

E9 205 Machine Learning for Signal Processing

Understanding Deep Networks

08-11-2019

Instructor - Sriram Ganapathy (sriramg@iisc.ac.in)




Understanding Deep Networks

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS VOL. 23, NO. 1, JANUARY 2017

Visualizing the Hidden Activity of Artificial Neural Networks

Paulo E. Rauber, Samuel G. Fadel, Alexandre X. Falcão, and Alexandru C. Telea

Understanding Deep Networks



Understanding Deep Networks

SVHN dataset



Understanding Deep Networks

CIFAR-10





Table 1. Test Set Accuracies for our Two Architectures

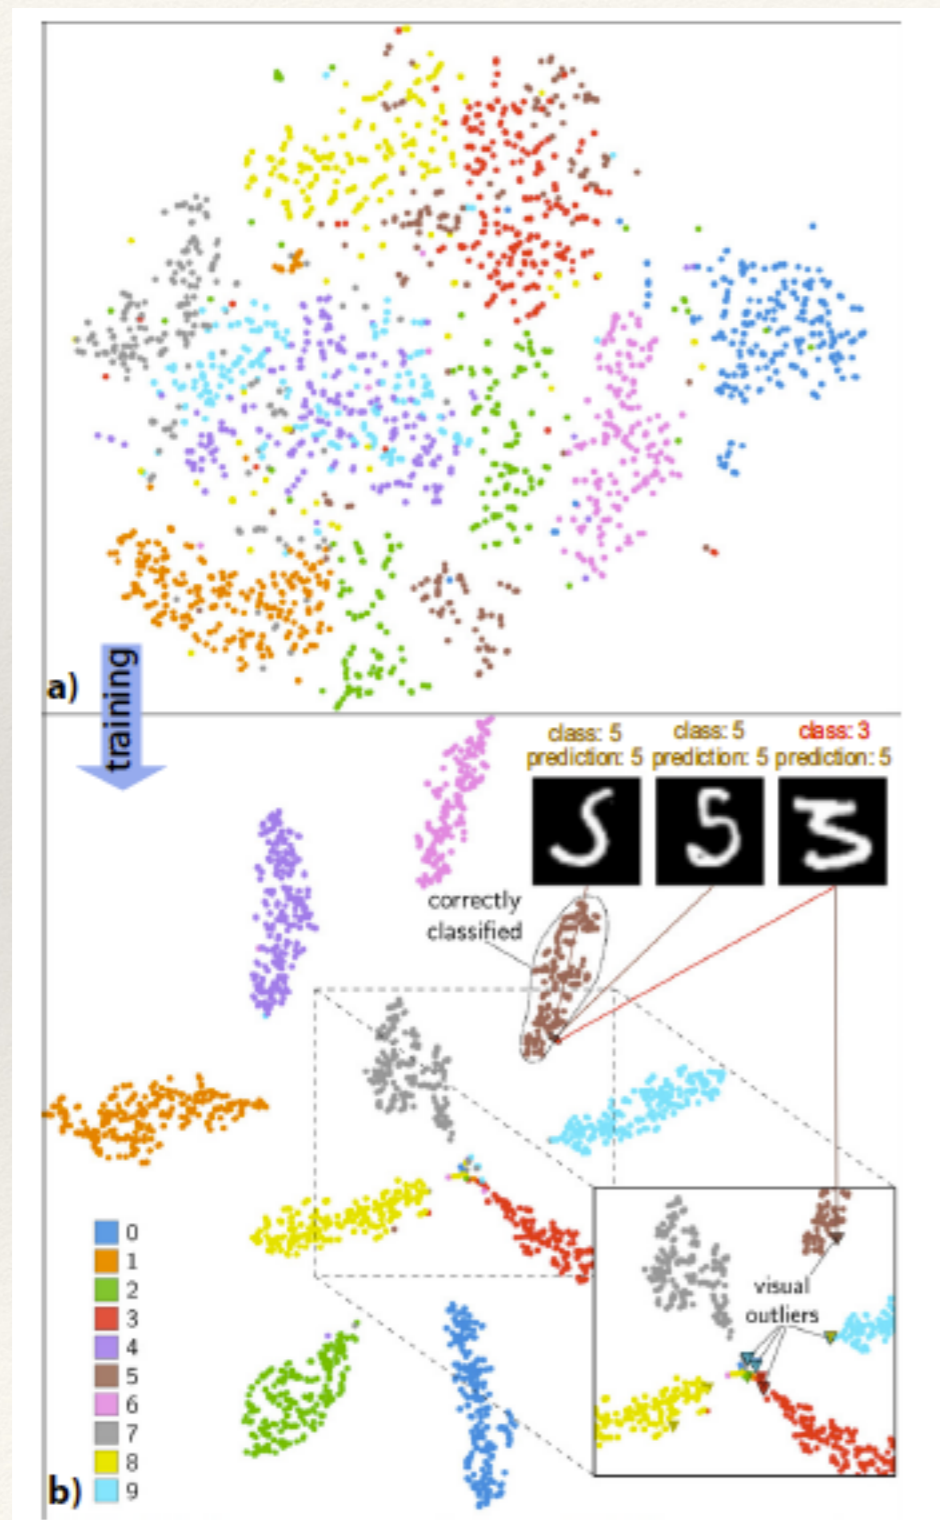
Model \ Dataset	MLP	CNN	State-of-the-art
MNIST	98.52%	99.62%	99.79% [47]
SVHN	77.38%	93.76%	98.08% [23]
CIFAR-10	52.91%	79.19%	91.78% [23]

Understanding Deep Networks



tSNE
projection
of MNIST
Images

Understanding Deep Networks



tSNE
projection
of last layer
of the neural network.

Fig. 3. Projection of the last MLP hidden layer activations, MNIST test subset. a) Before training (NH: 83.78%). b) After training (NH: 98.36%, AC: 99.15%). Inset shows classification of visual outliers.

Understanding Deep Networks



Fig. 4. Projection of the last MLP hidden layer activations before training, SVHN test subset (NH: 20.94%). Poor class separation is visible.

