Representing and Reasoning over Time in a Unified Cognitive Architecture

David J. Stracuzzi (david.stracuzzi@asu.edu)
Nan Li (nan.li.3@asu.edu)
Gary Cleveland (gary.cleveland@asu.edu)
Pat Langley (langley@asu.edu)
School of Computing and Informatics
Arizona State University
Tempe, AZ 85287 - 8809 USA

Abstract

Most cognitive architectures have an implicit representation of time. As a result, reasoning about specific temporal relationships among events is typically beyond their capability. In this paper, we describe an extension of the ICARUS architecture to include an episodic belief memory, an explicit representation of temporal relationships, and associated reasoning processes. We then demonstrate the ensuing reasoning capabilities on a task that involves recognizing football plays. Finally, we discuss the implications of our temporal representation and reasoning mechanisms for the larger architecture.

Keywords: cognitive architectures; episodic memory; temporal logic; event recognition

Introduction

The ability to remember and reason about events over time is fundamental to human cognition. Tulving (1983, 2002) describes episodic memory as a temporal or contextual store that captures an individual's experiences. This history can then be used to improve decision making by forming part of an internal model of the environment, by keeping track of long-term goals, or by improving behavior through learning. Many cognitive tasks, such as problem solving (Howes, 1993) and discourse comprehension (Kintsch, 1998), also depend on storing and recalling information about the past.

In spite of the broad applicability of episodic memory and temporal reasoning, few efforts at constructing computational models of such capabilities have been made. Kolodner's (1993) early work on case-based reasoning is particularly relevant to episodic memory. Here, a case typically describes the solution to a previously encountered problem which the system can then retrieve and adapt to new problems. However, case structures typically do not generalize well and are usually hand-crafted for specific tasks.

In the context of cognitive architectures, Altmann and John (1999) added an episodic memory to Soar, although it was task specific and was not integrated into the larger architecture. More recently, Nuxoll and Laird (2007) integrated a general-purpose episodic memory module into Soar, and then implemented cognitive capabilities such as learning from past successes and failures on top of the new module. ACT-R (Anderson & Lebiere, 1998) also supports a limited form of episodic memory. The architecture's chunking mechanism stores partial copies of working memory for subsequent retrieval, but does not support retrieval of temporally related items, and does not distinguish between memories of prior events and beliefs about the present. Systems like these tend to be similar in their focus on storing, retrieving, and using entire episodes in support of cognitive tasks.

None of the aforementioned systems provide an explicit language or inference mechanism that lets them reason about temporal relationships among individual events or entities. In this paper, we extend the ICARUS architecture (Langley & Choi, 2006) to (1) represent and retain beliefs about past experiences, (2) encode general temporal relationships in long-term conceptual memory, and (3) reason about temporal relationships based on past and present beliefs. Moreover, we show that these extensions fit naturally into the existing architecture, and that they expand its capabilities substantially without the addition of new or sophisticated modules. We begin our discussion with a brief review of ICARUS, after which we describe our changes to the architecture, demonstrate their effects, and discuss their implications.

A Brief Review of ICARUS

The objective of the ICARUS architecture is to qualitatively model results on human cognition. It incorporates many ideas from traditional work on cognitive modeling, and maintains that cognition is closely tied to perception and action so that a model must be linked to some external environment. Like Soar (Laird, Rosenbloom, & Newell, 1986) and ACT-R (Anderson, 1993), ICARUS makes theoretical commitments to formalisms for memories, knowledge representation, and cognitive processes. For example, ICARUS shares the distinction between short-term and long-term memories, and goal-driven but reactive execution with several other architectures, but also includes many novel features such as a commitment to separate storage for conceptual and skill knowledge, and indexing skills by the goals they achieve.

In this section we briefly review representation, inference, and execution in ICARUS to provide a basis for our discussion of temporal representation and reasoning. In particular, ICARUS maintains a tight integration between inference and execution processes,[1] thus qualifying it as an instance of a *unified* cognitive architecture (Newell, 1990). As we will see, this helps to expand the power of the temporal representation beyond the conceptual memory and the inference mechanisms without requiring substantial modification to other modules in the architecture, such as execution or learning.

[1]This tight integration also holds for problem solving and learning in ICARUS, though we do not discuss these here.

