Are robots stuck in time? Learning to represent actions in long horizon planning by learning to remember the past

Chad DeChant
Computer Science Department
Columbia University
chad.dechant@columbia.edu

Abstract: In order to do effective long horizon planning it will be necessary for robots to use representations of action sequences of various durations and levels of abstraction. Many of these will need to be of quite long duration, representing entire episodes or sequences of episodes of actions. Like the most successful image and text representations, these will need to be learned. Humans learn and use such representations of action episodes partly in taking actions but also in remembering their past actions. If robots are to develop the kind of expressive and flexible representations of action that humans have, they will also have to learn to represent and remember their past actions. These representations of past actions will, in turn, be able to serve as the basis for useful kinds of reasoning about the future, as they are thought to do in humans.

Keywords: Long horizon tasks, episodic memory, representation learning

1 Introduction

"Animals are cognitively stuck in time; that is, they have no sense of time and thus have no episodic memory or ability to anticipate long-range future events."
William Roberts, *"Are animals stuck in time?"* [1]

Like the animals mentioned above, robots are stuck in time. They occupy a small slice of the present and the few steps into the future over which they typically plan. They cannot represent sufficiently extended action sequences to do very long horizon planning and they typically have little to no true episodic memory. These are twin deficits and are very much related. An agent's ability to plan for long action sequences in the future would benefit from learning to represent the past in at least three ways.

First, and most importantly, the structure of the representations themselves and the algorithms to process them can and ought to be shared whether the episodes in question are in the past or the future. The factors that make a useful action representation, including what elements to represent and how, need not differ depending on whether the actions are being remembered or planned.

Second, stored episodic memories of past actions and sequences of actions can be used in planning future actions. In humans such memories are thought by some to be the building blocks for generating flexible plans and scenarios [2].

Third, learning to represent the past can help provide enough data to learn representations which can be used to reason about the future. It is not obvious how to structure or use representations of action, past or future. Rather than try to hand pick what we think ought to go into such representations, it will be necessary to learn them. In order to do that, large amounts of data will be needed, along with

Learning, Perception, and Abstraction for Long-Horizon Planning Workshop
6th Conference on Robot Learning (CoRL 2022), Auckland, New Zealand.

tasks which use that data to produce useful representations. Learning to represent, store, recall, and reason over episodes of past actions would be a significant source of training signal to shape and structure representations of such actions. It will likely be useful, even necessary, to jointly learn to represent past and future actions in order to benefit from the various demands that each task would make on such representations.

2 Learning representations of action episodes

What elements make up representations of actions and action sequences in the human brain? Cognitive scientists have many theories about that, some of which may serve as useful inspiration when designing memory-equipped learning architectures, but it is still very much an open research area. Rather than try to use these theories to manually pick and choose which elements might best make up an action representation, it would be better to use them as potentially helpful guides in answering a more tractable question: what data and tasks do we need in order to learn the structure of such representations?

2.1 Data

In order to learn rich representations of images and language that can be used in various downstream tasks, at least part of what is needed is now clear: large datasets of images and text, respectively. It is less clear what is needed to learn to represent extended action episodes for use in remembering the past and acting in the future. Therefore it will be necessary to capture as much data as possible in the real world or in very realistic simulations. The data will have to be exhaustive in two ways: it will need to cover episodes taking place over long periods of time and it will have to involve capturing as much data as a robot has access to, including its internal states. The latter is important because to exclude any available data would be to decide *a priori* what an episodic memory or plan representation should include. The actual memory and planning representations themselves will of course have to be significantly smaller in size than all of the data that might feed into them; the point of work in this direction will be learning what to compress and where the representational bottlenecks should be.

2.2 Tasks

Just as large amounts of data will be needed, so will a wide range of tasks be required. The variety and demandingness of the tasks used for training will shape the representations and influence how useful they are in planning new actions for novel tasks in unfamiliar environments. As many long horizon tasks as possible should be paired with an ever expanding set of tasks that explicitly require and test the formation of episodic memory abilities. These should be diverse enough to cover the range of ways that humans access their episodic memories.

