# A Year of Episodic Memory

**John E. Laird and Nate Derbinsky**

University of Michigan
2260 Hayward Street
Ann Arbor, MI 48109-2121
laird@umich.edu, nlderbin@umich.edu

**Abstract**

Endowing an intelligent agent with an episodic memory provides knowledge that is invaluable for acting in the present, as well as supporting a wide range of cognitive capabilities. However, affording fast episode retrievals over long agent lifetimes presents significant theoretical and practical challenges. In this paper, we propose developing episodic memory systems that can efficiently store and retrieve experiences over the course of a year.

## Introduction

Episodic memory is a long-term, contextualized store of specific events (Tulving 1983). Episodic memory provides an agent with the ability to recall past sensing, situational interpretation, planning, and action results. Retrospective reasoning and learning over these experiences affords an agent a multitude of cognitive capabilities that may be crucial to its efficacy (Nuxoll & Laird 2007). Even in novel environments, the memory of recent events and situations can help in recognizing the return to already visited places, while the memory of older, but related events and situations can be useful in informing decision making and planning ("this reminds me of the last time I dove into a lake and almost broke my neck.").

Although similar mechanisms have been studied in case-based reasoning (Kolodner 1992), usually case-based reasoning systems are designed for specific tasks or classes of tasks, while episodic memory should be task-independent and thus available for any and all problems an agent may face. Furthermore, the growth in an episodic store will typically be much faster than a case base, as an episodic store accumulates snapshots of agent experience over millions and tens of millions of episodes.

In the following sections, we characterize the functional requirements of a task-independent episodic memory. Furthermore, we discuss the computational challenges of maintaining efficient episode storage and retrievals over a year's worth of memories.



## Characterizing Episodic Memory

Episodic memory is distinguished from other memory models by a set of functional requirements (Nuxoll 2007). First, episodic memory is an *architectural* ability that does not change from task to task. Second, episode storage is *automatic*, and thus does not require deliberate action by the agent. Finally, episodes reflect agent *experience*: there is no pre-specified set of features to be stored, nor are there pre-specified constraints on the features available for retrieval. For an agent with rich and varied experiences, an episodic memory must support whatever structures the agent perceives and reasons with.

## The Challenge

We set one year of continual use of episodic memory as the challenge. One year is long enough to establish long-term viability and usefulness, while anything longer could introduce difficulties in testing and evaluation. To fulfill these goals, episodic memory must be embedded within an agent that is engaged in and learning about multiple challenging tasks with novel components in a rich, dynamic environment. This requires the co-creation of technologies that can support the construction of such an agent, which appear to be on the horizon. There are many possible environments for such an agent, from being a personal assistant (such as CALO; Myers et al. 2007), an internet assistant, a household robot, or even a character situated in a massively multiple online world, such as OpenSim, Second Life, or even World of Warcraft.

This is a grand challenge in that it is beyond the state of the art and fulfilling it will provide important functionality for autonomous agents. The value of episodic memory comes from its integration with reasoning, planning, decision making, learning, and other cognitive capabilities. To provide a focused and well-defined challenge, we have focused on the specific functionality afforded by episodic memory, which we maintain can be pursued in its own right. Thus, our proposal attempts to lay out a challenge that has specific criteria for achievement and that focuses on specific functional capabilities related to storing, maintaining and retrieving experiences.

## Efficiency Issues

The challenge of creating an episodic memory for an agent that exists for a year arises from the real-time constraints on the agent's decision making in a complex, dynamic environment. For episodic memory to be useful, episodic storage and retrievals must not interfere with the agent's ability to respond to the dynamics of the environment – relevant information must be available quickly enough for it to be useful, even in a dynamic environment. Furthermore, real-world agents face bounded storage: while relatively cheap and plentiful, memory is not unlimited.