Table 1: Non-temporal concepts from the football domain.

```
; Convert player activity perception into a belief
((agent-action ?agent ?action)
  :percepts    ((agent ?agent team OFFENSE action ?action)))

; Convert player direction of motion into a belief
((agent-direction ?agent ?dir)
  :percepts    ((agent ?agent direction ?dir team OFFENSE)))

; ?agent moved in direction ?dir
((moved ?agent ?dir)
  :relations   ((agent-action ?agent MOVE)
                (agent-direction ?agent ?dir)))
```

Like other architectures, ICARUS operates in cognitive cycles. A cycle begins when the agent perceives objects in the environment and places their descriptions into a short-term perceptual buffer. This initiates the inference process, which matches percepts against the structures stored in long-term conceptual memory, which contains a set of hierarchically organized logical rules. Each conceptual clause describes a class of environmental situations using a relational language similar to PROLOG; it includes a head, with the concept name and arguments, and a body that describes the conditions under which the concept applies.

The result of matching the percepts with concepts is a set of beliefs, which represent specific relational properties that hold in the current environment. The beliefs are stored in a short-term belief memory and then matched against other concept definitions in a bottom-up manner to produce new, higher-level (more abstract) beliefs. This process continues in a bottom-up manner until the architecture deduces all possible beliefs for the current environment state.

Table 1 shows three non-temporal concept definitions for football, which we use for illustration throughout the remainder of the paper. Symbols preceded by question marks indicate variables. The first two concepts extract information from the agent's perceptions of the current state, converting these into beliefs as shown in Table 2. For football, percepts include the identity, activity, and direction of each player on the field. The third concept recognizes the condition in which a given player is moving in a specific direction on the field by matching against lower-level beliefs about actions and directions. ICARUS does not maintain percepts and beliefs across cognitive cycles, which prevents it from reasoning about temporal events and relationships among players on the field.

After deducing the set of beliefs about the current state, ICARUS then uses its beliefs, combined with its goals and the structures contained in its long-term skill memory to determine which skills to apply in the environment to achieve these goals. Execution begins with a goal, which is a belief that the architecture wants to make true. Given a goal, the architecture finds a skill in long-term memory that both applies in the current state and achieves the goal.

Like conceptual memory, skill memory is organized hierarchically. Skills take a form similar to conceptual clauses; they have a head, which states the skill's objective, and a body, which states the environmental conditions required to initiate the skill, and the ordered actions or subgoals needed to achieve the skill's goal. After it finds an appropriate skill, the architecture must find a path through the subgoal hierarchy down to an executable action (atomic subgoal), ensuring that all of the intervening subgoals are applicable. If no such path exists, then ICARUS falls back on problem solving. We do not consider this case here, but Langley and Choi (2006) discuss problem solving in detail.

Table 2: Sample percepts with inferred non-temporal beliefs for the concepts shown in Table 1.

Percepts (cycle 1):
(agent QB team OFFENSE action WAIT direction 0)
Beliefs:
(AGENT-ACTION QB WAIT)
(AGENT-DIRECTION QB 0)

Percepts (cycle 2):
(agent QB team OFFENSE action MOVE direction S)
Beliefs:
(AGENT-ACTION QB MOVE)
(AGENT-DIRECTION QB S)
(MOVED QB S)

In the context of football, a skill with the head *(moved ?agent ?dir)* for achieving goals like *(MOVED QB S)* would refer directly to an executable action. More complex skills, such as those for running crossing receiver patterns (run n yards down field, then turn hard left or right), would use *moved* as a subgoal, thereby building complex behaviors from simpler ones.

Note the close correspondence between concepts and skills, as well as between beliefs and goals. This relationship figures centrally in the architecture's performance and learning mechanisms, and makes its various processes highly interdependent. For example, execution relies on inference to produce the beliefs that are matched against goals and skill preconditions. Thus, if ICARUS cannot infer that a specific temporal condition has been achieved, then it cannot determine whether a skill for achieving that temporal condition applies or has executed successfully. In the following section, we discuss an expansion of ICARUS' reasoning capabilities that begins to address this limitation.

Temporal Representation and Reasoning

Representing and reasoning over time plays an important role in a variety of cognitive tasks. For example, recognizing receiver patterns in football requires the ability to determine that certain events occurred in a specific order. However, past efforts at integrating episodic memories into cognitive architectures tended to result in either substantial modification of the existing modules or in the addition of entirely new modules (e.g. Nuxoll & Laird, 2007). Here, we outline a set of extensions that provide ICARUS with the ability to reason and execute over temporal structures. In particular, we draw at-

Table 3: Temporal concepts for the football domain.

```
; ?agent carried ?ball in the current time step
((possession ?agent ?ball)
  :percepts    ((ball ?ball carriedby ?agent)))

; ?agent caught ?ball
((caught-ball ?agent ?ball)
  :relations    (((possession AIR ?ball) ?air-start ?air-end)
                 ((possession ?agent ?ball) ?pos-start ?pos-end))
  :constraints ((eq ?air-end (- ?pos-end 1))))
```

tention to the ways in which the existing architecture provides a basis for temporal reasoning with only minor changes.

