



2020 Annual International Conference on Brain-Inspired Cognitive Architectures for Artificial Intelligence: Eleventh Annual Meeting of the BICA Society

## Towards sample-efficient policy learning with DAC-ML

Ismael T. Freire<sup>a,b,c,1,\*</sup>, Adrián F. Amil<sup>a,b,c,1</sup>, Vasiliki Vouloutsi<sup>a,c</sup>,  
Paul F.M.J. Verschure<sup>a,b,c,d</sup>

<sup>a</sup>*Institute for Bioengineering of Catalonia (IBEC), Av. d'Eduard Maristany 10-14, 08930 Barcelona, Spain*

<sup>b</sup>*Universitat Pompeu Fabra (UPF), Barcelona, Spain*

<sup>c</sup>*Barcelona Institute of Science and Technology (BIST), Barcelona, Spain*

<sup>d</sup>*Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain*

---

### Abstract

The sample-inefficiency problem in Artificial Intelligence refers to the inability of current Deep Reinforcement Learning models to optimize action policies within a small number of episodes. Recent studies have tried to overcome this limitation by adding memory systems and architectural biases to improve learning speed, such as in Episodic Reinforcement Learning. However, despite achieving incremental improvements, their performance is still not comparable to how humans learn behavioral policies. In this paper, we capitalize on the design principles of the Distributed Adaptive Control (DAC) theory of mind and brain to build a novel cognitive architecture (DAC-ML) that, by incorporating a hippocampus-inspired sequential memory system, can rapidly converge to effective action policies that maximize reward acquisition in a challenging foraging task.

© 2021 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the 2020 Annual International Conference on Brain-Inspired Cognitive Architectures for Artificial Intelligence: Eleventh Annual Meeting of the BICA Society

*Keywords:* cognitive architecture; sample-inefficiency problem; sequence learning; reinforcement learning; distributed adaptive control

---

---

<sup>1</sup> Both authors have contributed equally to this work.

\* Corresponding author. Tel.: +34 93 401 19 18.

E-mail address: [ifreire@ibecbarcelona.eu](mailto:ifreire@ibecbarcelona.eu)

## 1. Introduction

With the advent of Deep Reinforcement Learning (DRL), the last decade has been marked by several historical landmarks in the field of Artificial Intelligence (AI), in terms of both scientific and societal impact. Reaching human-level performance in historically considered uniquely-human intellectual domains (such as chess [1], Go [2], and real-time strategy video games [3]) has led to renewed claims about its resemblance to human cognition. Some researchers have even proposed that DRL could be a valid metaphor to understand the brain [4]. However, criticisms to this view have pointed out to at least one great fundamental difference between human learning and deep learning approaches: sample efficiency [5].

Sample efficiency refers to the number of data samples required for a learning system to attain a target level of performance (see [6] for a more detailed discussion on the topic). Current DRL models are sample inefficient because they require large amounts of training episodes to achieve human-level performance. An example of this limitation is the research on emergent tool-use in a hide-and-seek multi-agent game; where DRL agents took up to 380 million episodes to learn some of the more sophisticated behaviors [7].

Recent research in the field has sought to find solutions to the sample-inefficiency problem by introducing algorithms that make use of a memory system to bootstrap policy learning, such as in Episodic Reinforcement Learning (ERL) [8]. In ERL, stored episodes are retrieved from past interactions with the environment to support decision-making when the agent encounters similar states. However, states are treated as being independent of one another. Therefore, ERL cannot capitalize on the sequential nature of experience that gives rise to highly informative state-transition statistics.

In this paper, we present a novel cognitive architecture well suited for state-of-the-art Machine Learning (ML) environments, whereby policy learning and action selection are based on a sequential memory system inspired by the human hippocampus [9]. In particular, building upon the architectural principles of the Distributed Adaptive Control (DAC) theory of mind and brain, we derive a biologically inspired architecture (DAC-ML) that leverages the continuous nature of experience to bootstrap policy learning in a sample-efficient manner. We demonstrate the learning efficiency of DAC-ML in a challenging foraging task and show how the sequential feature of the memory system is key for achieving such fast-learning dynamics.

