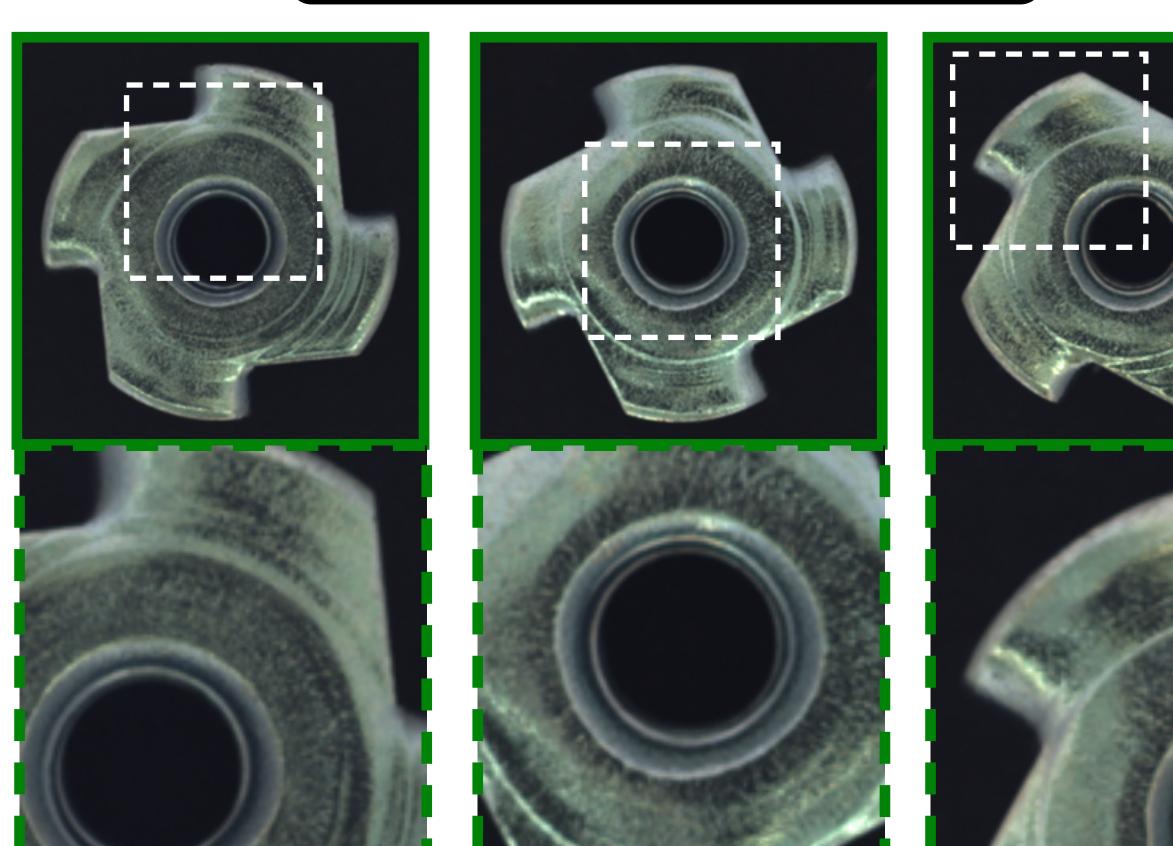
