

A Machine with Short-Term, Episodic, and Semantic Memory Systems

Kim, Taewoon; Cochez, Michael; François-Lavet, Vincent; Neerincx, Mark; Vossen, Piek

DOI

[10.1609/aaai.v37i1.25075](https://doi.org/10.1609/aaai.v37i1.25075)

Publication date

2023

Document Version

Final published version

Published in

AAAI-23 Technical Tracks 1

Citation (APA)

Kim, T., Cochez, M., François-Lavet, V., Neerincx, M., & Vossen, P. (2023). A Machine with Short-Term, Episodic, and Semantic Memory Systems. In B. Williams, Y. Chen, & J. Neville (Eds.), *AAAI-23 Technical Tracks 1* (pp. 48-56). (Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI 2023; Vol. 37, No. 1). American Association for Artificial Intelligence (AAAI). <https://doi.org/10.1609/aaai.v37i1.25075>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

A Machine with Short-Term, Episodic, and Semantic Memory Systems

Taewoon Kim¹, Michael Cochez¹, Vincent François-Lavet¹, Mark Neerincx², Piek Vossen¹

¹Vrije Universiteit Amsterdam

²Technische Universiteit Delft

{t.kim, m.cochez, vincent.francoislavet, p.t.j.m.vossen}@vu.nl, m.a.neerincx@tudelft.nl

Abstract

Inspired by the cognitive science theory of the explicit human memory systems, we have modeled an agent with short-term, episodic, and semantic memory systems, each of which is modeled with a knowledge graph. To evaluate this system and analyze the behavior of this agent, we designed and released our own reinforcement learning agent environment, “the Room”, where an agent has to learn how to encode, store, and retrieve memories to maximize its return by answering questions. We show that our deep Q-learning based agent successfully learns whether a short-term memory should be forgotten, or rather be stored in the episodic or semantic memory systems. Our experiments indicate that an agent with human-like memory systems can outperform an agent without this memory structure in the environment.

1 Introduction

The cognitive science theory suggests that humans have an explicit memory system, which is composed of semantic and episodic memory systems (Tulving 1983, 1985; Tulving and Thomson 1973). Semantic memory deals with general world knowledge, while episodic deals with one’s personal memory. For example, when asked: “Where are laptops *usually* located?”, you might be able to answer: “On a desk”, if you have successfully encoded and stored a relevant memory. However, you are unlikely able to recall when and where such memory (or knowledge) got to be stored in your memory system. Nonetheless, you were able to retrieve it. This is because this type of memory is semantic. Semantic memory contains (general) factual knowledge without momentary or time-bound information (i.e., where and when the event occurred). Let’s consider another question: “Where is Alice’s laptop?” and let us assume that you have observed where Alice’s laptop was. To answer this question, one revisits when and where this memory was encoded and stored. Retrieval of such a memory is the reconstruction process of the event. This type of memory is called episodic. These are more individual to you than semantic memories, include information about personal experiences, and often have an affective connotation.

Nowadays, machines have become reasonably good at answering factual knowledge questions. Search engines and

virtual assistants can retrieve relevant information (often harvested from the Internet) and answer factual and commonsense questions (i.e., those probing for semantic knowledge). However, these systems will not remember what they did yesterday (i.e., episodic memory). This is because they were not designed to answer such questions.

Consider, for example, a robot deployed to a nursery home whose job is to remind the elderly to take their pills every day. If such robot can successfully remember when and where it has seen them taking the pills, by observing the environment, it is a great help to them. However, it is not feasible for an agent to encode and store every observation as an episodic memory, as memory capacity is bounded. If the information is no longer stored, the agent can use commonsense knowledge (semantic memory) and say “You left your pills in the cabinet”.

Motivated by this type of examples, we have modeled an agent that has both semantic and episodic memory systems. The agent interacts with the environment and answers questions to maximize its reward. To make it more realistic, our agent can only partially observe the environment and hence it needs a memory system to be able to answer. Our hypothesis is that if it has two explicit memory systems, rather than one, it can answer the questions more successfully.

Knowledge graphs (KG) model data in the form of a graph, where entities are linked with directed relations. This way of representing data is not only readable / writable by both humans and machines, but also deductive and inductive reasoning can be performed. Moreover, once we convert them into knowledge graph embeddings (KGE), we can take advantage of neural networks to solve various problems (Hogan et al. 2022).