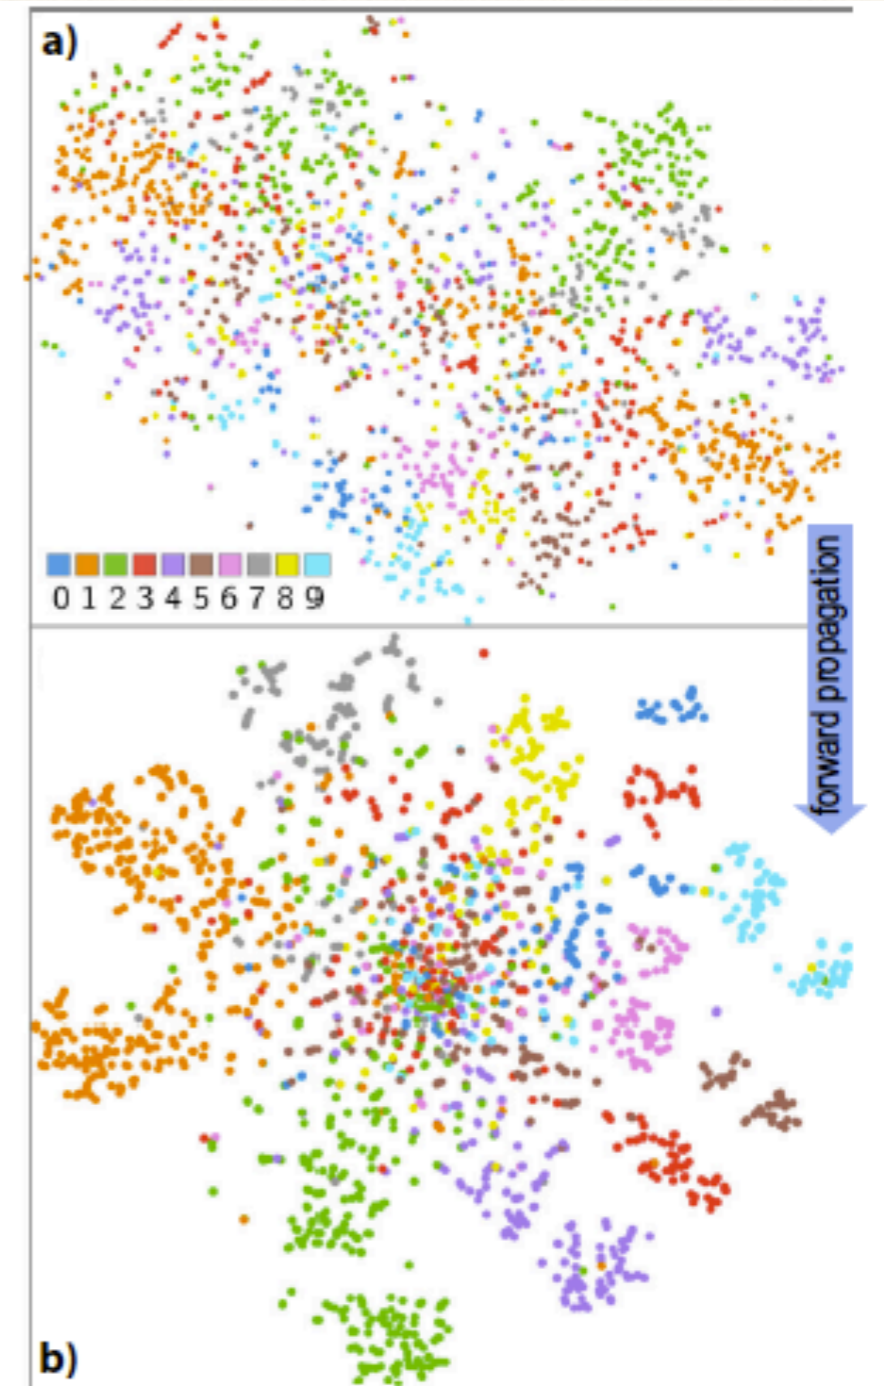


Fig. 5. Projection of the MLP hidden layer activations after training, SVHN test subset. a) First hidden layer (NH: 52.78%). b) Last hidden layer (NH: 67%).




Fig. 9. Projection of last CNN hidden layer activations after training, *CIFAR-10* test subset (NH: 53.43%, AC: 78.7%).

