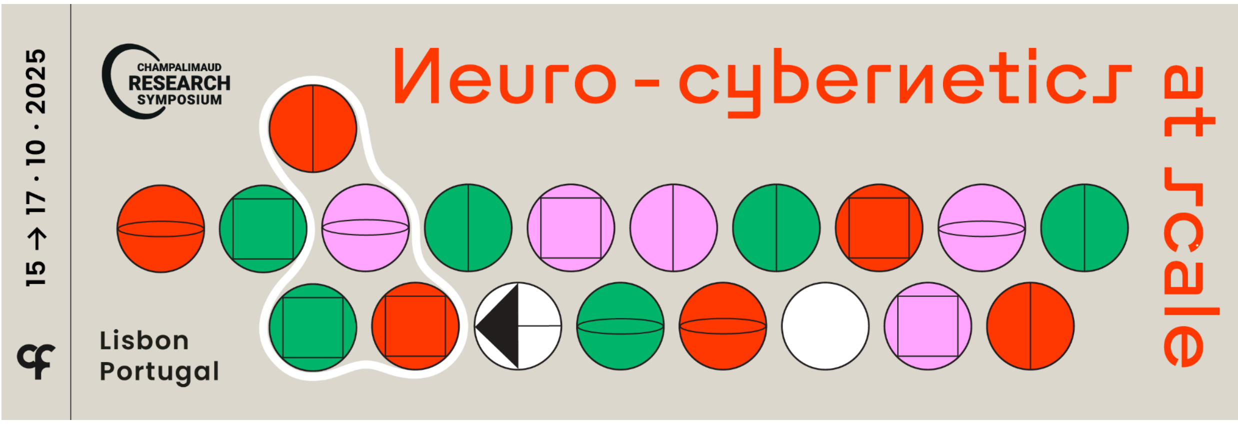

Evolution Imposes an Inductive Bias that Alters and Accelerates Learning Dynamics

Benjamin Midler 1, Alejandro Pan Vazquez 1
1. Princeton Neuroscience Institute, Princeton, USA

The learning dynamics of biological brains and artificial neural networks are of interest to both neuroscience and machine learning. A key difference between them is that neural networks are often trained from a randomly initialized state whereas each brain is the product of generations of evolutionary optimization, yielding innate structures that enable few-shot learning and inbuilt reflexes. Artificial neural networks, by contrast, require non-ethological quantities of training data to attain comparable performance. To investigate the effect of evolutionary optimization on the learning dynamics of neural networks, we combined algorithms simulating natural selection and online learning to form a new method for evolutionarily conditioning artificial neural networks, and applied it to both supervised and reinforcement learning problems. Evolutionarily conditioned networks learned both task types from a few dozen training examples—an order of magnitude less than randomly initialized baselines. Furthermore, the learning dynamics of evolutionarily conditioned networks differed significantly from those of conventionally trained networks, and did not respond to the same mathematical regularities in training data. These results indicate that evolutionary optimization of the population imposes an inductive bias on the learning dynamics of the individual, effectively tuning networks' initial configuration to the learning environment it will occupy.

evolution, learning, neural networks, inductive bias