INDIAN INSTITUTE OF SCIENCE
भारतीय विज्ञान संस्थान

LEAP

Visualization Tool Using t-SNE

Visualizing Data using t-SNE

Laurens van der Maaten LVDMAATEN@GMAIL.COM
TiCC
Tilburg University
P.O. Box 90153, 5000 LE Tilburg, The Netherlands

Geoffrey Hinton HINTON@CS.TORONTO.EDU
Department of Computer Science
University of Toronto
6 King's College Road, M5S 3G4 Toronto, ON, Canada

Dimensionality reduction $\{x_n\}_{n=1}^{N} \in \mathbb{R}^D$

$\{y_n\}_{n=1}^{N} \in \mathbb{R}^d$

$y_n = W^T (x_n - \mu)$

largest variance

$d << D$.

$d = 2$ for visualization.

Neighborhood Embedding

(x_i, x_j) are close in D dimensional space

$\{y_i, y_j\}$ to be close in d dimensional space

Dimensionality reduction

$\hookrightarrow$ preserving neighboorhood

probability model

and distance between probability distn

Visualization tool using t-SNE

$p_{ii} = 0$

Gaussian

D^n

width parameter

$$p_{ij} = \frac{\exp\left(-\|x_i - x_j\|^2/2\sigma^2\right)}{\sum_{k \neq l} \exp\left(-\|x_k - x_l\|^2/2\sigma^2\right)}$$

k,l

normalization

$\mathbb{R}^D$

student t-distribution

$\to 0$

$$q_{ij} = \frac{\left(1 + \|y_i - y_j\|^2\right)^{-1}}{\sum_{k \neq l}\left(1 + \|y_k - y_l\|^2\right)^{-1}}$$

k,l

$q_{ii} = 0 \;\; \forall i$

$\mathbb{R}^d$

$$C = KL(P\|Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}$$

KL divergence

$-\sum_{i,j} p_{ij} \log q_{ij}$

Goal in t-SNE

$$\underset{\{y_1, \dots y_N\}}{\arg\min}\ C. = \underline{y}_1^*, \underline{y}_2^* \dots \underline{y}_N^*$$

Gradient descent $\qquad y_i \in \mathbb{R}^d$

Iterative update of lower dimensional representation.

