THE CITY UNIVERSITY OF NEW YORK
College of Staten Island

Fibres of Failure

Mapper for Process Diagnostics

Mikael Vejdemo-Johansson
Leo Carlsson
Gunnar Carlsson

Before we start: ATMCS 2018

- Biennial TDA conference

- June 25-29, 2018
  IST-Austria, Klosterneuburg, outside Vienna

- Timed to match with
  Symposium on Computational Geometry /
  Computational Geometry Week in Budapest
  (also with a TDA session)

- http://atmcs8.appliedtopology.org

Outline

- Mapper review
- Prediction as a function
- Borrow from the future: prediction error fibres
- Fibres of Failure
- Example: CNN on corrupted MNIST

Topological background

- Consider:
  - Spaces
    X, Y
  - Continuous map
    $f: X \to Y$
  - Cover
    $Y = \cup Y_i$
- The cover pulls back to a cover
  $X=\cup f^{-1}Y_i$
- Refine cover to connected components
  $X=\cup X_j$; $X_j \in \pi_0 f^{-1}Y_i$
- If each X_j is contractible, Nerve lemma → nerve complex ≃ X.

Topological background

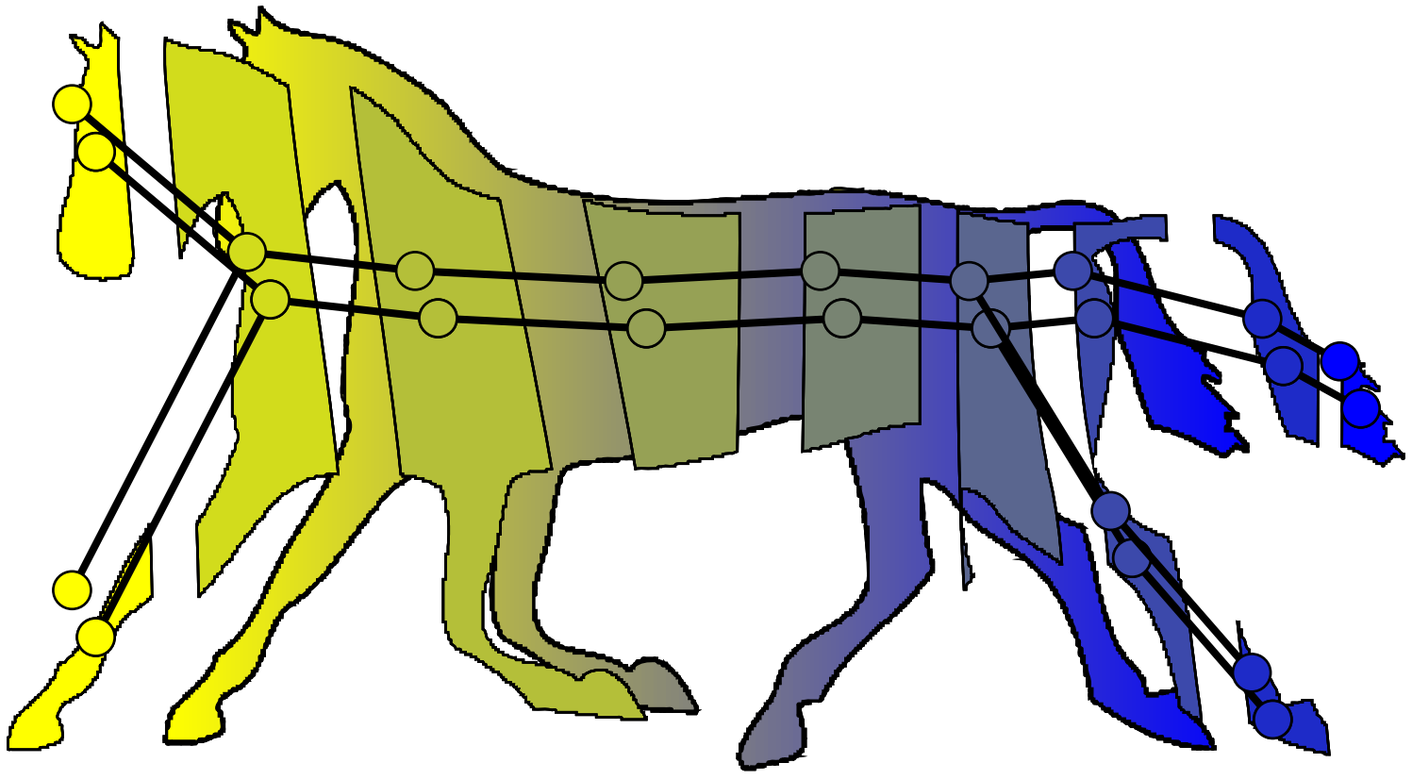

From topology to data: a dictionary

- Topological space
- Continuous map $X \to Y$
- Cover
- π_0
- Nerve complex

- Point cloud
- *Filter* function or *lens* $X \to \mathbb{R}^d$
- Partition with overlap
- Clustering wrt *metric*.
- Nerve complex

Choices: **lens**(es), **metric**,
(parameters for) **partition** and of **clustering** method.

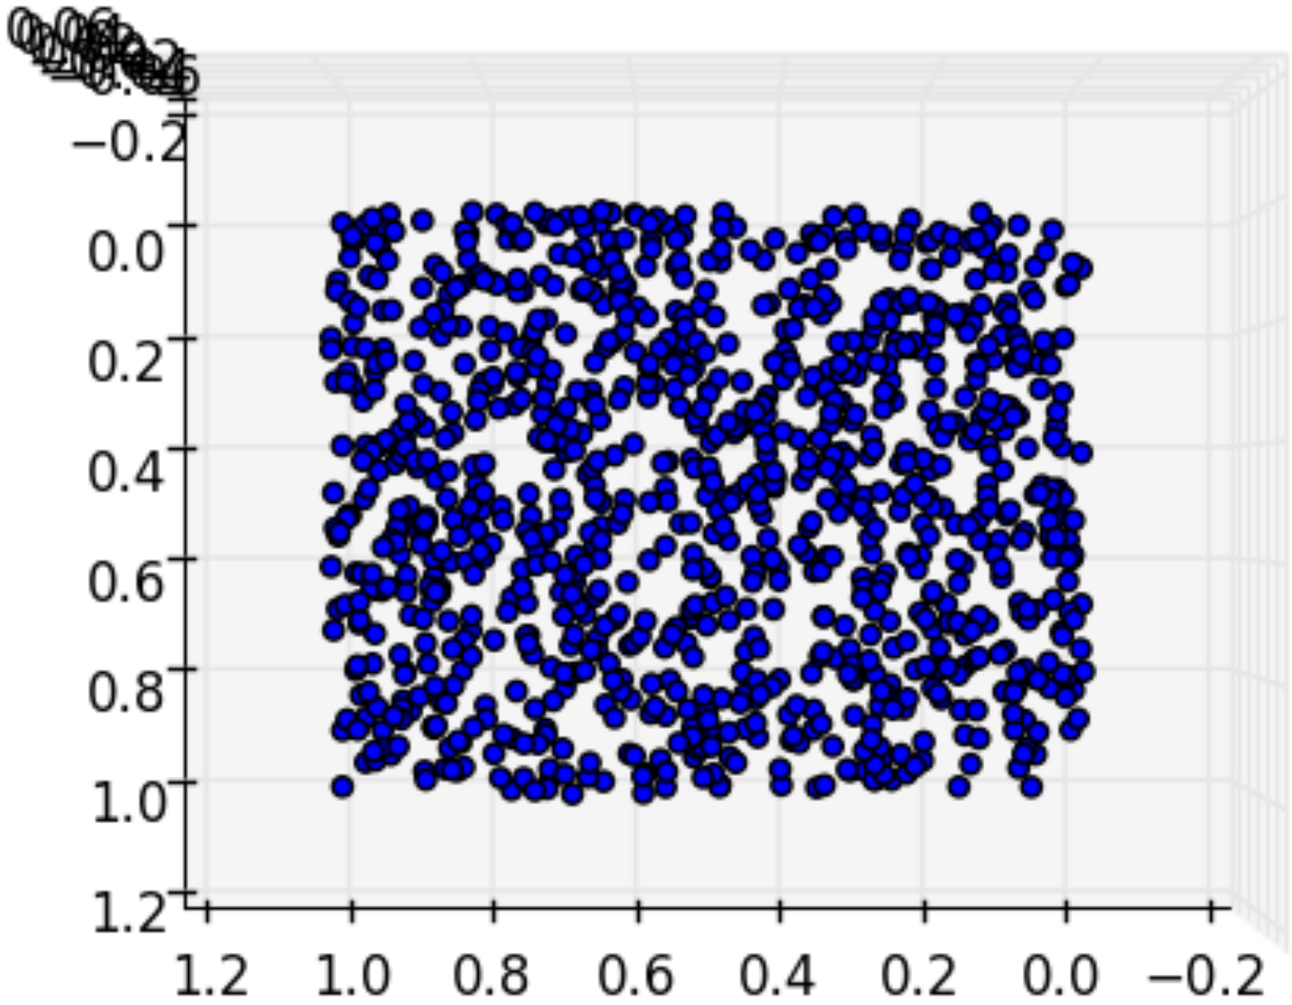
−0.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2