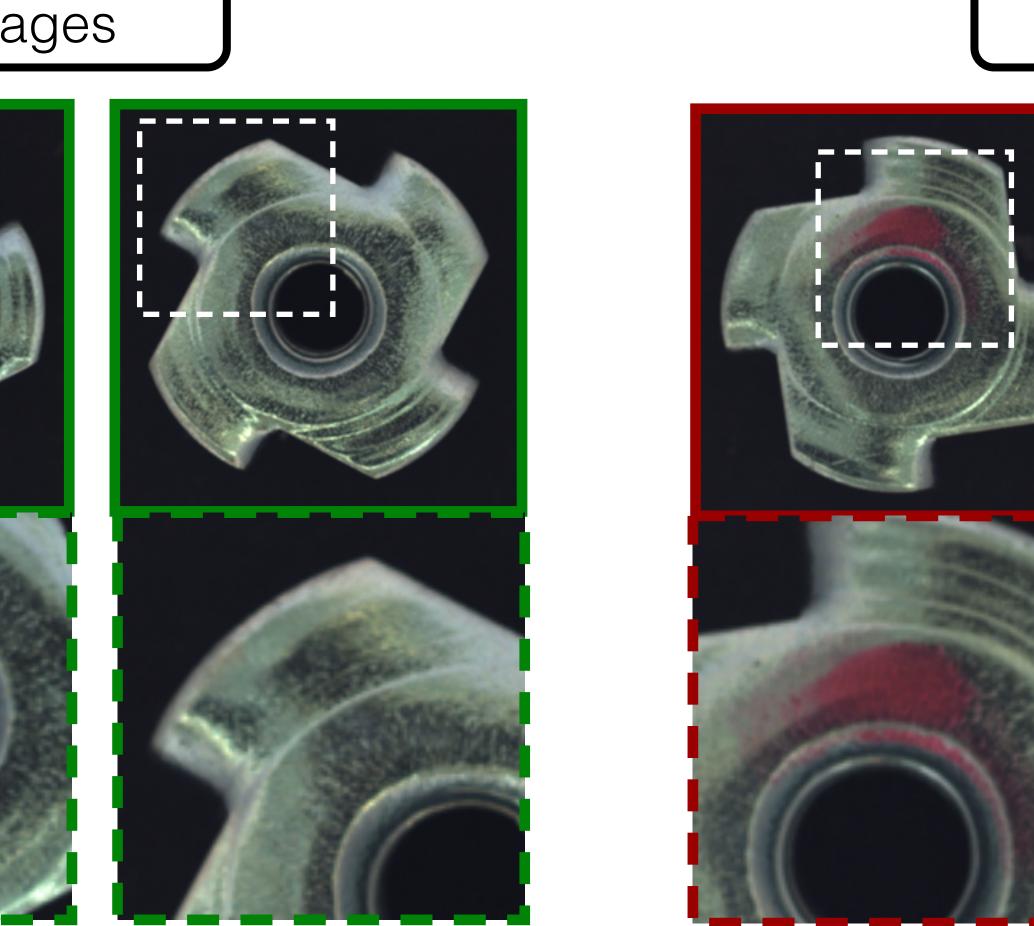


