
Contrastive Predictive Coding for Anomaly Detection and Segmentation

Puck de Haan¹ Sindy Löwe¹

Abstract

Reliable detection of anomalies is crucial when deploying machine learning models in practice, but remains challenging due to the lack of labeled data. To tackle this challenge, contrastive learning approaches are becoming increasingly popular, given the impressive results they have achieved in self-supervised representation learning settings. However, while most existing contrastive anomaly detection and segmentation approaches have been applied to images, none of them can use the contrastive losses directly for both anomaly detection and segmentation. In this paper, we close this gap by making use of the Contrastive Predictive Coding model (Oord et al., 2019). We show that its patch-wise contrastive loss can directly be interpreted as an anomaly score, and how this allows for the creation of anomaly segmentation masks. The resulting model achieves promising results for both anomaly detection and segmentation on the challenging MVTec-AD dataset.

1. Introduction

An anomaly (or outlier, novelty, out-of-distribution sample) is an observation that differs significantly from the vast majority of the data. Anomaly detection (AD) tries to distinguish anomalous samples from the samples that are deemed ‘normal’ in the data. It has become increasingly relevant to detect these anomalies to make machine learning methods more reliable and to improve their applicability in real-world scenarios, such as automated industrial inspections and medical diagnosis (Ruff et al., 2021). Typically, anomaly detection is treated as an unsupervised learning problem, since labelled data is generally unavailable and to allow for the development of methods that can detect previously unseen anomalies.

One promising direction involves the adaptation of con-

¹UvA-Bosch Delta Lab, University of Amsterdam, Netherlands. Correspondence to: Puck de Haan <pdehaan274@gmail.com>, Sindy Löwe <loewe.sindy@gmail.com>.

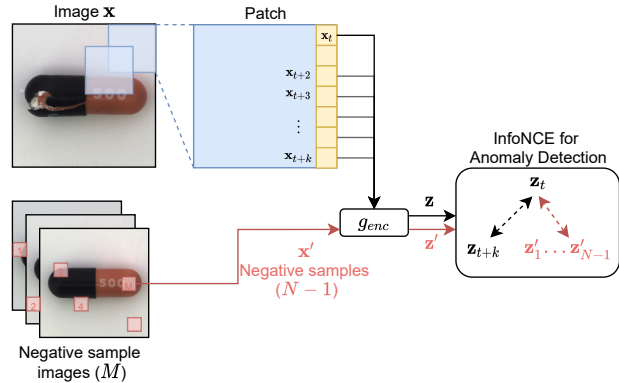


Figure 1. A schematic overview of Contrastive Predictive Coding for anomaly detection and segmentation in images. After extracting (sub-)patches from the input image, we contrast the encoded representations from within the same image (z_t, z_{t+k}) against $N - 1$ randomly matched representations (z_t, z_j). The resulting InfoNCE loss is used to determine whether sub-patch x_{t+k} is anomalous or not.

trastive learning approaches (Hjelm et al., 2019; Oord et al., 2019; Chen et al., 2020; He et al., 2020) to the anomaly detection setting (Tack et al., 2020; Winkens et al., 2020; Kopuklu et al., 2021; Qiu et al., 2021; Sohn et al., 2021). However, even though most of these approaches have been applied to image data, none of them can use the contrastive losses directly for both anomaly detection and segmentation.

In this paper, we demonstrate that Contrastive Predictive Coding (CPC) (Oord et al., 2019; Hénaff et al., 2020) can be applied to detect and segment anomalies in images. We show that the InfoNCE loss introduced by Oord et al. (2019) can be directly interpreted as an anomaly score. Since in this loss patches from within an image are contrasted against one another, we can further use it to create accurate anomaly segmentation masks. This results in a compact and straightforward anomaly detection and segmentation approach.

To improve the performance of the CPC model for anomaly detection, we introduce two adjustments. First, we adapt the setup of negative samples during testing such that anomalous patches can only appear within the positive sample. Second, we omit the autoregressive part of the CPC model. With these adjustments, our proposed method achieves promising performance on real-world data, such as the challenging MVTec-AD dataset (Bergmann et al., 2019).

2. Related Work

In this section, we will give an overview of contrastive learning approaches and different methods for anomaly detection.

