Representing and Reasoning over Time in a Symbolic 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

Abstract

Symbolic cognitive architectures typically have an implicit representation of time. Actions taken occur in a sequential order, but reasoning about specific temporal relationships among objects and events is typically beyond the architecture's capabilities. In this paper, we describe an extension of the Icarus architecture to include an episodic belief memory and an explicit representation of temporal relationships in long-term conceptual memory, along with the inference process that supports them. We then demonstrate the temporal reasoning system on an American-rules football play-recognition task. Finally, we discuss the implications of our temporal representation and reasoning system to the larger architecture.

Keywords: cognitive architectures; episodic memory; temporal logic

Introduction

The ability to remember and reason about objects and events over time is fundamental to human cognition. Tulving (1983, 2002) describes episodic memory as a temporal or contextual memory that captures the experiences of an entity. 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 determining efficient search strategies (Howes, 1993) and discourse comprehension (Kintsch, 1998), also depend on the ability to store and recall information about past experiences.

In spite of the broad applicability of episodic memory and temporal reasoning, few efforts at constructing computational models of such capabilities have been made. In the artificial intelligence community, case-based reasoning (Kolodner, 1993) is the work most 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 (2004, 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. All of these systems are similar in that they 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 allows them to reason about specific temporal relationships among individual events, entities or objects in the environment. In this paper, we extend the ICARUS architecture to (1) represent and retain beliefs about past experiences, (2) represent specific temporal relationships in long-term conceptual memory, and (3) reason about temporal relationships based on the stored beliefs about the past and present. Moreover, we show that these extensions fit naturally into the existing architecture, and that they expand ICARUS' capabilities substantially without the addition of new or sophisticated modules.

We begin our discussion with a brief review of ICARUS. We then describe the changes to the architecture required to support the representation of and reasoning over temporal structures. Next, we provide a demonstration of temporal reasoning in ICARUS on a football play-recognition task. We then move on to discuss the implications of the representational and reasoning capabilities on the broader architecture, including a discussion of related and future work. Finally, we conclude with a few remarks about the lessons that may be drawn from this work.

A Brief Review of ICARUS

ICARUS is a physical agent architecture whose objective is to qualitativly model results on human cognition. The architecture incorporates many ideas from traditional work on cognitive modeling, and also maintains that cognition is closely tied to perception and action so that a model must be linked to some external environment. ICARUS shares many features with other agent architectures like Soar (Laird, Rosenbloom, & Newell, 1986), ACT-R (Anderson, 1993), and Prodigy (Minton, 1990), such as a distinction between short-term and long-term memories, and goal-driven but reactive execution. ICARUS also includes many novel features including a commitment to separate storage for conceptual and skill knowledge, and indexing skills by the goals they achieve. In this section we review knowledge representation, inference and execution in ICARUS to provide a basis for our discussion of temporal representation and reasoning. Langley and Choi (2006) provide a more detailed description of the architecture, including its support for problem solving and learning.

Like other architectures, ICARUS operates in cognitive cycles. On each cycle, it perceives objects in the environment and places their descriptions into a short-term perceptual buffer. The architecture then performs inference based on these percepts by matching them against the hierarchical structures stored in long-term conceptual memory, thus letting the architecture deduce a set of beliefs consistent with both the environment and its conceptual knowledge. ICARUS then matches these beliefs against the knowledge structures stored in long-term skill memory to determine which skills to apply in order to achieve its current goal. After each executed action, ICARUS perceives a new state, and begins a new cycle.

Long-term conceptual memory in ICARUS contains a set of hierarchically organized logical structures. Each concept describes a class of environmental situations using a relational language similar to PROLOG. Each first-order clause includes a head, with the concept name and arguments, and a body that describes the conditions under which the concept applies. Inference then matches the concept definitions against the contents of perceptual memory to form beliefs, which are simply concept instances in which the arguments have been replaced with specific symbols from perceptual memory. These beliefs are then matched against other concept definitions in a bottom-up manner to produce new, higher-level beliefs. This process continues until the architecture deduces all possible beliefs for the current environment state.