Although KGs can be built by experts with their domain knowledge, they are more powerful when they are open and crowdsourced, since they can take advantage of collective human knowledge. Wikidata (Vrandečić and Krötzsch 2014) and ConceptNet (Speer, Chin, and Havasi 2017) are some of the examples. ConceptNet tries to capture the meanings of the words, resulting in capturing common sense knowledge (e.g., (laptop, AtLocation, desk)). We used part of their knowledge graph for our experiments.

Reinforcement learning (RL) has become increasingly popular over the past few years (Arulkumaran et al. 2017). Deep Q-learning (Mnih et al. 2013), an RL that uses deep




Figure 1: The memory systems of the agent. The long-term (explicit) memory systems consist of episodic and semantic memory systems.

learning, has showed us that an RL agent can learn how to play Atari games. It finds an optimal policy by finding the optimal action-value function. RL is especially a powerful method, when it is infeasible for us humans to label every best state-action pair that the agent can take.

The contributions of this paper are as follows. 1) Inspired by the cognitive science theory, we explicitly model an agent with human-like memory systems (i.e., short-term, episodic, and semantic), each of which is modeled with a knowledge graph. 2) We designed and released our own environment, compatible with OpenAI Gym (Brockman et al. 2016), where an agent has to learn how to encode, store, and retrieve memories to maximize its return by answering questions. The environment is open-sourced at <https://github.com/tae898/room-env>. 3) Our deep Q-learning based RL agent has successfully learned a memory management policy by interacting with the environment. The agent and its implementation are open-sourced at <https://github.com/tae898/explicit-memory>.

The rest of this paper is organized as follows. In Section 2, we propose our method, which allows an agent to learn how to encode observations, and to store and retrieve memories. In Section 3, we outline how we carried out our experiments and show the results in Section 4. In Section 5, we compare our work with other works, and show how ours differs from theirs. Finally, in Section 6, we conclude this paper and mention some possible future research directions.

2 Methodology

To perform our experiments, we have to design an agent which mimics the memory systems. Besides, we need an environment in which the agent operates, as well as define how the agent observes the environment.

2.1 The Room Environment

We model our environment with a discrete-event simulation (DES). In this simulation, there is a room with N_{humans} humans, $N_{objects}$ different types of objects, and $N_{object.locations}$ possible object locations. At each time step, every human can place an object at one of the possible object locations, within their capacity limit. For example, if the capacity of a desk is four, then only four objects can be placed on the desk at once. Each change in the room state is called an event. For example, let us assume that Alice was the only person in the room. If at $timestep = 1$, Alice’s lap-




Figure 2: An episodic memory and a semantic memory represented as a knowledge graph.

top was on the desk, but at $timestep = 2$, it is on her lap, then the event at $timestep = 2$ is the change of the state.

The humans in the simulation do not move other humans’ objects, but only theirs. They also have their own personal fixed routine. For example, Alice has her laptop on the desk two days in a row, and then the next three days she has it on her lap. She repeats this routine. The actual routines of the humans in our experiments are longer and involve more objects and object locations than this example.

The object locations on which each human places their object is not uniformly random. With a given probability $p_{commonsense}$, they place their objects on a commonsense location (e.g., a bowl on the cupboard), and with $1 - p_{commonsense}$ chance, they place them on a non-commonsense location (e.g., a bowl in the wardrobe). As ConceptNet (Speer, Chin, and Havasi 2017) can have multiple commonsense object locations per object, we chose *commonsense location* of an object as the one with the highest weight.

We make our environment compatible with OpenAI-Gym by creating a wrapper around the DES, where an agent only partially observes the entire state. The agent goes around the room and observes one human at a time (e.g., At $timestep = 42$, the agent observes that Alice has her laptop on her lap). After each observation, the agent is asked a question about the current state of the DES (e.g., Where is Alice’s laptop?). If it answers correctly, it gets a reward of +1 and if not, it gets a reward of 0.

The observations made are given as quadruples, $(h, r, t, timestamp)$. Here h , r , and t stand for head, relation, and tail, respectively. This means that the agent observes that the entity h has the relation r to the entity t at the given timestamp (e.g., (Bob’s laptop, AtLocation, desk, 42)). The question is given as a triple $(h, r, ?)$, such as (Bob’s laptop, AtLocation, ?).