Visualizing back propagation in tSNE

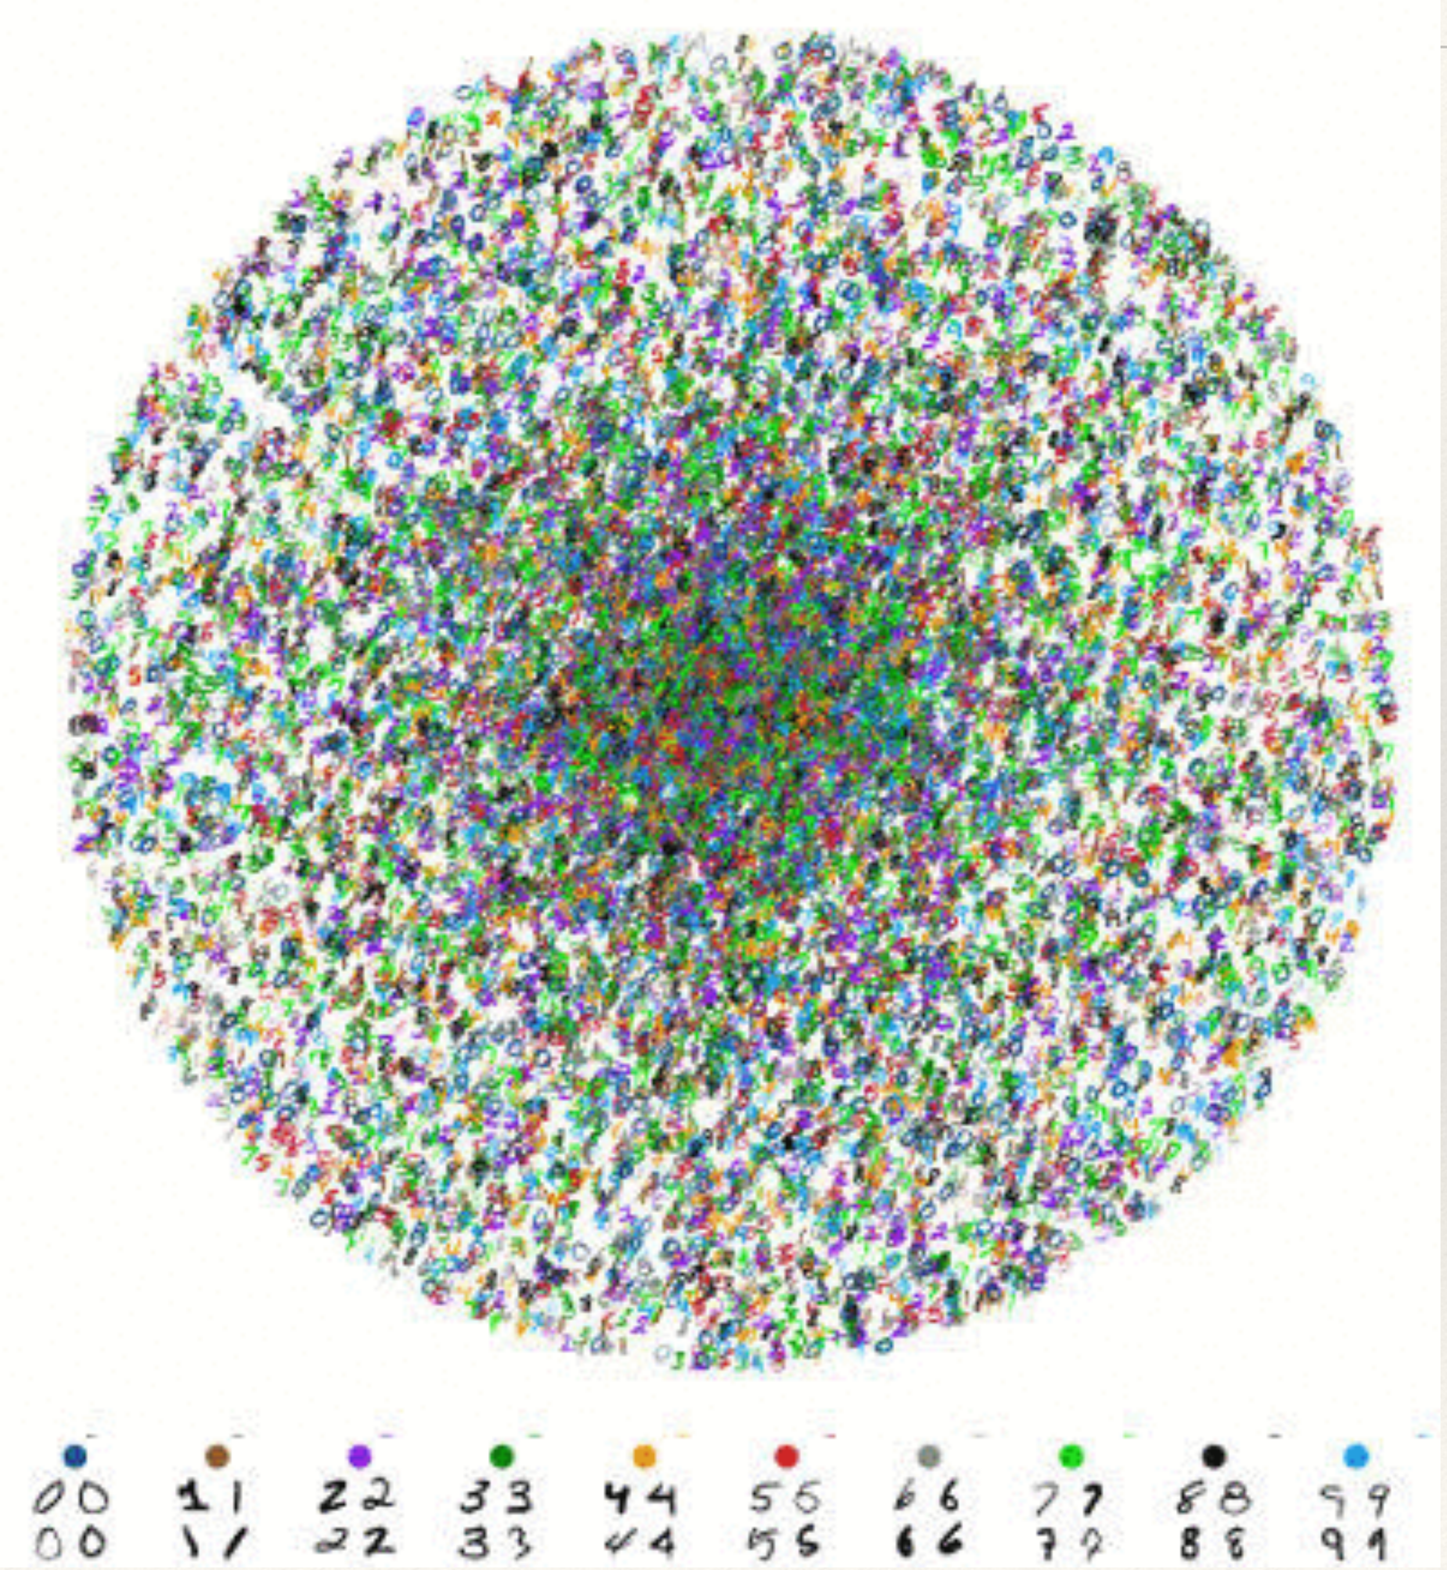

Understanding Deep Networks

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

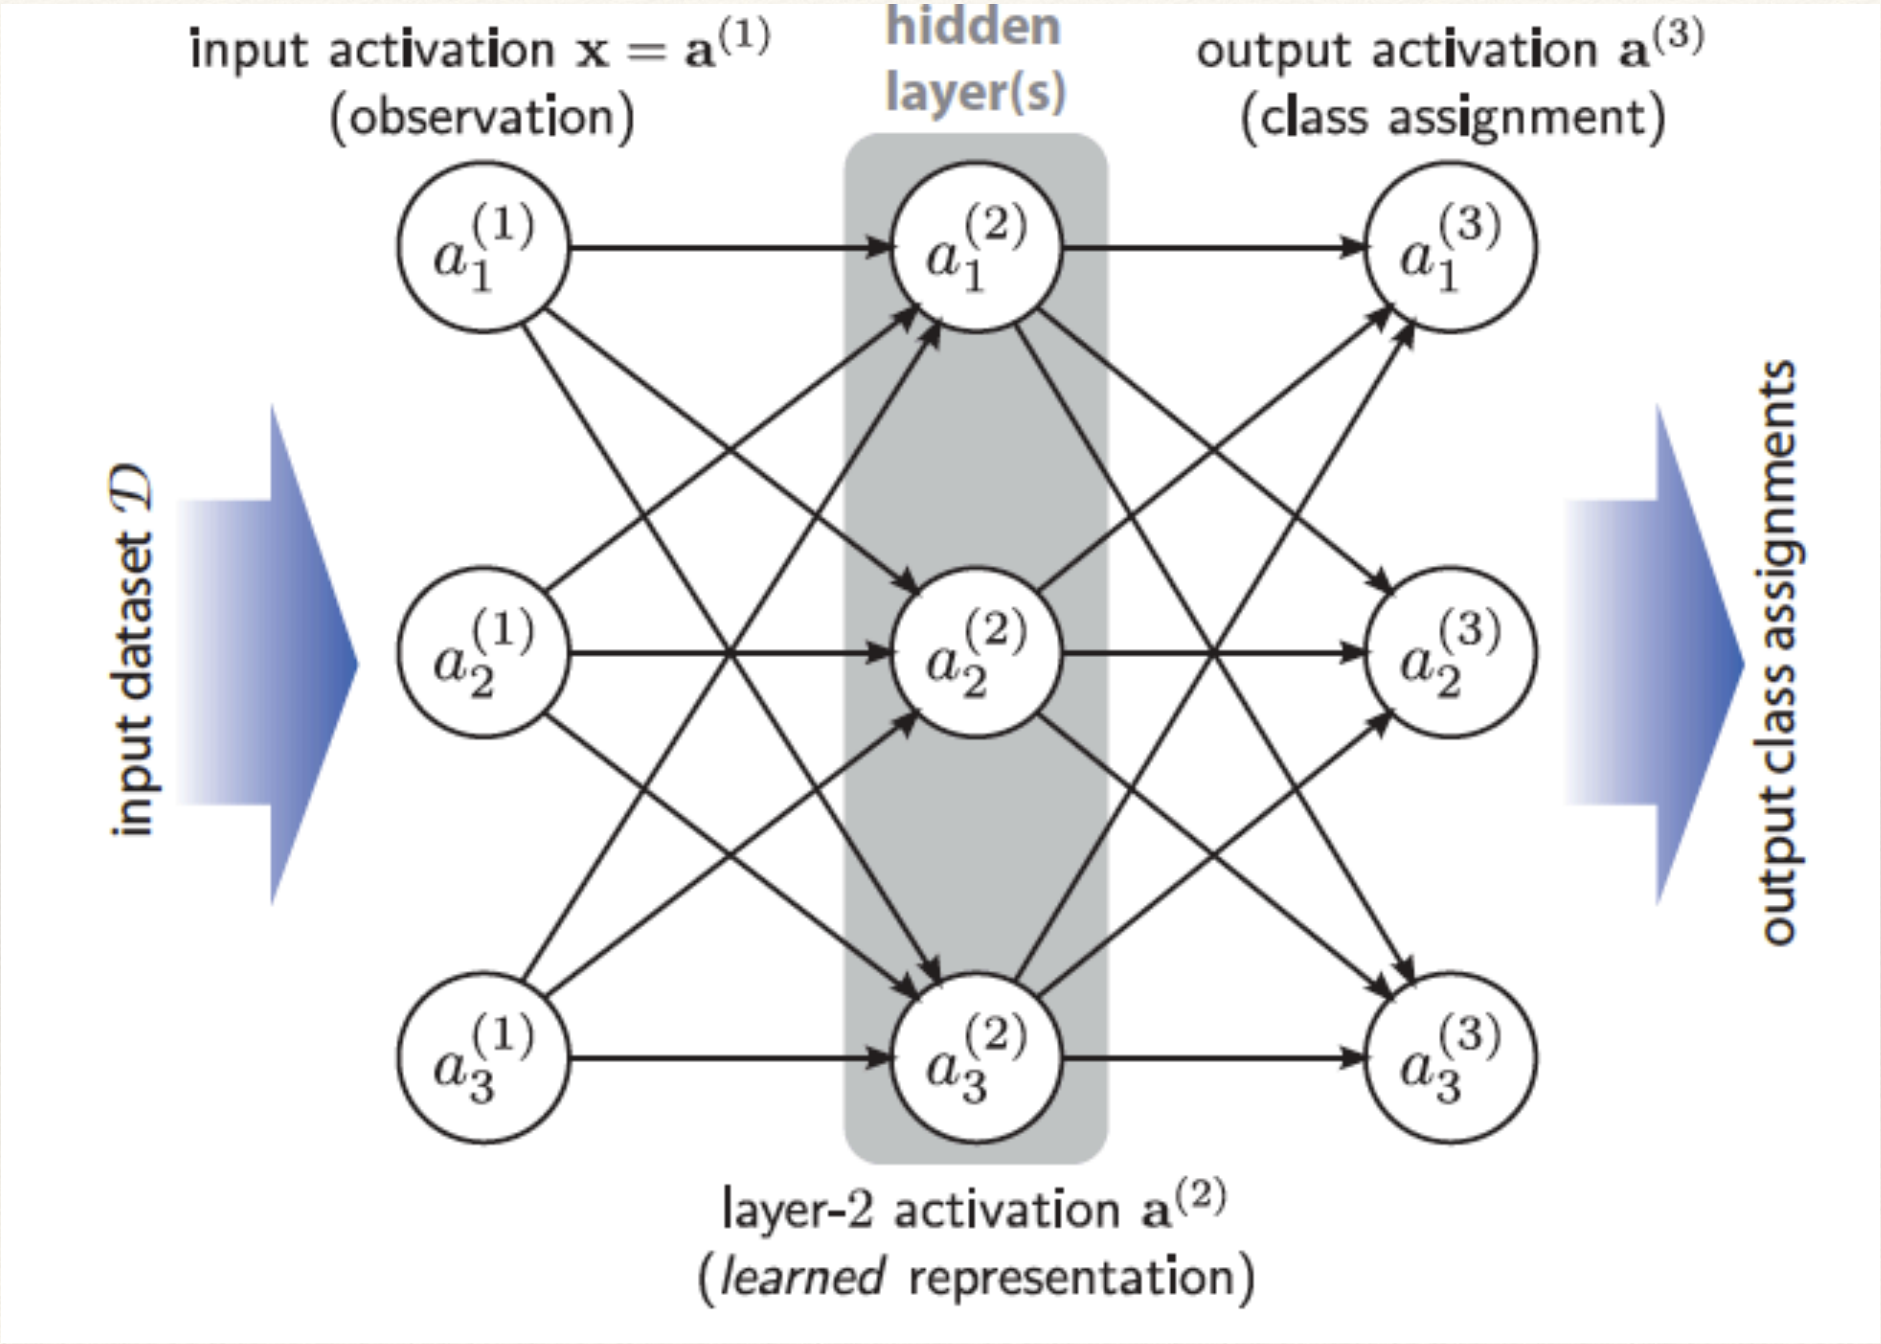

Understanding Deep Networks

SVHN dataset

Understanding Deep Networks

CIFAR-10

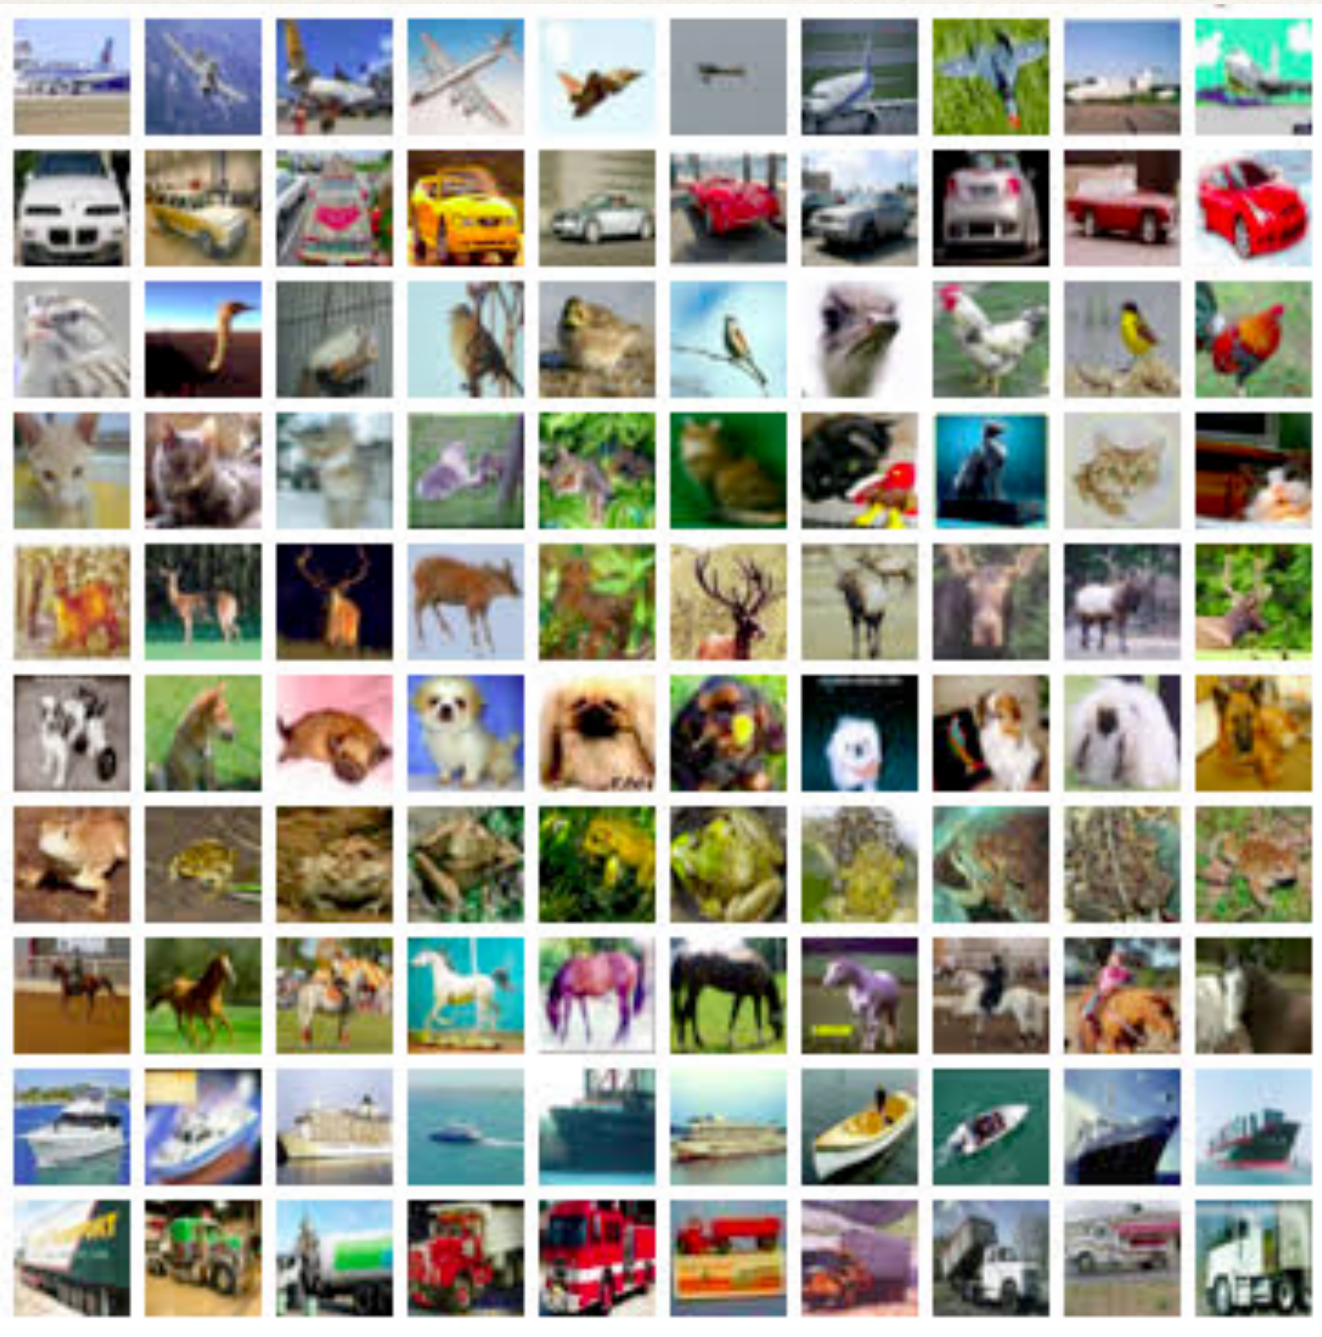

Table 1. Test Set Accuracies for our Two Architectures

Dataset \ Model	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

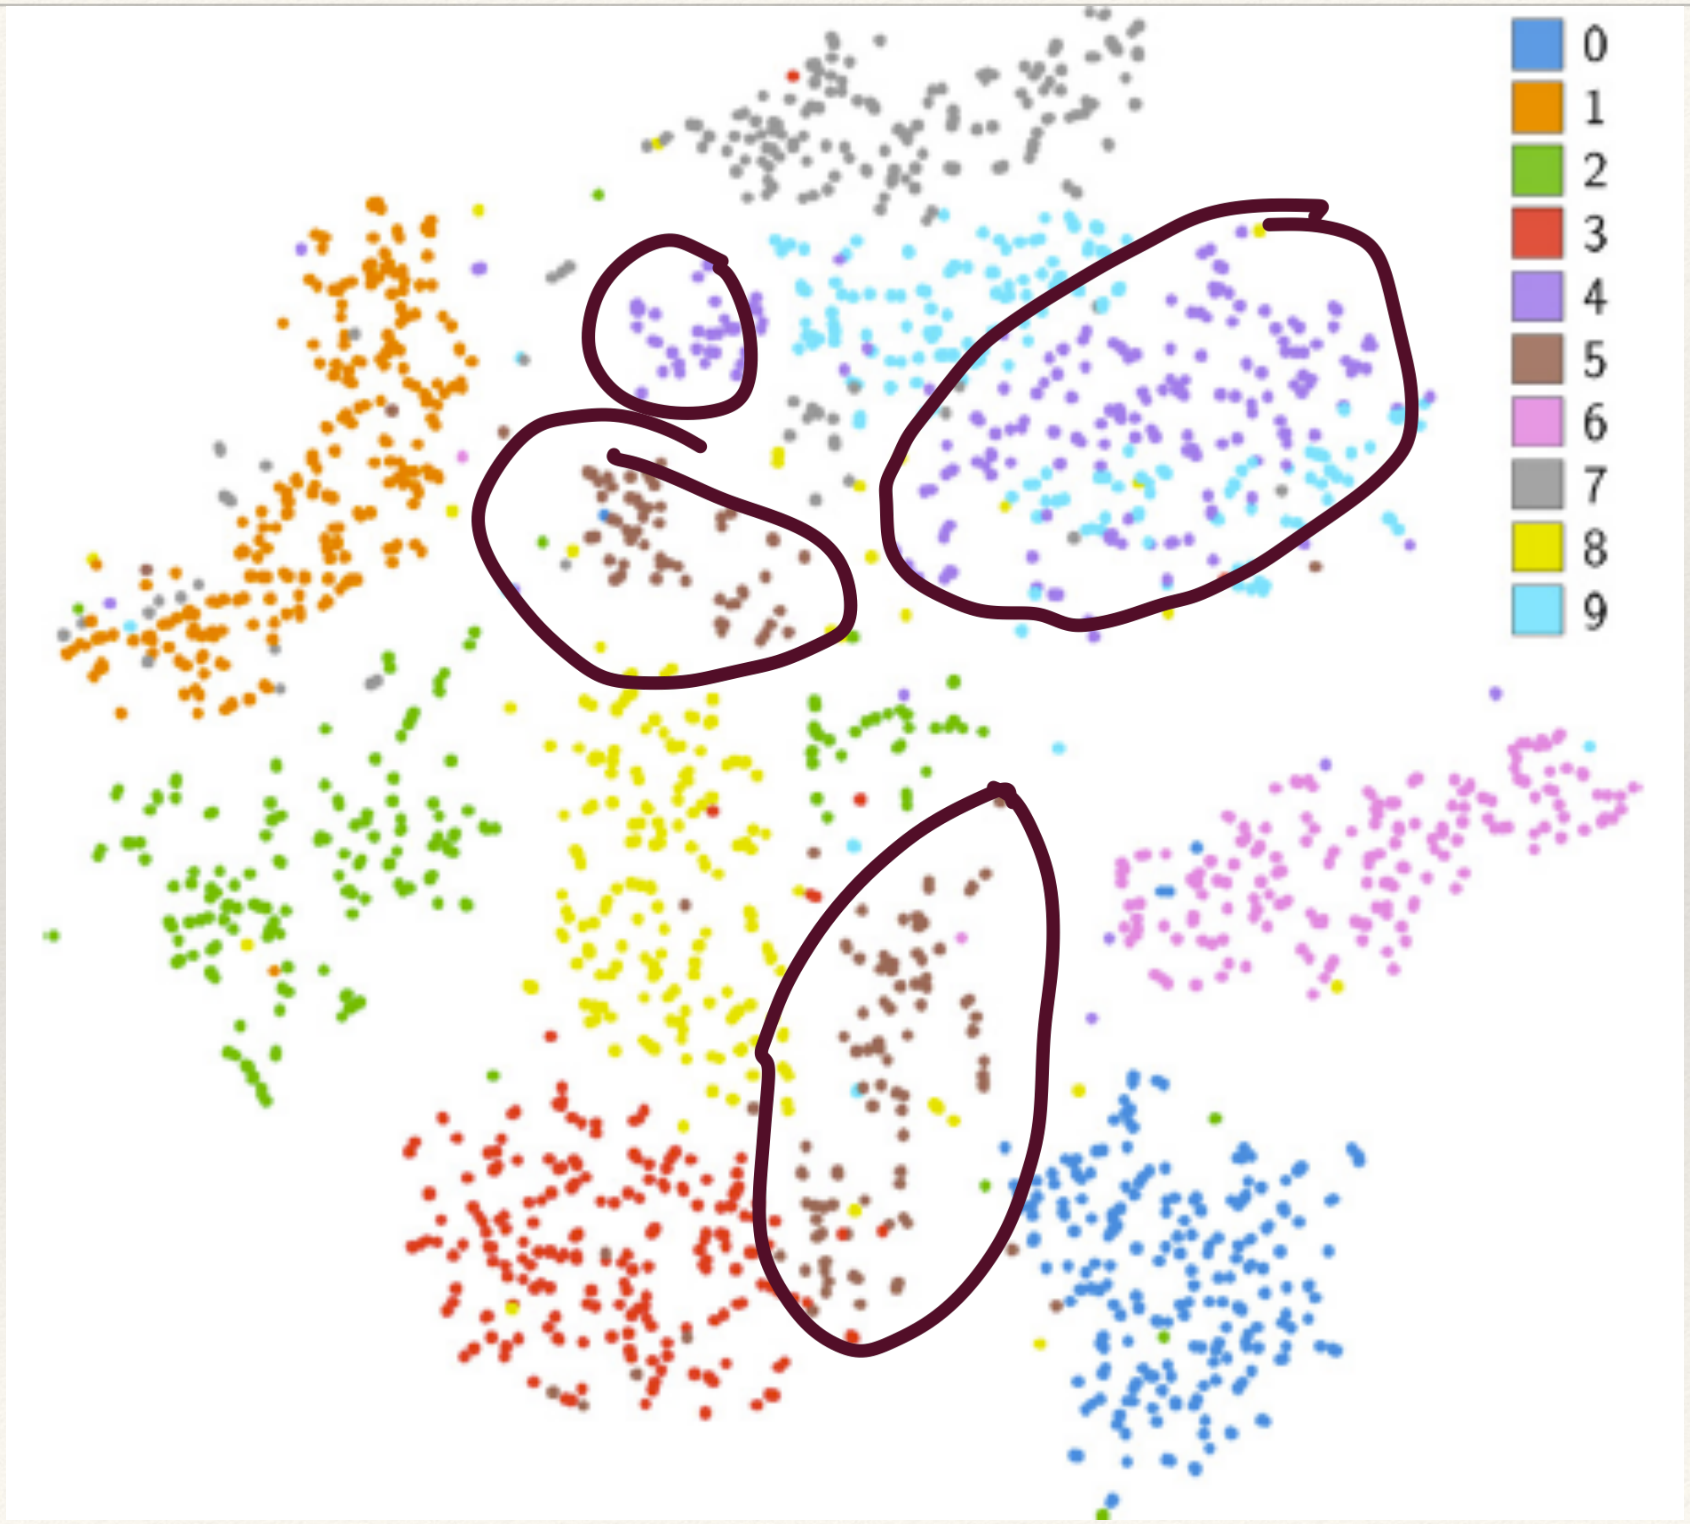

4 4

tSNE projection of MNIST Images

Understanding Deep Networks

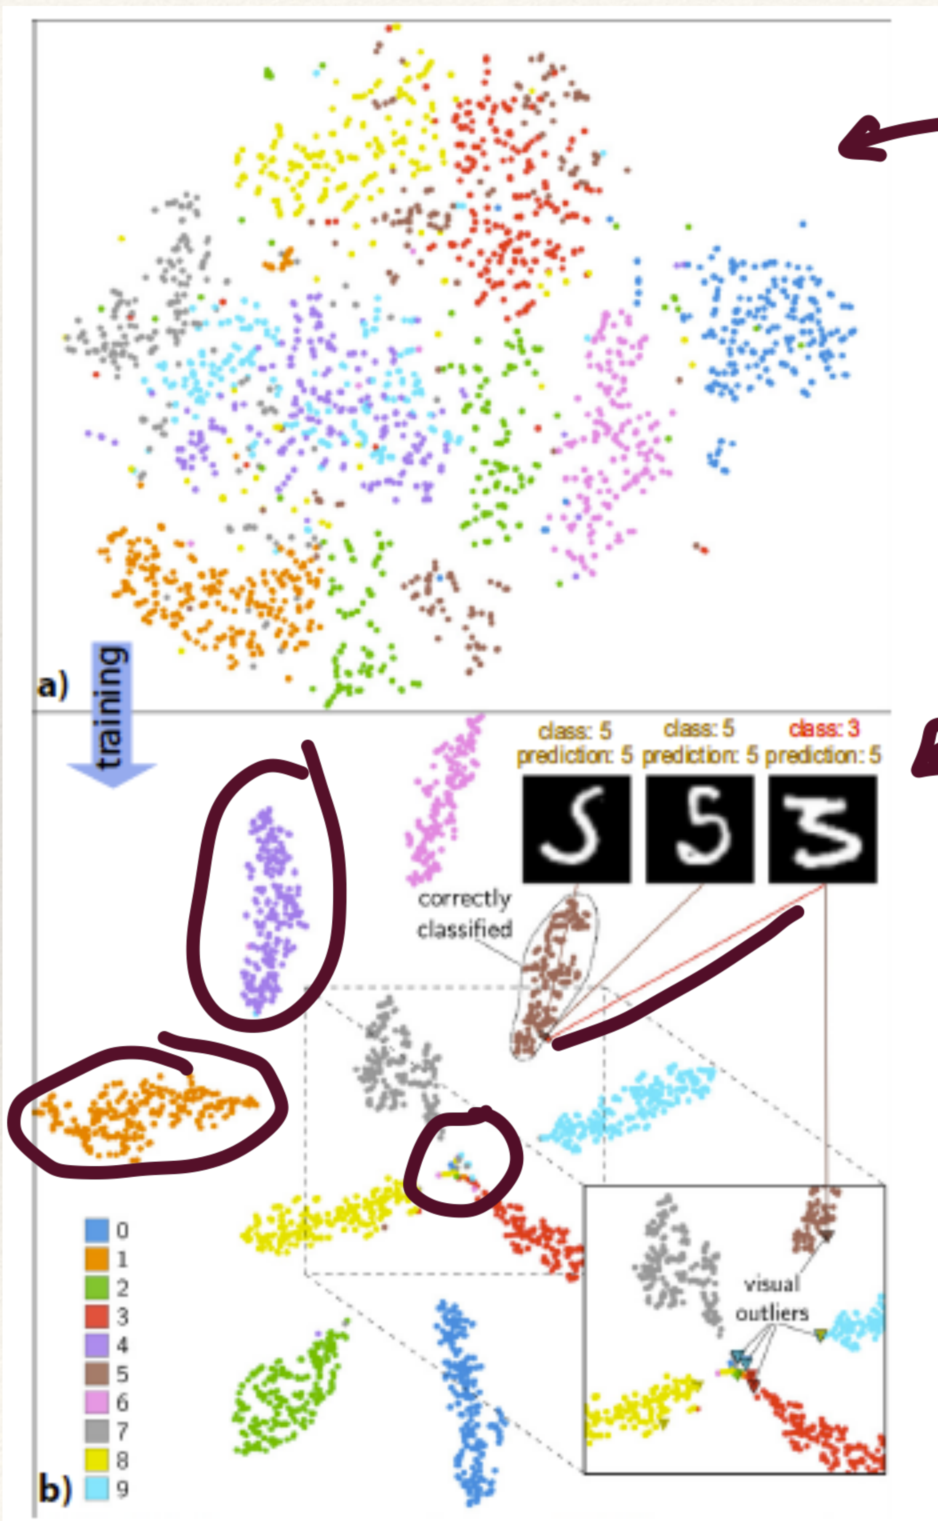

tSNE
projection
of last layer
of the neural network.

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.