Motivation

- The detection of regions that deviate from the training data is an important, yet challenging task.
- Little research is directed towards the unsupervised segmentation of anomalies in natural images.

Anomaly-free Images





Anomalous Images

Used for training and testing.

Used exclusively for testing.

- Existing methods focus on generative models such as autoencoders or generative adversarial networks.
- Pretrained feature extractors provide powerful embeddings that can be used for anomaly detection.

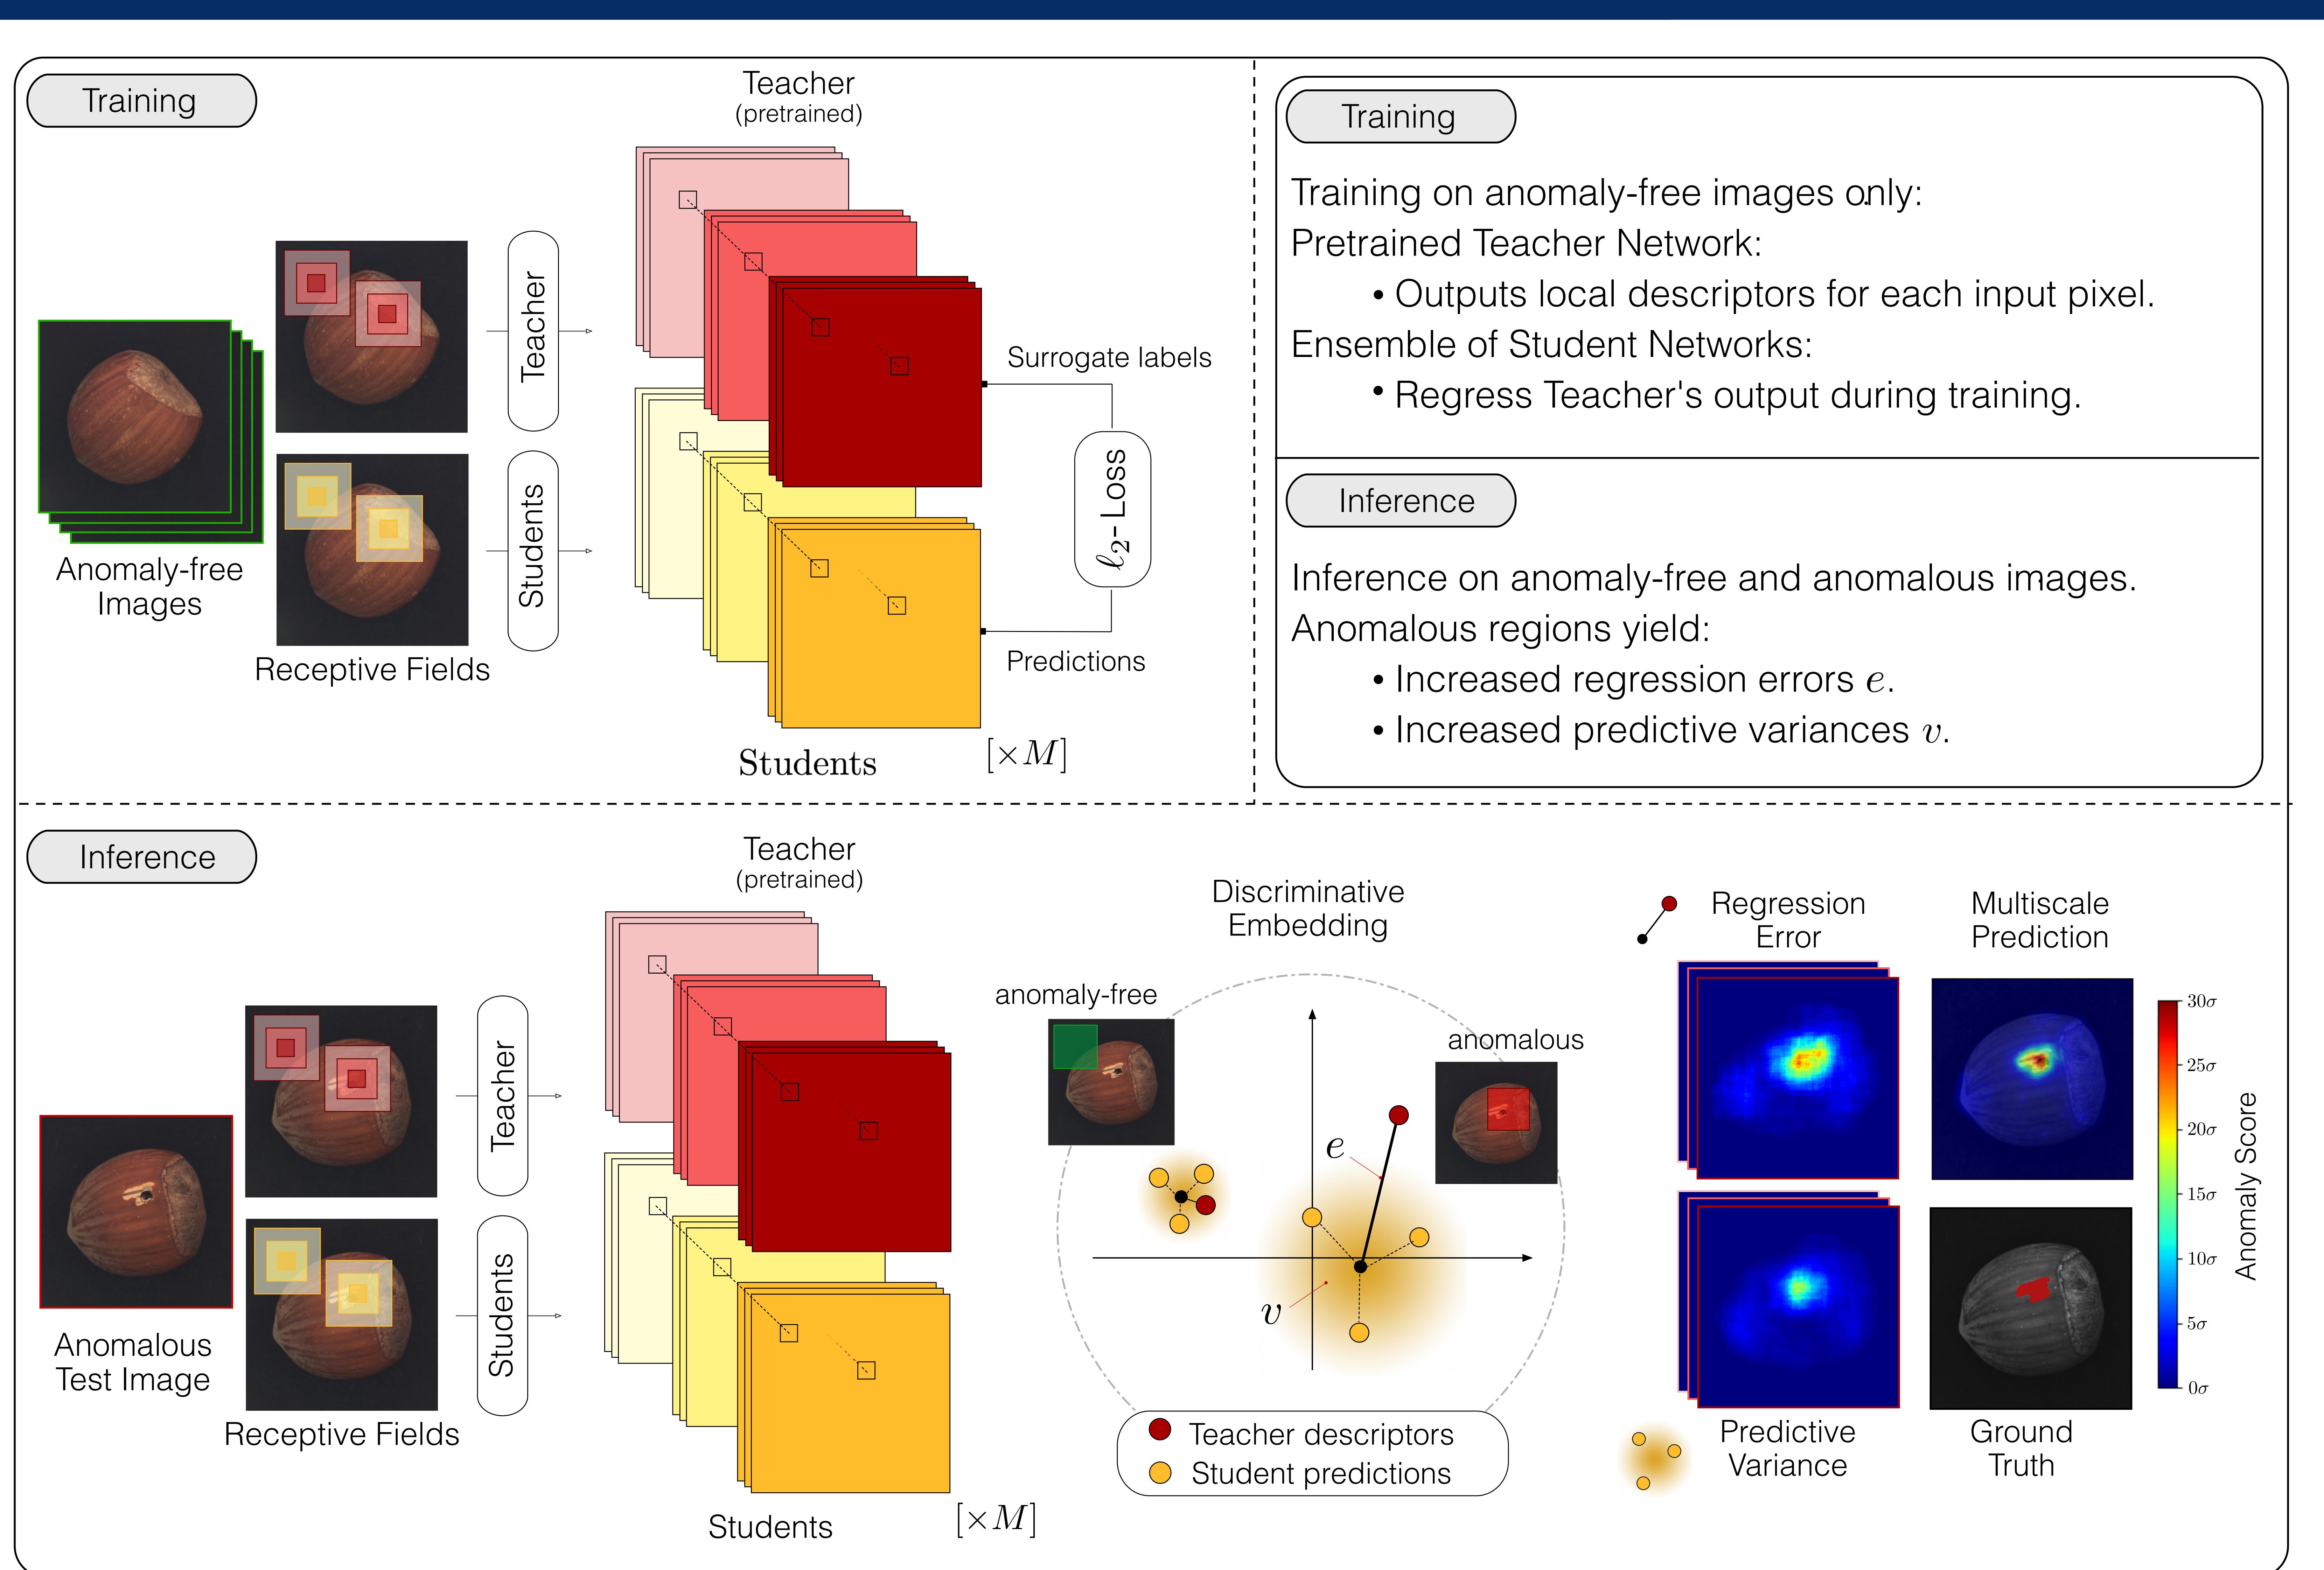
Key Contributions

- A novel framework for unsupervised anomaly detection based on student-teacher learning.
- A pretrained teacher network is constructed to output discriminative embedding vectors for local image regions.
- Anomaly scoring is based on the students' ability to mimic the teacher, and on their predictive variance.
- Explicit control over the networks' receptive fields allows for anomaly detection at multiple scales.

Teacher Pretraining

Uninformed Students: Student-Teacher Anomaly Detection with Discriminative Latent Embeddings Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger MVTec Software GmbH

