Elements of episodic memory: insights from artificial agents

Article:

Boyle, Alexandria ORCID: 0000-0001-8827-5479 and Blomkvist, Andrea (2024) Elements of episodic memory: insights from artificial agents. Philosophical Transactions of the Royal Society B: Biological Sciences. ISSN 0962-8436 (In Press)

Elements of Episodic Memory: Insights from Artificial Agents

Alexandria Boyle [1, 3]
ORCID iD: 0000-0001-8827-5479

Andrea Blomkvist [1, 2]
ORCID iD: 0000-0002-1997-592X

[1] London School of Economics and Political Science
[2] University of Glasgow
[3] CIFAR Azrieli Global Scholars Programme

Corresponding author: Alexandria Boyle (a.boyle2@lse.ac.uk)

Abstract: Many recent AI systems take inspiration from biological episodic memory. Here, we ask how these 'episodic-inspired' AI systems might inform our understanding of biological episodic memory. We discuss work showing that these systems implement some key features of episodic memory whilst differing in important respects, and appear to enjoy behavioural advantages in the domains of strategic decision-making, fast learning, navigation, exploration and acting over temporal distance. We propose that these systems could be used to evaluate competing theories of episodic memory's operations and function. However, further work is needed to validate them as models of episodic memory and isolate the contributions of their memory systems to their behaviour. More immediately, we propose that these systems have a role to play in directing episodic memory research by highlighting novel or neglected hypotheses as pursuit-worthy. In this vein, we propose that the evidence reviewed here highlights two pursuit-worthy hypotheses about episodic memory's function: that it plays a role in planning that is independent of future-oriented simulation, and that it is adaptive in virtue of its contributions to fast learning in novel, sparse-reward environments.

Key words: episodic memory; artificial intelligence; causal role function; mental time travel

1. Introduction

Forty years on from the publication of Elements of Episodic Memory (Tulving, 1983), researchers are still working to understand the nature of episodic memory, its role in cognition, its evolutionary origins, and its distribution in the animal kingdom. Recently, episodic memory has also begun to play a role in the development of AI systems: many new AI systems include processes inspired by episodic memory. It has been proposed that such AI systems might shed light on the nature of episodic memory itself (Hassabis et al., 2017). Here, we review 'episodic-inspired' work in AI with a view to evaluating whether, and in what ways, it could inform our understanding of episodic memory in biological systems. In focussing on episodic-inspired AI, our aim is to highlight how this emerging body of work in AI might add to the existing toolkit in episodic memory research, rather than to suggest that AI systems are in a privileged position relative to other computational models.[1]

In Sect 2., we explore similarities between biological and artificial memory systems, focussing on the presence of constructive memory processes and the use of hippocampus-inspired replay buffers in AI systems. We show that the term 'episodic memory' is used more liberally in AI than in other cognitive scientific domains, to characterise systems implementing some features of episodic memory whilst differing in significant respects. We propose that 'event memory' might be a more perspicuous label for the memory implemented in these systems. In Sect. 3, we consider the effects of event memory on behaviour, showing that event memory appears to contribute to strategic decision-making, fast learning, navigation, exploration and acting over temporal distance in AI systems. However, we suggest that further research is needed to isolate the precise contribution of event memory to these achievements. In Sect. 4, we consider how the implementation of event memory in AI might contribute to episodic memory research. We suggest that, given more rigorous testing to establish their validity as models, AI systems with event memory could be used to test hypotheses about the operations and functions of episodic memory. More immediately, studying event memory in AI may suggest new directions for episodic memory research by highlighting novel or neglected hypotheses as pursuit-worthy. In this vein, we suggest that the evidence reviewed here points toward at least two pursuit-worthy hypotheses about episodic memory's function: that episodic memory plays a role in planning that is unmediated by future-oriented simulation, and that episodic memory is adaptive in virtue of its contributions to fast learning in novel, sparse-reward environments.

2. Biological and artificial memory

In this section, we explore similarities and differences between biological episodic memory and episodic-inspired architectures in AI, considered at the algorithmic level (Marr, 1982). This level, which concerns the causal contributions of component processes to a system's performance on a task, is an appropriate level of analysis for two reasons. First, much existing research in cognitive science has focussed on characterising memory processes at the algorithmic level, leaving us well placed to evaluate similarities here. Second, similarities and differences at the algorithmic level are more informative in this

---

[1] See Norman et al. (2012) for a review of other computational approaches.

context than those found at Marr's other levels of analysis. Similarities at the computational level are somewhat trivial, as we often intentionally design algorithms to solve the same computational problems as humans. Conversely, we are unlikely to find meaningful similarities at the implementational level, as little emphasis has been put on mimicking biological systems at the level of individual neurons (but see (Camp et al., 2020)). Two themes will structure our discussion in this section: the presence of constructive memory processes in AI systems, and the use of hippocampus-inspired replay buffers to consolidate information.

2.1 Constructive episodic memory

The idea that human episodic memory is constructive originates in the 1930s (Bartlett, 1932), and has become increasingly prominent in both cognitive science (Schacter & Addis, 2007, 2020) and philosophy (Aronowitz, 2019; De Brigard, 2014; Michaelian, 2016). On this view, the memory system encodes information selectively, information can be altered as it is 'stored', it can be flexibly recombined with other information and altered upon recall, and new information may be generated. This contrasts with a classical 'preservative' view of memory in which a memory system is optimised to preserve information and memory is encoded and preserved as faithfully as possible, with the goal to minimise corruption.

Early architectures took a more preservative approach to memory. For example, the SOAR architecture (Nuxoll & Laird, 2012) comprises multiple memory systems, including modules for procedural memory, working memory and an 'episodic memory'. The agent can take virtual actions in the world specified as if-then rules in its procedural memory. Its working memory comprises its current state – including what it is currently 'sensing' (e.g., nearby objects) and potential actions (e.g., go north). Its 'episodic memory' stores past states copied from working memory, retrievable by cueing the memory store with a partial content. Here, there is no selective process operating on content to prioritise the encoding of certain content, there is no process to recombine content from different episodes, and there is no consolidation process which could potentially alter content. In short, it shares few constructive algorithmic features with human episodic memory.

But perhaps surprisingly, other systems contain processes resembling constructive memory (Bhat et al., 2014; Blundell et al., 2016; Hung et al., 2019). Designing a system in this way reduces memory storage, making it more economical. This method was exploited when designing the humanoid robotic system iCub (Mohan et al., 2014; Vernon et al., 2007), which can flexibly recombine information and produce novel solutions to problems, similarly to humans. In one experiment (Mohan et al., 2014), the robot was tasked with stacking the tallest tower using building blocks of various shapes, each of which had been encountered with only one other shape before. To solve the task optimally, the robot had to combine information from different past episodes into a novel representation of the solution in its 'episodic construction system', and the solution was forwarded to the action system. Interestingly, this resembles experimental paradigms elsewhere in cognitive science, in which human participants must imagine a novel scenario using content from disparate past remembered events, for example by taking the place in one episode and combining it with a person from a different episode (Addis et al., 2009). Several researchers propose that this ability plays an important role in our capacity for future-planning via the simulation of imagined future scenarios (Buckner & Carroll, 2007; Schacter et al., 2012; Schacter & Addis, 2007; Suddendorf & Corballis, 2007; Tulving, 2005).