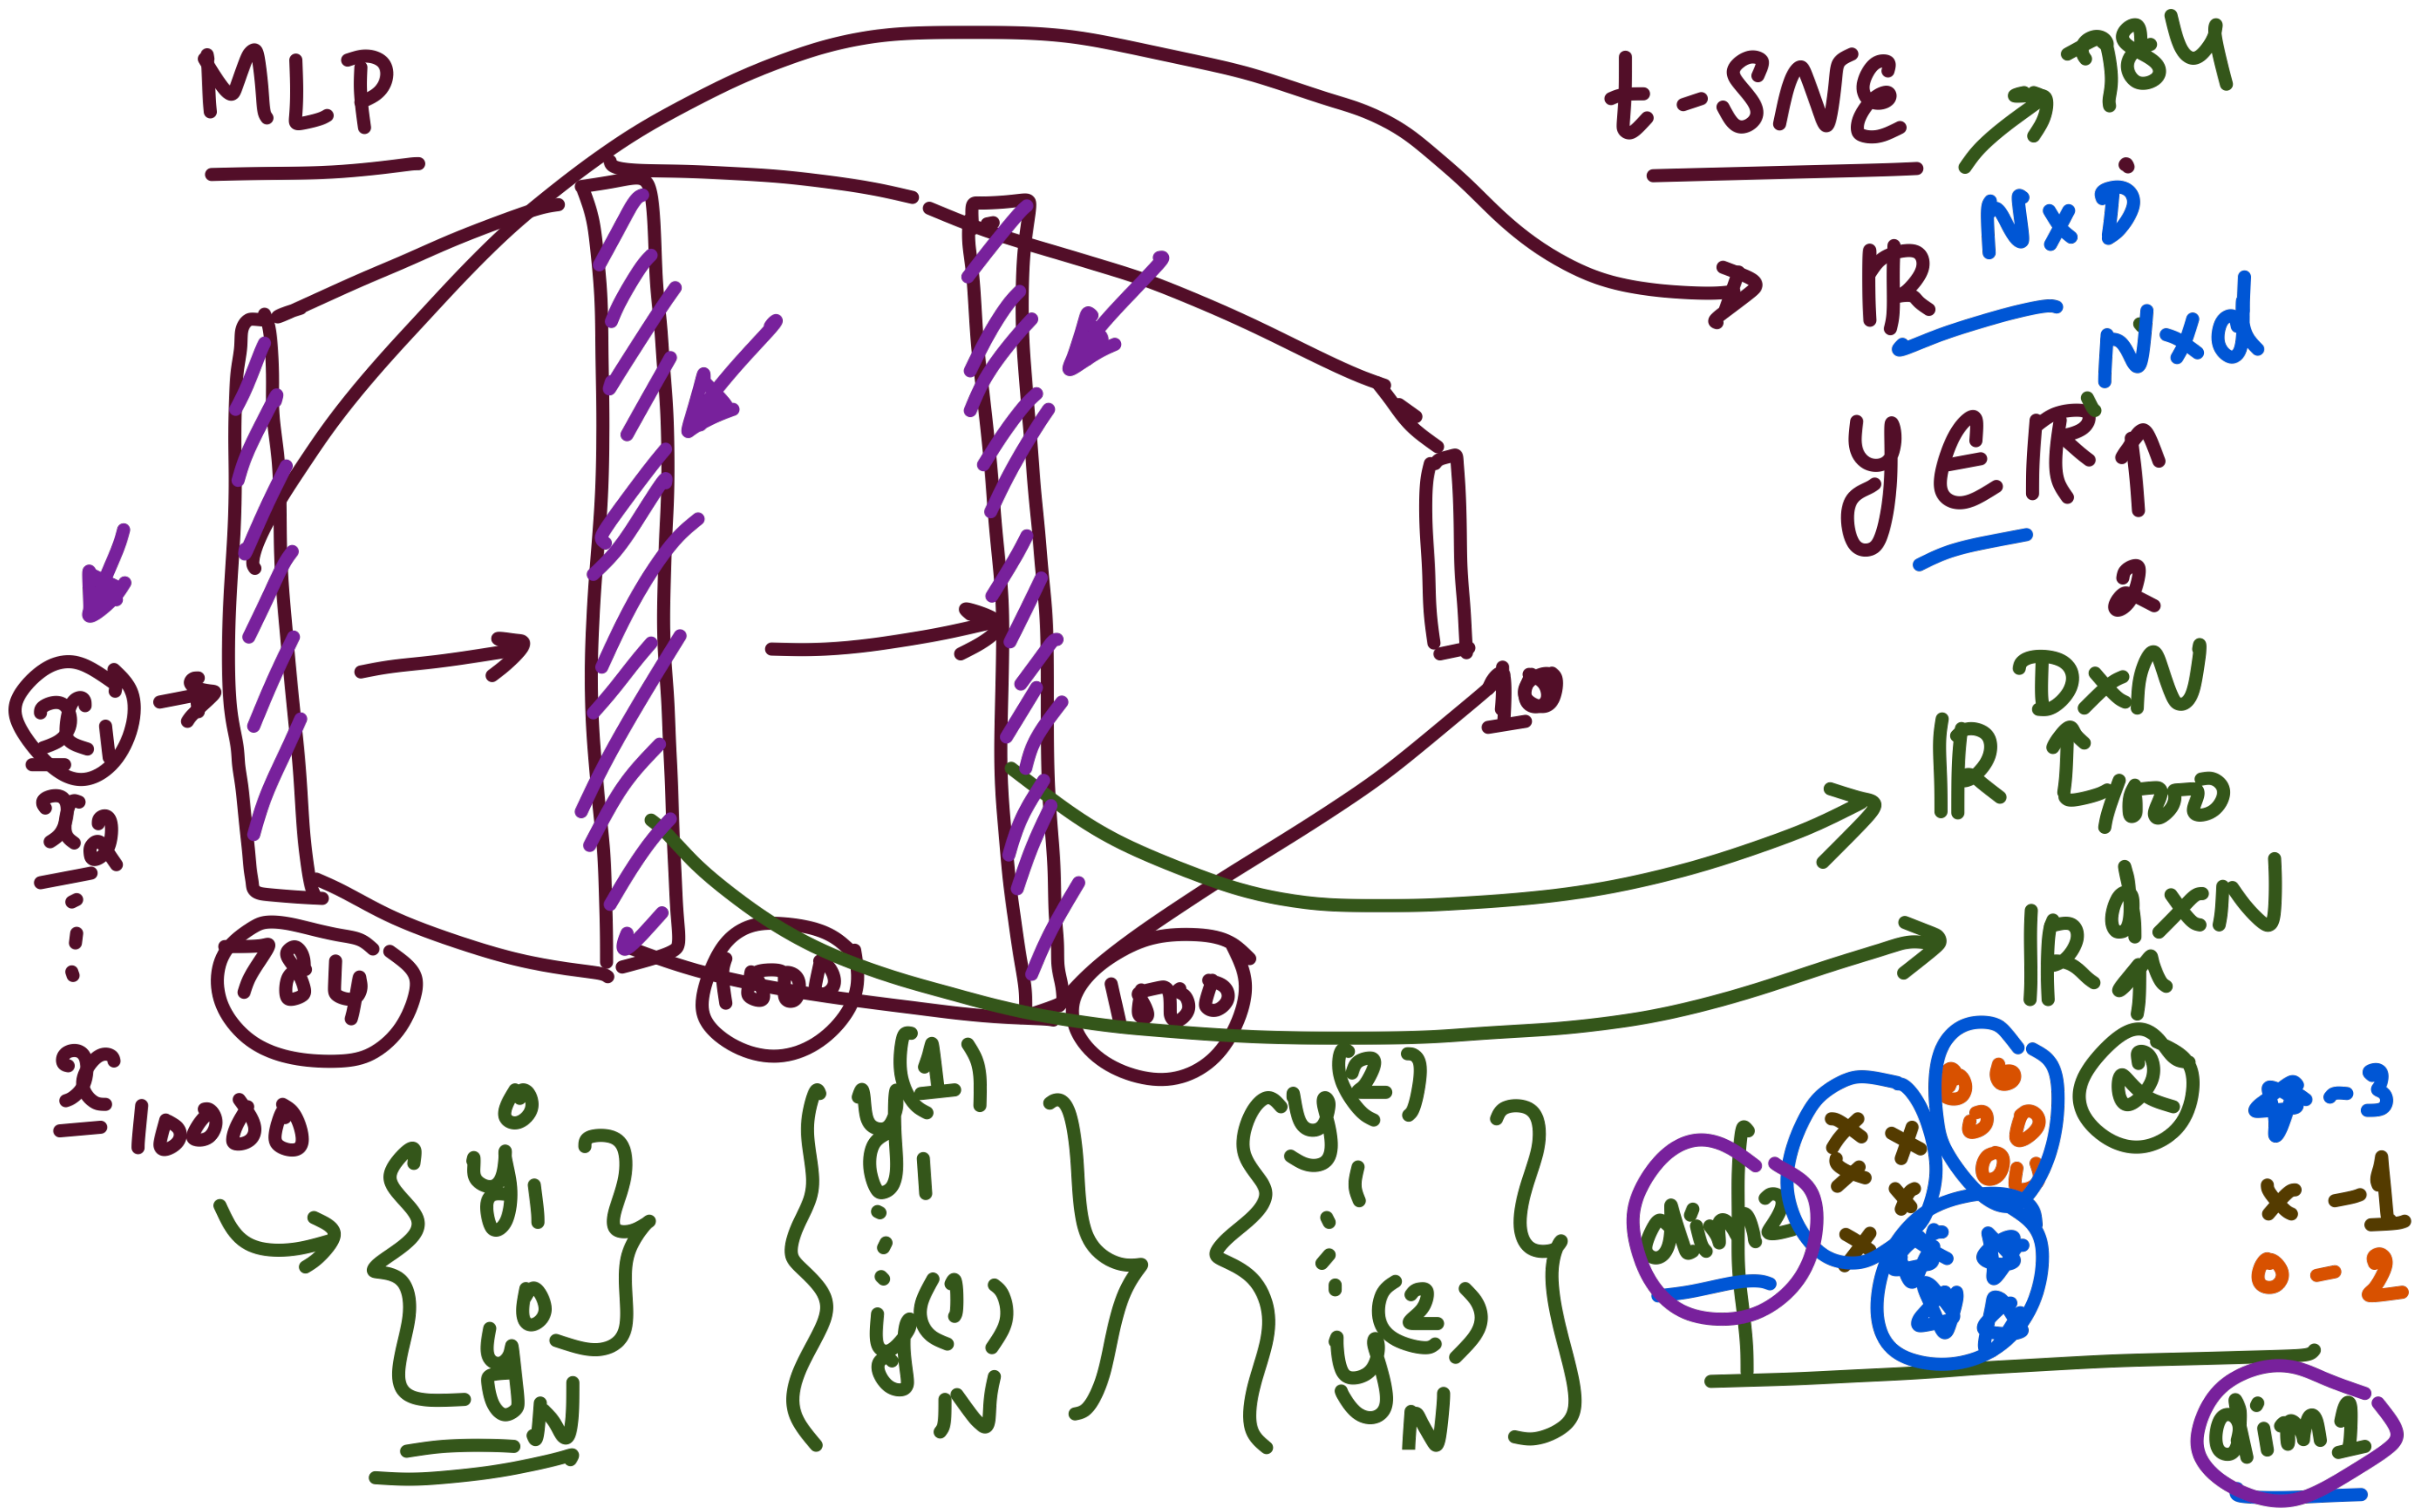

MLP
t-SNE
784
$\mathbb{R}^{N \times D}$
$y \in \mathbb{R}^{N \times d}$
2
$\mathbb{R}^{D \times N}$
1000
$\mathbb{R}^{d \times N}$
2
x_1
x_2
x_{10000}
784
1000
1000
10
$\{ y_1^0 \ldots y_N^0 \}$
$\{ y_1^{(1)} \ldots y_N^{(1)} \}$
$\{ y_1^{(2)} \ldots y_N^{(2)} \}$
dim2
dim1
* - 3
x - 1
o - 2

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.

SVHN

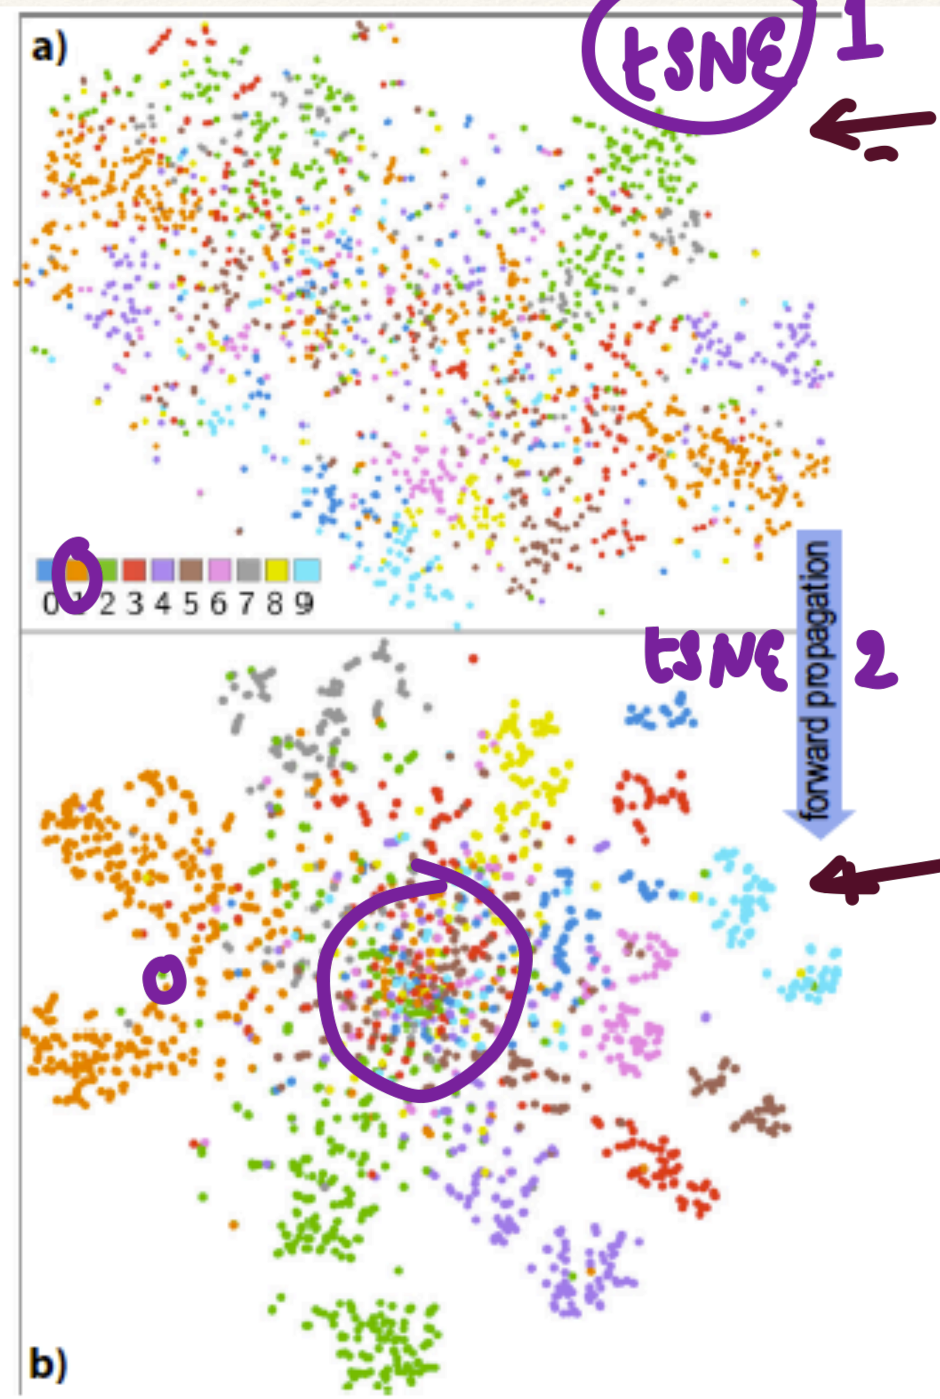

Fig. 5. Projection of the MLP hidden layer activations after training, SVHN test subset. a) First hidden layer (NH: 52.78%). b) Last hidden layer (NH: 67%).

CIFAR-10

flatten
CNN layer
↳ tSNE

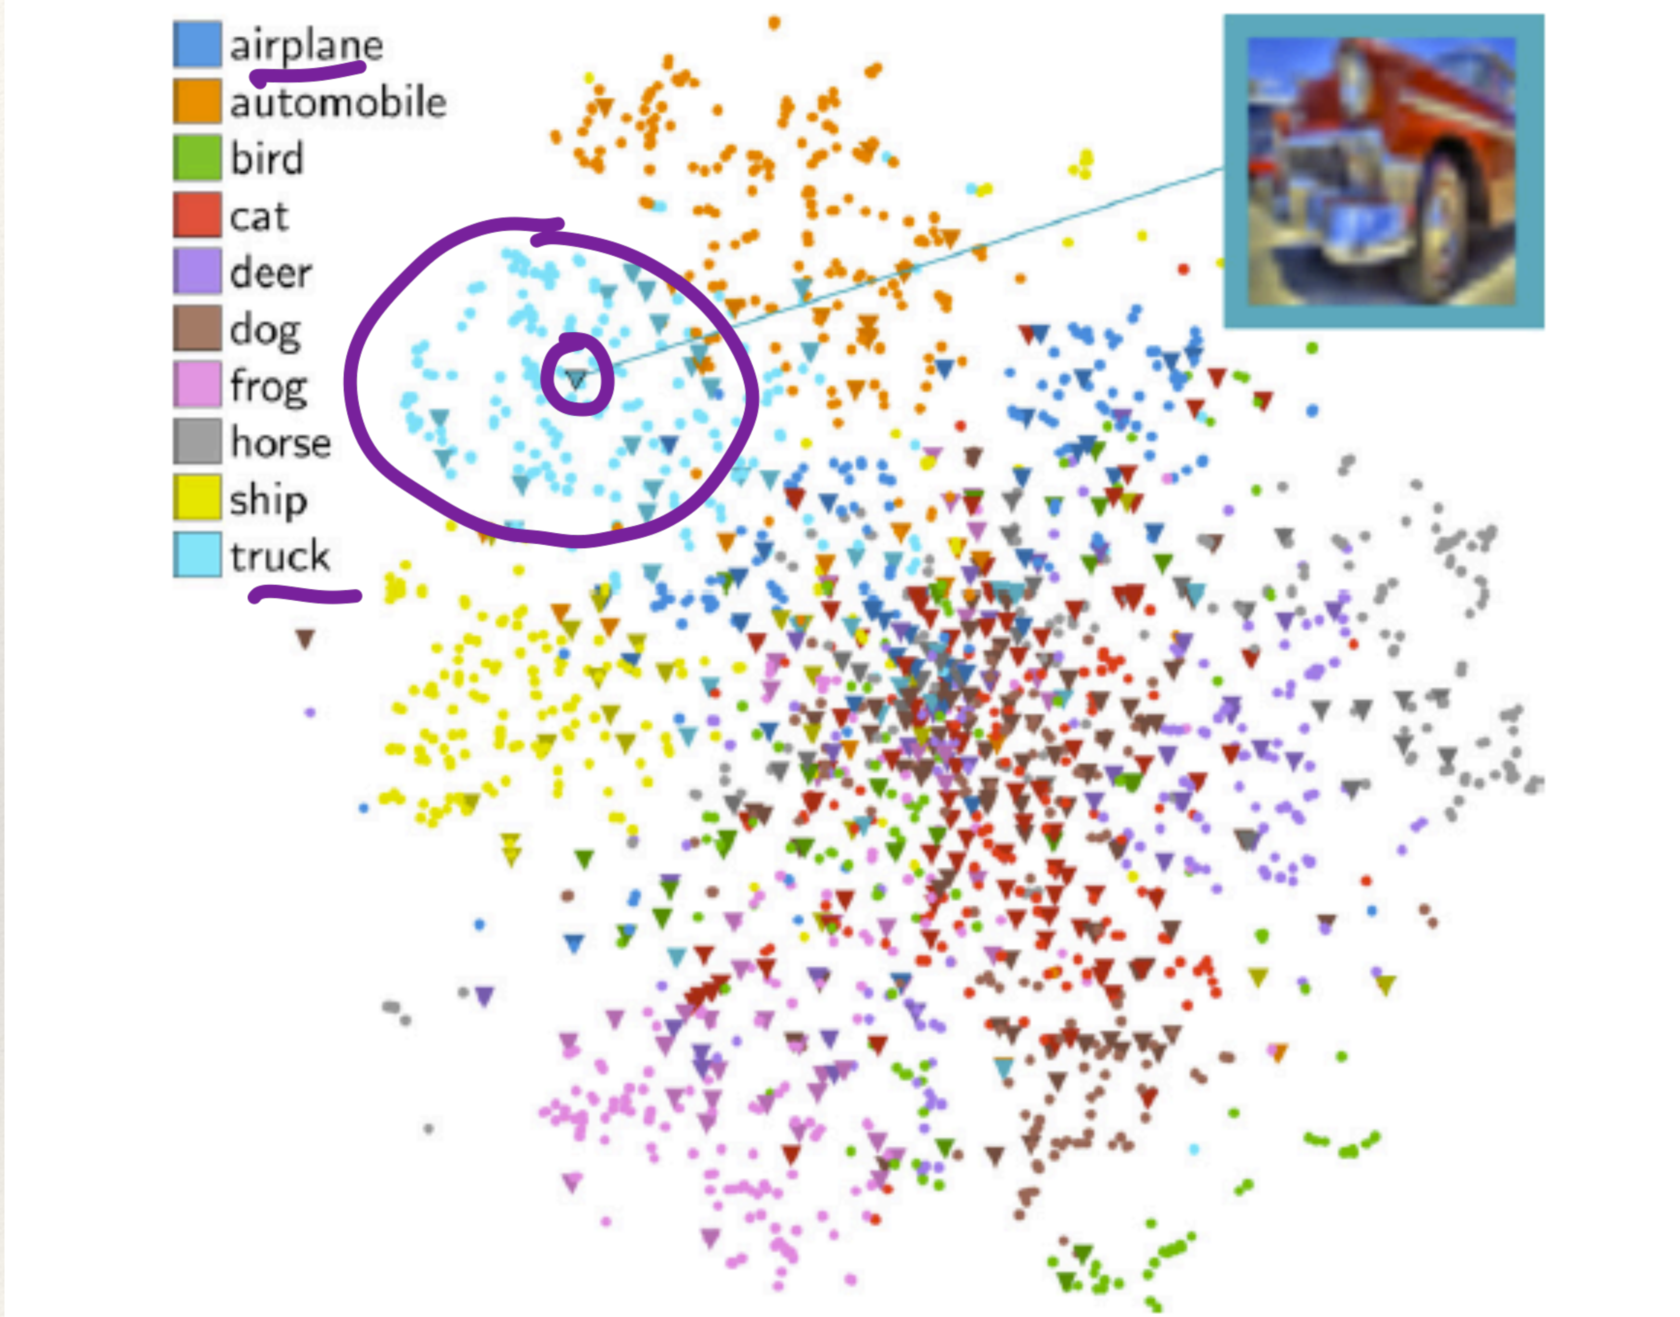

Fig. 9. Projection of last CNN hidden layer activations after training, *CIFAR*-10 test subset (NH: 53.43%, AC: 78.7%).

