

Algorithms for Scaling in a General Episodic Memory (Extended Abstract)

Nate Derbinsky

Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109-2121 USA
nlderbin@umich.edu

Justin Li

Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109-2121 USA
justinnh@umich.edu

John E. Laird

Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109-2121 USA
laird@umich.edu

ABSTRACT

Episodic memory endows autonomous agents with useful cognitive capabilities. However, for long-lived agents, there are numerous unexplored computational challenges in supporting useful episodic-memory functions while maintaining real-time reactivity. This paper presents and summarizes the evaluation of an algorithmic variant to the task-independent episodic memory of Soar that expands the class of tasks and cues the mechanism can support while remaining reactive over long agent lifetimes.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning

General Terms

Algorithms, Measurement, Performance, Experimentation

Keywords

Computational architectures for learning, Single agent learning

1. INTRODUCTION

Prior research has shown that autonomous agents with episodic memory, a task-independent, autobiographical store of prior experience [11], are more capable in problem solving, both individually [5][10] and with other agents [3][9]; better account for human psychological phenomena, such as those relating to memory blending [1] and emotional appraisal [4]; and are more believable as virtual characters [4] and long-term companions [8].

However, little work examines the computational challenges associated with maintaining effective and efficient access to experience over long periods of time. Most approaches to storing and retrieving episodic knowledge are task-specific (e.g. [9]) or apply to temporally limited problems (e.g. [5]).

By contrast, the episodic memory that is part of Soar [6] is task-independent and has been applied to complex, temporally extended tasks, such as action games [2] and mobile robotics [7]. To support effective and efficient episodic operation, the current mechanism makes specific design decisions within a space of algorithmic options. This paper presents and summarizes the evaluation of an algorithmic variant that expands the tasks and cues the mechanism can support while remaining reactive over long time periods, without adversely affecting performance.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4–8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2. EPISODIC MEMORY IN SOAR

Soar’s episodic memory [10] comprises three phases: (1) automatically *encoding* agent state; (2) *storing* this information as episodic knowledge; and (3) supporting *retrieval* at a later time.

The state of a Soar agent is represented as a connected di-graph. Episodic memory automatically encodes and permanently stores changes to this graph. Agents can later retrieve an episode by constructing a cue: a directed, connected, acyclic graph that specifies task-relevant relations and features. The *cue-matching* process identifies the “best” matching episode: the most recent episode that has the greatest number of structures in common with cue leaf nodes. Episodic memory then reconstructs this episode within a pre-specified region of the agent-state graph.

The cue-matching process (a) returns an episode if one exists that contains at least one feature in common with a cue leaf and (b) returns the “best” episode with respect to cue structure, cue leaves, and temporal recency. In the worst case, the encoding, storage, and retrieval operations scale at least linearly with state changes. However, exploiting regularities in state representation and dynamics may improve expected performance.

The current episodic-memory mechanism [2] exploits two regularities of agent state, both of which have been applied in the rule-matching literature. The first is *temporal contiguity*: agent-state changes between episodes will be few relative to agent-state size. The second is *structural regularity*: agent knowledge will reuse representational structure, and so over time, the number of distinct structures will be much smaller than the total experienced. Soar’s episodic memory exploits these assumptions. Episodic knowledge is captured in a dynamic-graph index, composed of (1) a global structure, termed the Working-Memory Graph (WMG), which captures all *distinct* graph edges that have been encoded, and (2) a set of temporal intervals that capture when each edge of the WMG was added to/removed from agent state. The cue-matching algorithm uses a subset of the WMG as a discrimination network (termed the DNF Graph), through which it streams relevant *changes*, such as to evaluate episodes relative to the cue.

3. NOVEL ALGORITHMIC VARIANT

Our algorithmic variant exploits a stronger form of the structural-regularity assumption: over long agent lifetimes, the number of distinct structures represented within a *single* episode is likely to be much smaller than the *total* number of distinct structures. The algorithm exploits this assumption by building the DNF Graph incrementally, adding edges when they are relevant and removing them as they become obsolete. We hypothesized that over long agent lifetimes, this algorithm would improve retrieval time, especially for cues that match relatively recent episodes; however, tradeoffs exist. First, extra computation is required to dynamically

maintain the DNF Graph, which may exceed any performance gains, especially for simple cues. Second, the storage process must encode additional information: the most recent episode during which each WMG edge was represented in agent state.

4. EVALUATION

We implemented our algorithmic variant in Soar v9.3.1 and evaluated agents that use episodic memory for hours to days of real time. We applied over 100 cues in numerous tasks, spanning word-sense disambiguation, 44 instances of 12 planning domains (e.g. *Grid* and *Logistics*), 3 video games, and mobile robotics.

To evaluate scaling, we measured the time to perform episodic operations. For cue matching, we instrumented Soar to perform this operation 100 times for each cue at regular intervals. All experiments were performed on a Xeon L5520 2.26GHz CPU with 48GB RAM running 64-bit Ubuntu v10.10.

