In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 163-168. Oxford, UK: Taylor & Francis/Psychology Press.

# Toward a Large-Scale Model of Language Comprehension in ACT-R 6

**Jerry Ball[1] (Jerry.Ball@mesa.afmc.af.mil)**
**Andrea Heiberg[2] (Andrea.Heiberg@mesa.afmc.af.mil)**
**Ronnie Silber[3] (Ronnie.Silber@mesa.afmc.af.mil)**
[1]Air Force Research Laboratory/[2]L-3 Communications/[3]Lockheed-Martin
6030 S. Kent St.
Mesa, AZ 85212 USA

## Abstract

We are developing a large-scale model of language comprehension in ACT-R 6 for use in the creation of a synthetic teammate that can function as the Air Vehicle Operator (AVO) in a three-person simulation of an Unmanned Aerial Vehicle (UAV) team performing a reconnaissance mission. The use of ACT-R 6 to implement the core components of this system reflects the strong cognitive modeling orientation of this research. However, the focus is on creating cognitively plausible linguistic and associated non-linguistic representations, rather than modeling the fine-grained time course of language processing as is more typical of ACT-R models. In this regard, an empirical study aimed at discovering evidence of linguistic representations is discussed. Beside the focus on linguistic and non-linguistic representations, the large-scale nature of this effort distinguishes it from typical cognitive modeling research.

## Introduction

We are using the ACT-R 6 cognitive architecture and modeling environment (Anderson et al., 2004; Anderson & Lebiere, 1998) for development of a synthetic teammate. The long-term goal of this research is to develop language-enabled synthetic entities which can be integrated into training simulations. To achieve this goal without detriment in training, the synthetic entities must be capable of closely matching human behavior, including human language behavior. The initial application is the creation of a synthetic teammate capable of performing the functions of a UAV AVO (or pilot) in the three-person Cognitive Engineering Research on Team Tasks (CERTT) simulation (Cooke & Shope, 2005).

This paper describes the current implementation of the language comprehension component of the system in ACT-R 6. The paper begins with a short description of the theory of linguistic representation and language processing underlying model development. It continues with a description of the CERTT lab team task simulation. Following this, the current version of the model is described. A brief discussion of modeling and development tools which facilitate development follows. The paper concludes with a discussion of model validation, describing an empirical study to validate the linguistic representations which are created during processing. Overall, we follow the approach to psycholinguistic research espoused in Crocker (2005) in building a large-scale, if qualitative, model, rather than focusing on the quantitative modeling of narrowly defined and often pathological language processing phenomena.

## Linguistic Representation

The language comprehension model is founded on basic principles of Cognitive Linguistics (Langacker, 1987, 1991; Talmy, 2000; Lakoff, 1987) and Construction Grammar (Fillmore, 1988; Fillmore and Kay, 1993; Goldberg, 1995). In Cognitive Linguistics, all grammatical elements have a semantic basis, including parts of speech, grammatical markers, phrases and clauses. Understanding of language is embodied and based on experience in the world (Lakoff & Johnson, 1999). Categorization is a key element of linguistic knowledge. Categories are seldom absolute, exhibiting, instead, effects of prototypicality, base level categories (Rosch, 1978), family resemblance (Wittgenstein, 1953), fuzzy boundaries, radial structure and the like (Lakoff, 1987). Our linguistic capabilities derive from basic cognitive capabilities—there is no autonomous syntactic component separate from the rest of cognition. Knowledge of language is for the most part learned and not innate. Abstract linguistic categories (e.g., noun, verb, referring expression) are learned on the basis of experience with multiple instances of words and expressions which are members of these categories, with the categories being abstracted and generalized from experience.

Construction Grammar is a linguistic theory based on the notion of *constructions*. "Constructions are stored pairings of form and function, including morphemes, words, idioms, partially lexically filled and fully general linguistic patterns…any linguistic pattern is recognized as a construction as long as some aspect of its form and function is not strictly predictable from its component parts" and even fully predictable constructions may be stored "as long as they occur with sufficient frequency" (Goldberg, 2003).

