Topological Methods for Deep Learning

Abel Symposium, Geiranger
June 5, 2018

Gunnar Carlsson
Stanford University and Ayasdi Inc.

What is Deep Learning?

- Methodology based on neural networks
- Has produced outstanding classification results for complex data
- Images
- Text
- Molecules (Guowei Wei)

Problems

- Adversarial examples
- General lack of transparency
- Limits usefulness in many key domains, financial regulation, health care
- Would like to be able to learn more complex models

Neural Networks

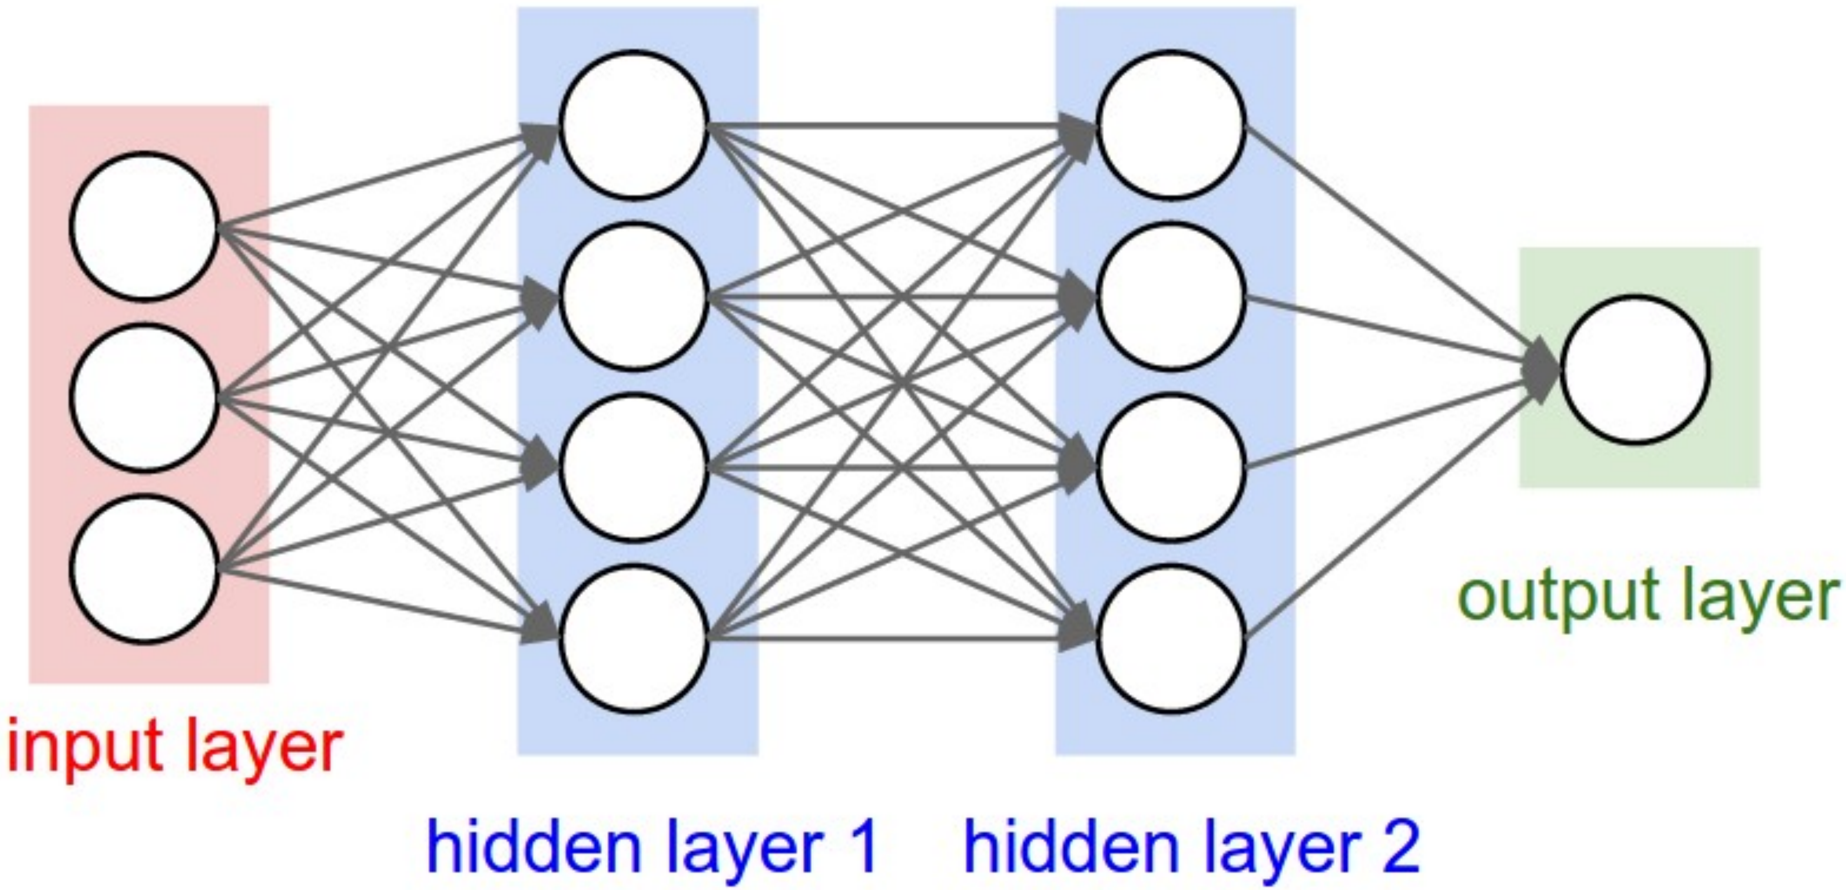

Neural Networks

- Given a data set and an output function, perhaps Boolean valued
- Weights are assigned to the directed edges of the network
- Activation at a node is computed using uniform function of activations of nodes connected to it
- Network is ”trained” by optimization algorithms acting on the set of weights
- Final output is a formula (very large) determined by the final set of weights

Convolutional Neural Networks

- Structure of network adapted to specific cases
- Images (2D rectangular arrays)
- Text (1D arrays)

