

Interpretation in a Cognitive Architecture

Harold Paredes-Frigolett

Center for the Study of Language and Information
Cordura Hall, Stanford University
Stanford, CA 94305, USA
harold.paredes-frigolett@csl.stanford.edu

Abstract

The work reported in this article presents a computational model of interpretation. The model proposes a cognitive architecture for intelligent agents to reason about competing analyses during interpretation and leverages the positive reinforcement principle.

1 Motivation

Interpretation of natural language involve the computational effort associated with repeatedly computing, interpreting and deindexing logical forms for ambiguous parses. In our view, interpretation can be construed as a negotiation process whereby lexical, structural, semantic, common-sense and world knowledge information and referential context are used to assign plausibilities to competing analyses.

The approach to interpretation taken here has been motivated by cognitive architectures for intelligent agents in the tradition of SOAR (Laird, Newell, and Rosenbloom, 1987; Laird, 1991), ACT-R (Anderson, 1993) and ICARUS (Langley et al., 2003).

In extending cognitive architectures in this tradition to deal with the problem of interpretation, agents carry both the “meaning” of competing analyses and the plausibilities associated with them, which we construe as the reward function of the agents. As more information becomes available from the input string, reward functions are updated and the analysis with the higher plausibility becomes the preferred interpretation.

2 The grammar formalism

Consider sentence (1):

(1) *The* \uparrow_a *spy* \uparrow_b *watched* \uparrow_c *the* \uparrow_d *cop* \uparrow_e *with* \uparrow_f *the*
 \uparrow_g *revolver* \uparrow_h \uparrow_i

The GPSG-like grammar fragment with semantic annotations in EPISODIC LOGIC (EL), a seman-

- A. DET[def] \leftarrow *The* ; *The*
- B. N \leftarrow *spy* ; *spy*
- C. N \leftarrow *cop* ; *cop*
- D. N \leftarrow *revolver* ; *revolver*
- E. N \leftarrow N PP ; $\lambda y[[y N'] \wedge [y PP']]$
- F. NP \leftarrow N[plur] ; (*K N'*)
- G. P[with-attrib] \leftarrow *with* ; *with-attrib*
- H. P[with-instr] \leftarrow *with* ; *with-instr*
- I. NP \leftarrow DET N ; (*DET' N'*)
- J. PP \leftarrow P NP ; (*P' NP'*)
- K. V[past, _NP] \leftarrow *watch* ; (*past watch*)
- L. ADVL[post-VP] \leftarrow PP[a-mod] ; (*adv-a PP'*)
- M. VP \leftarrow V[past, _NP] NP ; (*V' NP'*)
- N. VP \leftarrow V[past, _NP]
NP ADVL[a-mod] ; (*ADVL'(V'NP')*)
- O. S \leftarrow NP VP ; (*NP' VP'*)
- P. PUNC[tell] \leftarrow . ;
- Q. S \leftarrow S[full-decl] PUNC[tell] ; (*decl S'*)

Figure 1: GPSG-like grammar fragment

tic and knowledge representation language for general NLU (Hwang and Schubert, 1993), is shown in Figure 1.¹

3 Semantic analysis

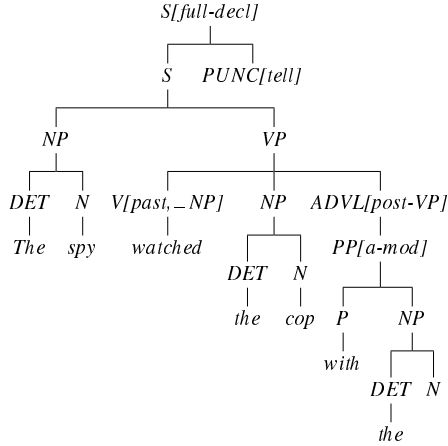
For sentence (1) the parser computes two initial analyses using not only structural, but also subcategorization and thematic role information as soon as the verb is encountered.

At point *g*, we could be already predicting several of the possible continuations. Based on subcategorization and thematic role information for the verb *watch*, there is a first analysis that results from applying rule **N** of our GPSG-like grammar fragment.

At point *g*, the first analysis is $APT\mathcal{T}_g$.

¹We refer the reader to (Hwang and Schubert, 1993) for a detailed description of EL.