Humans physically enact their memories nearly constantly, from the very simple and everyday (e.g. returning to a room to retrieve an object) to the more complicated and skilled (e.g. repeating a series of learned or observed movements in a workplace or while performing a dance). These very physical manifestations of memory are natural testing grounds for embodied robots. For example, after performing a long horizon task, a robot could be asked to repeat the action. This action replay could be either a verbatim reenacting of all of its actions in the sequence, including those it had to backtrack from or correct, or it could be a more abstract repetition, merely redoing the 'important' actions in the correct order. Both kinds of repetition would encourage and test different episodic memory skills which could be used to improve low and high level planning, respectively. Changing the order of such required repetitions, including performing actions in reverse order, could also strengthen episodic memory representations, and do so in ways which have been studied in humans in internally visualized replay of memories [3, 4]

Language and memory

Humans often convey their memories using language. Some researchers have suggested that language and its representational structure is an essential component of forming true autobiographical episodic memories, and that it is in part the lack of linguistic abilities which leads to the (hypothesized, but disputed [5]) absence of such memories in nonhuman animals [6]. Whether or not language has such an essential role in the formation of episodic memories, it is clear that for humans at least language plays a central role in how and why we remember events. Recounting our own past actions in conversation and stories is a major use of human language [7] and almost certainly shapes the structure of our episodic memories to enable them to be accessed for such interpersonal communication [8].

Language should therefore be used as one of the primary tools for guiding the development of representations of episodic memory. DeChant and Bauer [9] proposed that robots should learn to store and access episodic memories in order to summarize their past actions. Such a summary of past action corresponds to the highest level, most abstract form of episodic memory, described by researchers studying human memory as a 'gist' [10]. Language can also be used to encourage more detailed episodic memory by requiring an agent to produce narratives of its actions at various levels of granularity [11].

Episodic memories could also be significantly shaped by training an agent to answer questions about its past actions. In order to answer such questions, particularly unforeseen questions, an agent would have to learn to represent information which perhaps did not seem important during the execution of the action (such as the location of objects not interacted with). Aspects of an episode which seemed unimportant at the time could be significant in understanding corrective feedback or in relating the episode to a new situation. Developing representations of action sequences which enable question answering could, in turn, be useful when planning an action sequence; a robot could, for example, answer questions about its intended action plan before carrying it out.

2.3 Scaffolding of learning and data

Learning representations of action episodes to be used for both long horizon planning and episodic memory would seem to face something of a chicken and egg problem: how can we learn memories of long horizon tasks being performed before we have robots that can perform the kind of long horizon tasks we are interested in? It will be necessary to slowly build up to very long episodes and sequences of episodes.

Learning to make and use episodic memories can begin with episodes shorter in time than we would eventually like to be able to represent, including those involving tasks currently being worked on. These can also be artificially chained together to simulate very long sequences of action, though without the simultaneous ability to then use these representations for very long horizon planning.

It is possible that there can be some element of transfer learning from much simpler episodically structured data and tasks. Learning to represent long videos of action sequences might allow the initial development of representations capturing visual and event information similar to that which robots would need to represent. Even learning representations of long text narratives of action or operating in interactive text-based environments might provide some additional data.

Because data from an episode may be captured, especially in a virtual environment, and reused arbitrarily many times, it may initially be easier to train representations for episodic memory than for planning. Many attempts at training memory representations can be made on the same saved episode without having to re-run a simulation or reset a robot in the real world. In order to ensure that representations are equally useful for memory and for planning, this imbalance should be minimized if possible. However, there might be some justification for a slight imbalance: humans can recall our past actions many times but we can only take those actions once.

3 Related work

3.1 Memory in learning agents

Nematzadeh et al. [12] discuss the importance of memory of various kinds, including episodic memory, for humans and discuss how AI systems may better incorporate it. Some work has already been done on incorporating episodic memory into agents which learn to perform tasks, usually in virtual environments, often being trained using reinforcement learning. Most related to our proposal is work by Lampinen et al. [13] which proposes a hierarchical attention mechanism to access past events. Wayne et al. [14] develop an episodic memory architecture that stores compressed representations of episode states and enables an agent to complete various tasks which require memory and are inspired by experiments designed to test humans' memory skills. Santoro et al. [15] investigate the utility of episodic memory compared to more general representations of statistical regularities in an RL agent. Ritter et al. [16] equip an LSTM meta-learning agent with episodic memory.