The first architectural modification focuses on encoding temporal information with beliefs. The representation of beliefs expands to include start time stamps, which indicate the first time at which a belief held, and end time stamps, which indicates the last time at which it held. ICARUS already maintains an internal notion of time, based on cognitive cycles, that may be used to set these time stamps. A special symbol, `NOW`, indicates the current time and distinguishes beliefs about the present from beliefs about the past. Thus, when a new belief is inferred, it receives a start time corresponding to the current cycle number and an end time of `NOW` until the first cycle in which the belief no longer holds. At that time, the end time stamp reverts to a specific cycle number. Percepts are not similarly time stamped, as perceptual memory continues to represent perceptions on the current cycle.

This augmented belief representation lets ICARUS distinguish beliefs about past events from ones about the present. The next extension is then to expand the temporal scope of belief memory by retaining all of the beliefs held throughout an episode. This is equivalent to providing the architecture with an episodic belief memory, whereas previously belief memory included only those beliefs that held on the current cycle. All beliefs contained in the episodic memory are generated through inference, which is based on the agent's percepts, so belief memory maintains a record of experiences in the environment *from the agent's perspective.*

The importance of episodic memory is well established, but the memory alone provides little improvement to an architecture's capabilities. To exploit this memory, two minor changes to the concept language are required. First, the `:relations` field, which lists the lower-level concepts that support a higher-level definition, expands to reference the time stamps assigned to beliefs. Second, we add a new `:constraints` field that represents simple arithmetic tests over time values referenced in the `:relations` field. Thus, this field lets ICARUS use temporal constraints as antecedents to concepts.

The architecture's inference process also expands to support the changes in belief and conceptual memories. The fundamental mechanism, which computes in a bottom-up manner the deductive closure of conceptual memory with the belief and perceptual memories, remains unchanged. The only difference is that the time stamps and temporal constraints must be matched in addition to the percepts and relations fields. No new specialized control is required.

Table 4: Percepts and corresponding temporal belief memory for concepts shown in Table 3.

...		
Percepts (cycle 124:		
(ball BALL1 carriedby AIR)		
(agent RB direction E team OFFENSE action MOVE)		
Beliefs:		
(POSSESSION QB BALL1)	1	98
(POSSESSION AIR BALL1)	98	NOW
Percepts (cycle 125):		
(ball BALL1 carriedby RB)		
(agent RB direction E team OFFENSE action MOVE)		
Beliefs:		
(POSSESSION QB BALL1)	1	98
(POSSESSION AIR BALL1)	98	124
(POSSESSION RB BALL1)	125	NOW
(CAUGHT-BALL RB BALL1)	125	NOW
...		

Table 3 shows the temporal concept *caught-ball*, which holds if the ball is in the `AIR` (thrown by the passer) during one cycle, and then in the possession of a player during the next cycle. Here, the `:constraints` field relates the end time stamps of the two subconcepts (relations). Note that this definition of *caught-ball* only holds for the one cycle in which the receiver first gains possession of the ball (end times differ by one). An alternative definition could relate the end time stamp of the ball in the `AIR` with the start time of the receiver's possession, thereby letting the concept match on every cycle after the initial catch. Preference in definition depends on how the concept will be used by higher-level concepts.

Table 4 shows the results of inference over the temporal concepts for two cycles, including the beliefs inferred during previous cycles. Notice the compact and temporally descriptive form of the beliefs. Three beliefs describe the history of ball possession over 125 cognitive cycles. In general, a single temporal belief state is sufficient to describe an entire episode up to that point. Also note the transition of the end time from `NOW` to a cycle number for *(POSSESSION AIR BALL1).* The end time for *(CAUGHT-BALL RB BALL1)* will similarly revert to 125 in the next cycle.

Looking beyond inference, execution also requires only minor changes to support the new temporal representation. Skill syntax requires no changes, but we add the assumption that preconditions (beliefs) required for a skill to either start or continue execution must hold in the current time step (end time stamp equal to `NOW`). No further changes to skills are necessary because skill heads (goals) correspond to the heads of defined concepts. The concept definitions therefore contain the temporal constraints needed to determine whether a skill executed successfully. This is a key benefit of the close relationship between inference and execution in ICARUS.