\mathcal{APTT}_g :



We show here analysis \mathcal{APTT}_g and leave \mathcal{APTT}_g , the analysis in which the prepositional phrase attaches lower to the second NP, to the reader. As we will see, the cognitive architecture allows structural information, lexical factors and semantic biases to be used for on-line pruning of alternative parses, thus keeping the lid on the explosion of alternatives, and allowing human-like parsing behavior.²

For sentence (1) at point g , the semantic interpreter uses the semantic annotations associated with each syntactic rule in the grammar and applies compositional semantic interpretation rules to come up with the parameterized unscoped logical form \mathcal{PULF}_g .

\mathcal{PULF}_g :

$$\lambda P(\text{decl } [\langle \text{The } x:[x \text{ spy}] \rangle \\ ((\text{adv-a } (\text{with-instr } \langle \text{The } z:[z P] \rangle)) \\ (\langle \text{past watch} \rangle \langle \text{The } y:[y \text{ cop}] \rangle))])$$

We assume that (i) salient referents in the current discourse take wide scope over all other operators in the logical form and that they are scoped within speech act operators, (ii) tense operators are sentential operators scope within speech act operators and salient referents, (iii) tense operators take wide scope over nonsalient definites, (iv) nonsalient definites take wider scope over operators, and finally (v) existentials are scoped within all other operators in the logical form.

Using the above-mentioned heuristics on \mathcal{PILF}_g , the incremental scoper generates the following parameterized indexical logical form:

²As revealed by garden path phenomena and other psycholinguistic evidence.

\mathcal{PILF}_g :

$$\lambda P(\text{decl } (\text{past } (\text{The } x:[x \text{ spy}] \\ (\text{The } y:[y \text{ cop}] \\ (\text{The } z:[z P][x((\text{adv-a } (\text{with-instr } z)) \\ (\text{watch } y))))))))$$

If we propose semantic representations that are complete for the partial parse trees during incremental processing, we can use a slightly different version of the algorithm developed by Hwang and Schubert for incremental deindexing.³

The incremental deindexer yields the parameterized episodic logical form \mathcal{PELF}_g .

\mathcal{PELF}_g :

$$\lambda P(\exists u_1:[u_1 \text{ same-time } Now1] \wedge [u_0 \text{ precedes } u_1]) \\ [[\text{Speaker tell Hearer } (\text{That} \\ (\exists e_1:[e_1 \text{ before } u_1] \wedge [e_0 \text{ orients } e_1]) \\ [[(\text{The } x:[x \text{ spy}] \\ (\text{The } y:[y \text{ cop}] \\ (\text{The } z:[z P] \\ [[x | e_1] \\ ((\text{with-instr } z) \\ (\text{watch } y)))) \\))]] \\ ** e_1]]) \\ ** u_1])$$

The relation *orients* introduced in \mathcal{PELF}_g corresponds to a relation to be further particularized to a temporal, causal or *part-of* relation between situations. *Now1* corresponds to a term that refers to the speech time of the utterance. *Speaker* and *Hearer* stand for the speaker and the hearer of the utterance, respectively. e_0 corresponds to a prior episode described by the utterance situation u_0 . u_0 immediately precedes the utterance situation u_1 . e_1 is the situation being described by u_1 and occurs at about the same time as u_1 . *That* is a sentence nominalization operator that takes a sentence as argument and gives rise to a proposition-denoting term. The expression $[[x | e_1] (\text{with-instr } z)]$ corresponds to the action of the spy's watching modified so as to be performed with something. The function $|$ is a pairing function applicable to individuals and tuples. Thus $[x | e_1]$ is the action performed by x that gives rise to event e_1 . The operator \preceq is a metalogical operator that corresponds to the operator *coextensive-part-of* in EL. The expression $e_2 \preceq e_1$ indicates that situation e_2 is coextensive with situation e_1 , that is, e_1 and e_2 have the same spatiotemporal location.

³The reader is referred to (Hwang and Schubert, 1992) for details on the nonincremental deindexing rules.

information, referential context, and world knowledge in the calculation of plausibility for each analysis.

A naive approach to incremental interpretation would consist in exploiting this model to arrive at an episodic logical form and then consider all alternative equally plausible. We regard this alternative as implausible on the grounds of psycholinguistic results on control processes of inference in on-line text comprehension (Balota, Flores d'Arcais, and Rayner, 1990). As we will see, the model of incremental interpretation proposed is based on a multi-agent cognitive architecture in which agents are assigned competing interpretations.

Central to the architecture is the concept of agent reward. A reward function is calculated for each agent, each one of them representing an analysis. In general, the alignment between user utility function and agent reward function is one of the areas that is domain-dependent in this architecture. We will explore this alignment for our domain in the following sections.