Many other efforts have been made to incorporate memory into deep learning networks of various kinds, from LSTMs [17] and Neural Turing Machines [18] to the more recent Token Turing Machines [19] which learn how to read and write to memory in order to perform various tasks, including a real-world robot learning task [20].

3.2 Episodic memory in humans

Endel Tulving distinguished episodic from semantic memory and linked it to humans' ability to imaginatively project themselves into the future [21, 22]. Ranganath [23] and Xue [24] provide very recent overviews of research into episodic memory in humans. Szpunar et al. [25], Addis and Tanguay [26], and Wang et al. [27] review the role of episodic memory in so-called prospective cognition (thinking about the future) and deliberation. Klein [28] argues that although memory is usually thought of as being about the past, its role is primarily to anticipate and plan for the future.

Cox and Shiffrin [29] provide a survey of how computational models of episodic memory have been used in researching human memory; see Franklin et al. [30] for a recent example. How humans break up long tasks into shorter subgoals is likely influenced by how we break up memories into episodes, making work on such event segmentation particularly relevant [31, 32]. Memories are thought by many to be categorized and represented using learned *schema* templates which in turn shape how new events are understood (and, perhaps, planned) [33, 34].

4 Conclusion

Over the past two decades, cognitive scientists have found increasing evidence for the tight relationship between humans' ability to represent the past and reason about the future [26]. Episodic memory is thought to play a key role in various kinds of thinking about the future, including in planning future actions, simulating the future, forming and following intentions, and predicting the outcomes and value of various possible actions [25]. Brain imaging studies have established that imagining future actions recruits similar brain regions as remembering past actions. Damage which causes deficits in a person's memory often also lead to deficits in the ability to plan or imagine the future [35, 36].

Robots are currently limited in their ability to remember their past actions as well — not because of any unintended damage, but by design. And just as in humans, that inability may contribute to deficits in robotic planning.

It has been suggested that episodic memory is the form of memory which most recently evolved [37]. If so, it would be understandable and appropriate that it was previously largely overlooked by robot learning researchers. But because episodic memory is also conjectured to be the most flexible form of memory [37] and essential for human-like planning and reasoning about the future, it should be overlooked no longer.

References

[1] W. A. Roberts. Are animals stuck in time? *Psychological bulletin*, 128(3):473, 2002.

[2] D. L. Schacter, D. R. Addis, and R. L. Buckner. Episodic simulation of future events: Concepts, data, and applications. *Annals of the New York Academy of Sciences*, 1124(1):39–60, 2008.

[3] G. E. Wimmer, Y. Liu, N. Vehar, T. E. Behrens, and R. J. Dolan. Episodic memory retrieval success is associated with rapid replay of episode content. *Nature Neuroscience*, 23(8):1025–1033, 2020.

[4] Y. Liu, R. J. Dolan, Z. Kurth-Nelson, and T. E. Behrens. Human replay spontaneously reorganizes experience. *Cell*, 178(3):640–652, 2019.

[5] A. Boyle. Episodic memory in animals: Optimism, kind scepticism and pluralism. In *Current Controversies in Philosophy of Memory*, pages 189–205. Routledge.

[6] K. Nelson. The psychological and social origins of autobiographical memory. *Psychological science*, 4(1):7–14, 1993.

[7] D. Tannen et al. *Conversational style: Analyzing talk among friends*. Oxford University Press, 2005.

[8] K. Nelson and R. Fivush. The emergence of autobiographical memory: a social cultural developmental theory. *Psychological review*, 111(2):486, 2004.

[9] C. DeChant and D. Bauer. Toward robots that learn to summarize their actions in natural language: a set of tasks. In *Conference on Robot Learning*, pages 1807–1813. PMLR, 2022.

[10] G. Winocur and M. Moscovitch. Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5):766–780, 2011.

[11] C. DeChant and D. Bauer. Summarizing a virtual robot's past actions in natural language. *arXiv preprint arXiv:2203.06671*, 2022.