The use of convolutional neural networks to fill in missing information in memory may also be considered a constructive retrieval process (Fayyaz et al., 2022; Rothfuss et al., 2018). For example, the VQ-VEA network uses a Pixel Convolutional Neural Network (PixelCNN) which fills in missing pixels in photographs (Fayyaz et al., 2022). When a picture is fed to the network, the encoder compresses the image into an array of high-dimensional vectors, and an index of these vectors is created. Upon recall, the missing information in the compressed representation is filled in by PixelCNN. This not only trades on reconstructive processes in recall, but also on Hippocampal Indexing Theory (Teyler & DiScenna, 1986; Teyler & Rudy, 2007). This algorithm has been tested on an episodic memory paradigm and compared to results from human participants. In the human case, participants virtually experienced an apartment with congruent or incongruent objects in rooms – for instance, a toaster in a kitchen (congruent) or in a bathroom (incongruent). Correctly recalling an incongruent detail indicates that episodic memory is used. A parallel experiment was used to evaluate the algorithm, in which it was presented with digits on congruent or incongruent backgrounds, having been trained only on congruent cases. Results showed that when the attention level of the algorithm was high, it correctly recalled incongruent details, which could indicate the presence of something like an episodic memory capacity.

However, there are also differences between biological episodic memory and PixelCNN's memory architecture. PixelCNN is a probabilistic autoregressive generative model, meaning that the way it fills in missing information is by continuing a sequence of numbers, filling these in from left to right. Humans, by contrast, are likely to use a Bayesian process in reconstructive recall (Blomkvist, 2022; De Brigard, 2014; Hemmer & Steyvers, 2009). It is possible that these differences in algorithm could have functional consequences for how the system operates as a whole and what behaviour is ultimately produced. This is an empirical question which remains to be investigated. But to foreshadow issues regarding modelling discussed in sect. 4, it suffices to say here that a model need not replicate all aspects of the modelled system, such as the exact mathematical algorithm executed, in order for it to be useful in generating hypotheses about that system. For now, we conclude that, although there is a high-level similarity between humans and PixelCNN in that both employ reconstructive processes, the exact ways in which reconstruction is achieved are different.

2.2 Hippocampus-inspired replay buffers for consolidation

AI researchers have also taken inspiration from the role of the hippocampus in consolidation. Consolidation is thought to be underpinned by hippocampal replay, the reactivation of neural patterns during which spiking activity induces long-term changes in synapses during sleep or rest (Hayes et al., 2021). This process may determine which content is transferred into long-term memory, subsequently becoming dependent on the medial prefrontal cortex rather than the hippocampus for retrieval (Kitamura et al., 2017). Inspired by hippocampal replay, experience replay algorithms are trained on subsets of replayed ‘experienced episodes’ to encourage consolidation and subsequent learning.

Experience replay was prominently used in the DQN algorithm (Mnih et al., 2015), a reinforcement learning network exploiting uniform experience replay, which samples past episodes for replay uniformly at random from a memory buffer. Employing this algorithm when updating a network resulted in a new state of the art result on Atari 2600 games, commonly used as a benchmark for evaluating AI systems. Multiple algorithms have now been developed using alternative sampling methods to achieve superior performance (see Table 1 for an overview). For example, hindsight experience replay replays episodes with goals other than the one originally pursued by the agent (Andrychowicz et al., 2017), imaginary hindsight experience replay generates new data which is incorporated into replay (Mccarthy et al., 2023), and prioritized experience replay replays more informative transitions more frequently (Schaul et al., 2015).

But the degree to which these algorithms resemble hippocampal replay is questionable. There is biological precedent for the replay mechanism in general, lending biological plausibility to these algorithms (Hayes et al., 2021). Some of these algorithms also take note of the suggestion that the episodes which get replayed and consolidated in human memory are ones with high affective valence (Baran et al., 2012; Bowen et al., 2017; Hayes et al., 2021; Payne et al., 2008; Sakaki et al., 2013). For example, prioritized replay preferentially sample episodes that are more ‘surprising’, yielding much higher or lower than expected reward (Schaul et al., 2015). But human affect is much richer than pleasant and unpleasant surprise, incorporating states such as joy, sadness, and anger. So, whilst replay mechanisms in AI are influenced by both hippocampal replay and affect, the factors influencing replay in humans are likely to be more complicated. This is unsurprising, not only because of the many differences between AI systems and humans, but also because most AI researchers are not primarily attempting to emulate human cognition, but to solve specific computational problems, such as catastrophic forgetting or learning from sparse rewards. Again though, as discussed further in Sect. 4, some differences between models and modelled systems is tolerable. That is, models need not replicate the exact structure of modelled systems to be useful for hypothesis generation.

Type of experience replay algorithm	Replay algorithm	Description
Prioritized experience replay (PER)	Uniform experience replay (Mnih et al., 2015)	Samples past episodes for replay uniformly at random from a memory buffer.
	Prioritized experience replay (Schaul et al., 2015)	Replays episodes with higher expected learning progress.
	Episodic backward update (Lee et al., 2018)	Samples transitions episode by episode, and updates values recursively in a backward manner.
	Prioritized experience replay with successor representations (Y. Yuan & Mattar, 2021)	Augments PER with a "need term" inspired by replay patterns in biological organisms.
	Double prioritized state recycled experience replay (Bu & Chang, 2020)	Prioritizes certain experiences in training and storing, uses state recycling to reuse and update experiences based on old ones.
	Hindsight experience replay (Andrychowicz et al., 2017)	Replays episodes with goals other than the one originally pursued by the agent.
	Prioritized oversampled experience replay (Sovrano, 2019)	Replays the most important and useful experiences for an agent.
	Spatial structure and Frequency-weighted Memory Access (Zeng et al., 2023)	Sequential replay algorithm which reproduces various replay statistics observed in the rodent hippocampus.
Hindsight experience replay (HER)	Augmented curiosity-driven experience replay (Li et al., 2020)	Augments HER by adding a curiosity reward, and a dynamic initial state selection mechanism.
	Dynamic hindsight experience replay (Fang et al., 2019)	Assembles successful experiences from two relevant failures.
	Imaginary hindsight experience replay (Mccarthy et al., 2023)	Generates new data which is incorporated into replay.
	Cluster-based sampling in HER (Kim & Har, 2022)	Groups episodes with different achieved goals by using a cluster model and samples experiences to create the training batch.
	Hindsight relabelling experience replay (Packer et al., 2021)	Relabels experiences during meta-training to enable learning from sparse rewards.
Reverse experience replay (RER)	Reverse experience replay (Rotinov, 2019)	Replays data in the buffer in reverse temporal order.
	Introspective experience replay (Kumar & Nagaraj, 2022)	Selectively samples batches of data prior to surprising events.