The focus of this research is on the grammatical encoding of *Referential* and *Relational* meaning (Ball, 2005). In English, these two dimensions of meaning are typically encoded in distinct grammatical poles—a *referential pole* and a *relational pole*—with a *specifier* functioning as the locus of the referential pole and a *head* functioning as the locus of the relational pole. For example, in the expression

The pilot

the determiner "the" functions as a specifier and the noun "pilot" functions as the head. The specifier and head combine to form a *referring expression*, in this example an *object referring expression* (or nominal). Words in English divide into two basic classes: relation (verb, adjective, preposition, adverb) and object (noun, pronoun, proper noun). Relational words presume the existence of other words which express the arguments they relate. Most constructions center on some relational word (e.g., transitive verb construction, predicate adjective construction) which functions as the head of the construction, and is the locus for the encoding of relational meaning—with the construction as a whole constituting a *situation referring expression* (or clause).

Linguistic representations are perceptually grounded in non-linguistic representations of the objects and situations to which they refer. The representations of objects and situations are themselves learned from perceptual-motor (i.e., bodily) experience (cf. Barsalou, 1999). There are no purely abstract concepts that are devoid of perceptual grounding as is assumed in many cognitive theories (cf. Anderson et al., 2007). Concepts may be highly abstract, but they ultimately derive their meaning from a perceptual chain of experience (cf. Harnad, 1990)—in the limiting case perceptual experience of linguistic items themselves. A situation model (Kintsch, 1998; Zwann & Radvansky, 1998) is populated with instances of objects and situations activated by the linguistic input and non-linguistic context.

## Construction-Driven Language Processing

The processing mechanism is based on the *activation*, *selection* and *integration* of constructions corresponding to the linguistic input (Ball, 2007). Activation is a *parallel* process that biases or constrains the selection and integration of corresponding declarative memory (DM) elements into a linguistic representation. Based on the input and prior context, a collection of DM elements is activated via the parallel, spreading activation mechanism of ACT-R.

The selection mechanism is based on the *serial* retrieval mechanism of ACT-R—an alternative to the parallel *competitive inhibition* mechanism typical of connectionist models (cf. Vosse & Kempen, 2000). Retrieval occurs as a result of selection and execution of a production—only one production can be executed at a time—whose right-hand side provides a retrieval template that specifies which type of DM chunk is eligible to be retrieved. The single, most highly activated DM chunk matching the retrieval template—subject to random noise—is retrieved. The retrieval template varies in its level of specificity in accord with the production selected for execution. For example, when a production that retrieves a DM chunk of type *word* executes, the retrieval template may specify the form of the input (e.g., "airspeed") in addition to the DM type *word*. When a production that retrieves a DM chunk of type *part of speech* (POS) executes, the retrieval template may specify the word without specifying the POS—allowing the biasing mechanism to constrain POS determination. There is no assumption that humans use POS labels during language processing, but it is assumed that they categorize word into POS categories.

The retrieved DM chunk is matched on the left-hand side of another production which, if selected and executed, determines how to integrate the DM chunk into the representation of the preceding input. Production selection is driven by the matching of the left-hand side of the production against a collection of buffers (e.g., goal, retrieval, context, short-term working memory) which reflect the current goal, current input and previous context. The production with the highest utility—learned on the basis of prior experience—which matches the input and prior context, is selected for execution—subject to random noise. A default production which simply adds the retrieved DM chunk to a short-term working memory (ST-WM) stack (Ericsson & Kintsch, 1995) executes if no other production matches. The ST-WM stack—which is limited to four linguistic elements—constitutes part of the context for production selection and execution, and implements an extension to the ACT-R architecture.

A key element of the integration process is a mechanism of *context accommodation* which provides for *serial processing without backtracking*. According to Crocker (1999), there are three basic mechanisms of language processing: 1) serial processing with backtracking, 2) parallel processing, and 3) deterministic processing. Context accommodation is an alternative non-backtracking, serial processing mechanism. The basic idea behind this mechanism is that when the current input is unexpected with regard to the previously built structure, the structure is modified to accommodate the current input without backtracking. This mechanism is demonstrated using the example "no airspeed or altitude restrictions". The processing of the word "no" leads to retrieval of an object referring expression (ORE) construction containing the functional elements specifier and head (not all functional elements are shown):

$$[\texttt{spec head}]_{\texttt{obj-refer-expr}}$$

"No" is integrated as the specifier in this construction and expectations are established for the occurrence of the head:

$$[\textit{no}_{\texttt{spec}}\ \texttt{head}]_{\texttt{obj-refer-expr}}$$

This ORE construction is made available in the ST-WM stack to support subsequent processing. The processing of the noun "airspeed" leads to activation and selection of a head construction which contains the functional elements modifier and head, with "airspeed" functioning as the head:

$$[\texttt{mod airspeed}_{\texttt{head}}]_{\texttt{head}}$$

The head construction is integrated into the ORE construction.

