SAND2021-12343PE

Sandia National Laboratories

Anomaly Detection in Images

Presented By: Abigail Pribisova

Objective

- Automate anomaly detection in images
  1. Identify images that contain anomalies
  2. Identify locations of anomalies in images

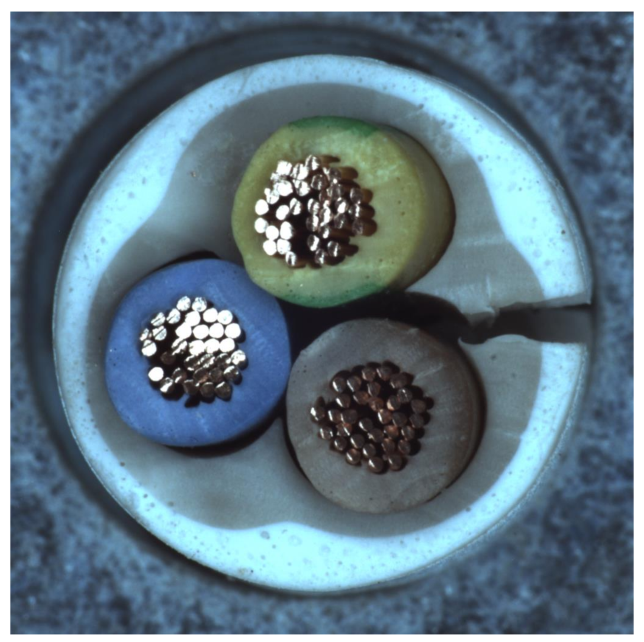

Image of cable with a cut in the outer insulation from the MVTEC dataset

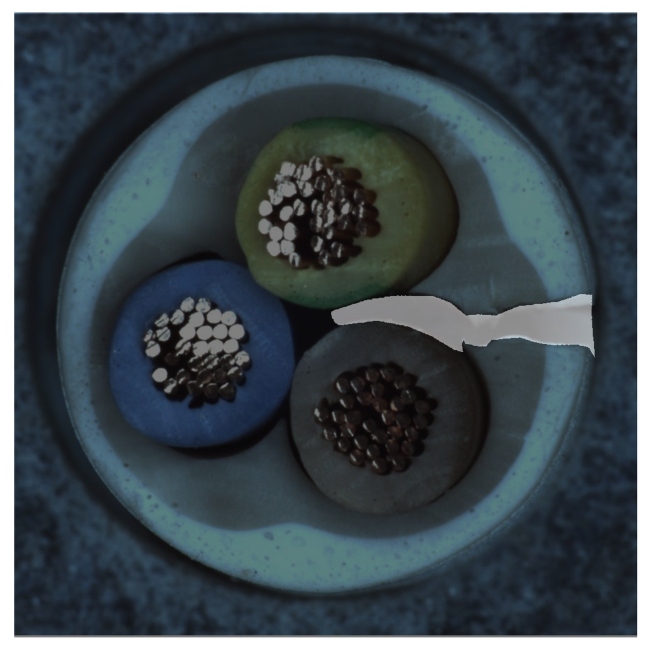

Image of cable with a cut in the outer insulation from the MVTEC dataset with mask showing the location of the anomaly