Table 1. Experience replay algorithms broadly building on biological principles of episodic memory.

2.3 Episodic memory in comparative cognitive science

There are both similarities and differences at the algorithmic level between human episodic memory and episodic memory inspired algorithms. What we find depends on how closely we look. For example, some artificial agents now have constructive algorithms, but their constructive processes are likely to differ from those used by humans to accomplish similar tasks. Similarly, biologically-inspired replay algorithms are widespread in reinforcement learning agents, but the way in which episodes are sampled for replay likely differ from the sampling methods employed in the brain.

Despite these differences, the term ‘episodic memory’ is used permissively in AI research to describe systems which take inspiration from human episodic memory. This terminological practice diverges from those adopted elsewhere in comparative cognitive science. In comparative cognition, for instance, many researchers prefer the label ‘episodic-like memory’ for the memory capacities of nonhuman animals (Clayton & Dickinson, 1998; Davies et al., 2022), reflecting uncertainty about whether animals’ experience the characteristic feeling of ‘reliving’ past events (see Boyle, 2020b for discussion). Other researchers adopt the label ‘event memory’ (Boyle, 2020a; Keven, 2016; Mahr & Csibra, 2018; Rubin & Umanath, 2015). This more inclusive term refers to explicit memories which integrate multiple spatiotemporal or perceptual details about past events, whilst perhaps lacking other features often associated with episodic memory, such as a narrative component, an understanding of the ‘pastness’ of the event, or an awareness of oneself as the author of the memory (sometimes referred to as ‘autonoetic awareness’ (Tulving, 1983)).

These divergent terminological practices risk giving the misleading impression that we can more confidently ascribe an episodic memory capacity to artificial agents than to nonhuman animals. This not only presents obstacles to interdisciplinary communication and research, but also signals that the memory capacities of artificial systems are more ‘human-like’ than we have reason to believe – a troubling result, given the role episodic memory plays in ethical frameworks and practices. In reality, whilst we can be confident that some episodic memory-inspired algorithms form explicit, integrated memories of past events, we can also be confident that they differ from biological episodic memory in important respects. The term ‘event memory’ would likely be used to characterise memories in nonhuman animals fitting this description, so – in line with recent recommendations to exercise caution when applying ‘rich psychological terms’ to AI systems (Shevlin & Halina, 2019) – we propose that it would be appropriate to apply the same term to such AI systems. Accordingly, in what follows we use the label ‘event memory’ to refer to memory systems in artificial agents which implement some aspects of episodic memory.

By proposing a new term here, we do not mean to imply that there are no similarities between event memory and human episodic memory. Just as researchers postulate similarities between episodic-like memory in non-human animals and human episodic memory (for instance, there is broad agreement that these record details pertaining to what-where-when), we hold that there are significant broad similarities between event memory and human episodic memory. Nevertheless, a shift in terminology is appropriate to parallel the more restricted use of ‘episodic memory’ elsewhere in cognitive science.

3. Event memory and behaviour in AI systems

How does event memory expand the behavioural repertoire of AI systems? One commonplace answer is that event memory can be used to overcome ‘catastrophic forgetting’ – a phenomenon wherein learning a new task leads the agent to forget a previously learned task (Hassabis et al., 2017). The use of event memory for this purpose is inspired by the Complementary Learning Systems theory (McClelland et al., 1995), which proposes that hippocampal replay is responsible for the gradual consolidation of rapidly learned information into neocortical long-term memory. Whilst replay methods are emerging as a promising solution to catastrophic forgetting, the ‘event’ element of event memory does not seem critical here: replay can be used in this way when the items being replayed are isolated datapoints (such as images in an image classification algorithm), rather than integrated, multimodal representations of past events. More complex event memories seem most likely to shed light on the roles episodic memory might play in biological systems (see Sect. 4).

Accordingly, we focus here on agents interacting with two- and three-dimensional real or simulated environments, learning directly from high-dimensional perceptual input and storing more complex event memories. Within this broad cluster, event memory has been implemented in a variety of ways, embedded in a range of architectures, and tailored to different types of tasks. This makes it difficult to offer a general answer to how event memory expands these agents’ capabilities. However, several themes emerge from the literature, many of which mirror functions that have been proposed for episodic memory in biological systems.

3.1 Discovering optimal strategies

One function that has been proposed for episodic memory in biological systems is that it enables agents to extract information from past experiences repeatedly and retrospectively, producing insights which might otherwise be elusive (Boyle, 2019; Brown, 2023). AI systems with event memory are similarly able to learn from events repeatedly and retrospectively. In some cases, AI systems using event memory in this way have been able to discover novel strategies for solving problems. For example, when augmented with event memory, SOAR (see Sect. 2.1) learned novel strategies in TankSoar– a complex two-dimensional environment in which the agent controls a tank with multiple sensors and must navigate around obstacles, collect objects and destroy enemies. The agent exhibited novel behaviours, including dodging missile attacks and moving out of sight of enemies, which were not robustly displayed by a comparison agent lacking event memory. Similarly, DQN learned a novel strategy in Atari Breakout, a game in which players eliminate rows of bricks using a ball manipulated with a movable paddle at the bottom of the screen. DQN spontaneously adopted an optimal ‘tunnelling’ strategy - knocking out a tunnel of bricks, enabling the ball to rapidly break several bricks by bouncing from the top of the screen without returning to the player (Kumaran & Hassabis, 2015).

3.2 Learning rates

Episodic memory is encoded rapidly in biological systems, allowing for one-shot learning. Similarly, event memory algorithms have displayed impressive learning speeds in complex environments. In particular, ‘episodic control’ algorithms (Lengyel & Dayan, 2008) use past experiences to guide online decision making, rather than simply replaying experiences for training purposes. In these algorithms, an event memory store encodes information about past actions. Agents query the store and emulate the most rewarding action previously taken in states similar to the one currently faced. In principle, such agents can exhibit ‘one-shot learning’ (Fei-Fei et al., 2006), since they need only take a rewarding action once in a given state to repeat that action when encountering a similar state again. However, exploiting early successes rather than exploring alternative actions may cause agents to get stuck in local optima, meaning that slower-learning algorithms may ultimately discover more effective policies.

For example, when tested on fifty-seven Atari games, the Neural Episodic Control algorithm (NEC) learned quickly, initially outperforming all comparison algorithms. It was eventually overtaken by prioritised experience replay (Pritzel et al., 2017). A similar pattern was observed when the Model-Free Episodic Control algorithm was compared with alternatives on a sample of five Atari games (Blundell et al., 2016). MFEC also learned quickly on a more complex task requiring it to locate a reward in the three-dimensional Labyrinth environment. Because rewards are very sparse in this environment, no comparison algorithm learned a policy with a positive expected reward on this task. By contrast, MFEC was able to learn from the very few instances in which a reward was obtained, a behaviour its developers describe as ‘akin to one-shot learning’ (Blundell et al., 2016, p. 7). These results have been taken to support the view that ‘episodic control could be used in the brain, especially in the early stages of learning in a new environment’, operating as one of several systems which the brain switches between according to context (Blundell et al., 2016, p. 8).