$$[\textit{no}_{\texttt{spec}}\ [\texttt{mod}\ \textit{airspeed}_{\texttt{head}}]_{\texttt{head}}]_{\texttt{obj-refer-expr}}$$

The processing of the conjunction (or disjunction) "or" leads to its addition to the ST-WM stack since the category of the first conjunct of a conjunction cannot be effectively determined until the linguistic element after the conjunction is processed—due to rampant ambiguity associated with conjunctions. Note that delaying determination of the category of the first conjunct until after processing of the linguistic element following the conjunction provides a form of deterministic processing reminiscent of Marcus's deterministic parser (1980). The processing of the noun "altitude" in the context of the conjunction "or" and the ORE "no airspeed" with head noun "airspeed" results in the accommodation of "altitude" such that the head of the ORE is modified to reflect the disjunction of the nouns "airspeed" and "altitude".

$$[\textit{no}_{\texttt{spec}}\ [\texttt{mod}\ (\textit{airspeed or altitude})_{\texttt{head}}]_{\texttt{head}}]_{\texttt{obj-refer-expr}}$$

The processing of "restrictions" in the context of the ORE "no airspeed or altitude" results in the accommodation of "restrictions" such that the current head "airspeed or altitude" becomes the modifier and "restrictions" becomes the head. The final representation has the form:

$$[\textit{no}_{\texttt{spec}}\ [(\textit{airspeed or altitude})_{\texttt{mod}}\ \textit{restrictions}_{\texttt{head}}]_{\texttt{head}}]_{\texttt{obj-refer-expr}}$$

This representation was arrived at using a serial processing mechanism without backtracking, despite the rampant local ambiguity of the utterance!

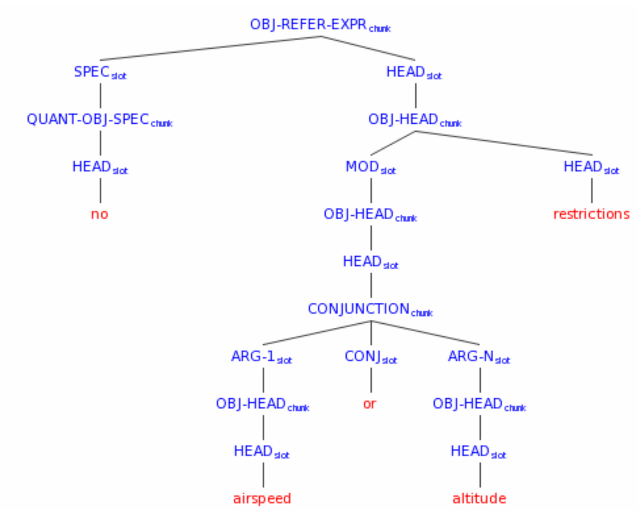


**Figure 1: No altitude or airspeed restrictions**

Context accommodation is a powerful serial processing mechanism which overcomes the limitations and cognitive implausibility of serial processing with algorithmic backtracking, full parallel processing, and full deterministic processing. Context accommodation is closely related to Lewis's notion of "limited repair parsing" (Lewis, 1998), although context accommodation is considered part and parcel of the normal processing mechanism and is not viewed as a repair mechanism. Regarding parallel processing, it is not cognitively plausible to carry forward more than a few possible representations at once, which means that a mechanism like context accommodation is needed to handle the case where the correct parse isn't in the parallel spotlight. Likewise, deterministic mechanisms require delaying integration of linguistic elements for indeterminate periods—requiring their separate representation—which is likely to exceed the limited capacity of ST-WM if used extensively.

