Sander Bohté

Scaling Biologically Plausible Deep Reinforcement Learning

Humans have a remarkable capacity for learning, yet neuronal learning is constrained to locality in time and space and limited feedback.
While neural learning rules have been designed that adhere to these principles and constraints, they exhibit difficulty in scaling to deep networks and complicated datasets. Here, we leverage insights from behavioural science by developing a curriculum that structures how samples are presented to a network to optimise learning. The key features of the curriculum involve progressively introducing new classes to the dataset based on performance metrics, and using a recency bias to protect recently acquired classes. We demonstrate that our curriculum approach makes feedback-based “BrainProp"-style learning robust and more rapid, while substantially improving classification accuracy. We also show the curriculum similarly improves performance for networks trained using error-backpropagation. Our results show the potential of curriculum learning in local learning settings with limited feedback and further bridges the gap between biologically plausible learning rules and error-backpropagation.