Understanding Deep Networks

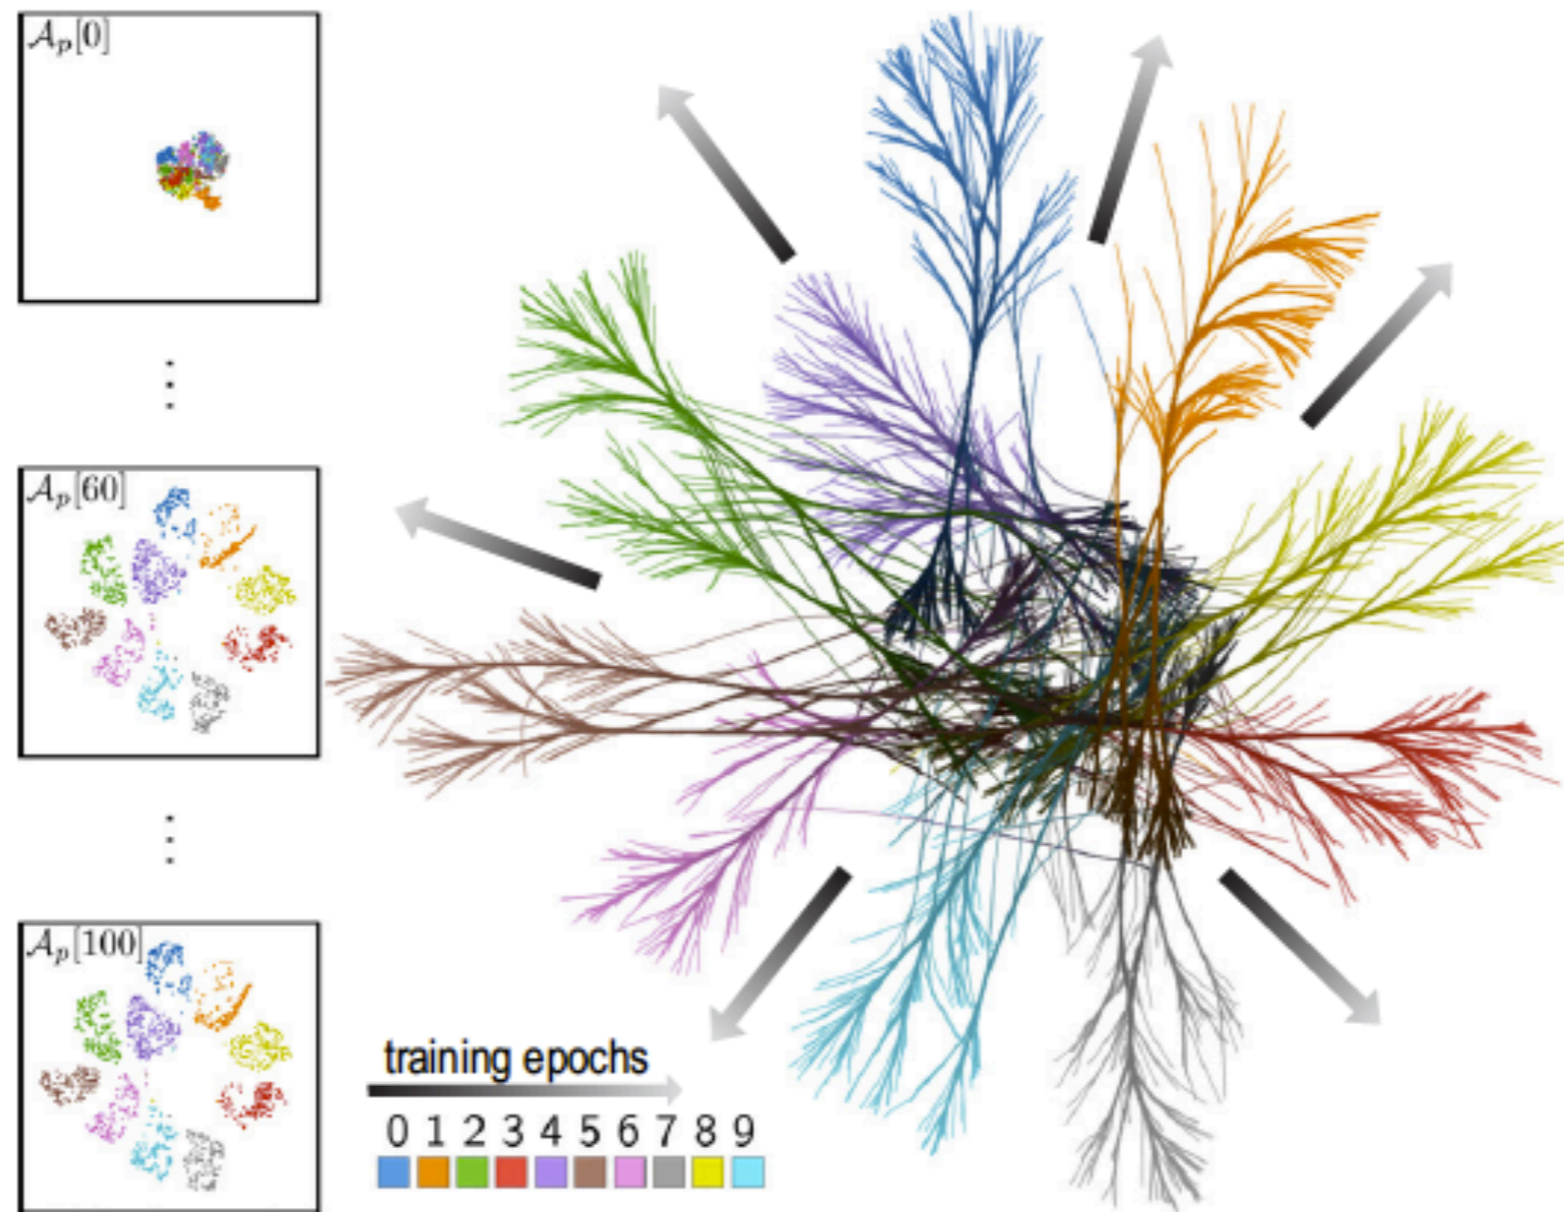


Fig. 11. Inter-epoch evolution, last CNN hidden layer, epochs 0-100, in steps of 20, *MNIST* test subset. Brighter trail parts show later epochs.