Understanding Deep Networks

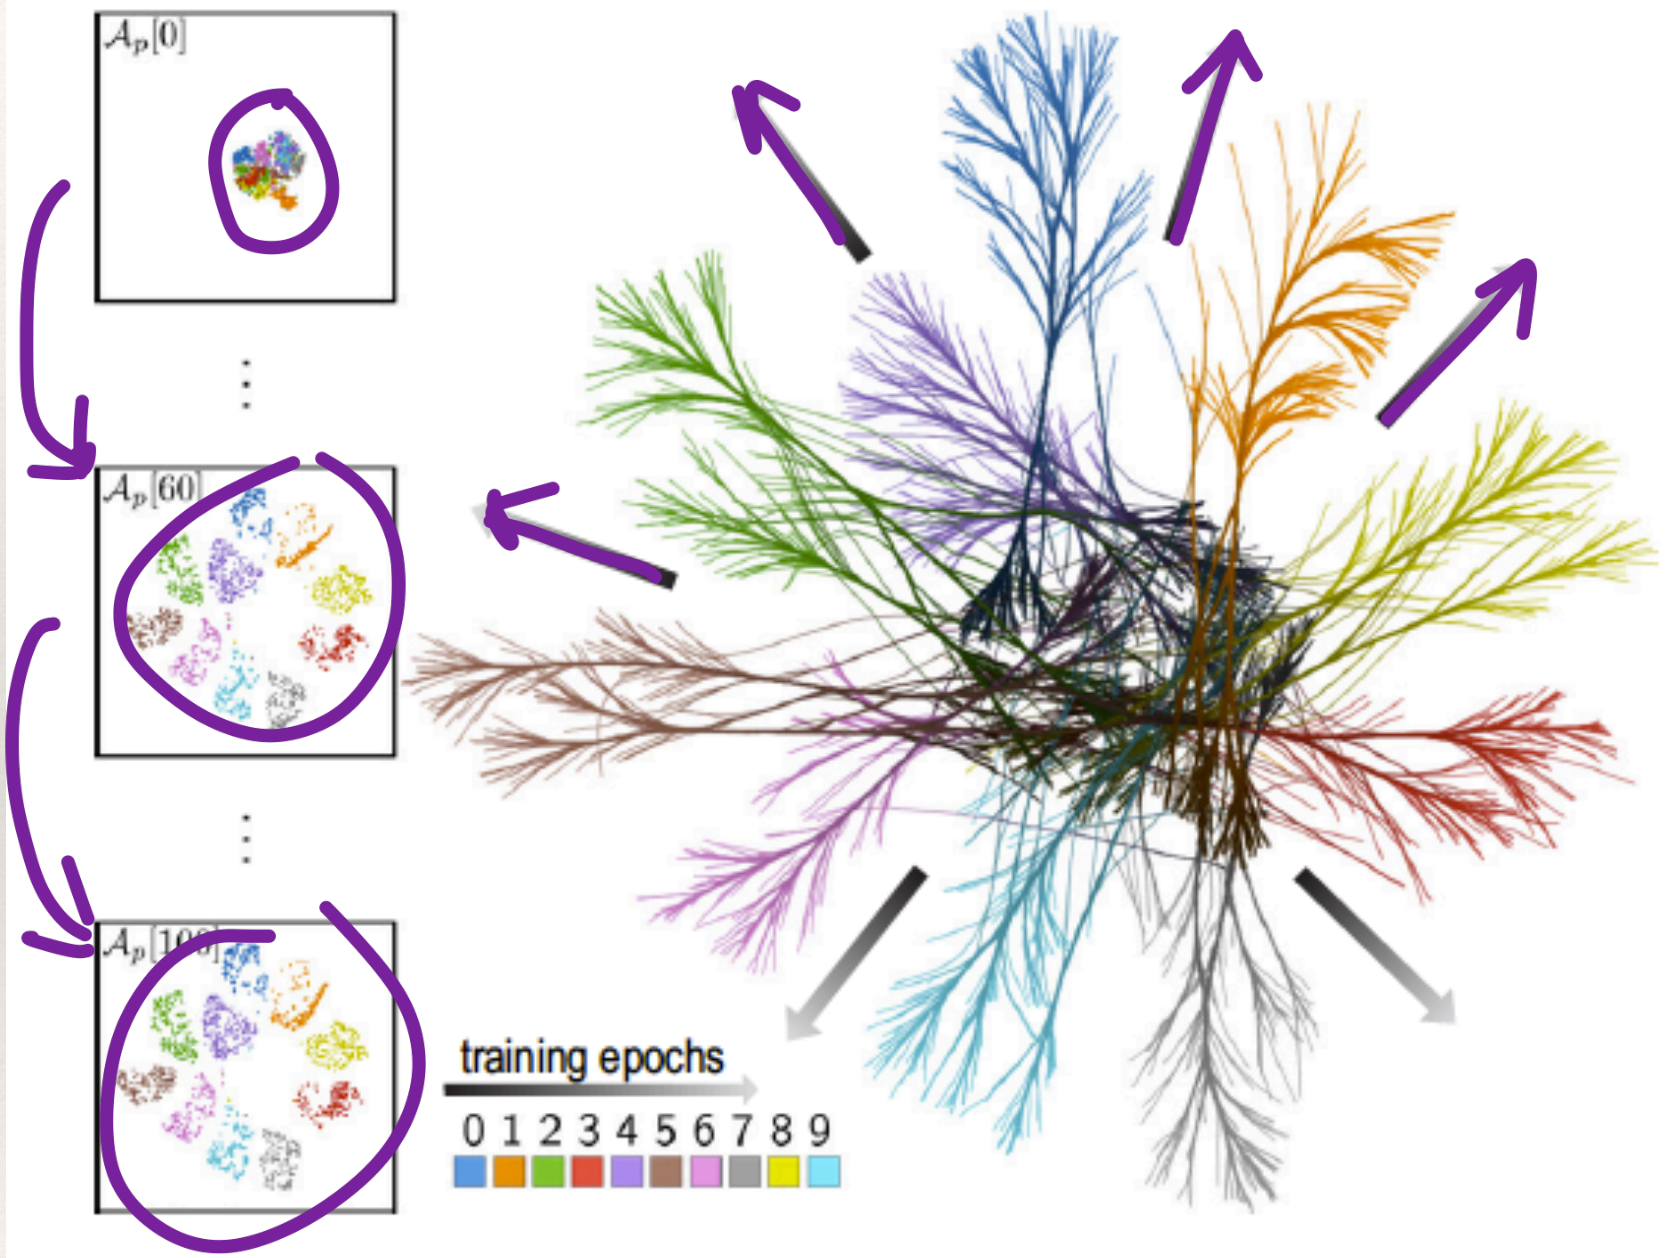

Fig. 11. Inter-epoch evolution, last CNN hidden layer, epochs 0-100, in steps of 20, *MNIST* test subset. Brighter trail parts show later epochs.