2.2 Human-Like Memory Systems

We model our agent with both a short-term and a long-term memory system. The long-term memory is split in two parts, the episodic and semantic memory. Every observation is initially stored in the short-term memory system. When that gets full, the agent must decide what to do with the oldest observation in short-term memory. It can take one of the three actions: 1) *forget* it completely, 2) move it to the *episodic part* of the long-term memory system, or 3) move it to the *semantic part* of the long-term memory system. A diagram of this memory structure can be found in Figure 1.

There are constraints on what type of information can be

Algorithm 1: Memory retrieval

Input: Question q , episodic memory system M_e , and semantic memory system M_s .

Output: Relevant memory m .

```
1: if  $\exists m_e \in M_e : m_e$  is relevant for  $q, \nexists m_{\hat{e}} \in M_e :$   
    $m_{\hat{e}}$  is relevant for  $q, m_{\hat{e}}$  is more recent than  $m_e$  then  
2:    $m = m_e$   
3: else if  $\exists m_s \in M_s : m_s$  is relevant for  $q, \nexists m_{\hat{s}} \in M_s :$   
    $m_{\hat{s}}$  is relevant for  $q, m_{\hat{s}}$  is stronger than  $m_s$  then  
4:    $m = m_s$   
5: else  
6:    $m = null$   
7: end if  
8: return  $m$ 
```

stored in each memory system; we closely follow the cognitive science theories behind them (Tulving 1983, 1985; Tulving and Thomson 1973). Both the short-term and episodic memory have the same quadruple format as the observation.

The facts in the semantic memory are also quadruples $(h, r, t, strength)$. Here, h , r , and t are the same as above, but rather than storing the timestamp, it has a *strength* term, which indicates how strong this memory is. An example could be $(laptop, AtLocation, desk, 3)$, where 3 stands for the number of times that the agent has decided to store this memory in its semantic memory system. This memory system has to do with the general world knowledge, and thus h is not individual-specific (e.g., `laptop` instead of `Bob's laptop`). Figure 2 shows an example of memories that could be stored.

Memory Storage All three memory systems (i.e., short-term, episodic, and semantic) are bounded in size. Each of them is a knowledge graph where the number of quadruples is limited. We use the cognitive science theory to manage the episodic and semantic memory systems. That is, if the episodic memory system is full, we “forget” the oldest memory stored. If the semantic memory system is full, we “forget” the weakest memory.

Dealing with the short-term memory system when it is full is more interesting and challenging. Although the cognitive science theory suggests that the memories in this system will be either moved to the long-term memory systems or forgotten, there is no clear explanation on how this choice is made. Besides, this gets complicated by the ever-changing environment the agent is in. Therefore, we decided to learn the best behavior. We use a deep Q-learning based agent that learns to do this by maximizing its return (Mnih et al. 2013).

Memory Retrieval In our scenario, each time an agent is asked a question, a relevant memory (i.e., quadruple) has to be retrieved from the knowledge graph to answer the question. For example, when the question $(Bob's\ laptop, AtLocation, ?)$ is asked at $timestep = 44$, the episodic memory $(Bob's\ laptop, AtLocation, desk, 42)$ might be relevant. However, when the agent has a more recent episodic memory $(Bob's\ laptop, AtLocation, cupboard, 43)$, it is likely that this

Algorithm 2: Converting a knowledge-graph-based memory system into a learnable knowledge graph embedding (KGE). $\parallel$ is the concatenation operation.

Input: A memory system M , the embedding function emb (implemented as a lookup table)

Output: The embedding of the memory system $KGE(M)$

```
1:  $M_{sorted} \leftarrow$  memory system  $M$  sorted by strength (re-  
   spectively timestamp)  
2:  $KGE(M) \leftarrow []$ , i.e., an empty list  
3: for all  $(h, r, t, value) \in M_{sorted}$  do  
4:   if  $M$  is semantic memory then  
5:     Append  $emb(h) \parallel emb(r) \parallel emb(t)$  to  $KGE(M)$   
6:   else if  $M$  is short-term or episodic memory then  
7:      $\{h_n, h_o\} \leftarrow$  split  $h$  into the human's name and the  
       object  
8:     Append  $(emb(h_n) + emb(h_o)) \parallel emb(r) \parallel emb(t)$   
       to  $KGE(M)$   
9:   end if  
10: end for  
11: return  $KGE(M)$ 
```

memory provides a more accurate answer to the question. Therefore, when there is more than one relevant memory in the episodic memory system, we'll have the agent choose the most recent one.

