3 Reasons Why Anomaly Detection is Hard

Maja Rudolph
UW-Madison

Hackathon 12/4/24

Anomaly Detection

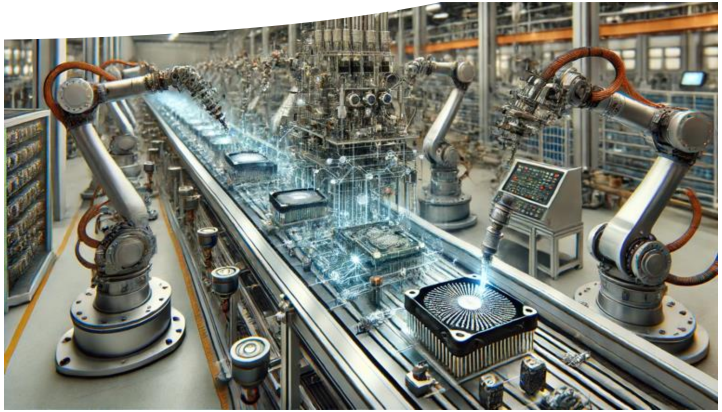

Visuals created with ChatGPT

Anomaly Detection

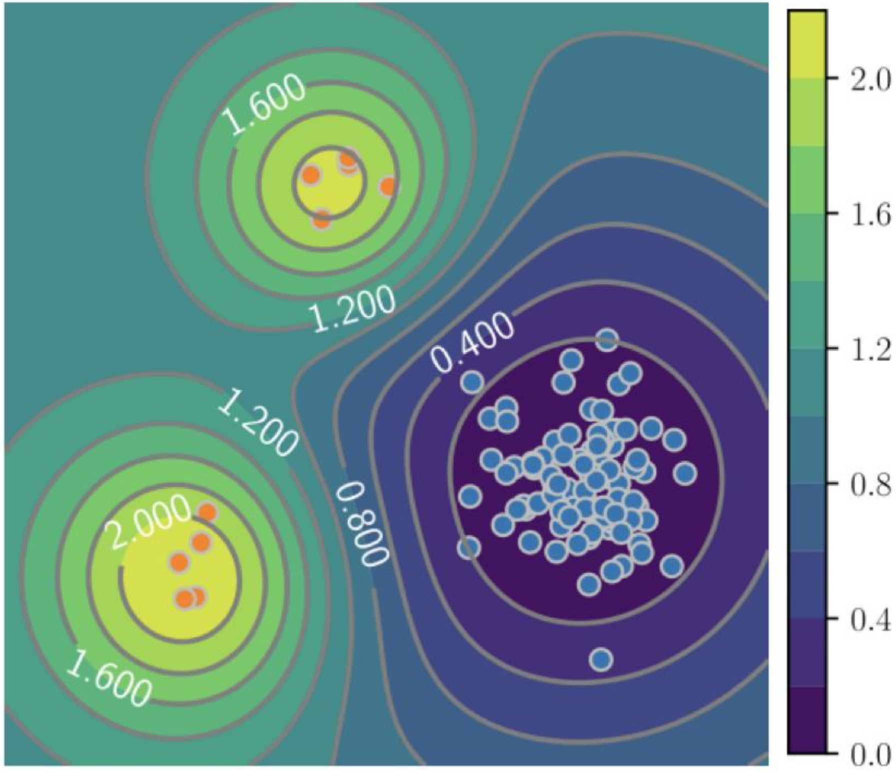

- most data representative of what is considered "normal"
- some sample differ in a systematic way, they are the "anomalies"
- given unlabeled data, goal is to learn an anomaly scoring function