Convolutional Neural Networks

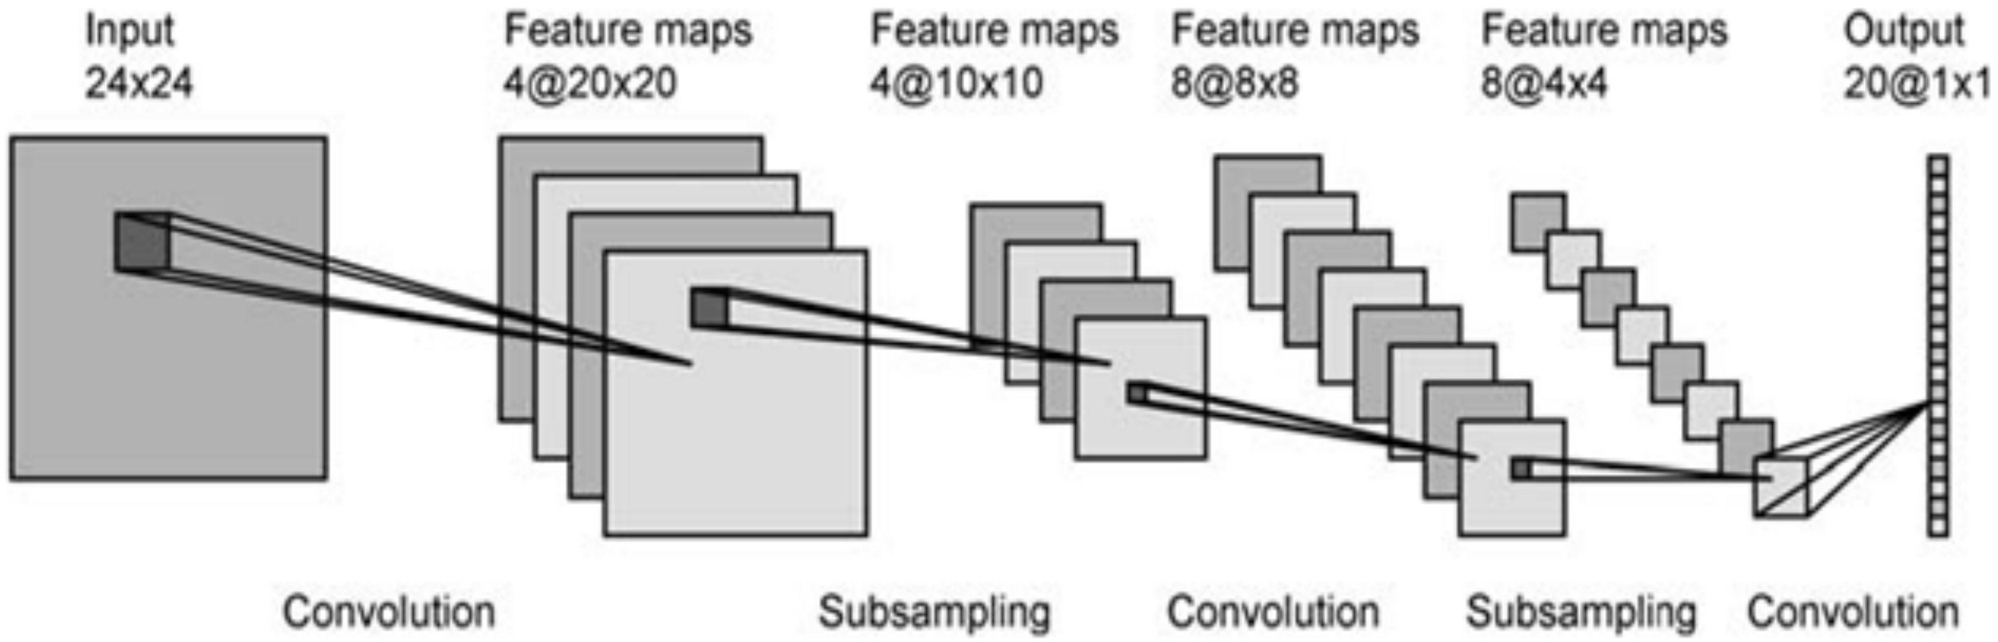

TDA and Deep Learning

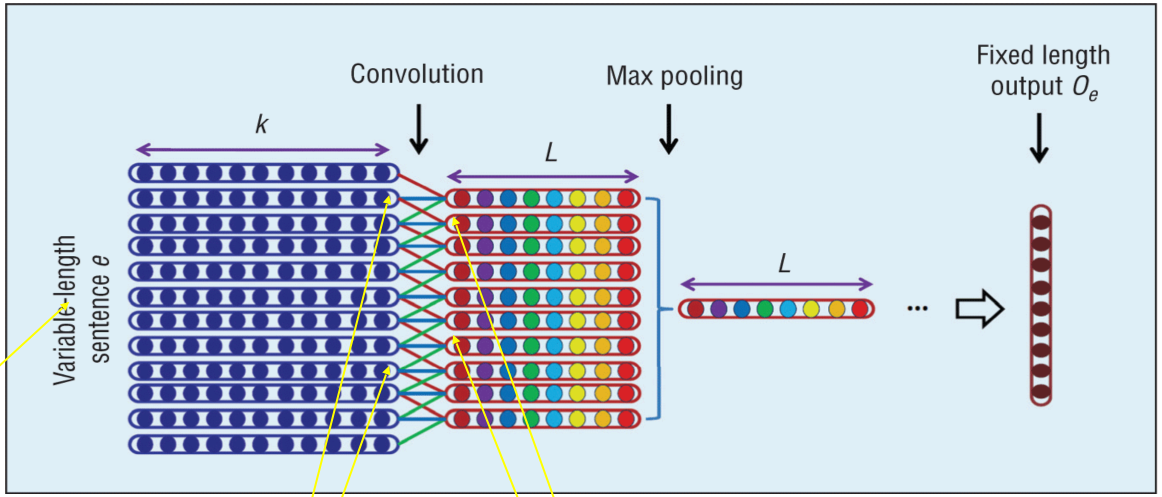

Feature generation

Weights

Activations

DL creation of
topological models

Mumford Data Set (De Silva, Ishkhanov, Zomorodian, C.)

- Analysis of a data set of 3 x 3 patches in natural images
- Studied only ”high variance” patches
- Studies only densest such patches (frequently occurring motifs) – density proxies of varying locality
- Motivated by goal of understanding how tuning of neurons in visual cortex is affected by statistics of natural images