## Synthetic Teammate

The CERTT Lab is a research facility for studying team performance and cognition in complex settings. CERTT's UAV-STE (Synthetic Task Environment) is a three-person task in which each team member is provided with distinct, though overlapping, training; has unique, yet interdependent roles; and is presented with different and overlapping information during the mission (Cooke & Shope, 2005). The overall goal is to fly the UAV to designated target areas and to take acceptable photos at these areas. The Air Vehicle Operator (AVO) controls airspeed, heading, and altitude, and monitors UAV systems. The Payload Operator (PLO) adjusts camera settings, takes photos, and monitors the camera equipment. The Data Exploitation, Mission Planning, and Communication Operator (DEMPC) oversees the mission and determines flight paths under various constraints. To successfully complete a mission, the team members need to share information with one another in a coordinated fashion.

Most communication is done via microphones and headsets, although some involves computer messaging. A set of initial speech transcripts has been collected from the UAV-STE for a number of teams. These transcripts are being used to guide development of the model. A portion of a transcript appears below.

PLO: AVO, can I please be about 3000 feet or higher, please? Cancel. Cancel.

AVO: Do I need to change my airspeed? I mean my altitude.

DEMPC: Once I get the first, uh, sequence figured out, I'll let you know. First waypoint LVN is an, uh, ROZ access point. There is no flight restrictions, but the, uh, radius is, uh, 2.5 miles. I'm pretty sure you can take a bearing towards H-area now. It looks like you're in within the 2.5 required for this entry point.

PLO: AVO, can I please, uh, keep, uh, altitude over 3000 feet for this picture, please? Can you give me a range?

DEMPC: The next target H-area has a range of 5 miles.

PLO: Copy.

AVO: Was that above 3000?

PLO: Yes, please. Can you also keep this current airspeed?

AVO: OK.

DEMPC: Next waypoint is H-area. There is no altitude restriction, but the speed restriction is between 50 & 200.

The language used by team members is not constrained; there is no special restrictive grammar. Over the three transcripts analyzed so far, the average number of utterances is 2300 per transcript; average utterance length is 7 words. There are 1300 unique words across transcripts, including 50 special vocabulary items related to the task. In each transcript, an average of 27% of the words are unique to that transcript.

The transcripts contain a number of grammatical features that are challenging from a language processing perspective:

- Multiword expressions: "Picking up the pace."
- Complex object referring expression (or nominal):
  - "First waypoint LVN is an, uh, ROZ access point."
  - "Do you have any additional altitude or speed restrictions that I need to get from you?"
  - "Can I get the 2, 3, 4, and 5 current setting ranges?"
- Anaphora:
  - "PLO, is this a photo?" (current waypoint)
  - "Does that make sense?" (previous statement)
- Complex verb argument structure: "Can I keep altitude over 3000 feet for this picture?"
- Corrections:
  - "DEMPC to PLO, effective radius is, uh, 2, uh, 5 miles, sorry about that."
  - "Do I need to change my airspeed? I mean my altitude."
- Ambiguous closed-class words, e.g., "that":
  - Complementizer: "They just told me that there's gonna be a priority target in this area that we're entering."
  - Object referring expression: "You already told me that."
  - Determiner: "Got a good photo on that one."

## The Current Model

The language comprehension model is currently capable of processing a range of grammatical constructions attested in the transcripts, including:

- Intransitive verb: "You can go."
- Transitive verb: "We already hit [$_{OBJ}$ ROW]."
- Ditransitive verb: "You can give [$_{IOBJ}$ me] [$_{OBJ}$ R-STE]."
- Verb taking clausal complement: "You told [$_{IOBJ}$ me] [$_{SITCOMP}$ the altitude restriction was below 3000 feet]."
- Auxiliary verb: "I would have had a wrong picture."
- Predicate nominal: "First waypoint is LVN."
- Predicate adjective: "Altitude is stable."
- Predicate preposition: "We are in those constraints."
- Attributive adjective modifier: "It's a good picture."
- Adverbial modifier: "Our altitude still should be fine."
- Complex nominal: "The next photographic target point is M-STR."
- Nominal conjunction: "We will maintain current airspeed and altitude."
- Sentence conjunction: "The entry is KGM and the exit is FRT."

The model creates a linguistic representation of the input, but doesn't yet map that representation to the corresponding objects and situations in the situation model.

