Modeling the Distribution of Normal Data in Pre-Trained Deep Features for Anomaly Detection

Oliver Rippel*, Patrick Mertens* & Dorit Merhof

*equal contributions

Institute of Imaging and Computer Vision

RWTH AACHEN UNIVERSITY

Introduction

Anomaly Detection (AD) identifies patterns in data that deviate from a (learned) concept of normality.
AD is defined by the following, essential properties:

- Anomalies are rare events.
- There exists limited knowledge about the anomaly distribution, i.e. it is not well-defined.
- Nominal/Normal data is easy to sample.

Main Challenge in AD:

- How to learn discriminative features with limited access to anomalies?

Hypothesis

1. Learning features for AD from scratch is difficult $\Rightarrow$ Transfer Learning features generated from ImageNet training.
2. Inspired by linkage between Generative and Discriminative models shown in [1]: Nominal/normal class may also follow a Multivariate Gaussian (MVG) in transfer learning setting.

[1] Kimin Lee et al. "A simple unified framework for detecting out-of-distribution samples and adversarial attacks". In: Advances in Neural Information Processing Systems. 2018, pp. 7167–7177

Multivariate Gaussian Modelling

MVG is given by:

$$\varphi_{\boldsymbol{\mu},\Sigma}(\mathbf{x}) := \frac{1}{\sqrt{(2\pi)^D |\det \Sigma|}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x}-\boldsymbol{\mu})}. \tag{1}$$

Under a Gaussian distribution a sensible distance measure between a particular point $\mathbf{x} \in \mathbb{R}^D$ and the distribution is called the Mahalanobis distance and defined as

$$M(\mathbf{x}) = \sqrt{(\mathbf{x} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})}. \tag{2}$$