Understanding Deep Networks

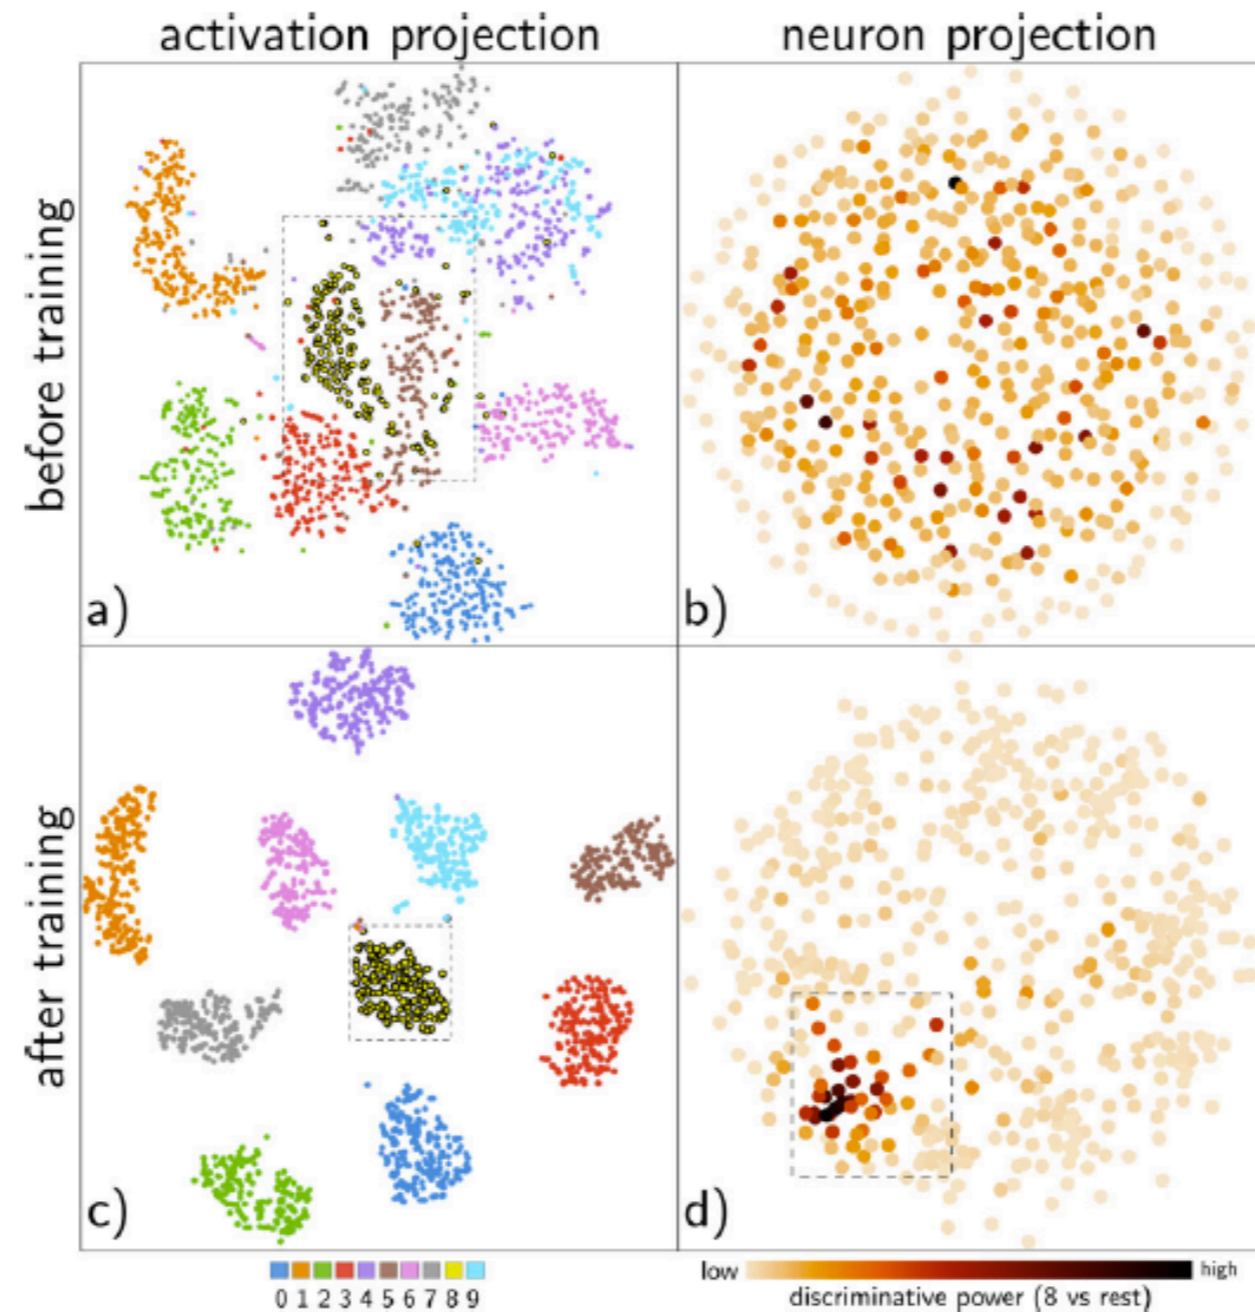


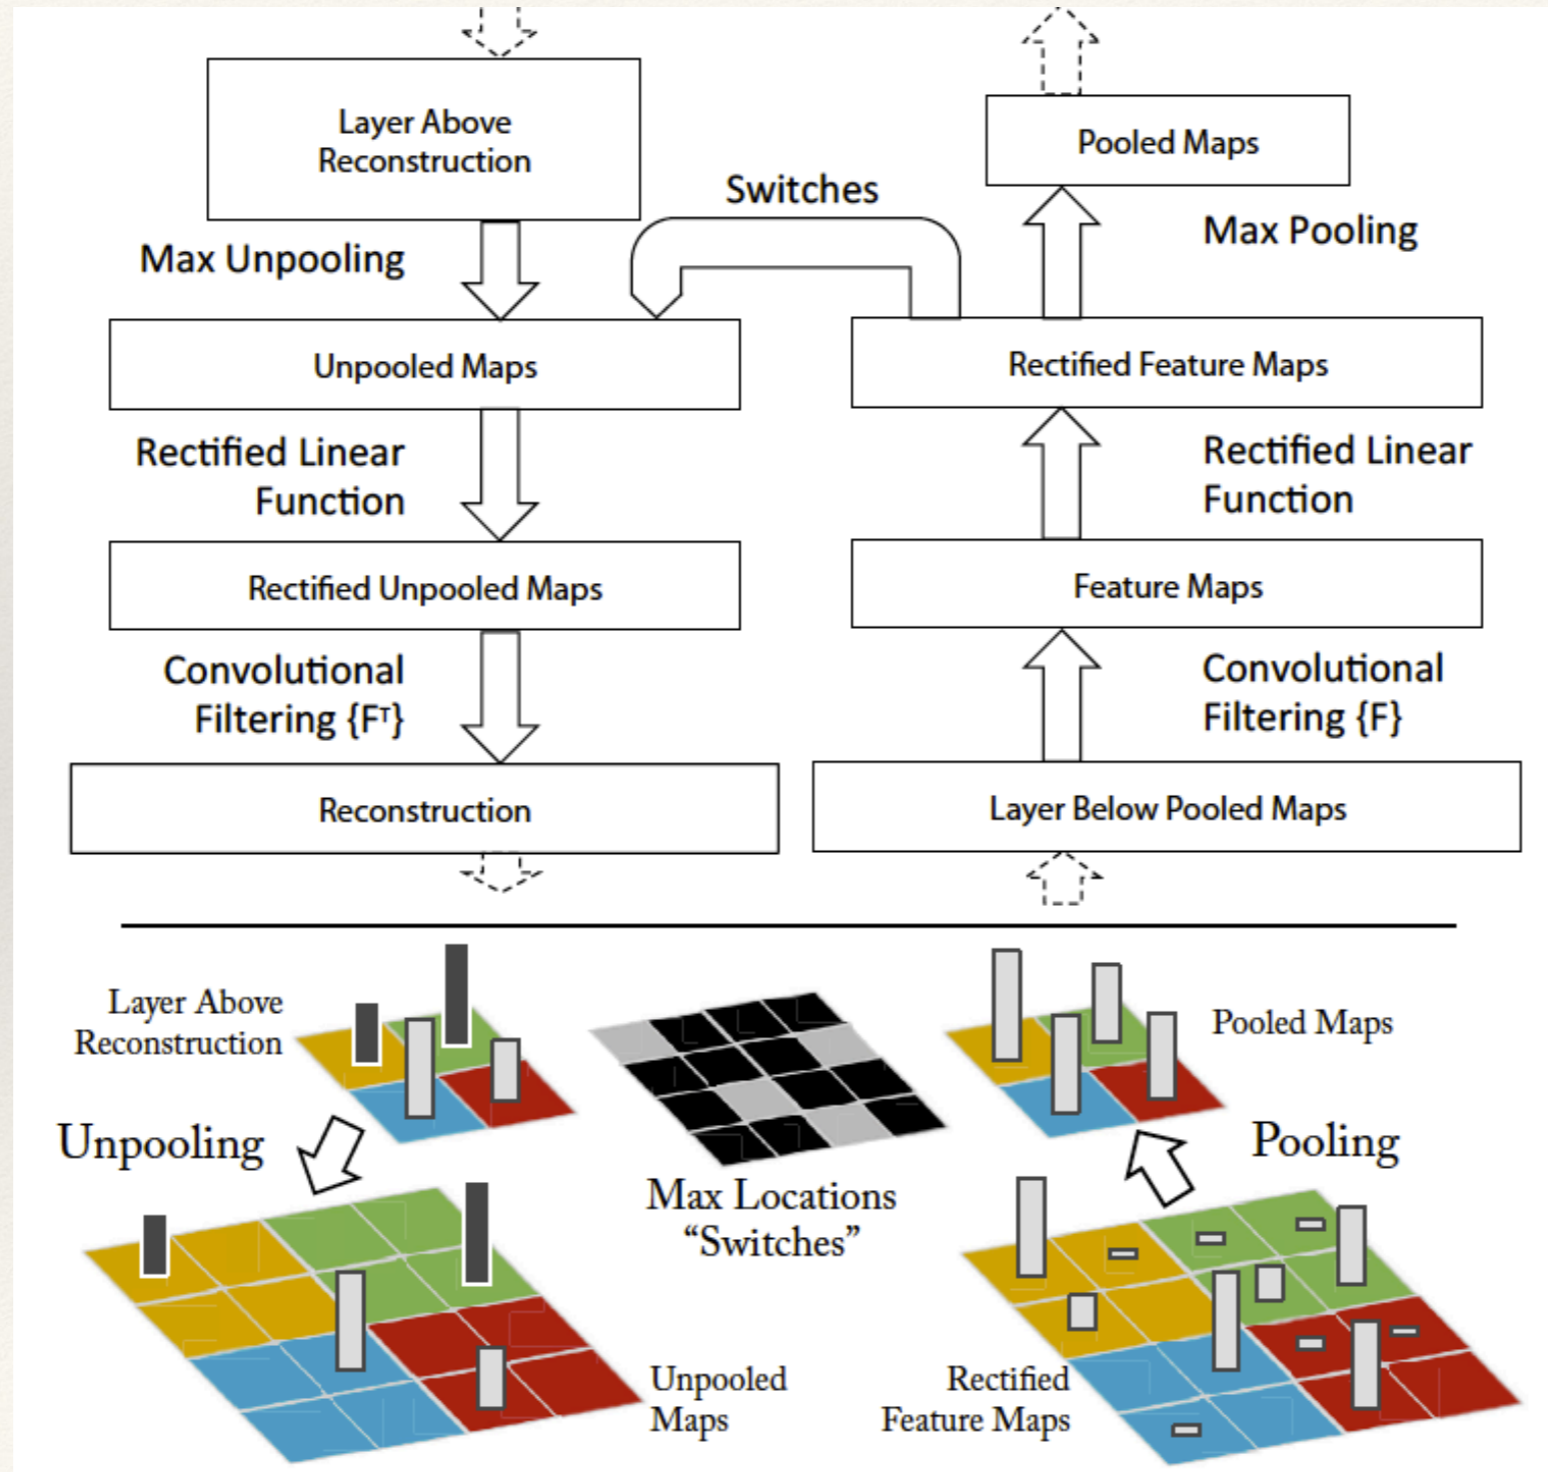
Fig. 12. Activation and neuron projections of last CNN hidden layer activations before and after training, *MNIST* test subset. Neuron projection colors show the neurons' power to discriminate class 8 vs rest.