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

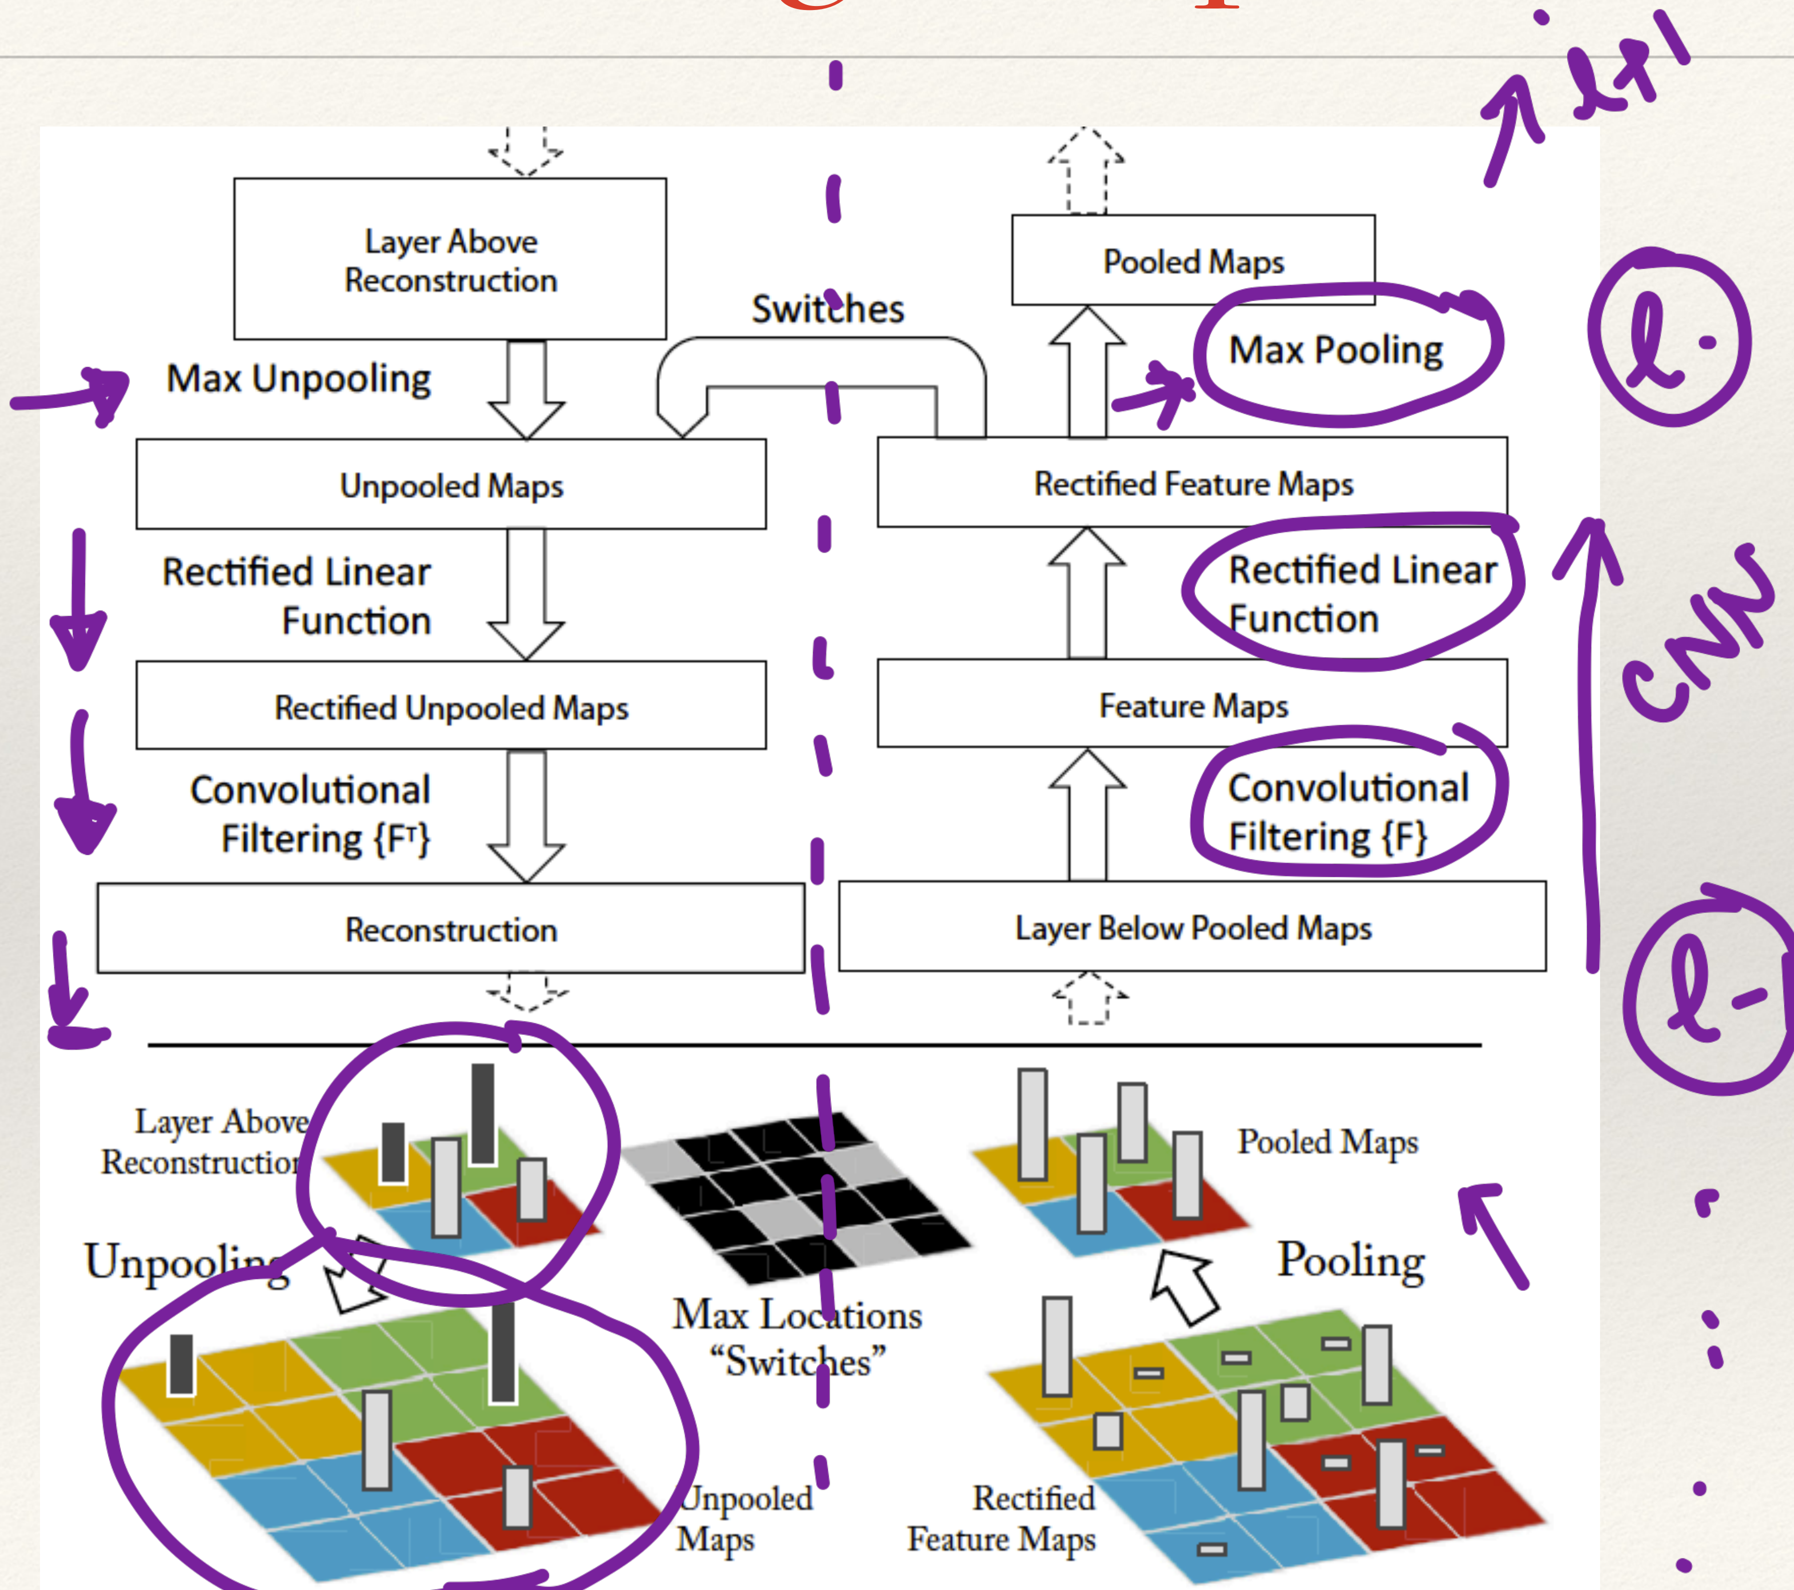

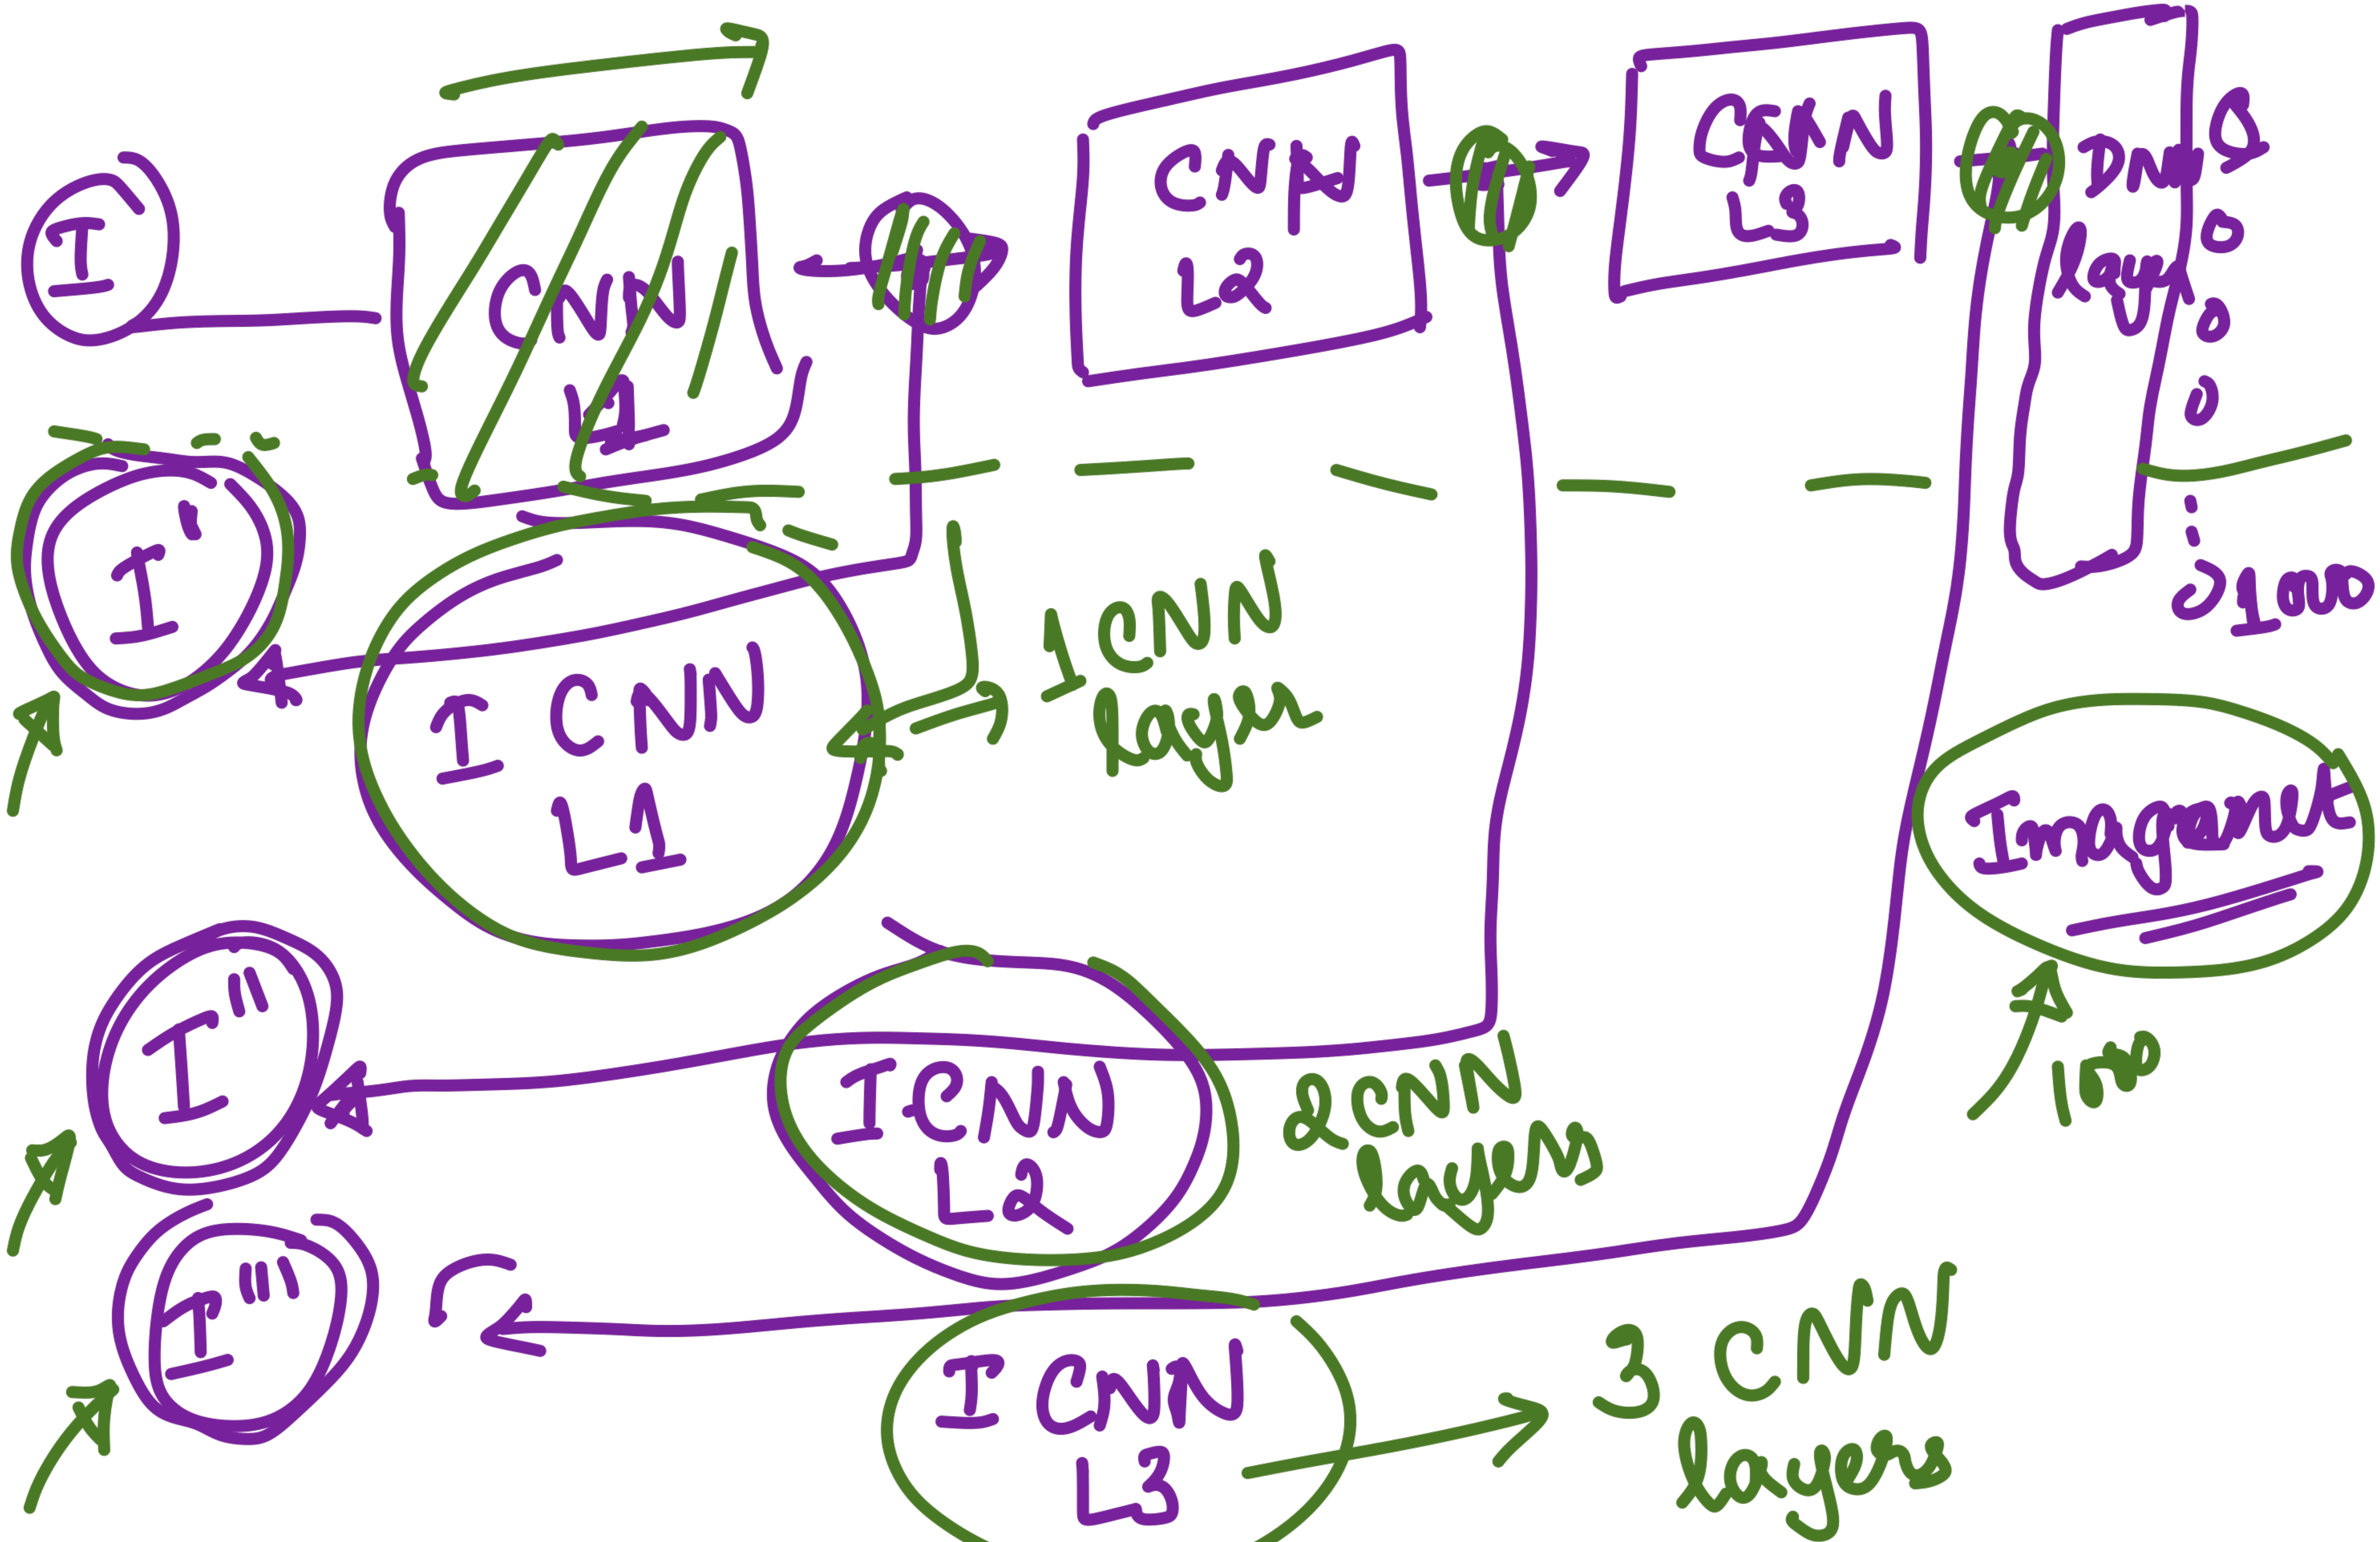
I
CNN L1
CNN L2
CNN L3
DNN layer
1000
I'
I CNN L1
1 CNN layer
I''
I CNN L2
2 CNN layers
I'''
I CNN L3
3 CNN layers
Imagenet
1000

Understanding Deep Networks

Recon

original image.