$$S(\boldsymbol{x}; \theta)$$

Anomaly Detection

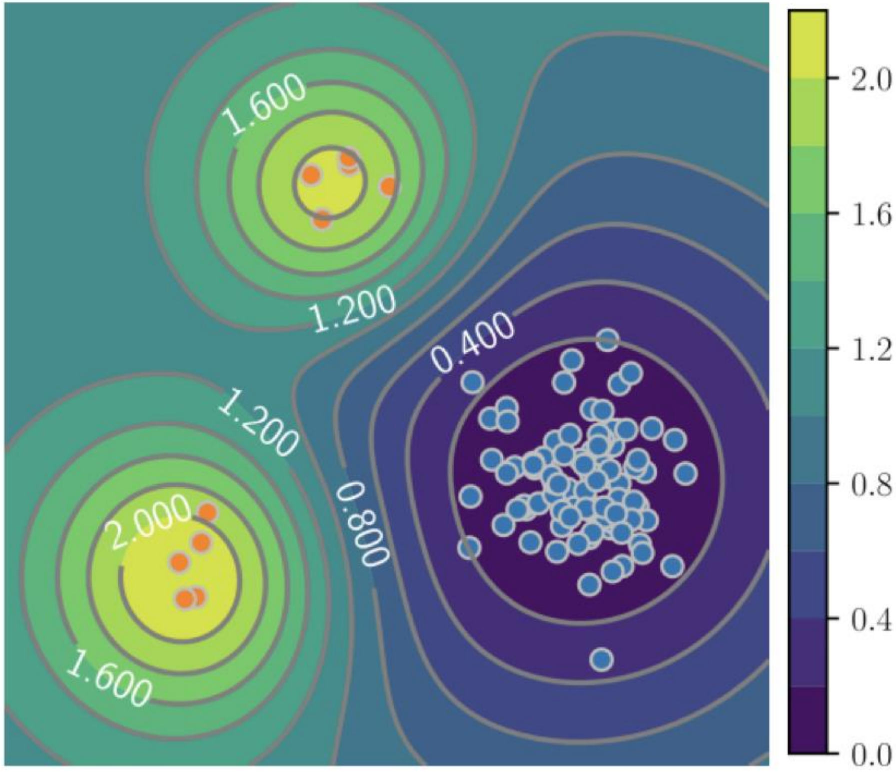

AD is statistically scary!

The Three Challenges

- Distribution Shifts
- High Dimensions
- Limited Labels

Challenge 1: Distribution Shifts

- Distribution of normal data

$$x \sim p(x)$$

- Distribution of anomalies

Distribution of anomalies keeps shifting

$$x \sim p_a(x, t)$$

Challenge 1: Distribution Shifts

- Training data

$$\mathcal{D}_{train} = \{x | x \sim p(x)\}$$

- At inference time

$$\mathcal{D}_{test}(t) = \{(x, y) | y \sim Bern(\alpha), x \sim y p_a(x, t) + (1 - y) p($$