The language comprehension model is approaching a scale and complexity atypical of most cognitive models. Verifying that the model generates theoretically motivated linguistic representations is an important on-going aspect of the project. Inputs to the model are comprised of actual utterances from the UAV-STE transcripts and a set of canonical phrases and sentences. The verification strategy includes running the model against this set of inputs, and testing that the model produces the expected output.

The model generates linguistic representations which include such information as phrase constituency, predicate/argument relations, head/modifier relations, and head/specifier relations. Linguistic representations are complex structures of DM chunks. For testing, the DM chunk structure is converted into a graphical representation (automatically generated with phpSyntaxTree, Eisenbach & Eisenbach, 2006) shown in Figure 2 (below).

At a gross level, testing is fully automated. The complex output structure (e.g., Figure 2) is traversed in left-to-right order, and the terminal symbols are reassembled into a string (e.g., "I increased the airspeed"). This output string is compared to the input string; any mismatches are flagged for further investigation. At a more detailed level of testing, the output representation is hand-checked to ensure its validity. Valid output representations are stored as the known-good baseline. A capability to dynamically visualize the evolving DM representation during the processing of each word in an input text also exists (Heiberg, Harris & Ball, 2007). Any further changes to the model may be easily regression tested by regenerating the outputs, and comparing them to the known-good baseline with an automated file comparison tool. This set of methodologies taken together helps facilitate the development of a large scale and complex model.

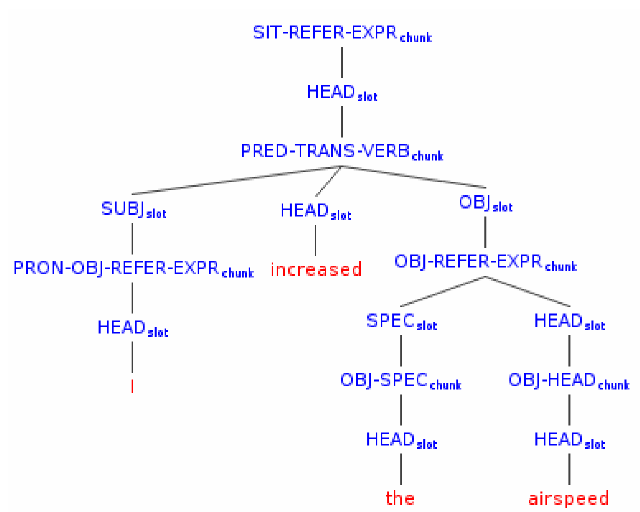


**Figure 2 Graphical Representation**

## Model Validation

We are committed to the development of a cognitively plausible model of language comprehension. However, we are not modeling the fine-grained time course of processing during language comprehension. As Just and Carpenter (1987) note, "in most cases…syntactic and semantic analyses occur concurrently with other processes that are longer and more variable in duration". It is only in the processing of unusual texts like Garden-Path and center embedded sentences that syntactic and semantic influences on processing are exposed. Instead, our focus is on validating the linguistic and non-linguistic representations that are generated during processing of more ordinary texts as reflected in our UAV team task corpus. Our approach aligns with Crocker (2005), who argues for "an alternative approach to developing and assessing theories and models of sentence comprehension" in which "a model's coverage should not be limited to a few 'interesting' construction types, but must also extend to realistically large and complex language fragments, and must account for why most processing is typically rapid and accurate".

To validate the model, we have devised a multi-part empirical study to identify the kinds of linguistic representations that humans create during language processing. Some preliminary data from a pilot study involving 20 subjects are reported.

In one part of this study subjects are asked to determine if paired expressions differ in meaning. For example, does "the man bit the dog" differ in meaning from "the dog bit the man" (all 20 subjects said "yes")? Does "this book" differ in meaning from "that book" (all 20 subjects said "yes")? Does "the old house on the hill" differ from "an old house on a hill" (14 subjects said "yes")? A difference in meaning indicates that either the different lexical items or the different structural arrangement of the lexical items in the paired expressions affects the meaning.