Figure 1: The pass play observed by ICARUS with annotations indicating actions taken by individual players.

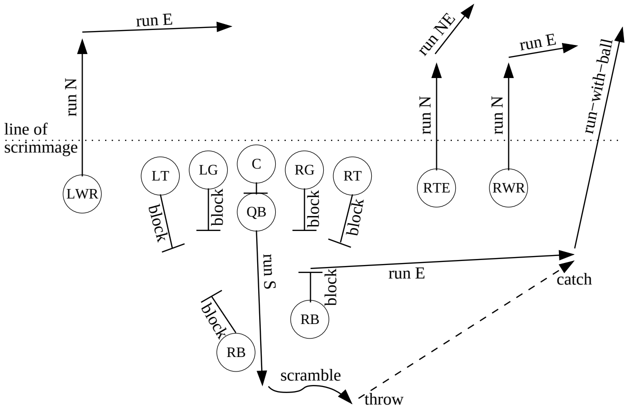

Figure 2: Diagram of observed play with annotations indicating higher-level goals of individual players and player units.

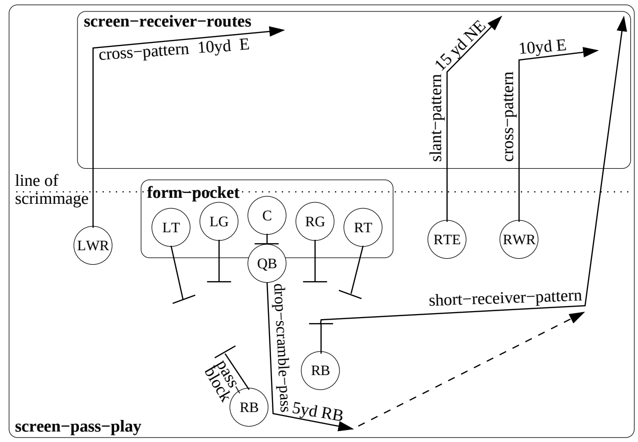

The representation of temporal beliefs in ICARUS is consistent with recent studies which suggest that neurogenesis in the hippocampus plays a role in encoding temporal information as a form of time stamp (Aimone, Wiles, & Gage, 2009). The extensions also increase the correspondence of ICARUS' belief memory with the notion of working memory, which typically includes items from the past and present that are currently being manipulated. Although the detailed mechanisms that we use to represent time in ICARUS are not psychologically plausible, they are consistent with our objective of qualitatively modeling human cognitive abilities. Specifically, our implementation of time stamps provides a more precise temporal reasoning ability, but it does not provide representational power beyond that of humans.

We have implemented the above modifications to the architecture and tested them extensively in the football setting. The problem solving and learning modules have not yet been fully revised to support temporal concepts and beliefs, so we do not focus on them here. However, we summarize the issues that arise in the discussion section. In the next section, we demonstrate the use of temporal concepts, beliefs, and inference in ICARUS by applying them to recognize football plays observed from video footage.

An Illustrative Example

The ability to remember past experiences and to relate them temporally to other experiences is critical in recognizing complex behaviors. Here we demonstrate ICARUS' ability to recognize such behaviors as they unfold over time. Specifically, we apply the architecture to interpret three football plays as observed in video footage from a college football game. The goal is for ICARUS to interpret the behavior of the players, both individually and as a team.

Figure 1 shows a play diagram of one of the offensive passing plays presented to ICARUS. Notice the sequential nature of the individual player behaviors, such as the running back (RB, right side) who first blocks, then runs east, catches the ball, and finally runs north with the ball until tackled. Also note the coordinated aspects of the play, such as between the quarterback (QB) and the running back, who perform very different activities, but time their activities such that the ball is caught as the running back completes his run to the east. Figure 2 shows a higher-level view of player behavior, and illustrates the type of interpretation that ICARUS must produce.

ICARUS assumes that low-level perceptual information, such as pixel-based video footage, has already been processed into a symbolic format. All domain objects must be described by some combination of symbolic and numeric attributes. We therefore rely on the results of video post-processing procedures (Hess & Fern, 2007; Hess, Fern, & Mortenson, 2007) to serve as the percepts. In this case, ICARUS perceives the identity, role (such as quarterback or running back), team (offense or defense), location, direction and current activity (such as moving or blocking) of each of the 22 players on the field in each video frame ($1/30^{\text{th}}$ second), along with information about the ball carrier.