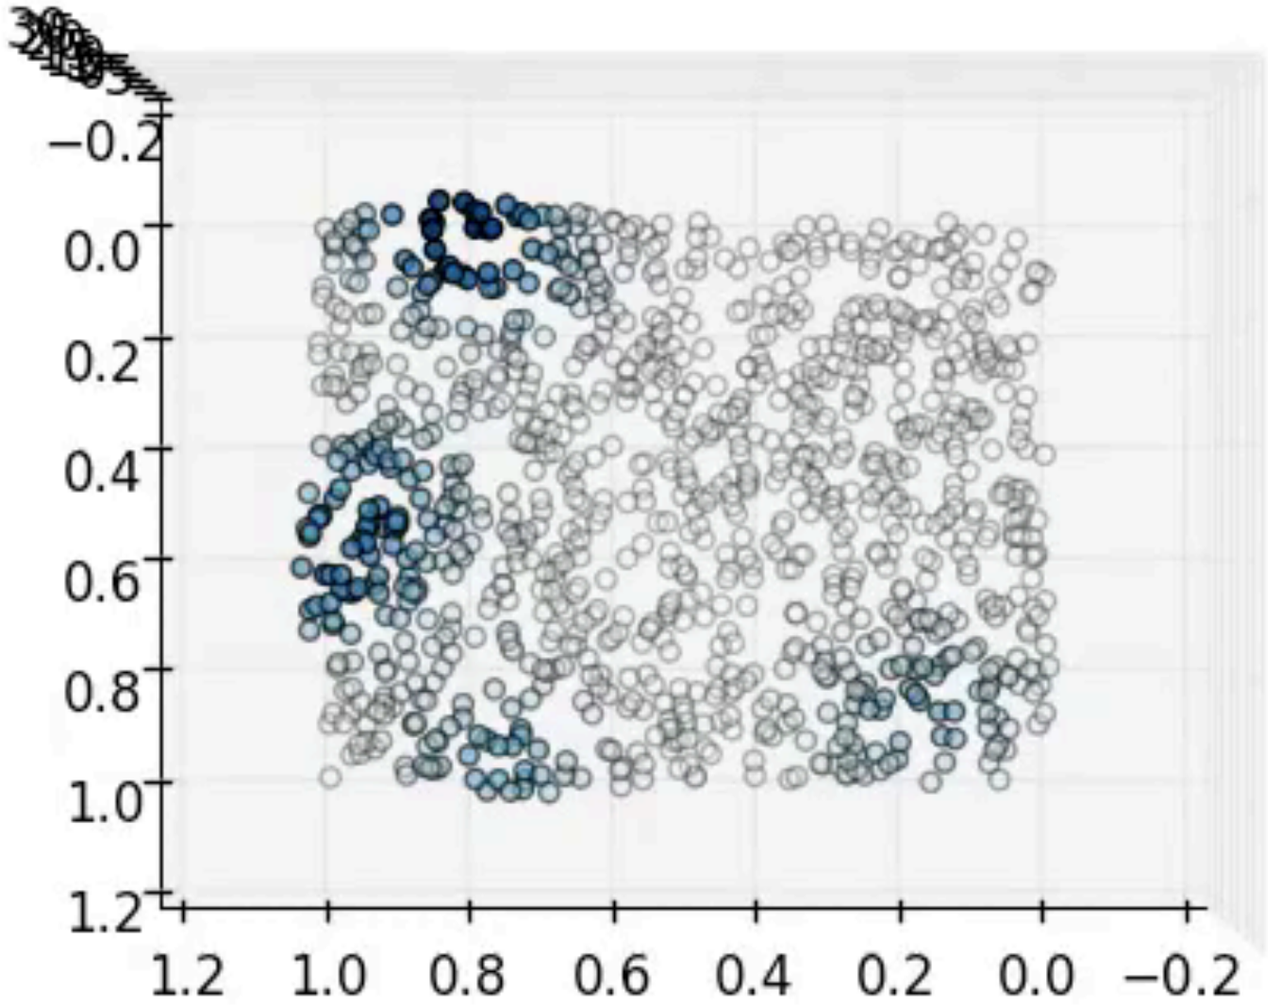
−0.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2

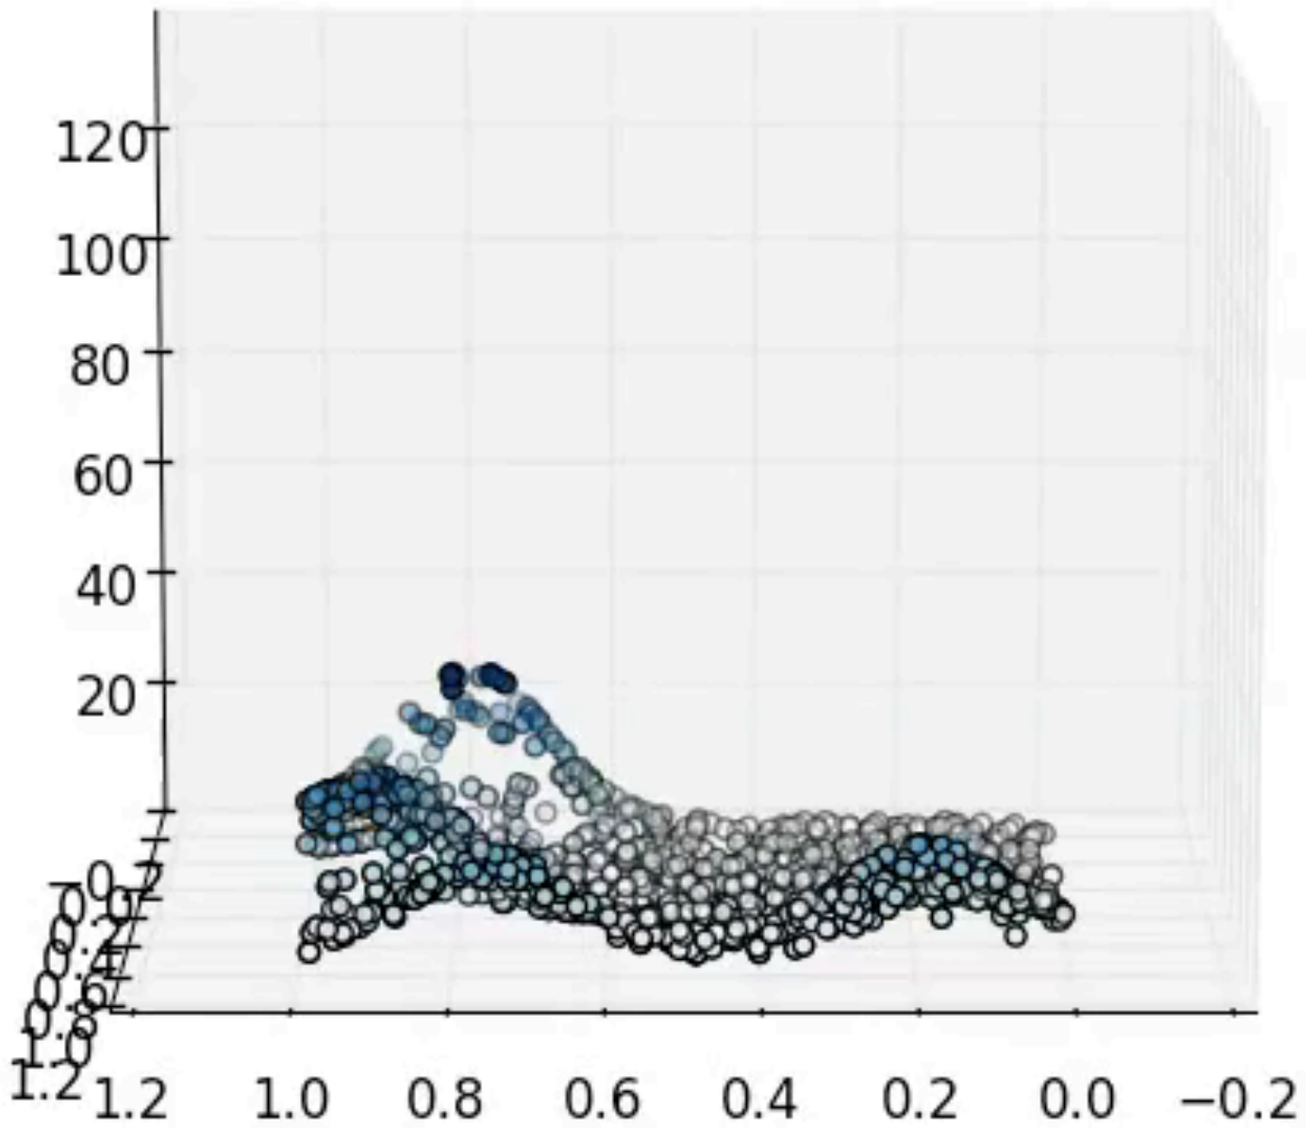
120
100
80
60
40
20
0
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2