In another part, subjects are asked to group expressions into meaningful units of various sizes given the overall meaning of the linguistic expression. A key question for this part is whether the preferred representation of a clause aligns with the Subject-Predicator$_{Head}$-Object construction put forward in most theories of Functional Grammar (cf. Halliday & Matthiessen, 2004), the S → NP VP (i.e., Subject-Predicate) construction put forward in Generative Grammar (Chomsky, 1965) or the ReferencePoint-PredSpec-Predication construction specific to Double R Theory (Ball, 2005), but related to the Mood-Residue construction of Halliday & Matthiessen (2004). Also of interest is whether this preference varies from clause to clause and from subject to subject. For example, in "he is kicking the ball" do subjects prefer to group "is" with "kicking" (as part of the Predicator$_{Head}$ "is kicking"), with "kicking the ball" (as part of the Predicate "is kicking the ball") or with "he" (as part of the ReferencePoint-PredSpec "he is")? Preliminary results indicate the identification of the entire sentence as a meaningful group by 13 subjects. The second most common grouping, "kicking the bucket", was only identified by 4 subjects. "He is" was identified as a group by 3 subjects. The group "is kicking the ball" was not identified by any subjects. Overall, subjects tend not to include function words (e.g., "the", "is") in meaningful groups, making it difficult to assess these results, and perhaps making it necessary to revise the methodology.

In a third part subjects are asked to rank the relative contributions of various words to the overall meaning of the expression. In conjunction with the grouping task, the ranking task will allow us to identify the semantically most important word in a meaningful group, which we take to be the head of the group. For sentences containing transitive verbs, subjects identified the head of the subject as most meaningful, the main verb as second most meaningful and the head of the object as third most meaningful. The ranking of heads of subjects as more important than main verbs suggests that relational structure is not the only dimension of meaning which influences this decision.

Finally, in a fourth part, subjects are asked to identify the part of speech (POS) of a word in different expression contexts, for example, "running" in "the man is running" vs. "the running man". This part is intended to get at whether or not words are separately represented in the mental lexicon for each possible grammatical function they can fulfill (e.g., head, modifier, complement, specifier), as indicated by the POS labels they are assigned. For example, if "running" is treated as a verb in "the man is running" but as an adjective in "the running man" this might indicate separate representations in the mental lexicon. If, on the other hand, "running" is labeled a verb in both uses, a single entry in the mental lexicon is suggested. Preliminary results suggest that subjects tend to treat words in different functions as having the same part of speech. For example, 18 subjects call "boy" in "the boy" a noun; 11 subjects called "altitude" in "no altitude restrictions" a noun (6 subjects called it an adjective); 16 subjects called "home" in "he went home" a noun (only 1 subject called it an adverb); and 14 subjects called "president" in "George Bush is president" a noun (3 subjects called it an adjective).

To the extent that the linguistic representations generated by the language comprehension model are consistent with the results of this empirical study, the linguistic representations will have more validity. To some extent this validity hinges on whether or not humans have explicit knowledge of the linguistic representations they generate during language comprehension, and whether the empirical study is successful in tapping into that knowledge. It is a general assumption of the empirical study that humans have explicit knowledge of the linguistic representations they create. This assumption is motivated by the model's creation of linguistic representations composed of DM chunks, which suggests that these representations can be explicitly attended to and cognitively manipulated, and by rejection of the *autonomy of syntax* assumption of generative grammar, with its informationally encapsulated (and hence implicit) syntax module. Although the mechanisms by which linguistic representations are constructed may be largely implicit, the resulting representations are declarative and explicit. For example, humans explicitly know what "the man bit the dog" means. They explicitly know that "the man" refers to a man, "the dog" refers to a dog, "bit" establishes a relation of biting between the man and the dog, with the expression as a whole referring to a situation in which it is the man who is doing the biting, and the dog who is being bitten.

## Conclusion

We are using the ACT-R cognitive architecture in the development of a language-enabled synthetic teammate intended to closely match human behavior. To date, the research has focused on the development of the language comprehension component of the system. This component has approached a scale at which the need for development and testing tools has become important. The goal for the project is to maintain cognitive plausibility by adhering to well-established theoretical constraints from cognitive linguistics and cognitive psychology, as the system grows. We believe these constraints will actually facilitate development of a functional system (Ball, 2006).

## Acknowledgements

This research is funded by the Warfighter Readiness Research Division, Human Effectiveness Directorate, Air Force Research Laboratory. We thank Nancy Cooke and Steve Shope of CERI for providing access to the CERTT lab UAV-STE and transcripts.