Challenge 2: High-dimensional Data






Challenge 3: Limited Labels

- Training data unlabeled
- Mostly normal
- Possibly contaminated

Challenge 3: Limited Labels

- Training data unlabeled

Tip1: Try Self-Supervised AD

Self-Supervised Anomaly Detection

unsupervised learning
$p(x)$

supervised learning
$p(y|x)$

Self-Supervised Anomaly Detection

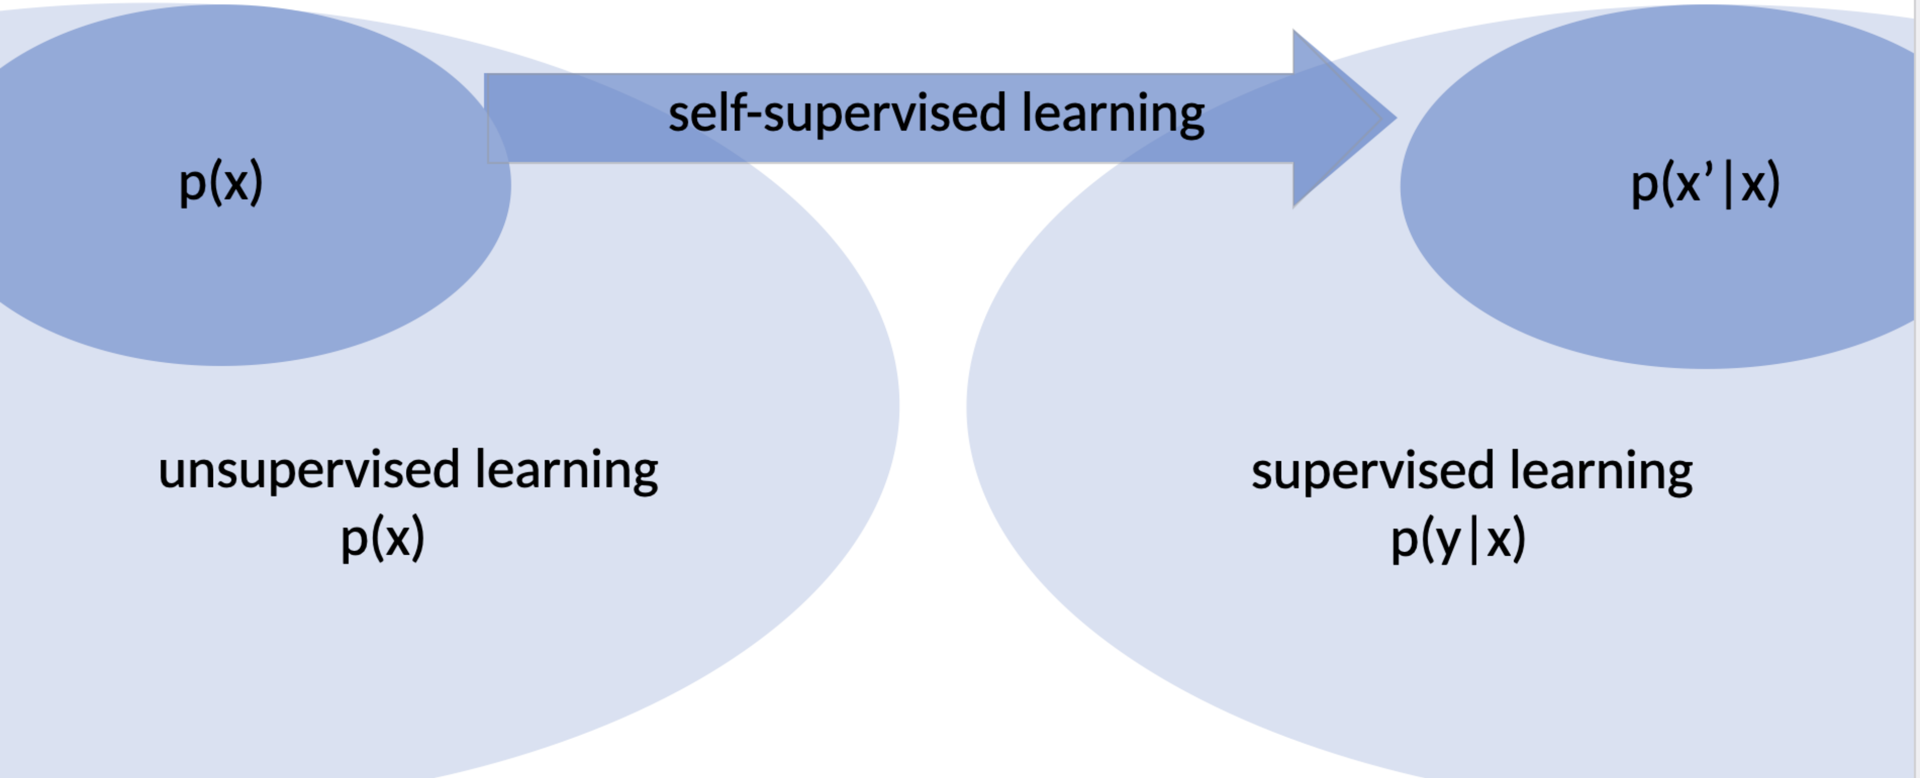

Self-Supervised Anomaly Detection

- performance on auxiliary task as anomaly scoring function

$$S(\boldsymbol{x}_i;\theta) \equiv \ell(\boldsymbol{x}_i;\theta)$$

- training objective

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N}\ell(\boldsymbol{x}_i;\theta)$$