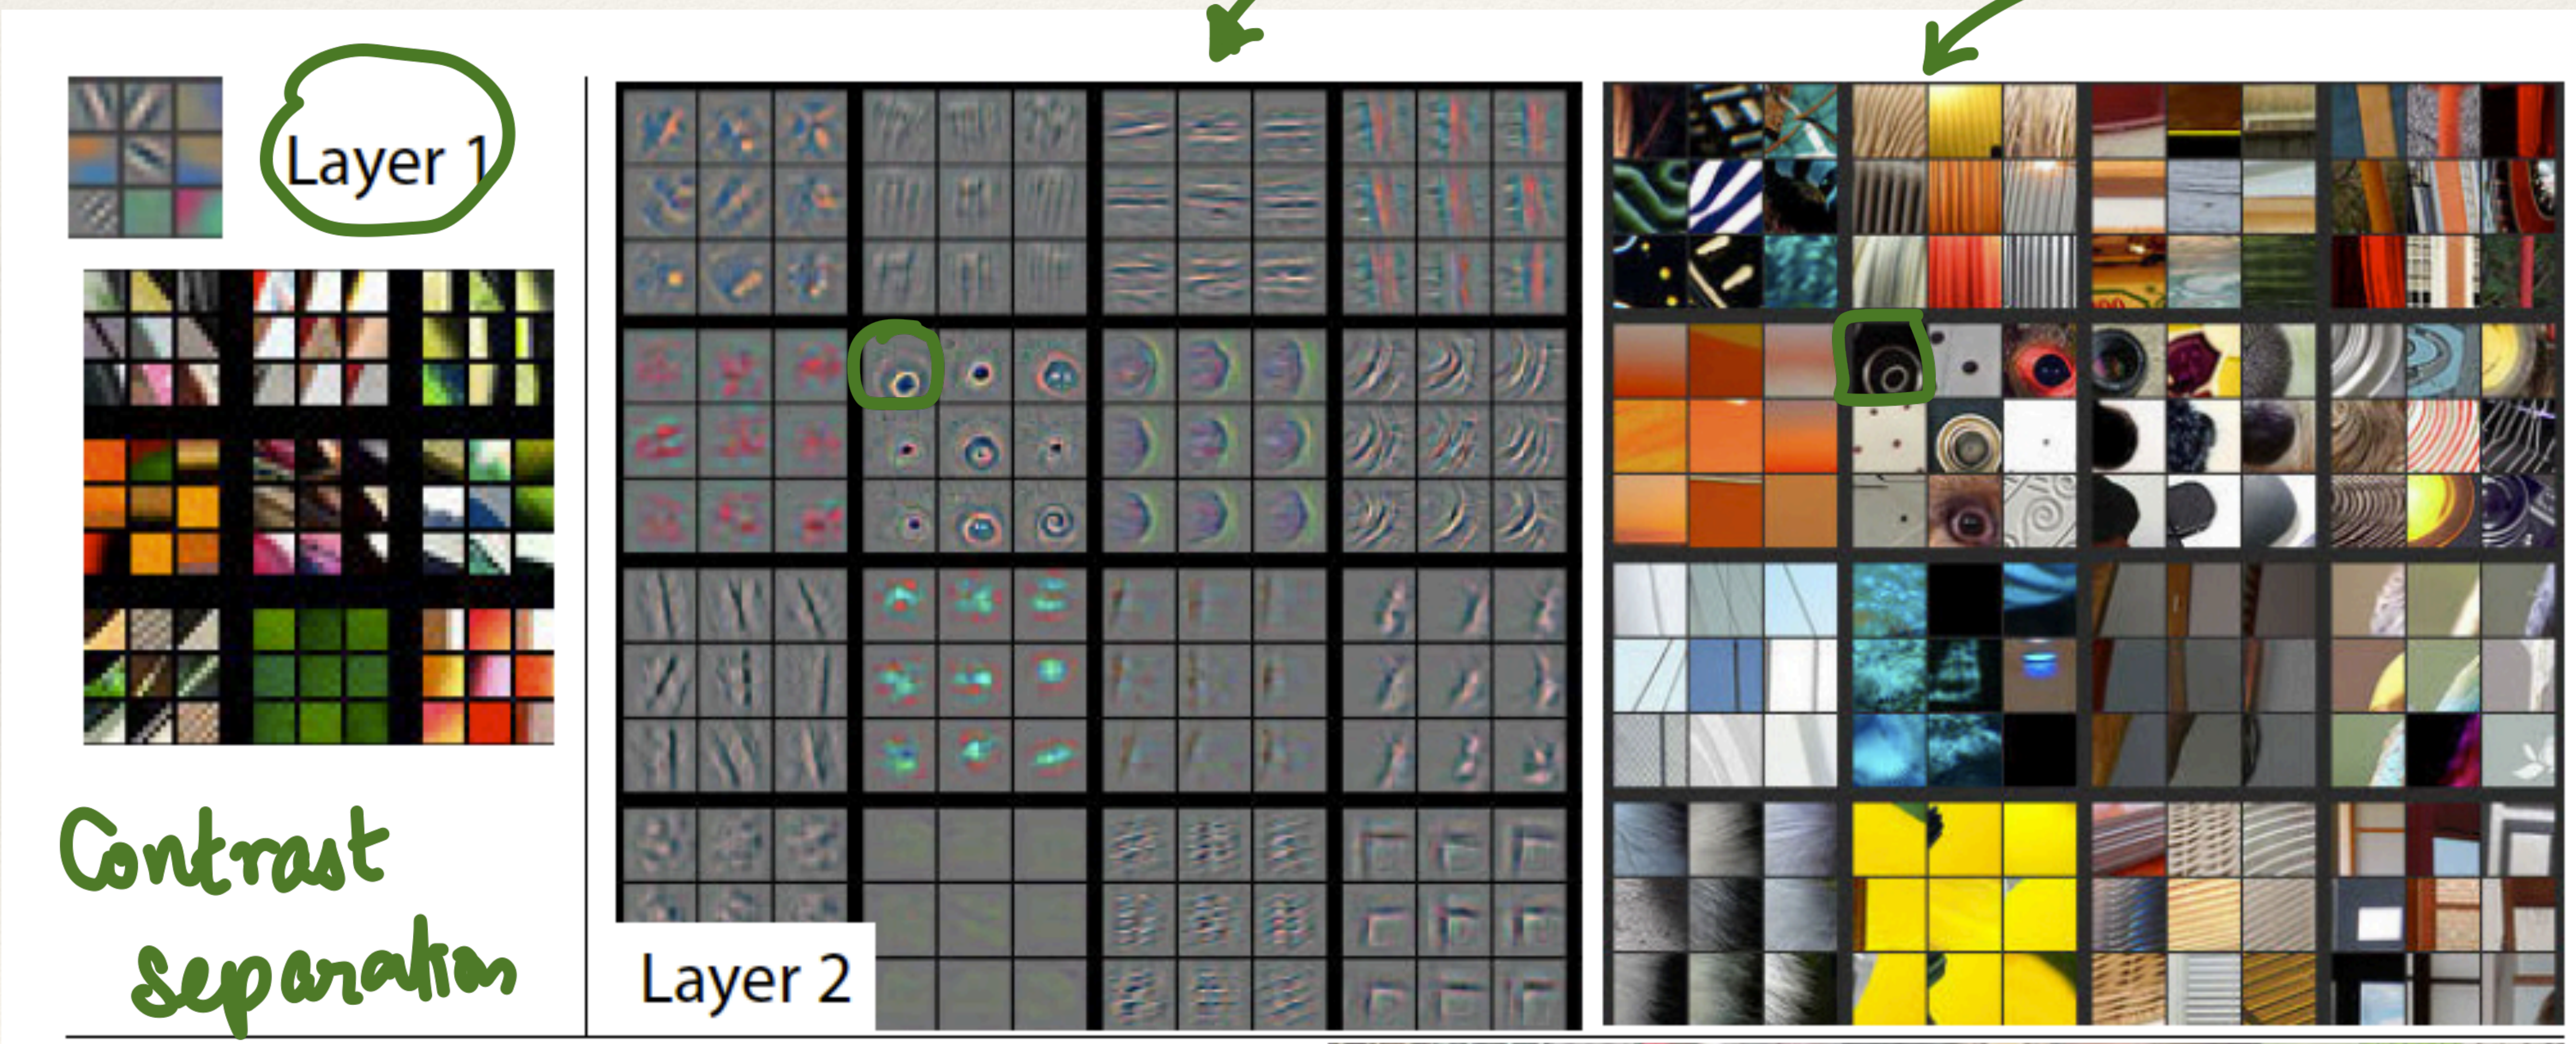

Understanding Deep Networks

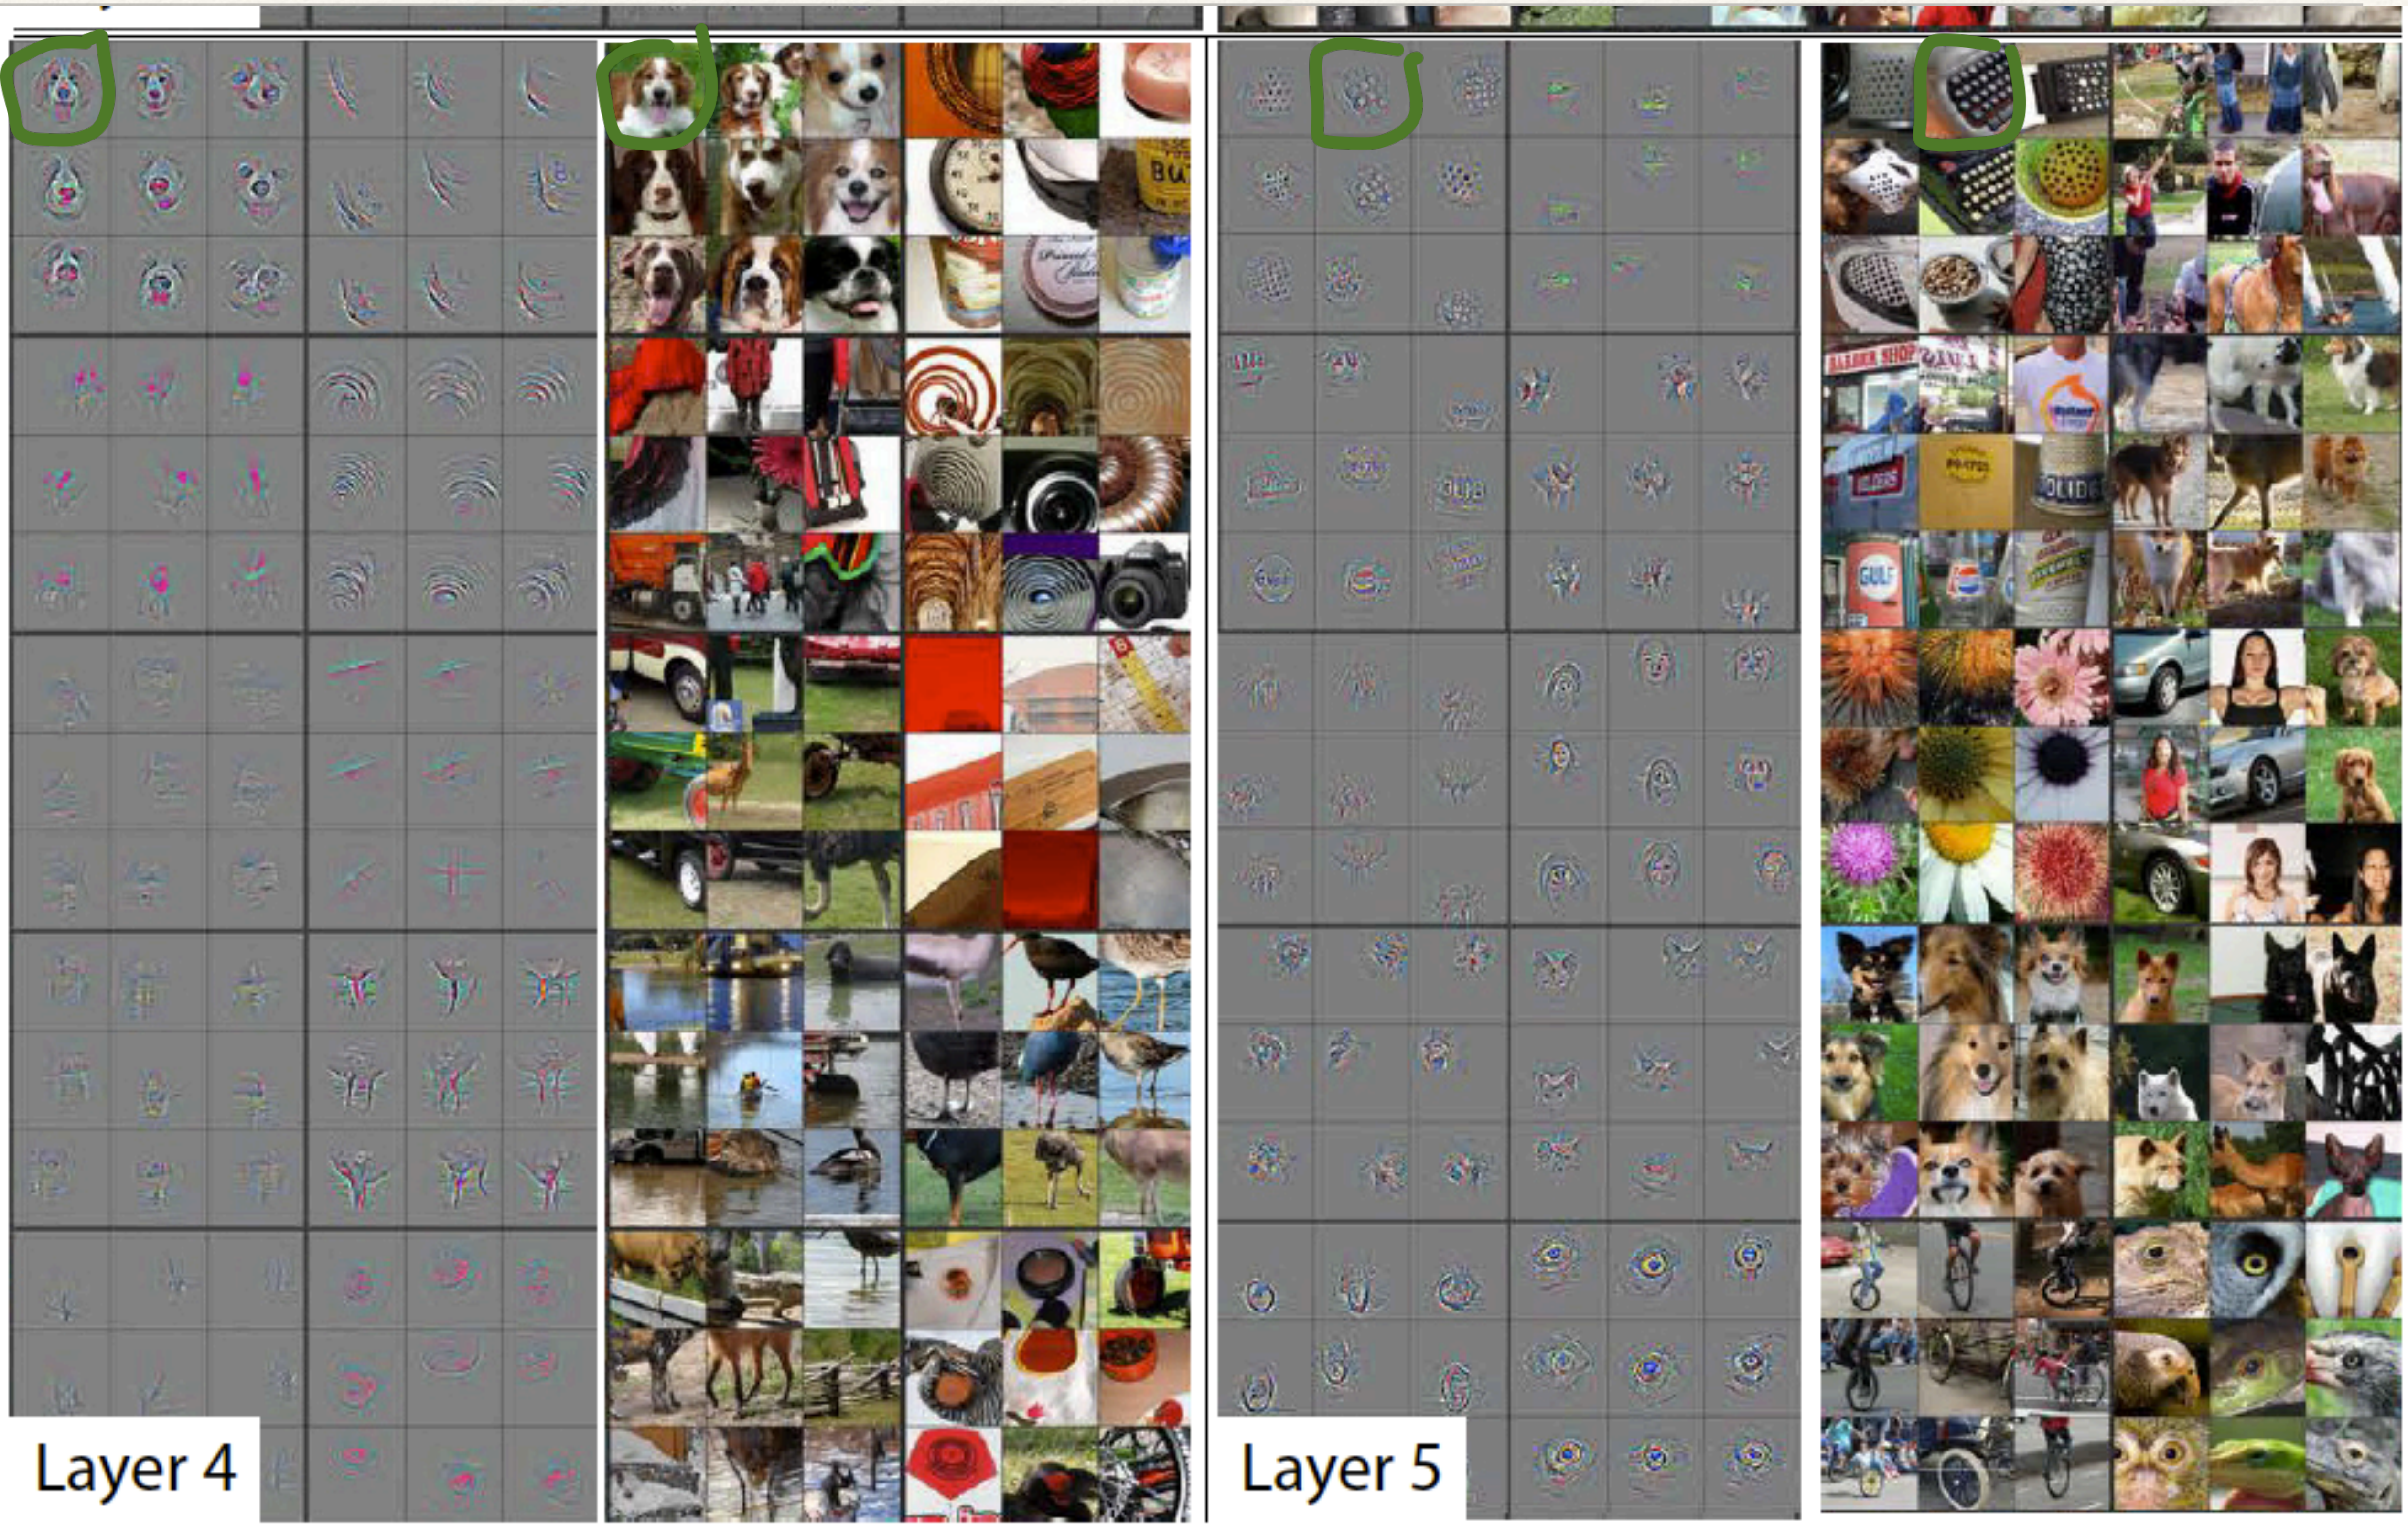

Representation Learning in Deep Networks

[Zeiler, 2014]

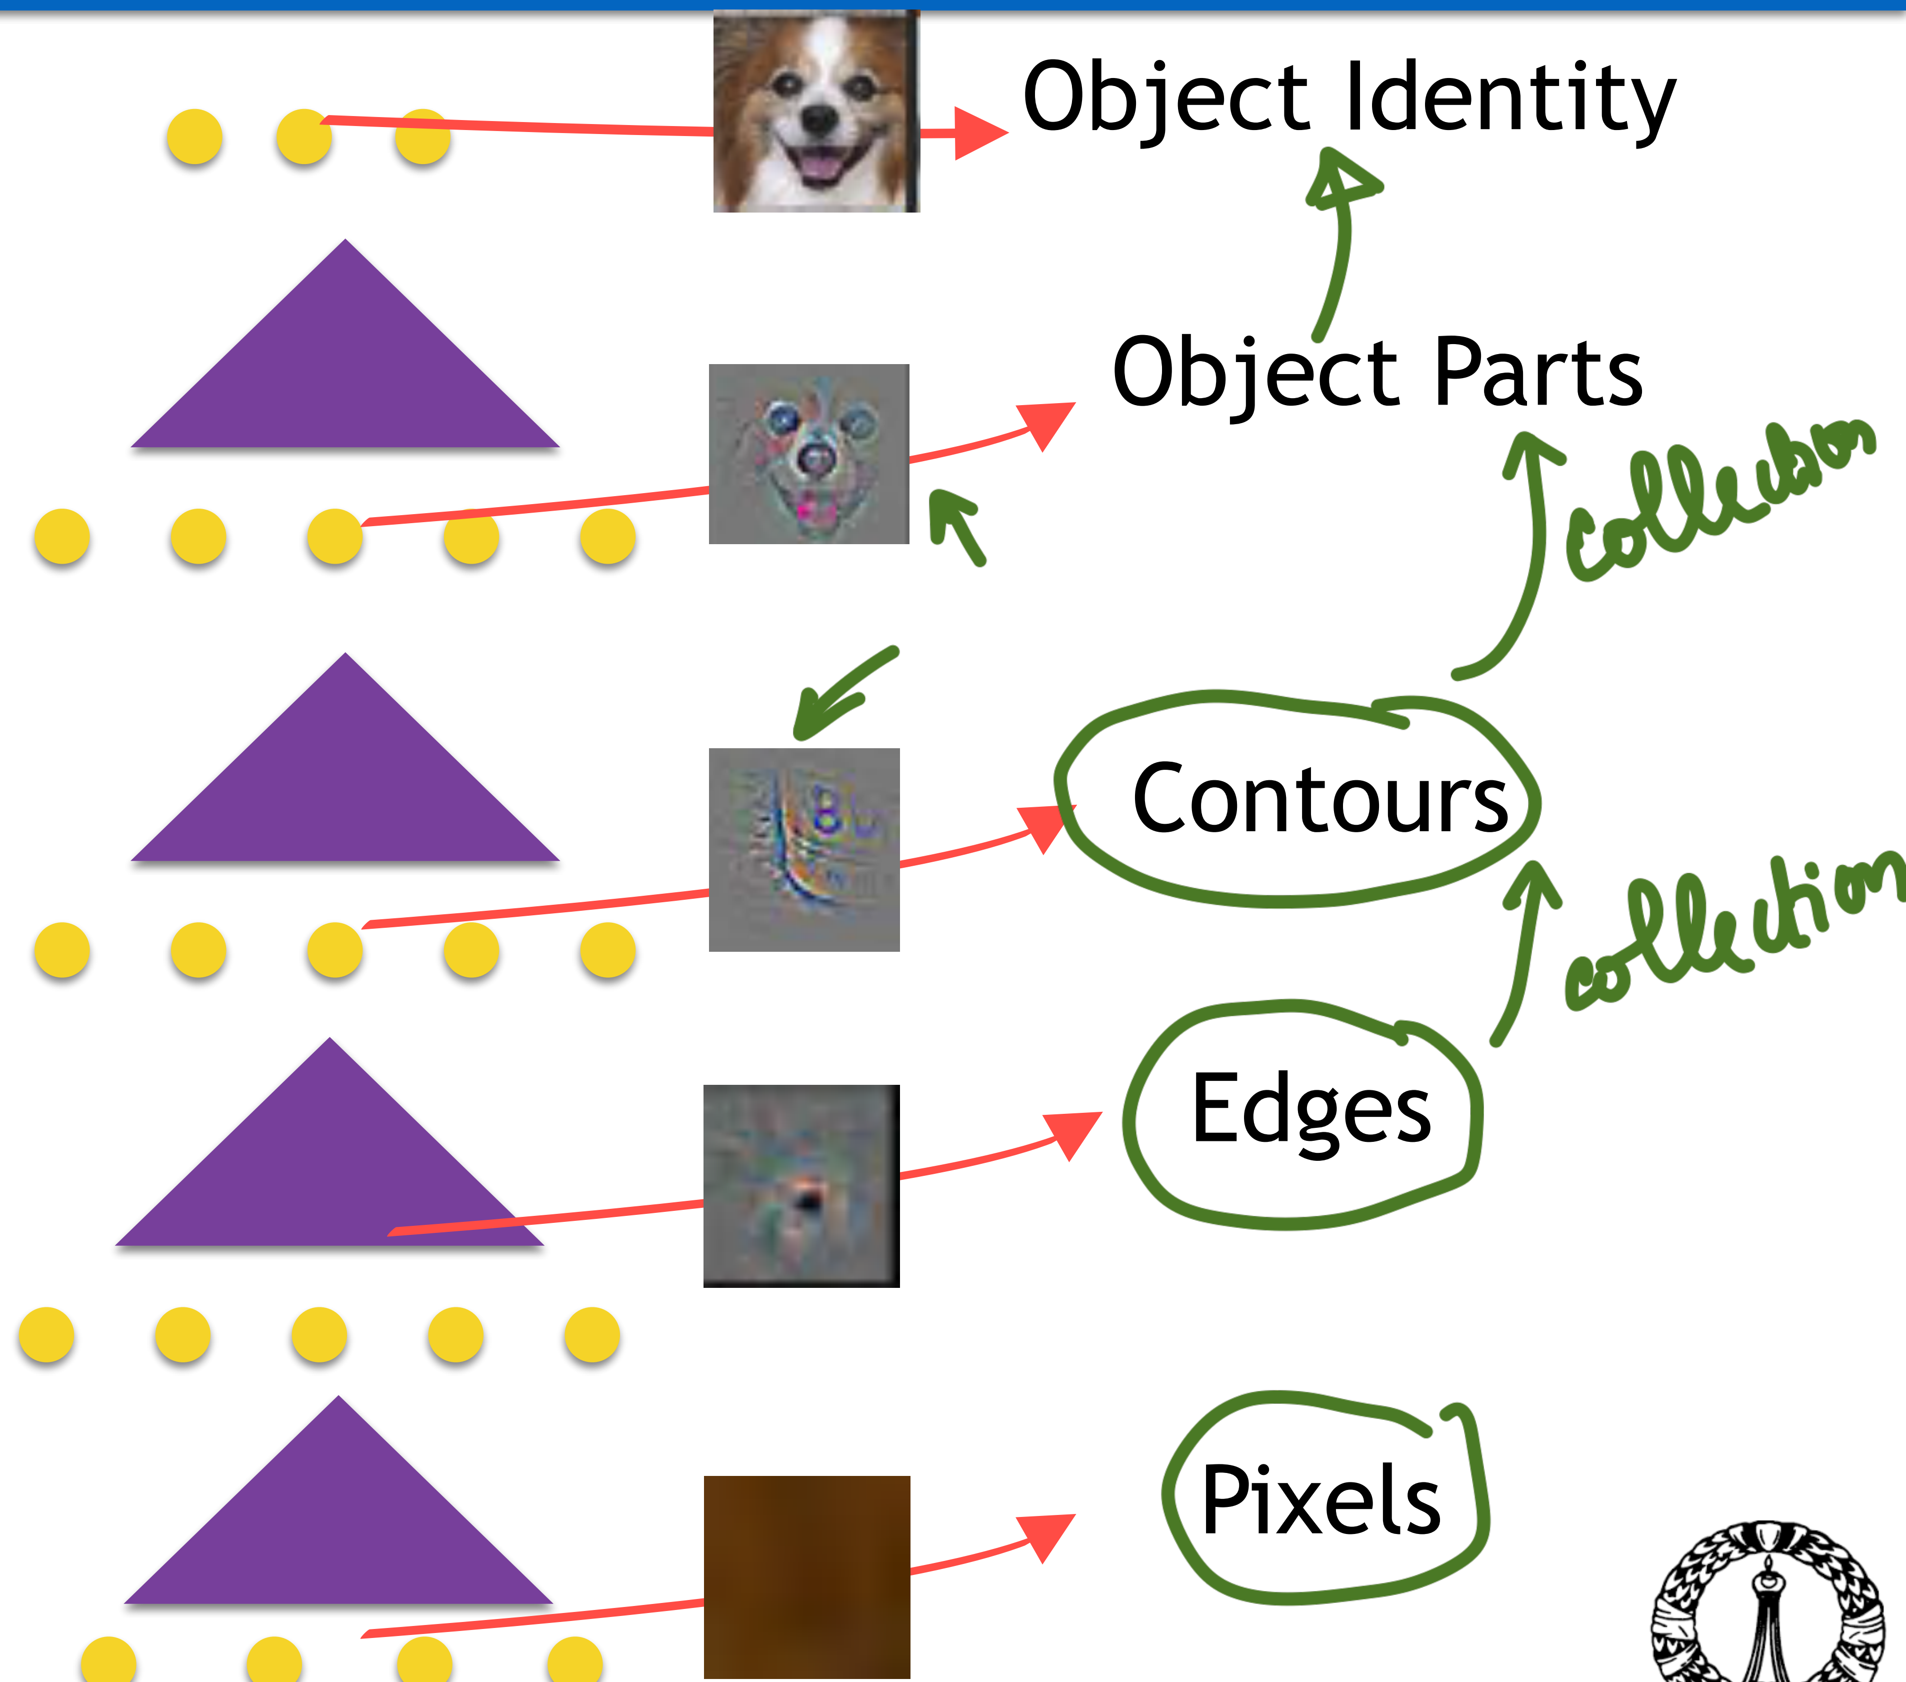

UNDERSTANDING HOW DEEP BELIEF NETWORKS PERFORM ACOUSTIC MODELLING

Garcia-Romero, Daniel, et al. "Speaker diarization using deep neural network embeddings." *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.