3.3 Exploration

Whilst exploiting event memory can lead agents to become stuck in local optima, event memory has also been combined with other techniques to improve exploratory behaviour in artificial agents. In one case, event memory was augmented with ‘subjective timescale models’ (STM), a constructive memory mechanism in which the rate of memory accumulation is adaptively adjusted as a function of prediction error (Zakharov et al., 2020). This model skips frames when recording memory sequences, writing experiences to memory only when a threshold of prediction error is exceeded. The result is that salient or surprising events are ‘squeezed together’ in memory, and less interesting periods of time are ‘compressed’. This agent exhibits more pronounced exploratory behaviour than comparable agents without STM, because its temporally distorted memories cause it to predict that it will encounter rewards more often, incentivising exploration. The agent outperforms agents with more ‘accurate’ memories in sparsely rewarded three-dimensional foraging tasks, indicating a possible advantage to this type of memory error – in keeping with discussions of the adaptive benefits of constructive mechanisms in biological systems (Schacter et al., 2011). Other models combining experience replay with ‘intrinsic rewards’, a method which incentivises exploration by rewarding experiences with high prediction error, have also produced improved exploratory behaviour and greater sample efficiency in robotics tasks (Mccarthy et al., 2023) and Atari (Sovrano, 2019).

3.4 Navigation

As noted above (Sect 2.2), event memory often takes inspiration from theories of hippocampal function. The hippocampus is thought to be critical to spatial memory and navigation; indeed, some propose this to be the primary function of the hippocampus (O'Keefe & Nadel, 1978; Schiller et al., 2015). With that in mind, it is perhaps no surprise that event memory algorithms have performed well on navigation tasks.

For example, SOAR (see Sect. 2.1) can use event memory to construct cognitive maps, enabling it to represent and plan routes through parts of the environment it is not currently sensing (Nuxoll & Laird, 2012). Similarly, both Hierarchical Chunk Attention Memory (HCAM) (Lampinen et al., 2021) and the Episodic Planning Network (EPN) (Ritter et al., 2021) perform significantly better than strong baselines in the One-Shot Street Learn task. In this task, agents navigate through a three-dimensional Google StreetView environment to reach as many goals as possible in a novel neighbourhood in a short space of time. Strong performance requires agents to learn the environment quickly at an early stage in order to plan the most efficient route. On a similar task in which agents must navigate to a cued goal from random starting points as many times as possible in a short space of time, a 'learning to reinforcement learn' agent (L2RL) (based on Wang et al., 2016) augmented with event memory significantly outperformed the maximum reward achievable without event memory, navigating near optimally following the first exposure to a goal location (Ritter et al., 2018). Event memory has also been used to facilitate robot navigation in real world environments, enabling robots to construct cognitive maps of novel environments used to plan future trajectories (Chin et al., 2019; Liu et al., 2017; J. Yuan et al., 2022).

3.5 Acting over temporal distance

Some systems with event memory can solve problems involving 'long-term temporal credit assignment' – effectively, those requiring action over extended periods of time. This includes scenarios in which optimal behaviour depends on earlier actions or observations, as well as complex hierarchical tasks in which attaining a goal requires performing a sequence of unrewarded or minimally rewarded actions.

For example, the Temporal Value Transport algorithm (TVT) (Hung et al., 2019) encodes compressed memories of events, retrieves them to guide action selection, and retrospectively re-values earlier actions based on later return. In one task requiring long-term temporal credit assignment, TVT is required to fetch a key, complete a distractor task involving the collection of small rewards, use the key to open a door to see a coloured square, complete a second distractor task, and then press a button matching the colour behind the door. TVT reliably solves this and similar tasks, outperforming a baseline agent with a similar memory architecture which lacked the retrospective revaluation mechanism. Agents with event memory have also successfully completed complex hierarchical tasks. For example, Dreamer v3, a model-based reinforcement learning agent with an event memory buffer, is the first agent to mine a diamond in Minecraft without human data or curricula (Hafner et al., 2023). This is a significant milestone because mining a diamond requires completing a sequence of twelve sparsely rewarded interim steps. Similarly, NEC (see Sect 3.2) learns to forego small rewards in favour of multi-step processes leading to greater reward in Atari.

Interestingly, these temporally extended and hierarchical tasks resemble the 'spoon test' proposed by Tulving (2005) to investigate episodic memory and mental time travel in animals and children. In a spoon test, a subject must anticipate their need for a tool in a location to be visited in future, pick up the appropriate tool when the opportunity arises, and take the tool to the relevant location. This is a temporally extended hierarchical task in which success requires an unrewarded early action, and has been used to investigate the episodic memory capacities of several great ape species (see Scarf et al., 2014 for review) – reflecting the idea that a function of episodic memory in biological agents is to support this kind of temporally extended action.

3.6 Event memory's contributions to behaviour

AI systems with event memory have exhibited strong performance in a range of areas, including discovering optimal strategies, learning quickly, exploring, navigating, and acting at a temporal distance (see Table 2 for an overview). These are all areas in which episodic memory and its component processes are thought to play a role in biological systems.[2]

Whilst it is tempting to attribute these achievements to the use of event memory, caution is required in interpreting these results. It is often difficult to isolate the contribution of event memory to an agents' behaviour because algorithms are frequently evaluated in terms of a comparison with other prominent algorithms on a certain test or benchmark. Such comparisons are difficult to interpret because the algorithms being compared may differ in ways other than the presence or absence of event memory. Clearer evidence about the influence of event memory on an agent's behaviour can be obtained through ablation experiments, in which the intact algorithm's performance is compared with the performance of the same algorithm lacking the event memory component. Whilst ablation methods are not uncommon (e.g. Ritter et al., 2021; Zakharov et al., 2020), they are used less frequently than comparisons with other algorithms. More widespread use of ablation methods targeting event memory would facilitate interdisciplinary insights.

Even in those cases where event memory is demonstrably critical to an agent's strong performance on a task, it may not be the only critical factor. First, the way that event memory is embedded into the overall architecture is important: different algorithms exploit event memory in different ways and have different capabilities as a result. For instance, 'episodic control' architectures, which consult buffers of stored past events to

---

[2] The absence of future-oriented simulation from this list may appear striking in this context, given the role episodic memory is thought to play in this process in humans (see Sect. 2.1). No conclusions should be drawn from this omission, however, as it merely reflects a lack of evidence about the relationship between event memory and future-oriented simulation in AI systems. This would be a fruitful avenue for future research.

inform online decision making, exhibit much faster learning than architectures like DQN, in which memory buffers are used only to train policy networks. Second, event memory may facilitate some capabilities only in the presence of other factors. For example, a key reason NEC can discover policies in Atari which elude DQN is the way the algorithms represent rewards: DQN clips reward values to between 1 and -1, whereas NEC tracks the true reward. Similarly, whilst event memory makes a key contribution to the success of EPN in navigation, its employment of a self-attention module is crucial. And TVT's performance in tasks involving long-term temporal credit assignment requires event memory, but also critically depends on a mechanism enabling the retrospective revaluation of past events in light of later reward.