Execution in ICARUS begins with a goal, which is simply a belief that the architecture wants to make true.[1] Given a goal, the architecture attempts to find a skill in long-term memory that both applies in the current state and achieves the goal. Skills take a form similar to conceptual clauses; they have a head, which states the skill's objective and corresponds to the head of some concept, and a body, which states the conditions that must be present in order to initiate or continue executing the skill along with the actions that must be taken or subgoals that must be achieved in order to achieve the skill's goal. Like conceptual memory, skill memory is organized hierarchically. After an appropriate skill is found, the architecture must find an applicable path through the subgoal hierarchy down to an executable action, ensuring that all of the intervening subgoals have their start conditions satisfied. If no such path exists, then ICARUS falls back on problem solving, which we do not consider here.

Note the close correspondence between concepts and skills, and between beliefs and goals. This relationship figures centrally in the architecture's performance and learning systems, and makes the goal processing, execution and inference procedures highly interdependent. Execution relies on goal processing to determine when goals have been achieved. Both procedures rely on inference to produce the beliefs that the architecture matches against goals and skill conditions. The tight integration of the inference and execution modules thus qualifies ICARUS as 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 procedures without requiring substantial modification to other modules in the architecture, such as execution or learning.

[1]Note that unlike the inferred beliefs in memory, a goal may leave arguments unbound.

Temporal Representation and Reasoning

Representing and reasoning over time, particularly in the context of an episodic memory, plays an important role in a variety of cognitive tasks. However, past efforts at including episodic memories into cognitive architectures tended to result in either substantial modification of the existing modules or in the addition of entirely new architectural modules as in Soar (Nuxoll & Laird, 2004). In the following, we outline a set of natural extensions that provide ICARUS with the ability to represent and reason over time. In particular, we draw attention to the ways in which the representation and mechanisms of the existing architecture made such a major expansion in capability possible given only minor changes to the architecture.

Representing and reasoning over time in ICARUS requires several minor changes to the architecture. First, ICARUS must maintain some notion of the current time in the environment. Second, the belief representation expands to (1) annotate beliefs with the time periods over which they held (2) retain beliefs held in the past. Next, the concept language is augmented to exploit the temporal properties of beliefs, and to encode temporal constraints as concept antecedents. Finally, the inference process must expand to account for the added representational complexity.

Our initial implementation of temporal support in ICARUS assumes that the architecture receives the current time as a percept. The architecture itself does not maintain an internal sense of passing time. The perceived time may trivially correspond to the cycle number, but this is not required. The only requirement is that the perceived time increase monotonically from one cycle to the next.

Given the availability of the current time, the next step is to expand the representation of beliefs to include start and end time stamps. The start time stamp indicates the first time at which a belief held true, while the end time stamp indicates the last time at which the belief held continuously. The architecture uses a special symbol, `NOW` to indicate the current time. Notice that the time stamps do not indicate the period over which an event occurred, but only time over which the belief in that event held. Percepts are not similarly time stamped, as perceptual memory continues to represent the ICARUS' perceptions in the current time step only.

This augmented belief representation allows ICARUS to distinguish beliefs about past events from beliefs about the present. Now we can expand belief memory to retain all of the beliefs held by ICARUS throughout an episode. This is equivalent to providing the architecture with an episodic belief memory, whereas previously belief memory was updated

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

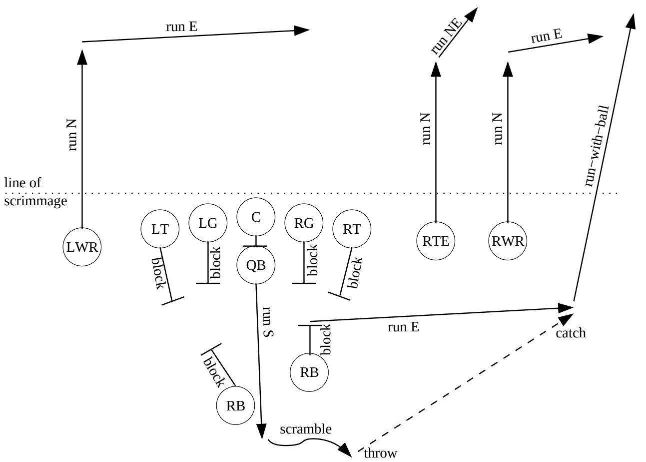

Figure 2: Diagram of observed play with annotations indicating higher-level goals of individual players and player units.