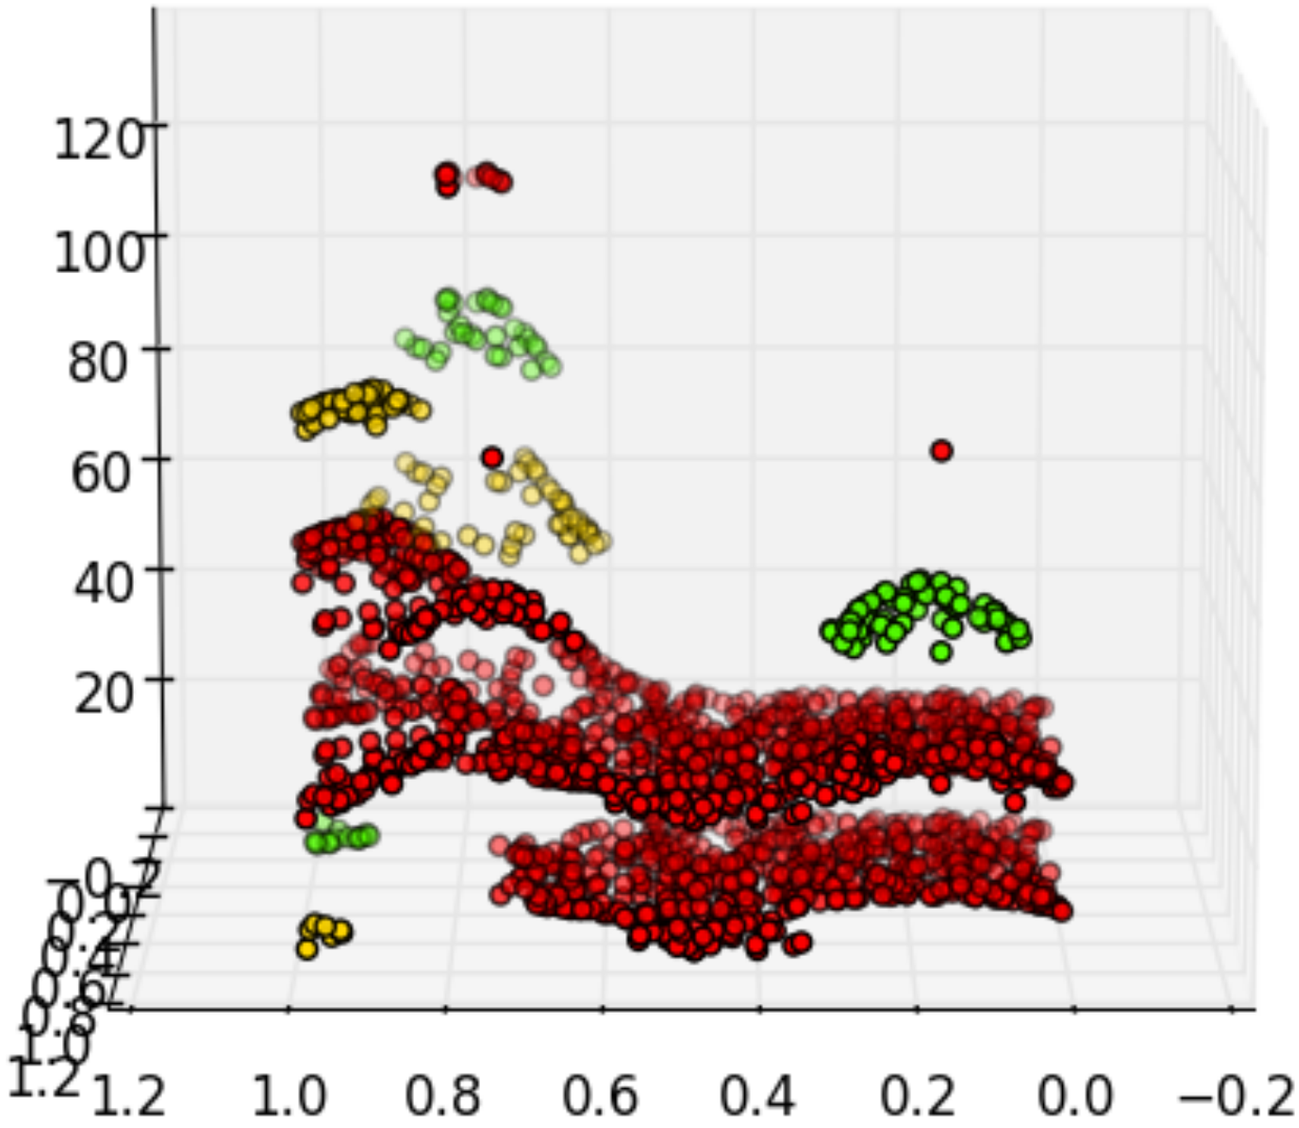
120
100
80
60
40
20
0
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2

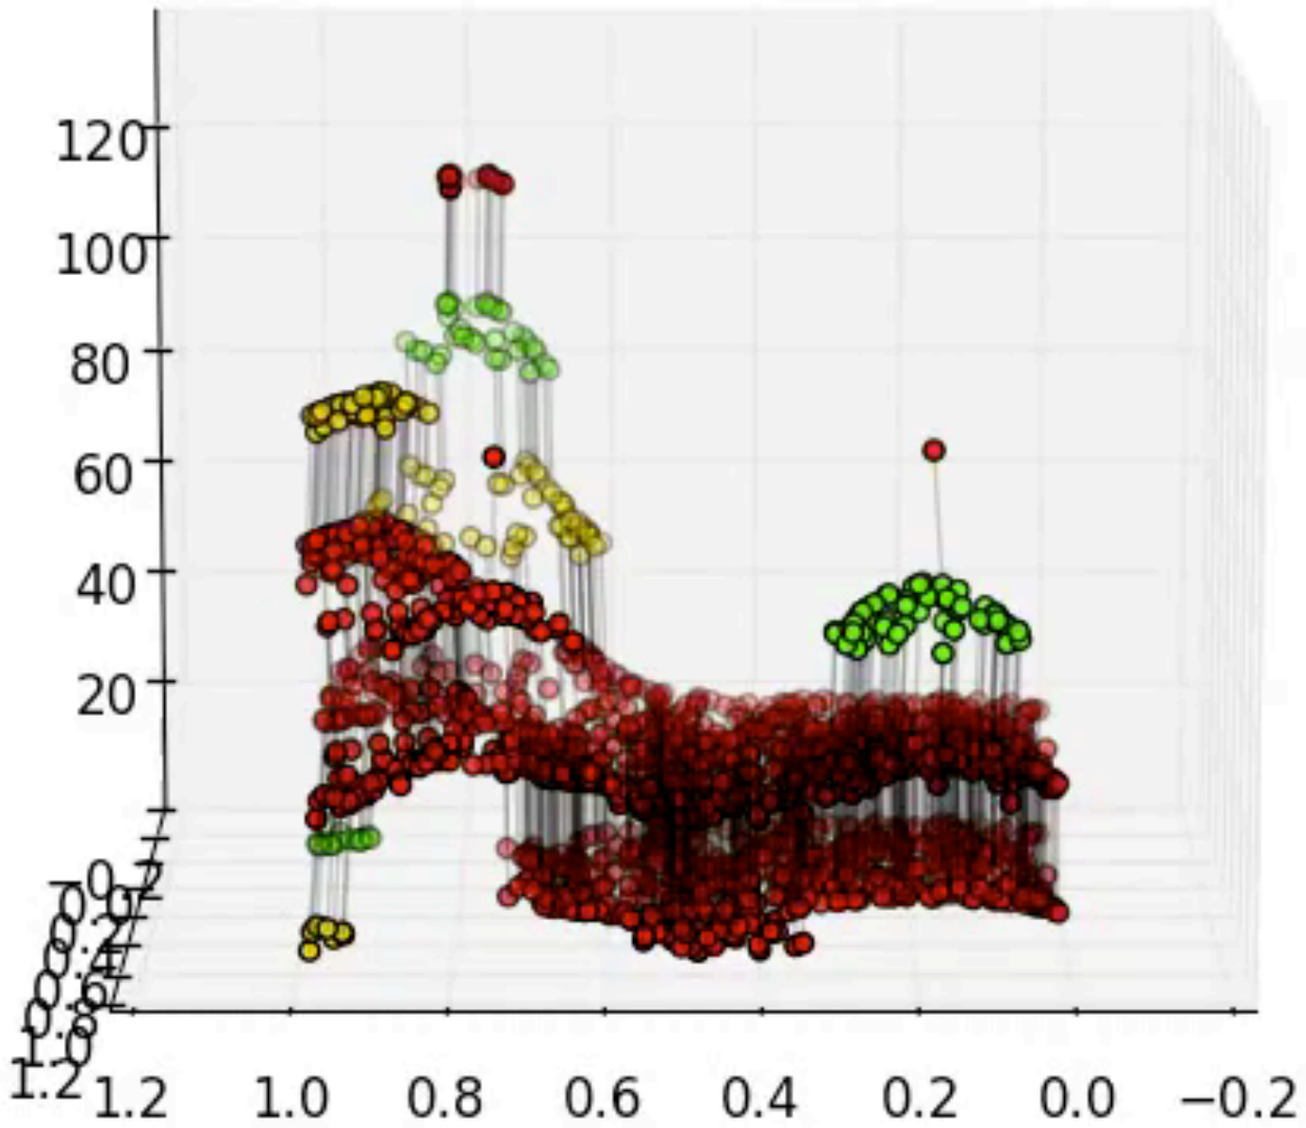
120
100
80
60
40
20
0
1.2
1.0
0.8
0.6
0.4
0.2
0.0
−0.2