Anomaly Detection Method

1. Train a generative model to represent the probability density function of a **non-anomalous** image distribution

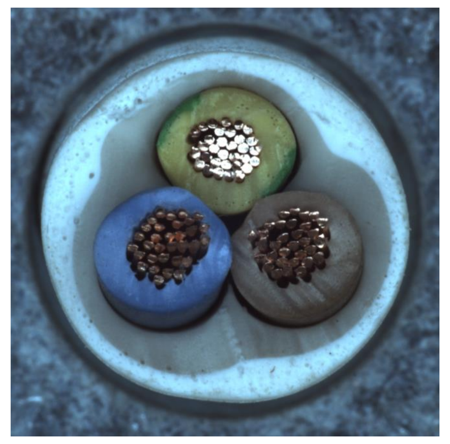

Non-anomalous image of a cable from the MVTEC dataset

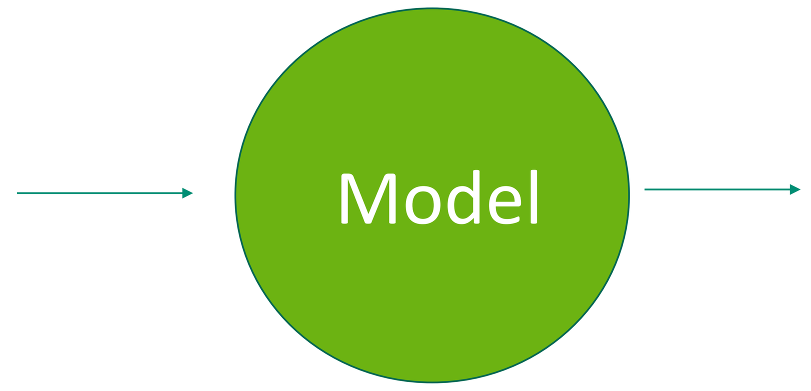

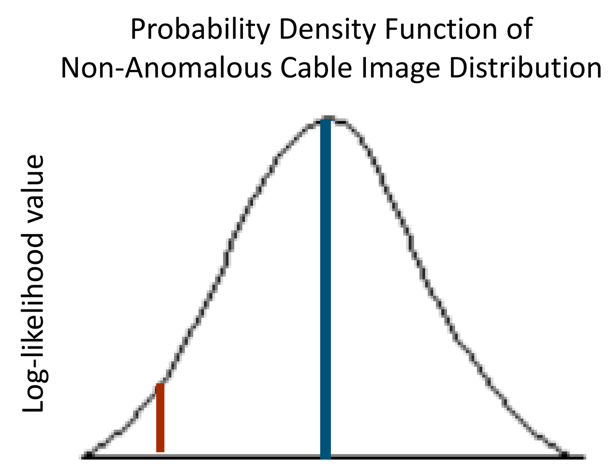

Input image
(red – out-of-distribution input;
blue – in-distribution input)

2. Query the probability density function with non-anomalous and anomalous images and compare their log-likelihood values

Anomalous Images:	Non-Anomalous Images:
• Out-of-distribution • Low log-likelihood value	• In-distribution • High log-likelihood value

Problem

- After training the model on non-anomalous images (Fashion-MNIST), anomalous images (MNIST) are being assigned higher log-likelihood values than non-anomalous images
- Backgrounds between anomalous and non-anomalous images are similar – **dominates** the log-likelihood value

Sample image from MNIST dataset

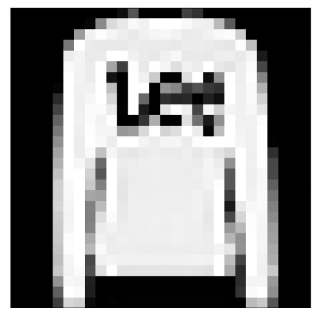

Sample image from Fashion-MNIST dataset

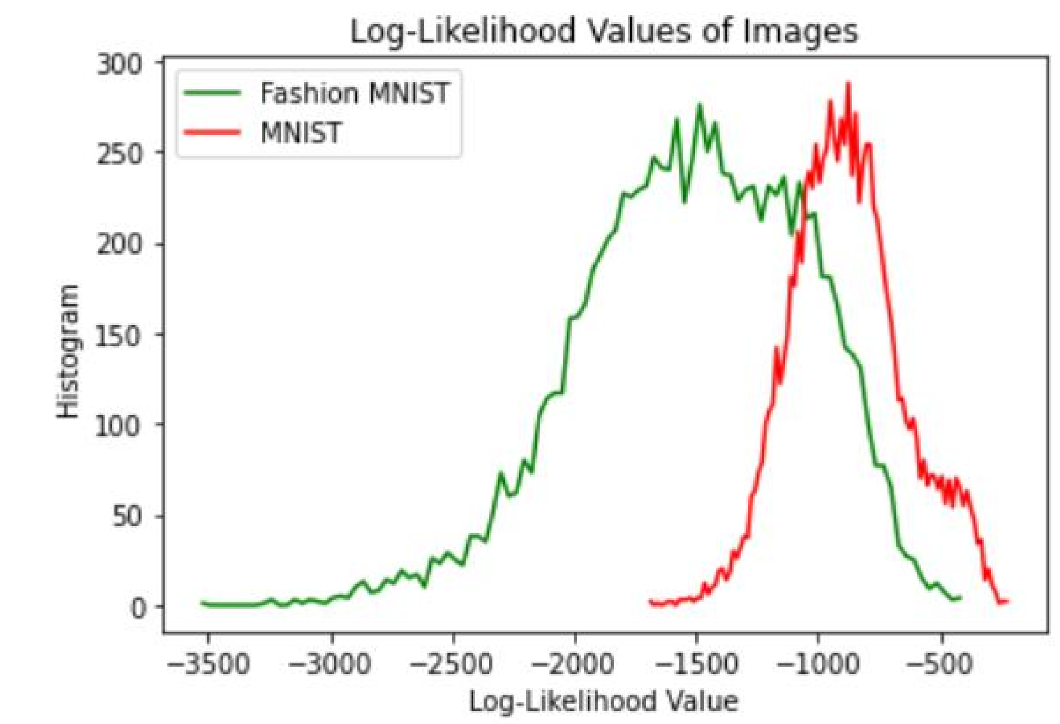

Log-Likelihood Ratio Method

1. Train **two** generative models for the:
   - **non-anomalous, un-mutated** image distribution
   - **non-anomalous, mutated** image distribution
2. Compare log-likelihood values, which is the product of:
   - Background log-likelihood value (p_θ (x_B))
   - Semantic log-likelihood value (p_θ (x_S))
3. Since the background is dominant, the mutation **does not** affect the background log-likelihood value
4. Therefore, the **ratio** between the un-mutated (θ) and mutated (θ_0) log-likelihood values cancels out the background log-likelihood value – only the semantic log-likelihood ratio is considered