## References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S, Lebiere, C, and Qin, Y. (2004). An Integrated Theory of the Mind. *Psychological Review* 111, (4). 1036-1060.

Anderson, J. R. & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: LEA.

Anderson, J. R., Qin, Y., Jung, K-J, & Carter, C. (2007). Information-processing modules and their relative modality specificity. *Cognitive Psychology*, 54, 185-217.

Ball, J. (2005). A Bi-Polar Theory of Nominal and Clause Structure and Function. In *Proceedings of the 27th Annual Cognitive Science Society*.

Ball, J. (2006). Can NLP Systems be a Cognitive Black Box? In *Papers from the AAAI Spring Symposium, Tech Report SS-06-02*, 1-6, Menlo Park, CA: AAAI Press.

Ball, J. (2007). Construction Driven Language Processing. In *Proceedings of the 2nd European Cognitive Science Conference.*

Barsalou, L. (1999). Perceptual Symbol Systems. In *Behavioral and Brain Sciences* 22, 577-609.

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

Cooke, N. & Shope, S. (2005). Synthetic Task Environments for Teams: CERTT's UAV-STE. *Handbook on Human Factors and Ergonomics Methods*. 46-1-46-6. Boca Raton, FL: CLC Press, LLC.

Crocker, M. (1999). Mechanisms for Sentence Processing. In Garrod, S. & Pickering, M. (eds), *Language Processing*, London: Psychology Press.

Crocker, M. (2005). Rational models of comprehension: addressing the performance paradox. In A. Culter (ed) *Twenty-First Century Psycholinguistics: Four Cornerstones*, LEA, Hillsdale.

Eisenbach, A. & Eisenbach, M. (2006). phpSyntaxTree tool, http://ironcreek.net/phpsyntaxtree.

Ericsson, K & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102 211-245.

Fillmore, C. (1988). The Mechanisms of Construction Grammar. *BLS* 14: 35-55.

Fillmore, C. and Kay, P. (1993). *Construction Grammar Coursebook*. Berkeley, CA: Copy Central.

Goldberg, A. (1995). *A Construction Grammar Approach to Argument Structure*. Chicago: The University of Chicago Press.

Goldberg, A. (2003). Constructions: a new theoretical approach to language. *TRENDS in Cognitive* Sciences. Vol. 7, No. 5, pp. 219-224.

Halliday, MAK & Matthiessen, C (2004). *An Introduction to Functional Grammar*. NY: Oxford University Press.

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D*, 42: 335-346.

Heiberg, A., Harris, J. & Ball, J. (2007). Dynamic Visualization of ACT-R Declarative Memory Structure. In *Proceedings of the 8th International Conference on Cognitive Modeling*.

Just, A. & Carpenter, P. (1987). *The Psychology of Reading and Language Comprehension*. Boston: Allyn and Bacon, Inc.

Kintsch, W. (1998). *Comprehension, a Paradigm for Cognition*. New York, NY: Cambridge Univ Press.

Langacker, R. (1987, 1991). *Foundations of Cognitive Grammar*, Volumes 1 & 2. Stanford, CA: Stanford University Press.

Lakoff, G. (1987). *Women, Fire and Dangerous Things*. Chicago: The University of Chicago Press

Lakoff, G. & Johnson, M. (1999). *Philosophy in the Flesh: The embodied mind and its challenge to Western thought*. New York: Basic Books.

Lewis, R.L. (1998). Leaping off the garden path: Reanalysis and limited repair parsing. In Fodor, J.D. and Ferreira, F. (eds), *Reanalysis in Sentence Processing*. Boston: Kluwer Academic.

Marcus, M. (1980). *A Theory of Syntactic Recognition for Natural Language*. Cambridge, MA: The MIT Press.

Rosch, E. (1978). Principles of Categorization. In *Cognition and Categorization*. Edited by E. Rosch & B. Lloyd. Hillsdale, NJ: LEA

Talmy, L. (2000). *Toward a Cognitive Semantics, Vols I and II*. Cambridge, MA: The MIT Press

Vosse, T. & Kempen, G. (2000). Syntactic structure assembly in human parsing. *Cognition*, 75, 105-143.

Wittgenstein, L. (1953). *Philosophical Investigations*. New York: MacMillan.

Zwann, R., and Radvansky, G. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123, 162-185.