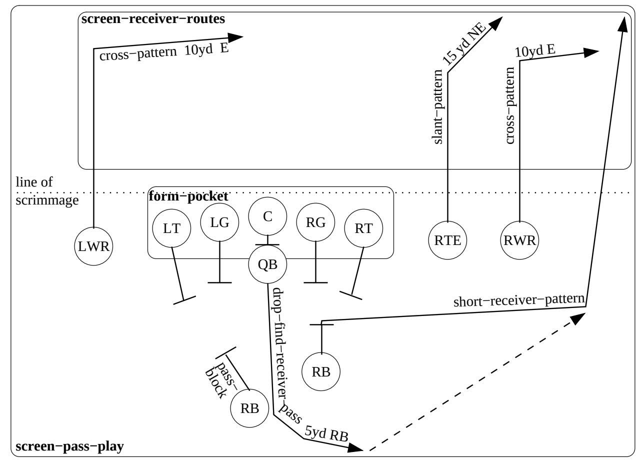

on each cycle to include only those beliefs that held on the current cycle. All of the beliefs contained in the episodic memory are generated through inference, which is based on the agent's percepts, so belief memory maintains a record of experiences in the environment *from ICARUS' perspective.*

The importance of episodic memory is well established, but the memory alone provides little improvement to an architecture's capabilities. To exploit the episodic 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 mathematical operations on and comparisons of specific time stamp values referenced in the `:relations` field. Thus, the field lets ICARUS use time constraints as antecedents to concepts.

The final modification to the architecture is the expansion of its inference process to support the changes to belief and conceptual memory. The fundamental procedure, which is a bottom-up computation of the deductive closure of conceptual long-term memory with the belief and perceptual memories, remains unchanged. Likewise, the matching process used to determine whether a particular concept instance (belief) should be inferred also remains unchanged. The only differences between the existing inference process and the revised process are that (1) the time stamps and temporal constraints must be matched in addition to the percepts, relations and tests fields, and (2) temporally adjacent instances of the same belief are merged into a single temporal belief that covers both. No new specialized control is required.

Looking beyond inference, the execution module also requires only minor changes to support the new concept and belief representations. Skill syntax requires no changes, but we add the assumption that conditions (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 define the content of the skill body. This is a key benefit of the close relationship between inference and execution in ICARUS.

The problem solving and learning modules have not yet been revised to support temporal concepts and beliefs, so we do not comment on them here. However, we discuss the issues that arise in the discussion section. In the following section, we demonstrate temporal concepts, beliefs and inference in ICARUS by using 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 the temporal inference module's ability to recognize complex behaviors as they unfold over time. Specifically, we apply ICARUS to interpret the offensive passing play diagrammed in Figure 1 as executed by an eleven player college football team and observed as video footage from an overhead camera.

Figure 1 shows the sequence of actions performed by individual players. ICARUS perceives information about each player on the field (including the defense, which is not shown in the diagram) in each video frame, or every $1/30^{\text{th}}$ of a second. The goal is for ICARUS to interpret the behavior of the players, both individually and as a team. 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. Specifically, 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, rather than the

Table 1: Sample concepts from the football domain.

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

; 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) ?act-start NOW)
                ((agent-direction ?agent ?dir) ?dir-start NOW)))

; ?agent has possession of ?ball
((possession ?agent ?ball)
  :percepts    ((ball ?ball carriedby ?agent)))

; Some ?agent had possession of ?ball in a previous time step
((possession-before-now ?ball)
  :relations   (((possession ?agent ?ball) ?start ?end))
  :constraints ((< ?start NOW)))

; ?agent received the snap of ?ball. True if:
; (1) ?agent has just received possession of ?ball, and
; (2) no other player has had possession of ?ball
((snap ?agent ?ball)
  :relations   (((possession ?agent ?ball) NOW ?p-end)
                (not (possession-before-now ?ball))))

; ?passer dropped back after receiving the snap
((dropped-back ?passer ?ball)
  :relations   (((snap ?passer ?ball) ?snap-start ?snap-end)
                ((possession ?passer ?ball) ?poss-start NOW)
                ((moved ?passer S) ?mov-start NOW))
  :constraints ((≤ ?snap-end ?poss-start)))
```

Table 2: Percepts and inferred beliefs given Table 1 concepts.

