

A model-free deep reinforcement learning approach to tackle neural control problems

July 31 2018, by Ingrid Fadelli



Left: example of an adjacency matrix with approximate block-diagonal structure. Assuming a linear mixture model of neuronal interactions, this network structure will induce an approximately block diagonal covariance of similar structure. Right: the principal components associated with the adjacency matrix on the left. Credit: Mitchell & Petzold

Brian Mitchell and Linda Petzold, two researchers at the University of California, have recently applied model-free deep reinforcement learning to models of neural dynamics, achieving very promising results.

Reinforcement learning is an area of machine learning inspired by behaviorist psychology that trains algorithms to effectively complete particular tasks, using a system based on reward and punishment. A



prominent milestone in this area has been the development of the Deep-Q-Network (DQN), which was initially used to train a computer to play Atari games.

Model-free reinforcement learning has been applied to a variety of problems, but DQN is generally not used. The primary reason for this is that DQN can propose a limited number of actions, while physical problems generally require a method that can propose a continuum of actions.

While reading existing literature on <u>neural control</u>, Mitchell and Petzold noticed the widespread use of a classical paradigm for solving neural control problems with machine learning strategies. First, the engineer and experimenter agree on the objective and design of their study. Then, the latter runs the experiment and collects data, which will later be analyzed by the engineer and used to build a model of the system of interest. Finally, the engineer develops a controller for the model and the device implements this controller.





Results of the experiment controlling oscillation in the phase space defined by a single principal component. The first plot from the top is a plot of the input into the actuated cell over time; the second plot from the top is a plot of the spikes of the entire network, where different colors correspond to different cells; the third plot from the top corresponds to the membrane potential of each cell over time; the fourth from the top plot shows the target oscillation; the bottom plot shows the observed oscillation. The policy, despite delivering input to only a single cell, is able to approximately induce the target oscillation in the observed phase space. Credit: Mitchell & Petzold

"This work flow ignores recent advances in model-free control (e.g. AlphaGo AlphaGo Zero), which could make the design of controllers



more efficient," Mitchell told *Tech Xplore*. "In a model-free framework, steps b, c, and d are combined into a single step and no explicit model is ever built. Rather, the model-free system repeatedly interacts with the neural system and learns over time to achieve the desired objective. We wanted to fill this gap to see if model-free control could be used to quickly solve new problems in neural control."

The researchers adapted a model-free reinforcement learning method called "deep deterministic policy gradients" (DDPG) and applied it to models of low-level and high-level neural dynamics. They specifically chose DDPG because it offers a very flexible framework, which does not require the user to model system dynamics.

Recent research has found that model-free methods generally need too much experimentation with the environment, making it harder to apply them to more practical problems. Nonetheless, the researchers found that their model-free approach performed better than current model-based methods and was able to solve more difficult <u>neural dynamics</u> problems, such as the control of trajectories through a latent phase space of an under actuated network of neurons.

"For the problems we considered in this <u>paper</u>, model-free approaches were quite efficient and didn't require much experimentation at all, suggesting that for neural problems, state-of-the-art controllers are more practically useful than people might have thought," said Mitchell.



a



b





Summary results of 10 synchronization experiments. (a) Depicts the mean and standard deviation of the global synchronization, (i.e. q from equation 16), against the number of training periods of the controller. (b) Shows histograms demonstrating the synchronization level of all network oscillators with the reference oscillator (i.e. qi from equation 16). That is, a point on either the blue or green curves demonstrates the probability of having a given value for qi. The blue histogram shows counts before training while the green histogram shows counts after training. The average synchronization with the reference, qi, is much higher than global synchronization, q, which is explained by the fact that synchronization with the reference is easier to induce than global synchronization. Credit: Mitchell & Petzold

Mitchell and Petzold carried out their study as a simulation, hence important practical and safety aspects need to be considered before their method can be introduced within clinical settings. Further research that incorporates models into model-free approaches, or that poses limits to model-free controllers, could help to enhance safety before these methods enter clinical settings.

In future, the researchers also plan to investigate how neural systems adapt to control. Human brains are highly dynamic organs that adapt to their surroundings and change in response to external stimulation. This could cause a competition between the brain and the controller, particularly when their objectives are not aligned.

"In many cases, we want the controller to win and the design of controllers that always win is an important and interesting problem," said Mitchell. "For example, in the case where the tissue being controlled is a diseased region of the brain, this region may have a certain progression that the <u>controller</u> is trying to correct. In many diseases, this progression



may resist treatment (e.g. a tumor adapting to expel chemotherapy is a canonical example), but current model-free approaches don't adapt well to these kinds of changes. Improving model-free controllers to better handle adaptation on the part of the brain is an interesting direction that we're looking into."

The research is published in Scientific Reports.

More information: B. A. Mitchell et al. Control of neural systems at multiple scales using model-free, deep reinforcement learning, *Scientific Reports* (2018). DOI: 10.1038/s41598-018-29134-x

© 2018 Tech Xplore

Citation: A model-free deep reinforcement learning approach to tackle neural control problems (2018, July 31) retrieved 27 April 2024 from <u>https://techxplore.com/news/2018-07-model-free-deep-approach-tackle-neural.html</u>

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.