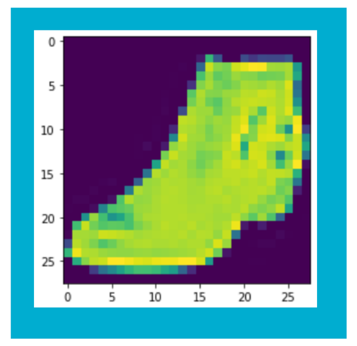

Non-anomalous, un-mutated image

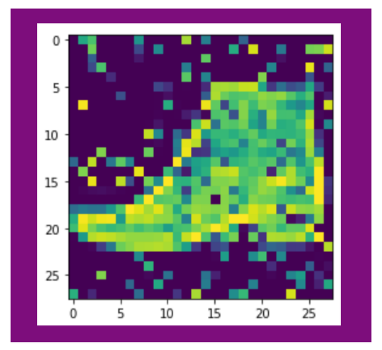

Non-anomalous, mutated image

$$\mathrm{LLR}(x) = \log \frac{\cancel{p_\theta(x_B)}\,\boxed{p_\theta(x_S)}}{\cancel{p_{\theta_0}(x_B)}\,\boxed{p_{\theta_0}(x_S)}},$$

Log-likelihood ratio value calculation – since the background log-likelihood value is cancelled out, only the ratio between the semantic log-likelihood values is considered

Log-likelihood ratio score:
High = in-distribution (Fashion-MNIST)
Low = out-of-distribution (MNIST)

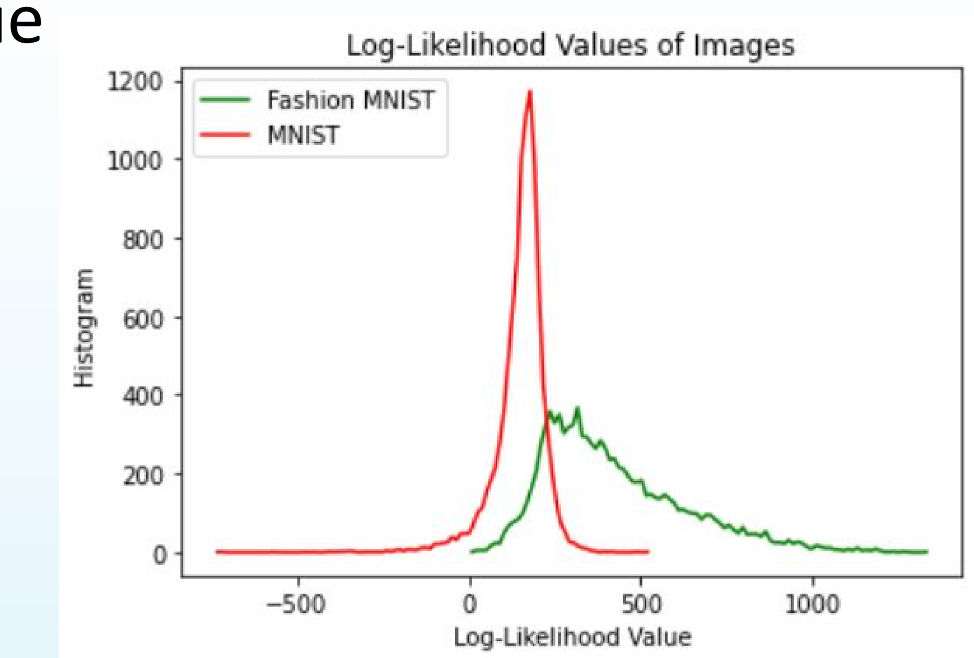

Application

- Given images or CT scans of high-reliability materials, can provide reliable automated anomaly detection

Works Cited

Ren, J. et al. (2019). Likelihood Ratios for Out-of-Distribution Detection. Retrieved July 13, 2020 from https://arxiv.org/pdf/1906.02845.pdf.

https://www.mvtec.com/company/research/datasets/mvtec-ad/

https://machinelearningmastery.com/how-to-develop-a-convolutional-neural-network-from-scratch-for-mnist-handwritten-digit-classification/