Below we decompose the computational challenges facing a task-independent episodic memory system into encoding, storage and retrieval. Before jumping into the analysis, we observe that in a year, a human is awake approximately 21 million seconds (assuming 16 hours of waking activity). Although this number seems small when compared to gigabytes of memory and gigahertz of processing cycles, for an ongoing agent with a non-trivial amount of distinctive data in its short-term memories, these numbers provide a significant challenge.

### Encoding

Encoding involves the storage of volatile data from an agent's short-term memory into the long-term, persistent episodic store. It involves determining when to store an episode (episode initiation), what aspects of the current situation should be stored, and how it is represented (episode determination).

**Episode Initiation.** Frequency of episode storage, with respect to agent lifetime, directly determines the number of episodes. Broadly, more episodes provides the agent with a finer-grain history of its experiences, but also implies more data to process in the store, and thus greater challenges to achieving system efficiency.

To fully reflect an agent's experiences, an episodic memory system must capture all structural and feature changes that have taken place since the last recorded episode. In the worst case, environmental changes are dramatic and an episode must reproduce *all* structure/features. Frequency of episode initiation, with respect to agent lifetime, is a linear multiplier of this cost. In situations of frequent storage, efficient systems must attempt to identify and exploit data regularities, so as to reduce redundancy, both in time and space, of stored episodic data. Here we briefly discuss *structural* and *temporal* regularities.

To facilitate reasoning and achieve productive behavior in data-rich environments, intelligent agents must exploit structural data regularities across time and tasks. Episodic memory systems that recognize and adapt to these regularities can expend less storage and time recording data that is common across many episodes.

In real-world environments, features tend to change slowly and locally. Episodic memory systems that process only feature *changes* can realize dramatic performance gains by focusing computation on important changes, as opposed to static features (Nuxoll & Laird 2007).

Although it is unlikely there is one universal frequency of episode storage that applies to all agents and all tasks, we can get some indication of upper and lower bounds from human cognitive processing. For humans, the basic "cycle time" of deliberate behavior is on the order of 50ms. That provides an upper bound on the amount of episodic memory that must be stored in a year of 20 * 21 million episodes = 420 million. It is difficult to determine a lower bound, at least for humans, as on internal reflection it seems that the resolution of our episodic memory is significantly coarser – except for intense experiences. Our own experience with mobile robots falls somewhere between 2 and 10 episodes/second, which translates to between 42 million and 210 million episodes for a year.

These numbers also provide a limit as to how much processing is required in storing an episode. If we assume the availability of a dedicated processor for episodic memory, then it is well within the realm of today's processors to process between 1 and 20 episodes/second (under the assumptions of episode content laid out in the next section).

**Episode Determination.** In addition to the number of stored episodes, the contents of each episode can pose significant performance challenges. Here we consider the implications of episode size, its representation, as well as memory element distinctiveness.

In a rich environment, an agent will need to maintain a potentially large representation of the current situation. Clearly the larger the representation, the larger the storage requirements and the higher the computational cost of determining the appropriate episode during retrieval. This is one area that is wide open in terms of estimating upper and lower bounds. Attempting to base a lower-bound on the famous $7^{+}/-2$ size of human short-term memory (Miller 1956) is fraught with error as those results were based on memory tests of lists of items of similar "types," which did not take into account the human ability to remember rich representations when items have rich semantic content and distinctiveness. Episodes appear to capture rich representations of situations that require 10's and 100's or even 1,000's of elements in current computer models. Our own experience with episodic memory for non-toy tasks is that somewhere between 100 and 1,000 elements are required per episode.

Real-world environments contain rich relational feature descriptions and thus agent episodes must be sufficiently expressive to represent arbitrarily complex structures. Thus, we impose the requirement that the underlying representation cannot be merely propositional, but must support relational representations. This level of expressiveness, however, can introduce significant performance challenges in the context of efficient storage and retrieval, such as tracking feature changes, indexing episodes for efficient retrievals, and reconstruction of episodes for agent use. An open question is whether non-symbolic representations are also required, such as images or other perception-based representations.