2.1. Contrastive Learning

Lately, impressive results have been achieved with self-supervised methods based on contrastive learning (Wu et al., 2018; Oord et al., 2019; Hjelm et al., 2019; He et al., 2020; Chen et al., 2020; Li et al., 2021b). Overall, these methods work by making a model decide whether two (randomly) transformed inputs originated from the same input sample, or from two samples that have been randomly drawn from across the dataset. Different transformations can be chosen depending on the domain and downstream task. For example, on image data, random data augmentation such as random cropping and color jittering has proven useful (Chen et al., 2020; He et al., 2020). In this paper, we use the Contrastive Predictive Coding model (Oord et al., 2019; Hénaff et al., 2020), which makes use of temporal transformations. Generally, these approaches are evaluated by training a linear classifier on top of the created representations and by measuring the performance that this linear classifier can achieve on downstream tasks.

2.2. Anomaly Detection

Anomaly detection methods can roughly be divided into three categories: density-based, reconstruction-based and discriminative-based methods (Ruff et al., 2021). Density-based methods predict anomalies by estimating the probability distribution of the data (e.g. GANs, VAEs, or flow-based models) (Schlegl et al., 2017; Winkens et al., 2020; Liu et al., 2020); reconstruction-based methods are based on models that are trained with a reconstruction objective (e.g. autoencoders) (Zhou & Paffenroth, 2017; Bergmann et al., 2018; Luo et al., 2020); discriminative-based methods learn a decision boundary between anomalous and normal data (e.g. SVM, one-class classification) (Ruff et al., 2020; Tack et al., 2020; Liznerski et al., 2021; Li et al., 2021a). The method proposed in this paper can be seen as a density-based method with a discriminative one-class objective.

Several previous works investigate the use of contrastive learning for AD. Tack et al. (2020); Winkens et al. (2020); Sohn et al. (2021) make use of the SimCLR framework (Chen et al., 2020) to learn representations of the data. Then, they calculate a separate anomaly score by using these representations for density estimation, one-class classification, or by applying metric measures like the cosine similarity and the norm of the representations. The downsides of this approach are that it requires extensive data augmentations and multiple different measures, or multiple models. Another comparable contrastive learning AD method (Kopuklu et al.,

2021) uses noise contrastive estimation for training, similar to our method. Differently to our method, they map the samples to multiple latent spaces and use anomalous samples as negatives during training. This results in a more complex model with a supervised training phase. NeuTraL AD (Qiu et al., 2021) makes use of a contrastive loss with learnable transformations, and reuses this loss as an anomaly score. In contrast to our method, their approach has been evaluated on time-series and tabular data.

3. Contrastive Predictive Coding

Contrastive Predictive Coding (Oord et al., 2019) is a self-supervised representation learning approach that leverages the structure of the data and enforces temporally nearby inputs to be encoded similarly in latent space. It achieves this by making the model decide whether a pair of samples is made up of temporally nearby samples or randomly assigned samples. This approach can also be applied to static image data by splitting the images up into patches, and interpreting each row of patches as a separate time-step.

The CPC model makes use of a contrastive loss function, coined InfoNCE, that is based on Noise-Contrastive Estimation (Gutmann & Hyvärinen, 2010) and is designed to optimize the mutual information between the latent representations of patches (z_t) and their surrounding patches (c_{t+k}):

$$\mathcal{L}_k = -\mathbb{E}_X \left[\log \frac{\exp(z_{t+k}^\top \mathbf{W}_k c_t)}{\sum_X \exp(z_j^\top \mathbf{W}_k c_t)} \right], \quad (1)$$

where $z_t = g_{\text{enc}}(\mathbf{x}_t)$ and g_{enc} represents a non-linear encoder, $c_t = g_{\text{ar}}(\mathbf{z}_{\leq t})$ and g_{ar} represents an autoregressive model. Furthermore, \mathbf{W}_k describes a linear transformation used for predicting k time-steps ahead. The set of samples X consists of one positive sample ($\mathbf{x}_t, \mathbf{x}_{t+k}$) and $N-1$ negative samples ($\mathbf{x}_t, \mathbf{x}_j$) for which \mathbf{x}_j is randomly sampled from across the current batch.