The episodic memory system alone might not contain the answer to every question. For example, if there are no relevant memories in the episodic memory system to answer the above question, the agent can look up its semantic memory system where it might find something like $(laptop, AtLocation, cupboard, 2)$. Using this general world knowledge, the agent might be able to answer the question. However, there might also be another relevant memory in the system such as $(laptop, AtLocation, desk, 3)$. Since this memory is stronger than the other, it is likely that this memory provides the correct answer to the question. These rules form an algorithm for memory retrieval with pseudocode presented in Algorithm 1.

2.3 The Deep Q-learning Reinforcement Learning Agent

As mentioned in Section 2.2, we learn the task of choosing what to do with the oldest short-term memory with an RL approach. That is, the memory storage policy is learned by RL, while the memory retrieval policy is not learned but done by Algorithm 1. We have chosen deep Q-learning as our RL algorithm, because of its proven track record on other problems.

We define the state s as (M_o, M_e, M_s) , where the elements stand for the short-term, episodic, and semantic memory systems, respectively. The size of the discrete action space $|\mathcal{A}|$ is three: forget the oldest short-term memory in M_o completely, move it to the episodic memory M_e , or move it to the semantic memory M_s . Every time the agent takes an action, it also answers a given question, and a cor-




Figure 3: The Q-network diagram, where the short-term M_o , episodic M_e , and semantic M_s memory systems are given as the initial input. The LSTM blocks output the last hidden states of the LSTMs. The final output is the state-action values of the Q function.

responding reward is given immediately. This reward is 1 if the answer is correct, 0 otherwise.

For every entity and relation in each of the memory systems (i.e., knowledge graphs), we maintain a learnable embedding vector in the lookup table emb . Now, in order to use a memory system as an input for the Q-network, we need to create an embedding (i.e., a numeric representation) for the complete memory system itself. We do this by ordering the quadruples of the memory system by strength for semantic memories and timestamp for short-term and episodic. The embedding of the memory system is then the sequence of embeddings of quadruple in the same order. To get the embedding of a quadruple, we retrieve and concatenate the embeddings of its head, relation, and tail. Algorithm 2 outlines the procedure. Note that for the short-term and episodic, we also compute embeddings for the human actors, so that the agent can learn their individual characteristics. Figure 4 gives an example.

There are many sophisticated neural networks. For our Q-network we use an LSTM-based (Hochreiter and Schmidhuber 1997) neural network because this has shown to work well in many cases. The Q-network is written out in Equation 1 and the diagram of it is visualized in Figure 3. The output dimension of MLP_{all} is three, and each of them represents the state-action value (Q-value) of the Q function.

$$\begin{aligned}
 \text{Q-network} = & \\
 & \text{MLP}_{all} \left(\text{MLP}_o \left(\text{LSTM}_o \left(\text{KGE}_o(M_o) \right) \right) \parallel \right. \\
 & \text{MLP}_e \left(\text{LSTM}_e \left(\text{KGE}_e(M_e) \right) \right) \parallel \\
 & \left. \text{MLP}_s \left(\text{LSTM}_s \left(\text{KGE}_s(M_s) \right) \right) \right)
 \end{aligned} \tag{1}$$

where $\parallel$ denotes the vector concatenation operator.




Figure 4: An example of converting an episodic memory system into a KGE representation.

3 Experimental Setup

3.1 The Environment and Agent Parameters

In our experimental environment, we set $N_{humans} = 64$, $N_{objects} = 16$, and $N_{object.locations} = 28$. We determined these values to be sufficiently challenging. As for $p_{commonsense}$, we set it as 0.5, to avoid bias by giving equal weights to commonsense and non-commonsense probabilities, and thus avoiding bias.

The environment terminates after 128 steps, meaning that the agent will make 128 observations and be asked a total of 128 questions. This makes the minimum and maximum total episode rewards 0 and 128, respectively.

We further simplify the agent's observation, $(h, r, t, timestamp)$ because we only use AtLocation for the relation r . Since this would be the same for all the memory systems, we set the relation embedding to 0.

The agent observes the humans in a fixed order (e.g., Alice $\rightarrow$ Bob $\rightarrow$ Alice $\rightarrow$ Bob), this ensures that it observes all of them roughly equally often. The questions about object locations that the agent has observed are sampled uniformly at random. This means that an agent with infinite and perfect memory would be able to always answer these questions perfectly.

