Available online at www.sciencedirect.com

SciVerse ScienceDirect

Cognitive Systems Research 17–18 (2012) 34–48

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Enhancing intelligent agents with episodic memory

Action editor: Vasant Honavar

Andrew M. Nuxoll *, John E. Laird

University of Michigan, 2260 Hayward Street, Ann Arbor, MI 48109-2121, USA

Available online 31 October 2011

Abstract

For a human, episodic memory is a memory of past experiences that one gains over a lifetime. While episodic memory appears critical to human function, researchers have done little to explore the potential benefits for an artificially intelligent agent. In this research, we have added a task-independent, episodic memory to a cognitive architecture. To frame the research, we propose that episodic memory supports a set of cognitive capabilities that improve an agent's ability to sense its environment, reason, and learn. We demonstrate that episodic memory enables agents created with our architecture to employ these cognitive capabilities.

Keywords: Episodic memory; Cognitive architecture; Soar

1. Introduction

One advantage that humans have over current Artificial Intelligence (AI) systems is a personal history of specific events that they can draw upon to improve their decision making and learning. This episodic memory was first described in depth by Tulving (1983, 2002). Tulving's focus was phenomenological and in particular distinguished episodic memory from semantic memory. Episodic memory provides humans with the ability to remember where they have been, what they have sensed, and what actions they have taken in various situations. This knowledge of the past is invaluable for acting in the present. Episodic memory supports a wide range of cognitive capabilities from keeping track of the world outside immediate perception, to allowing retrospective learning on previously encountered situations. Certainly, there is evidence that human cognition is severely crippled by the loss of episodic memory and the difficulties that amnesiacs face have been documented (Tulving, 2002) and were dramatically portrayed in the movie Memento (Nolan, 2000).

As in any learning system embedded in a performance system, episodic memory involves: the capturing and encoding of experience in an internal format; storing that experience in a knowledge base for future use; retrieving knowledge in the future when given an appropriate cue. In addition to supporting these fundamental operations, there are additional functional requirements that Tulving (1983) identified that distinguish episodic memory from other memory and learning mechanisms:

- *Automatic*: The system creates new memories automatically without the agent deciding to do so. The underlying assumptions are that: (a) deliberately deciding which situations to remember can interfere with task-based reasoning and (b) it is unlikely that the agent can effectively determine which experiences will be relevant to future decisions.
- *Autonoetic*: A retrieved memory is distinguished from current sensing, so that an agent does not confuse a retrieved memory with the current situation.
- *Temporally indexed*: Because an episodic memory describes a particular, unique moment in time, some temporal information is a part of any episodic memory and

* Corresponding author. Present address: University of Portland, MSC 145, 5000 N Willamette Blvd., Portland, OR 97203, USA. Tel.: +1 503 943 7688; fax: +1 503 943 7316.

E-mail addresses: nuxoll@up.edu (A.M. Nuxoll), laird@umich.edu (J.E. Laird).

1389-0417/$ - see front matter
doi:10.1016/j.cogsys.2011.10.002

can also be part of an episodic memory cue. This need not be an exact time but it should convey a sense of the relative time of the episode with respect to other episodes.

This paper presents our progress toward creating a general purpose episodic memory within a cognitive architecture that supports the creation of general AI agents, that is agents that use large bodies of knowledge, continually learn from experiences in their environment, pursue multiple diverse tasks, and exist for extended periods of time. Although there has been sporadic research on episodic memory within AI in the past, there has not been research on task-independent episodic memories that support a wide variety of cognitive capabilities within a cognitive architecture.

Thus, our research involves determining the requirements for an episodic memory; designing, implementing, and integrating an episodic memory system within a cognitive architecture; and exploring the capabilities supported by such an integration. The emphasis of our research has been to create a computational system with the most important features of episodic memory so that we can develop and evaluate not just an episodic memory module, but an integration of that module within a cognitive architecture in which we can build agents. This paper describes our progress to date on this work. While we have made considerable progress, our episodic memory architecture is far from complete. For example, it does not include memory consolidation, forgetting, interference, priming, generalization across episodes or specific models of time. Our architecture does support effective encoding, storage and retrieval and we have used it to create agents for a variety of tasks. Our research suggests that episodic memory enhances the performance of AI agents and may be a "missing link" in current cognitive architectures, enabling a gamut of cognitive capabilities.

2. Cognitive capabilities

The focus of our research is to investigate whether episodic memory can support high-level cognitive capabilities beyond the simple recall of previous events. We define a cognitive capability as a beneficial ability or pattern of behavior that can be indirectly observed via the actions of intelligent agents. We do not claim that episodic memory is the sole means for achieving these capabilities but we claim that episodic memory can play an integral role is realizing these capabilities. Our research has focused on demonstrating multiple capabilities with a single architecture in a limited set of tasks set within two distinct environments. Our hypothesis is that episodic memory is integral to supporting the following cognitive capabilities, and in Section 7, we support this claim with example implementations.

- *Action modeling*: If an agent can recall the outcomes of its past actions, it can use that information to predict the outcomes of those same actions in the present. By recalling situations similar to the present from its past, an agent can use episodic memory to predict the immediate changes in its environment that will result from a given action.
- *Decision-making based on past experiences*: A history of prior successes and failures can be used during planning to guide decision making. Episodic memory can provide access to memories of similar situations that can help direct the agent to its goal and away from failures.
- *Retroactive learning*: Often, it is not possible to learn while an event is occurring because the agent lacks the specific information or resources that it needs to learn. For example, an agent in a real-time environment may not have time to apply an iterative learning algorithm while it is performing a task. Episodic memory allows previous experiences to be relived or rehearsed once the resources are available.
- "*Boost*" *other learning mechanisms*: An episodic memory store provides a wealth of data for training other learning mechanisms potentially allowing the agent to speed the rate at which it learns in new situations and even learn how to behave in situations it has never encountered before.
- *Virtual sensing*: In general, an agent's sensing is limited to the current situation: what is available directly from perception. Memory of recent events or situations can greatly expand an agent's ability to interact with the world outside. Episodic memory expands an agent's sensing by providing memories of areas beyond its immediate perceptions.

There are potentially many more cognitive capabilities that episodic memory enables, such as prospective memory (Kliegel, McDaniel, & Einstein, 2007) where an agent decides to perform future activity, such as stopping at the store on the way home from work, and then uses episodic memory to recall that plan at an appropriate time in the future. One of our goals for future research is to investigate more of the capabilities that may be possible through the addition of episodic memory to a cognitive architecture.

3. Related work

Within Artificial Intelligence, episodic memory research is closely related to case-based reasoning (CBR) (Kolodner, 1993; Schank, 1999). In a typical CBR system, each case describes a problem that the agent faced and a specific solution to that problem. When a new situation arises, the agent retrieves a stored case and adapts its solution to the new situation. Cases are usually described by a fixed number of task dependent fields, which are designed by a human, making them task dependent. Nonetheless, research in CBR highlights some important research that is equally relevant to episodic memory. In particular, Goodman (1993) describes using a CBR system's case library to predict future events. This prediction is analo-