4. CPC for Anomaly Detection

We propose to apply the CPC model for anomaly detection and segmentation (Fig. 1). In order to improve the performance of the CPC model in this setting, we introduce two adjustments to its architecture: (1) We omit the autoregressive model g_{ar} . As a result, our loss function changes to:

$$\mathcal{L}_k = -\mathbb{E}_X \left[\log \frac{\exp(z_{t+k}^\top \mathbf{W}_k z_t)}{\sum_X \exp(z_j^\top \mathbf{W}_k z_t)} \right]. \quad (2)$$

This formulation is equivalent to the loss used in the Greedy InfoMax model (Löwe et al., 2019). This adjustment results in a simpler model, which is still able to learn useful latent representations – according to preliminary results. (2) We

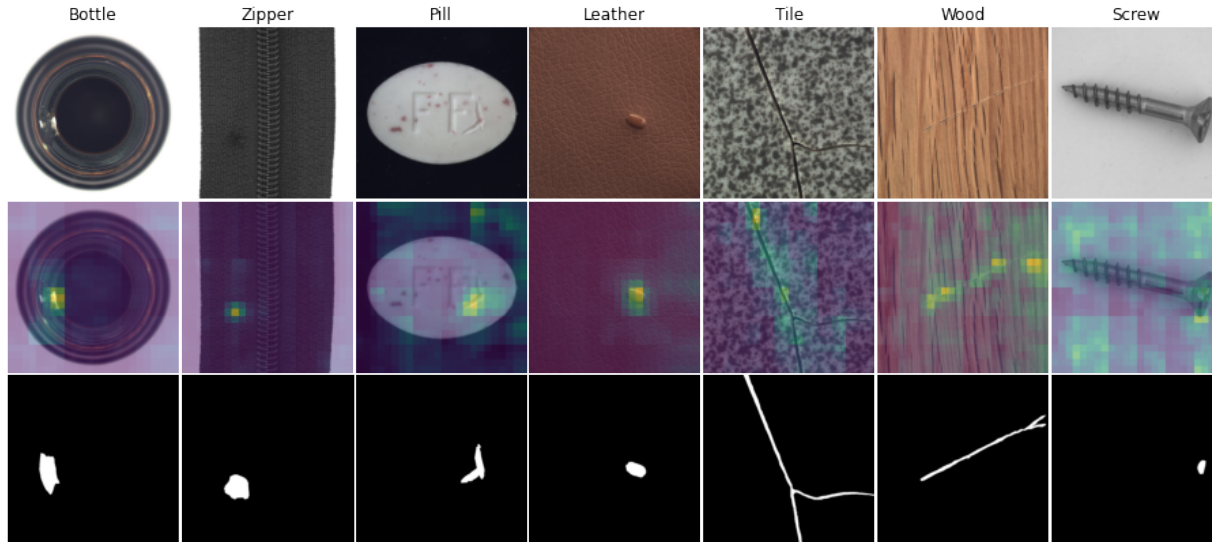


Figure 2. Localization of anomalous regions for different classes in the MVTEC-AD dataset. The top row shows the original input image, the mid row depicts the superimposition of the image and corresponding InfoNCE loss values (brighter colors represent higher loss values) and the bottom row shows the ground truth annotation. We find that our model consistently highlights anomalous regions across many classes. One notable exception is the screw class (right), for which the model assigns high loss values to the background in many cases.

change the setup of the negative samples during testing. Previous implementations of the CPC model use random patches from within the same test-batch (Hénaff et al., 2020) as negative samples. However, this may result in negative samples containing anomalous patches, which could make it harder for the model to detect anomalous patches in the positive sample based on the contrastive loss. To avoid this, during testing, we draw negative samples from the (non-anomalous) training data.

In the test-phase, we use the loss function in Eq. (2) to decide whether an image patch x_{t+k} can be classified as anomalous:

$$x_{t+k} = \begin{cases} \text{anomalous,} & \text{if } \mathcal{L}_k \geq \tau \\ \text{normal,} & \text{if } \mathcal{L}_k < \tau \end{cases} \quad (3)$$

The threshold value τ remains implicit, since we use the area under the receiver operating characteristic curve (AUROC) as performance measure. While we can create anomaly segmentation masks by making use of the anomaly scores per patch x_{t+k} , we can also apply our approach to decide whether a sample is anomalous – either by averaging over the scores of all patches within an image, or by examining the patch with the highest score.