We train RL agents with different memory capacities. The short-term memory capacity is fixed at 1, while we try episodic and semantic capacity from 1, 2, 4, 8, 16, and to 32. For example, the agent with the biggest memory capacity is $|M_o| = 1$, $|M_e| = 32$, and $|M_s| = 32$, where M_o , M_e , and M_s stand for the short-term, episodic, and semantic memory systems, respectively. In the experiments, we give an equal capacity for the episodic and semantic memory systems. We will refer to the agent's total memory capacity




Figure 5: Training, validation, and test results of the agents with the memory capacity of 32. All of them were run five times with different random seeds. The standard deviation is shown together with the average. As for training, instead of total rewards per episode, we report loss over time since the varying epsilon value affects the rewards. The maximum total rewards per episode are 128, since in total of 128 questions are asked per episode.

as $|\mathcal{M}_e| + |\mathcal{M}_s|$.

We train two variants of our agents: one variant starts with an empty semantic memory, while the other variant starts with general world knowledge from ConceptNet in its semantic memory system. The agent with its semantic memory system prefilled can be considered as a “pretrained” agent. With these variants, we want to investigate whether agents that already know about the world would learn differently from those that start from scratch.

At test time, the policy of a trained agent given a state is to choose the action that has the highest Q-value (i.e., the one the agent considers to be most “valuable”). We compare our two trained agents with the three baselines. The first one always moves the short-term memory contents to episodic memory. The second, on the other hand, always chooses to move to semantic memory. The third one is a random baseline that chooses one of the three actions (i.e., including dropping the memory) uniform at random.

3.2 The Deep Q-learning and Other Training Parameters

Our training objective follows that of the original deep Q-learning: to minimize the temporal difference loss (See Equation 2). Some of the few differences are that we use the Huber loss (Huber 1964), instead of the squared error loss, and that we use Adam (Kingma and Ba 2014) as an optimizer. We found these to achieve slightly better performance.

$$L(\theta) = \text{HuberLoss}(Q_{\text{target}}, Q(s, a; \theta)) \quad (2)$$

$$Q_{\text{target}} = r + \gamma \max_{a'} Q(s', a'; \theta^{\text{target}})$$

where s, a, s', a', r , and γ denote state, action, next state, next action, reward, and discount factor, respectively. θ^{target} is the Q-network parameters from k steps before.

We use PyTorch Lighting as our deep learning framework (Falcon 2019). See Appendix A for the complete list of the hyperparameters.

4 Results

4.1 Deep Q-learning Performance by Agent

We know that the agent observes 16 different kinds of objects in the environment. Therefore, we first look at the agents with a memory capacity of 32. Ideally, these agents will use their semantic memory capacity of 16 to learn the general world knowledge and the other 16 to remember individual-specific episodic memories. The episodic-only and semantic-only baseline agents also have a memory capacity of 32, which is either all episodic or all semantic memory. To account for the stochasticity of the questions that the environment asks, after each training epoch (i.e., episode) we run it on a validation environment. After 16 epochs of training, we choose the model that has the highest total rewards per episode on validation. This model is then run again on a test environment to report the final numbers. We ran each experiment 5 times. Figure 5 shows the results.

We observe that the RL agents with episodic and semantic memory systems outperform all the three baselines. It is also interesting to see the performance gap between the two RL agents. The one with its semantic memory system pre-filled not only learns faster, but also converges to a better optimum (116 vs. 109 for average total test rewards). This is comparable to transfer learning in neural networks (Pan and Yang 2010), except that the transferred knowledge in our setting




Figure 6: Average total test rewards of the agents with different memory capacities. All of them were run five times with different random seeds. The error bars indicate the standard deviation. The maximum total rewards per episode are 128, since in total of 128 questions are asked per episode. The detailed numbers can be found in Appendix B.

is a symbolic knowledge graph, not neural network parameters.

Next, we investigate the varying size of the memory capacity. That is, we want to see how the agents perform if it can store less or more than 32 memory triples, while all other parameters are kept the same. Figure 6 shows the results. As the memory capacity grows, the episodic-only agent can answer more questions, resulting in higher total rewards. If it has high enough memory capacity (e.g., 64), it can indeed answer all 128 questions. As for the agent with only semantic memory system, it performs better than the others in the low memory capacity settings. This suggests that it is better to focus only on the general world knowledge rather than remembering individuals, when the memory capacity is too limited. However, this semantic-only agent’s performance plateaus after a certain memory capacity. For example, the total test rewards for the memory capacity of 32 and 64 are both about 92. This is because 32 general world knowledge triples were already enough to answer the general knowledge related questions, and having more memories has only a minimal impact.