Department of Computer Science, University of Toronto

Speech Recognition

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

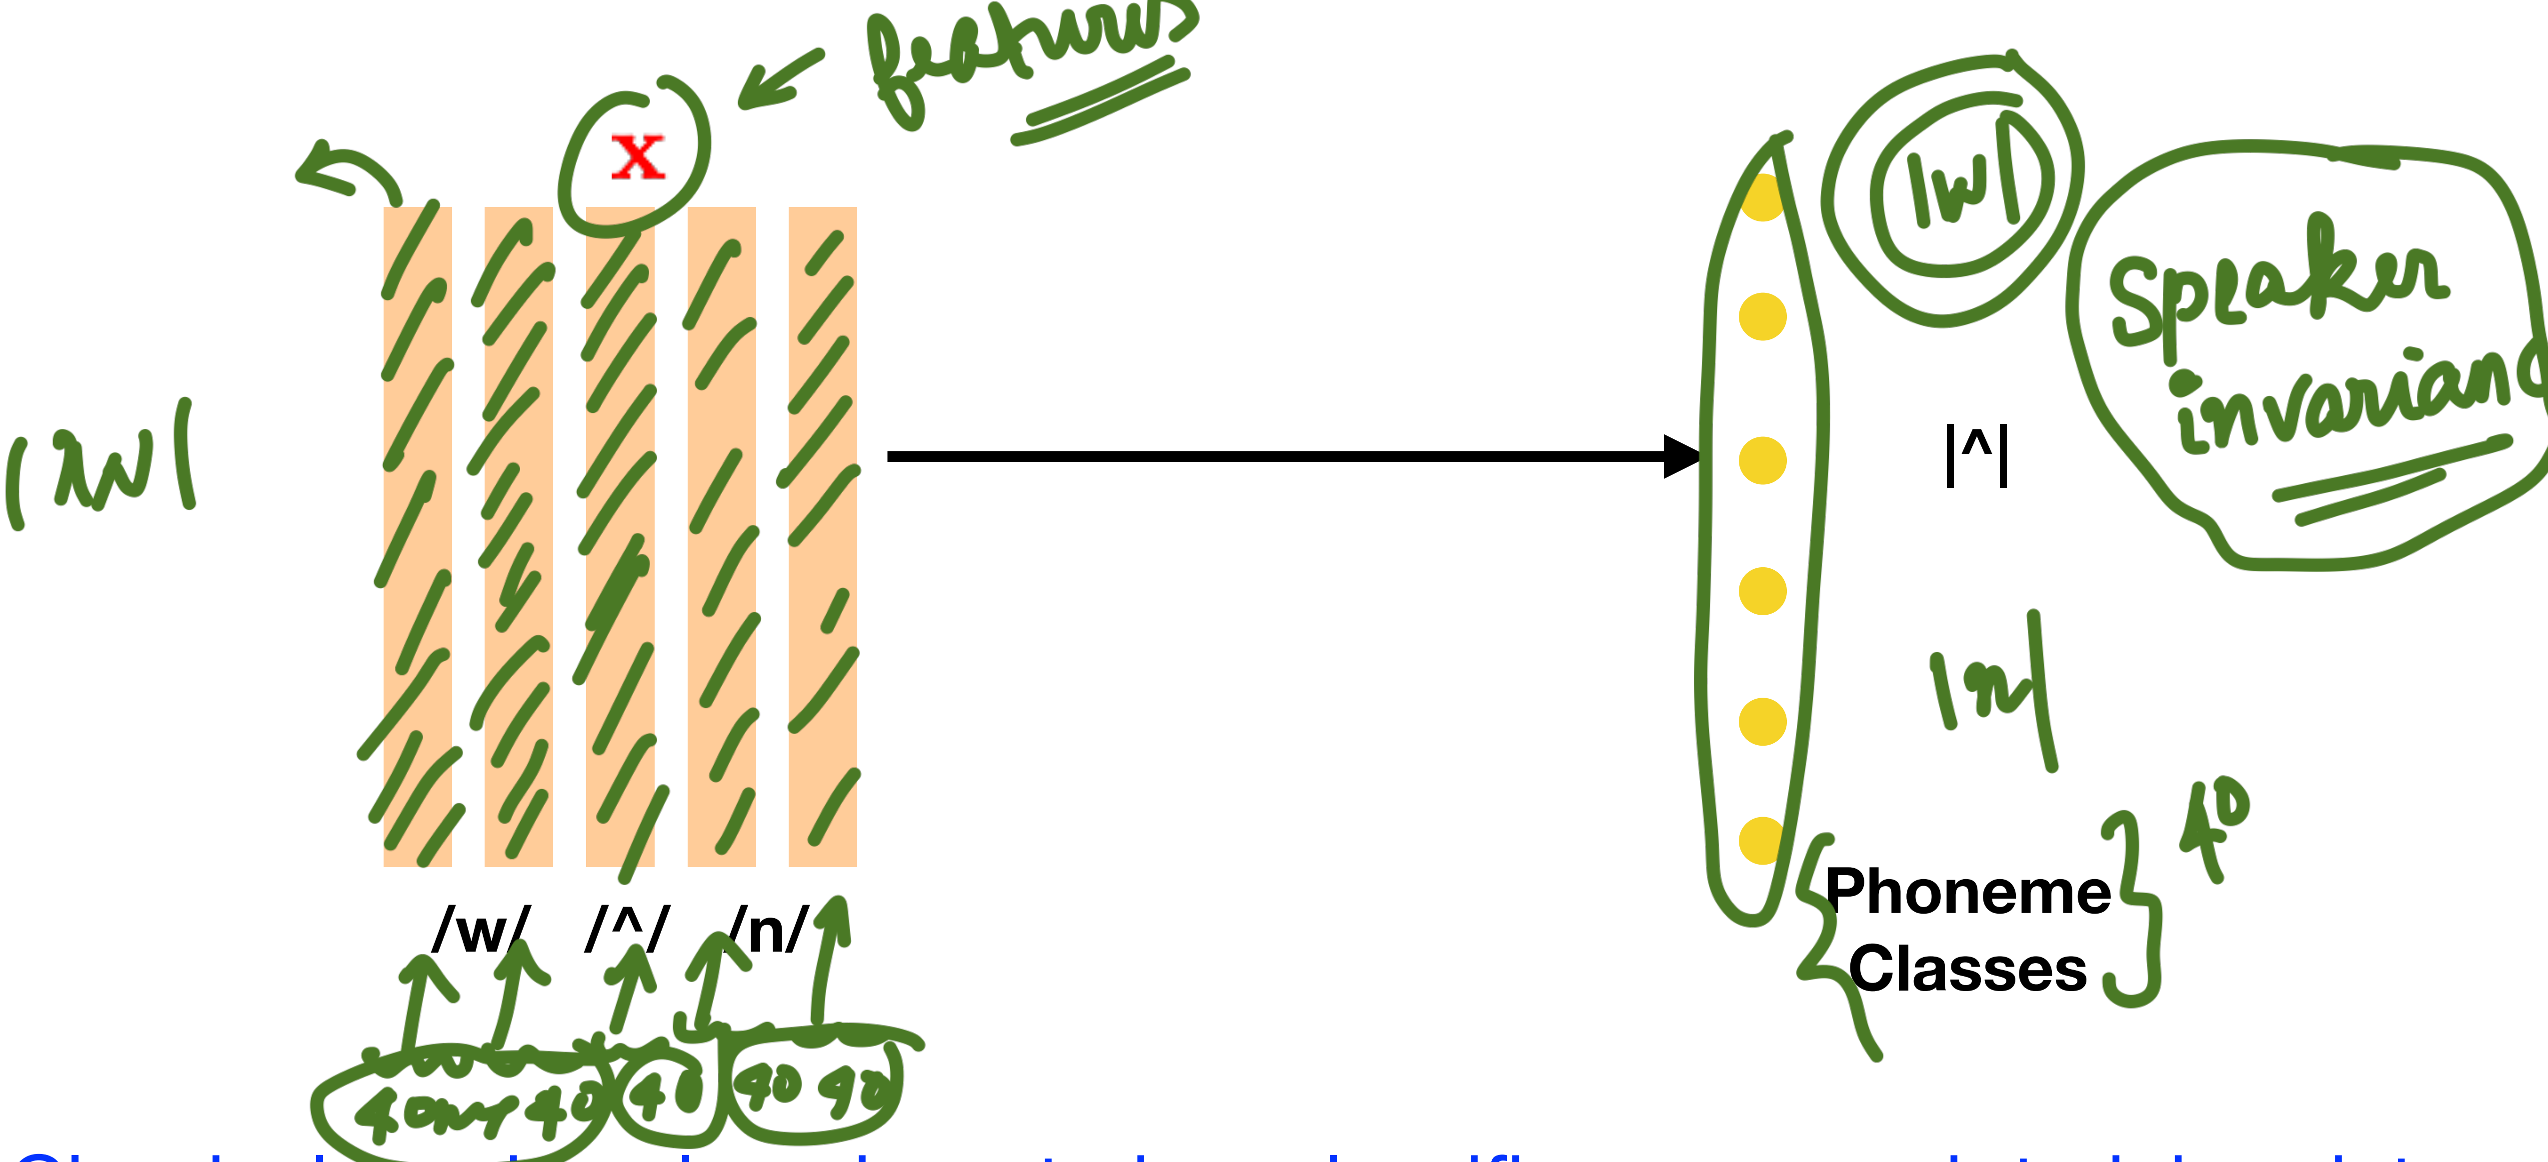

- Classical machine learning - train a classifier on speech training data that maps to the target phoneme class.

Understanding DNNs for Speech

2-D projection of 1st layer DNN

[Abdel Rahman, 2012]

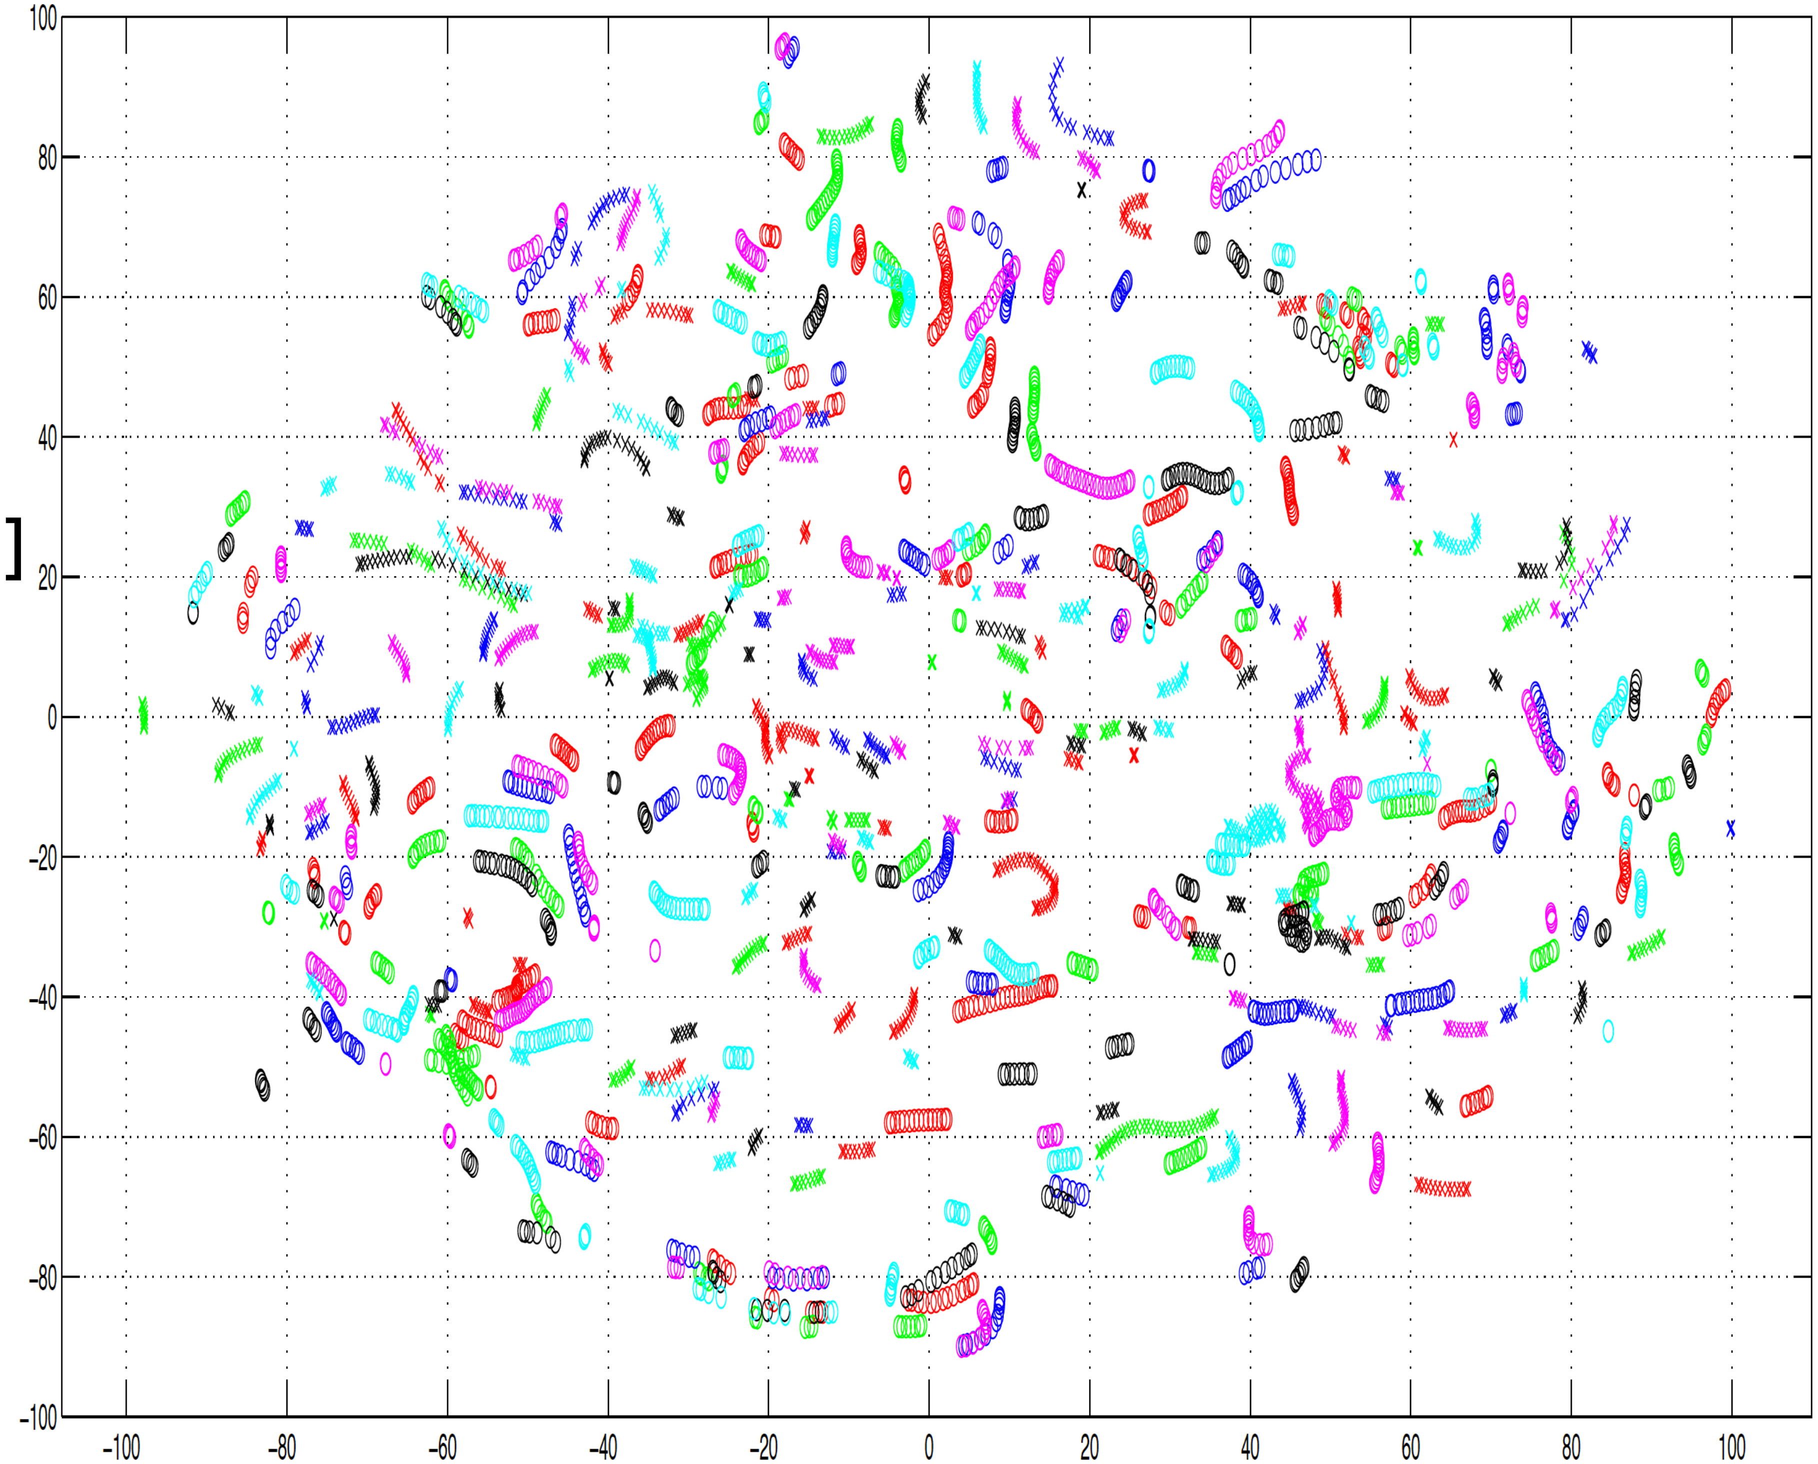

Established 1911

Understanding DNNs for Speech

[Abdel Rahman, 2012]

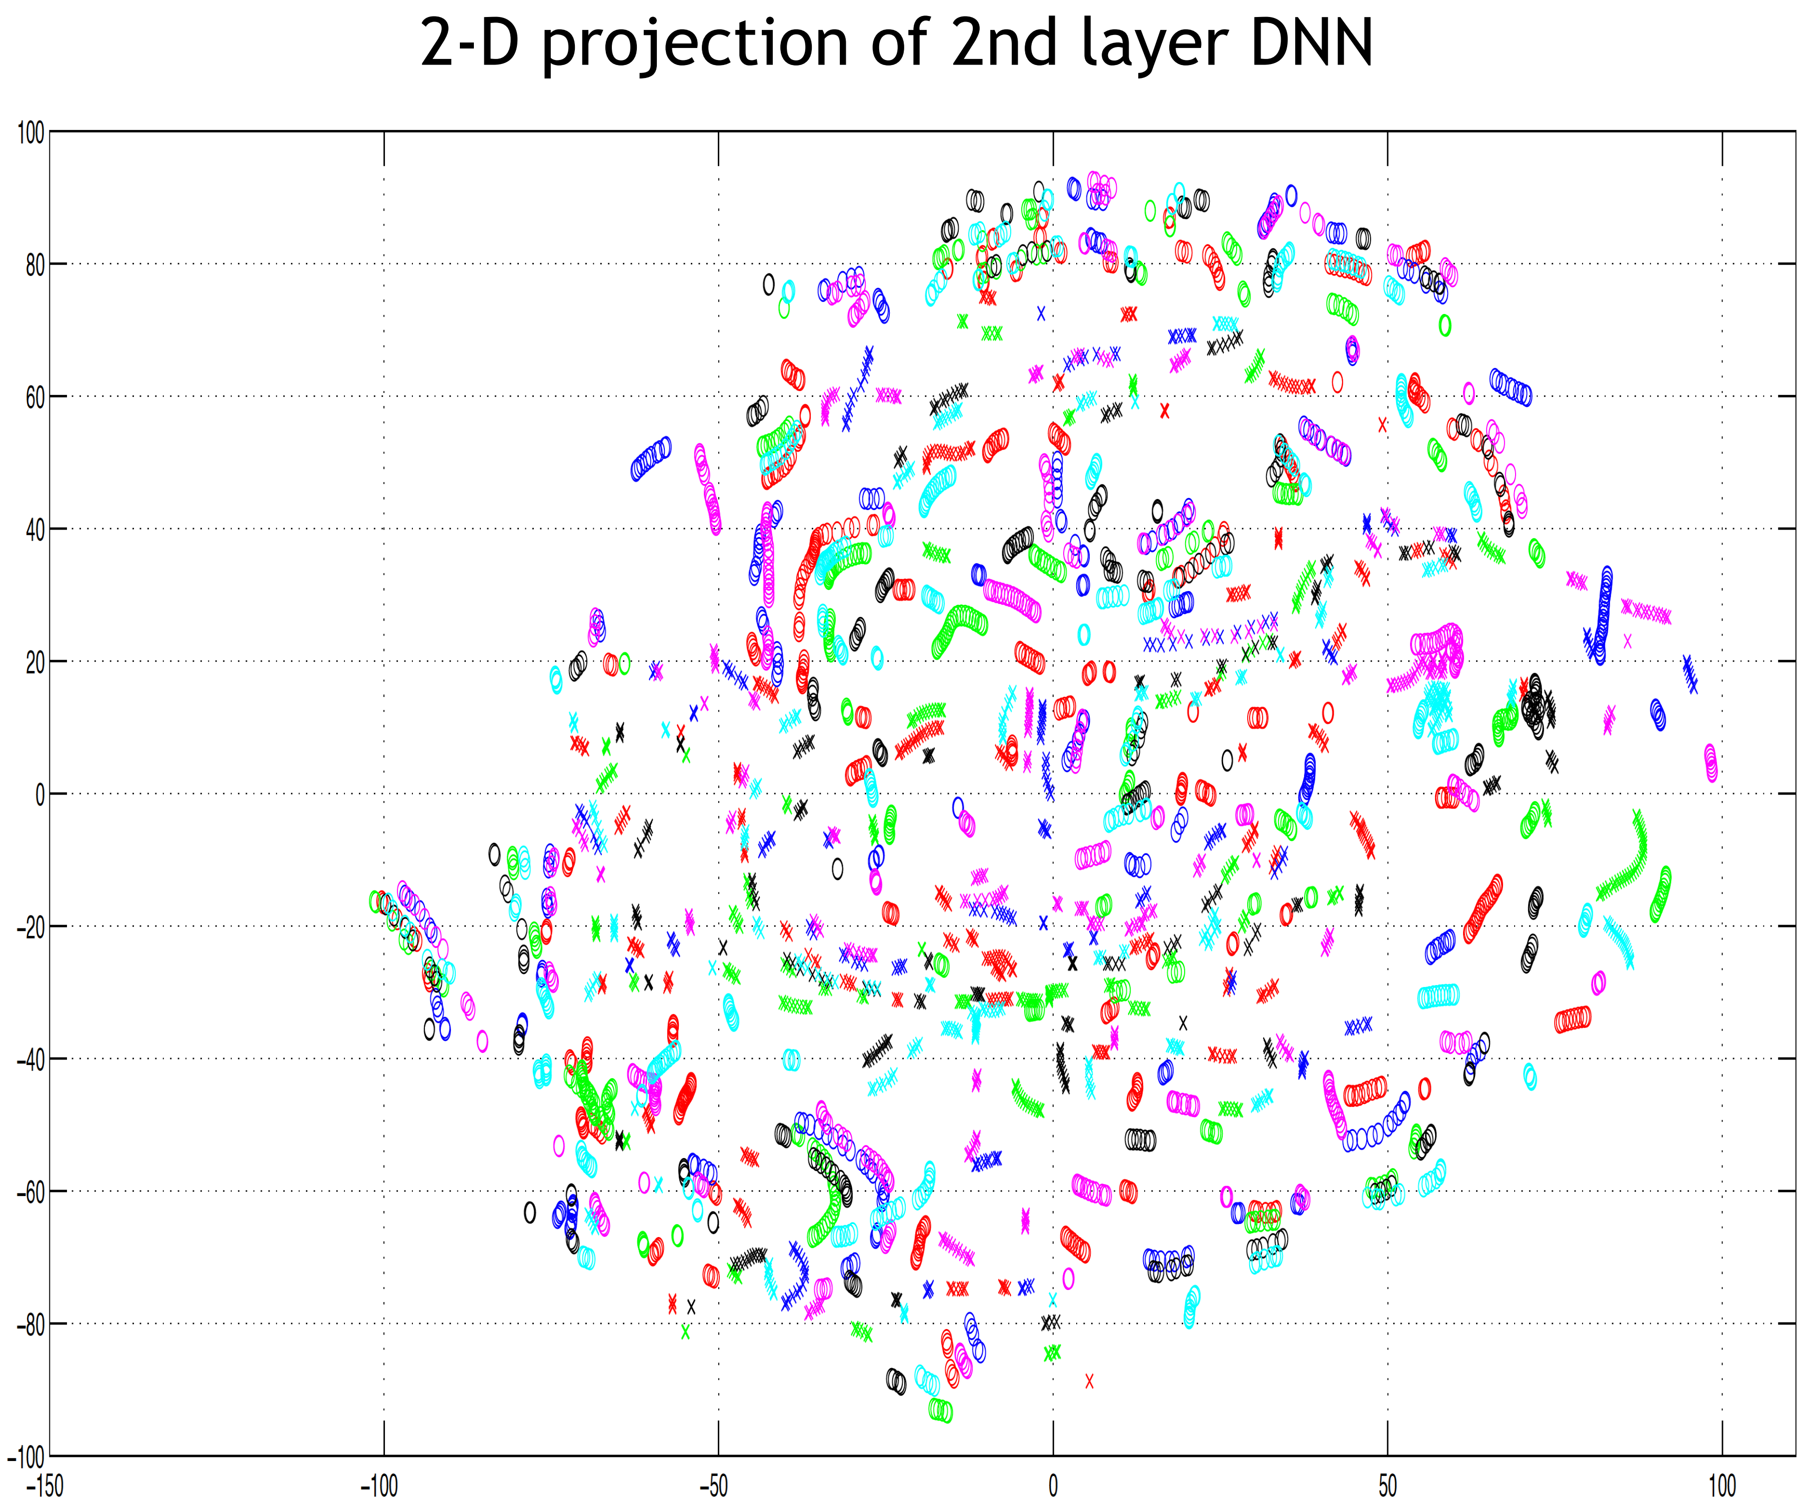

Established
1911

Understanding DNNs for Speech

Cat drink milk

/k/

[Abdel Rahman, 2012]

