

# Rethinking Assumptions in Deep Anomaly Detection

Lukas Ruff<sup>1\*</sup> Robert A. Vandermeulen<sup>2\*</sup> Billy Joe Franks<sup>3</sup> Klaus-Robert Müller<sup>2,4,5</sup> Marius Kloft<sup>3</sup>

## Abstract

Though anomaly detection (AD) can be viewed as a classification problem (nominal vs. anomalous) it is usually treated in an unsupervised manner since one typically does not have access to, or it is infeasible to utilize, a dataset that sufficiently characterizes what it means to be “anomalous.” In this paper we present results demonstrating that this intuition surprisingly seems not to extend to deep AD on images. For a recent AD benchmark on ImageNet, classifiers trained to discern between normal samples and just a few (64) random natural images are able to outperform the current state of the art in deep AD. Experimentally we discover that the multiscale structure of image data makes example anomalies exceptionally informative.

## 1. Introduction

Anomaly detection (AD) (Chandola et al., 2009) is the task of determining if a sample is anomalous compared to a corpus of data. Recently there has been a great interest in developing novel deep methods for AD (Ruff et al., 2021; Pang et al., 2021). Some of the best performing new AD methods for images were proposed by Golan & El-Yaniv (2018) and Hendrycks et al. (2019b). These methods, like most previous works on AD, are performed in an *unsupervised* way: they only utilize an unlabeled corpus of mostly nominal data. While AD can be interpreted as a classification problem of “nominal vs. anomalous,” it is typically treated as an unsupervised problem due to the rather tricky issue of finding or constructing a dataset that somehow captures *everything different* from a nominal dataset.

One often has, in addition to a corpus of nominal data, access to some data which is known to be anomalous. There exist deep methods for incorporating anomalous data to aug-

\*Equal contribution <sup>1</sup>Aignostics, Germany (majority of work done while with TU Berlin) <sup>2</sup>TU Berlin, Germany <sup>3</sup>TU Kaiserslautern, Germany <sup>4</sup>Korea University, Seoul, South Korea <sup>5</sup>MPII, Saarbrücken, Germany. Correspondence to: Lukas Ruff <contact@lukasruff.com>.

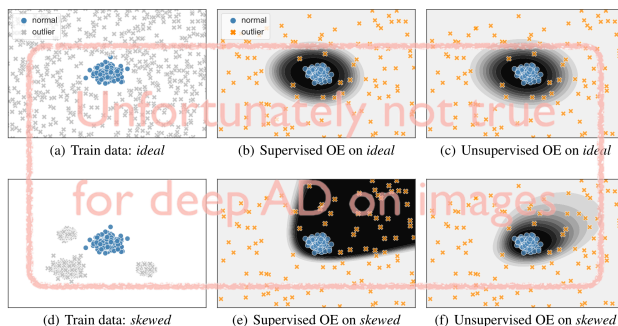


Figure 1. The decision boundaries of a supervised OE method (neural net with binary cross entropy) and an unsupervised OE method (neural net with hypersphere loss) on two toy data settings: *ideal* ((a)–(c)) and *skewed* ((d)–(f)). The unsupervised OE method ((c) + (f)) learns compact decision regions of the nominal class. A supervised OE approach ((b) + (e)) learns decision regions that do not generalize well on this toy AD task. Our results suggest that this intuition does not hold for a deep approach to image AD, where supervised OE performs remarkably well.

ment unsupervised AD (Hendrycks et al., 2019a; Ruff et al., 2020). This setting has also been called “semi-supervised” AD (Görnitz et al., 2013; Ruff et al., 2020). In Hendrycks et al. (2019a) it was noted that, for an image AD problem, one has access to a virtually limitless amount of random natural images from the internet that are likely not nominal, and that such data should be utilized to improve unsupervised methods. They term the utilization of such data *outlier exposure* (OE). The state-of-the-art method presented in Hendrycks et al. (2019b) utilizes tens of thousands of OE samples combined with a modified version of the self-supervised method from Golan & El-Yaniv (2018) and is one of the best performing AD method to date on standard image AD benchmarks. For clarity, we here delineate the following three basic approaches to anomaly detection:

Unsupervised: Methods trained on (mostly) nominal data. This is the classic and most common approach to AD.

Unsupervised OE: Adaptations of unsupervised methods that incorporate auxiliary data that is not nominal. Elsewhere this is also called “semi-supervised” AD (Görnitz et al., 2013; Ruff et al., 2020).

Supervised OE: The approach of simply applying a standard classification method to discern between nominal data

and an auxiliary dataset that is not nominal.