Image Patch Analysis: Primary Circle

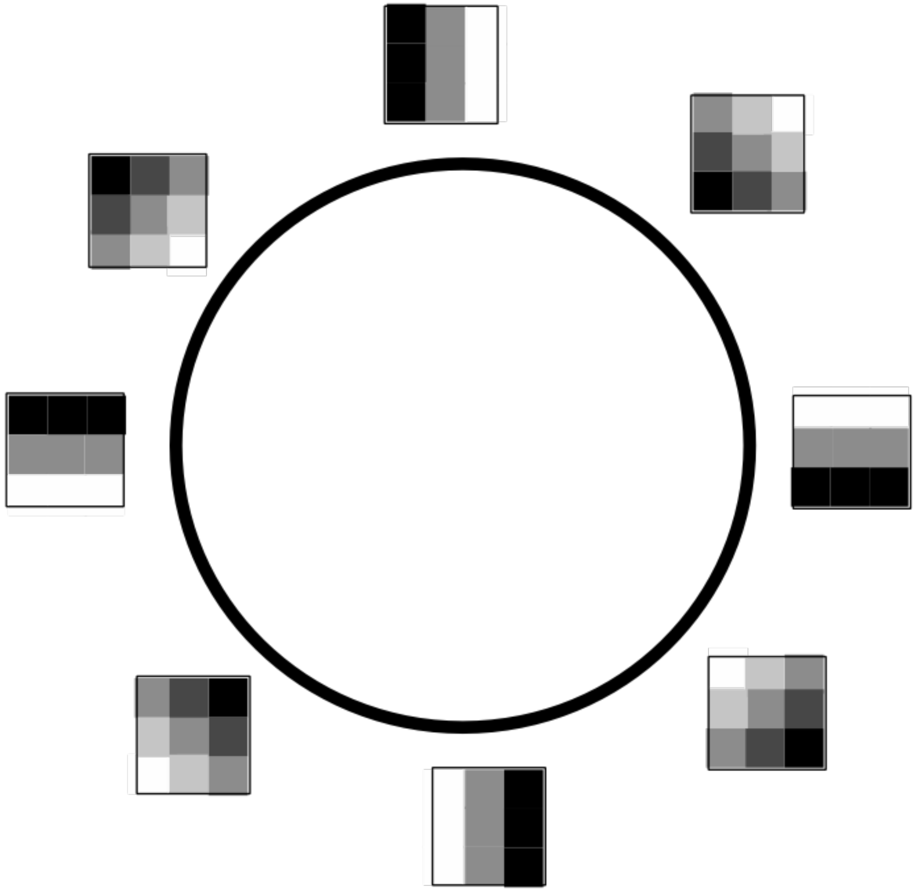

Highest density high variance patches – non-local density measure

Image Patch Analysis: Three Circle Model

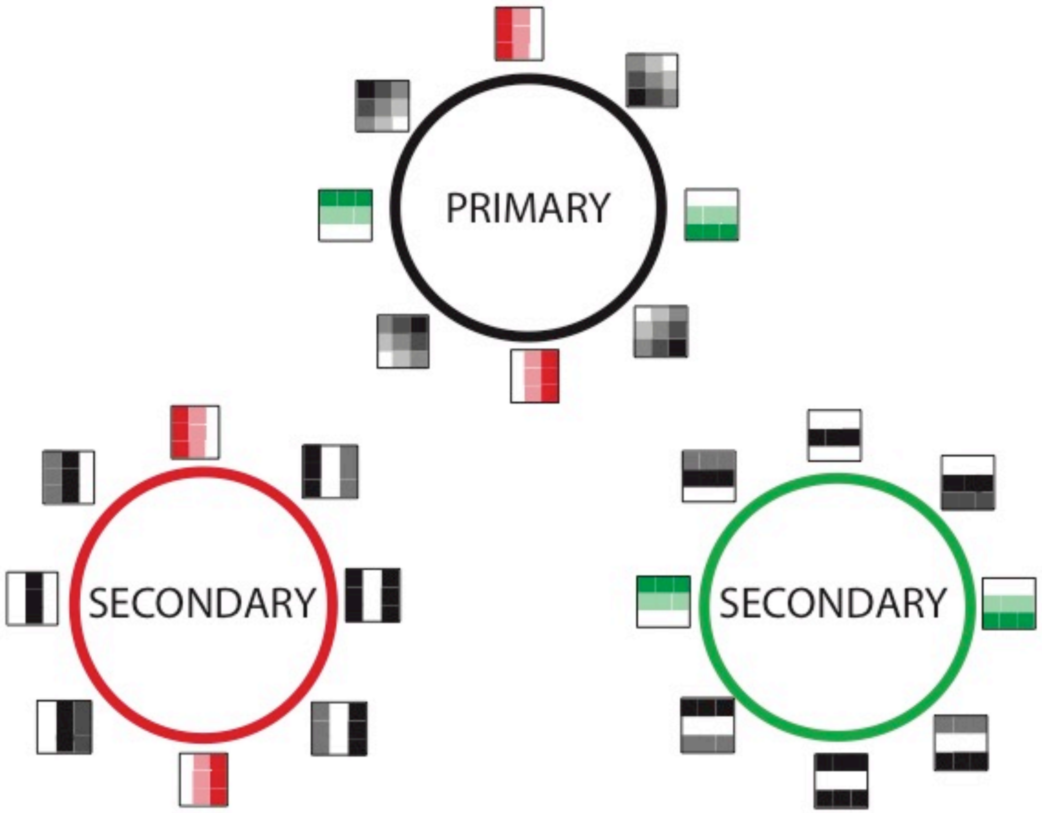

More local density measure

Image Patch Analysis: Klein Bottle

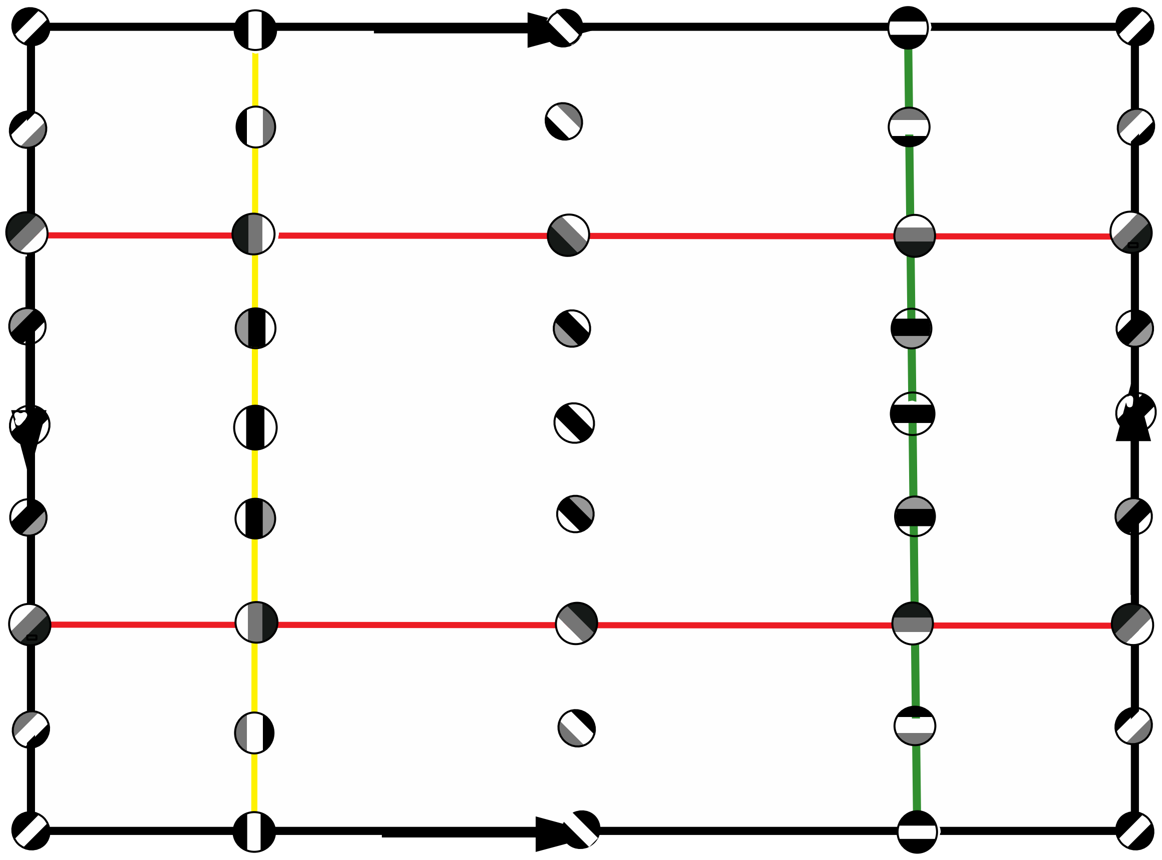

Still weaker threshold

Primary Visual Cortex

- Primary visual cortex (V1) lowest level processing beyond retina
- Higher levels (V2,V4, LGN, etc.) perform more abstract tasks
- Hubel-Wiesel show that individual neurons detect edges and lines
- Consistent with idea of compression of frequent signals

Visual Pathway

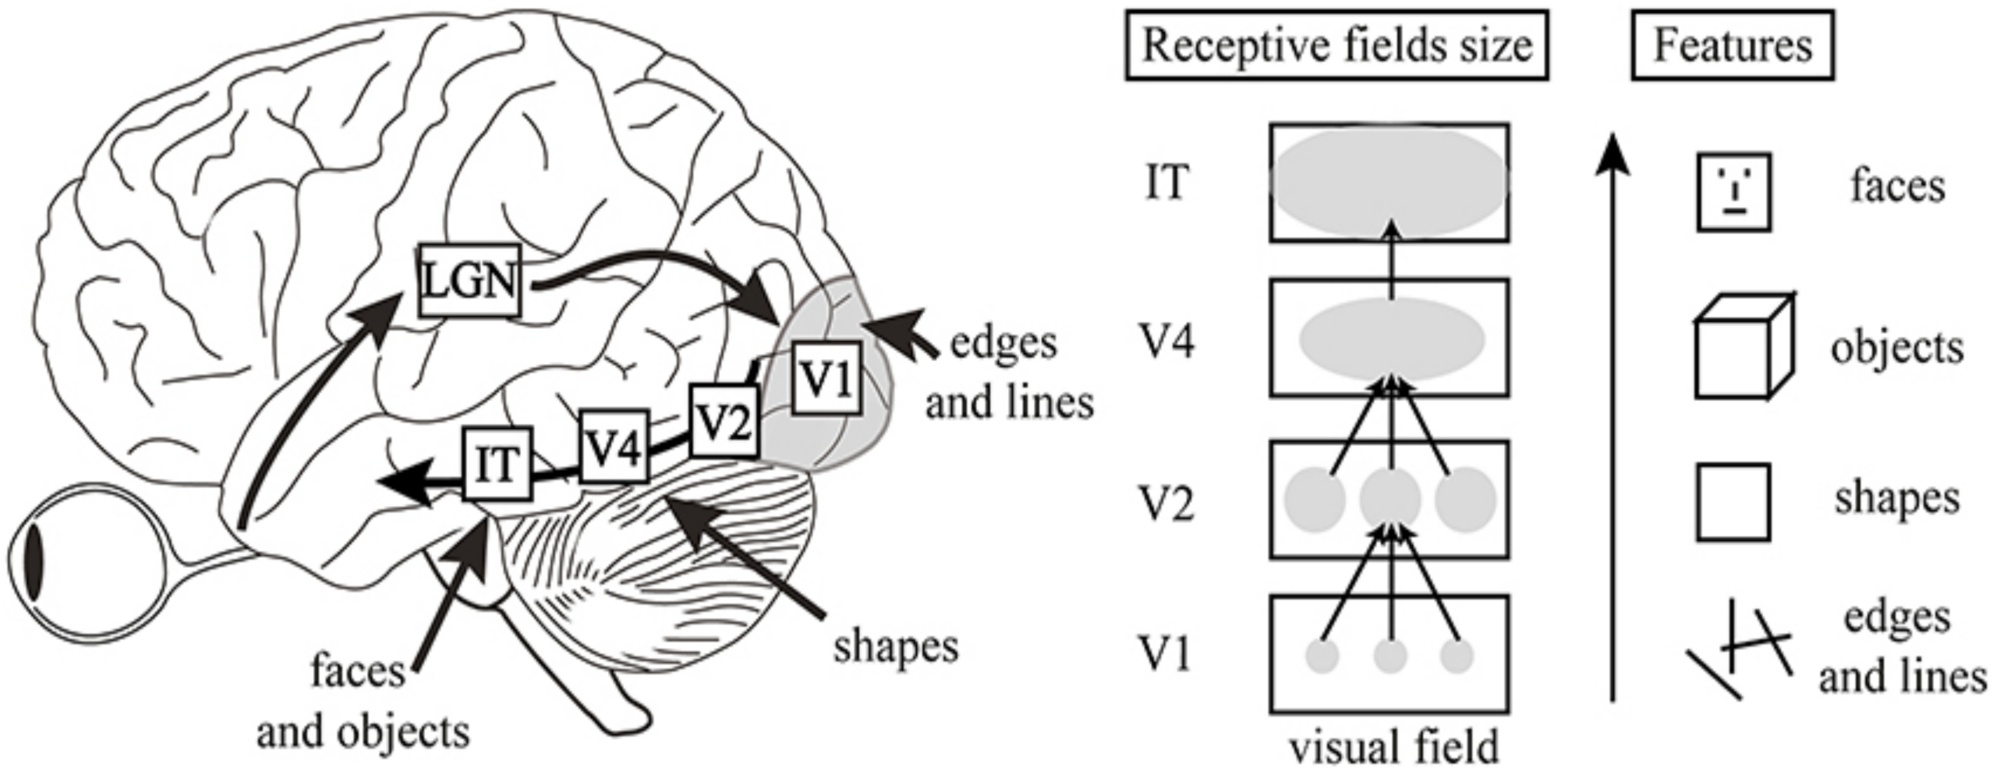

Does Learning by CNN's Behave Like Human Learning?

Joint work with Rickard Brüel Gabrielsson

What Do We Want to Know?

- Can we see similarities to what we have found in the image patch data?
- What happens as the network learns?
- What are the “responsibilities” of the various layers?

Data Sets

- MNIST - hand drawn images
- Cifar10 – Images of various objects, airplanes, cars, etc.
- ImageNet – pretrained network VGG16

How to Build Networks for Data Sets

- Apply a projection to the data set
- Use the projection to bin the data into overlapping bins
- Cluster each bin using a fixed clustering method (requires data equipped with metric)
- Create a node for each “partial cluster”
- Create an edge between any two nodes whose corresponding clusters overlap

How to Build Networks for Data Sets

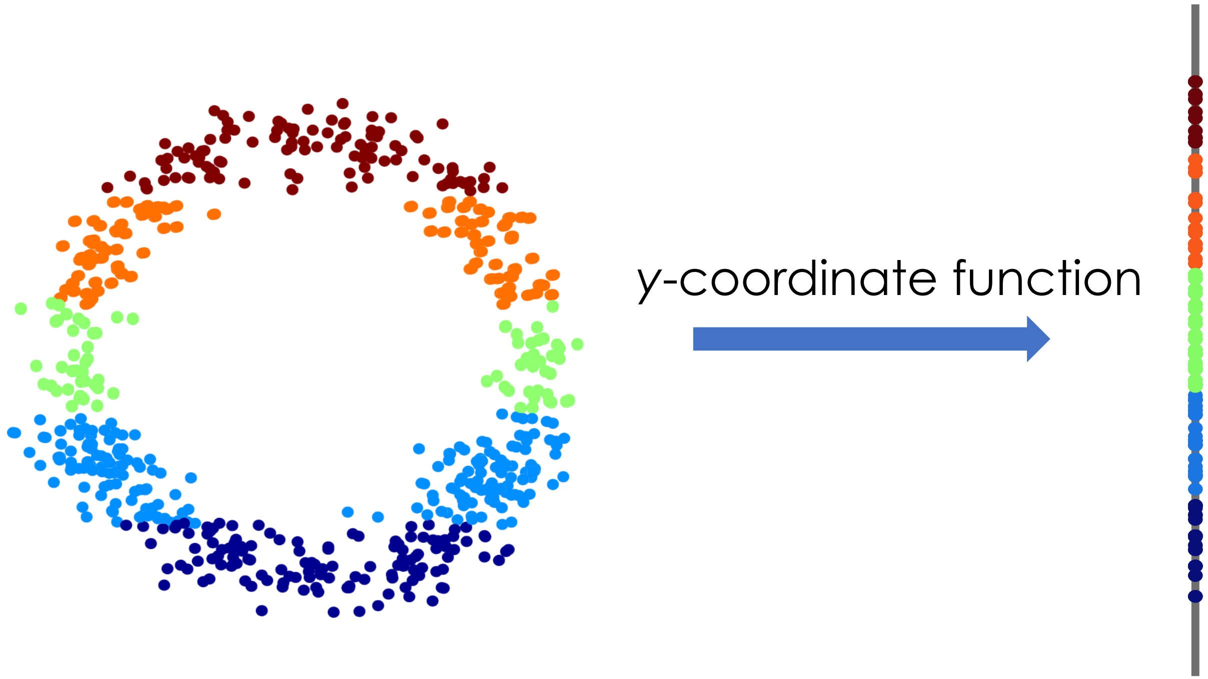

- Apply a projection to the data set
- Use the projection to bin the data into overlapping bins
- Cluster each bin using a fixed clustering method (requires data equipped with metric)
- Create a node for each “partial cluster”
- Create an edge between any two nodes whose corresponding clusters overlap