```
Percepts:
   (time 115)
   (ball BALL1 carriedby QB)
   (agent QB direction S team OFFENSE action RUN)
Beliefs:
   (AGENT-ACTION QB MOVE)                115   NOW
   (AGENT-DIRECTION QB S)                115   NOW
   (MOVED QB S)                          115   NOW
   (POSSESSION QB BALL1)                 115   NOW
   (SNAP QB BALL1)                       115   NOW
   (DROPPED-BACK QB BALL1)               115   NOW

Percepts:
   (time 116)
   (ball BALL1 carriedby QB)
   (agent QB direction S team OFFENSE action RUN)
Beliefs:
   (AGENT-ACTION QB MOVE)                115   NOW
   (AGENT-DIRECTION QB S)                115   NOW
   (MOVED QB S)                          115   NOW
   (POSSESSION QB BALL1)                 115   NOW
   (POSSESSION-BEFORE-NOW BALL1)         116   NOW
   (SNAP QB BALL1)                       115   115
   (DROPPED-BACK QB BALL1)               115   NOW
   ...

Percepts:
   (time 177)
   (ball BALL1 carriedby QB)
   (agent QB direction S team OFFENSE action RUN)
Beliefs:
   (AGENT-ACTION QB MOVE)                115   NOW
   (AGENT-DIRECTION QB S)                115   NOW
   (MOVED QB S)                          115   NOW
   (POSSESSION QB BALL1)                 115   NOW
   (POSSESSION-BEFORE-NOW BALL1)         116   NOW
   (SNAP QB BALL1)                       115   115
   (DROPPED-BACK QB BALL1)               115   NOW

Percepts:
   (time 178)
   (ball BALL1 carriedby QB)
   (agent QB direction S team OFFENSE
        action PREPARE-TO-PASS)
Beliefs:
   (AGENT-ACTION QB MOVE)                115   177
   (AGENT-ACTION QB SCRAMBLE)            178   NOW
   (AGENT-DIRECTION QB S)                115   NOW
   (MOVED QB S)                          115   177
   (POSSESSION QB BALL1)                 115   NOW
   (POSSESSION-BEFORE-NOW BALL1)         116   NOW
   (SNAP QB BALL1)                       115   115
   (DROPPED-BACK QB BALL1)               115   177
   ...
```

raw video itself. 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 running or blocking) of each of the 22 players on the field, along with information about the ball carrier and current time.

We also provided ICARUS with a set of 44 temporal concept definitions sufficient for interpreting the observed play. Table 1 lists a subset of the concepts which are sufficient for interpreting the quarterback's (QB) drop-back after receiving the snap. This corresponds to the first leg of his motion in Figures 1 and 2.

Notice that there are two general types of temporal constraints used in these concepts. One type establishes that some condition holds in the present, such as the passer having possession of the ball and moving south in *dropped-back*. The second type of constraint establishes an ordering among events, such as in the definition of *dropped-back*, in which the passer must receive the ball snap prior to taking possession of the ball. This precludes the center (C), who snaps the ball to the quarterback, from taking the ball himself and moving backward. Other types of constraints, such as requiring specific events to occur at specific times, are possible but we have found them to be less useful in interpreting behavior.

Table 2 shows the available percepts and beliefs produced by running inference over the concepts shown in Table 1. Notice that a single temporal belief state is sufficient to reconstruct the sequence of events that led to the current state, although not all of the details would necessarily be available. To interpret the entire 6.73 second play (202 frames) based on the 44 provided concepts, ICARUS requires 55.4 CPU seconds to generate a total of 185 temporal beliefs. This is 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 many viewers rely on help from announcers and instant replay to see the details of a given play. We revisit the question of efficiency in

Table 3: Concept definition for *drop-back-completed*.