Predictive Processes

- Regression (continuous predictions)
  Classification (discrete predictions)

- Functions from data to probability distribution or summary statistic

- P_θ(input) → prediction

Predictive Processes

- All observed inputs: point cloud, sampled from all possible inputs

- Observation of prediction and ground truth yields: (input, prediction, outcome) tuples

- Training data set

Clairvoyant Mapper

- Train a Mapper model using
    - Only input as data
    - Prediction error (and prediction/ground truth) as filter
- Separates inputs on the errors they eventually make
- New inputs can be matched against Mapper model

Fibres of Failure

(input, prediction, outcome)

Mapper
input as data
prediction-outcome as filter

Mapper model
Identify high error flares (failure modes)

Quantitative

Adjusted predictive process:
$Q_\theta(x) = P_\theta(x)$ + flare adjustments

Qualitative

Flare investigation:
what characterizes a failure mode?
Feedback to predictive modeling

Experiment:
MNIST Digit Recognition

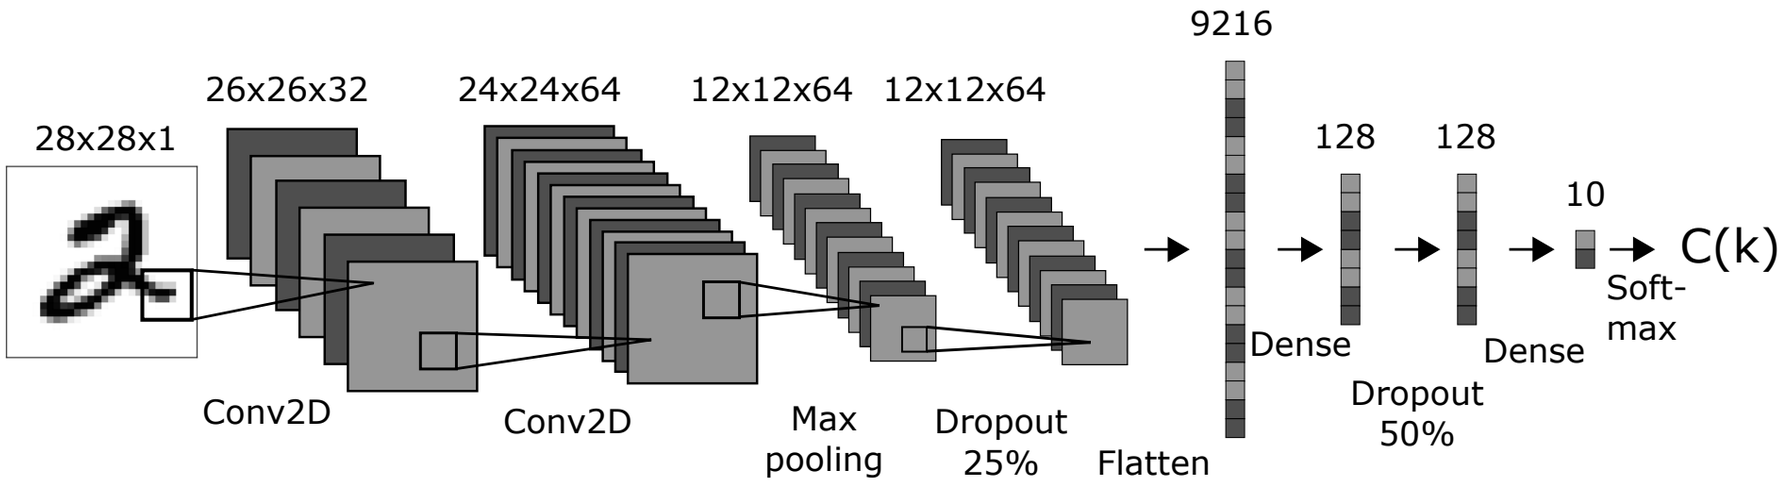

- Trained a simple CNN on recognizing hand-written digits

- CNN accuracy 99% on test (new) data

Let’s make it more difficult