In any representation, an episodic memory system must contend with the spectrum of feature distinctiveness. At one extreme is the problem of efficiently storing and searching episodes containing many qualitatively different features. For instance, consider searching over continuous sensor features in a robotics domain. Episodic memory systems in these domains must evaluate policies amongst efficient, but ineffective, qualitative matches and informative, but potentially costly, quantitative search (Stottler et al. 1989).

Real-world environments also challenge episodic memory systems with sets of qualitatively identical objects. Object identification can vastly reduce storage requirements by recognizing and compressing redundant object descriptors, but at the cost of computationally expensive comparison algorithms. A common approach during retrieval is a two-stage matching algorithm (Gentner & Forbus 1991; Tecuci & Porter 2007), bounding expensive search operations over a small/constant number of candidate episodes.

## Storage

**Episode Structure.** The simplest and most expensive model of episodic memory is to store a distinct snapshot in each cognitive cycle. Based on our earlier worst-case analysis, this would translate to one episode (1,000 elements) every 50ms (20 times/second) for 1 year (21 million seconds), which equals 420 billion items. Our best-case scenario would be one episode (100 elements) every 500ms (2 times/second) for 1 year (21 million seconds), which equals 4.2 billion items. There is an additional multiplier for representing an item, which we estimate to be between 10 and 100 bytes, giving a range of memory requirements of between 42 GB and 42 TB.

More sophisticated approaches to encoding episodes, such as taking advantage of temporal continuity and redundant structures, can decrease these demands, at the cost of additional indexing. In our work, we have seen a minimum of a factor of 2 decrease from compression. Thus, a reasonable estimate is between 21 GB and 21 TB using current approaches. Both of these are within range of today's commodity secondary storage systems. The lower end of the range is possible with primary memory for servers, which appear to have a current practical limit of 256 GB. This analysis suggests that memory alone is probably not a limiting factor for achieving year-long episodic memory.

**Episode Dynamics.** The continued growth of episodes introduces computational challenges both in the amount of memory for storage and the computational resources required for retrieval. To decrease the growth or even bound the size of the episodic store, a forgetting or consolidation mechanism could be introduced. Such a mechanism could maintain statistics about the episodic store (retrieval frequency, age, usefulness, redundancy, etc) and, upon reaching maximum capacity, make a decision about which episode(s)/episode feature(s) to remove or consolidate. An important issue is whether the gain in computational performance through forgetting outweighs the potential negative impact on reasoning that arises from the loss of episodes that could be useful in future reasoning.

## Retrieval

Retrieving a memory involves matching a cue against the stored episodes, determining a best match, and reconstructing the episode so it can be used in reasoning. Cue matching over a growing episodic store has the potential to be the most costly episodic operation as a complete search of the store is at least linear in the number of episodes, and because the episodes have complex relational structures.

Moreover, the frequency of retrievals and source of retrieval cues can significantly affect the computational cost of the implementation. For instance, supporting spontaneous cues (i.e. episodic retrieval over the entirety of the present state) demands an implementation that is continually searching episodic memory using cues with relatively large numbers of features. Although spontaneous retrieval is potentially beneficial, we take as a starting point for this challenge that retrieval is based on deliberately constructed cues that are significantly smaller than individual episodes.

Within the context of a broader agent, a retrieval does not necessarily have to be tied directly to the time scale of primitive deliberation. In humans, an episodic retrieval can be on a time scale of many primitive cycles (50 ms). However, a memory will lose its utility if it has not been retrieved within some limited amount of time after cue initiation. Clearly this depends on the dynamics of the task. To give a ballpark figure, we will assume 1 second, which corresponds to 20 primitive cycles. We also assume that a retrieval can be executed asynchronously using its own processing resources, making 1 second a fixed bound.

Consider a linear search of the best case for a year's worth of memories (21 GB) stored sequentially in [an extended] primary memory given modern hardware (2GHz CPU). Ideal conditions would require at least 10 seconds for the processor to read every memory element. This estimate will increase by multiple orders of magnitude when search comparisons are taken into account. And this is the best case with the worst case requiring 1,000 times as much data to be processed.