5.1 Value alignment

In general, the objective of each agent is to maximize its reward function. How well they optimize the user utility function will depend on the alignment between the user utility function and the agent reward function.

Notice that, in our domain, there is a disconnect between the objectives of the agent and those of the user, respectively. The agent commits to an analysis and in so doing its fate is already sealed. How well they end up optimizing the user utility function will depend on variables the agent can only partially control as this process ultimately depends on the information not yet absorbed from the input string and on the behavior of the other agents in the architecture.

The interesting point to note here is that although this might at first glance appear as an undesirable feature of the architecture, it actually leads to a model in which the different agents cooperate towards the ultimate goal of optimizing the user utility function. Thus, instead of competing analyses we might as well refer to them as co-operating analyses.

5.2 Agent reward function

Reward functions for the agents are defined based on the principle of positive reinforcement.

5.2.1 Positive reinforcement principle

An analysis is preferred over another to the extent that it satisfies the constraints of its immediate ref-

erential context and to the extent that the inferences triggered in the knowledge base are more consistent, more specific and more numerous.

5.2.2 Heuristics

Based on the principle, our model uses the following sets of heuristics for assigning a reward function to agents:

1. Give referential context highest precedence;
2. Give consistency of inferences drawn in the knowledge base precedence over specificity of inferences drawn in the knowledge base;
3. Give specificity of inferences drawn in the knowledge base preference over subcategorization information;
4. Give subcategorization information precedence over the amount of inferences drawn in the knowledge base, and
5. Consider only inferences with a minimum level of "interestingness."⁴

The list above is not exhaustive, but it gives us an initial set of heuristics to define the reward function for the agents. The choice of some precedences in the heuristics above has been psycholinguistically motivated, as shown in (Altmann and Steedman, 1988).⁵ The approach to interpretation followed here is based on the assumption that information from different sources enters the interpretation process at different times and that they concurrently restrain the number of potential analyses, as suggested in recent psycholinguistic theories of human sentence comprehension (Spivey-Knowlton and Sedivy, 1995).

5.3 Interpretation as learning

The process of finding a preferred interpretation at a given time t is the result of a process of entropy reversal through information expressed in terms of a set of heuristics that govern the agent reward in this cognitive architecture. The heuristics above are a distillation of the information required for this entropy reversal process.

⁴With interestingness measured as a threshold on the conditional probability that results in an inference chain through world knowledge axioms expressed as probabilistic conditionals (Bacchus, 1990).

⁵Altmann and Steedman dealt with referential context only. To our knowledge, there is no psycholinguistic studies that give us a more complete picture on the precedences stated above.

5.4 An example

Let us illustrate the process of agent-based interpretation using our example. When processing sentence (1) up to point g , we do have two analyses.

Skolemizing E_1/e_1 , E_2/e_2 , E_3/e_3 , E_4/e_4 , E_5/e_5 , U_1/u_1 , X/y , Y/y and Z/z , the set of inferences drawn at point g is as follows:

Agent1:

- F_1 [E_1 before U_1]
- F_2 [[[X | E_1] (with-instr Z)] \wedge [X watch Y]] ** E_1]
- F_3 [X spy]
- F_4 [Y cop]
- F_5 [$E_2 \preceq E_1$]
- F_6 [[X watch Y] ** E_2]
- F_7 [[Z = (K ((nn viewing) instrument))]]
- F_8 [$E_1 \preceq E_3$]
- F_9 [[[X | E_3] (with-instr P)] \wedge [X see Y]] ** E_3]

Facts F_1 through F_4 are directly obtained in the knowledge base by asserting \mathcal{ELF}_g after splitting conjunctions and top-level skolemization is performed on \mathcal{ELF}_g . E_1 is a situation fully described by the action of the spy watching the cop and being modified so as to be performed with “something.”

Facts F_5 and F_6 are directly obtained by meaning postulate $\mathcal{MP} 1$. E_2 , co-extensive with E_1 , is fully described by the action of the spy watching the cop. $\mathcal{MP} 3$ accounts for triggering fact F_7 , thus setting the expectation in *Agent1*'s discourse model that the incoming referent is a viewing instrument. Facts F_8 and F_9 are obtained using meaning postulate $\mathcal{MP} 2$. E_3 , co-extensive with E_1 , is fully described by the action of the spy seeing the cop with “something.” Notice that using facts F_7 and F_9 and world knowledge axiom $\mathcal{WK} 1$, we would also be setting the uncertain prediction that the spy sees the cop clearly.