5. Experiments

We evaluate the proposed Contrastive Predictive Coding model for anomaly detection and segmentation on the MVTEC-AD dataset (Bergmann et al., 2019). This dataset

contains high-resolution images of ten objects and five textures with pixel-accurate annotations and provides between 60 and 391 training images per class. During training we randomly crop every image to $[0.8 - 1]$ times the original dimensions. Then, both train and test images are resized to 768×768 pixels. The resulting image is split into patches of size 256×256 , where each patch has 50% overlap with its neighbouring patches. These patches are further divided into sub-patches of size 64×64 , also with 50% overlap. These sub-patches are used in the InfoNCE loss (Fig. 1) to detect anomalies. The cropped and resized images are horizontally flipped with a probability of 50% during training.

We use a ResNet-18 v2 (He et al., 2016) up until the third residual block as encoder g_{enc} . We train a separate model from scratch for each class with a batch size of 16 for 150 epochs using the Adam optimizer (Kingma & Ba, 2015) with a learning rate of $1.5e-4$. As proposed by Oord et al. (2019), we train and evaluate the model on grayscale images. For both training and evaluation we use 16 negative samples in the InfoNCE loss. To increase the accuracy of the InfoNCE loss as an indicator for anomalous patches, we apply four separate models in four different directions – predicting patches using context from above, below, left and right – and combine their losses in the test-phase.

5.1. Anomaly Detection

To evaluate our model’s performance for detecting anomalies, we average the top-5% InfoNCE loss values across all sub-patches within an image and use this value to calculate

Table 1. Anomaly detection AUROC score on the MVTec-AD test-set per category. We find that the proposed **CPC-AD** approach substantially outperforms the kernel density estimation model (**KDE**) and the autoencoding model (**Auto**) presented by Kauffmann et al. (2020). It is outperformed by the **CutPaste** model (Li et al., 2021a), which relies heavily on dataset-specific augmentations for its training.

	KDE	Auto	CutPaste	CPC-AD
Bottle	0.833	0.950	0.983	0.998
Cable	0.669	0.573	0.806	0.880
Capsule	0.562	0.525	0.962	0.641
Carpet	0.348	0.368	0.931	0.809
Grid	0.717	0.746	0.999	0.983
Hazelnut	0.699	0.905	0.973	0.996
Leather	0.415	0.640	1.000	0.990
Metal nut	0.333	0.455	0.993	0.845
Pill	0.691	0.760	0.924	0.921
Screw	0.369	0.779	0.863	0.897
Tile	0.689	0.518	0.934	0.957
Toothbrush	0.933	0.494	0.983	0.878
Transistor	0.724	0.512	0.955	0.925
Wood	0.947	0.885	0.986	0.803
Zipper	0.614	0.350	0.994	0.993
Mean	0.636	0.631	0.952	0.901

the AUROC score. In Table 1, we compare against previously published works from peer-reviewed venues that do not make use of pre-trained feature extractors. We find that our proposed **CPC-AD** model substantially improves upon a kernel density estimation model (**KDE**) and an autoencoding model (**Auto**) as presented in Kauffmann et al. (2020). We also improve upon the contrastive learning approach combined with a KDE model (avg. AUROC: 0.865) as proposed by Sohn et al. (2021). The performance of our model lags behind the **CutPaste** model (Li et al., 2021a). However, we argue that CPC-AD provides a more generally applicable approach for anomaly detection. The CutPaste model relies heavily on randomly sampled artificial anomalies that are designed to resemble the anomalies encountered in the dataset. As a result, it is not applicable to a k -classes-out task, where anomalies differ semantically from the normal data. For comparison, the current state-of-the-art model on this dataset which makes use of a pre-trained feature extractor achieves 0.979 AUROC averaged across all classes (Defard et al., 2020).

5.2. Anomaly Segmentation

For the evaluation of the proposed CPC-AD model’s anomaly segmentation performance, we up-sample the sub-patch-wise InfoNCE loss values to match the pixel-wise ground truth annotations. To do so, we average the In-

foNCE losses of overlapping sub-patches and assign the resulting values to all affected pixels. This allows us to create anomaly segmentation masks at the resolution of half a sub-patch (32×32 pixels) that are of the same dimensions as the resized images (768×768).