Visualizing and Understanding Convolutional Networks


Matthew D. Zeiler and Rob Fergus

Dept. of Computer Science,
New York University, USA
{zeiler,fergus}@cs.nyu.edu

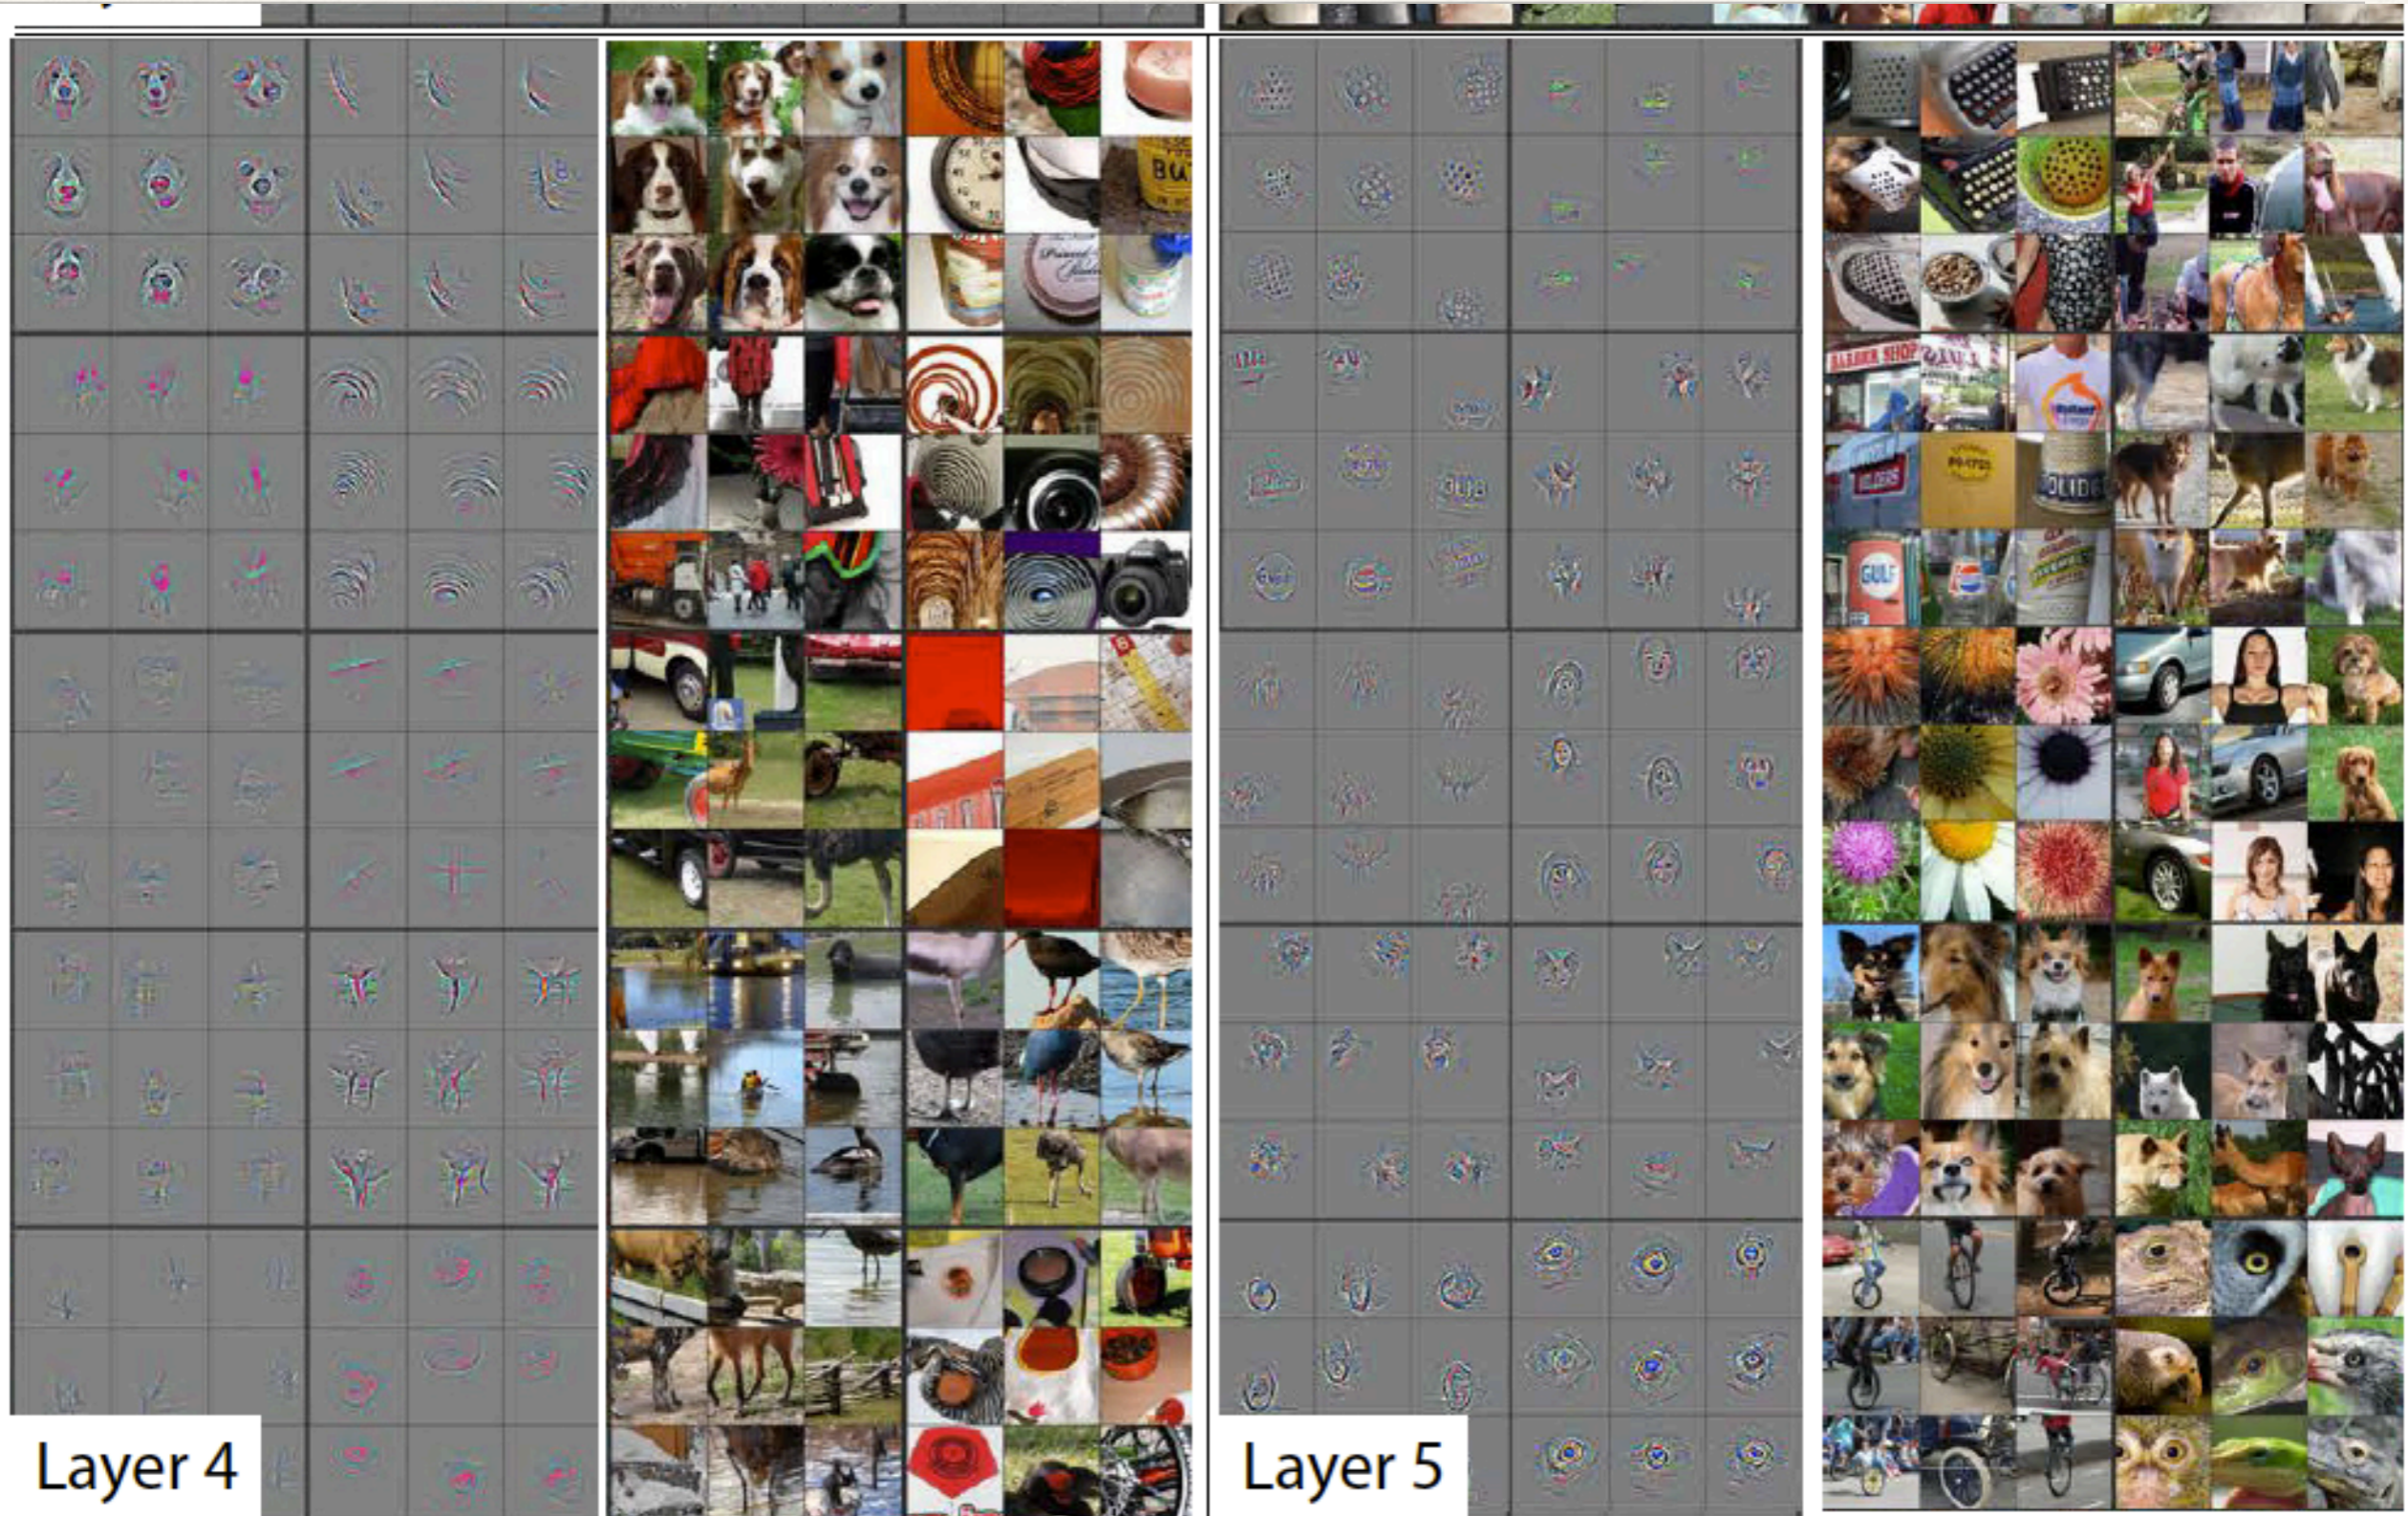
Understanding Deep Networks



Understanding Deep Networks




Understanding Deep Networks



Representation Learning in Deep Networks

[Zeiler, 2014]