Using unsupervised OE rather than supervised OE to discern between the nominal data and known anomalies seems intuitive since the presented anomalies likely do not completely characterize “anomalousness.” This is illustrated in Figure 1. This intuition and the benefits of the unsupervised OE approach when incorporating known anomalies has also been observed previously (Tax, 2001; Görnitz et al., 2013; Ruff et al., 2020).

In this paper, we present experimental results that challenge the assumption that deep AD on images needs an unsupervised approach (with or without OE). We find that, using the same experimental OE setup as Hendrycks et al. (2019b), a standard classifier is able to outperform current state-of-the-art AD methods on the one vs. rest AD benchmarks on MNIST and CIFAR-10. The one vs. rest benchmark has been recommended as a general approach to experimentally validate AD methods (Emmott et al., 2013). This benchmark applied to the aforementioned datasets is used as a litmus test in virtually all deep AD papers published at top-tier venues; see for example (Ruff et al., 2018; Deecke et al., 2018; Golan & El-Yaniv, 2018; Hendrycks et al., 2019b; Akcay et al., 2018; Abati et al., 2019; Perera et al., 2019; Wang et al., 2019a; Ruff et al., 2020; Bergman & Hoshen, 2020; Kim et al., 2020). Additionally, we find that remarkably few OE examples are necessary to characterize “anomalousness.” With 128 OE samples a classifier is competitive with state-of-the-art unsupervised methods on the CIFAR-10 one vs. rest benchmark. With only 64 OE samples a classifier outperforms unsupervised methods (with or without OE) on the ImageNet one vs. rest benchmark (Hendrycks et al., 2019b). This test was recently proposed as a more challenging successor to the CIFAR-10 benchmark.

Our results seem to contradict the following pieces of common wisdom in deep learning and AD:

- Many (thousands) samples are needed for a deep method to understand a class (Goodfellow et al., 2016).
- Anomalies are unconcentrated and thus inherently difficult to characterize with data (Steinwart et al., 2005; Chapelle et al., 2006).

These points should imply that classification with few OE samples should be ineffective at deep AD. Instead, we find that relatively few random OE samples are necessary to yield state-of-the-art detection performance. In all of our experiments, the nominal and OE data available during training are exactly those used in (Hendrycks et al., 2019b) and do not contain any representatives from the ground-truth anomaly classes. The OE data is not tailored to be representative of the anomalies used at test time. Based on the presence of information at multiple spatial scales in images (Olshausen & Field, 1996), which is one key difference between classic AD and deep image AD, we hypothesize that

each OE image contains multiple features at different scales that present informative examples of anomalousness.

## 2. Deep One-Class Classification

Deep one-class classification (Ruff et al., 2018), which learns (or transfers) data representations such that normal data is concentrated in feature space, has been introduced as a deep learning extension of the one-class classification approach to anomaly detection (Schölkopf et al., 2001; Tax, 2001; Ruff et al., 2021). Specifically, the Deep SVDD method (Ruff et al., 2018) is trained to map nominal samples close to a center  $c$  in feature space. For a neural network  $\phi_\theta$  with parameters  $\theta$ , the Deep SVDD objective is given by

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \|\phi_\theta(\mathbf{x}_i) - c\|^2. \quad (1)$$

In Ruff et al. (2020), an extension of Deep SVDD that incorporates known anomalies is proposed, called *Deep Semi-supervised Anomaly Detection* (Deep SAD). Deep SAD trains a network to concentrate nominal data near a center  $c$  and maps anomalous samples away from that center. This is therefore an unsupervised OE approach to AD. We here present a principled modification of Deep SAD based on cross-entropy classification that concentrates nominal samples. We call this method *hypersphere classification* (HSC). We found this modification to significantly improve upon the performance of Deep SAD and use it in our experiments as a representative of the unsupervised OE approach to AD.

Let  $\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$  be a dataset with  $\mathbf{x}_i \in \mathbb{R}^d$  and  $y \in \{0, 1\}$  with  $y = 1$  denoting nominal and  $y = 0$  anomalous data points. Let  $\phi_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^r$  be a neural network and  $l : \mathbb{R}^r \rightarrow [0, 1]$  be a function which maps the output to a probabilistic score. Then, we can formulate the cross-entropy loss as

$$-\frac{1}{n} \sum_{i=1}^n y_i \log l(\phi_\theta(\mathbf{x}_i)) + (1 - y_i) \log (1 - l(\phi_\theta(\mathbf{x}_i))). \quad (2)$$

