integrated models of several important cognitive capacities into a composite synthetic teammate model. In addition to language comprehension and situation modeling, these capacities include the ability to perform the UAV piloting task, and language generation and dialog modeling capabilities.

Although we do not report a direct comparison of model results to human data, Cassimatis, Bello & Langley (2009) argue that models of higher-level cognitive processes, such as language comprehension, may be better evaluated on model breadth, parsimony, and functionality. Ball (2008) provides similar arguments for a functional approach, but makes a stronger commitment to cognitive plausibility. The synthetic teammate is capable of receiving text communications from a teammate, reading the text, producing linguistic representations of the text, and mapping the representations into a situation model. Based

on the contents of the situation model, the synthetic teammate then interacts with its task environment, or responds to communications with its own text messages. We believe that this demonstrates the functionality and capability of the presented language comprehension model.

References

- Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, 93, 203-231.
- Altmann, G. (1998). Ambiguity in sentence processing. *Trends in Cognitive Sciences*, 2(4), 146-152.
- Altmann, G. & Mirkovic, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, 222, 583-609.
- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191-238.
- Anderson, J. (2007). *How Can the Human Mind Occur in the Physical Universe?* NY: Oxford University Press.
- Anderson, J. R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
- Ball, J. (2007). A bi-polar theory of nominal and clause structure and function. *Annual Review of Cognitive Linguistics*, 5, 27-54.
- Ball, J. (2008). A naturalistic, functional approach to modeling language comprehension. *Papers from the AAI Fall 2008 Symposium, Naturally Inspired Artificial Intelligence*. Menlo Park, CA: AAI Press.
- Ball, J. (2010a). Context Accommodation in Human Language Processing. *Proceedings of the Natural Language Processing and Cognitive Science Workshop*. Lisbon: INSTICC Press.
- Ball, J. (2010b). Projecting grammatical features in nominals: Cognitive Processing Theory and Computational Implementation. *Proceedings of the 19th Behavior Representation in Modeling and Simulation Conference*.
- Ball, J. (2010c). Simplifying the mapping from referring expression to referent in a conceptual semantics of reference. *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society*.
- Ball, J., Heiberger, A. & Silber, R. (2007). Toward a large-scale model of language comprehension in ACT-R 6. In R. Lewis, T. Polk & J. Laird (Eds.) *Proceedings of the 8th International Conference on Cognitive Modeling* (pp. 173-179). NY: Psychology Press.
- Ball, J., Myers, C. W., Heiberger, A., Cooke, N. J., Matessa, M., & Freiman, M. (2009). The Synthetic Teammate Project. *Proceedings of the 18th Annual Conference on Behavior Representation in Modeling and Simulation*. Sundance, UT.
- Bever, T. (1970). The cognitive basis for linguistic structures. In J. Hayes (Ed.), *Cognition and the development of language*, 279-362. NY: Wiley.
- Binder, K., Pollatsek, A., & Rayner, K. (1999). Extraction of information to the left of the fixated word in reading. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 1162-1172.
- Carver, R. (1973). Understanding, information processing and learning from prose materials. *Journal of Educational Psychology*, 64, 76-84.
- Carver, R. (1973). Effect of increasing the rate of speech presentation upon comprehension. *Journal of Educational Psychology*, 65, 118-126.
- Cassimatis, N., Bello, P. & Langley, P. (2008). Ability, breadth, and parsimony in computational models of higher-order cognition. *Cognitive Science*, 32, 1304-1322.
- Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht, Holland: Foris.
- Culicover, P. (2009). *Natural Language Syntax*. NY: Oxford University Press.
- Douglass, S., Ball, J. & Rodgers, S. (2009). Large declarative memories in ACT-R. *Proceedings of the 9th International Conference on Cognitive Modeling 2009*, Manchester, UK.
- Freiman, M., & Ball, J. (2008). Computational cognitive modeling of reading comprehension at the word level. *Proceedings of the 38th Western Conference on Linguistics*, 34-45. Davis, CA: University of California, Davis.
- Freiman, M. & Ball, J. (submitted). Improving the reading rate of Double-R-Language.
- Grosz, B., Joshi, A. & S. Weinstein (1995). Centering: A framework for modelling the local coherence of discourse. University of Pennsylvania: IRCS Technical Report Series.
- Guerrera, C. (2004). Flexibility and constraint in lexical access: Explorations in transposed-letter priming. Unpublished dissertation, Department of Psychology, University of Arizona.
- Gibson, E., & Pearlmutter, N. (1998). Constraints on sentence comprehension. *Trends in Cognitive Sciences*, 2(7), 262-268.
- Lewis, R. L. (1998). Leaping off the garden path: Reanalysis and limited repair parsing. In J. D. Fodor, & F. Ferreira (Eds.), *Reanalysis in Sentence Processing*. Boston: Kluwer Academic.
- McClelland, J., & Rumelhart, D. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88(5), 375-407.
- McConkie, G. W., & Rayner, K. (1975). The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics*, 17, 578-586.
- Paap, K., Newsome, S., McDonald, J. & Schvaneveldt, R. (1982). An activation-verification model of letter and word recognition: the word-superiority effect. *Psychological Review*, 89, 573-594.
- Perea, M., & Lupker, S. J. (2003). "Does judge activate COURT? Transposed-letter similarity effects in masked associative priming". *Memory and Cognition*, 31, 829- 841.
- Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, 7, 65-81.
- Rayner, K. (1986). Eye movements and the perceptual span in beginning and skilled readers. *Journal of Experimental Child Psychology*, 41, 211-236.
- Seidenberg, Mark S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96(4), 523-568.
- Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., & Sedivy, J. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632-1634.
- Vosse, T. & Kempen, G. (2000). Syntactic structure assembly in human parsing: A computational model based on competitive inhibition and a lexicalist grammar. *Cognition*, 75, 105-143.
- Zwann, R. & Radvansky, G. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162-185.