4.2 Understanding What the Agent Learns

We display the Q-values during the test time in Figure 7. We can see that the semantic-scratch agent starts both of its episodic and semantic memory systems from scratch, while the semantic memory system of the semantic-pretrained agent is already prefilled with the general world knowledge from the beginning. There is one same strategy between the two agents. Both of them think it is most “valuable” to store a not-so-common short-term memory (e.g., (Frank’s train, AtLocation, circus, 86)) in the episodic memory system. This is understandable, since such an observation is more individual-specific than generalizable.

But we observe one fundamental difference between them. Whenever the semantic-scratch agent observes something that is generalizable (e.g., (Ann’s bowl,


AtLocation, cupboard, 2)), it stores it in the semantic memory system. However, the semantic-prefilled agent never decides to store any short-term memory in the semantic memory system. It always forgets generalizable short-term memories completely. This makes its semantic memory system untouched. As it is filled in the beginning, it stays the same until the end. What we observe is that it retains the general world knowledge it was given. Then, it either stores short-term memory in the episodic one or forgets it completely. In our experiments, this strategy indeed results in receiving higher rewards, since we prefilled its semantic memory system with the perfect world knowledge. As for the semantic-scratch agent, although it manages to learn most of the general world knowledge, it also makes mistakes. For example, it thinks that (train, AtLocation, kitchen) is generalizable, while actually it should have been (train, AtLocation, zoo). This kind of mistakes results in answering some questions incorrectly, leading to lower total rewards per episode.

5 Related Work


Our work stands somewhere between cognitive architecture and machine learning with memory systems. Cognitive architectures focus on the theory of the human mind, while machine learning with memory systems try to take advantage of memory systems to solve problems such as partial observability and catastrophic forgetting. Here we list the works that are the most relevant to us.

ACT-R (Anderson 2007) and Soar (Laird 2012) study the human memory systems comprehensively and provide opportunities for AI agents. However, they lack in computational methods, such as training an RL agent with memory systems.

Gorski and Laird (Gorski and Laird 2011) have applied RL within the Soar cognitive architecture. While our work focuses on learning a memory storage policy, they have focused on learning episodic memory retrieval. Li (Li 2020)



(a) The agent with its semantic memory from scratch.



(b) The agent with its semantic memory prefilled.

Figure 7: The Q-values of each of the three possible actions during the test time. The short-term, episodic, and semantic memory systems at time step of 2 and 86 are shown together. Figure 7(a) is the agent that starts with its semantic memory from scratch, while Figure 7(b) is the one where its semantic memory is prefilled with the general world knowledge.

has extended their work by applying it to bigger data.

The Tensor Brain (Tresp et al. 2021) focuses on the computational theory of an agent’s perception, episodic memory, semantic memory. They introduced the model named bilayer tensor network (BTN) where perception, episodic memory, and semantic memory are realized. They used the VRD dataset (Lu et al. 2016) to carry out their experiments. Although this dataset includes image modality, it is similar to ours in that it is a form of triples that can form a knowledge graph. In their experiments, they used supervised learning to predict the relationships between the entities.

Hemmer and Steyvers (Hemmer and Steyvers 2009) have studied how episodic and semantic memory systems play a role in recall of objects, but their experiments were human-based empirical results, which does not scale as well as our computational method.

Memory-based RL (Meng, Gorbet, and Kulić 2021; Lin

et al. 2018) adds memory components to their classical RL problems to address the partial observability problem and to achieve sample efficiency, respectively. Episodic Memory Deep Q-Networks (Lin et al. 2018) uses the term “episodic memory” that is used to regularize their deep Q-network. They use the term episodic memory to refer to the state-action pairs that the agent has taken along the training episode.

Episodic Memory Reader (Han et al. 2019) trains an RL agent that learns what to remember in their memory system, and they also use question-answering to evaluate their method. Their work differs from ours in that they have only one memory system and that it is not made of a knowledge graph of triples, but rather numeric embeddings, which are harder to interpret.

When a deep neural network trained to solve task A is re-trained to solve task B, it often forgets what it has learned previously, which is called *catastrophic forgetting*. Continual learning (also known as lifelong learning) is concerned with training a neural network on a sequence of tasks without forgetting them. Continual learning with an episodic memory system (Guo et al. 2020; Lopez-Paz and Ranzato 2017) addresses this issue by storing some previous task examples in their episodic memory system. Adjusting the neural network weights for the current task also accounts for the previous task examples, to avoid forgetting them.

