Neuroprospecting with DeepRL agents

Satpreet H. Singh
University of Washington (Seattle)
satsingh@uw.edu

Abstract

A virtuous cycle between neuroscience and deep reinforcement learning is emerging and the AI community can do much to enable and accelerate it.

Introduction: There is a rich history of cross-pollination between the technological pursuit of developing artificial intelligence (AI) and the scientific endeavor to understand biological intelligence [1, 2, 3, 4]. Deep Neural Networks (DNNs), which are now ubiquitous in AI systems, were initially inspired by the brain. Coming full circle, DNNs now form the basis of the leading models of neuronal systems (e.g. perceptual, cognitive, motor control) and have eclipsed traditional mathematical models in their ability to explain experimental data [5]. While these studies have predominantly used supervised learning to train their networks, a small but steadily growing set of recent studies investigating complex naturalistic behavior and high-level cognitive functions such as meta-learning, decision-making and control [6, 7, 8, 9, 10, 11] have instead used used deep reinforcement learning (DRL) [12, 13] to train their neural network models in the *agent* setting. Recent perspectives, written for an audience with a background in neurobiology, discuss the opportunities that this alternative framework offers to neuroscience [14, 15]. Here we describe some of the current technological and algorithmic challenges in this emerging niche that AI researchers could help address, and also highlight some potential opportunities for cross-pollination with AI.

Design and analysis of DeepRL agents: Like many other scientific domains in recent years, neuroscience has embraced the use of DNNs to learn models that mimic the function of a specific neuronal system. This trend was accelerated by the availability of cheaper computing power, large labeled training datasets (such as in computer vision or NLP) and user-friendly deep-learning frameworks. Similarly, recent neuroscientific studies using DRL are benefiting from the ongoing development and standardization of DRL training frameworks [16] and physics simulators [17, 18, 19]. However, much work remains to be done to realize the true potential of this new paradigm.

As summarized in Figure 1a, we believe that there are three dimensions along which the currently AI-focused DRL ecosystem can be developed further in the aid of such neuroscientific studies. First, biophysical complexity can be incorporated into DRL agents across several modeling levels. Unlike current studies in decision making that assume action spaces that directly imply decisions, more realistic motor-control could emerge by embodying agents in a biomechanical body model [20, 21, 22, 23]. Artificial neural networks (ANNs) are coarse approximations of real biological circuits that do not account for information stored in individual neuronal spikes or on the timescale of spike timing differences. Using spiking neural networks, along with emerging techniques for reward-based training [24, 25, 26] could lead to more biologically plausible models. Furthermore, constraining network wiring using connectomes from model organisms [27, 28, 29] could provide a fine-grained correspondence between biological circuits and trained ANNs.

Second, tasks and simulators tailored to neuroscience could produce agents with more biologically plausible neural computations and richer behavioral repertoires. Unlike most current studies that focus on one sensory modality, providing agents with multiple senses in flexible or procedural high-fidelity 3D virtual-reality environments [30, 31] could further capture further biological complexity and simulate more realistic perception. Similarly, while current studies train networks on individual tasks [32, 33, 34, 35], training agents on multiple tasks [36, 21] or using auxiliary objectives that might not

NeurIPS 2021 AI for Science Workshop.

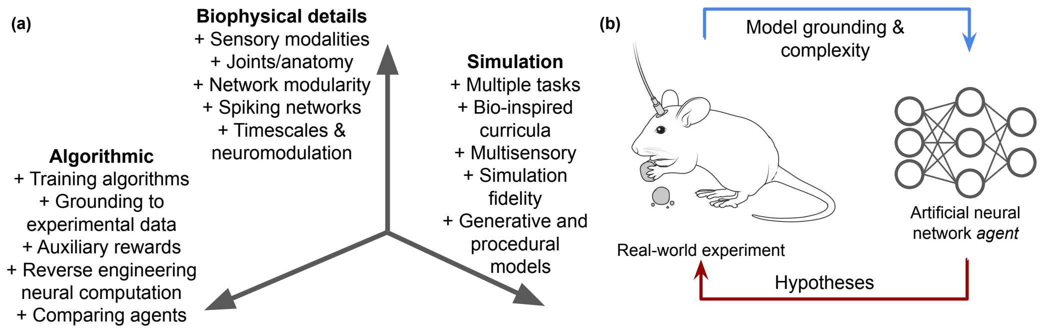

Figure 1: (a) Dimensions along which DRL frameworks and agents can be improved to support neuroscience. (b) Closing the loop between real-world experiments and artificial agent models (Open source rodent experiment clipart credit: [61]).

be aligned with the agent's primary task (such as minimizing metabolic cost [37, 34]), could produce agents with not just a richer repertoire of behaviors, but also more complex neural activity structures with shared task-representations, shared task-dynamics and task-specific adaptations [38, 39, 40, 41].