We provided ICARUS with a set of 67 temporal concept definitions sufficient for interpreting the observed plays. Table 5 shows the results of applying ICARUS to the three plays. In all three cases, the architecture produced a set of beliefs consistent with the play, including the top-level classification of the entire coordinated sequence, such as *screen-pass-play* in Figure 2. The processing times are clearly slower than humans, although even human performance in this task is highly variable. Coaches and broadcast announcers can often interpret plays in real time, but most viewers rely on help from announcers and instant replay to see the details of a given play. We revisit the question of efficiency in the next section.

Table 5: Temporal inference results for three football plays.

Play	Frames	Duration	Beliefs	CPU
1	149	4.97 s	619	321.46 s
2	200	6.67 s	624	539.23 s
3	202	6.73 s	661	484.20 s

Discussion

Our use of numeric time stamps on beliefs has several benefits for ICARUS. First, it provides a unique form of episodic memory based on individual temporal beliefs instead of entire episodes. Second, time stamped beliefs support retrieval of temporally related items, rather than entire episodes. Third, the approach provides an explicit mechanism for reasoning about temporal relationships among individual beliefs. Finally, these capabilities allow for recognition and execution of event sequences not previously possible. For example, the patterns run by pass receivers in football require execution of a sequence of simpler motions with specific temporal relations among them. Previously, even if ICARUS executed such a sequence, it could not evaluate whether it had done so successfully either during or after execution. The revised concept language and belief memory provides support for both.

Numeric time stamps on beliefs also has implications for the representational power of ICARUS' concept language and inferred beliefs. First, they let the architecture avoid distinguishing between instantaneous and extended events, while still allowing agents to make such determinations. For example, an agent can have a concept that tests the equality of two time stamps to determine whether some belief held for only one cycle. Likewise, the architecture does not distinguish between ongoing and completed events, but an agent can do so simply by testing the value of a belief's end time stamp.

This approach is a strict departure from past work on time in agent architectures. For example, Allen (1984) relies on temporal intervals that are compared relationally, but does not specify specific times as end points. He maintains that this is important because it supports the notion that intervals and events are infinitely decomposable. While mathematically true, this idea is cognitively implausible. Limits on human perception imply that support for such capabilities at the architectural level is unnecessary. ICARUS' use of the cognitive cycle to determine time stamp values implements exactly this restriction in a qualitative manner.

Providing architectural support for time in ICARUS has so far been about generalizing the existing architecture, rather than about adding new modules and mechanisms. The knowledge representation expanded to accommodate temporal information in beliefs and concepts, and the belief memory expanded to include beliefs about the past, but no new structures or memories were required. Likewise, the revised inference process performs additional steps, but relies on the same fundamental procedures. The execution module requires no modification, relying instead on information passed through concepts and beliefs to achieve temporal goals.

Looking deeper into the architecture, the next steps of integrating temporal capacity into the learning and problem solving modules should similarly be matters of generalization. Each module depends on both concepts and skills, so the parts of the modules that depend on concepts must be modified to use the information contained in the temporal constraints. Specifically, these constraints will inform the partial order in which subgoals should be considered (problem solving) or stored (skill learning). Aspects of problem solving and learning that depend on skills should not require substantial change. Further research is needed to determine the details of the integration, but we do not anticipate any major changes to the content of the architecture.

The relatively uncomplicated integration of temporal representation and reasoning capabilities into ICARUS suggests that some of the architecture's other assumptions and commitments are also beneficial. In particular, the distinction between conceptual and skill memories substantially simplifies the integration by separating the potentially complex temporal constraints and associated reasoning issues from the skill knowledge that uses the inferred beliefs. Likewise, the close relationship between the two types of knowledge, and the strong interdependence between inference and execution allows both modules to exploit the temporal information.

A noted earlier, one temporal belief state is sufficient to reconstruct the sequence of events that led to that state within the limits of the concept hierarchy. This is consistent with Bartlett's (1932) theory of reconstructive memory, which states that only some information about the past is available in memory and the mind reconstructs the missing parts. ICARUS' ability to remember perfectly *all* beliefs is not psychologically plausible, and one area of future work is to add a mechanism for forgetting. Bartlett's theory suggests that detailed beliefs (lower-level in the context of ICARUS) tend to be lost and reconstructed while the more abstract, big-picture beliefs that form the core of an experience are retained. Such a process in ICARUS would let symbols in the arguments of high-level beliefs flow down through the hierarchy toward the lower-levels. However, this may not bind symbols to all low-level concept arguments, so additional reasoning would be required.

