

A generative memory approach to enable lifelong reinforcement learning

12 March 2019, by Ingrid Fadelli



Credit: Raghavan, Hostetler & Chai.

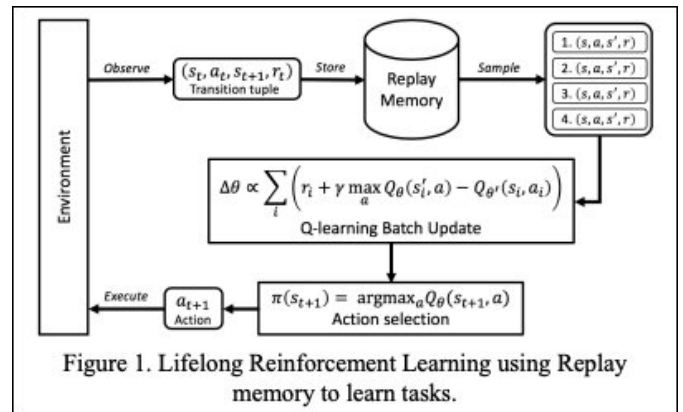
A key limitation of existing artificial intelligence (AI) systems is that they are unable to tackle tasks for which they have not been trained. In fact, even when they are retrained, the majority these systems are prone to 'catastrophic forgetting,' which essentially means that a new item can disrupt their previously acquired knowledge.

For instance, if a model is initially trained to complete [task A](#) and then subsequently retrained on task B, its performance on task A could decline considerably. A naïve solution would be to infinitely add more neural layers to support additional tasks or items being trained, but such an approach would not be efficient, or even functionally scalable.

Researchers at [SRI international](#) have recently tried to apply biological [memory](#) transfer mechanisms to AI systems, as they believe that this could enhance their performance and make them more adaptive. Their study, [pre-published on arXiv](#), draws inspiration from mechanisms of memory transfer in humans, such as long-term and [short-term memory](#).

"We are building a new generation of AI systems that can learn from experiences," Sek Chai, a co-

PI of the [DARPA Lifelong Learning Machines \(L2M\)](#) project, told TechXplore. "This means that they can adapt to new scenarios based their experiences. Today, AI systems fail because they are not adaptive. The DARPA L2M project, led by Dr. Hava Siegelmann, seeks to achieve paradigm-changing advances in AI capabilities."



Credit: Raghavan, Hostetler & Chai.

Memory transfer entails a complex sequence of dynamic processes, which allow humans to easily access salient or relevant memories when thinking, planning, creating or make predictions about future events. Sleep is thought to play a critical role in the consolidation of memories, particularly REM sleep, the stage in which dreaming occurs most frequently.

In their study, Chai and his SRI colleagues developed a generative memory mechanism that can be used to train AI systems in a pseudo-rehearsal manner. Using replay and [reinforcement learning](#) (RL), this mechanism allows AI systems to learn from salient memories throughout their lifetime, and scale with a large number of training tasks or items. The generative memory approach

developed by Chai and his colleagues uses an encoding method to separate the latent space. This allows an AI system to learn even when tasks are not well-defined or when the number of tasks is unknown.

"Our AI system does not directly store raw data, such as video, audio, etc.," Chai explained. "Rather, we use generative memory to generate or imagine what it has experienced previously. Generative AI systems have been used to create art, music, etc. In our research, we use them to encode generative experiences that can be used later with reinforcement learning. Such an approach is inspired by biological mechanisms in sleep and dreams, where we recall or imagine fragments of experiences that are reinforced in our long-term memories."

In the future, the new generative memory approach introduced by Chai and his colleagues could help to address the issue of catastrophic forgetting in neural network-based models, enabling lifelong learning in AI systems. The researchers are now testing their approach on computer-based strategy games that are commonly employed to train and evaluate AI systems.

"We are using real-time strategy games such as StarCraft2 to train and study our AI agents on lifelong learning metrics such as adaptation, robustness, and safety," Chai said. "Our AI agents are trained with surprises injected into the game (e.g. terrain and unit capability change)."

More information: Generative memory for lifelong reinforcement learning. arXiv:1902.08349 [cs.LG]. arxiv.org/abs/1902.08349

www.sri.com/work/projects/arti...m-continually-learns

www.darpa.mil/program/lifelong-learning-machines

© 2019 Science X Network

APA citation: A generative memory approach to enable lifelong reinforcement learning (2019, March 12) retrieved 23 March 2019 from <https://techxplore.com/news/2019-03-memory-approach-enable-lifelong.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.