Third, the core algorithms and strategies for training agents could also be adapted to the goals of this scientific niche. Learning algorithms that respect biological constraints (e.g. excitation-inhibition balance, Dale's law) have received attention in the supervised-learning setting [42, 43] but not in the reinforcement-learning setting. Other directions for methodological innovation in agent training include using meta-algorithms like evolution [44], exploiting environmental symmetries or other regularities in training [45], using offline methods to directly ground agent behavior in recorded animal behavior [46], using algorithms constrained by morphology [47], and using complex biologically-inspired training curricula [48, 49, 50].

Neuroscientists train task-solving DNNs to generate hypotheses for how neural systems implementing distributed computation solve problems in perception, cognition or other psychological capabilities [2]. Trained task-solving ANNs and agents often exhibit behavioral features or neural dynamics that resemble their biological counterparts. This makes them useful as tools for generating hypotheses of behavior and neural computation in conditions where neural recording, or experiments might be challenging or expensive. Such hypotheses could guide future real-world experiments (Figure 1b). See [51, 52] for striking recent examples of such artificial-network inspired neuroscience experiments. Moreover, the process of designing an optimization objective for training such networks provides insights into systems at an abstract level that descriptive approaches typically cannot [53, 54]. Currently popular methods for neural network understanding [55, 56] are focused on comparing the responses of feedforward networks to a library of stimuli, and, by design, do not take into account the order or amount of stimuli previously seen by the network. Though alternative approaches developed for use with RNNs do take stimulus history into consideration [57, 58, 59], they do not account for the network being deployed in a closed-loop agent setting where the network's current state determines it's next action, which in turn affects the next observation seen by it. New theoretical development is required for comparing RNN-based agents, specifically to deal with small differences in how agents react to stimuli, that result in compounding differences over behavior trajectories [60].

Neuroprospecting: The well-established practice of searching in nature for molecules, microorganisms, plants, and other species containing biochemical or genetic information of commercial value is known as bioprospecting. In a similar spirit, we use the term *neuroprospecting* to refer to the search for mechanisms, architectures and algorithms implemented in neurobiological systems that could inspire new approaches in AI. As artificial agent designs and analyses become increasingly sophisticated, more are more opportunities arise for extracting biological *middleware* that could inspire advances in artificial intelligence. This could potentially allow artificial systems to behave more like biological creatures, that are known to be born with sophisticated innate capabilities [62], and be more vastly more robust and swift to adapt in complex real-world environments [63, 64]. A virtuous cycle between real-world experiments and increasingly complex yet biologically grounded artificial models could benefit both the neuroscience and machine intelligence communities [65, 66, 4, 67].

References

[1] Surya Ganguli. The intertwined quest for understanding biological intelligence and creating artificial intelligence. https://hai.stanford.edu/news/intertwined-quest-understanding-biological-intelligence-and-creating-artificial-intelligence, 2018.

[2] Timothy T Rogers and James L McClelland. Parallel distributed processing at 25: Further explorations in the microstructure of cognition. *Cognitive science*, 38(6):1024–1077, 2014.

[3] Adam H Marblestone, Greg Wayne, and Konrad P Kording. Toward an integration of deep learning and neuroscience. *Frontiers in computational neuroscience*, 10:94, 2016.

[4] Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017.

[5] Nikolaus Kriegeskorte. Deep Neural Networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, 1:417–446, 2015.

[6] Nicolas Heess, Greg Wayne, Yuval Tassa, Timothy Lillicrap, Martin Riedmiller, and David Silver. Learning and transfer of modulated locomotor controllers. *arXiv preprint arXiv:1610.05182*, 2016.

[7] Ari Weinstein and Matthew M Botvinick. Structure learning in motor control: A deep reinforcement learning model. *arXiv preprint arXiv:1706.06827*, 2017.

[8] Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer L Hicks, Serge Levine, Christopher Atkeson, and Scot Delp. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *bioRxiv*, 2020.

[9] H Francis Song, Guangyu R Yang, and Xiao-Jing Wang. Reward-based training of recurrent neural networks for cognitive and value-based tasks. *Elife*, 6:e21492, 2017.

[10] Jane X Wang, Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Demis Hassabis, and Matthew Botvinick. Prefrontal cortex as a meta-reinforcement learning system. *Nature neuroscience*, 21(6):860–868, 2018.

[11] Mathew Botvinick, Sam Ritter, Jane X Wang, Zeb Kurth-Nelson, Charles Blundell, and Demis Hassabis. Reinforcement learning, fast and slow. *Trends in Cognitive Sciences*, 2019.