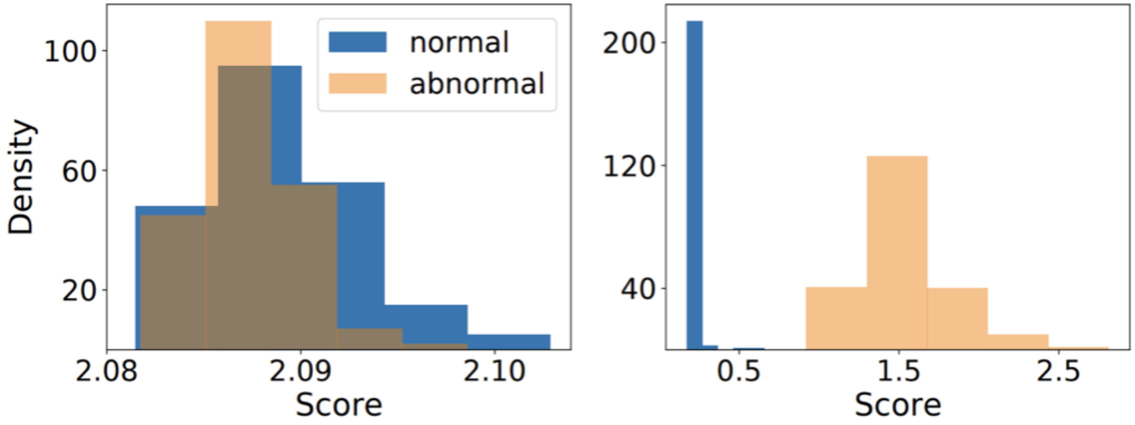

(a) histogram before training (b) histogram after training

Self-Supervised Anomaly Detection

- Neural Transformation Learning (ICML '21)

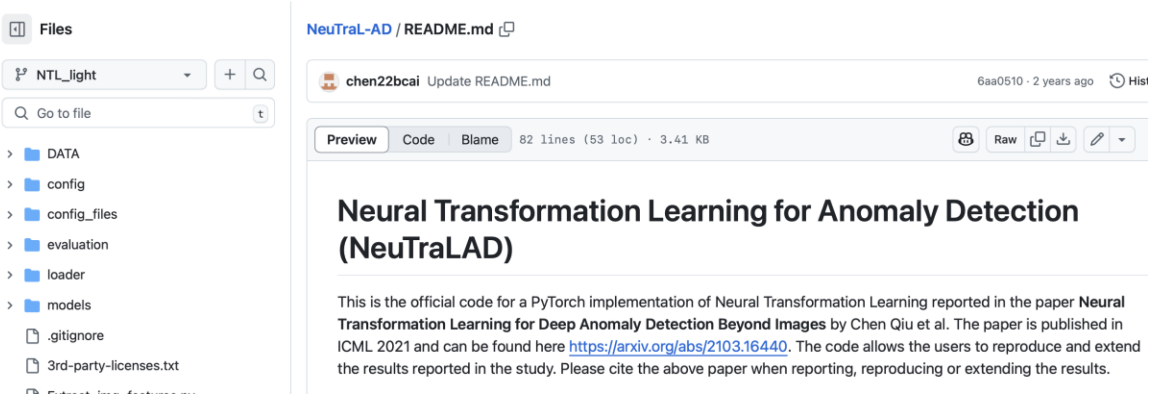

- Robust detection performance for many different data types

A Revealing Large-Scale Evaluation of Unsupervised Anomaly Detection Algorithms

Maxime Alvarez*, Jean-Charles Verdier*, D'Jeff K. Nkas
Marc Frappier, Pierre-Martin Tardif, Froduald Kabanz

GRIC, Université de Sherbrooke
Sherbrooke, QC, Canada
{maxime.alvarez,jean-charles.verdi
marc.frappier,pierre-martin.t

"[NTL], the transformation-based approach, offers consistently above-average performance across all datasets"

Abst

Anomaly detection has many applications ranging fr ... ion a...
cyber-threat detection to equipment maintenance and health monito ... However,

Challenge 3: Limited Labels

- Training data unlabeled
- Mostly normal
- Possibly contaminated

Tip 1: Try ~~Self-supervised AD~~ a Foundation Model

Anomaly Detection with Foundation Models

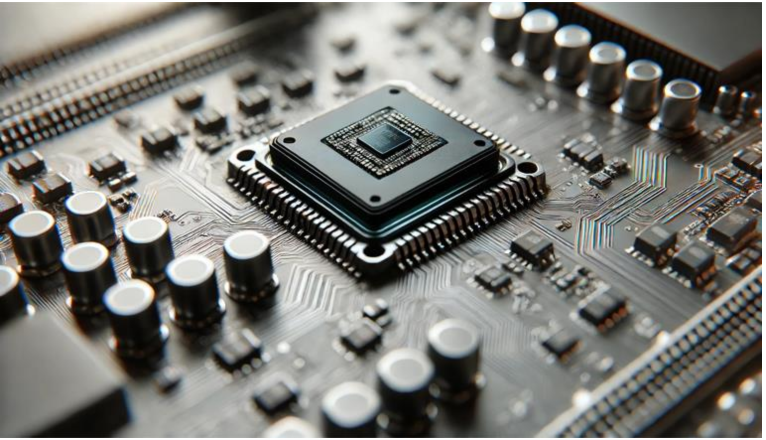

“A close-up of a damaged chip”