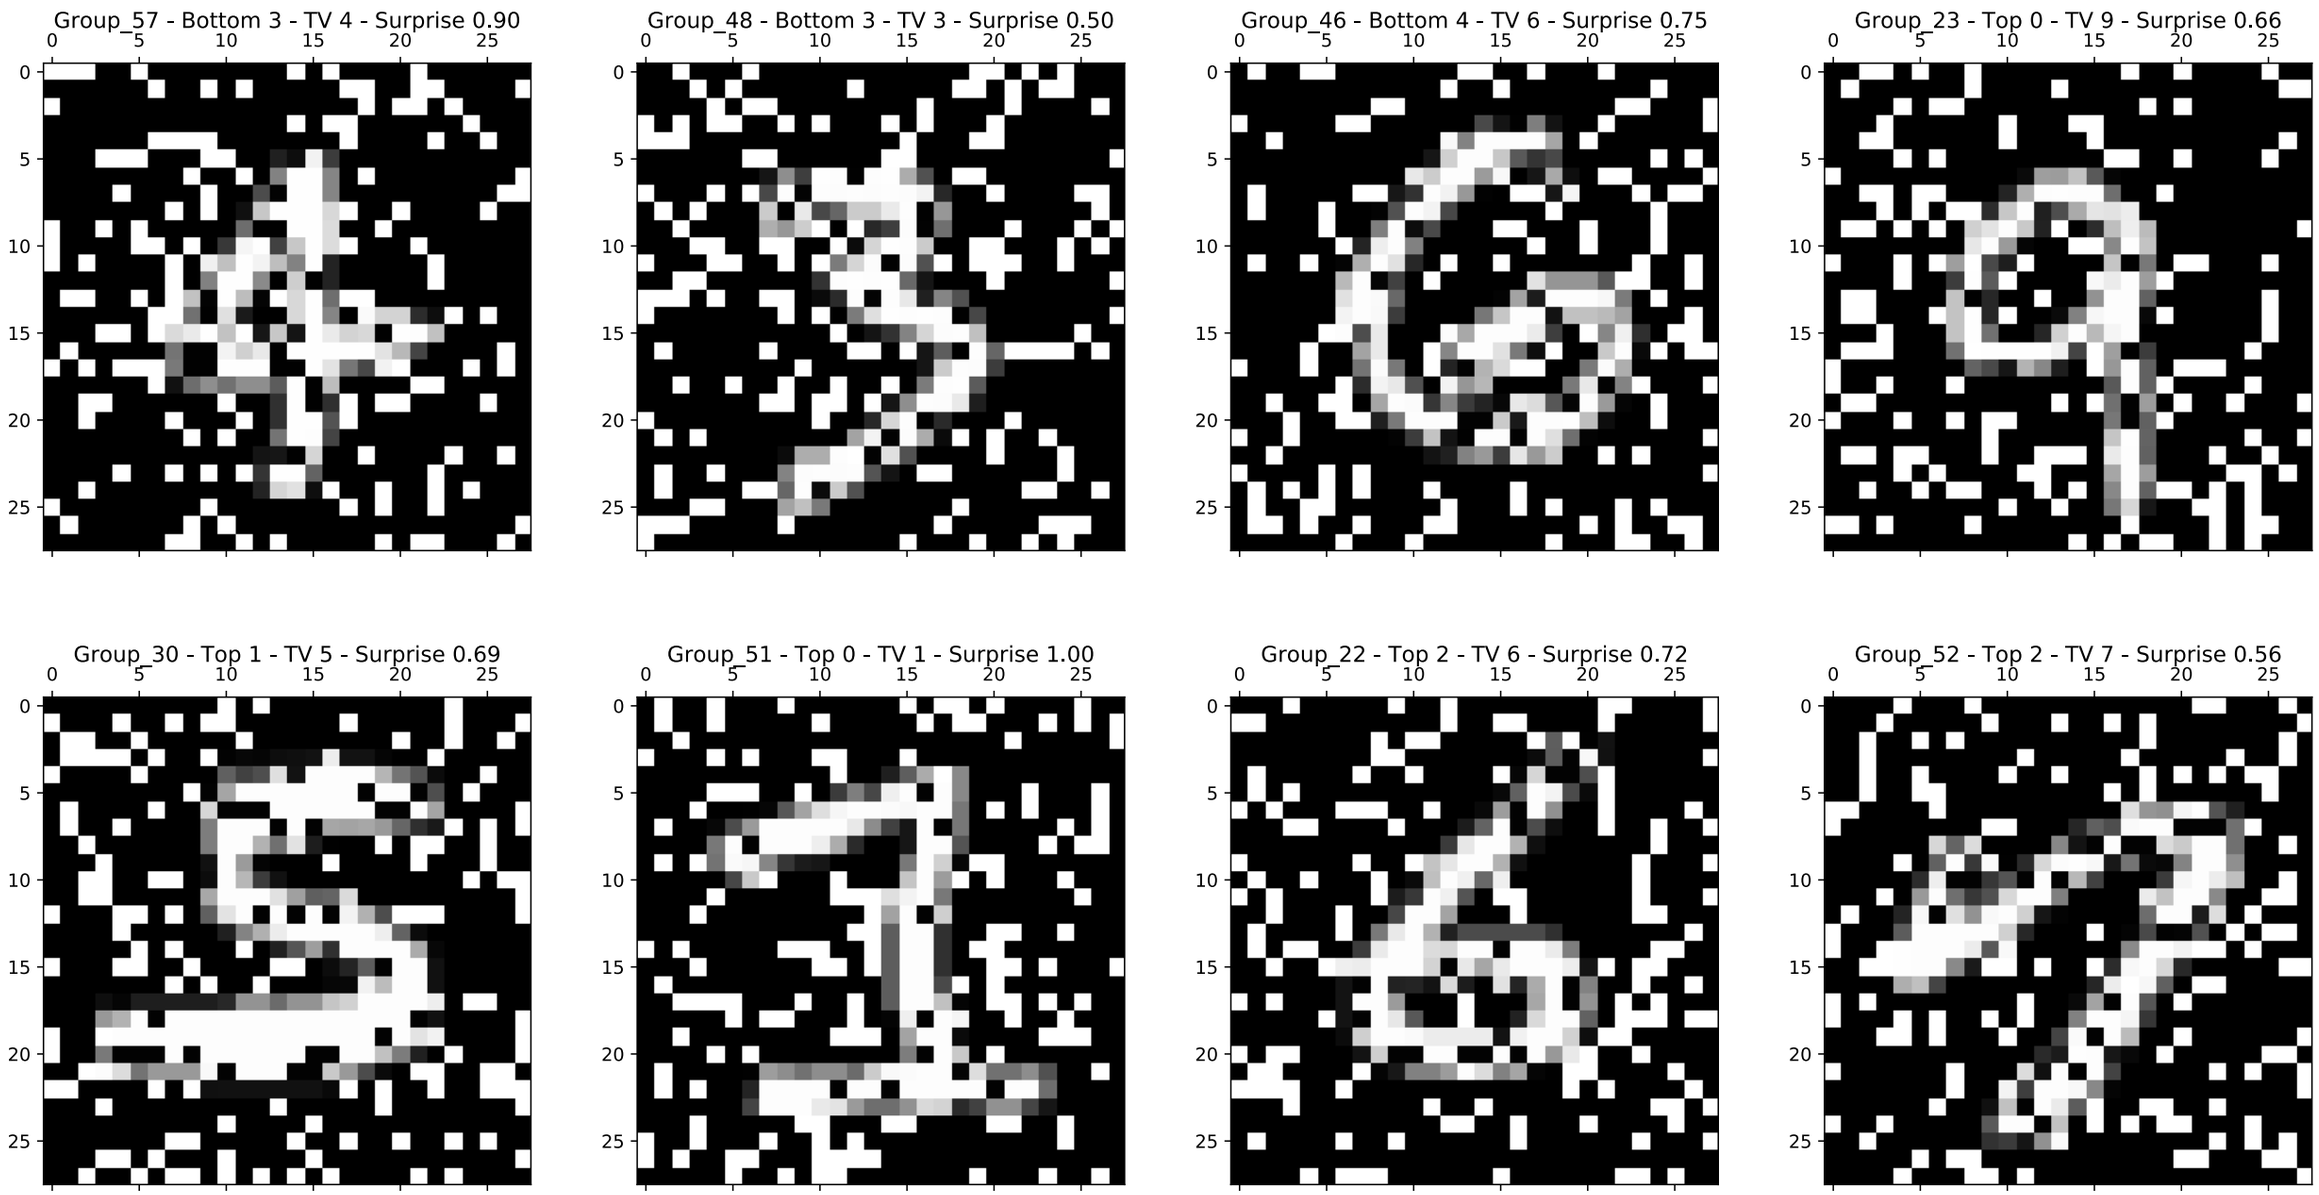

- Add 25% salt/pepper noise: flip pixels to pure black or white

- CNN accuracy 40.9% on corrupted data

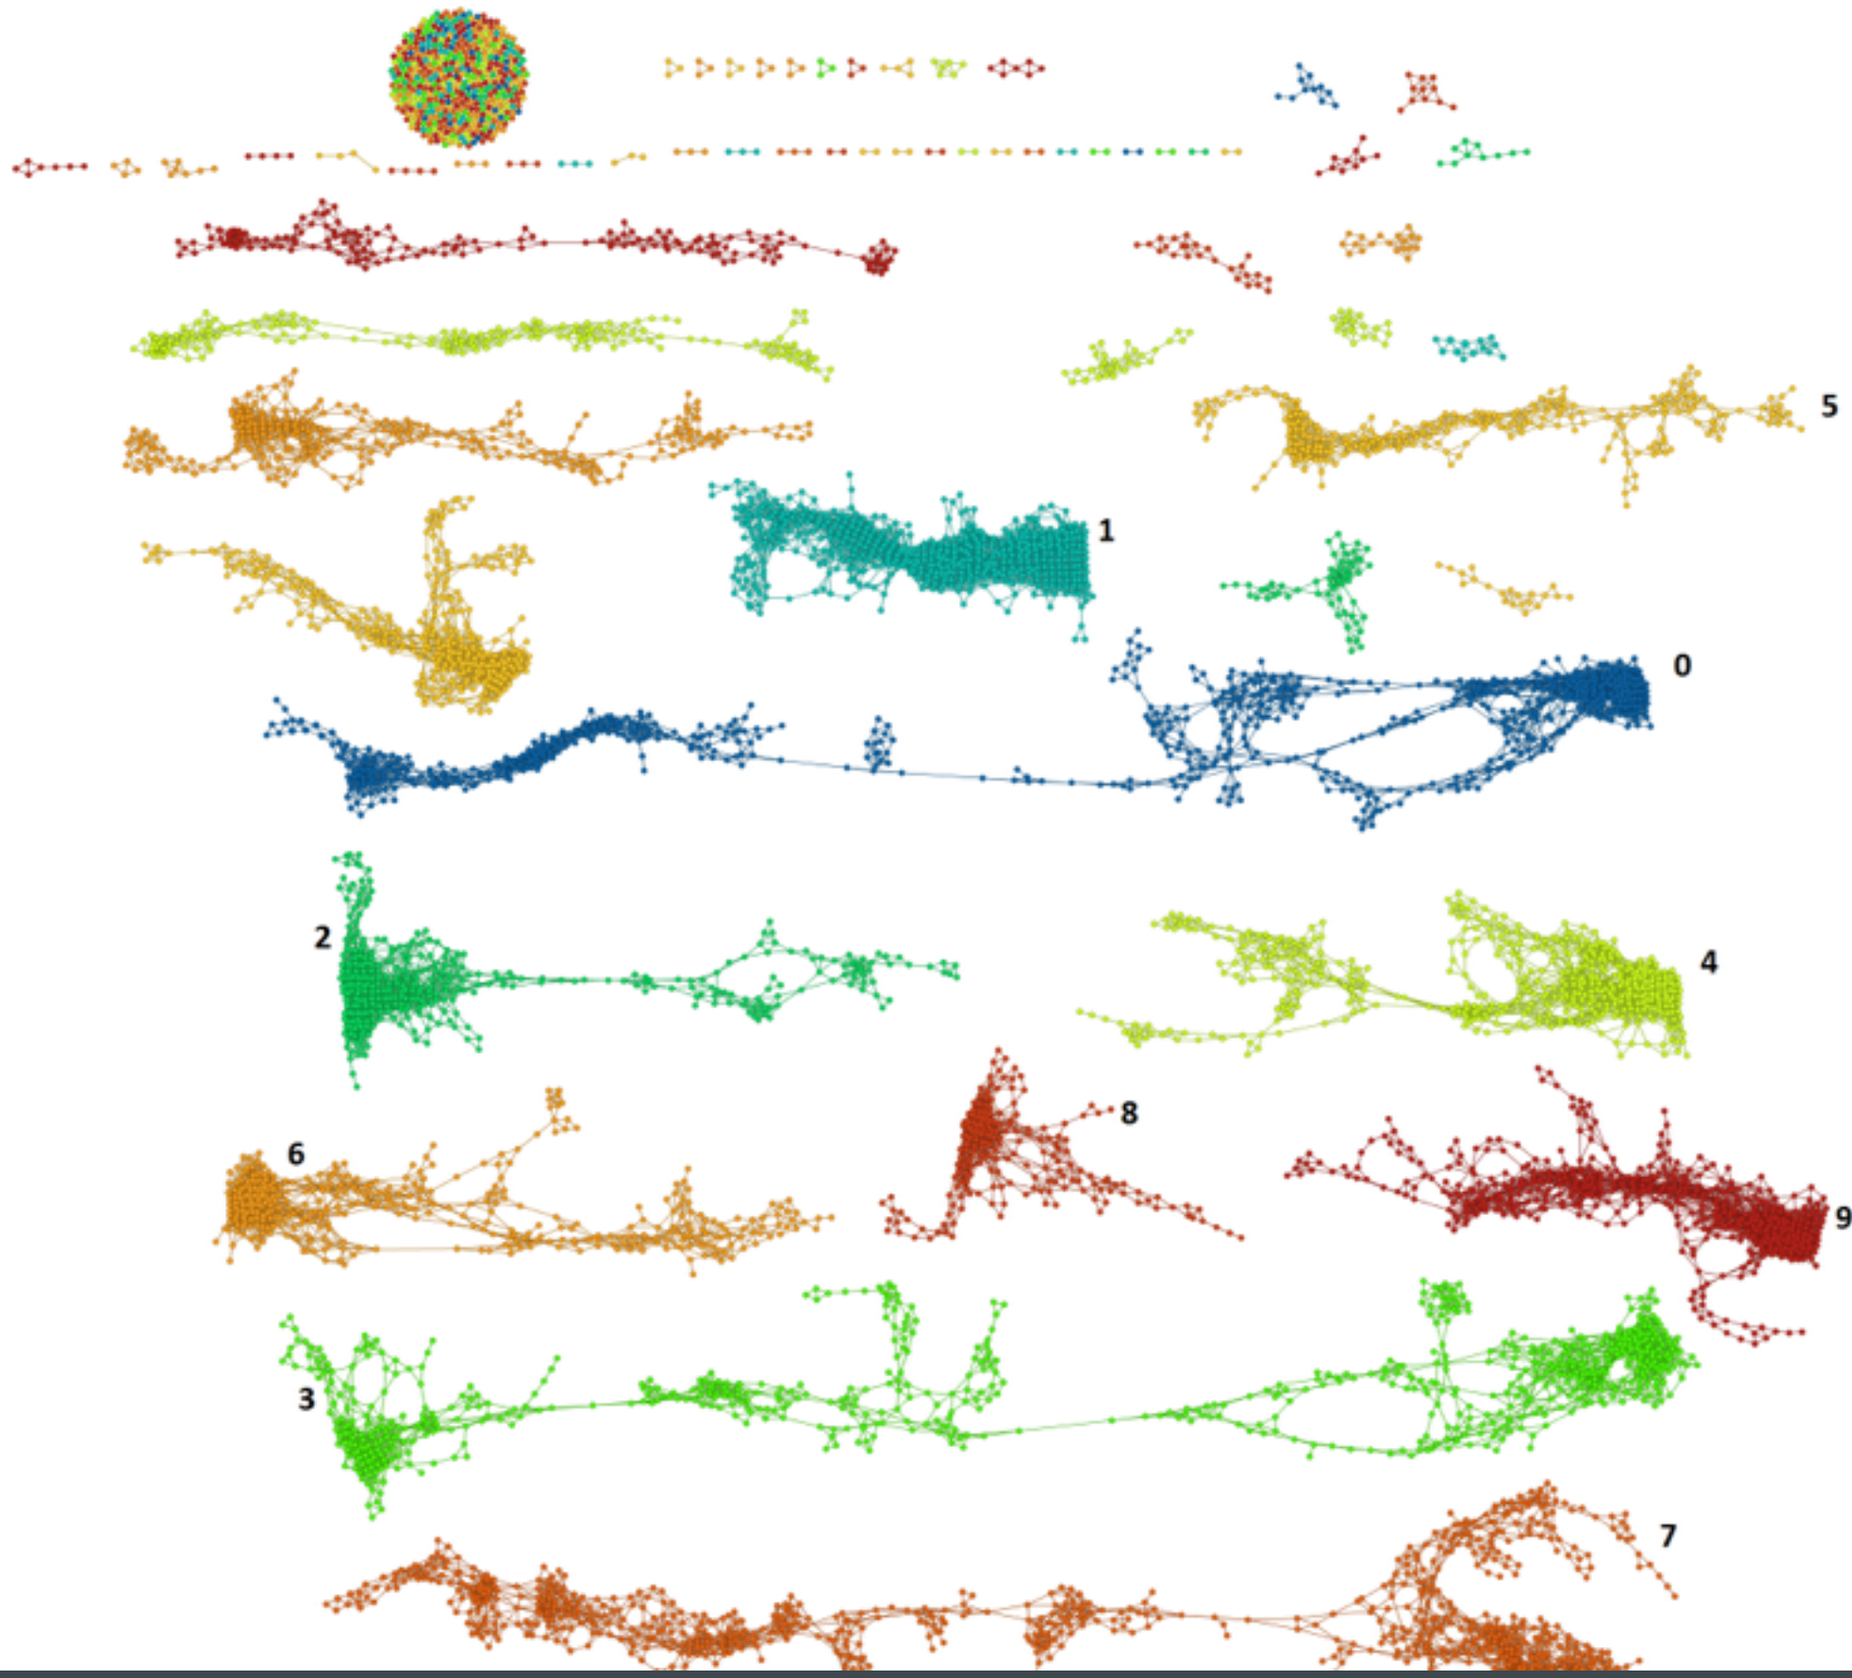
5
1
0
2
4
8
6
9
3
7

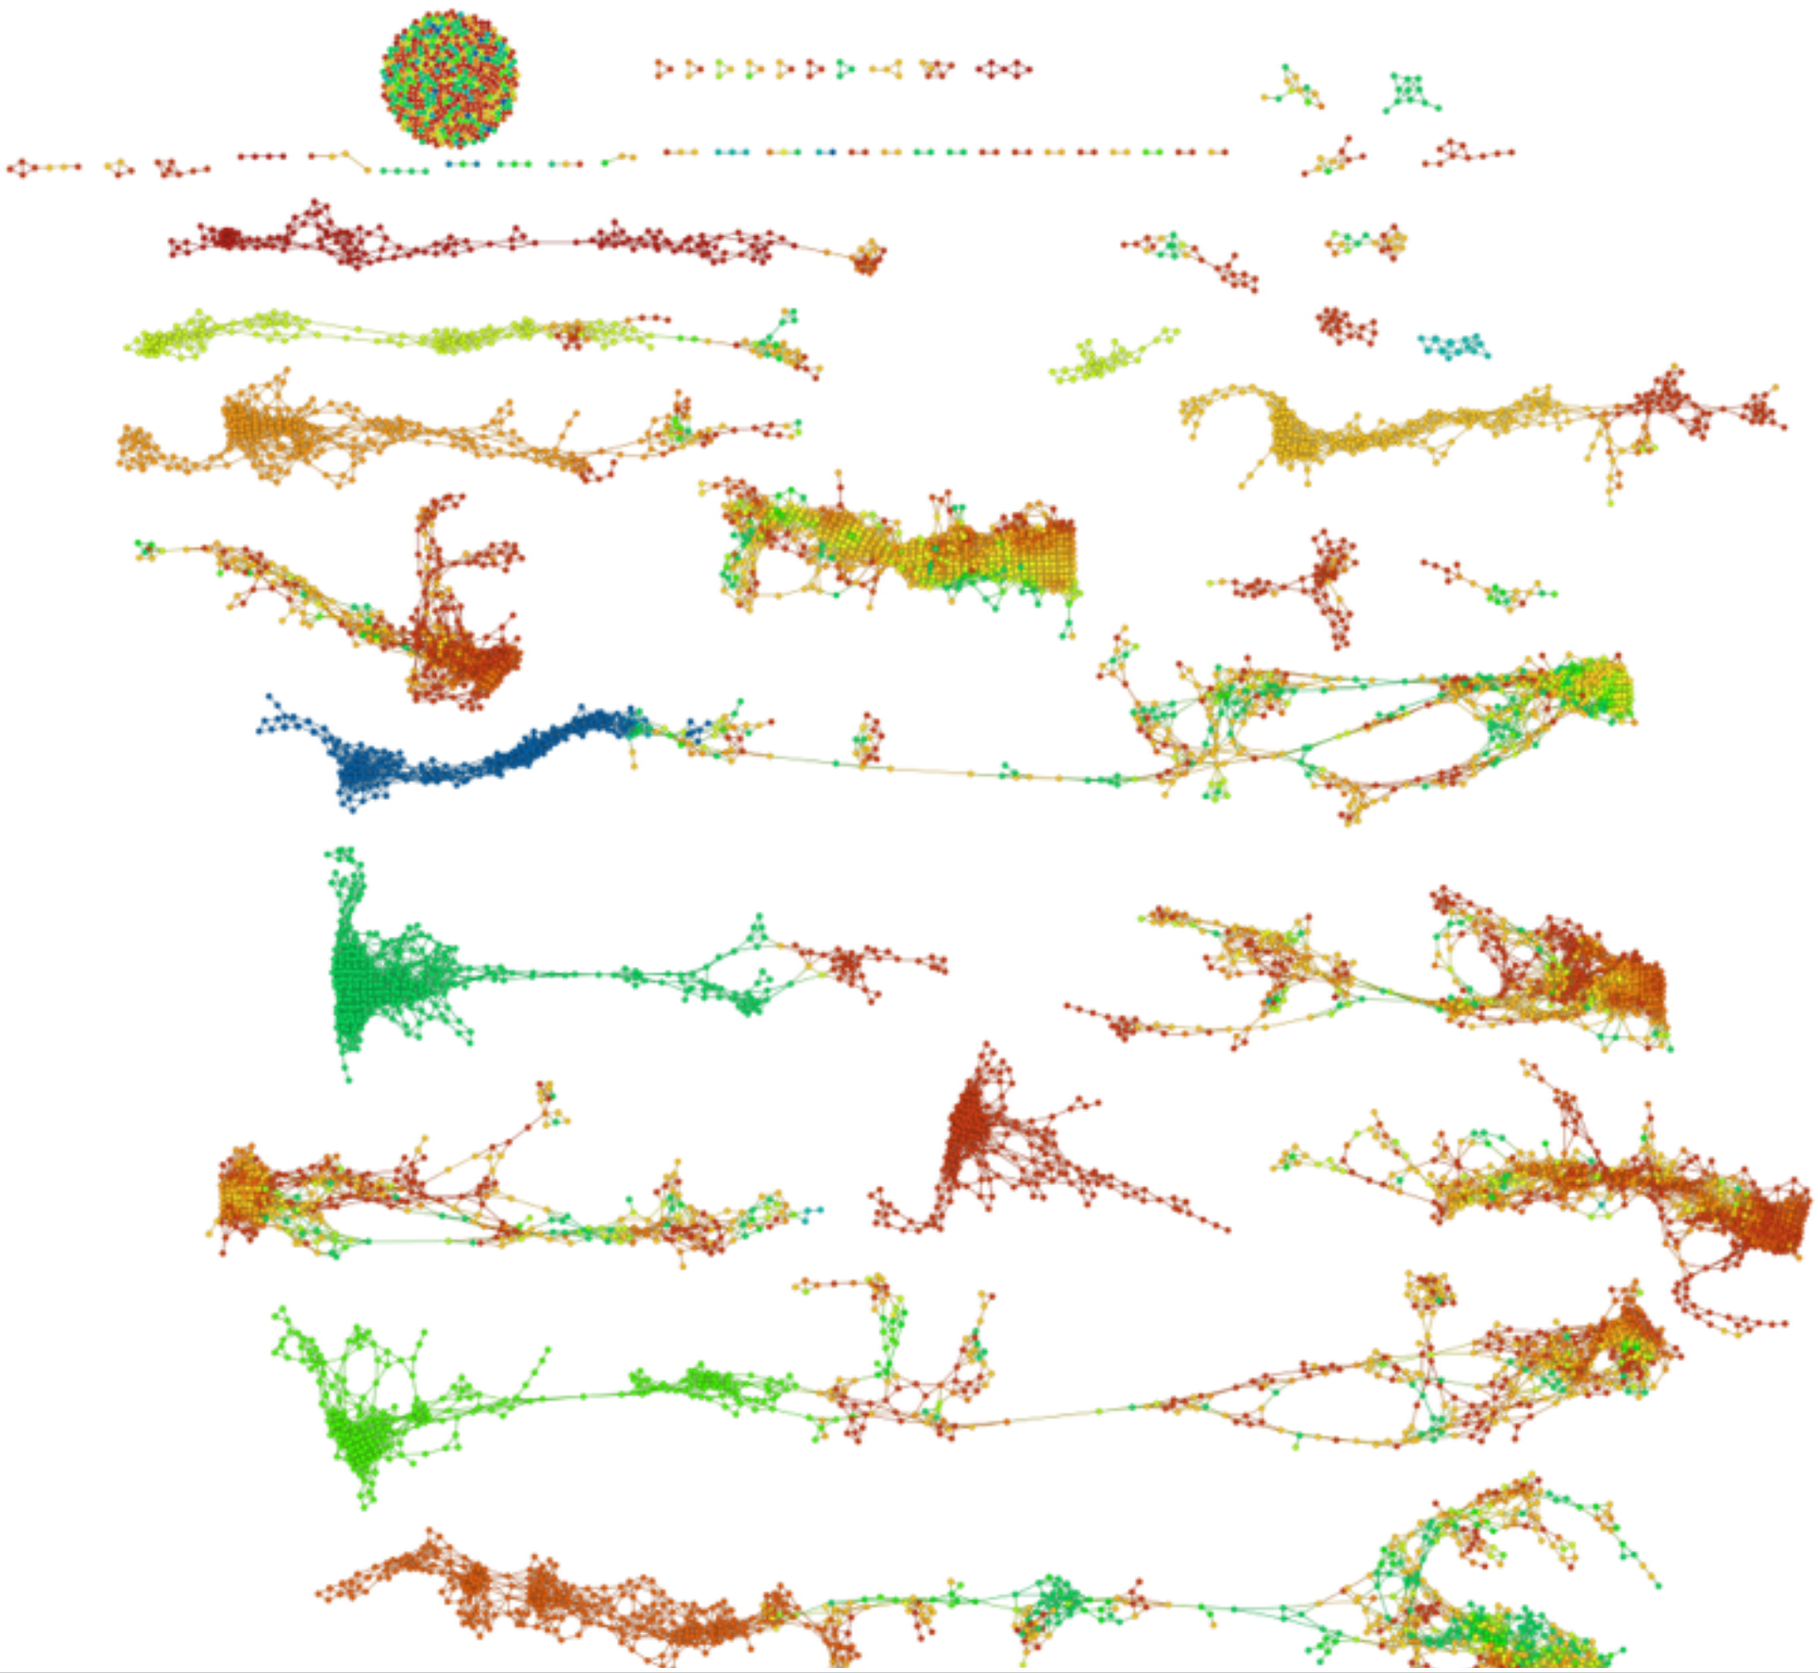

Quantitative

- Identified 39 high error groups
  Consistent ground truth within each group
  Cover ~30% of all corrupted images