## 2. Cognitive architecture

DAC-ML is a biologically inspired cognitive architecture composed of three layers (Reactive, Adaptive, and Contextual, see Fig. 1). It is built based on the architectural principles of DAC of layered and hierarchical organization, that predicate that cognition is hierarchically distributed in different layers of control [10–12].

The Reactive layer represents the agent's prewired reflexes and provides predefined responses that allow for simple but robust interaction with the environment. In DAC-ML, the Reactive layer incorporates a pseudo-random-walk algorithm that drives the agent to explore the environment and facilitates the input sampling initially required by the learning in the upper layers of the architecture (see Fig. 1, bottom panel).

The Adaptive Layer, where the perceptual learning takes place, contains a convolutional autoencoder that uses the sensory input sampled by the Reactive layer to build compressed representations of the state space (see Fig. 1, middle panel). These state representations are then linked with the agent's action at every step to build the state-action couplets to be conveyed to the Contextual layer for policy learning. Moreover, the transmission of the couplets to the Contextual layer depends on the compressed representations surpassing a quality threshold, assessed by the convolutional autoencoder's corresponding reconstruction error.

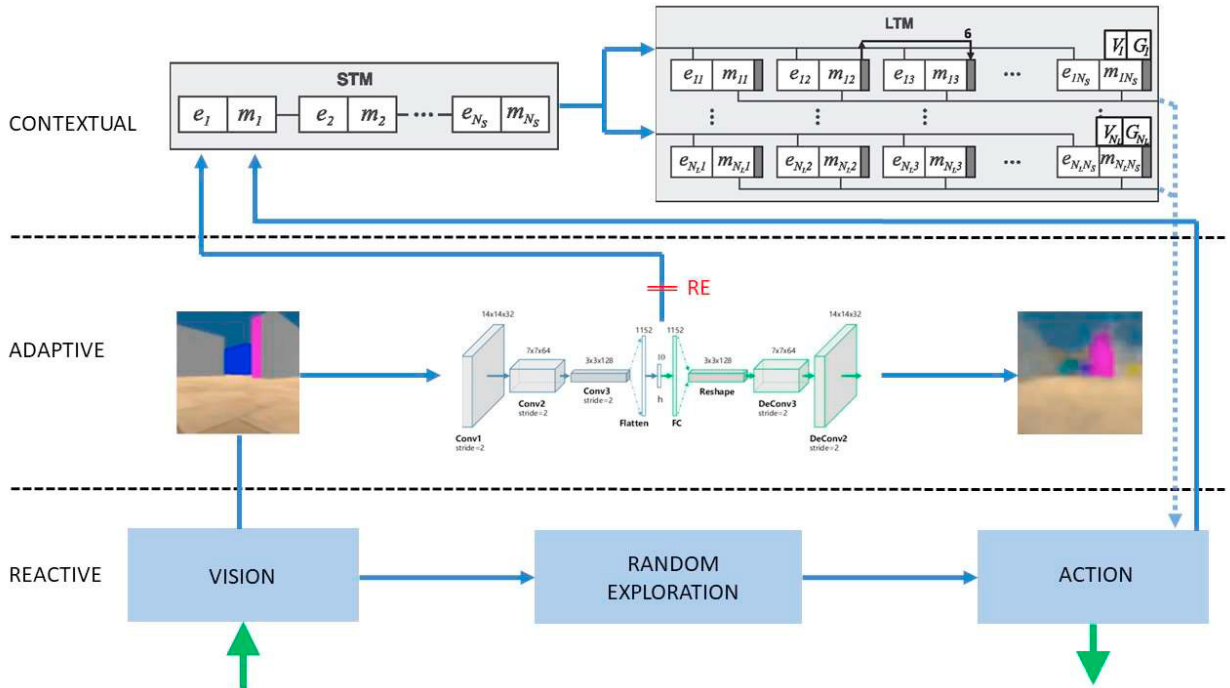


Fig. 1. (a) Representation of the three layers of the DAC-ML architecture: Reactive (bottom), Adaptive (center), and Contextual (top). The Reactive layer implements a pseudo-random-walk generator that allows an efficient random exploration of the environment. The Adaptive layer comprises a convolutional autoencoder that generates compressed representations of the visual input to produce the state-action couples to be conveyed to the STM buffer in the Contextual layer. The Contextual layer incorporates the STM and LTM buffers, along with an action selection algorithm, that allows the agent to perform policy learning and decision-making based on the stored episodic sequences. RE: reconstruction error; e: compressed state vectors; m: action; V: reward value; G: episodic sequence.

