

Evaluating Pretrained Visual Models for Industrial Visual Anomaly Detection

Visual foundation models (VFM) have recently become a focus in computer vision research due to their ability to learn general-purpose representations from large-scale, diverse datasets. These models are designed for broad transferability across tasks such as classification, segmentation, and detection, and support zero-shot or few-shot learning. Examples include CLIP, SAM, and DINOv2, which have demonstrated strong performance in open-domain settings. In industrial visual inspection, VFMs offer potential advantages for anomaly detection, particularly in scenarios with limited labeled data and high variability in defect types. However, their suitability for real-time or resource-constrained environments is limited by factors such as inference latency, memory footprint, and cost of deployment. For this reason, models like ResNet or EfficientNet also remain relevant as feature-extractors due to their efficiency and ease of integration.

This thesis will evaluate pretrained visual models for visual anomaly detection using the MVTec AD dataset [Bergmann et al., CVPR 2019], a benchmark for industrial defect detection. Existing libraries for anomaly detection such as anomalib can be used to integrate and compare different approaches.

A specific focus will be the implementation and evaluation of a recent method for **uncertainty quantification** using Gaussian Processes on neural activations [Bergna et al., arXiv:2502.20966]. The approach involves:

- Using a pre-trained model (e.g., ResNet or a lightweight transformer) as a feature extractor.
- Adding a (binary) classification head to distinguish, for example, normal from anomalous samples.
- Modeling the activations of intermediate layers with Gaussian Processes to propagate uncertainty through the network and leveraging the uncertainty measure in the anomaly detection.

The thesis will:

- Implement this uncertainty-aware method on a selected visual model.
- Evaluate its performance on MVTec AD using metrics such as AUROC, uncertainty, and computational resource requirements.
- Compare it against state-of-the-art anomaly detection methods, e.g., based on anomalib.

This work aims to assess the trade-offs between model complexity, accuracy, and uncertainty calibration in industrial contexts, and to explore how uncertainty modeling can enhance decision reliability in visual inspection systems.

Experience and interest in probabilistic / Bayesian methods, and pytorch-programming is recommended.

References

- Yang et al., [A Survey on Foundation-Model-Based Industrial Defect Detection](#), arXiv:2502.19106, 2025
- Li et al., [A Survey of Deep Learning for Industrial Visual Anomaly Detection](#), Artif Intell Rev 58, 2025
- Bergmann et al., [MVTec AD - A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection](#), 2019
- Bergna et al., [Post-Hoc Uncertainty Quantification in Pre-Trained Neural Networks via Activation-Level Gaussian Processes](#), arXiv:2502.20966, 2025

Contact

Prof. Dr. Volker Tresp (volker.tresp@lmu.de)