For our second agent, we would have the following set of inferences in the knowledge base:

Agent2:

- F_1 [E_1 before U_1]
- F_2 [X spy]
- F_3 [Y cop]
- F_4 [Y with-attrib Z]
- F_5 [[[X watch Y]] ** E_1]

Facts F_1 to F_5 are directly obtained in the knowledge base by asserting \mathcal{ELF}_g . E_1 is a situation

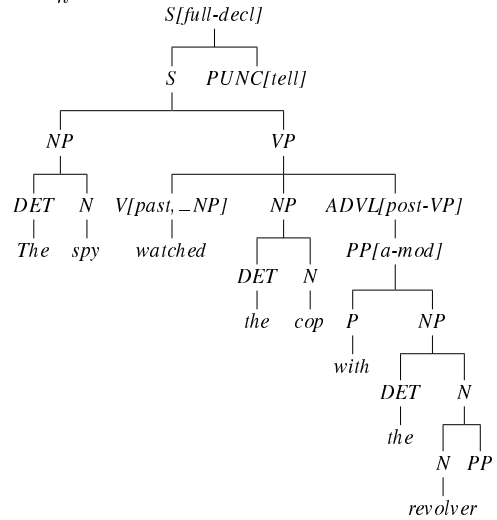
fully described by the action of the spy watching a cop carrying “something.”

5.4.1 Positive reinforcement: Scenario 1

In this first scenario, we assume that the discourse model is initially empty. Applying our positive reinforcement principle at point g , *Agent1* is the most plausible one. The verb *to watch* subcategorizes for an instrumental argument and the analysis pursued by *Agent1* is initially preferred. This analysis also leads to more inferences in the knowledge base, including the certain prediction that the incoming NP introduces a viewing instrument in *Agent1*'s discourse model and the uncertain prediction that the spy sees the cop clearly. Referential context does not play a role yet since there were no discourse referents introduced initially under this scenario.

The analysis pursued up to point h by *Agent1* is shown in \mathcal{PPT}_h .

\mathcal{PPT}_h :



At point h , we obtain \mathcal{ELF}_h :

\mathcal{ELF}_h :

- $(\exists u_1: [[u_1$ same-time *Now1*] \wedge
- $[u_0$ immediately-precedes $u_1]]$
- [[Speaker tell Hearer (That
- $(\exists e_1: [[e_1$ before u_1] \wedge [e_0 orients e_1]]
- [[The x : [x spy]
- (The y : [y cop]
- (The z : [[z revolver] \wedge
- $[z$ P]]
- [[x | e_1]
- [[(with-instr z)
- (watch y)
-]])))]
-]])))]
- ** e_1]]))]
- ** u_1])

After asserting \mathcal{ELF}_h and using type-hierarchical knowledge, the following additional facts can be triggered in *Agent1*'s discourse model:

Agent1:

F_{10} [Z revolver]

F_{11} [Z weapon]

F_{12} [Z instrument]

Fact F_7 is inconsistent with facts F_{10} through F_{11} above. Thus, by a process of "hierarchy climbing" in the knowledge base, it turns out that the variable z introduced in \mathcal{ELF}_h is not subsumed by a generic term denoting a viewing instrument in the knowledge base, as expected at point g . By our positive reinforcement heuristics at point h , this analysis is not positively reinforced. It turns out to be inconsistent with *Agent1*'s prior referential context.

Agent2's discourse model at point h has led to the following discourse model:

Agent2:

F_1 [E_1 before U_1]

F_2 [X spy]

F_3 [Y cop]

F_4 [Z revolver]

F_5 [[[X watch Y]] ** E_1]

F_6 [Y with-attrib Z]

Using our heuristics, the analysis preferred under this first scenario turns out to be the one pursued by the second agent. Notice that the agent-based cognitive architecture will get "garden-pathed" as the analysis initially preferred on the grounds of subcategorization information and specificity and interestingness of the inferences drawn in the knowledge base proves anomalous by referential context.

5.4.2 Positive reinforcement: Scenario 2

In this second scenario, we assume the discourse model initially consisting of three referents, a spy and two cops.