The Contextual layer uses stored sequential episodic memories to perform policy learning and to guide real-time decision-making. It is composed of short-term memory (STM) and long-term memory (LTM) buffers, and an action selection algorithm (see Fig. 1, top panel). The STM transiently stores the most recent sequence of state-action couples generated by the agent's interaction with the environment, and it is updated following a first-in, first-out (FIFO) rule. Once a goal is encountered, the current sequence maintained in the STM is conveyed and stored in the LTM with its associated reward value. In the current implementation, the STM has a fixed length of 50 state-action couples, while the capacity of LTM is limited to 100 sequences.

### 2.1. Action selection with a sequential memory bias for rapid policy learning.

The Contextual layer's action selection algorithm, while being invariant in itself, selects an action based on the recent history of observed states and performed actions. Concretely, at every time step, the observed compressed state is compared, based on a similarity metric (i.e., Euclidean distance), with all the state representations stored in LTM. The similarity scores are then weighted by an additional memory system storing the trigger values, which keep track of the recent history of selected state-action couples (as in [13]). This procedure increases the eligibility of the couples that come immediately after the recently selected ones. Hence, the resulting eligibility scores (i.e., similarity scores weighted by the trigger values) effectively bias the action selection towards continuing on recently visited sequences in LTM.

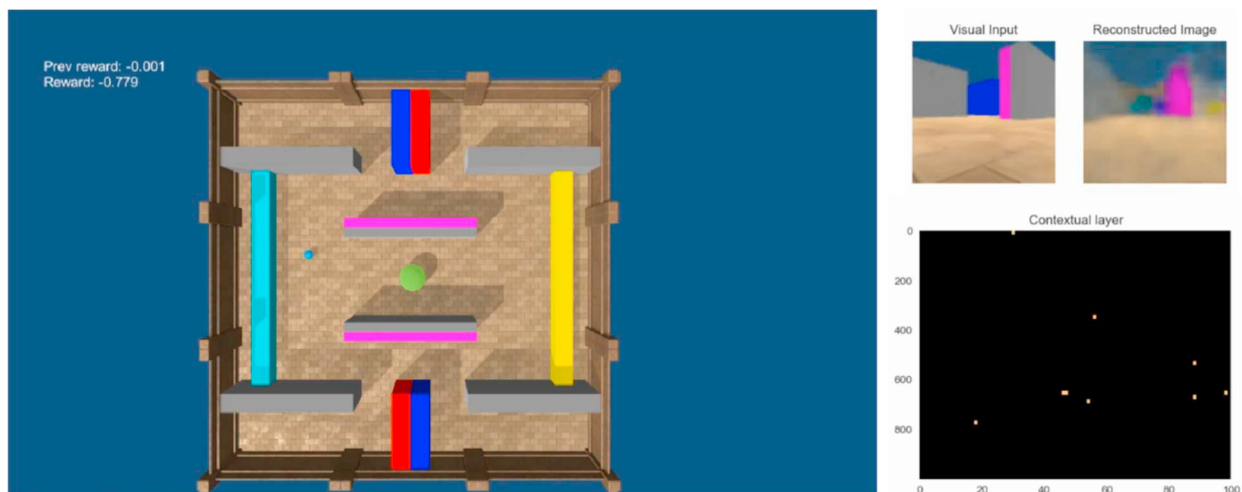


Fig. 2. Experimental setup. Left panel: a top-view of the foraging task modeled in Animal-AI. The small blue sphere represents the agent. The big green sphere at the center represents the reward. Top right panels: first-person view of the agent (left) and the reconstructed image generated by the Adaptive layer (right). Bottom right panel: representation of the selected couplets of the LTM in the Contextual layer.