How to Build Networks for Data Sets

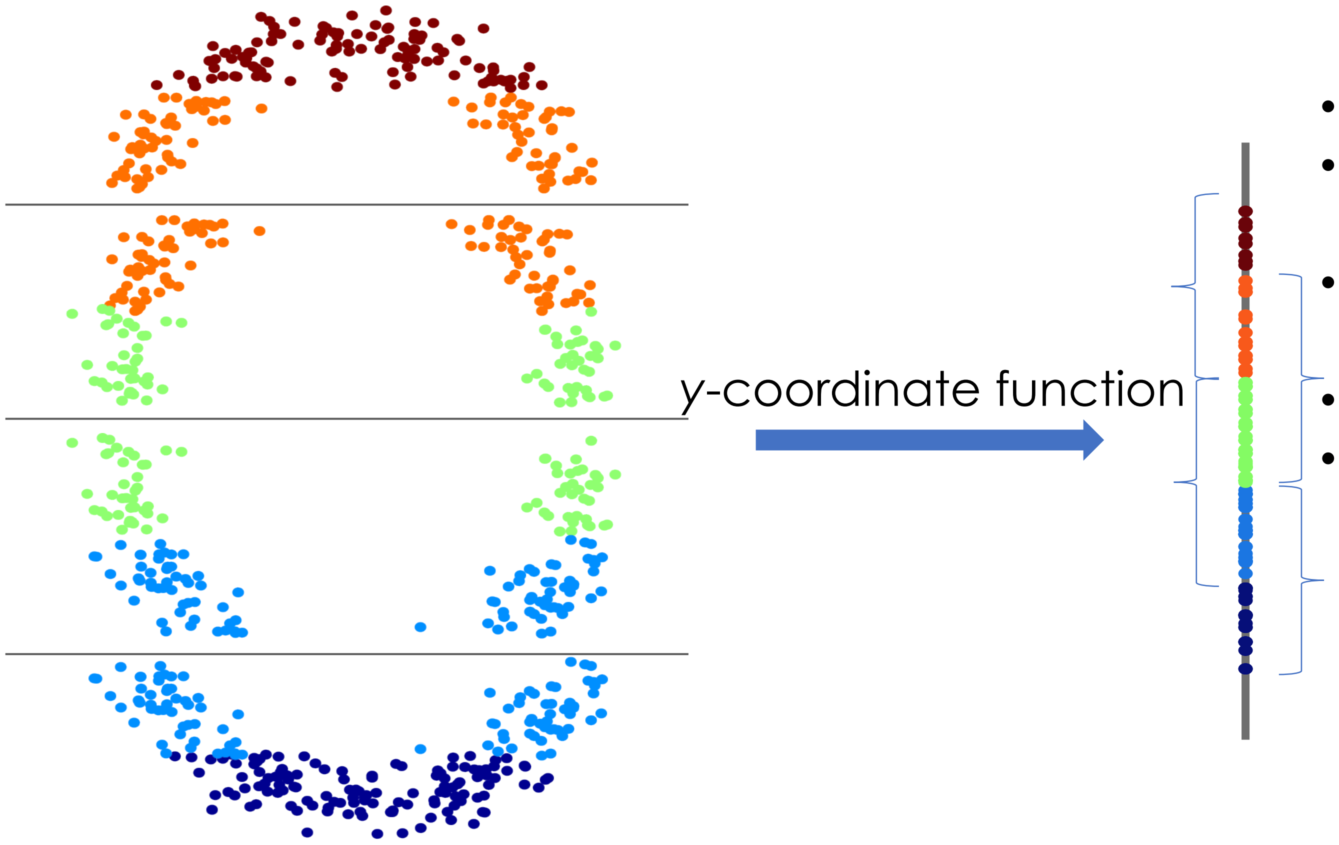

- Apply a projection to the data set
- Use the projection to bin the data into overlapping bins
- Cluster each bin using a fixed clustering method (requires data equipped with metric)
- Create a node for each “partial cluster”
- Create an edge between any two nodes whose corresponding clusters overlap

How to Build Networks for Data Sets

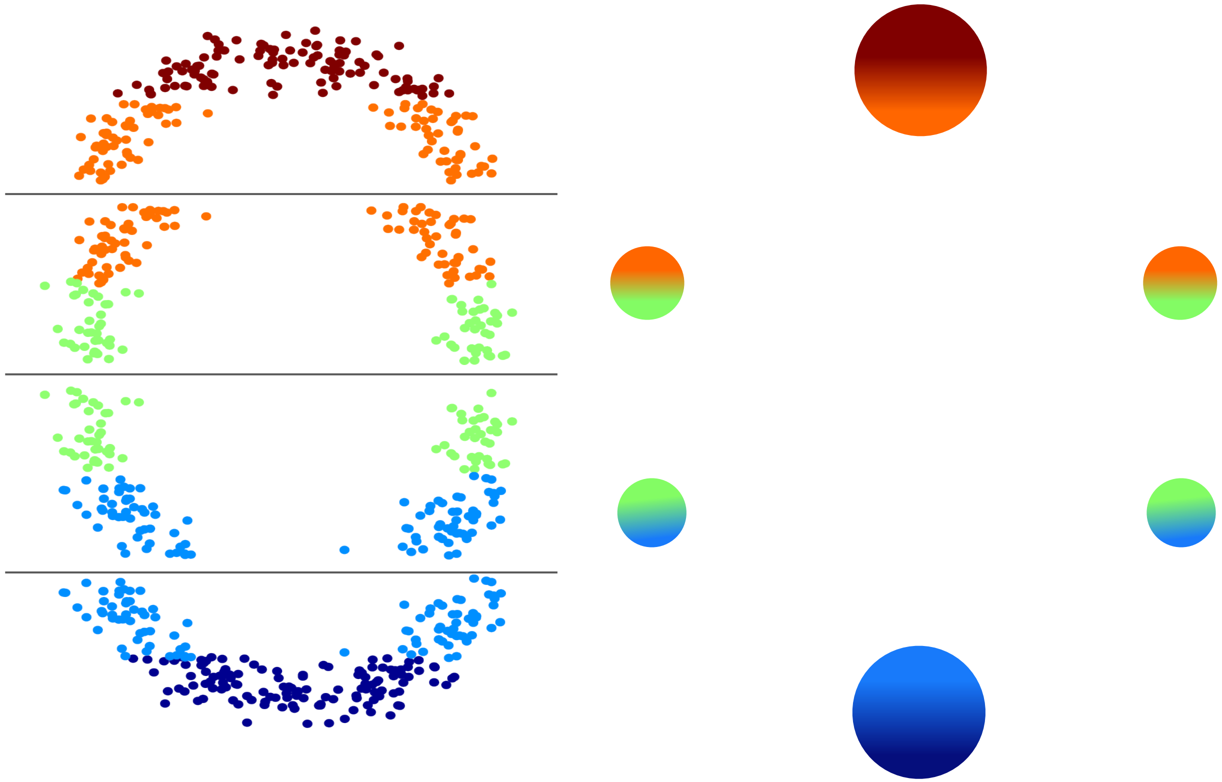

- Apply a projection to the data set
- Use the projection to bin the data into overlapping bins
- Cluster each bin using a fixed clustering method (requires data equipped with metric)
- Create a node for each “partial cluster”
- Create an edge between any two nodes whose corresponding clusters overlap

How to Build Networks for Data Sets

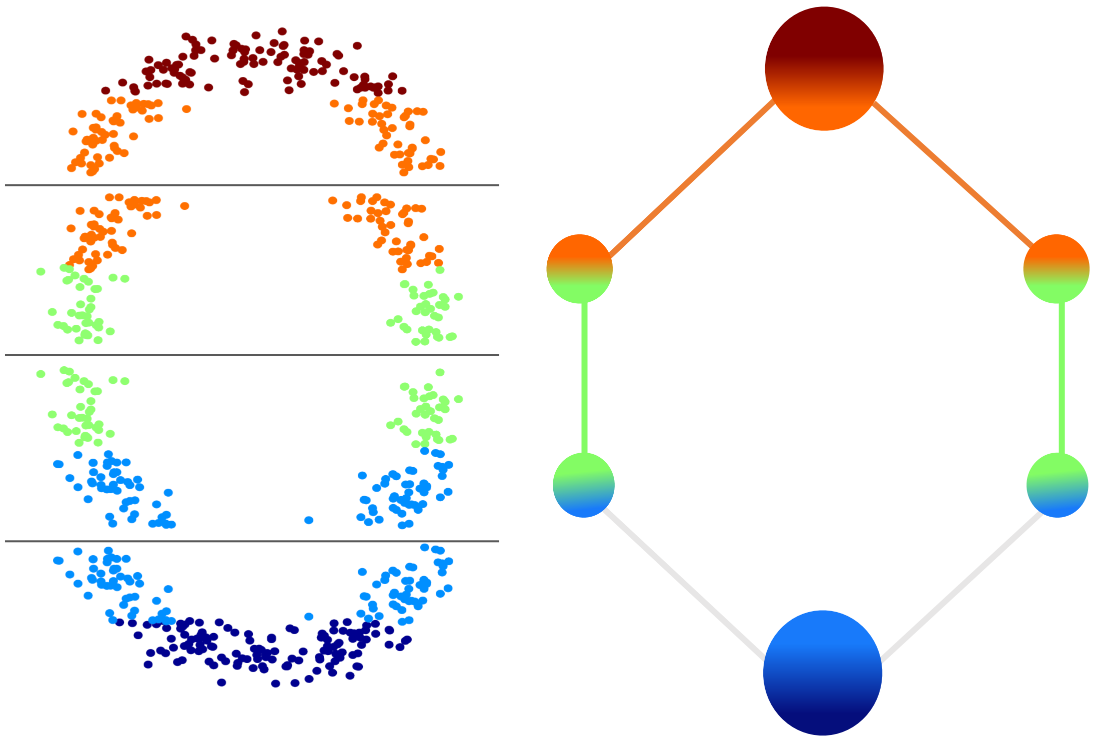

- Apply a projection to the data set
- Use the projection to bin the data into overlapping bins
- Cluster each bin using a fixed clustering method (requires data equipped with metric)
- Create a node for each “partial cluster”
- Create an edge between any two nodes whose corresponding clusters overlap