“A close-up of a functioning chip”

Liznerski, P.; Ruff, L.; Vandermeulen, R.A.; Franks, B.J.; Muller, K.R.; Kloft, M. Exposing Outlier Exposure: What Can Be Learned From Few, One, and Zero Outlier Images. Transactions on Machine Learning Research, 2022

Jeong, J.; Zou, Y.; Kim, T.; Zhang, D.; Ravichandran, A.; Dabeer, O. WinCLIP: Zero-/few-shot anomaly classification and segmentation. In CVPR, 2023

Challenge 3: Limited Labels

- Training data unlabeled
- Mostly normal
- Possibly contaminated
- **Model selection is key!!**
  (But impossible without labels?)

Tip 2: Create Synthetic Validation Data!

Generating Anomalies with Foundation Models

ChatGPT 4o >

Generate a picture of a circuit board with a flaw

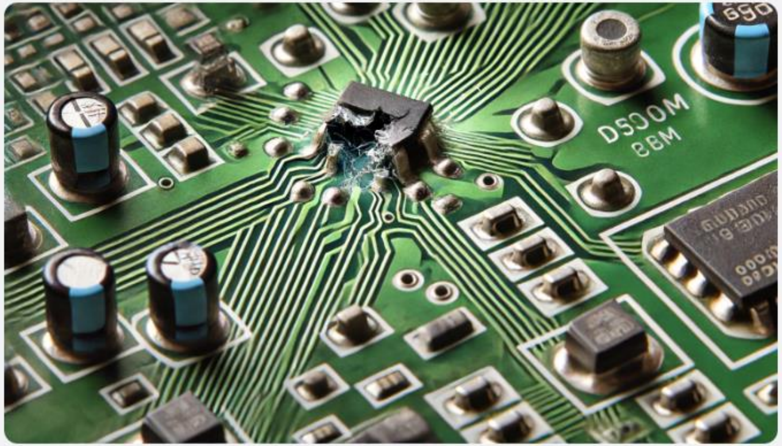

Model Selection with Synthetic Anomalies

Diffusion-guided anomaly generation

- Assumes access to a validation dataset of normal samples
- No training or fine-tuning
- No custom prompts

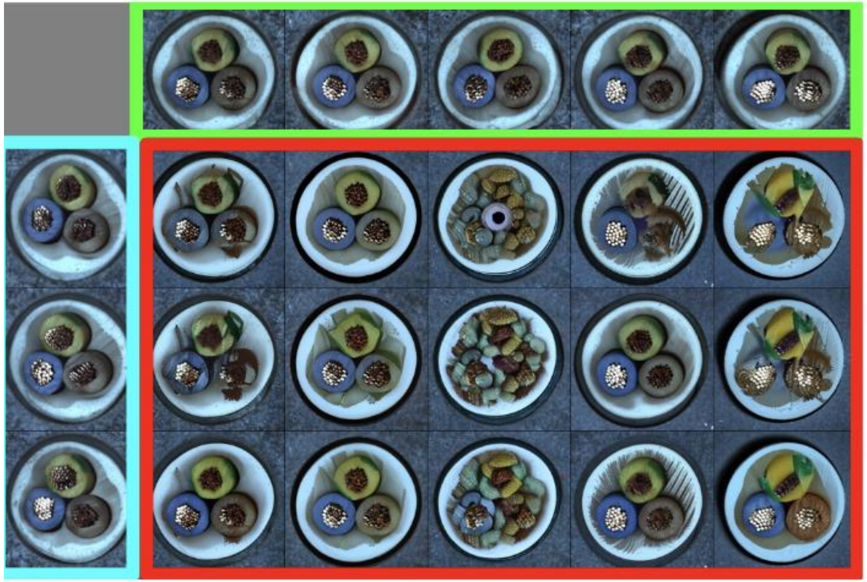

(b) MVTec-AD cable

Fung, Qiu, Li, **Rudolph**, [Model Selection of Anomaly Detectors in the Absence of Labeled Validation Data](https://arxiv.org/pdf/2310.10461), https://arxiv.org/pdf/2310.10461

Summary

- Distribution Shifts
- High Dimensions
- Limited Labels

Tip 3: Don't be Scared.
Have Fun!! :)

Thank You!

contact: maja.rudolph@wisc.edu