Applying our positive reinforcement principle at point g for *Agent1*, we have that the verb *to watch* subcategorizes for an instrumental argument and the analysis pursued by *Agent1* would be initially preferred based on subcategorization information for

this verb. *Agent1* would also get positively reinforced since its interpretation leads to more inferences in the knowledge base and to make the certain prediction that the incoming NP introduces a viewing instrument and the uncertain prediction that the spy sees the cop clearly in *Agent1*'s discourse model, as we have already seen. But according to our set of heuristics, referential context preempts these preferences, as the need to resolve the anaphoric reference *the cop* in the discourse model immediately preceding sentence (1) takes precedence over the other criteria at point g . The analysis pursued by *Agent1* does not contribute to resolving this anaphoric reference.

On the other hand, the analysis pursued by *Agent2* at point g raises the expectation that the cop will be further "particularized" so as to resolve this anaphoric reference. Given the heuristics, this expectation takes precedence over *Agent1*'s interpretation and is preferred in this cognitive architecture. It is interesting to note that in this second scenario, subjects are not being led down the garden path when given the referential context in which the need for resolving the anaphoric reference introduced by the second NP arises at point g . Our model predicts this behavior accordingly.

At point i , the analysis pursued by *Agent 2* is the preferred one. The resulting episodic logical form at point i is \mathcal{ELF}_i , as shown below.

\mathcal{ELF}_i :

$$\begin{aligned}
 &(\exists u_1: [[u_1 \text{ same-time } Now1] \wedge \\
 &\quad [u_0 \text{ immediately-precedes } u_1]] \\
 &\quad [[\text{Speaker tell Hearer (That} \\
 &\quad \quad (\exists e_1: [[e_1 \text{ before } u_1] \wedge [e_0 \text{ orients } e_1]] \\
 &\quad \quad \quad [[(\text{The } x: [x \text{ spy}] \\
 &\quad \quad \quad \quad (\text{The } y: [[y \text{ cop}] \wedge \\
 &\quad \quad \quad \quad \quad (\text{The } z: [z \text{ revolver}] \\
 &\quad \quad \quad \quad \quad \quad [y \text{ with-attrib } z]]) \\
 &\quad \quad \quad \quad \quad \quad \quad [x \text{ watch } y])]]) \\
 &\quad \quad \quad \quad \quad \quad \quad \quad ** e_1)])]) \\
 &\quad \quad \quad \quad \quad \quad \quad \quad ** u_1])
 \end{aligned}$$

\mathcal{ELF}_i leads to the following discourse model.

Agent2:

F_1 [E_1 before U_1]

F_2 [X cop]

F_3 [Y spy]

F_4 [Z revolver]

F_5 [Y with-attrib Z]

F_6 [[[X watch Y]] ** E_1]

6 Discussion

The work presented in this article puts forth an approach to interpretation using a cognitive architecture for intelligent behavior. Our work has so far consisted in defining agent reward based on the positive reinforcement principle. For the initial implementation of the principle, we have followed a heuristics-based approach.

Though some of the information used in the plausibility computation is probabilistic (Bacchus, 1990), the heuristics are non-probabilistic in our model. In defining the heuristics, we have incorporated recent results in psycholinguistic studies of human sentence processing. In our view of the interpretation process, agents contribute to arriving at a “preferred” interpretation by maintaining a “more plausible” analysis—and its associated discourse model—as the most salient one, while other less plausible analyses are kept in memory for a given period of time by other agents. By a process of heuristics-based plausibility computation, the “most plausible analysis” remain active in this architecture and take the lead during the interpretation process.

This cognitive architecture gives a plausible account of some of the issues that pervade human sentence processing such as garden-path phenomena. In so doing, we depart from serial first-analysis approaches to sentence comprehension in the tradition of the garden-path theory of sentence processing (Frazier and Fodor, 1978; Frazier and Clifton, 1996) and endorse more recent psycholinguistic accounts of this problem which view the interpretation process as a concurrent negotiation of information from syntactic, semantic and pragmatic sources by several agents (Spivey-Knowlton and Sedivy, 1995).

We also aim to bridge the gap between models of interpretation in the tradition of the garden-path theory, which are related to symbolic approaches to NLP, and subsymbolic approaches in the tradition of parallel theories of sentence processing. Our model benefits from the “niceties” of the former approach to arrive at semantic and knowledge representations for alternative analyses while also leveraging a cognitive architecture that is suited to implement a parallel approach to interpretation.

7 Future work

Our future work will focus on studying the role agents will have in learning or refining new heuristics. As a matter of fact, we believe that the architecture is well-suited to mine the context-sensitive information that makes an analysis more plausible than another in a given discourse situation. We see

this as a machine learning process by which agents contribute to the common goal of “entropy reversal” by learning new heuristics and applying them during the incremental interpretation process.