Topological Analysis of Weight Spaces (MNIST)

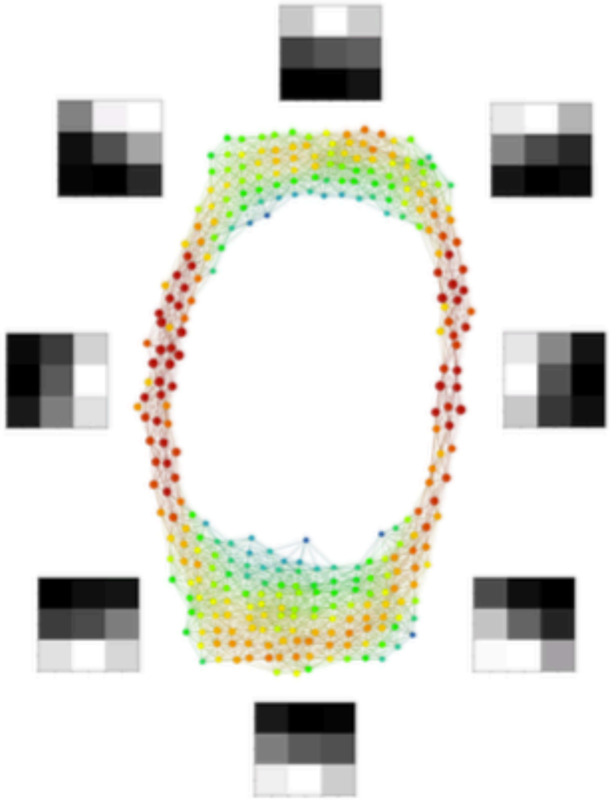

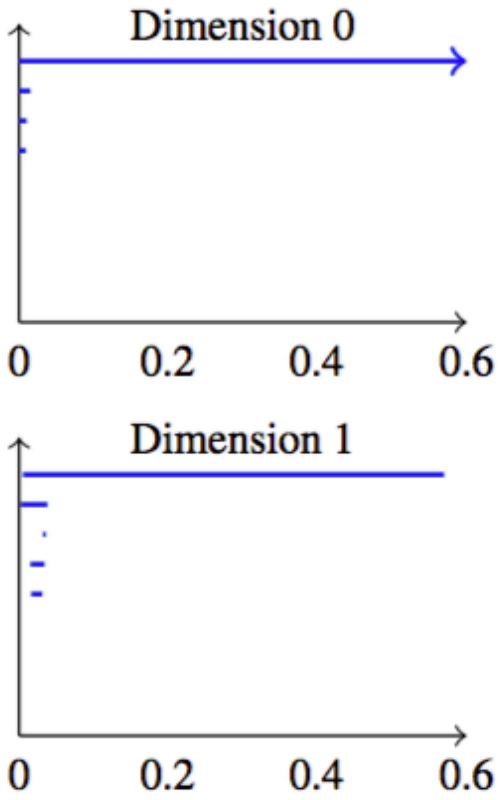

Non-local density thresholding for layer 1 of depth 2 net

Topological Analysis of Weight Spaces (MNIST)

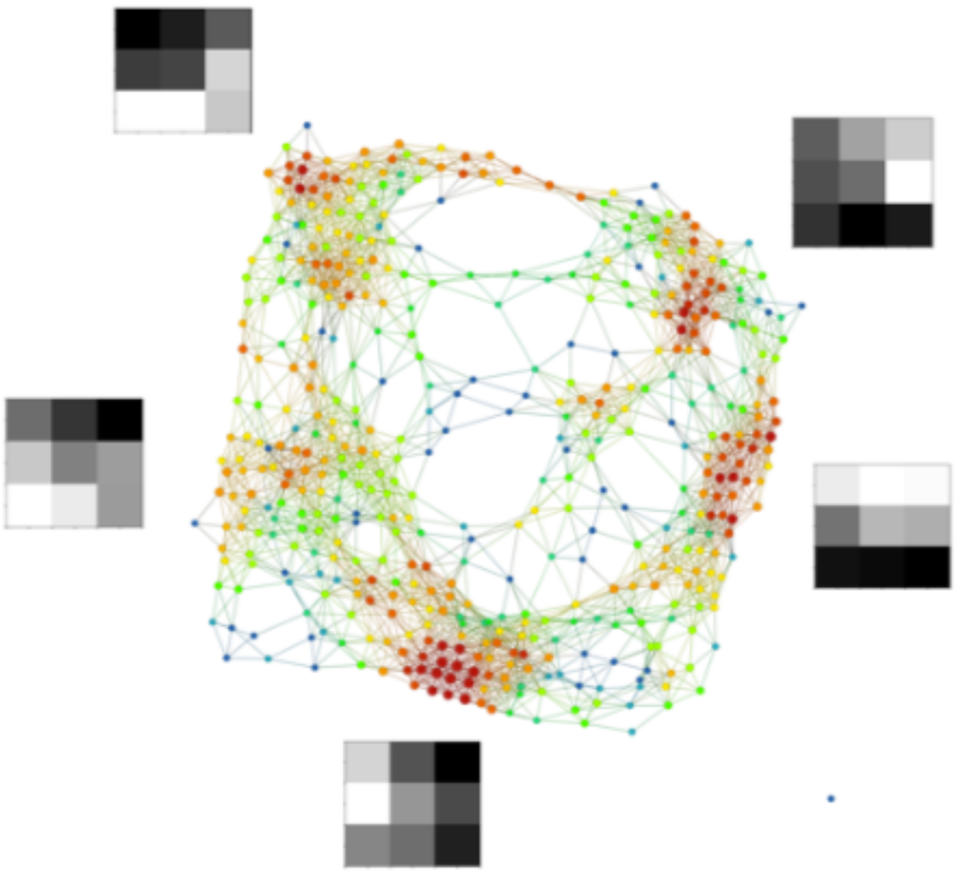

More localized density measurement for layer 1 of depth 2 net

Topological Analysis of Weight Spaces (Cifar10)

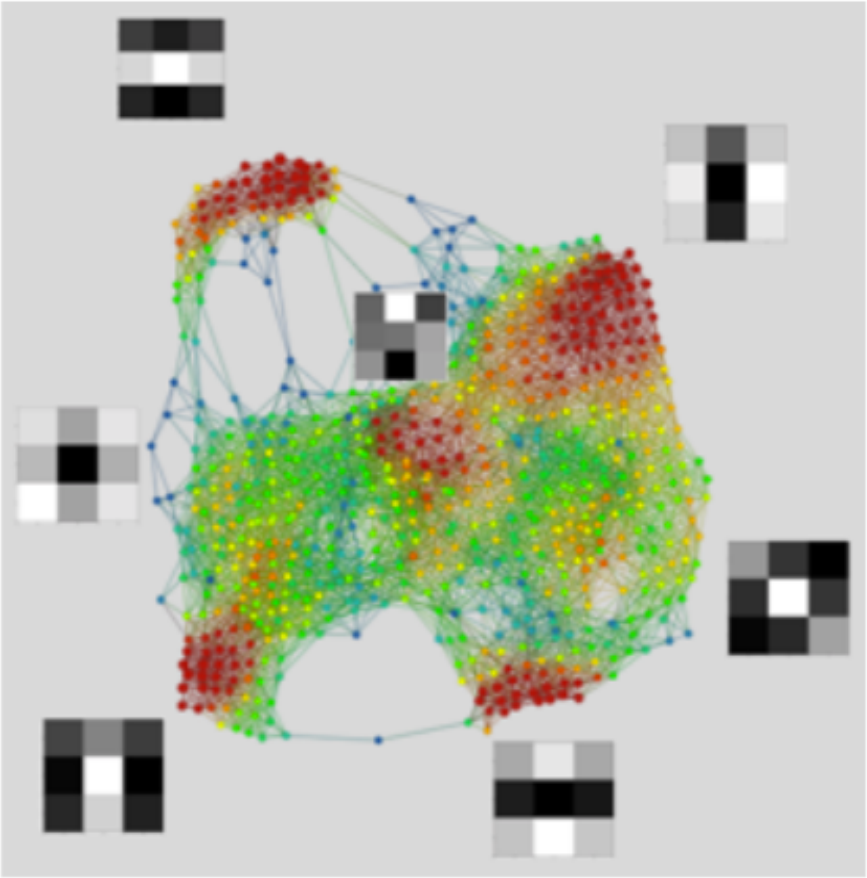

1st Layer of CNN for this data set, reduced to gray scale

Topological Analysis of Weight Spaces (Cifar10)