Similar to our work, CLS-ER (Arani, Sarfraz, and Zonooz 2022) also takes advantage of the dual memory system (i.e., episodic and semantic) for their continual learning problem. Their episodic memory system stores data samples, while the semantic stores previous model weights.

6 Conclusions

Current machine learning agents with memory systems cannot answer all types of questions. They are especially weak at answering personal questions. To tackle this problem, we investigated an agent with human-like memory systems, inspired by the cognitive science theory. Our agent has short-term, episodic, and semantic memory systems, each of which is modeled with a knowledge graph.

To experiment with this agent, we have designed and released our own environment, “the Room”, compatible with OpenAI Gym. We successfully trained an agent that has interpretable memory systems, with a deep Q-learning based reinforcement learning algorithm. This experiment confirms that a machine with human-like memory systems outperforms the ones without this structure, when both general world knowledge and individual-specific questions are asked.

In the future, we would like to make our environment not only more complex (e.g., include more relations in its triples), but also multimodal (e.g., images). The questions given by the environment can extend to natural language with ambiguity, not just triples. We would also like to add other kinds of human-like memory systems to our agent (e.g., working memory, implicit memory, etc.). It would also be interesting to include real humans in the loop, so that the agent can ask them a question when it is uncertain of the actions that it can take.

Maximum number of epochs (episodes)	16
Batch size	1024
Epoch length	1024×128
Replay size	1024×128
Warm start size	1024×128
Epsilon start value	1.0
Epsilon end value	0
Epsilon last step	128×16
Gamma (discount factor)	0.65
Learning rate	0.001
DQN sync rate	10
Loss function	Huber
Optimizer	Adam
Neural network architecture	LSTM + MLP
Embedding dimension	32
LSTM hidden size dimension	64
Number of LSTM layers	2
Processing unit	1 CPU
Floating point precision	32 bits
Number of iterations for evaluation	10
Number of runs	5

Table 1: The full list of the hyperparameters used.

A Hyperparameters

Table 1 lists all the hyperparameters used. Since every episode terminates after 128 steps (questions) and we are training for 16 epochs, the total number of training steps is $128 \times 16 = 2048$. One epoch corresponds to one episode (i.e., 128 steps), where every training step includes computing one temporal difference loss from one batch of 1024 data samples. And then Adam (Kingma and Ba 2014) makes one weight update.

We set epoch length, replay size, and warm start size as 1024×128 . This corresponds to the total number of data samples that the optimizer sees in one epoch. The epsilon value of DQN decays linearly. It starts with 1.0 and ends with 0. The time step when it reaches 0 is 128×16 , which corresponds to the last (16th) epoch. We tried different values for gamma (discount factor), and we found 0.65 to be the best. Q_{target} is synced with Q_{θ} every 10 steps.

As for the Q-network, the embedding dimension was set to 32, and the LSTM hidden size dimension was set to 64. The input feature size of the linear layers was also set to 64, so that last hidden state of the LSTM can be directly fed into them. ReLU (Nair and Hinton 2010) was used for the non-linearity between the linear layers in the MLP. The number of LSTM layers used is 2.

Since our neural network is relatively small (i.e., the number of parameters is 265,000), we just used one CPU (i.e., Intel® Core™ i7-10875H) with 64 GB of RAM, instead of using GPUs or ASICs. It took about 3 hours to train one semantic-scratch agent with the memory capacity of 32.

For every validation and test, we ran the model 10 times to get the average of total rewards per episode. We did this because the environment is stochastic, and we wanted to get the best estimate. This means that the final reported total re-

Memory capacity		2	4
Agent memory type	8	16	32
	64	128	256
	512	1024	2048
	4096	8192	16384
	32768	65536	131072
<i>Episodic-only</i>	17.4 ± 1.0	23.6 ± 2.7	
<i>Semantic-only</i>	22.4 ± 4.5	36.8 ± 3.0	
<i>Both, random</i>	15.4 ± 2.6	25.0 ± 3.0	
<i>Both, RL, semantic-scratch</i>	23.9 ± 1.4	37.6 ± 1.2	
<i>Both, RL, semantic-pretrained</i>	25.5 ± 0.6	38.2 ± 2.5	
	35.0 ± 2.3	56.8 ± 4.1	90.6 ± 5.0
	61.6 ± 4.6	84.2 ± 3.3	91.6 ± 4.4
	39.4 ± 3.8	54.6 ± 4.5	76.0 ± 12.5
	60.4 ± 3.1	85.3 ± 3.9	109.1 ± 4.0
	59.2 ± 4.0	86.8 ± 2.3	116.5 ± 0.9

Table 2: The average total test rewards per episode for all capacities. The standard deviations are shown after $\pm$.

wards of the test environment are the average of 10×5 numbers.