Another aspect we will be focusing on in future work is a process we call “signaling,” which we shall illustrate using sentence (2):

(2) *Every ten minutes a man gets mugged in New York.*

Based on our interpretation algorithm, the most plausible analysis would be the one with the following representation.

Agent1:

\mathcal{ELF} :

$$\begin{aligned}
 &(\exists e:[e \text{ ((attr periodic)(plur episode))}] \wedge \\
 &(\forall et:[et \text{ member-of } e] \\
 &\quad [[(\text{period-of } et) = (\text{K} ((\text{num } 10) \\
 &\quad \quad \quad \quad (\text{plur minute})))] \wedge \\
 &\quad (\exists etl:[[etl \neq et] \wedge [etl \text{ overlap } et]]]) \wedge \\
 &\quad [et \text{ in-loc New York}] \wedge \\
 &\quad (\exists x:[x \text{ person}] \\
 &\quad \quad (\exists y:[y \text{ man}][x \text{ mug } y])) \\
 &\quad ** et])
 \end{aligned}$$

In the absence of any referential context that might indicate otherwise, our model does assign narrow scope to the existentially quantified expression introducing “a man” in the discourse model. Applying the heuristics, *Agent1* carries the most plausible interpretation in which there is an episode e consisting of a collection of periodically, non-overlapping subepisodes et , each one of them introducing a different individual getting mugged in New York, none of whom is salient in the immediate referential context in which sentence (2) is uttered.

Suppose that this fragment continues with sentence (3):

(3) *We are in New York today to interview him.*

As a result, our cognitive architecture gets “jungled” after processing sentence (3). In this case, *Agent1* is forced to come up with a single salient referent in its discourse model, corresponding to the poor individual who gets mugged every ten minutes in New York. *Agent1* is unable to provide such a referent.

By a process of signaling, agents cannot only be leveraged to keep a given analysis and corresponding interpretation active in memory for a given period of time, but also to “send” information, including referents, to other agents that might “request”

this information during their own interpretation process. We will be studying how this signaling process can be used to resolve anaphoric references and ensure discourse coherence. Our approach will consist in implementing Schubert's dynamic skolemization mechanism using this cognitive architecture.

Acknowledgments

Many thanks to Pat Langley for discussions on earlier versions of this paper, to Dan Shapiro for discussions on value alignment theory, and to Len Schubert for his continuous support teaching me EL.

References

- Altmann, G. and M. Steedman. 1988. Interaction with context during human sentence processing. *Cognition*, 30:191–238.
- Anderson, J. R. 1993. *Rules of the Mind*. Lawrence-Erlbaum, Hillsdale, NJ.
- Bacchus, F. 1990. *Representing and Reasoning with Probabilistic Knowledge*. MIT Press, Cambridge, MA.
- Balota, D.A., G.B. Flores d'Arcais, and K. Rayner. 1990. *Comprehension Processes in Reading*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Frazier, L. and C. Clifton. 1996. *Construal*. MIT Press, Cambridge, MA.
- Frazier, L. and J. A. Fodor. 1978. The sausage machine: A new two-stage parsing model. *Cognition*, 6:291–325.
- Hwang, C.H. and L.K. Schubert. 1992. Tense trees as the fine structure of discourse. In *Proceedings of the 30th Annual Meeting of the American Association for Computational Linguistics*, pages 232–240.
- Hwang, C.H. and L.K. Schubert. 1993. Episodic logic, a comprehensive, natural representation for language understanding. *Minds and Machines*, 3:381–419.
- Laird, J. E., editor, 1991. *Integrated cognitive architectures*, volume 2 (4) of *SIGART Bulletin*, pages 12–184. ACM.
- Laird, J. E., A. Newell, and P.S. Rosenbloom. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33:1–64.
- Langley, P., D. Shapiro, M. Aycinena, and M. Siliski. 2003. A value-driven architecture for intelligent behavior. In *Proceedings of the IJCAI-2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*, pages 10–18.
- Schaeffer, S., C.H. Hwang, J. de Haan, and L.K. Schubert. 1991. EPILOG: The computational system for episodic logic. Technical report, University of Alberta, Edmonton.
- Spivey-Knowlton, M.J. and J.C. Sedivy. 1995. Resolving attachment ambiguities with multiple constraints. *Cognition*, 55:227–267.