Institute of Imaging and Computer Vision
RWTH AACHEN UNIVERSITY

Feature Extraction

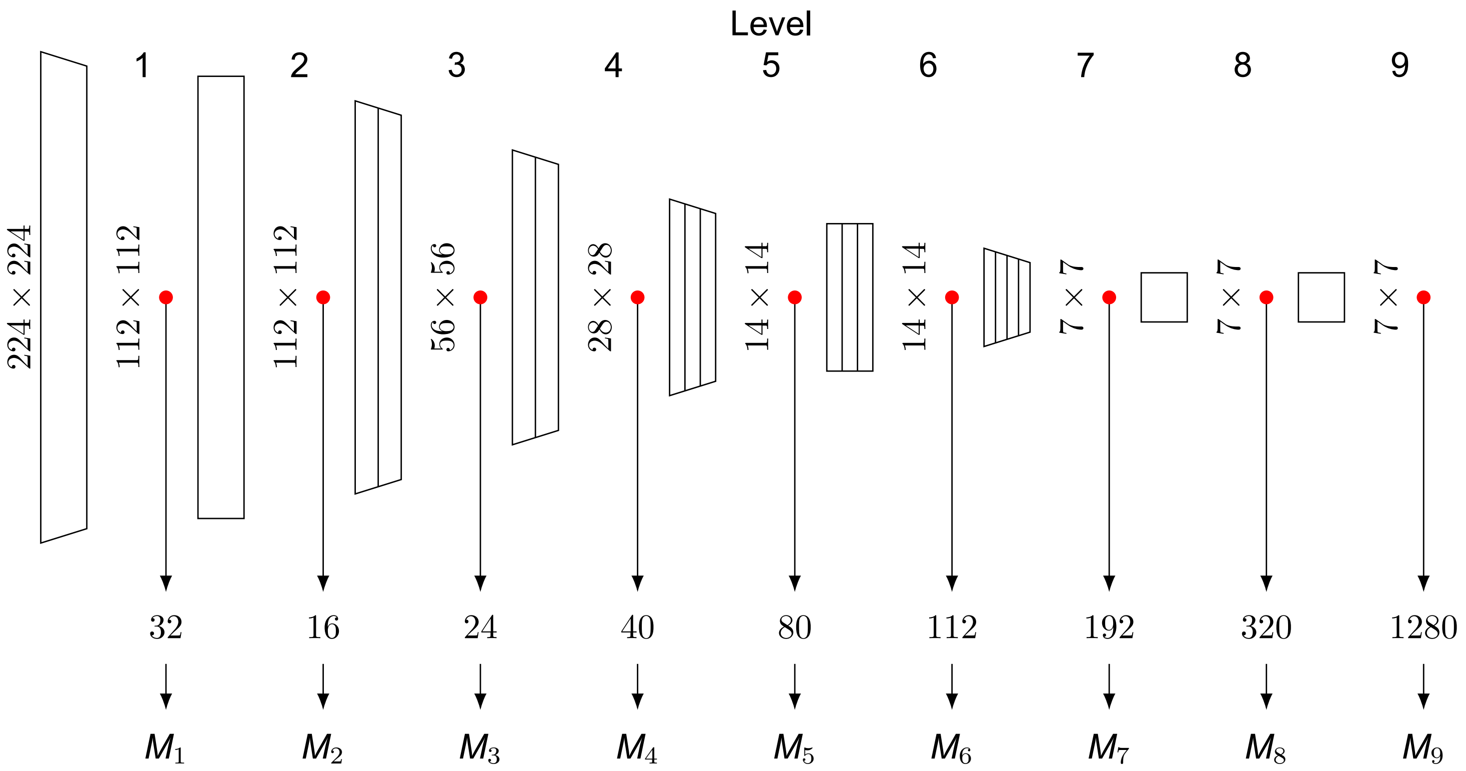

$$\sum_{i=1}^{9} M_i$$

Dataset

MVTec Anomaly Detection Dataset

- First open dataset for industrial AD
- 5354 RGB images of 15 objects and textures
- Small semantic differences (as opposed to “cats” vs. “dogs”)