A lesson here is that extracting insights about biological episodic memory from implementations of event memory in AI requires attending not only to the similarities and differences between the relevant memory systems (Sect. 2), but also to those between the cognitive systems in which they're embedded. To facilitate such comparisons, more systematic behavioural comparisons between artificial and biological systems would be fruitful. Methodological frameworks from comparative cognition would be highly applicable here. For instance, 'signature testing' frameworks (Taylor et al., 2022) could be used to evaluate similarities between artificial and biological systems by investigating fine-grained behavioural patterns, biases and limitations.

Function	Brief Explanation	Example
Discovering optimal strategies	Finding solutions to problems not uncovered by comparable agents	DQN uncovering the 'tunnelling' strategy in Atari Breakout
Fast learning	Learning more quickly than comparable agents, even approaching one-shot learning	MFEC uncovering a rewarding strategy in the Labyrinth environment.
Exploration	Spending time exploring the environment or option space, rather than exploiting the current policy	STM exploring more than comparable agents in a 3D foraging task, leading it to obtain more rewards.
Navigation	Navigating or planning routes more effectively, particularly in novel environments.	HCAM and EPN outperforming baselines in One-Shot Street Learn, in which agents must reach goals in a short time in a novel environment.
Acting over temporal distance	Selecting optimal behaviour given earlier actions or observations, or performing hierarchical tasks involving sequences of unrewarding actions	Dreamer v3 mining a diamond in Minecraft.

Table 2. Overview of areas in which AI systems with event memory have exhibited strong performance.

4. Learning about episodic memory from AI systems

Despite differences between the two, there are several ways in which event memory in AI systems might shed light on biological episodic memory.

First, event memory algorithms could be designed to more closely capture our best theories about the operations of biological episodic memory, enabling these theories to be tested. A particular benefit is that experimental methods which raise ethical challenges when performed on animals, such as lesion studies, would be unproblematic when performed on AI systems. Such studies could be informative about the contributions of particular processes to the behaviour of the episodic memory system and the agent. These methods could also assist in uncovering the causal role function of episodic memory processes in humans (De Brigard, 2014) – that is, determining how episodic memory processes interact with other cognitive processes, what kind of inputs and outputs they take, and what behavioural patterns they produce. Currently, we rely largely on clinical populations to test hypotheses about the causal role function of episodic memory processes (Balota et al., 1999; Budson et al., 2003; Ciaramelli et al., 2006; Melo et al., 1999), meaning that the hypotheses we can test are constrained by the particular lesions of individuals, which are rarely clear-cut. Studying event memory implemented in artificial systems would grant more flexibility, as we could cleanly lesion models to target hypotheses about episodic memory's function without any ethical concern.

However, to enable this development more cross-disciplinary collaboration would be necessary. Algorithms are often designed with a focus on solving problems which will advance the AI industry, rather than on developing models of biological cognitive systems. Similarly, benchmarks are typically established with reference to other AI systems' performance on the same tasks, rather than the performance of humans or other animals. Currently, only a minority of studies include any cross-system comparison (Fayyaz et al., 2022; Gershman & Daw, 2017). This makes it difficult to determine whether event memory architectures are suitably biologically grounded. What is critical is whether the degree and type of similarity between these systems and biological episodic memory is suitable, given the inferences we are interested in drawing (Stinson, 2020). At present, we are poorly positioned to answer such questions. To establish the utility of AI systems with event memory as models of biological episodic memory, a practice of rigorously testing algorithms on a wide range of standard memory tasks is required.

These considerations are particularly acute when we consider the potential for these systems to provide 'how-actually' explanations of biological episodic memory (see Box 1) – that is, to support theories about the structure, operations or function of episodic memory. For artificial event memory to play this evidential role, we would need high confidence that it resembled biological episodic memory in relevant respects. What counts as a relevant respect depends on what we are trying to model. For example, if we are interested in modelling the workings of pyramidal cells within the hippocampus in memory consolidation, it would be relevant that our model captures the organization and activities of these cells. But it is not clear that this would be a necessary feature to model if we were interested in modelling how episodic memories are semanticized over time.

Knowing which details are relevant to capture is a complex question which starts from the aim of the model.

Even when we are not confident in their resemblance to biological episodic memory, artificial event memory can play a productive role in our theorising by providing 'how-possibly' explanations. Whereas how-actually explanations aim to explain how a system actually operates, how-possibly explanations are more loosely constrained conjectures about ways a system might operate or contribute to a larger system. When systems with event memory are used to construct how-possibly explanations, the immediate intention is not to provide evidence about the way episodic memory actually works. Rather, the aim is to provide 'proof of concept' for ideas about the operations or functions of episodic memory, and highlight some hypotheses as 'pursuit-worthy' (Nyrup, 2020). How-possibly explanations can be a useful heuristic tool, particularly in the face of uncertainty, as generating several how-possibly explanations enables us to explore the hypothesis space. We may then test the more promising of these how-possibly explanations against each other to get closer to a how-actually explanation (Craver, 2006). Of course, not all how-possibly explanations are equally promising: in cases where we know that an event memory system differs from episodic memory in relevant respects, the how-possibly explanations it generates will be implausible. But in many cases, it may be an open empirical question whether the event memory system sufficiently resembles episodic memory, or whether any differences we've identified matter. In these cases, the event memory system does not provide how-actually explanations, but we can productively treat the how-possibly explanations it generates as candidates for further investigation.

Some researchers are already beginning to use event memory in AI systems to generate how-possibly explanations about the functions of episodic memory. Zeng et al. (2023) compare three AI systems to investigate episodic memory's role in spatial learning, comparing an episodic control algorithm, a replay algorithm and an 'online learning' algorithm with no event memory. These are used to model, respectively, the retrieval of episodic memory to guide online behaviour, episodic memory replay supporting offline learning, and a hippocampal lesion in which episodic memory is abolished. When tested in a series of spatial navigation tasks, both event memory algorithms outperformed the online learning algorithm, whose performance was highly variable. The episodic control algorithm learned more quickly than the replay algorithm, but attained lower asymptotic performance. The researchers suggest that episodic memory may play two complementary roles in spatial learning, which are difficult to distinguish in biological systems due to their shared neural basis. Additionally, the researchers compared different sampling methods for experience replay, and found that a biologically-inspired sequential replay method outperformed random sampling when the number of replays is limited. On this basis, they hypothesise that the sequential nature of hippocampal replay may be explained by the brain's limited computational resources constraining replay. Importantly, the researchers do not claim that the brain implements the very same processes as any of these algorithms and caution that their results are not sufficient to establish strong (how-actually) claims about biological episodic memory, but point toward 'potential explanations' – how-possibly explanations – about aspects of episodic memory.