[12] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.

[13] Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.

[14] Matthew Botvinick, Jane X Wang, Will Dabney, Kevin J Miller, and Zeb Kurth-Nelson. Deep reinforcement learning and its neuroscientific implications. *Neuron*, 2020.

[15] Samuel J Gershman and Bence P Ölveczky. The neurobiology of deep reinforcement learning. *Current Biology*, 30(11):R629–R632, 2020.

[16] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. Stable baselines. https://github.com/hill-a/stable-baselines, 2018.

[17] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.

[18] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

[19] Erwin Coumans. Bullet physics simulation. In *ACM SIGGRAPH 2015 Courses*, page 1. 2015.

[20] Victor Lobato Ríos, Pembe Gizem Özdil, Shravan Tata Ramalingasetty, Jonathan Arreguit, Stéphanie Clerc Rosset, Graham Knott, Auke Jan Ijspeert, and Pavan Ramdya. Neuromechfly, a neuromechanical model of adult drosophila melanogaster. *bioRxiv*, 2021.

[21] Josh Merel, Diego Aldarondo, Jesse Marshall, Yuval Tassa, Greg Wayne, and Bence Ölveczky. Deep neuroethology of a virtual rodent. *arXiv preprint arXiv:1911.09451*, 2019.

[22] Florian Fischer, Miroslav Bachinski, Markus Klar, Arthur Fleig, and Jörg Müller. Reinforcement learning control of a biomechanical model of the upper extremity. *Scientific Reports*, 11(1):1–15, 2021.

[23] Łukasz Kidziński, Sharada P Mohanty, Carmichael F Ong, Jennifer L Hicks, Sean F Carroll, Sergey Levine, Marcel Salathé, and Scott L Delp. Learning to run challenge: Synthesizing physiologically accurate motion using deep reinforcement learning. In *The NIPS'17 Competition: Building Intelligent Systems*, pages 101–120. Springer, 2018.

[24] Shuncheng Jia, Tielin Zhang, Xiang Cheng, Hongxing Liu, and Bo Xu. Neuronal-plasticity and reward-propagation improved recurrent spiking neural networks. *Frontiers in Neuroscience*, 15:205, 2021.

[25] Mengwen Yuan, Xi Wu, Rui Yan, and Huajin Tang. Reinforcement learning in spiking neural networks with stochastic and deterministic synapses. *Neural computation*, 31(12):2368–2389, 2019.

[26] Charlotte Frenkel, David Bol, and Giacomo Indiveri. Bottom-up and top-down neural processing systems design: Neuromorphic intelligence as the convergence of natural and artificial intelligence. *arXiv preprint arXiv:2106.01288*, 2021.

[27] Louis K Scheffer, C Shan Xu, Michal Januszewski, Zhiyuan Lu, Shin-ya Takemura, Kenneth J Hayworth, Gary B Huang, Kazunori Shinomiya, Jeremy Maitlin-Shepard, Stuart Berg, et al. A connectome and analysis of the adult drosophila central brain. *Elife*, 9:e57443, 2020.

[28] Steven J Cook, Travis A Jarrell, Christopher A Brittin, Yi Wang, Adam E Bloniarz, Maksim A Yakovlev, Ken CQ Nguyen, Leo T-H Tang, Emily A Bayer, Janet S Duerr, et al. Whole-animal connectomes of both caenorhabditis elegans sexes. *Nature*, 571(7763):63–71, 2019.

[29] Alexandros Goulas, Fabrizio Damicelli, and Claus C. Hilgetag. Bio-instantiated recurrent neural networks: Integrating neurobiology-based network topology in artificial networks. *Neural Networks*, 2021.

[30] Matthew Crosby, Benjamin Beyret, and Marta Halina. The Animal-AI olympics. *Nature Machine Intelligence*, 1(5):257–257, 2019.

[31] Matthew Crosby. Building thinking machines by solving animal cognition tasks. *Minds and Machines*, pages 1–27, 2020.

[32] Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474):78–84, 2013.

[33] Alexander JE Kell, Daniel LK Yamins, Erica N Shook, Sam V Norman-Haignere, and Josh H McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.

[34] Christopher J Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In *International Conference on Learning Representations*, 2018.

[35] Christopher J Cueva, Peter Y Wang, Matthew Chin, and Xue-Xin Wei. Emergence of functional and structural properties of the head direction system by optimization of recurrent neural networks. In *International Conference on Learning Representations*, 2019.

[36] Michael Crawshaw. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.

[37] Abdullahi Ali, Nasir Ahmad, Elgar de Groot, Marcel AJ van Gerven, and Tim C Kietzmann. Predictive coding is a consequence of energy efficiency in recurrent neural networks. *bioRxiv*, 2021.