In Table 2 in the Appendix, we compare the anomaly segmentation performance of the proposed CPC-AD method against previously published works from peer-reviewed venues. The best results on the MVTec-AD dataset are achieved with extensive models that are pre-trained on ImageNet, such as **FCDD** and **PaDiM** (Liznerski et al., 2021; Defard et al., 2020), or make use of additional artificial anomalies and ensemble methods, such as **CutPaste** (Li et al., 2021a). Our model is trained from scratch and uses merely the provided training data, making for a less complex and more general method. The proposed **CPC-AD** approach is further outperformed by one autoencoding approach (**AE-SSIM**) and a partially contrastive approach (**DistAug**), but is on par with another autoencoding approach (**AE-L2**). Our proposed method outperforms the GAN-based approach (**AnoGAN**) (Bergmann et al., 2019; Schlegl et al., 2017). Interestingly, the CPC-AD model scores relatively well on textures, compared to similar models.

Nonetheless, although the quantitative results achieved with CPC-AD are not state-of-the-art, the model succeeds in generating accurate segmentation masks for most classes (Fig. 2). Even for classes with a low pixelwise AUROC score, such as pill, it can be seen that the created segmentation masks correctly highlight anomalous input regions, although there is some background noise. This corresponds with the comparatively high detection performance that the CPC-AD method achieves for this class (Table 1). These results indicate that part of the low segmentation scores (compared to the detection scores) could be due to small spatial deviations from the ground truth. This effect might be exacerbated by the relatively low resolution of the segmentation masks that our patch-wise approach creates. Nonetheless, we argue that this resolution would be sufficient in practice to provide interpretable results for human inspection. Overall, CPC-AD provides a promising first step towards anomaly segmentation methods that are based on contrastive learning.

6. Conclusion

Overall, the CPC-AD model shows that contrastive learning can be used not just for anomaly detection, but also for anomaly segmentation. The proposed method performs well on the anomaly detection task, with competitive results for a majority of the data. Additionally the generated segmentation masks provide a promising first step towards anomaly segmentation methods that are based on contrastive losses.

References

- Bergmann, P., Löwe, S., Fauser, M., Sattlegger, D., and Steger, C. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *arXiv preprint arXiv:1807.02011*, 2018.
- Bergmann, P., Fauser, M., Sattlegger, D., and Steger, C. Mvtec ad – a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Bergmann, P., Fauser, M., Sattlegger, D., and Steger, C. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4183–4192, 2020.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 13–18 Jul 2020.
- Defard, T., Setkov, A., Loesch, A., and Audigier, R. Padim: a patch distribution modeling framework for anomaly detection and localization. *arXiv preprint arXiv:2011.08785*, 2020.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Teh, Y. W. and Titterton, M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 297–304, 13–15 May 2010.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Hénaff, O. J., Srinivas, A., Fauw, J. D., Razavi, A., Doersch, C., Eslami, S. M. A., and van den Oord, A. Data-efficient image recognition with contrastive predictive coding. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4182–4192, 13–18 Jul 2020.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- Kauffmann, J., Ruff, L., Montavon, G., and Müller, K.-R. The clever hans effect in anomaly detection. *arXiv preprint arXiv:2006.10609*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *Proceedings of the 7th International Conference on Learning Representations*, 2015.
- Kopuklu, O., Zheng, J., Xu, H., and Rigoll, G. Driver anomaly detection: A dataset and contrastive learning approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 91–100, January 2021.
- Li, C.-L., Sohn, K., Yoon, J., and Pfister, T. Cutpaste: Self-supervised learning for anomaly detection and localization. *arXiv preprint arXiv:2104.04015*, 2021a.
- Li, J., Zhou, P., Xiong, C., and Hoi, S. C. H. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2021b.
- Liu, W., Li, R., Zheng, M., Karanam, S., Wu, Z., Bhanu, B., Radke, R. J., and Camps, O. Towards visually explaining variational autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Liznerski, P., Ruff, L., Vandermeulen, R. A., Franks, B. J., Kloft, M., and Müller, K.-R. Explainable deep one-class classification. *Proceedings of the 7th International Conference on Learning Representations*, 2021.
- Löwe, S., O’Connor, P., and Veeling, B. Putting an end to end-to-end: Gradient-isolated learning of representations. In *Advances in Neural Information Processing Systems*, pp. 3039–3051, 2019.
- Luo, W., Gu, Z., Liu, J., and Gao, S. Encoding structure-texture relation with p-net for anomaly detection in retinal images. In *European conference on computer vision*, pp. 360–377. Springer, 2020.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2019.
- Qiu, C., Pfrommer, T., Kloft, M., Mandt, S., and Rudolph, M. Neural transformation learning for deep anomaly detection beyond images. *arXiv preprint arXiv:2103.16440*, 2021.