Box 1: how-actually vs how-possibly

How-actually explanations: Explanations describing how episodic memory actually operates or functions. Based on systems which are known to resemble episodic memory in relevant respects.

How-possibly explanations: Loosely constrained conjectures about how episodic memory might operate or function. Based on systems whose resemblance to episodic memory is uncertain. Used to explore the hypothesis space and identify promising hypotheses to test.

A similar approach might be used to investigate other ideas about episodic memory's causal role function. As noted above (Sect 2.1), a prominent view is that episodic memory's primary function is to support mental time travel into the future, which in turn facilitates future planning. A key piece of evidence for this view is that neural activation when these capacities are exercised partially overlap, and the capacities are impaired by similar patterns of brain damage. But this same evidence makes it difficult to precisely determine the relationship between the capacities: for instance, is episodic memory's role in planning entirely mediated by future-oriented simulation, or could it support planning on its own? The TVT algorithm (sect. 3.5) is relevant here: it learns to engage in planning-like behaviour, fetching a key that will be needed later in a task resembling the 'spoon test'. TVT's event memory is critical to its success, but as a model-free algorithm performing no forward rollouts it is unlikely that it engages in something like future-oriented simulation. This does not show that future-oriented simulation plays no role in planning in biological organisms, but it does suggest that a system like episodic memory could support planning even in the absence of future-oriented simulation. This possibility is one we should consider when constructing and testing theories about episodic memory's function.

Similarly, in the parallel debate about episodic memory's evolution, it has been proposed that episodic memory is an evolutionary by-product with no adaptive function of its own (Schulz & Robins, 2023). But the many benefits apparently conferred on AI systems by event memory provide how-possibly explanations of how a system like episodic memory could be adaptive. For instance, as noted above, the MFEC algorithm displayed very fast learning in a complex environment with few repeated states and sparse rewards, leading the researchers to propose that using event memories for online decision making confers advantages at the early stages of learning in novel environments. Whilst this cannot straightforwardly be treated as evidence about the function of episodic memory in biological organisms, it does show that a system like episodic memory could confer learning advantages which would be adaptive under certain conditions – a possibility that should be fully considered in developing an account of episodic memory's evolution. More generally, we suggest that each of the advantages apparently conferred on AI systems by event memory discussed in Sect. 3 suggests a corresponding hypothesis about episodic memory's evolutionary and cognitive role functions worthy of further pursuit.

5. Conclusion

AI systems with event memory, inspired by human episodic memory, appear to enjoy behavioural advantages in several domains. Such systems could serve as an important tool for evaluating competing hypotheses about episodic memory's operations and functions in biological systems, enabling us to test these hypotheses in an ethical way. There are two obstacles to this at present. First, event memory systems in AI implement only some aspects of episodic memory, departing from it in crucial respects. Second, existing research does not typically use methods capable of isolating the contribution of event memory and its component processes to the enclosing system's behaviour. Further research, ideally involving cross-disciplinary collaborations, would be needed to establish these systems as models of episodic memory and draw robust conclusions. Nevertheless, we suggest that even at this early stage, event memory in AI systems stands to play an important role in the development of theories about episodic memory's causal role function by highlighting the pursuit-worthiness of novel or neglected hypotheses. As a preliminary step in this direction, we propose that the evidence reviewed here motivates giving greater consideration to the possibility that episodic memory has evolutionary and causal role functions in the domains of learning, decision-making, exploration, navigation and temporally extended action, in addition to – and perhaps independently of – its role in future-oriented simulation.

Acknowledgments

Both authors were supported by a UKRI Future Leaders Fellowship (grant number MR/W00741X/1) awarded to Alexandria Boyle. Andrea Blomkvist was also supported by a British Academy Postdoctoral Fellowship (grant number PFSS23\230115).

Data Access

There are no data associated with this article.

Open Access

References

Addis, D. R., Pan, L., Vu, M. A., Laiser, N., & Schacter, D. L. (2009). Constructive episodic simulation of the future and the past: Distinct subsystems of a core brain network mediate imagining and remembering. *Neuropsychologia*, *47*(11), 2222–2238. https://doi.org/10.1016/j.neuropsychologia.2008.10.026

Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., & Zaremba, W. (2017). Hindsight experience replay. *Advances in Neural Information Processing Systems*, *30*. https://goo.gl/SMrQnI.

Aronowitz, S. (2019). Memory is a modeling system. *Mind and Language*, *34*(4), 483–502. https://doi.org/10.1111/mila.12220

Balota, D. A., Cortese, M. J., Duchek, J. M., Adams, D., Roediger, H. L., McDermott, K. B., & Yerys, B. E. (1999). Veridical and false memories in healthy older adults and in dementia of the Alzheimer's type. *Cognitive Neuropsychology*, *16*(3–5), 361–384. https://doi.org/10.1080/026432999380834

Baran, B., Pace-Schott, E. F., Ericson, C., & Spencer, R. M. C. (2012). *Behavioral/Systems/Cognitive Processing of Emotional Reactivity and Emotional Memory over Sleep*. https://doi.org/10.1523/JNEUROSCI.2532-11.2012

Bartlett, F. C. (1932). *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press.

Bhat, A. A., Mohan, V., Rea, F., Sandini, G., & Morasso, P. (2014). "Connecting experiences": Towards a biologically inspired memory for developmental robots. *IEEE ICDL-EPIROB 2014 - 4th Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*, 359–365. https://doi.org/10.1109/DEVLRN.2014.6983007

Blomkvist, A. (2022). Imagination as a skill: A Bayesian proposal. *Synthese*, *200*(2), 1–23. https://doi.org/10.1007/S11229-022-03550-Z/TABLES/3

Blundell, C., Uria, B., Pritzel, A., Li, Y., Ruderman, A., Leibo, J. Z., Rae, J., Wierstra, D., & Hassabis, D. (2016). *Model-free episodic control*. http://arxiv.org/abs/1606.04460

Bowen, H. J., Kark, S. M., & Kensinger, E. A. (2017). *NEVER forget: negative emotional valence enhances recapitulation*. https://doi.org/10.3758/s13423-017-1313-9

Boyle, A. (2019). Learning from the past: Epistemic generativity and the function of episodic memory. *Journal of Consciousness Studies*, *26*(5–6), 242–251. https://doi.org/10.17863/CAM.35867

Boyle, A. (2020a). Remembering events and representing time. *Synthese*. https://doi.org/10.1007/s11229-020-02896-6

Boyle, A. (2020b). The impure phenomenology of episodic memory. *Mind & Language*, *35*(5), 641–660. https://doi.org/10.1111/mila.12261

Brown, S. A. B. (2023). Episodic Memory and Unrestricted Learning. *Philosophy of Science*, 1–29. https://doi.org/DOI: 10.1017/psa.2023.16

Bu, F., & Chang, D. E. (2020). *Double Prioritized State Recycled Experience Replay*. https://doi.org/10.1109/ICCE-Asia49877.2020.9276975

Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, *11*(2), 49–57. https://doi.org/10.1016/j.tics.2006.11.004

