Mammalian-like visual learning by spatially modulating learning rate in deep belief networks

Michael Lam, Jeff Orchard

Summary: The mammalian visual system has several interesting organizational properties: a retinotopic mapping, centre-surround receptive fields, and orientation selective partitioning in V1. However, unsupervised training of deep belief networks has failed to generate all of these behaviours in a biologically feasible manner. We implemented a biologically motivated variation of these learning algorithms, where nearby neurons more quickly interconnect than distant neurons. Simulations of visual learning in a deep belief network show that this simple learning variant successfully produces retinotopy, as well as centre-surround receptive fields and orientation partitioning typical of the LGN and primary visual cortex, respectively. This study points to a general organizing principle that might play a key role in the development of biological visual systems.

Introduction: The preliminary processing areas of the mammalian visual system have several basic, yet interesting organizational properties: a retinotopic mapping, centre-surround receptive fields in the lateral geniculate nucleus (LGN), and orientation selective partitioning in the primary visual cortex (V1).

Many attempts have been made to model the visual system using neural networks. More recently, a class of neural networks called deep belief networks (DBNs) has been employed for the purpose of studying visual learning. Training methods designed to use these networks, however, have yet to generate all of the basic behaviours in a biologically plausible manner.

Methods: We propose a simple and biologically plausible variation to the standard DBN learning algorithm, motivated by results from developmental biology and neurophysiology. In a standard DBN, each neuron is indiscriminately connected to every neuron in the next layer, regardless of location. To better model the biology of the visual system, we add a proximity constraint in the network's learning process. Neurons that are closer together can become connected more quickly than neurons that are far apart. We implement this concept by modulating the learning rate between pairs of neurons according to their proximity in the retinotopic space. The learning rate for the connection between neurons is multiplied by a Gaussian function that peaks when the neurons are in the same retinotopic location, but falls off asymptotically to zero as the two neurons become retinotopically further apart.

Experiments: We trained neural networks using our proximity-modulated method, as well as the standard learning algorithm (contrastive divergence). We ran 12 network training experiments by applying both learning algorithms to each of six different simulated visual environments, including natural images, spontaneous retinal waves, and an orientation-constrained environment.

Results: After training each network, we analyzed the resulting receptive fields of the various layers. The networks trained using our proximity-modulated method showed centre-surround receptive fields, orientation partitioning, and retinotopy. Networks trained using the standard DBN method were unable to produce these results. Moreover, our results are consistent with observations made on animals reared in an orientation-constrained environment.

Conclusions: Despite its simplicity, proximity-modulated learning yields a spectrum of behaviours remarkably similar to those observed in neurophysiology studies, such as centre-surround receptive fields, and orientation partitioning in V1. Other biologically plausible learning algorithms do not exhibit all of these behaviours. We are currently investigating whether this simple constraint can yield other visual system properties, such as ocular dominance columns and opponent-colour encoding.