```
; ?passer has finshed dropping back after receiving the snap
((drop-back-completed ?passer ?ball ?n-steps)
  :relations    (((snap ?passer ?ball) ?snap-start ?snap-end)
                 ((possession ?passer ?ball) ?poss-start ?poss-end)
                 ((moved-distance ?passer ?n-steps S)
                  ?mov-start ?mov-end))
  :constraints ((≤ ?snap-end ?poss-start)
                 (≤ ?mov-end ?poss-end)))
```

the next section.

Just as there are different types of temporal constraints that than can be used in concept definitions, there are also different types of beliefs that arise. Some beliefs indicate the occurrence or completion of events. For example, *snap* indicates the cycle on which the ball snap was completed. Other beliefs indicate the time period over which some event occurred. For example, *dropped-back* indicates the period of time during which the quarterback had possession of the ball and moved backward from the line of scrimmage. The difference often arises from the specific constraints employed, although additional relations may also be required to recognize that some activity or event has completed.

This distinction is important in the context of execution, however. Skills represent methods for achieving particular beliefs, and ICARUS stops executing a skill when its goal is achieved. This means that if ICARUS were to execute a skill for *dropped-back*, then the skill would stop executing on cycle 115, when the quarterback first takes possession of the ball and starts moving backward. In practice, the skill must continue executing until the quarterback completes his drop-back. A separate concept, *drop-back-completed*, which is shown in Table 3 and only evaluates to true after the quarterback completes the entire drop-back, is required to serve as the head of the skill.

Discussion

The integration of temporal memory and reasoning into ICARUS is more a question of generalizing the existing architecture than of adding new modules and mechanisms. The knowledge representation expands to accommodate temporal information, but no new structures or memories are required. Likewise, the revised inference module 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, integrating the new temporal capacity into the learning and problem solving modules should similarly be questions of generalization. Each module depends on both concepts and skills. In both cases, the parts of the modules that depend on concepts must be modified to use the information contained in the temporal constraints. Specifically, the constraints will inform the partial order in which subgoals should be stored (skill learning) or considered (problem solving). The portions of the modules that depend on skills will 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.

As noted earlier, a single temporal belief state is sufficient to reconstruct the sequence of events that led to the current state, although some details may be missing. 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. However, ICARUS' ability to remember perfectly *all* beliefs throughout an episode is not psychologically plausible. One area of future work then is to add a mechanism for forgetting to belief memory. Bartlett's theory suggests that more detailed beliefs tend to be lost and reconstructed while the more abstract, big-picture beliefs that form the core of an experience are retained.

There are several other avenues for future work with respect to temporal beliefs and concepts in ICARUS. One such area relates to the intentions of an agent with respect to execution. Currently , ICARUS does not have access to goals that were either achieved or abandoned in the past. Allowing the execution engine to generate new temporal beliefs that represent the intentions of the agent lets ICARUS both know and reason about past goals. The additions of time stamped intentions to belief memory also makes a new class of goals available to the architecture. For example, the goal *work on homework until dinner is ready*, states that ICARUS 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*.

A second line of future work regards the retrieval of beliefs from the episodic memory. Currently, ICARUS uses the same pattern matching process that it used prior to switch to temporal belief memory. In practice, the temporal belief memory holds far more information than in earlier versions. As a result, the cost of matching (inferring) concepts grows steadily as beliefs get added to the memory. Soar employs a recency-biased retrieval mechanism that helps to reduce the amount of computation required for determining whether an episode is relevant (Nuxoll & Laird, 2007). A similar mechanism may be beneficial for ICARUS, although the matching details would be different since ICARUS retrieves individual beliefs rather than entire episodes.

A related issue concerns the architecture's current 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 of the details of a given 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 beliefs stored and retrieved is significantly increased. This suggests that ICARUS requires a utility-based inference process whereby concepts with higher utility are applied first, and low utility concepts only get applied to the current stimulus if time permits.

Concluding Remarks

The ability to remember past experiences and to reason about relationships over time is a fundamental cognitive capability 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. ICARUS' temporal belief memory is equivalent to an episodic memory, and the architecture's ability to refer to past beliefs individually rather than only in the context of a larger episode makes our implementation of episodic memory more flexible than others. We also argued that the relatively simple integration of temporal reasoning into ICARUS suggests that other aspects of the architecture are also particularly beneficial. Substantial evaluation will be required to confirm these points, but our initial tests and demonstrations are encouraging. Finally, the integration of temporal reasoning capabilities into ICARUS opens a wide variety directions for future research on the architecture.

Acknowledgments

The authors recognize Alan Fern, Robin Hess, William Brendel and Jervis Pinto for their efforts in converting video into the symbolic descriptions that we used in our demonstrations. 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

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 *Preceedings of the twenty-first national conference on artificial intelligence.* Boston: AAAI Press.

Minton, S. N. (1990). Quantitative results concerning the utility of explanation-based learning. *Artificial Intelligence*, *42*, 363–391.

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

Nuxoll, A., & Laird, J. E. (2004). A cognitive model of episodic memory integrated with a general cognitive architecture. In *Proceedings of the sixth international conference on cognitive modeling* (pp. 220–225). Pittsburgh, PA: Erlbaum.

Nuxoll, A., & Laird, J. E. (2007). Extending cognitive architecture with episodic memory. In *Preceedings 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.