B Total Test Rewards for All Capacities

Table 2 shows the average total test rewards and their standard deviations per episode for all capacities.

Acknowledgments

This research was (partially) funded by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organization for Scientific Research, <https://www.hybrid-intelligence-centre.nl/>.

References

- Anderson, J. R. 2007. *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press. ISBN 9780195324259.
- Arani, E.; Sarfraz, F.; and Zonooz, B. 2022. Learning Fast, Learning Slow: A General Continual Learning Method based on Complementary Learning System. In *International Conference on Learning Representations*.
- Arulkumaran, K.; Deisenroth, M.; Brundage, M.; and Bharath, A. 2017. A Brief Survey of Deep Reinforcement Learning. *IEEE Signal Processing Magazine*, 34.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. Cite arxiv:1606.01540.
- Falcon, W. 2019. PyTorch Lightning: The lightweight PyTorch wrapper for high-performance AI research. *GitHub*.
- Gorski, N. A.; and Laird, J. E. 2011. Learning to use episodic memory. *Cognitive Systems Research*, 12(2): 144–153. The 9th International Conference on Cognitive Modeling. Manchester, UK, July 2009.
- Guo, Y.; Liu, M.; Yang, T.; and Rosing, T. 2020. Improved Schemes for Episodic Memory-based Lifelong Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 1023–1035. Curran Associates, Inc.

- Han, M.; Kang, M.; Jung, H.; and Hwang, S. J. 2019. Episodic Memory Reader: Learning What to Remember for Question Answering from Streaming Data. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4407–4417. Florence, Italy: Association for Computational Linguistics.
- Hemmer, P.; and Steyvers, M. 2009. Integrating Episodic and Semantic Information in Memory for Natural Scenes. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation*, 9(8): 1735–1780.
- Hogan, A.; Blomqvist, E.; Cochez, M.; D’amato, C.; Melo, G. D.; Gutierrez, C.; Kirrane, S.; Gayo, J. E. L.; Navigli, R.; Neumaier, S.; Ngomo, A.-C. N.; Polleres, A.; Rashid, S. M.; Rula, A.; Schmelzeisen, L.; Sequeda, J.; Staab, S.; and Zimmermann, A. 2022. Knowledge Graphs. *ACM Computing Surveys*, 54(4): 1–37.
- Huber, P. J. 1964. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1): 73 – 101.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Laird, J. E. 2012. *The Soar Cognitive Architecture*. The MIT Press. ISBN 0262122960.
- Li, J. 2020. Integrating Declarative Long-Term Memory Retrievals into Reinforcement Learning. In *Eighth Annual Conference on Advances in Cognitive Systems*.
- Lin, Z.; Zhao, T.; Yang, G.; and Zhang, L. 2018. Episodic Memory Deep Q-Networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI’18*, 2433–2439. AAAI Press. ISBN 9780999241127.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient Episodic Memory for Continual Learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, 6470–6479. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Lu, C.; Krishna, R.; Bernstein, M.; and Fei-Fei, L. 2016. Visual Relationship Detection with Language Priors. In *European Conference on Computer Vision*.
- Meng, L.; Gorbet, R.; and Kulić, D. 2021. Memory-Based Deep Reinforcement Learning for POMDPs. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5619–5626. IEEE Press.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. A. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR*, abs/1312.5602.
- Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *ICML 2010*, 807–814.
- Pan, S. J.; and Yang, Q. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10): 1345–1359.
- Speer, R.; Chin, J.; and Havasi, C. 2017. ConceptNet 5.5: An Open Multilingual Graph of General Knowledge. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1).
- Tresp, V.; Sharifzadeh, S.; Li, H.; Konopatzki, D.; and Ma, Y. 2021. The Tensor Brain: A Unified Theory of Perception, Memory and Semantic Decoding. *CoRR*, abs/2109.13392.
- Tulving, E. 1983. *Elements of Episodic Memory*. Oxford University Press.
- Tulving, E. 1985. Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26(1): 1.
- Tulving, E.; and Thomson, D. M. 1973. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80: 352–373.
- Vrandečić, D.; and Krötzsch, M. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM*, 57(10): 78–85.