Furthermore, a set of state-action couplets in LTM is selected based on their eligibility scores surpassing both absolute and proportional thresholds. This procedure enforces a winner-takes-all mechanism that has been shown to operate in the hippocampus [14]. The selected set of state-action couplets are then assigned the reward values associated with their corresponding sequences, weighted by an exponential decay based on each couplet's normalized distance to the end of the sequence. This mechanism favors the selection of actions that are closer in time to bigger rewards. Moreover, the reward values assigned to each couplet are normalized to the maximum reward associated with the selected couplets, thus implementing a relative reward valuing function. Lastly, the final action to be performed is randomly drawn from the probability distribution generated over the discrete action space by normalizing the sums of all the relative reward values associated with the actions of each selected couplet. Then, the trigger values are updated based on the selected couplets so that the couplets selection (and thus action selection) in the next time step is biased towards following the same sequences. This biasing mechanism is further reinforced by behavioral feedback [12], whereby the input sampling caused by the selected action favors the storage of similar sequences of state-action couplets. At the same time, our implementation of the action selection as a probability distribution over the discrete action space keeps the Bayesian optimality feature of previous DAC models operating in a continuous action space [11].

### 3. Results

We used the Animal-AI framework [15] to create the simulation environment for testing the capabilities of the DAC-ML cognitive architecture. Animal-AI is an AI experimentation and evaluation platform specially designed to allow for cognitive testing in reinforcement learning and spatial navigation tasks. This framework imposes several realistic constraints regarding the type of input an agent can receive since the agents have to learn to navigate the 3D environment only guided by first-person observations in the form of raw pixel data.

To test the efficacy of DAC-ML in obtaining rewards by acquiring stable behavioral sequences, we tested it in a foraging task that requires the sequential encoding and retrieval of adequate visual cues in order to reach the reward location in an optimal amount of time. The environment is a colored maze (see Fig. 2) with only one reward source in the center. The walls are painted with different colors to serve as visual cues for the agent. At the beginning of each episode, the agent is randomly allocated in one of the four corners of the maze and must reach the center to

acquire the reward. The reward value is initially set to 3, and it decays linearly with time until the end of the episode. An episode finishes either when the agent reaches the reward or when the time limit is reached.