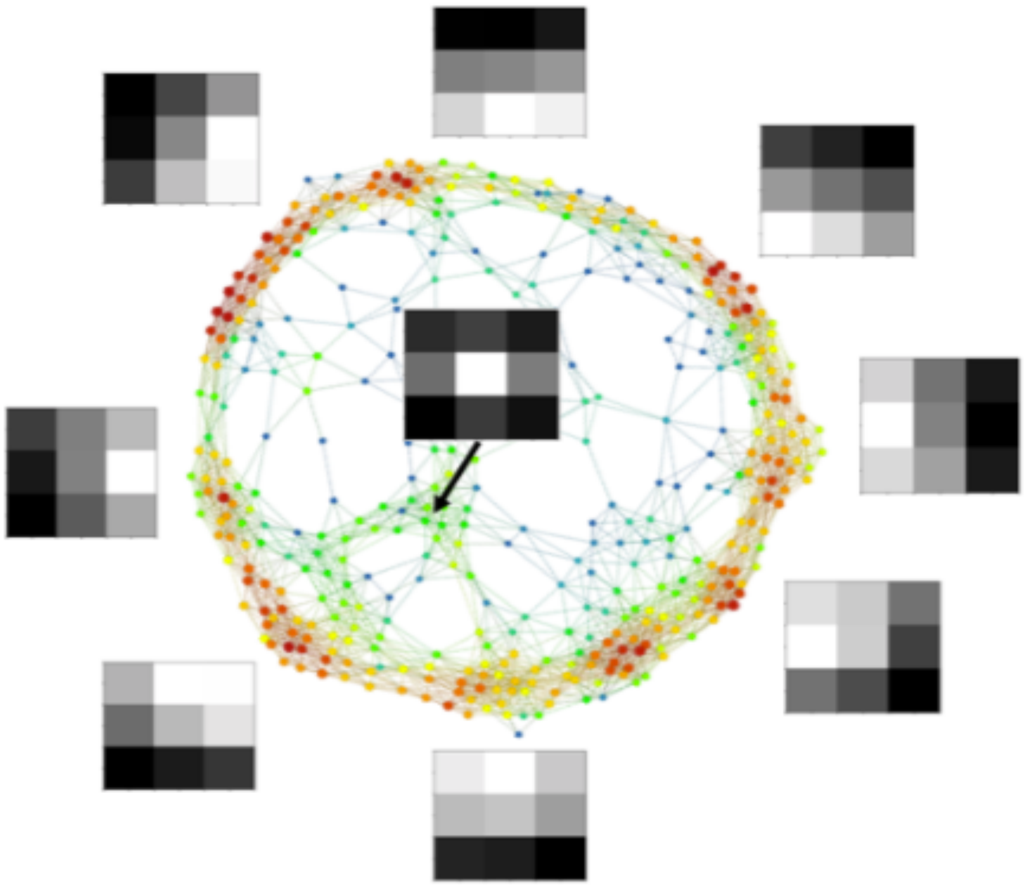

2nd Layer for single CNN, reduced to gray scale

Topological Analysis of Weight Spaces (Cifar10)

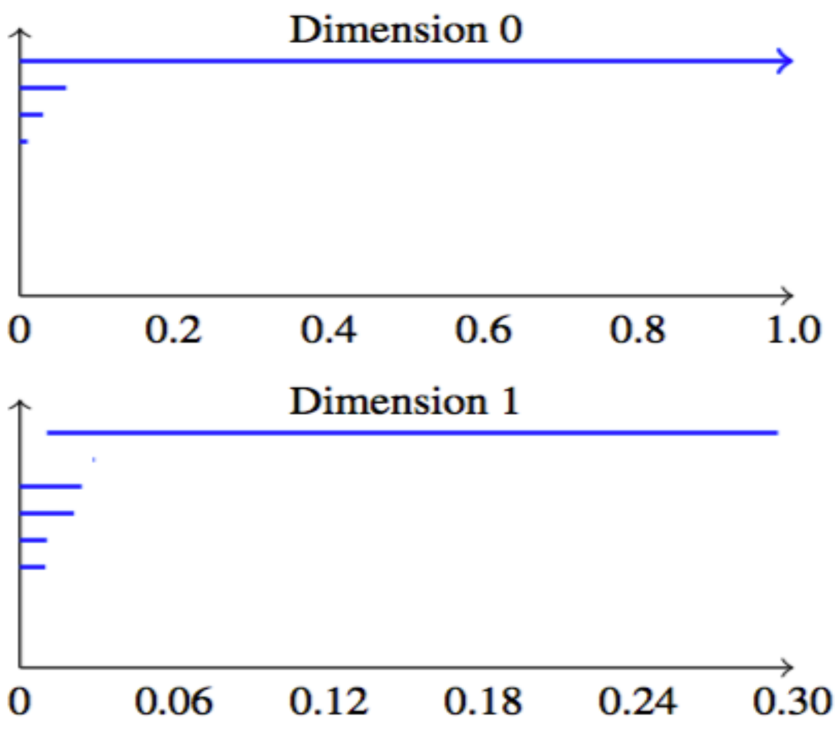

1D barcode for tightly thresholded data set of for 2nd layers , reduced to gray scale

Topological Analysis of Weight Spaces (Cifar10)

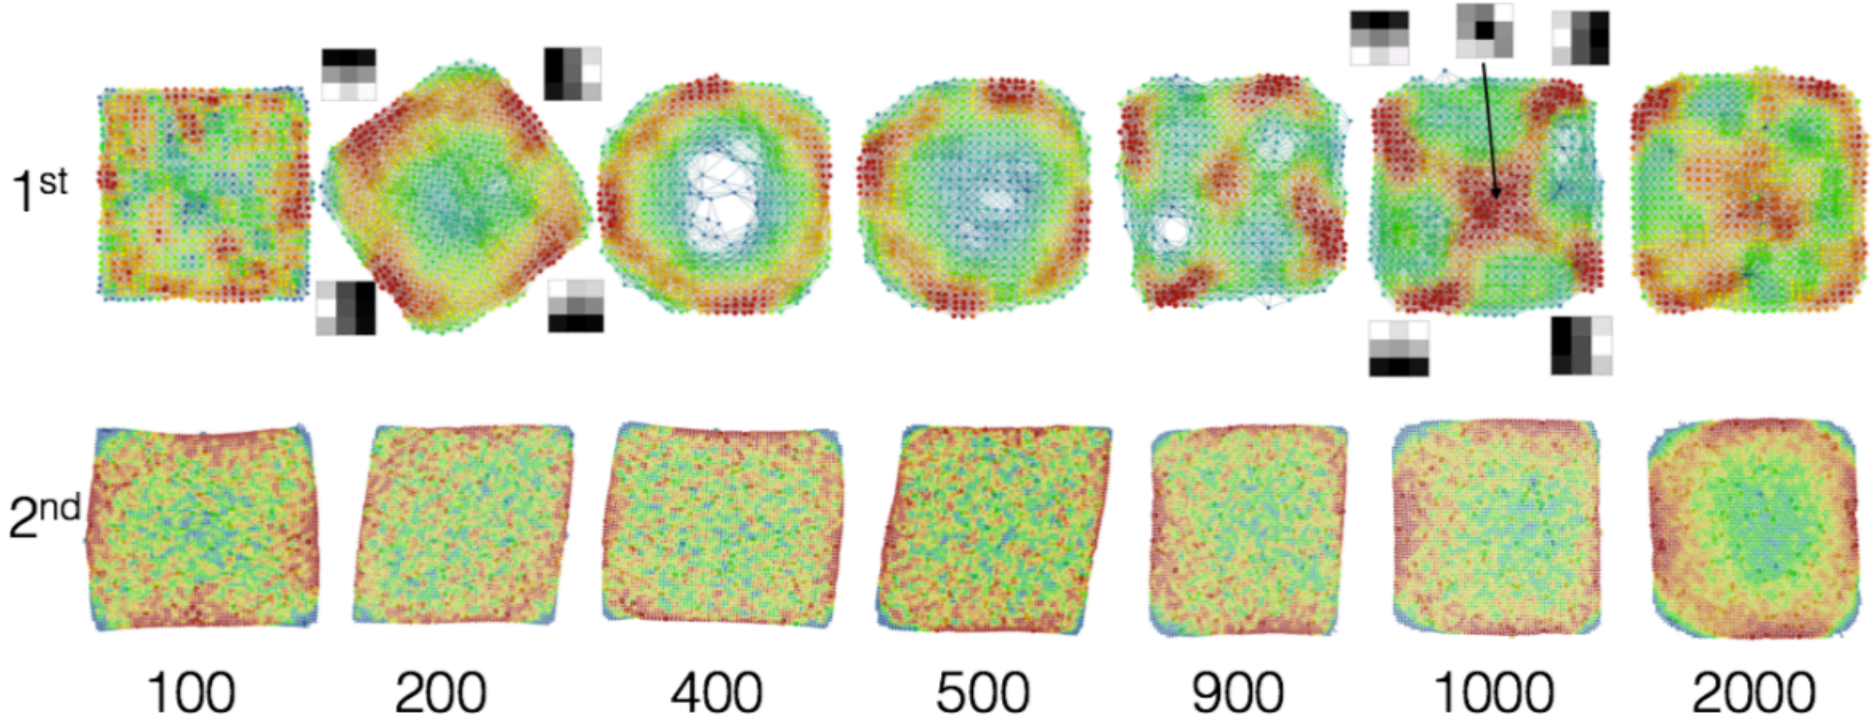

Mapper representations over the number of iterations, tightly density thresholded, gray scale reduced

Topological Analysis of Weight Spaces (Cifar10)