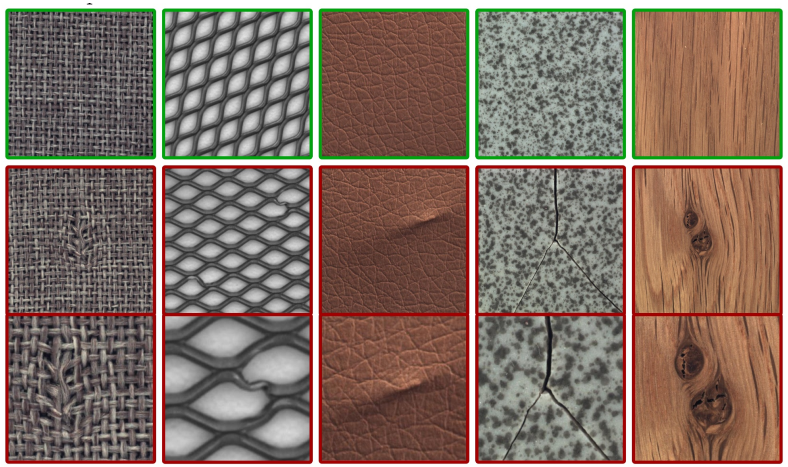

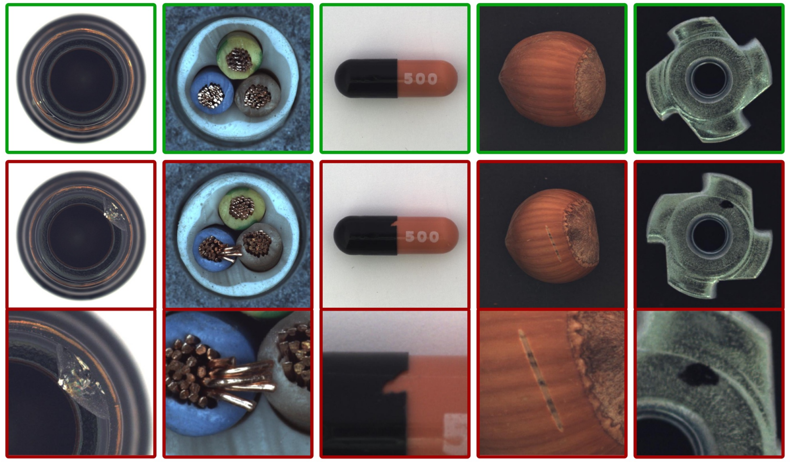

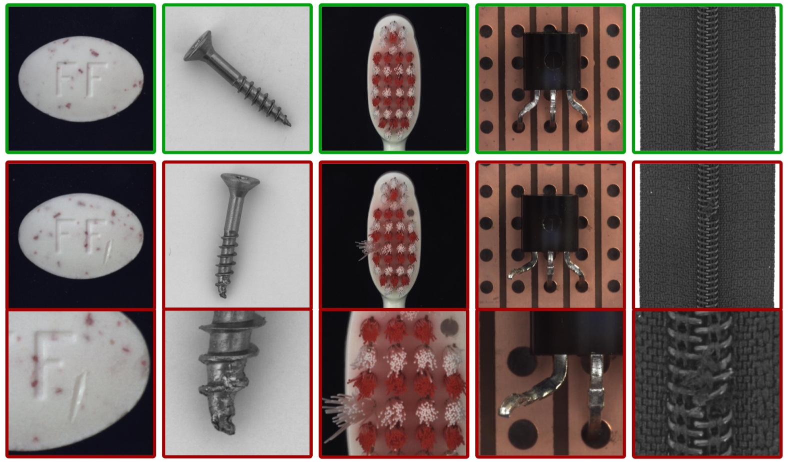

Figures from Paul Bergmann et al. “MVTec AD – A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection”. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). June 2019

Results

Transfer Learning MVG

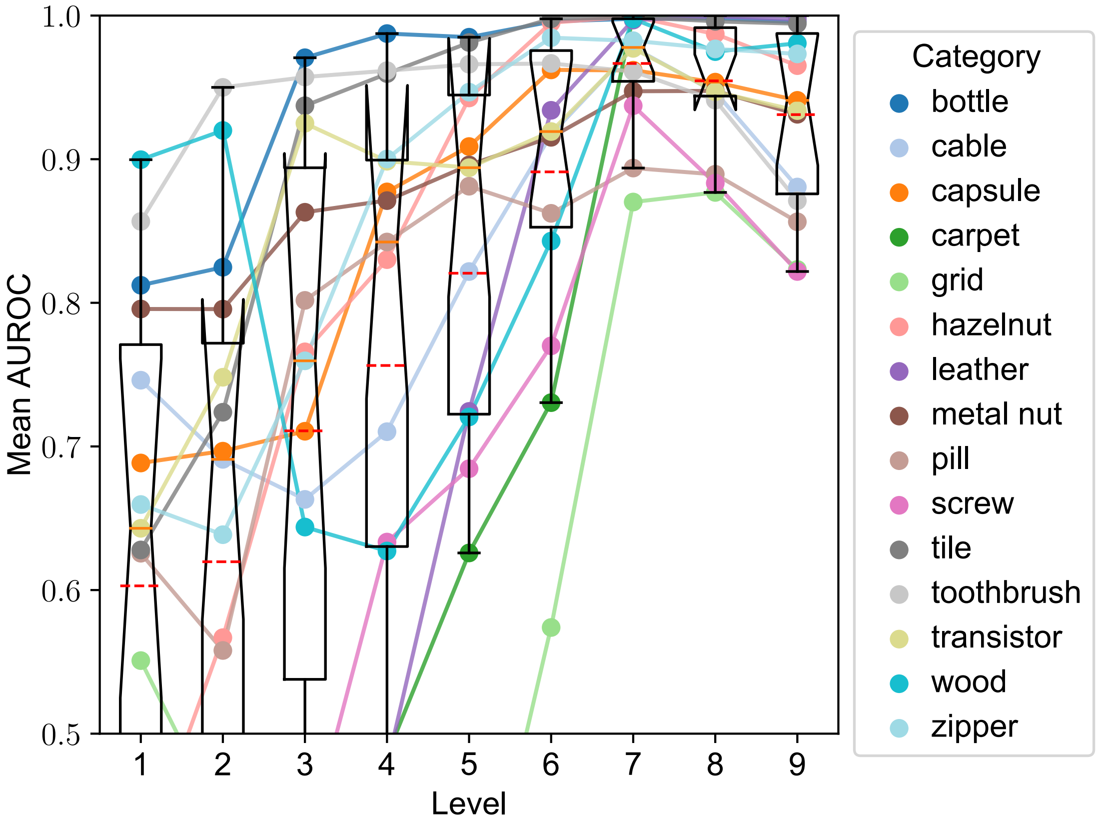

Figure: EfficientNet-B4 with Mahalanobis distance

Table: EfficientNet-B4

Level	L_2		SED		Mahalanobis	
	Mean	SEM	Mean	SEM	Mean	SEM
1	44.5	4.8	51.6	5.7	60.3	6.1
2	47.3	5.3	48.1	5.1	62.0	6.4
3	58.1	5.9	59.2	6.3	71.1	5.4
4	59.7	4.7	61.5	5.1	75.6	5.5
5	62.6	4.8	66.1	5.0	82.1	4.6
6	71.7	4.4	74.3	4.3	89.1	3.1
7	82.9	4.3	85.1	4.0	96.7	1.0
8	83.2	3.7	85.2	3.4	95.5	1.1
9	83.3	3.7	87.8	3.0	93.1	1.7
Sum	72.5	5.2	78.9	6.6	94.8	1.6

Why does Transfer Learning + MVG work so well?