[12] A. Nematzadeh, S. Ruder, and D. Yogatama. On memory in human and artificial language processing systems. In *Proceedings of ICLR Workshop on Bridging AI and Cognitive Science*, 2020.

[13] A. Lampinen, S. Chan, A. Banino, and F. Hill. Towards mental time travel: a hierarchical memory for reinforcement learning agents. *Advances in Neural Information Processing Systems*, 34:28182–28195, 2021.

[14] G. Wayne, C.-C. Hung, D. Amos, M. Mirza, A. Ahuja, A. Grabska-Barwinska, J. Rae, P. Mirowski, J. Z. Leibo, A. Santoro, et al. Unsupervised predictive memory in a goal-directed agent. *arXiv preprint arXiv:1803.10760*, 2018.

[15] A. Santoro, P. W. Frankland, and B. A. Richards. Memory transformation enhances reinforcement learning in dynamic environments. *Journal of Neuroscience*, 36(48):12228–12242, 2016.

[16] S. Ritter, J. Wang, Z. Kurth-Nelson, S. Jayakumar, C. Blundell, R. Pascanu, and M. Botvinick. Been there, done that: Meta-learning with episodic recall. In *International conference on machine learning*, pages 4354–4363. PMLR, 2018.

[17] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[18] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

[19] M. S. Ryoo, K. Gopalakrishnan, K. Kahatapitiya, T. Xiao, K. Rao, A. Stone, Y. Lu, J. Ibarz, and A. Arnab. Token turing machines. *arXiv preprint arXiv:2211.09119*, 2022.

[20] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.

[21] E. Tulving. Episodic and semantic memory. 1972.

[22] E. Tulving. Episodic memory: From mind to brain. *Annual review of psychology*, 53(1):1–25, 2002.

[23] C. Ranganath. Oxford handbook of memory: Episodic memory (ranganath). 2022.

[24] G. Xue. From remembering to reconstruction: The transformative neural representation of episodic memory. *Progress in Neurobiology*, page 102351, 2022.

[25] K. K. Szpunar, R. N. Spreng, and D. L. Schacter. A taxonomy of prospection: Introducing an organizational framework for future-oriented cognition. *Proceedings of the National Academy of Sciences*, 111(52):18414–18421, 2014.

[26] D. R. Addis and A. Tanguay. Prospective cognition and its links with memory. 2022.

[27] S. Wang, S. F. Feng, and A. M. Bornstein. Mixing memory and desire: How memory reactivation supports deliberative decision-making. *Wiley Interdisciplinary Reviews: Cognitive Science*, 13(2):e1581, 2022.

[28] S. B. Klein. The temporal orientation of memory: It's time for a change of direction. *Journal of Applied Research in Memory and Cognition*, 2(4):222–234, 2013.

[29] G. E. Cox and R. M. Shiffrin. Computational models of event memory. 2022.

[30] N. T. Franklin, K. A. Norman, C. Ranganath, J. M. Zacks, and S. J. Gershman. Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review*, 127(3): 327, 2020.

[31] G. A. Radvansky and J. M. Zacks. Event boundaries in memory and cognition. *Current opinion in behavioral sciences*, 17:133–140, 2017.

[32] C. Baldassano, J. Chen, A. Zadbood, J. W. Pillow, U. Hasson, and K. A. Norman. Discovering event structure in continuous narrative perception and memory. *Neuron*, 95(3):709–721, 2017.

[33] A. Gilboa and H. Marlatte. Neurobiology of schemas and schema-mediated memory. *Trends in cognitive sciences*, 21(8):618–631, 2017.

[34] N. Varga, N. Morton, and A. Preston. Schema, inference, and memory. 2022.

[35] D. R. Addis, A. T. Wong, and D. L. Schacter. Age-related changes in the episodic simulation of future events. *Psychological science*, 19(1):33–41, 2008.

[36] D. Kwan, N. Carson, D. Addis, and R. Rosenbaum. Deficits in past remembering extend to future imagining in a case of developmental amnesia. *Neuropsychologia*, 48(11):3179–3186, 2010.

[37] T. Suddendorf and M. C. Corballis. The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and brain sciences*, 30(3):299–313, 2007.