- Trained one-vs-rest linear classifier ensemble to recognize failure modes

- Replace prediction with known group ground truth

    - Overall accuracy: 64.5% (up from 40.9%)

    - CNN accuracy on recognized failure mode members: 16.1%

    - Group ground truth accuracy: 70%-90%

Qualitative

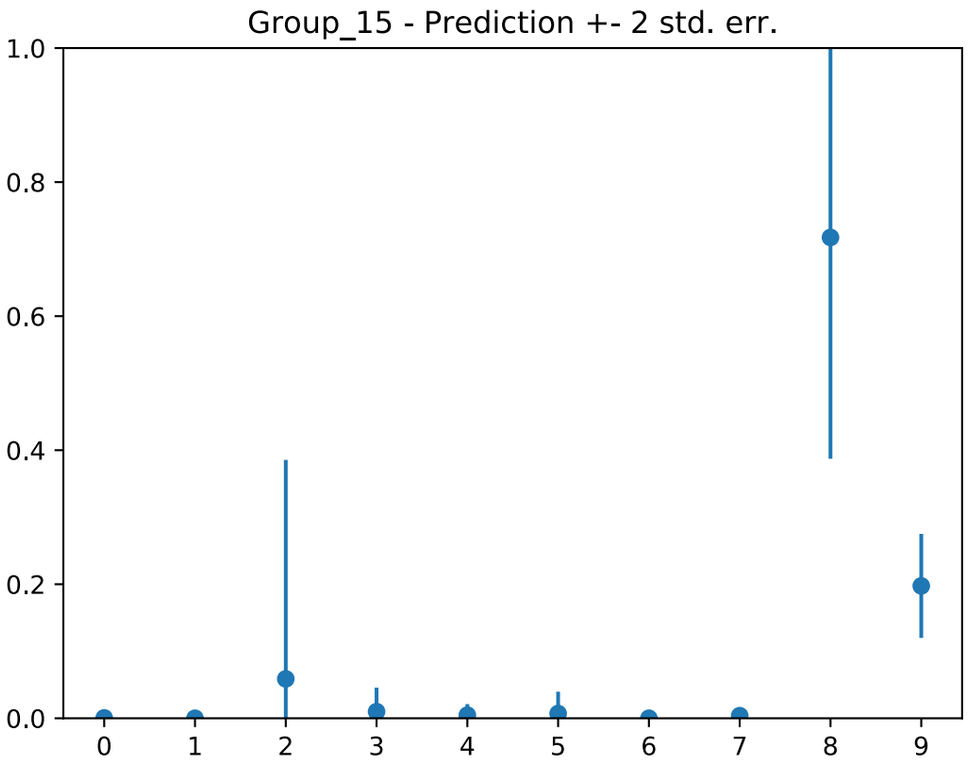

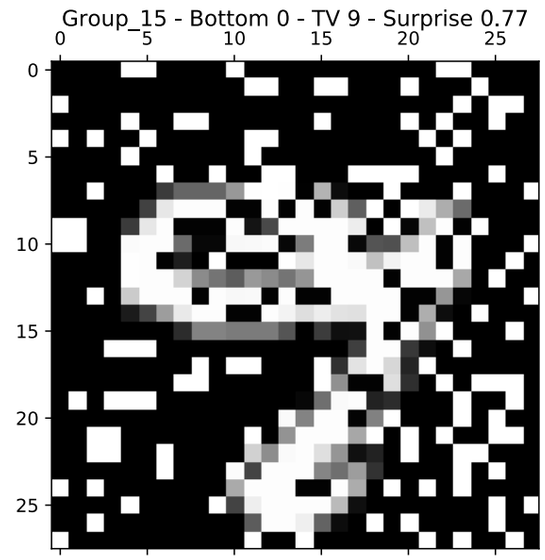

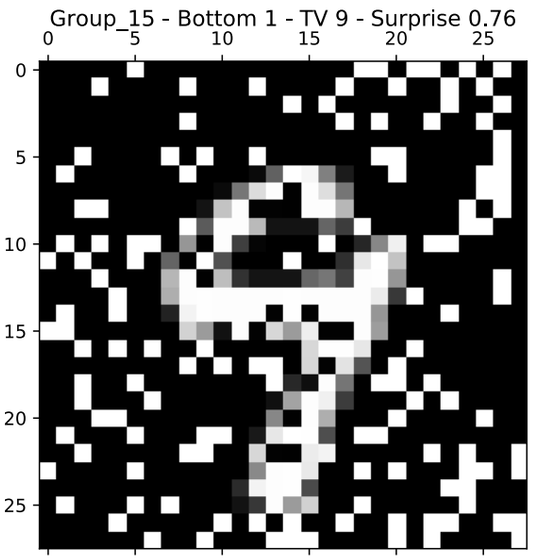

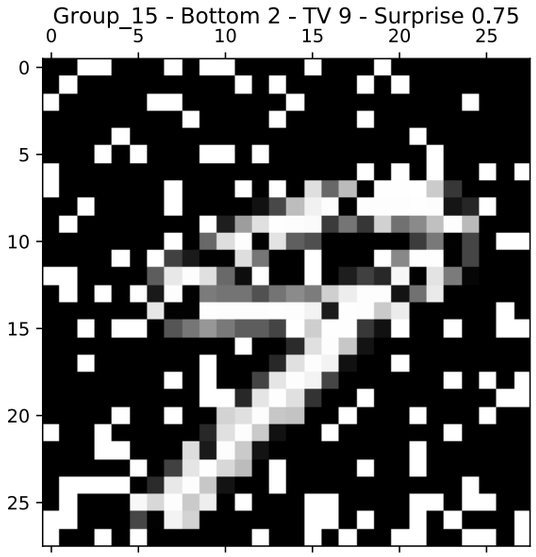

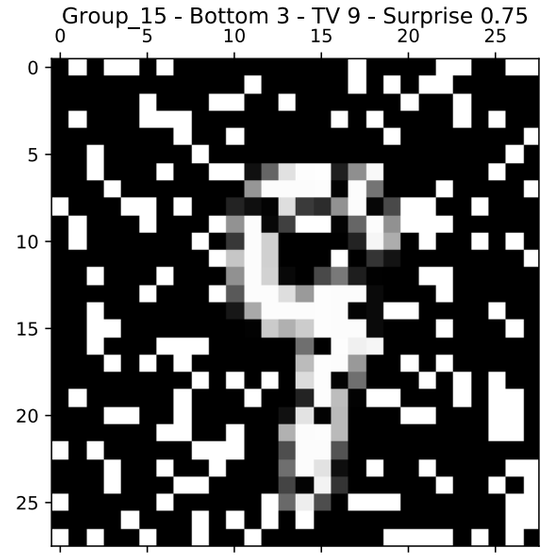

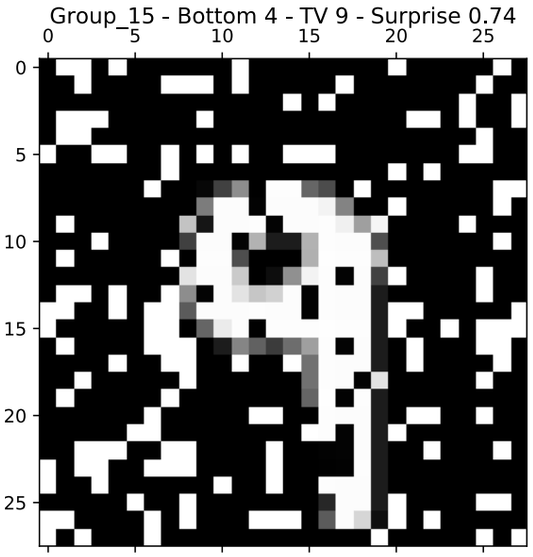

Thank you for listening

- **Fi**bres of **Fa**ilure: Classify failure modes
    - Mapper with failure measure as a filter function
    - Identify high failure flares
        - Inspect failure modes qualitatively
        - Generate ensemble classifier to adjust original predictions