/k/ /a/ /t/

100ms

200ms

/k/ /a/ /t/

2-D projection of final layer DNN

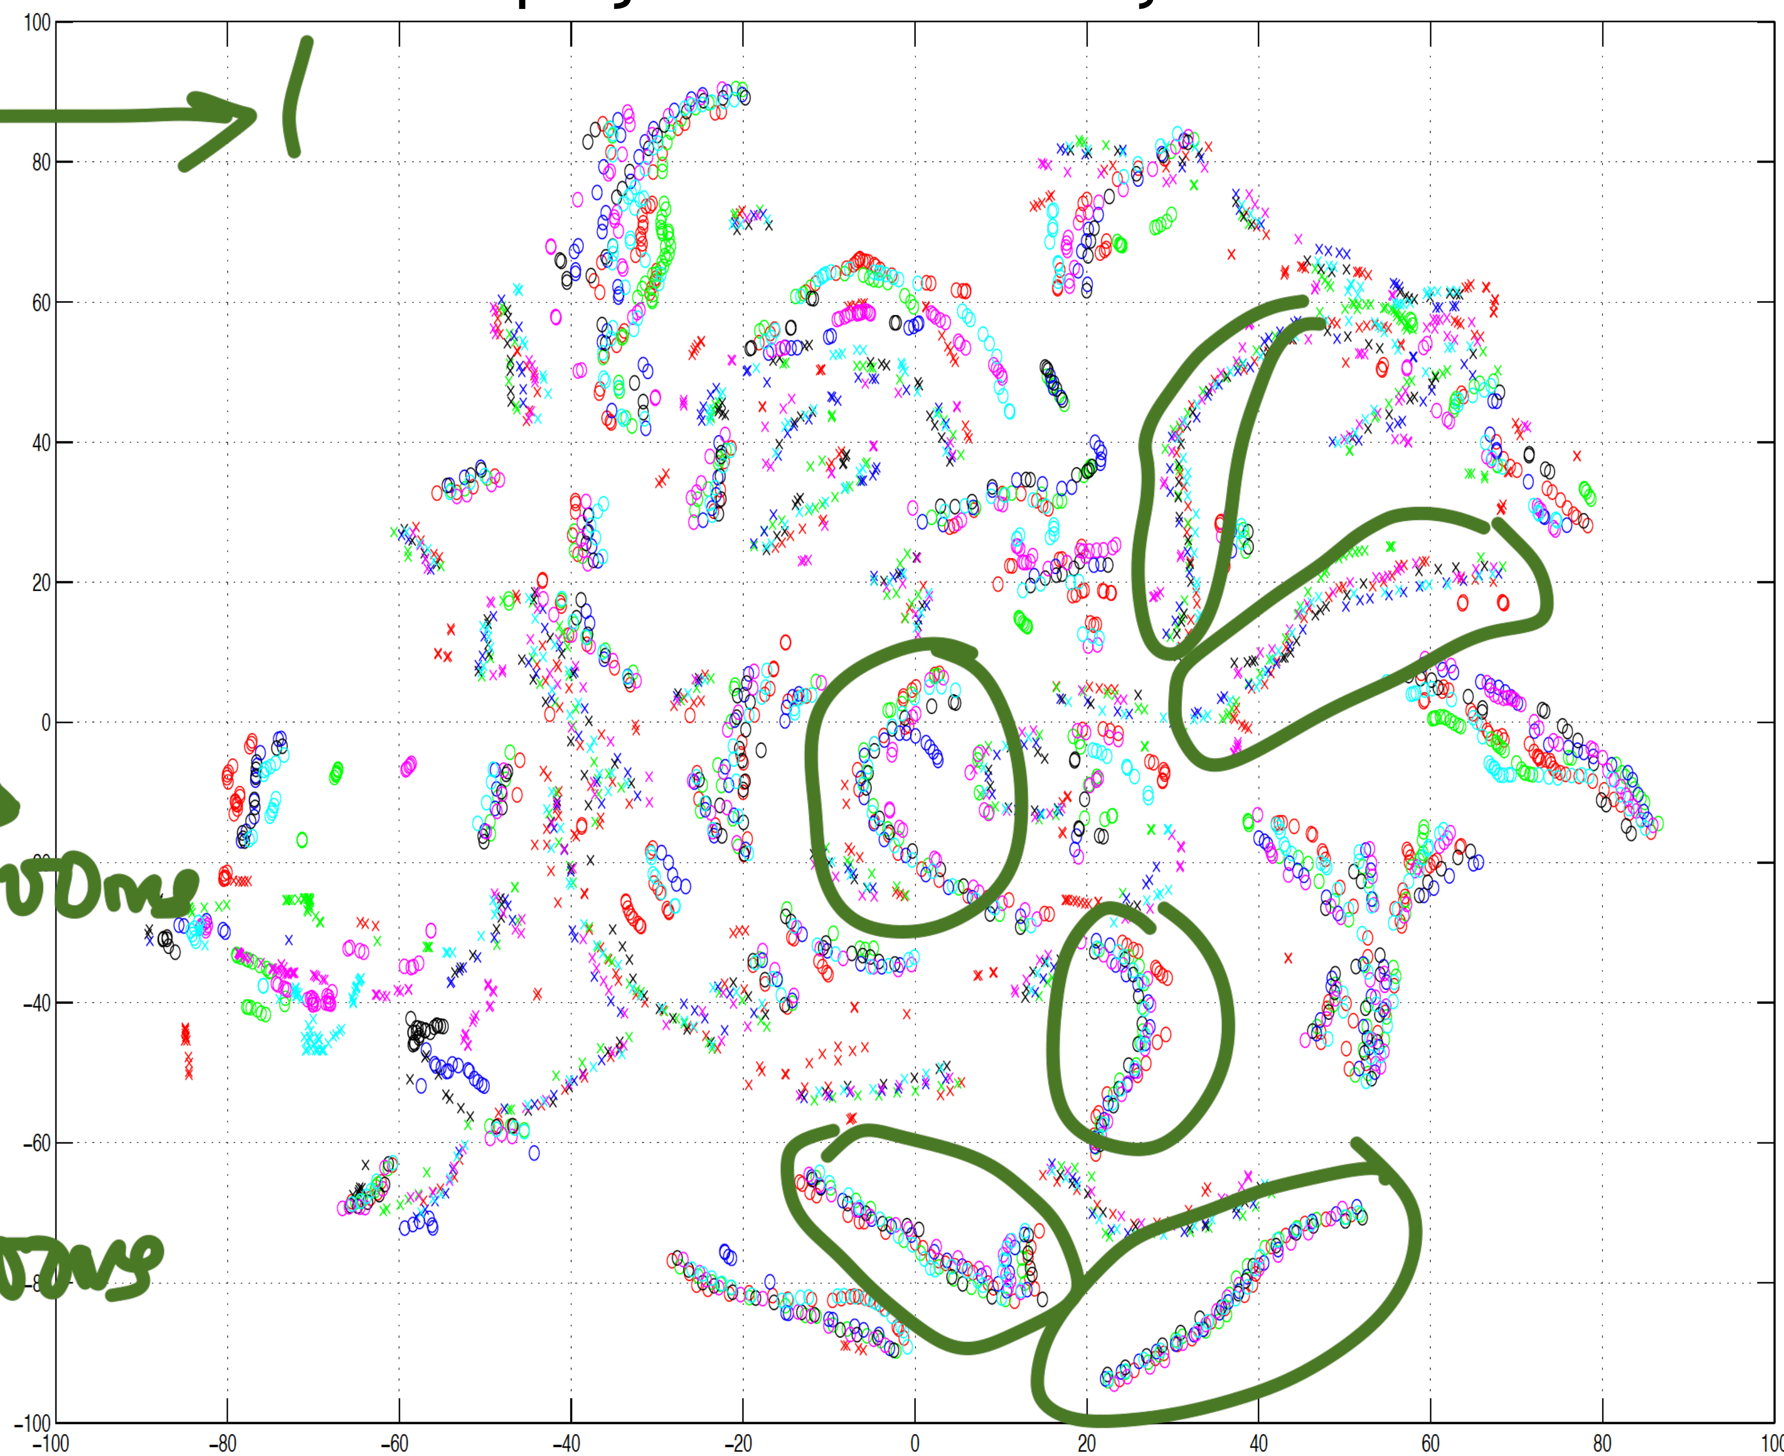

Established 1911

Understanding DNNs for Speech

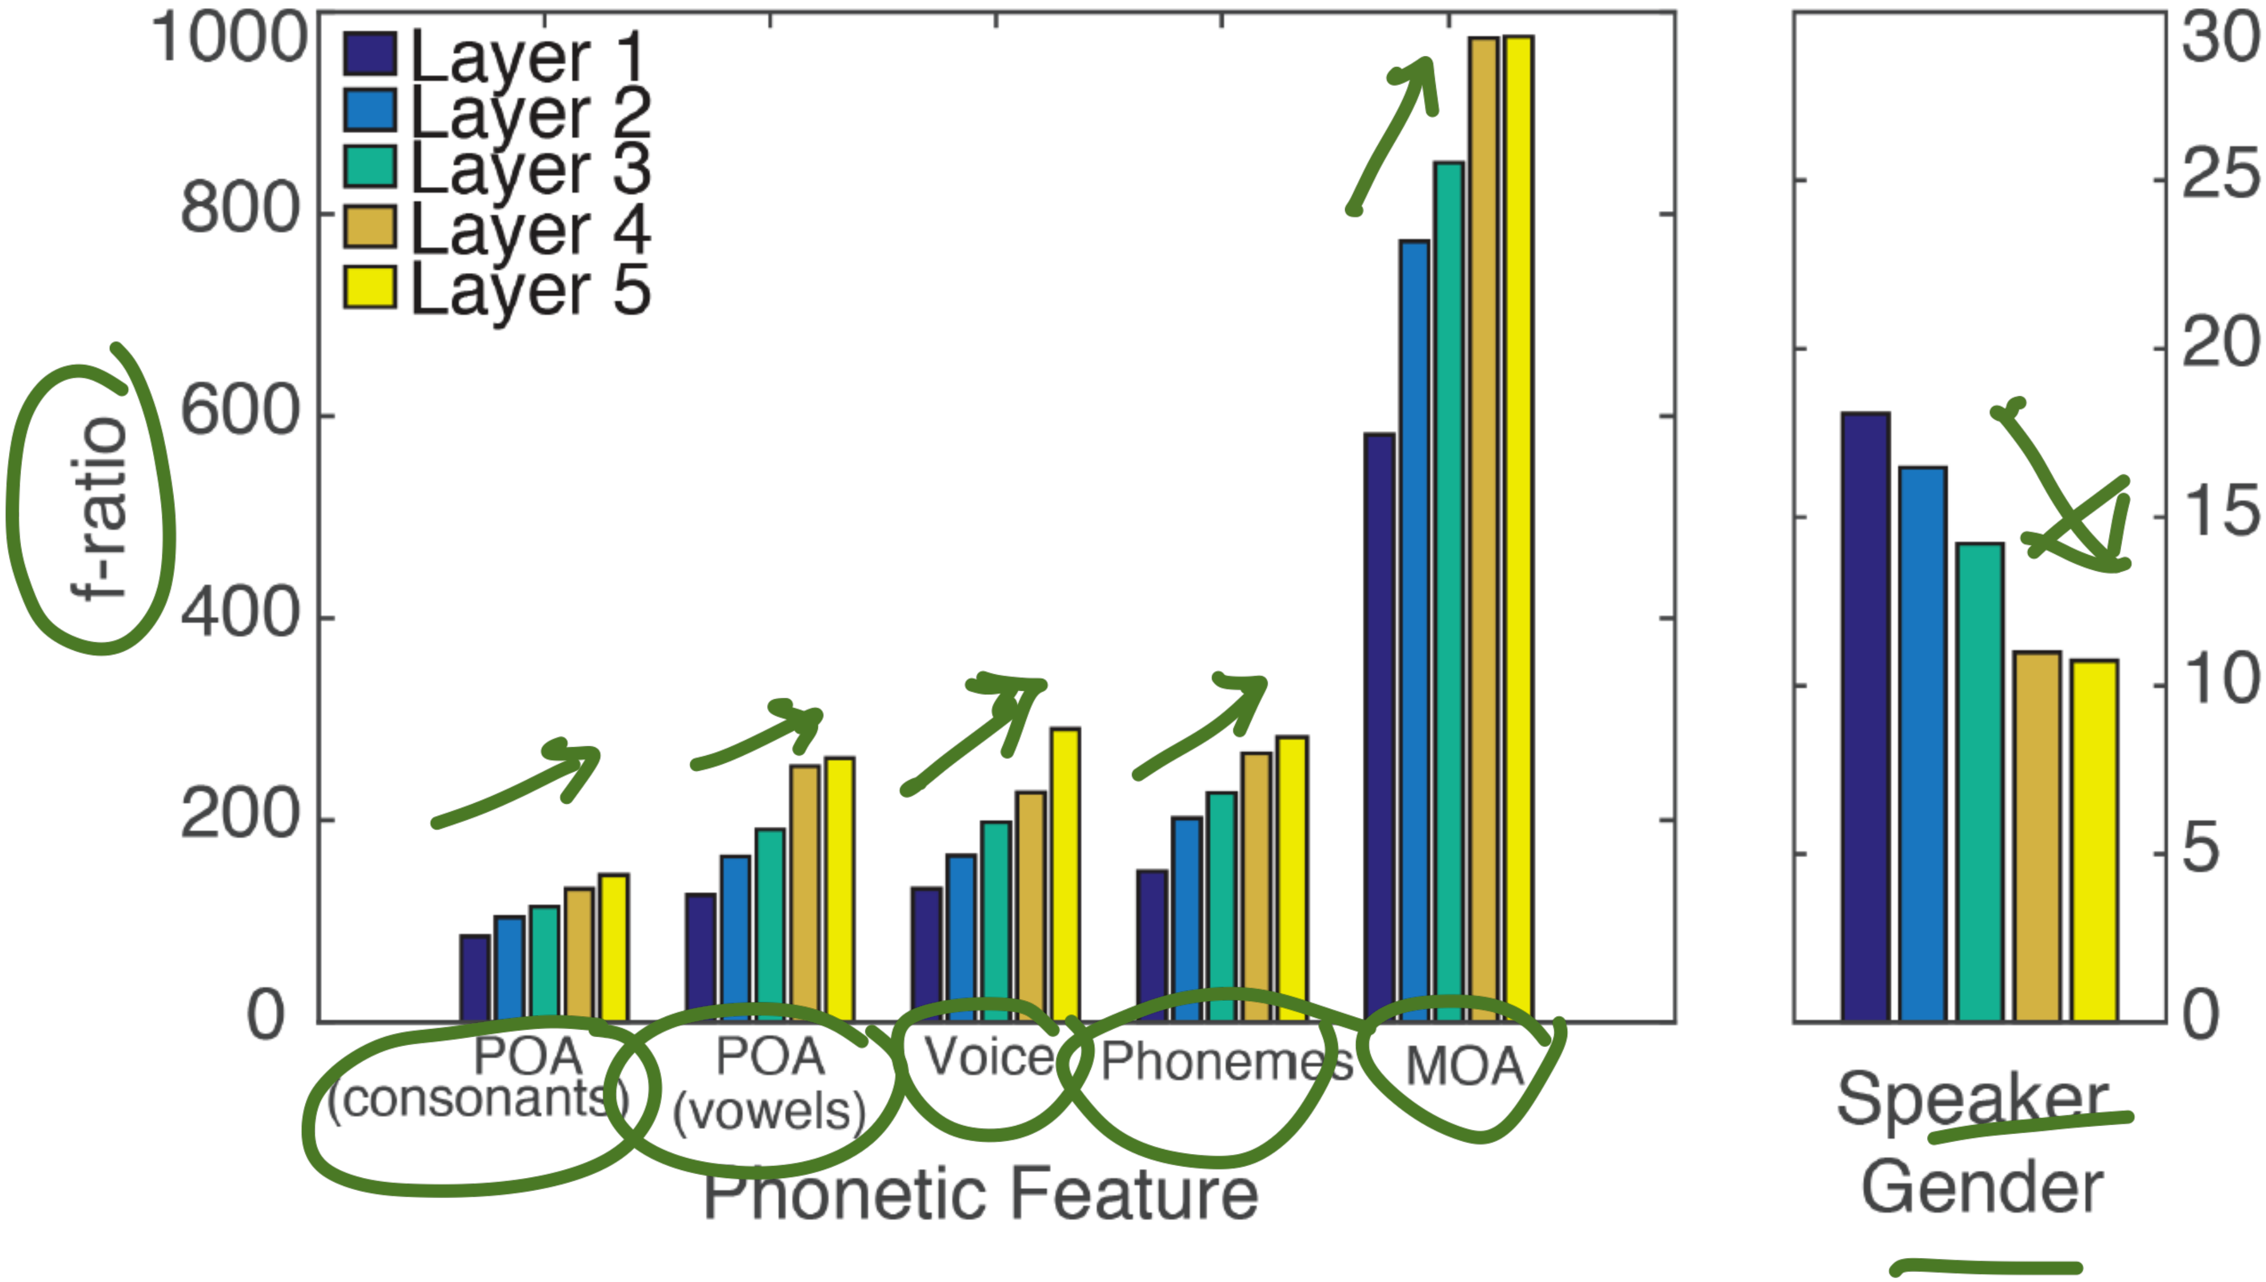

[Nagamine, 2015]

Established 1911

INDIAN INSTITUTE OF SCIENCE
भारतीय विज्ञान संस्थान