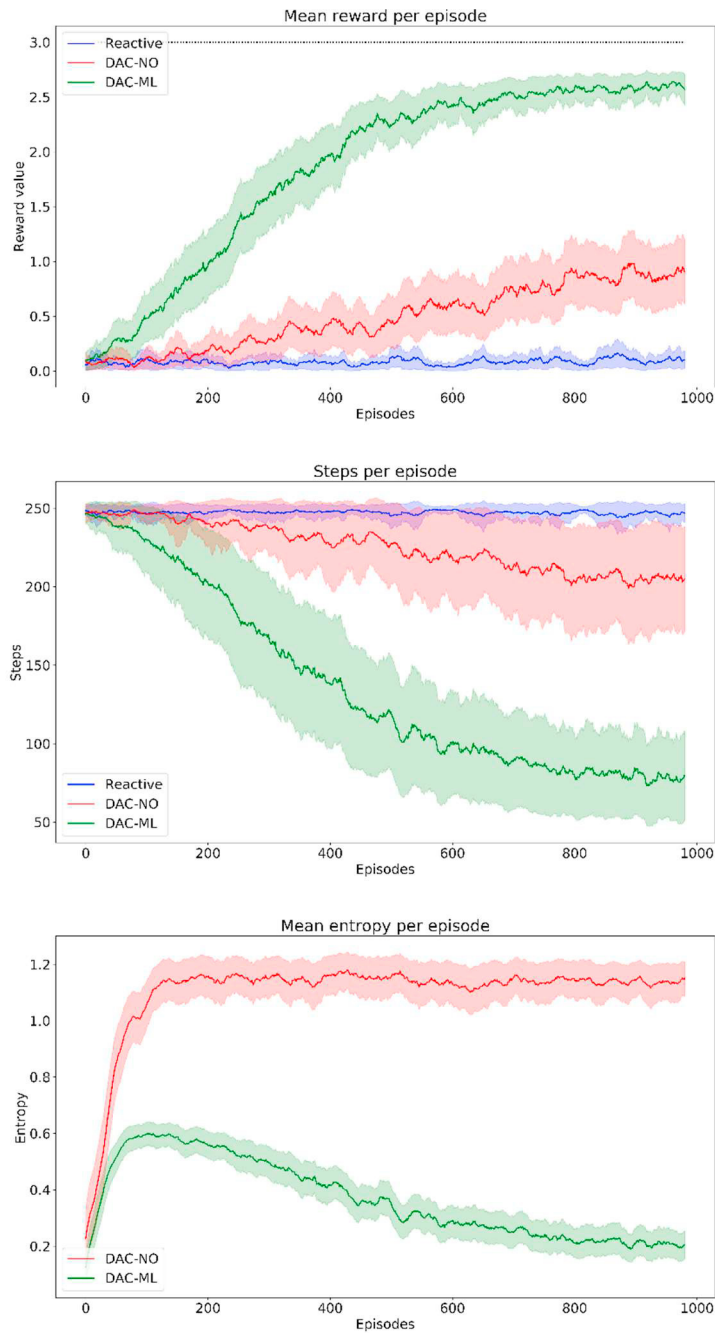


Fig. 3. DAC-ML rapidly stabilizes and maximizes reward acquisition in a foraging task. Top panel: mean reward per episode scores of DAC-ML (green), DAC-ML with no sequencing bias (red), and a reactive agent (blue). Center panel: mean steps per episode of the same agents. Bottom panel: mean entropy per episode for both DAC-ML agents. Entropy is computed over the probability distribution of actions at every step so that lower values indicate policy convergence and stability. All metrics have been computed by an averaged sliding window of 20 episodes over time.

We ran 20 simulations of 1.000 episodes each and compared the results of DAC-ML against two control agents: (1) a Reactive agent that randomly explores the environment, and (2) a version of DAC-ML with the same set of parameters but without incorporating the sequential memory bias, and thus treating the state-action couplets as independent units. The results show the capacity of DAC-ML to efficiently acquire rewards in the foraging task in relatively few episodes, as compared with the control agents. Fig. 3-top shows the difference in reward acquisition between the three agents (DAC-ML, DAC-ML without the sequential memory bias, and the Reactive agent). We show that DAC-ML is able to maximize the reward acquisition in less than 1.000 episodes. In addition, Fig. 3-center shows the number of steps (time) taken by the agents in every episode to reach the reward. These results demonstrate that DAC-ML can optimize the time it requires to reach the reward and that it generates rapid learning dynamics compared to a more standard cognitive architecture that does not incorporate a memory system based on episodic sequences. Finally, we also estimated the stability over time of the action selection (i.e., policy learning) by computing the mean entropy across the probability distributions over the action space throughout each episode. In the bottom panel of Fig. 3, we observe how the policy entropy also steadily decreases over time for DAC-ML (but not for the control agent), indicating convergence towards a stable action policy. Overall, these results demonstrate the advantage of incorporating a sequential memory bias in achieving rapid and stable policy learning in a cognitive architecture optimized by reinforcement learning.

#### 4. Discussion and Future Work

In this paper, we have presented DAC-ML, a novel cognitive architecture based on the organizational principles of the Distributed Adaptive Control theory, that can efficiently maximize reward acquisition in a challenging foraging task. DAC-ML is a memory-guided decision-making architecture inspired by the hippocampus, that stores complete episodic sequences of state-action couplets to guide action selection. We showed that this biologically inspired architectural bias promotes the rapid generation of stable and efficient goal-directed behavior, thus affording sample-efficient policy learning.

In contrast with classical RL where the policy is directly updated by the learning algorithm, in DAC-ML, policy learning is achieved through the interaction of the different components of its layered architecture. Since action-selection in DAC-ML follows a fixed rule, the main driver of policy learning is the constant acquisition and retrieval of new sequential memories. In other words, the behavioral policy of a DAC-ML agent is implicitly updated by the memories it forms in its interaction with the environment. Moreover, DAC-ML emphasizes the learning of complete behavioral sequences, whereas ERL usually treats its memory units as separate disconnected events. The incorporation of this hippocampus-inspired architectural bias bootstraps the formation of a behavioral feedback loop [12], allowing DAC-ML to rapidly transition from an initial exploration phase to a more exploitative phase in which the behavior of the agent stabilizes over time while maximizing reward acquisition.