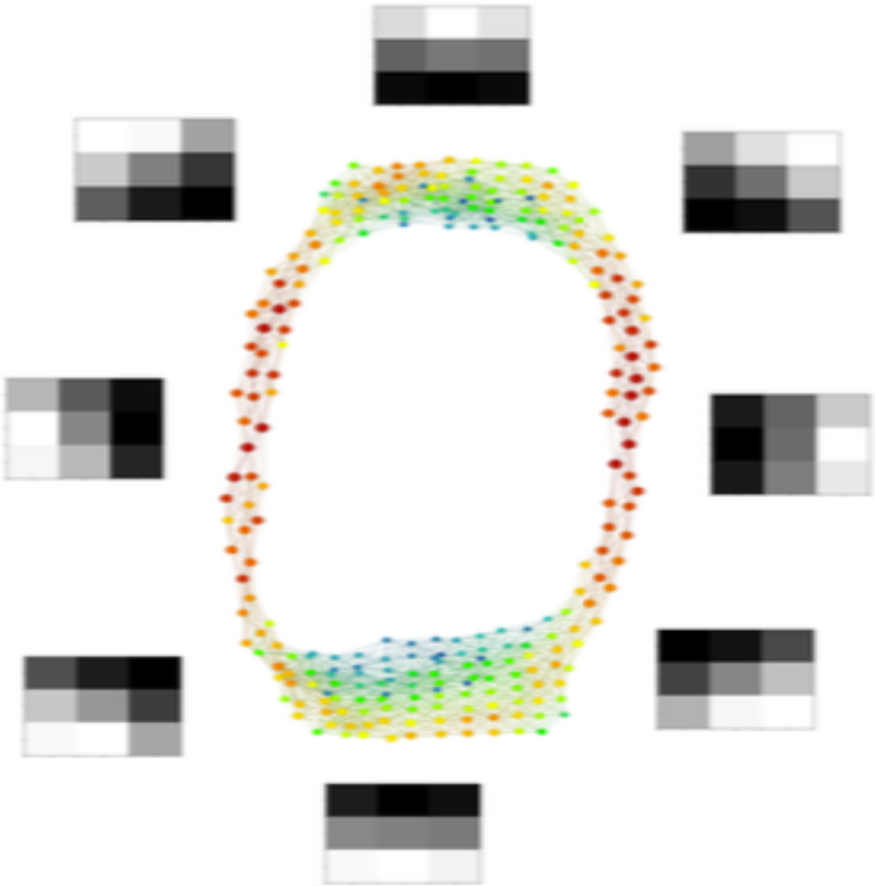

1st layer, Coarse density thresholding, color retained

Topological Analysis of Weight Spaces (Cifar10)

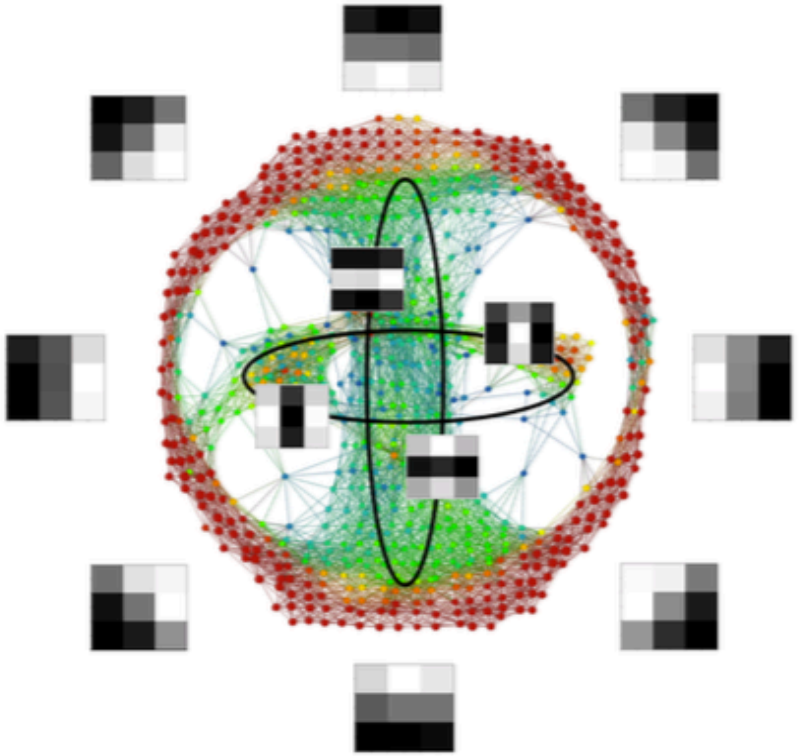

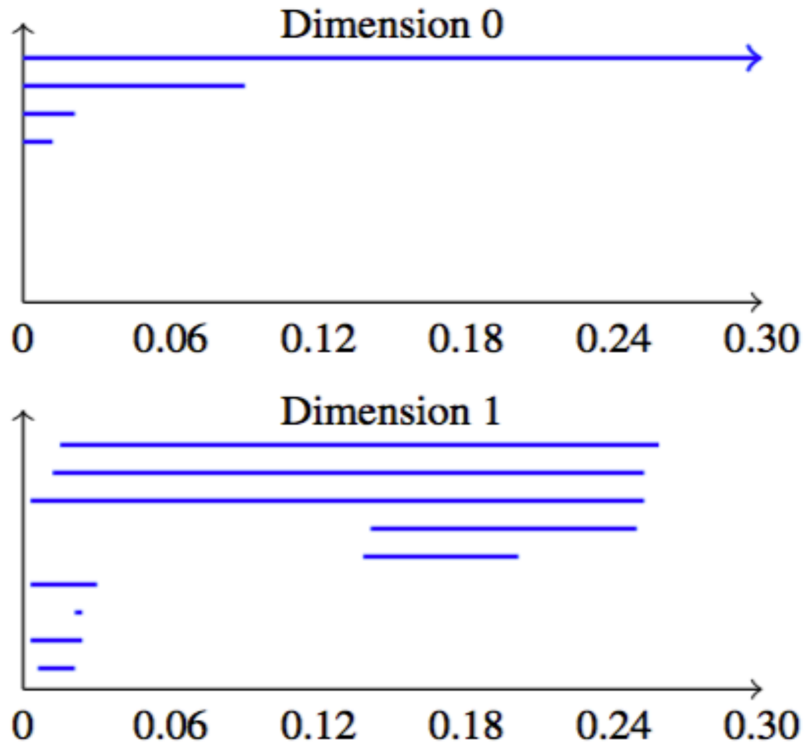

1st layer, looser density thresholding, more localized density estimator, color retained

Topological Analysis of Weight Spaces (Cifar10)

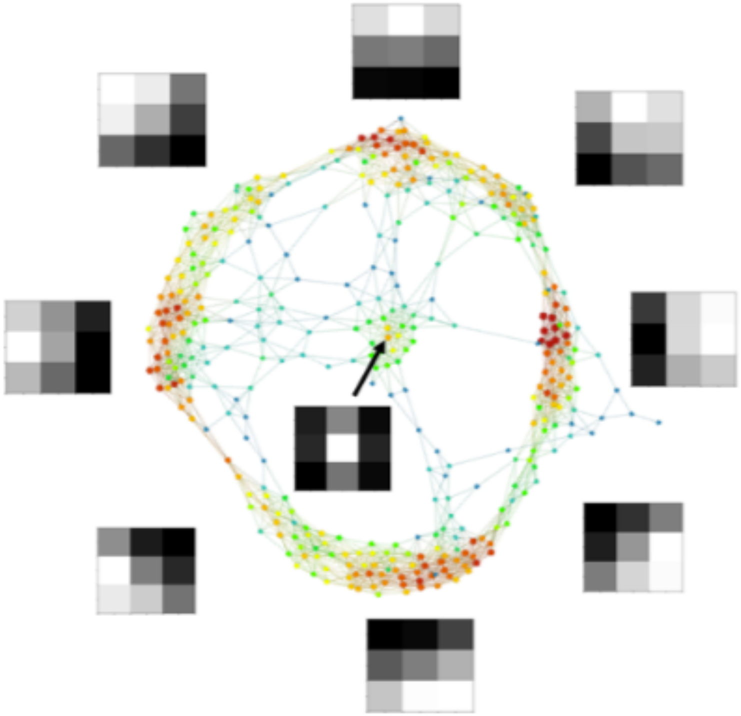

2nd layer, fine density thresholding, color retained

Topological Analysis of Weight Spaces (VGG16)