Table: EfficientNet-B4

Level	Full		PCA 99%		PCA 95%		NPCA 1%		NPCA 0.1%		NPCA 0.01%	
	Mean	SEM	Mean	SEM	Mean	SEM	Mean	SEM	Mean	SEM	Mean	SEM
1	60.3	6.1	50.8	6.1	45.8	5.4	70.2	4.0	67.8	6.6	69.9	6.2
2	62.0	6.4	53.3	5.9	48.6	5.7	67.6	6.3	68.0	5.9	67.3	5.7
3	71.1	5.4	65.5	6.5	59.7	6.2	71.6	4.9	68.1	4.2	65.7	3.7
4	75.6	5.5	69.5	6.1	63.2	6.4	76.1	5.1	73.1	4.6	69.4	4.0
5	82.1	4.6	76.2	5.3	66.6	6.5	82.5	4.0	78.7	3.6	72.3	3.7
6	89.1	3.1	83.3	4.8	77.3	5.7	90.2	2.5	88.2	2.4	83.6	2.9
7	96.7	1.0	93.4	2.1	87.1	4.0	96.1	1.0	94.5	1.3	89.6	2.5
8	95.5	1.1	91.9	2.1	88.6	3.1	94.8	1.2	93.8	1.4	90.6	2.3
9	93.1	1.7	91.3	2.1	88.6	2.9	93.3	1.6	91.2	2.1	86.3	3.0
Sum	94.8	1.6	89.6	3.4	82.2	6.0	95.8	1.1	95.5	1.2	94.0	1.6

Results

Comparison with SotA

Approach	Architecture	Mean	SEM
GeoTrans[1] (source:[2])	Wide-ResNet	67.2	4.7
GANomaly[3] (source:[2])	DCGAN	76.1	1.6
ITAE[2]	Custom	83.9	2.8
SPADE[4]	Wide-ResNet50-2	85.5	—
MSE AE			
CCA	ResNet-18	81.8	3.4
Pre-Trained Classifier			
Fully-Supervised Fine-tune	EfficientNet-B4	96.3	1.0
Mahalanobis (ours)			
All Features	EfficientNet-B4	95.2	1.5
NPCA 1%	EfficientNet-B4	**95.8**	1.2

[1] Izhak Golan et al. “Deep anomaly detection using geometric transformations”. In: Advances in Neural Information Processing Systems. 2018, pp. 9758–9769

[2] Chaoqing Huang et al. “Inverse-Transform AutoEncoder for Anomaly Detection”. In: arXiv preprint arXiv:1911.10676 (2019)

[3] Samet Akçay et al. “GANomaly: Semi-supervised anomaly detection via adversarial training”. In: Asian Conference on Computer Vision. Springer. 2018, pp. 622–637

[4] Niv Cohen et al. “Sub-Image Anomaly Detection with Deep Pyramid Correspondences”. In: arXiv preprint arXiv:2005.02357 (2020)

Conclusion

1. Learning features from scratch is indeed difficult, as discriminative features account for little variance in normal (defect-free) images.
2. MVG based modeling of PDF in a Transfer Learning setting achieves SotA with 95.8% AUROC.
3. Code is available at: `https://github.com/ORippler/gaussian-ad-mvtec`

Outlook:

- Enforce Gaussian distribution ⇒ self-normalizing neural networks[5]
- Multi-modal distributions e.g. Gaussian Mixture Model (GMM) or Normalizing Flows
- Apply PDF Modeling to Anomaly Segmentation[6]

[5] Günter Klambauer et al. “Self-Normalizing Neural Networks”. In: Advances in Neural Information Processing Systems 30. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 971–980. URL: http://papers.nips.cc/paper/6698-self-normalizing-neural-networks.pdf

[6] Congrats to Defard for achieving this and publishing the results at the Industrial Machine Learning Workshop at ICPR2020. Thomas Defard et al. “PaDiM: a Patch Distribution Modeling Framework for Anomaly Detection and Localization”. In: arXiv preprint arXiv:2011.08785 (2020)

Thank you for your attention!