Our algorithm did not greatly impact performance in most tasks; however, it did enable a new, general capability for long-running agents using episodic memory: the management of long-term goals. This capability is best illustrated in the mobile-robotics domain, which has been used in prior work both in simulation and on physical hardware [7]. The agent perceives both physical perception data and symbolic representations of objects, rooms, and doorways. The agent’s task is to explore a building, consisting of 100 offices, and then execute a fixed patrol pattern. While performing these tasks, the agent builds an internal map, which it uses for path planning and navigation. One cue in this domain asks, “*When was my desired destination doorway #5?*” The agent could examine episodes that followed this retrieval to recall progress made towards that goal. We ran the agent for 12 hours of real-time operation and measured performance every 300K episodes (~10 min.).

Figure 1 shows timing data to evaluate this cue, comparing maximum cue-matching time (in msec.) between Soar’s current algorithm (“baseline”) and our algorithmic “variant.” Both algorithms exhibit growth in cue-matching time because some features in the cue are relevant with each new goal the agent encodes, but the goal of interest is increasingly distant in time. The first difference between the data sets is the number of episodes encoded over the 12-hour period: whereas the baseline algorithm encoded over 58 million, our variant encoded nearly 109 million. This difference has to do with optimizations we implemented in the incremental episodic-encoding algorithm, which resulted in an average of more than 50% improvement in encoding/storage speed in this task. However, both algorithms exhibit a common shift in behavior when the agent has finished exploring the building and proceeds to execute a patrol (~8M




Figure 1. Timing comparison for goal-management cue in the mobile-robotics evaluation task.

episodes for baseline, ~12M for variant). Before this point, the agent encodes new navigation goals much more frequently than after, and so the maximum search time grows more slowly after this point. Before the shift, our variant grows 3.6x slower than the baseline, and after it grows 4.9x slower. Furthermore, in fewer than 12 hours, the maximum computation time for the baseline algorithm grows above 50 msec., a level of reactivity that has been established in numerous domains, including video games, robotics, and HCI. By contrast, given the rate of growth in this task, our variant algorithm can continue to provide reactive real-time cue-matching performance for nearly 694M episodes (> 3 days of real time). The goal-management cue in the mobile-robotics domain is just one instance in a class of episodic cues and tasks in which there is a growth of distinct structures over the agent’s lifetime. This data provides evidence that our algorithmic variant expands the problems in which Soar’s episodic memory can support useful operation while agents remain reactive in dynamic environments over long time periods.

5. ACKNOWLEDGMENTS

The authors acknowledge the funding support of the Air Force Office of Scientific Research under contract FA2386-10-1-4127.

6. REFERENCES

- [1] Brom, C., Burkert, O., Kadlec, R. 2010. Timing in Episodic Memory for Virtual Characters. In Proc. of the IEEE 2010 Conf. on Computational Intelligence and Games, Copenhagen, Denmark, 305-312.
- [2] Derbinsky, N., Laird, J. E. 2009. Efficiently Implementing Episodic Memory. In Proc. of the 8th Int. Conf. on Case-Based Reasoning, Seattle, WA, USA, 403-417.
- [3] Deutsch, T., Gruber, A., Lang, R., Velik, R. Episodic Memory for Autonomous Agents. In Proc. of the 5th Int. Conf. on Human Sys. Interaction, Krakow, Poland, 621-626.
- [4] Gomes, P. F., Martinho, C., Paiva, A. 2011. I’ve Been Here Before! Location and Appraisal in Memory Retrieval. In Proc. of the 10th AAMAS, Taipei, Taiwan, 1039-1046.
- [5] Kuppaswamy, N. S., Cho, S., Kim, J. 2006. A Cognitive Control Architecture for an Artificial Creature using Episodic Memory. In Proc. of the 3rd SICE-ICASE Int. Joint Conf., Busan, Korea, 3104-3110.
- [6] Laird, J. E. 2012. *The Soar Cognitive Architecture*. MIT Press, Cambridge.
- [7] Laird, J. E., Derbinsky, N., Voigt, J. R. 2011. Performance Evaluation of Declarative Memory Systems in Soar. In Proc. of the 20th Behavior Representation in Modeling & Simulation Conf., Sundance, UT, USA, 33-40.
- [8] Lim, M. Y., Aylett, R., Vargas, P. A., Ho, W. C., Dias, J. 2011. Human-like Memory Retrieval Mechanisms for Social Companions. In Proc. of the 10th AAMAS, Taipei, Taiwan, 1117-1118.
- [9] Macedo, L., Cardoso, A. 2004. Exploration of Unknown Environments with Motivational Agents. In Proc. of the 3rd AAMAS, New York, NY, USA, 328-335.
- [10] Nuxoll, A. M., Laird, J. E. 2011. Enhancing Intelligent Agents with Episodic Memory. *Cognitive Systems Research* (In Press).
- [11] Tulving, E. 1983. *Elements of Episodic Memory*. Clarendon Press, Oxford.