UNDERSTANDING HOW DEEP BELIEF NETWORKS PERFORM ACOUSTIC MODELLING

Abdel-rahman Mohamed, Geoffrey Hinton, and Gerald Penn

Department of Computer Science, University of Toronto

Speech Recognition


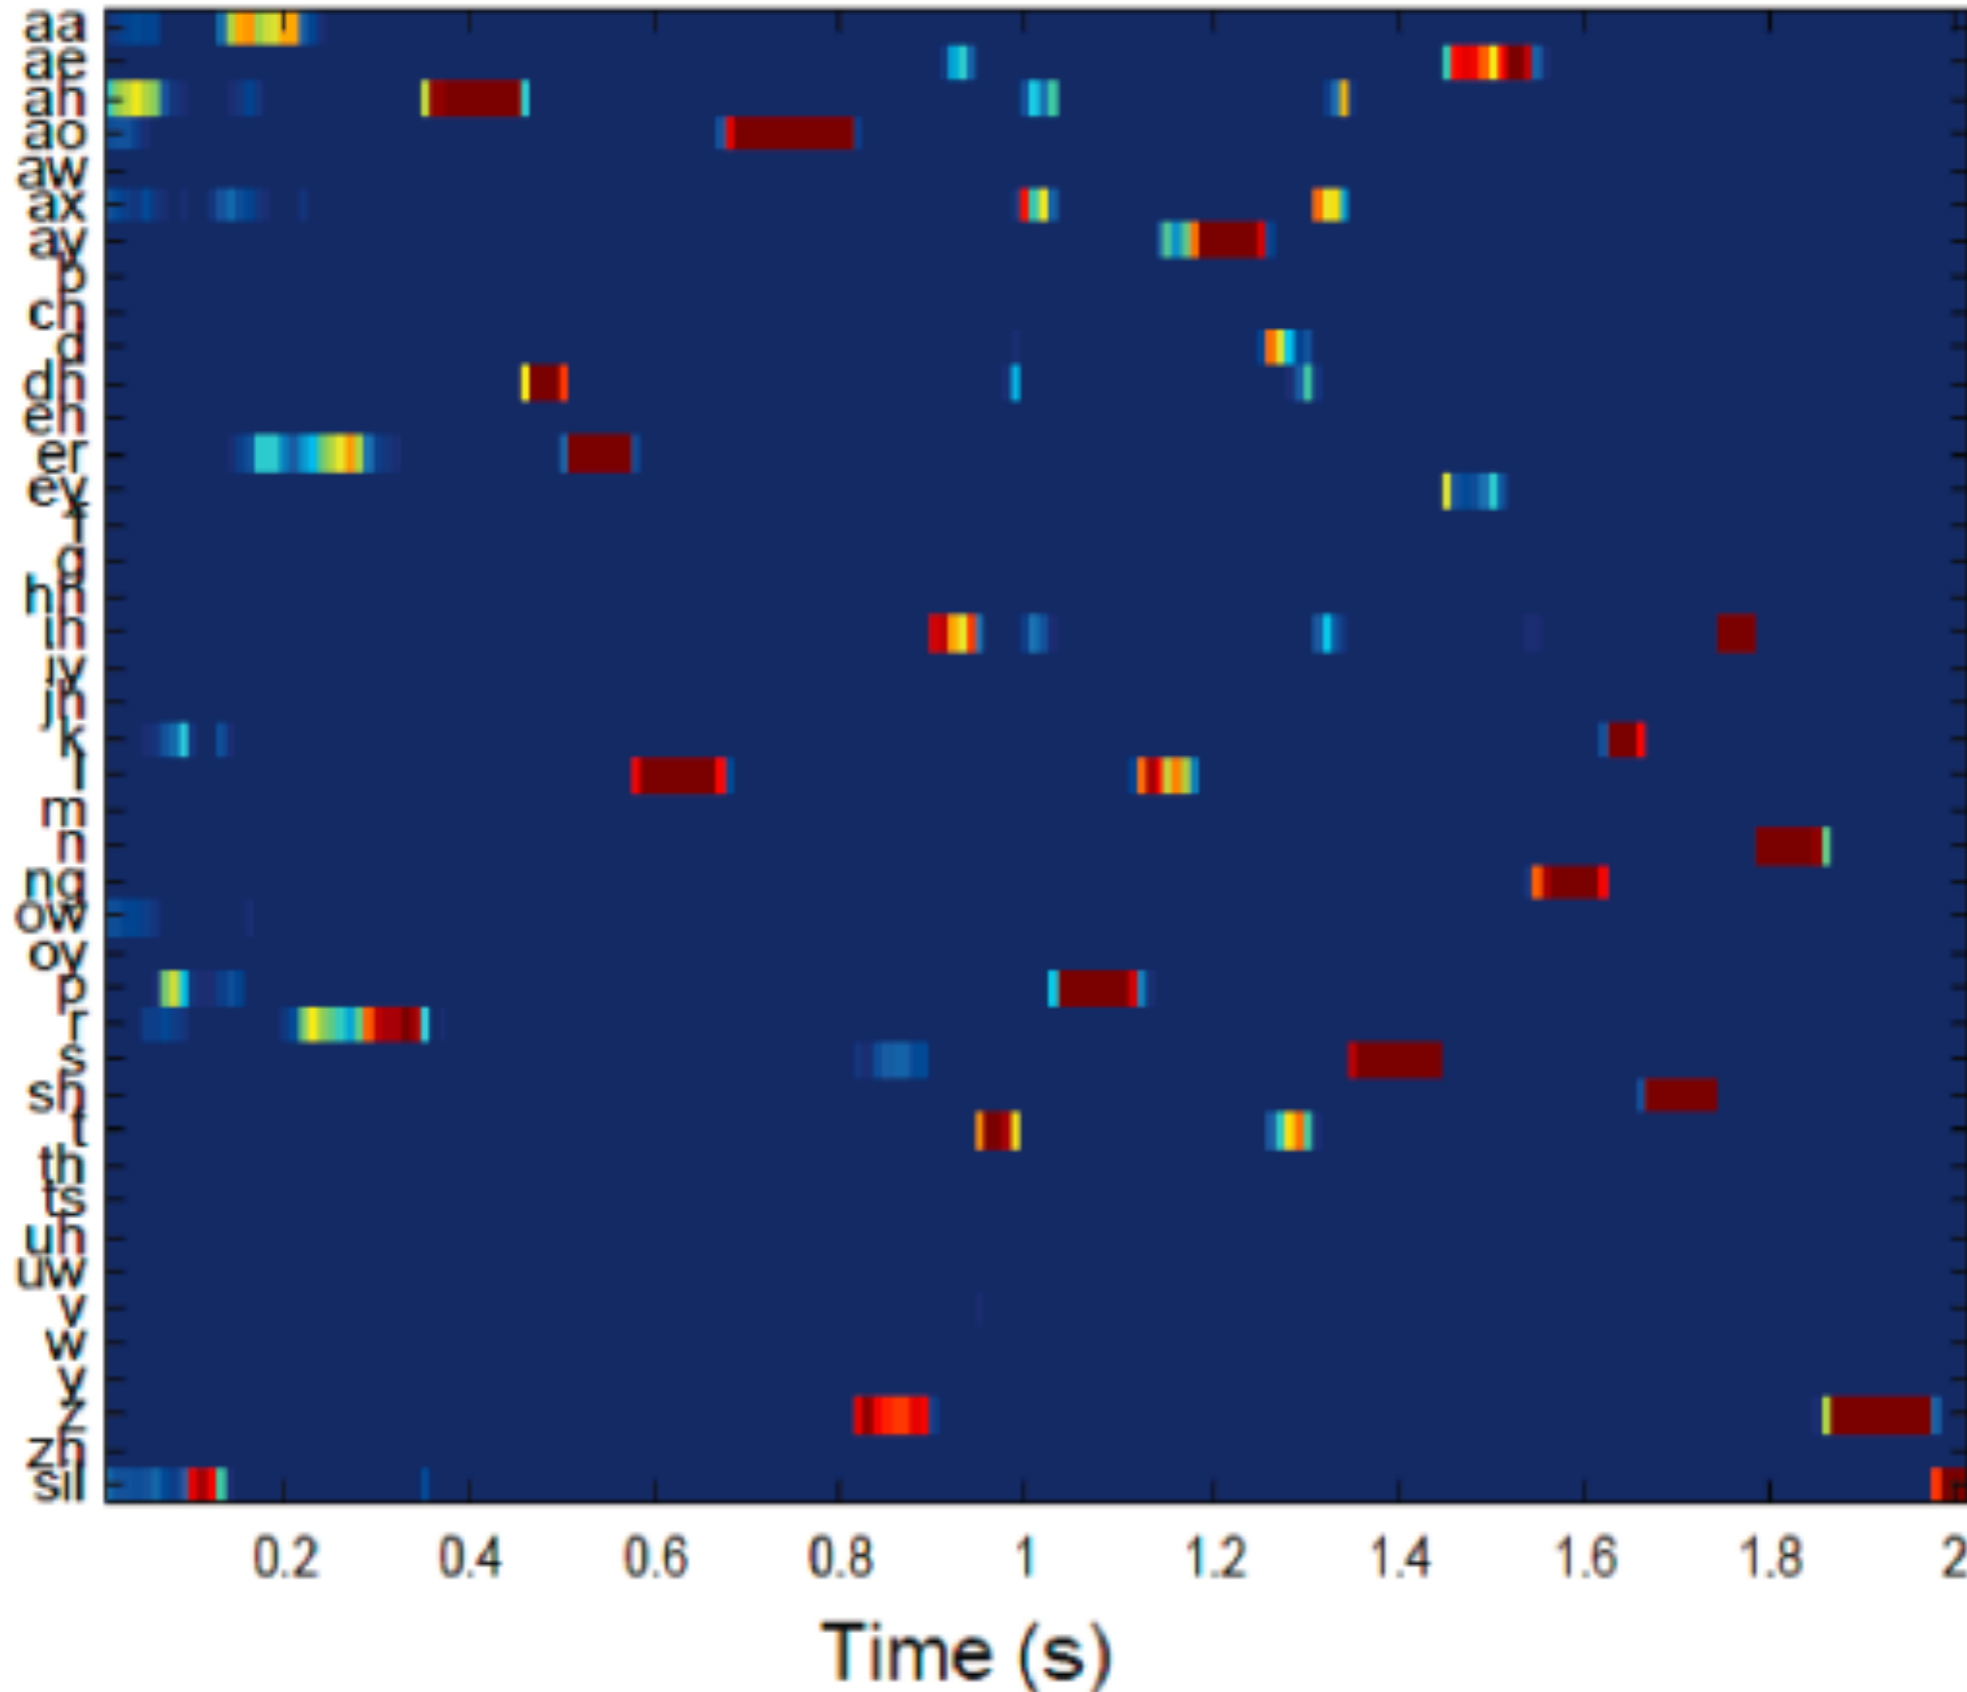
- Map the features to phone class. Using phone labelled data.



- Classical machine learning - train a classifier on speech training data that maps to the target phoneme class.


Back to Speech Recognition

Ma



Understanding DNNs for Speech


2-D projection of 1st layer DNN



[Abdel Rahman, 2012]

Understanding DNNs for Speech


2-D projection of 2nd layer DNN



[Abdel Rahman, 2012]