This defines the real challenge: how to organize episodic data, as it is learned over a year, so that it can be searched in bounded time. There are two basic approaches. The first is to develop algorithms and data structures that speed the matching processes but still guarantee a best match. Although these could potentially achieve high performance in the average case, in the worst case they must examine every episode; however, with massive parallel processing, it might be possible to still achieve the requisite reactivity.

An alternative is to develop heuristic approaches that no longer guarantee a best match, but can provide guaranteed performance (possibly through an anytime algorithm). Some possible approaches include history compression (Schmidhuber 1992) and query caching/optimization (Chaudhuri 1998; Gupta et al. 1997). If this challenge is accepted, it will be necessary not only to define computational performance requirements, but also metrics for assessing the quality of the retrieved episodes.

## Conclusions

Prior work has shown that endowing an agent with an episodic memory provides knowledge to support a vast array of cognitive abilities crucial for intelligent behavior. We have proposed the challenge of developing episodic memory systems that can capture the experience of an agent that exists for a year and runs in real time. We have attempted to provide bounds on the computational resources required for such a system. These are very broad, spanning multiple orders of magnitude and providing answers that are more specific, unfortunately, awaits the development of agents that actually exist for a year.

It appears that the computational resources required to encode and store the resulting episodes are within reach. Developing algorithms that can support sufficiently fast bounded retrievals is the real challenge. Meeting that challenge will involve drawing on recent work in large scale relational databases, but given the open-endedness of the structures that must be stored, the challenge will require significant extensions to existing work. In addition to research in data structures and algorithms, we foresee a considerable need for empirical evaluation across a range of tasks and environments. We surmise that these performance benchmarks will draw on and contribute to a variety of experienced-based reasoning research.

## References

Chaudhuri, S. 1998. *An Overview of Query Optimization in Relational Systems*. In Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 34-43.

Gentner, D., Forbus, K.D. 1991. *MAC/FAC: A Model of Similarity-based Retrieval*. In Proceedings of the 13th Cognitive Science Conference, 504-509.

Gupta, H., Harinarayan, V., Rajaraman, A., Ullman, J.D. 1997. *Index Selection for OLAP*. Conference on Data Engineering, 208-219.

Kolodner, J. 1992. *An Introduction to Case-Based Reasoning*. Artificial Intelligence Review, 6(1): 3-34.

Miller, G.A. 1956. *The Magical Number Seven, Plus or Minus Two*. The Psychological Review 63(2): 81-97.

Myers, K., Berry, P., Blythe, J., Conley, K., Gervasio, M., McGuinness, D., Morley, D., Pfeffer, A., Pollack, M., and Tambe, M. 2007, *An Intelligent Personal Assistant for Task and Time Management*, AI Magazine, 28(2):47-61.

Nuxoll, A. 2007. Enhancing Intelligent Agents with Episodic Memory. Ph.D. diss., Univ. of Michigan, Ann Arbor.

Nuxoll, A., Laird, J.E. 2007. *Extending Cognitive Architecture with Episodic Memory*. In Proceedings of the 21st National Conference on Artificial Intelligence, 1560-1564.

OpenSim, http://opensimulator.org

Schmidhuber, J. 1992. *Learning Complex, Extended Sequences Using the Principle of History Compression*. Neural Computation, 4(2), 234-242, MIT Press.

Second Life, http://secondlife.com

Stottler, R., Henke, A., King, J. *1989*. *Rapid Retrieval Algorithms for Case-Based Reasoning*. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 233-237.

Tecuci, D., Porter, B. 2007. *A Generic Memory Module for Events*. In Proceedings of the Twentieth International FLAIRS Conference, 152-157.

Tulving, E. 1983. *Elements of Episodic Memory*. Oxford: Clarenden Press.

World of Warcraft, http://www.worldofwarcraft.com