With this work, we seek to contribute to the available proposed solutions to the sample-inefficiency problem in reinforcement learning by taking inspiration from empirical and theoretical research from the cognitive sciences. In future work, we plan to extend the results presented in this paper by directly comparing DAC-ML with state-of-the-art ERL models, and test it in multi-agent environments, such as in game-theoretic dilemmas [16,17].

#### Acknowledgements

This research received funding from the H2020-EU project HR-Recycler, ID: 820742.

#### References

- [1] Silver D, Hubert T, Schrittwieser J, Antonoglou I, Lai M, Guez A, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* (80- ) 2018;362:1140–4. <https://doi.org/10.1126/science.aar6404>.
- [2] Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. *Nature* 2016;529:484–9. <https://doi.org/10.1038/nature16961>.

- [3] Vinyals O, Babuschkin I, Czarnecki WM, Mathieu M, Dudzik A, Chung J, et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 2019;575:350–4.
- [4] Marblestone AH, Wayne G, Kording KP. Toward an integration of deep learning and neuroscience. *Front Comput Neurosci* 2016;10:94.
- [5] Marcus G. Deep Learning: A Critical Appraisal. *ArXiv E-Prints* 2018.
- [6] Botvinick M, Ritter S, Wang JX, Kurth-Nelson Z, Blundell C, Hassabis D. Reinforcement learning, fast and slow. *Trends Cogn Sci* 2019;23:408–22.
- [7] Baker B, Kanitscheider I, Markov T, Wu Y, Powell G, McGrew B, et al. Emergent Tool Use From Multi-Agent Autocurricula. *ArXiv E-Prints* 2019.
- [8] Pritzel A, Uria B, Srinivasan S, Badia AP, Vinyals O, Hassabis D, et al. Neural Episodic Control. *Proc. 34th Int. Conf. Mach. Learn., JMLR.org*; 2017, p. 2827–36.
- [9] Buzsaki G, Tingley D. Space and time: The hippocampus as a sequence generator. *Trends Cogn Sci* 2018;22:853–69.
- [10] Verschure PFMJ, Pennartz CMA, Pezzulo G. The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philos Trans R Soc B Biol Sci* 2014;369:20130483. <https://doi.org/10.1098/rstb.2013.0483>.
- [11] Verschure P. A real-world rational agent: unifying old and new AI. *Cogn Sci* 2003;27:561–90.
- [12] Verschure PFMJ, Voegtlin T, Douglas RJ. Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature* 2003;425:620–4.
- [13] Marcos E, Ringwald M, Duff A, Sánchez-Fibla M, Verschure PFMJ. The Hierarchical Accumulation of Knowledge in the Distributed Adaptive Control Architecture. *Comput. Robot. Model. Hierarchical Organ. Behav., Berlin, Heidelberg: Springer Berlin Heidelberg*; 2013, p. 213–34. [https://doi.org/10.1007/978-3-642-39875-9\\_10](https://doi.org/10.1007/978-3-642-39875-9_10).
- [14] de Almeida L, Idiart M, Lisman JE. A second function of gamma frequency oscillations: an E%-max winner-take-all mechanism selects which cells fire. *J Neurosci* 2009;29:7497–503.
- [15] Crosby M, Beyret B, Halina M. The animal-ai olympics. *Nat Mach Intell* 2019;1:257.
- [16] Freire IT, Moulin-Frier C, Sanchez-Fibla M, Arsiwalla XD, Verschure PFMJ. Modeling the formation of social conventions from embodied real-time interactions. *PLoS One* 2020;15:e0234434.
- [17] Freire IT, Arsiwalla XD, Puigbò J-Y, Verschure P. Modeling theory of Mind in multi-agent games using adaptive feedback control. *ArXiv E-Prints* 2019.