[38] Christian David Marton, Guillaume Lajoie, and Kanaka Rajan. Efficient and robust multi-task learning in the brain with modular task primitives. *arXiv preprint arXiv:2105.14108*, 2021.

[39] Guangyu Robert Yang, Madhura R Joglekar, H Francis Song, William T Newsome, and Xiao-Jing Wang. Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2):297–306, 2019.

[40] Lea Duncker, Laura Driscoll, Krishna V Shenoy, Maneesh Sahani, and David Sussillo. Organizing recurrent network dynamics by task-computation to enable continual learning. *Advances in Neural Information Processing Systems*, 33, 2020.

[41] Wiktor F Młynarski and Ann M Hermundstad. Efficient and adaptive sensory codes. *Nature Neuroscience*, pages 1–12, 2021.

[42] Charles B Delahunt, Jeffrey A Riffell, and J Nathan Kutz. Biological mechanisms for learning: a computational model of olfactory learning in the manduca sexta moth, with applications to neural nets. *Frontiers in computational neuroscience*, 12:102, 2018.

[43] Daniel B Ehrlich, Jasmine T Stone, David Brandfonbrener, Alexander Atanasov, and John D Murray. Psychrnn: An accessible and flexible python package for training recurrent neural network models on cognitive tasks. *Eneuro*, 8(1), 2021.

[44] Kenneth O Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. Designing neural networks through neuroevolution. *Nature Machine Intelligence*, 1(1):24–35, 2019.

[45] Elise van der Pol, Daniel Worrall, Herke van Hoof, Frans Oliehoek, and Max Welling. Mdp homomorphic networks: Group symmetries in reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.

[46] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

[47] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *arXiv preprint arXiv:2102.02202*, 2021.

[48] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.

[49] Banafsheh Rafiee, Sina Ghiassian, Raksha Kumaraswamy, Richard Sutton, Elliot Ludvig, and Adam White. Prediction problems inspired by animal learning. *arXiv e-prints*, pages arXiv–2011, 2020.

[50] Paul Cisek. Resynthesizing behavior through phylogenetic refinement. *Attention, Perception, & Psychophysics*, 81(7):2265–2287, 2019.

[51] Adam S Lowet, Qiao Zheng, Sara Matias, Jan Drugowitsch, and Naoshige Uchida. Distributional reinforcement learning in the brain. *Trends in Neurosciences*, 2020.

[52] Misha B Ahrens. Zebrafish neuroscience: Using artificial neural networks to help understand brains. *Current Biology*, 29(21):R1138–R1140, 2019.

[53] Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.

[54] Paul Cisek. Beyond the computer metaphor: Behaviour as interaction. *Journal of Consciousness Studies*, 6(11-12):125–142, 1999.

[55] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *NIPS*, 2017.

[56] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pages 3519–3529. PMLR, 2019.

[57] David Sussillo and Omri Barak. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural computation*, 25(3):626–649, 2013.

[58] Niru Maheswaranathan, Alex Williams, Matthew Golub, Surya Ganguli, and David Sussillo. Universality and individuality in neural dynamics across large populations of recurrent networks. In *Advances in neural information processing systems*, pages 15629–15641, 2019.

[59] Niru Maheswaranathan, Alex Williams, Matthew Golub, Surya Ganguli, and David Sussillo. Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics. In *Advances in Neural Information Processing Systems*, pages 15696–15705, 2019.

[60] Erika Puiutta and Eric MSP Veith. Explainable reinforcement learning: A survey. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 77–95. Springer, 2020.

[61] Ethan Tyler and Lex Kravitz. mouse eating fibre, July 2020.

[62] Anthony M Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature communications*, 10(1):1–7, 2019.

[63] Sam Kriegman. Why virtual creatures matter. *Nature Machine Intelligence*, 1(10):492–492, 2019.

[64] Edgar Bermudez-Contreras, Benjamin J Clark, and Aaron Wilber. The neuroscience of spatial navigation and the relationship to artificial intelligence. *Frontiers in Computational Neuroscience*, 2020.

[65] Shashank Srikant and Una-May O'Reilly. Can cognitive neuroscience inform neuro-symbolic inference models? In *Is Neuro-Symbolic SOTA still a myth for Natural Language Inference? The first workshop*, 2021.

[66] Edward S Boyden and Adam H Marblestone. Architecting discovery: A model for how engineers can help invent tools for neuroscience. *Neuron*, 102(3):523–525, 2019.

[67] Scott W Linderman and Samuel J Gershman. Using computational theory to constrain statistical models of neural data. *Current opinion in neurobiology*, 46:14–24, 2017.