A second avenue for future work relates to the intentions of an agent with respect to execution. Currently, the inference process does not have access to current goals or to those that were achieved or abandoned in the past. Generating new temporal beliefs that represent the intentions would let ICARUS reason about past goals and current goals. The addition of time-stamped intentions to belief memory would make a new class of goals available to the architecture. For example, the goal *work on homework until dinner is ready* states that the agent should maintain the intention to complete homework (which implies execution of skills for completing homework) until a specific event is satisfied. This is distinctly less restrictive of an agent's behavior than a goal of *complete homework before dinner*.

The retrieval of beliefs from the episodic memory is another possible line for future development. Currently, ICARUS uses the same pattern-matching process that it uses for temporal beliefs. In practice, the temporal belief memory holds far more information than in earlier versions. As a result, the cost of matching (inferring) concepts grows with the number of temporally distinct beliefs added to the memory, although this is far less than the cross-product of beliefs with states. Soar reduces the computation required to determine the relevance of an episode by focusing first on its most recent experiences (Nuxoll & Laird, 2007). A similar mechanism may be beneficial for ICARUS, although the matching details would be different since it would retrieves individual beliefs rather than entire episodes.

Finally, a related issue concerns the architecture's approach of processing each perceptual state in its entirety, regardless of the amount of processing time available. In the case of play recognition, even coaches may be unable to recognize all details of a play in real time. Instead, they process the most salient features of the play during the initial viewing, and then focus on finding more detailed behaviors during subsequent reviews. Time-sensitive application of conceptual knowledge and inference is particularly important in the context of a temporal belief memory, as the volume of information available is large. This suggests that we incorporate a utility-based inference process that focuses on concepts with higher utility first, while low utility concepts receive attention only if time permits.

Concluding Remarks

Remembering past experiences and reasoning about relationships over time are a fundamental cognitive abilities that humans rely on for a variety of tasks. However, very few cognitive models or intelligent systems have been developed to model this capability. In this paper, we showed how to integrate an explicit representation of time and a temporal reasoning mechanism into the ICARUS architecture. The resulting temporal belief memory serves as an episodic store, and the architecture's ability to refer to past beliefs individually supports a finer-grained episodic memory than other accounts.

We also argued that our approach is functionally adequate, and that the relatively simple integration of temporal reasoning into ICARUS suggests that other aspects of the architecture are also beneficial. Substantial evaluation will be required to confirm these points, but our initial tests and demonstrations have been encouraging. Finally, the integration of temporal reasoning capabilities into ICARUS opens a wide variety directions for future research on the architecture.

Acknowledgments

This material is based on research sponsored by DARPA under agreement FA8750-05-2-0283. The U. S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed on implied, of DARPA or the U. S. Government.

References

Aimone, J. B., Wiles, J., & Gage, F. H. (2009). Computational influence of adult neurogenesis on memory encoding. *Neuron*, *61*(2), 187–202.

Allen, J. F. (1984). Towards a general theory of action and time. *Artificial Intelligence*, *23*, 123–154.

Altmann, E. M., & John, B. (1999). Episodic indexing: A model of memory for attention events. *Cognitive Science*, *23*, 117–156.

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge, UK: Cambridge University Press.

Hess, R., & Fern, A. (2007). Improved video registration using non-distinctive local image features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recogntion.* Minneapolis, MN: IEEE Press.

Hess, R., Fern, A., & Mortenson, E. (2007). Mixture-of-parts pictorial structures for objects with variable part sets. In *Proceedings of the Eleventh IEEE International Conference on Computer Vision.* Rio de Janeiro, Brazil: IEEE Press.

Howes, A. (1993). Recognition-based problem solving. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society* (pp. 551–556). Hillsdale, NJ: Erlbaum.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.

Kolodner, J. (1993). *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, *1*, 11–46.

Langley, P., & Choi, D. (2006). A unified cognitive architecture for physical agents. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence.* Boston: AAAI Press.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Nuxoll, A., & Laird, J. E. (2007). Extending cognitive architecture with episodic memory. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence* (pp. 1560–1565). Vancouver, Canada: AAAI Press.

Tulving, E. (1983). *Elements of episodic memory*. Oxford: Clarendon Press.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, *53*, 1–25.