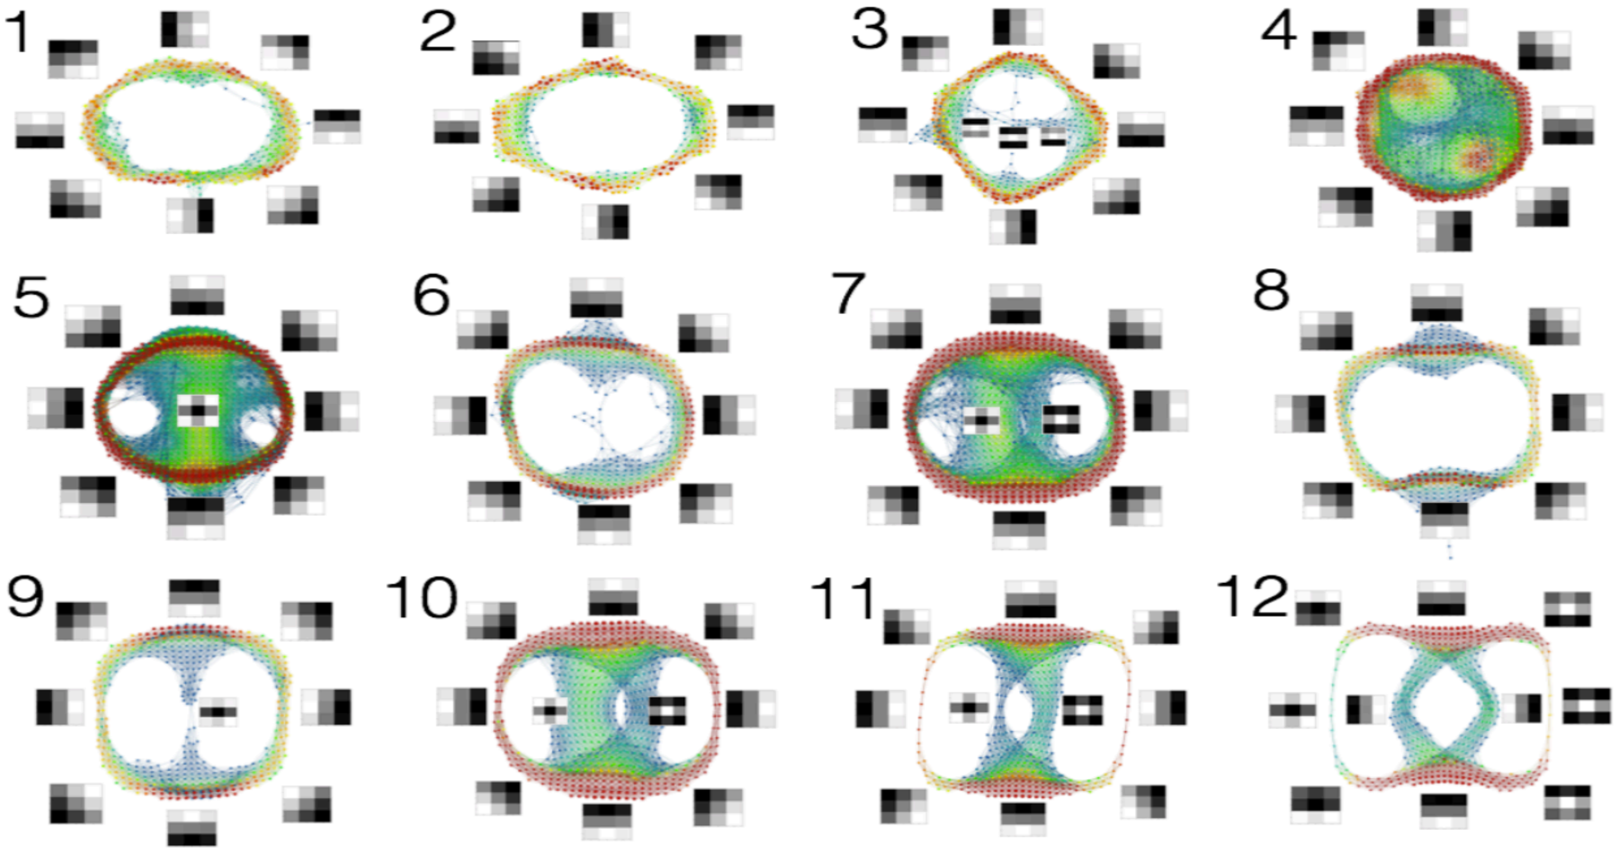

Mapper findings from each of 13 layers, same density thresholding, relatively local estimator

Topological Analysis of Activations – P. Musial

Convolutional Neural Network (NN)

pre-trained softmax classifier with ReLU units, dropout
test accuracy 99.39%

layer:	dimensions:
output	10
fully connected	1024
max pooling	7 x 7 x 64
convolutional	14 x 14 x 64
max pooling	14 x 14 x 32
convolutional	28 x 28 x 32
input	28 x 28

Adversarial Examples on MNIST

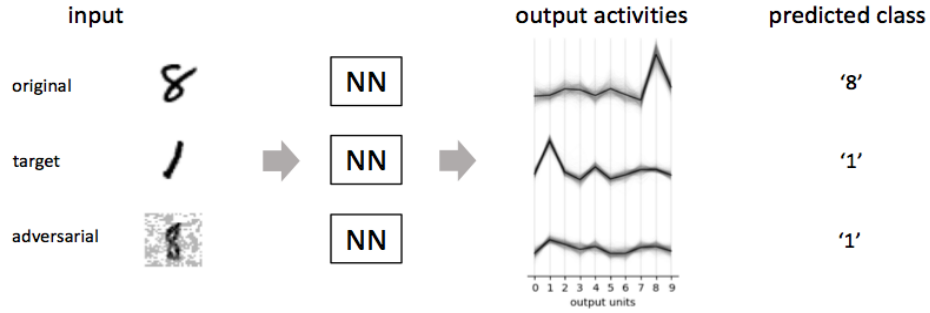

Topological Data Analysis

Activities at output layer (Euclidean L2 metric, Neighborhood Lenses 1,2)

predicted class: blue - original ('8'), teal – target ('1')

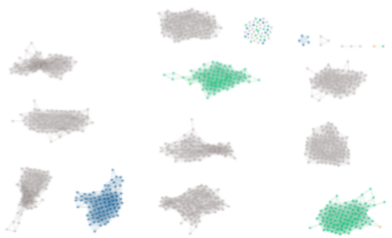

red – adversarial examples

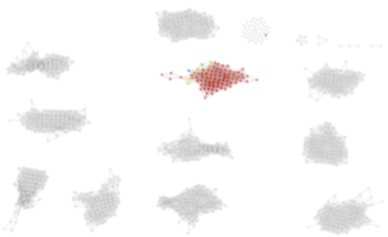

Detects ”adversarial” behavior

Remarks

- Analysis of the black box is a data analysis problem in its own right
- Density critical - what is common and what is not
- Can begin to understand what happens in more abstract layers
- Can study behavior over number of iterations in optimization step
- Adversarial behavior can be detected

Feature Space Modeling

- Given a data matrix, one can also consider the transpose matrix
- The rows of the transpose are the features of the data set
- When there are many features, very useful to create Mapper models
- Compresses and recognizes correlations among features
- Each row of original matrix gives a function on feature set, and on nodes of the topological model

Microarray Analysis of Breast Cancer

Cohort B

Cohort C

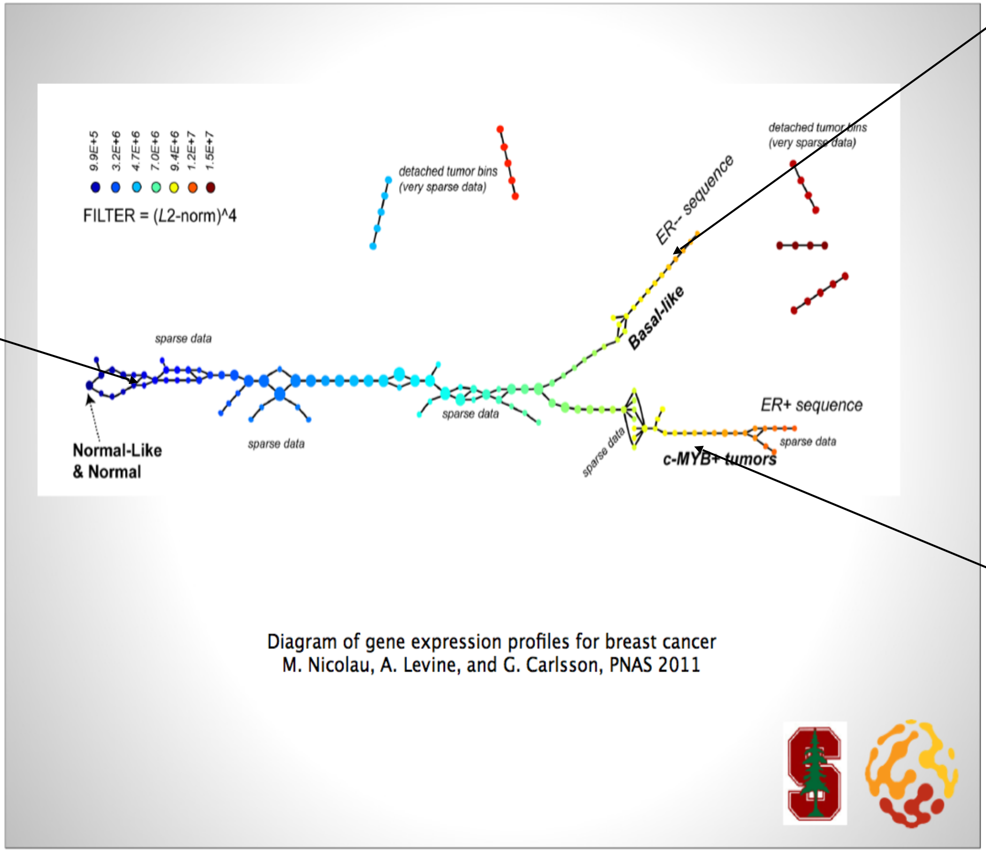

Diagram of gene expression profiles for breast cancer
M. Nicolau, A. Levine, and G. Carlsson, PNAS 2011

Cohort A

Explaining the Different Cohorts

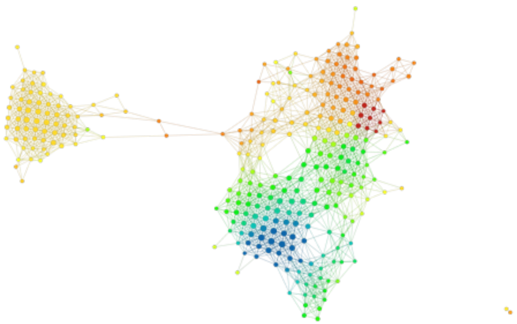

Cohort A

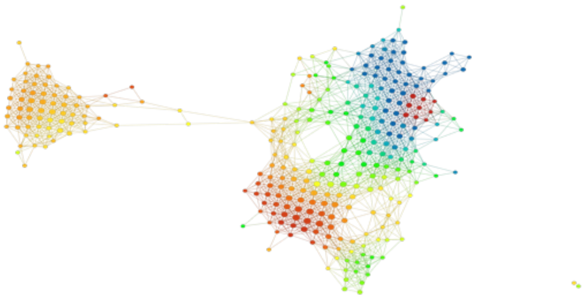

Cohort B

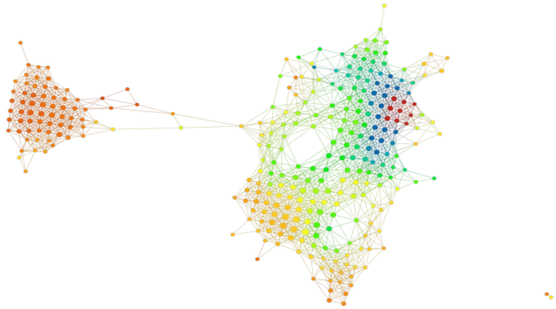

Cohort C

UCSD Microbiome

Feature Space Modeling

- Gives direct representation of high dimensional data sets
- Can be viewed as a smoothing operation
- Treat any data analytic problem as an imaging problem
- Natural to plug into CNN’s