Budson, A. E., Sullivan, A. L., Daffner, K. R., & Schacter, D. L. (2003). Semantic versus phonological false recognition in aging and Alzheimer's disease. *Brain and Cognition*, *51*(3), 251–261. https://doi.org/10.1016/S0278-2626(03)00030-7

Camp, B., Mandivarapu, J. K., & Estrada, R. (2020). *Continual Learning with Deep Artificial Neurons*. http://arxiv.org/abs/2011.07035

Chin, W. H., Toda, Y., Kubota, N., Loo, C. K., & Seera, M. (2019). Episodic memory multimodal learning for robot sensorimotor map building and navigation. *IEEE Transactions on Cognitive and Developmental Systems*, *11*(2), 210–220. https://doi.org/10.1109/TCDS.2018.2875309

Ciaramelli, E., Ghetti, S., Frattarelli, M., & Làdavas, E. (2006). When true memory availability promotes false memory: Evidence from confabulating patients. *Neuropsychologia*, *44*(10), 1866–1877. https://doi.org/10.1016/j.neuropsychologia.2006.02.008

Clayton, N., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*(6699), 272–274. https://doi.org/10.1038/255243a0

Craver, C. F. (2006). When mechanistic models explain. *Synthese*, *153*(3), 355–376. https://doi.org/10.1007/S11229-006-9097-X/METRICS

Davies, J. R., Garcia-Pelegrin, E., Baciadonna, L., Pilenga, C., Favaro, L., & Clayton, N. S. (2022). Episodic-like memory in common bottlenose dolphins. *Current Biology*, *32*(15), 3436-3442.e2. https://doi.org/10.1016/j.cub.2022.06.032

De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, *191*(2), 155–185. https://doi.org/10.1007/s11229-013-0247-7

Fang, M., Zhou, C., Shi, B., Gong, B., Xu, J., & Zhang, T. (2019). *DHER: Hindsight experience replay for dynamic goals*. https://github.com/mengf1/DHER.

Fayyaz, Z., Altamimi, A., Zoellner, C., Klein, N., Wolf, O. T., Cheng, S., & Wiskott, L. (2022). A Model of Semantic Completion in Generative Episodic Memory. *Neural Computation*, *34*(9), 1841–1870. https://doi.org/10.1162/neco_a_01520

Fei-Fei, L., Fergus, R., & Perona, P. (2006). One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *28*(4). https://doi.org/10.1109/TPAMI.2006.79

Gershman, S. J., & Daw, N. D. (2017). Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework. *Annual Review of Psychology*, *68*, 101–128. https://doi.org/10.1146/annurev-psych-122414-033625

Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T., & Diamond, M. (2023). *Mastering diverse domains through world models*. https://arxiv.org/pdf/2301.04104v1.pdf

Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, *95*(2), 245–258. https://doi.org/10.1016/j.neuron.2017.06.011

Hayes, T. L., Krishnan, G. P., Bazhenov, M., Siegelmann, H. T., Sejnowski, T. J., & Kanan, C. (2021). Replay in Deep Learning: Current Approaches and Missing Biological Elements. *Neural Computation*, *33*(11), 2908–2950. https://doi.org/10.1162/NECO_A_01433

Hemmer, P., & Steyvers, M. (2009). A Bayesian Account of Reconstructive Memory. *Topics in Cognitive Science*, *1*, 189–202. https://doi.org/10.1111/j.1756-8765.2008.01010.x

Hung, C. C., Lillicrap, T., Abramson, J., Wu, Y., Mirza, M., Carnevale, F., Ahuja, A., & Wayne, G. (2019). Optimizing agent behavior over long time scales by transporting value. *Nature Communications*, *10*(1). https://doi.org/10.1038/s41467-019-13073-w

Keven, N. (2016). Events, narratives and memory. *Synthese*, *193*(8), 2497–2517. https://doi.org/10.1007/s11229-015-0862-6

Kim, T., & Har, D. (2022). *Failed Goal Aware Hindsight Experience Replay*. http://arxiv.org/abs/2208.14741

Kitamura, T., Ogawa, S. K., Roy, D. S., Okuyama, T., Morrissey, M. D., Smith, L. M., Redondo, R. L., & Tonegawa, S. (2017). Engrams and circuits crucial for systems consolidation of a memory. *Science*, *365*(6333). https://www.science.org

Kumar, R., & Nagaraj, D. (2022). *Introspective Experience Replay: Look Back When Surprised*. http://arxiv.org/abs/2206.03171

Kumaran, D., & Hassabis, D. (2015, February 25). *From pixels to actions: Human-level control through deep reinforcement learning*. Google Research Blog. https://blog.research.google/2015/02/from-pixels-to-actions-human-level.html?m=1

Lampinen, A. K., Chan, S. C. Y., Banino, A., & Hill, F. (2021). Towards mental time travel: a hierarchical memory for reinforcement learning agents. *Advances in Neural Information Processing Systems*, *34*.

Lee, S. Y., Choi, S., & Chung, S.-Y. (2018). *Sample-Efficient Deep Reinforcement Learning via Episodic Backward Update*. http://arxiv.org/abs/1805.12375

Lengyel, M., & Dayan, P. (2008). Hippocampal contributions to control: The third way. *Advances in Neural Information Processing Systems 20 - Proceedings of the 2007 Conference*.

Li, B., Lu, T., Li, J., Lu, N., Cai, Y., & Wang, S. (2020). ACDER: Augmented Curiosity-Driven Experience Replay. *Proceedings - IEEE International Conference on Robotics and Automation*, 4218–4224. https://doi.org/10.1109/ICRA40945.2020.9197421

Liu, D., Cong, M., Zou, Q., & Du, Y. (2017). A biological-inspired episodic cognitive map building framework for mobile robot navigation. *International Journal of Advanced Robotic Systems*, *14*(3). https://doi.org/10.1177/1729881417705922/ASSET/IMAGES/LARGE/10.1177_1729881417705922-FIG14.JPEG

Mahr, J. B., & Csibra, G. (2018). Why do we remember? The communicative function of episodic memory. *Behavioral and Brain Sciences*, *41*(E1), 1–16. https://doi.org/10.1017/S0140525X17000012

Marr, D. (1982). *Vision: A Computational Approach*. Freeman & Co.

Mccarthy, R., Wang, Q., & Redmond, S. J. (2023). *Imaginary hindsight experience replay: Curious model-based learning for sparse reward tasks*. https://arxiv.org/pdf/2110.02414.pdf

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457.