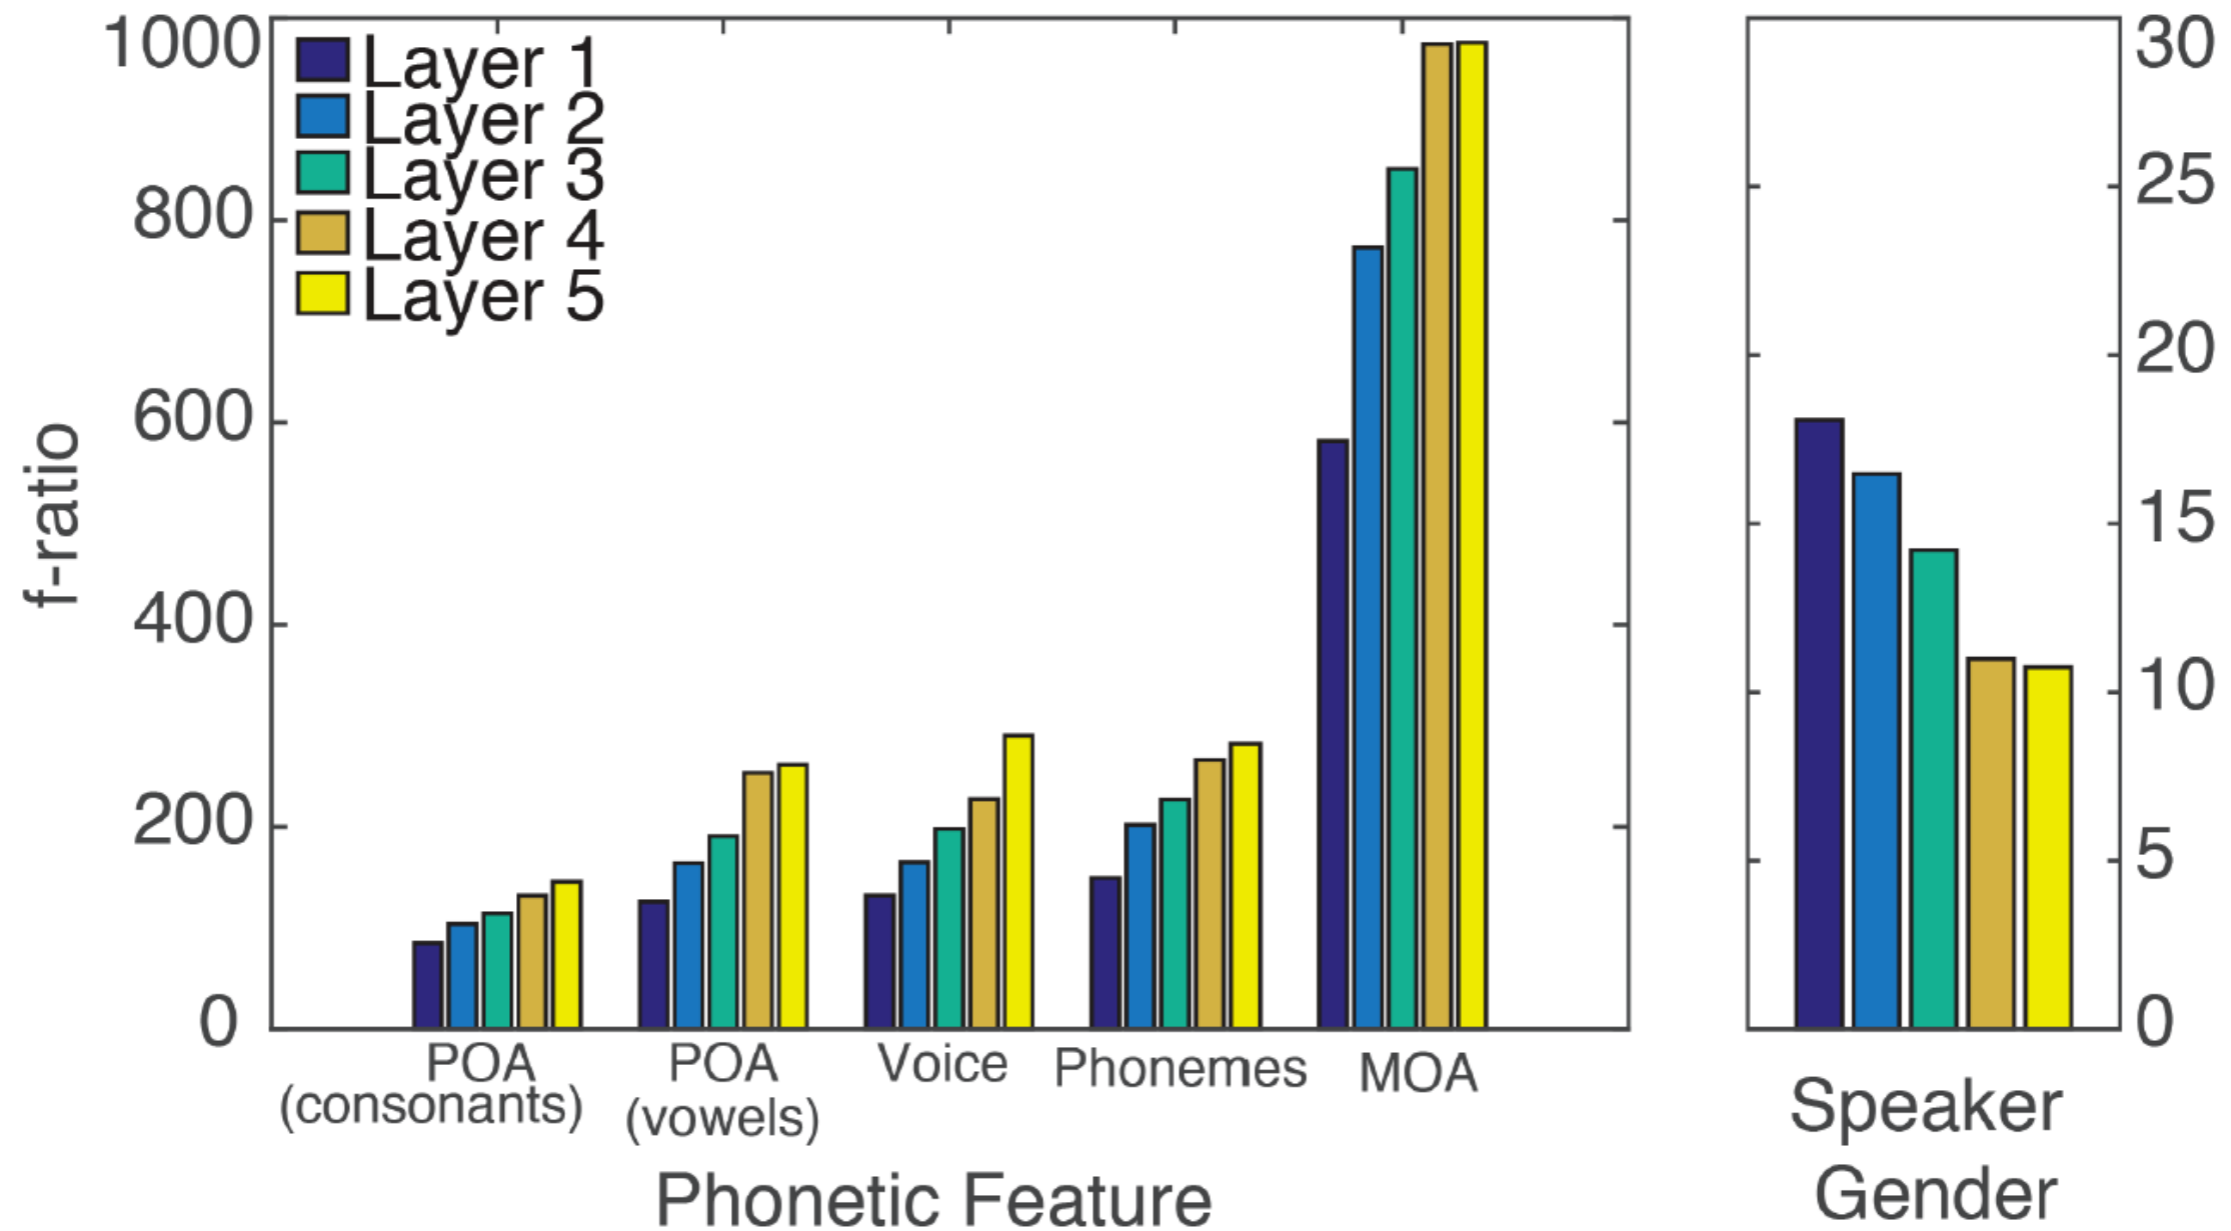
Understanding DNNs for Speech

2-D projection of final layer DNN



[Abdel Rahman, 2012]

Understanding DNNs for Speech



[Nagamine, 2015]