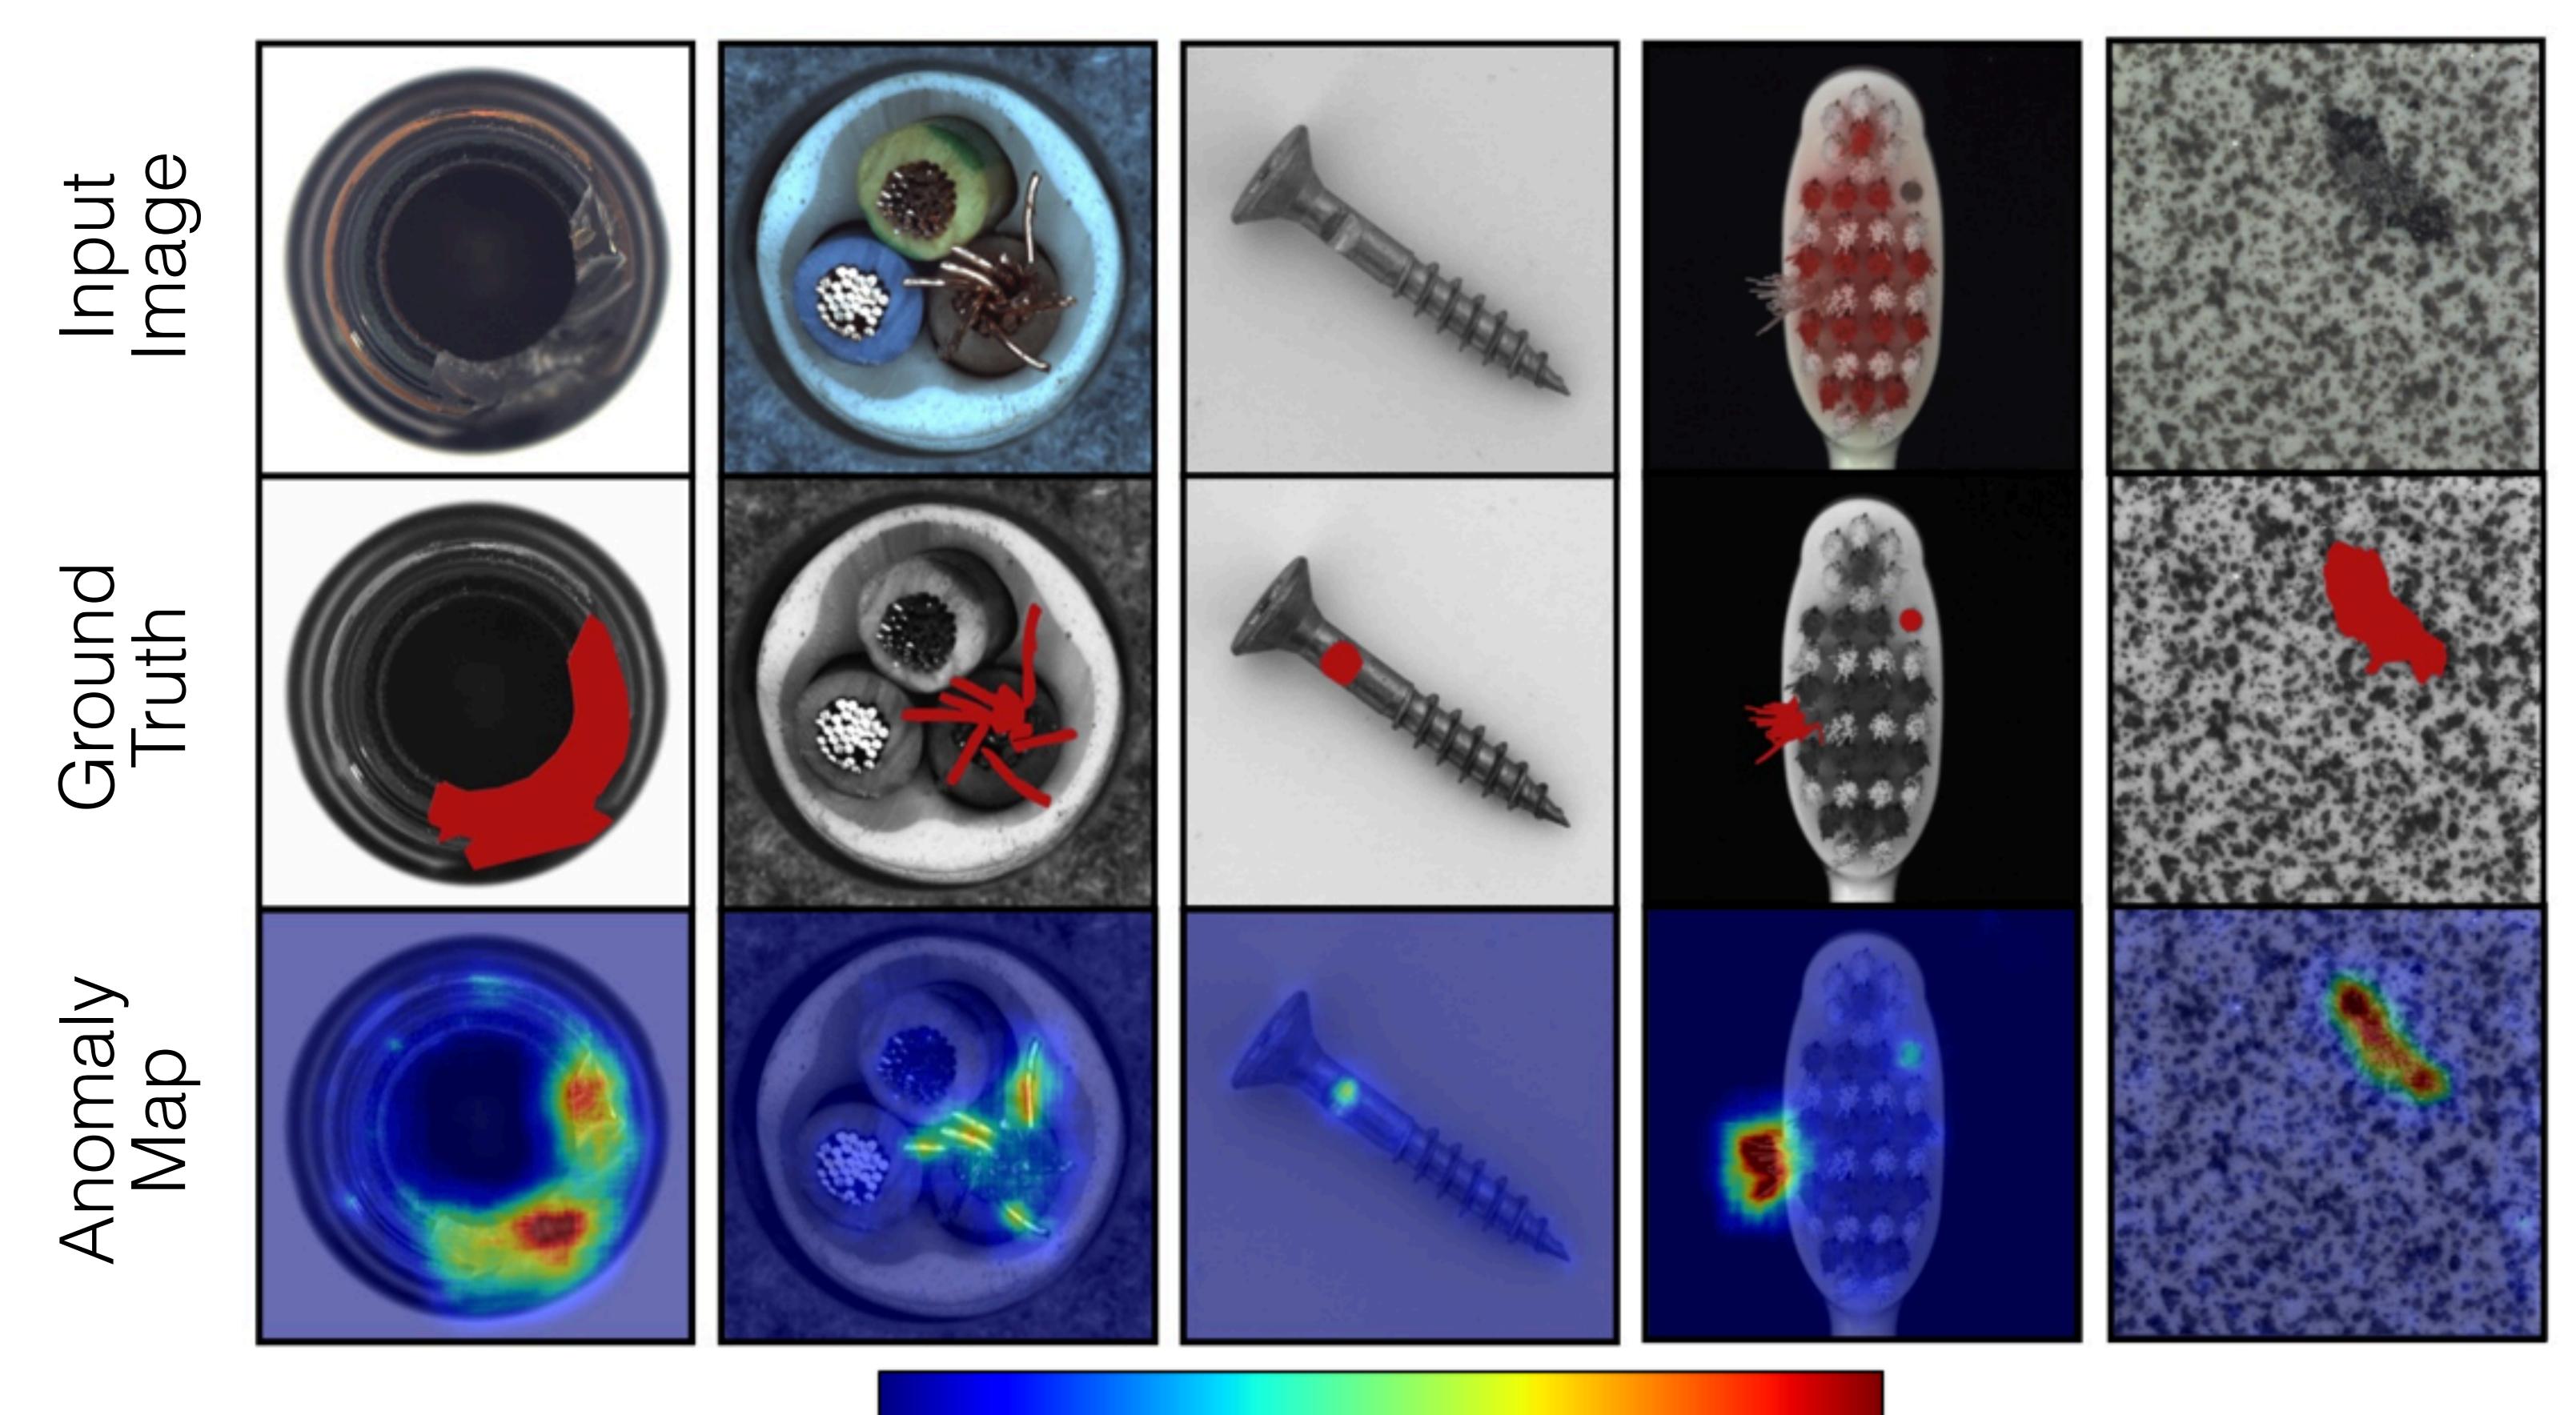
Method Overview



Multiscale Anomaly Detection

CVPR SEATTLE WASHINGTON JUNE 16-18 2020

Qualitative Results



 0σ 5σ 10σ 15σ 20σ 25σ 30σ Anomaly Score

Quantitative Results

| | Textures | Objects |
|-------------------------------|----------|---------|
| AnoGAN | 0.274 | 0.533 |
| SSIM-AE | 0.567 | 0.758 |
| 1-NN | 0.502 | 0.710 |
| OC-SVM | 0.369 | 0.535 |
| K-Means | 0.372 | 0.450 |
| ℓ_2 -AE | 0.696 | 0.838 |
| VAE | 0.572 | 0.673 |
| CNN Feature Dictionary | 0.542 | 0.503 |
| Ours (multiscale) | 0.927 | 0.908 |

 Anomaly segmentation results on the MVTec AD dataset. (CVPR 2019 - https://www.mvtec.com/company/research/datasets)

- The normalized area under the Per-Region Overlap (PRO) curve is computed up to a FPR of 30%.
- Our approach outperforms the state-of-the-art.