

Deep Neural Networks with Biologically Inspired Local Connectivity and Topographic Organization

Author: Filip Rechterík, M.Sc.
Supervisor: Ján Antolík, Ph.D.

Faculty of Mathematics and Physics,
Charles University

Motivation

Can biology inspire a better computer vision model?

A standard architecture in computer vision is the convolutional neural network, or CNN. CNNs organize visual features, such as edges, corners, or even eyes and ears into separate channels (third dimension). Real brains arrange features on two-dimensional sheets and cluster similar features near each other (Figure 1).




Figure 1: Each colour represents a different visual feature—in this case, edge orientation. Similar colours indicate similar orientations.

What if we arrange visual features more like the brain? Even a naive implementation requires custom implementation and performs badly. This thesis develops a way to make it work.

Results

Compared with CNNs of similar size, our model achieves comparable classification accuracy with **54× fewer computational operations (FLOPs)**. See Figure 2 for exact values.

Model	Accuracy	Params	FLOPs
ResNet-18 (large model)	47.1%	11.16 M	2221.0 M
Ours	18.4%	0.10 M	4.1 M
CNN (similar size)	18.2%	0.15 M	221.0 M

Figure 2: Comparison of our model with a CNN of similar size and ResNet-18, which has substantially more parameters. Results on Tiny ImageNet.

Model

Moving away from CNNs significantly increases the number of weights. Each neuron has to remember all of its connections in a weight matrix.

We replace this weight matrix with a low-dimensional descriptor of the neuron's visual feature. A neural network then generates the weight between two neurons from their descriptors (Figure 3). Several such layers are stacked to form the full architecture (Figure 4).




Figure 3: Illustration of one *LocalLinear* layer. Instead of training 225 parameters per neuron, we only train 2, plus one shared neural network. During a forward pass, the weight matrix gets generated.




Figure 4: The architecture of our model. *LocalLinear* denotes the layer described in Figure 3.

Contributions

- A novel layer, *LocalLinear*, allowing arbitrary-size sheet connection and two-dimensional representations. Custom CUDA implementation **reducing memory requirements by 99.8%**.
- A novel architecture utilizing **weight-generating hypernetworks** to reduce the number of parameters, achieving comparable accuracy with substantially fewer floating-point operations.
- Emergent **cortical-like clustering** of similar visual features.