Ruff, L., Vandermeulen, R. A., Görnitz, N., Binder, A., Müller, E., Müller, K.-R., and Kloft, M. Deep semi-supervised anomaly detection. *arXiv preprint arXiv:1906.02694*, 2020.

Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., and Müller, K.-R. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5):756–795, 2021.

Schlegl, T., Seeböck, P., Waldstein, S. M., Schmidt-Erfurth, U., and Langs, G. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging*, pp. 146–157. Springer, 2017.

Sohn, K., Li, C.-L., Yoon, J., Jin, M., and Pfister, T. Learning and evaluating representations for deep one-class classification. In *International Conference on Learning Representations*, 2021.

Tack, J., Mo, S., Jeong, J., and Shin, J. Csi: Novelty detection via contrastive learning on distributionally shifted instances. In *34th Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

Winkens, J., Bunel, R., Roy, A. G., Stanforth, R., Natarajan, V., Ledsam, J. R., MacWilliams, P., Kohli, P., Karthikesalingam, A., Kohl, S., Cemgil, T., Eslami, S. M. A., and Ronneberger, O. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*, 2020.

Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

Zhou, C. and Paffenroth, R. C. Anomaly detection with robust deep autoencoders. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 665–674, 2017.

A. Additional Results

A.1. Anomaly Segmentation

In Table 2, we compare the anomaly segmentation performance of the proposed CPC-AD method against previously published works from peer-reviewed venues.

Table 2. Anomaly segmentation – pixelwise AUROC score on the MVTec-AD test-set per category. AE-SSIM and AE-L2 (Bergmann et al., 2020; Liznerski et al., 2021), AnoGAN (Schlegl et al., 2017; Liznerski et al., 2021), DistAug in combination with KDE (Sohn et al., 2021) and CutPaste (Li et al., 2021a), as well as our proposed CPC-AD model do not use a pre-trained feature extractor. FCDD (Liznerski et al., 2021) and PaDiM (Defard et al., 2020) do make use of pre-training to improve their results (denoted with an *).

	AE-SSIM	AE-L2	AnoGAN	DistAug	CutPaste	FCDD*	PaDiM*	CPC-AD
Bottle	0.93	0.86	0.86	-	0.98	0.97	0.98	0.89
Cable	0.82	0.86	0.78	-	0.9	0.9	0.97	0.84
Capsule	0.94	0.88	0.84	-	0.97	0.93	0.99	0.72
Carpet	0.87	0.59	0.54	-	0.98	0.96	0.99	0.74
Grid	0.94	0.9	0.58	-	0.98	0.91	0.97	0.80
Hazelnut	0.97	0.95	0.87	-	0.97	0.95	0.98	0.81
Leather	0.78	0.75	0.64	-	1.0	0.98	0.99	0.94
Metal nut	0.89	0.86	0.76	-	0.93	0.94	0.97	0.76
Pill	0.91	0.85	0.87	-	0.96	0.81	0.96	0.77
Screw	0.96	0.96	0.8	-	0.97	0.86	0.99	0.65
Tile	0.59	0.51	0.5	-	0.91	0.91	0.94	0.82
Toothbrush	0.92	0.93	0.9	-	0.98	0.94	0.99	0.81
Transistor	0.9	0.86	0.8	-	0.93	0.88	0.98	0.90
Wood	0.73	0.73	0.62	-	0.96	0.88	0.95	0.82
Zipper	0.88	0.77	0.78	-	0.99	0.92	0.99	0.95
Mean	0.87	0.82	0.74	0.90	0.96	0.92	0.98	0.82
Texture mean	0.78	0.70	0.58	0.83	0.97	0.93	0.97	0.83
Object mean	0.91	0.88	0.83	0.94	0.96	0.91	0.98	0.81