Curriculum learning and experience replay in a model of a cognitive decision-making task.

Austin Gibson, Jasmine Stone, Ashok Litwin-Kumar
Center for Theoretical Neuroscience, Zuckerman Institute, Columbia University

Introduction: Humans can learn and easily perform a wide variety of cognitive tasks, and recent research (1) establishes recurrent neural networks (RNNs) as a modeling technique in cognitive neuroscience. Both humans and networks can learn faster when trained on tasks in a specific order, a method termed 'curriculum learning' (3). However, when trained on multiple tasks sequentially, networks' performance can decrease on previously learned tasks, a phenomenon termed 'catastrophic forgetting'. Techniques such as replay can alleviate catastrophic forgetting. Here we investigate the effect of training order and replay on network performance of multiple cognitive tasks commonly studied in neuroscience.

Methods: A type of RNN, a long short-term memory (LSTM) network, was trained on two similar cognitive tasks (2) in various conditions. In task 1, the network is tested on which of two stimuli is larger on average over time; task 2 is identical except the network must respond after a delay. Baseline performance was established by training networks separately on each task. Baseline performance was compared with networks trained first on task 1 followed by task 2 or the reverse. Networks trained on both tasks sequentially were trained both with and without replaying instances of the first-learned task during training on the second-learned tasks.

Results: When the network was trained only on task 1, it reached high performance with fewer training iterations than when the network was trained only on task 2 (Fig. 1, green and magenta dashed lines, respectively)**.** Networks trained only on task 2 rarely performed above chance levels. The networks consistently learned and reached high levels of performance when the network was trained on a block of task 1 followed by a block of task 2 (Fig. 1), though replays were included to avoid catastrophic forgetting. The network could achieve similar performance when trained on task 2 first if replays of task 1 were interleaved (Fig. 2).

Figure 1: Task performance when trained on task 1 before task 2, with replays.

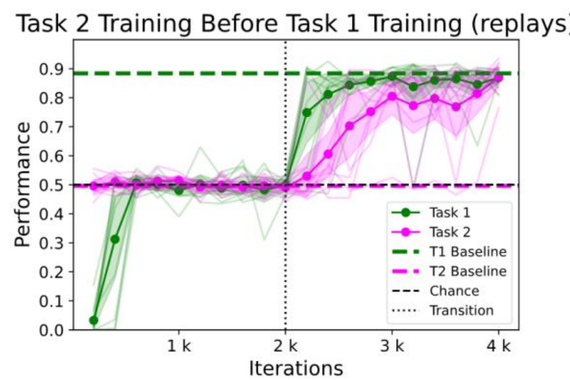

Figure 2: Task performance when trained on task 2 before task 1, with replays.

Conclusion: RNNs can learn multiple cognitive tasks, however some tasks are harder than others for RNNs to learn. In this study, a task with a delay was harder for an RNN to learn than one without a delay. Performance on the delay task was improved via curriculum learning (3): The network's performance was maximized when it was presented with and learned the easier task first (task 1) followed by the more difficult task (task 2). Future studies should investigate the effect of curriculum learning when networks are trained on dissimilar tasks in sequence. Future studies could also investigate performance of a network trained on more than two tasks.

References:

1. Yang GR. et al. Nat Neuro. 2019;22:297-306.
2. Molano-Mazon M. et al. PsyArXiv. 2022.
3. Bengio Y. et al. ICML. 2009: 41-48.