Melo, B., Winocur, G., & Moscovitch, M. (1999). False recall and false recognition: An examination of the effects of selective and combined lesions to the medial temporal lobe/diencephalon and frontal lobe structures. *Cognitive Neuropsychology*, *16*(3–5), 343–359. https://doi.org/10.1080/026432999380825

Michaelian, K. (2016). *Mental time travel: Episodic memory and our knowledge of the personal past*. MIT Press.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. https://doi.org/10.1038/nature14236

Mohan, V., Sandini, G., & Morasso, P. (2014). A neural framework for organization and flexible utilization of episodic memory in cumulatively learning baby humanoids. *Neural Computation*, *26*(12), 2692–2734. https://doi.org/10.1162/NECO_a_00664

Norman, K. A., Detre, G., & Polyn, S. M. (2012). Computational Models of Episodic Memory. *The Cambridge Handbook of Computational Psychology*, 189–225. https://doi.org/10.1017/CBO9780511816772.011

Nuxoll, A. M., & Laird, J. E. (2012). Enhancing intelligent agents with episodic memory. *Cognitive Systems Research*, *17–18*, 34–48. https://doi.org/10.1016/j.cogsys.2011.10.002

Nyrup, R. (2020). Of water drops and atomic nuclei: Analogies and pursuit worthiness in science. *British Journal for the Philosophy of Science*, *71*(3), 881–903. https://doi.org/10.1093/BJPS/AXY036

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

Payne, J. D., Stickgold, R., Swanberg, K., & Kensinger, E. A. (2008). *Sleep Preferentially Enhances Memory for Emotional Components of Scenes*.

Pritzel, A., Uria, B., Srinivasan, S., Badia, A. P., Vinyals, O., Hassabis, D., Wierstra, D., & Blundell, C. (2017). Neural episodic control. *34th International Conference on Machine Learning, ICML 2017*, *6*.

Ritter, S., Faulkner, R., Sartran, L., Santoro, A., Botvinick, M., & Raposo, D. (2021). Rapid task-solving in novel environments. *ICLR 2021 - 9th International Conference on Learning Representations*.

Ritter, S., Wang, J. X., Kurth-Nelson, Z., Jayakumar, S. M., Blundell, C., Pascanu, R., & Botvinick, M. (2018). Been there, done that: Meta-learning with episodic recall. *35th International Conference on Machine Learning, ICML 2018*, *10*.

Rothfuss, J., Ferreira, F., Aksoy, E. E., Zhou, Y., & Asfour, T. (2018). Deep Episodic Memory: Encoding, Recalling, and Predicting Episodic Experiences for Robot Action Execution. *IEEE Robotics and Automation Letters*, *3*(4), 4007–4014. https://doi.org/10.1109/LRA.2018.2860057

Rotinov, E. (2019). *Reverse Experience Replay*. http://arxiv.org/abs/1910.08780

Rubin, D. C., & Umanath, S. (2015). Event memory: A theory of memory for laboratory, autobiographical, and fictional events. *Psychological Review*, *122*(1), 1–23. https://doi.org/10.1037/a0037907

Sakaki, M., Fryer, K., & Mather, M. (2013). Emotion strengthens high-priority memory traces but weakens low-priority memory traces. *Https://Doi.Org/10.1177/0956797613504784*, *25*(2), 387–395. https://doi.org/10.1177/0956797613504784

Scarf, D., Smith, C., & Stuart, M. (2014). A spoon full of studies helps the comparison go down: a comparative analysis of Tulving's spoon test. *Frontiers in Psychology*, *5*(AUG). https://doi.org/10.3389/FPSYG.2014.00893

Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*, 773–786. https://doi.org/10.1098/rstb.2007.2087

Schacter, D. L., & Addis, D. R. (2020). Memory and Imagination: Perspectives on Constructive Episodic Simulation. *The Cambridge Handbook of the Imagination*, 111–131. https://doi.org/10.5860/choice.50-4715

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The Future of Memory: Remembering, Imagining, and the Brain. *Neuron*, *76*(4), 677–694. https://doi.org/10.1016/j.neuron.2012.11.001

Schacter, D. L., Guerin, S. A., & St. Jacques, P. L. (2011). Memory distortion: an adaptive perspective. *Trends in Cognitive Sciences*, *15*(10), 467. https://doi.org/10.1016/J.TICS.2011.08.004

Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2015). *Prioritized Experience Replay*. http://arxiv.org/abs/1511.05952

Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., & Ranganath, C. (2015). Memory and space: Towards an understanding of the cognitive map. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *35*(41), 13904–13911. https://doi.org/10.1523/JNEUROSCI.2618-15.2015

Schulz, A. W., & Robins, S. (2023). Episodic memory, simulated future planning, and their evolution. *Review of Philosophy and Psychology*, *14*(3). https://doi.org/10.1007/s13164-021-00601-1

Shevlin, H., & Halina, M. (2019). Apply rich psychological terms in AI with care. *Nature Machine Intelligence*, *1*(4), 165–167. https://doi.org/10.1038/s42256-019-0039-y

Sovrano, F. (2019). Combining experience replay with exploration by random network distillation. *IEEE Conference on Computatonal Intelligence and Games, CIG*, *2019-August*. https://doi.org/10.1109/CIG.2019.8848046

Stinson, C. (2020). From implausible artificial neurons to idealized cognitive models: Rebooting philosophy of artificial intelligence. *Philosophy of Science*, *87*(4), 590–611. https://doi.org/10.1086/709730

Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences*, *30*(3), 299–351. https://doi.org/10.1017/S0140525X07001975

Taylor, A. H., Bastos, A. P. M., Brown, R. L., & Allen, C. (2022). The signature-testing approach to mapping biological and artificial intelligences. *Trends in Cognitive Sciences*, *26*(9), 738–750. https://doi.org/10.1016/J.TICS.2022.06.002

Teyler, T. J., & DiScenna, P. (1986). The Hippocampal Memory Indexing Theory. *Behavioral Neuroscience*. https://doi.org/10.1037/0735-7044.100.2.147

Teyler, T. J., & Rudy, J. W. (2007). The Hippocampal Indexing Theory and Episodic Memory: Updating the Index. *Hippocampus*, *17*, 1158–1169.

Tulving, E. (1983). *Elements of Episodic Memory*. Oxford University Press.

Tulving, E. (2005). Episodic memory and autonoesis: uniquely human? In H. S. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition* (pp. 4–56). Oxford University Press.

Vernon, D., Metta, G., & Sandini, G. (2007). *The iCub Cognitive Architecture: Interactive Development in a Humanoid Robot*.

Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., & Botvinick, M. (2016). *Learning to reinforcement learn*. https://arxiv.org/abs/1611.05763v3

Yuan, J., Guo, W., Zha, F., Li, M., & Sun, L. (2022). Method of robot episode cognition based on hippocampus mechanism. *IEEE Access*, *10*, 42386–42395. https://doi.org/10.1109/ACCESS.2021.3119530

Yuan, Y., & Mattar, M. G. (2021). *Improving Experience Replay with Successor Representation*. http://arxiv.org/abs/2111.14331

Zakharov, A., Crosby, M., & Fountas, S. (2020). *Episodic memory for learning subjective-timescale models*. https://arxiv.org/abs/2010.01430v1

Zeng, X., Diekmann, N., Wiskott, L., & Cheng, S. (2023). Modeling the function of episodic memory in spatial learning. *Frontiers in Psychology*, *14*, 1160648. https://doi.org/10.3389/FPSYG.2023.1160648/BIBTEX