For standard binary deep classifiers,  $l$  is a linear layer followed by the sigmoid activation and the decision region for the mapped samples  $\phi_\theta(\mathbf{x}_1), \dots, \phi_\theta(\mathbf{x}_n)$  is a half-space  $S$ . In this case the preimage of  $S$ ,  $\phi_\theta^{-1}(S)$ , is not guaranteed to be compact. In order to enforce the preimage of our nominal decision region to be compact, thereby encouraging the mapped nominal data to be concentrated in a way similar to Deep SAD, we propose  $l$  to be a radial basis function. To construct a spherical decision boundary we let  $l(z) := \exp(-\|z\|^2)$ . In this case, (2) becomes

$$\frac{1}{n} \sum_{i=1}^n y_i \|\phi_\theta(\mathbf{x}_i)\|^2 - (1 - y_i) \log (1 - \exp(-\|\phi_\theta(\mathbf{x}_i)\|^2)).$$

If there are no anomalies, the HSC loss simplifies to  $\frac{1}{n} \sum_{i=1}^n \|\phi_\theta(\mathbf{x}_i)\|^2$ . For  $c = 0$ , we thus recover Deep

SVDD (1) as a special case. Similar to Deep SVDD/SAD, we define our anomaly score as  $s(\mathbf{x}) := \|\phi_\theta(\mathbf{x})\|^2$ .

Motivated by robust statistics (Hampel et al., 2005; Huber & Ronchetti, 2009) we also considered replacing  $l$  with other radial functions where the squared-norm is replaced with a robust alternative. We found that using a pseudo-Huber loss (Charbonnier et al., 1997)  $l(\mathbf{z}) = \exp(-h(\mathbf{z}))$  that interpolates between squared and absolute value penalization yielded the best results:  $h(\mathbf{z}) = \sqrt{\|\mathbf{z}\|^2 + 1} - 1$ . We include a sensitivity analysis comparing various choices of norms for the hypersphere classifier in Appendix D.

### 3. Experiments

**One vs. Rest Benchmark** The one vs. rest evaluation procedure is a ubiquitous benchmark in the deep AD literature as mentioned in the introduction (Section 1). This benchmark constructs AD settings from classification datasets (e.g., MNIST) by considering the “one” class (e.g., digit 0) as being nominal and the “rest” classes (e.g., digits 1–9) as being anomalous at test time. In each experiment, we train a model using only the training set of the nominal class as well as random samples from an OE set (e.g., EMNIST-Letters) which is disjoint from the ground-truth anomaly classes of the benchmark. We use the same OE auxiliary datasets as suggested in previous works (Hendrycks et al., 2019a;b). To evaluate detection performance, we use the common Area Under the ROC curve (AUC) on the one vs. rest test sets. This is repeated over classes and multiple random seeds.

#### Datasets

**MNIST:** The ten MNIST classes are used as our one vs. rest classes. For OE we use the EMNIST-Letters dataset (Cohen et al., 2017) which shares no common classes with MNIST.

**CIFAR-10:** The ten CIFAR-10 classes are used as our one vs. rest classes. For OE we use 80 Million Tiny Images (80MTI) (Torralla et al., 2008) with CIFAR-10 and CIFAR-100 images removed. This follows the experimental setup in Hendrycks et al. (2019b). In one ablation experiment on OE diversity, we alternatively use CIFAR-100 for OE.

**ImageNet:** 30 classes from the ImageNet-1K (Deng et al., 2009) dataset are used as the one vs. rest classes. These are the classes as proposed in Hendrycks et al. (2019b). For OE we use the ImageNet-22K dataset with the ImageNet-1K removed also following Hendrycks et al. (2019b).

**Methods** We present results from methods that achieve state-of-the-art performance on the one vs. rest benchmarks.

**Unsupervised:** We use Deep SVDD (Ruff et al., 2018), GT (Golan & El-Yaniv, 2018), GT+ (Hendrycks et al., 2019b), and IT (Huang et al., 2019) as shorthands for the state-of-the-art unsupervised methods. We also report results of the shallow SVDD (Tax, 2001; Schölkopf et al., 2001) baseline.

**Unsupervised OE:** We implement the HSC from Section

2 and Deep SAD (Ruff et al., 2020) as unsupervised OE methods. We also report the results from the state-of-the-art unsupervised OE GT+ variant (Hendrycks et al., 2019b).

**Supervised OE:** We consider a standard binary cross-entropy classifier (BCE). Moreover, we implement the Focal loss classifier (Lin et al., 2017), a BCE variant that specifically addresses class imbalance, which was also presented in Hendrycks et al. (2019b). We indicate the results from (Hendrycks et al., 2019b) with an asterisk as Focal\*. We set  $\gamma = 2$  as recommended in (Lin et al., 2017).

We provide more background on the above methods as well as network architecture and optimization details in Appendices A and F respectively. Due to space constraints, we report the mean AUC performance over all classes and seeds of the competitive methods in the main paper and report individual results per class and method in Appendix G.

**Varying the OE Size** For HSC and BCE, we also present extensive experiments showing the performance as the OE training set size is varied on a log scale starting from just  $2^0 = 1$  sample to using the maximal amount of OE data such that no OE sample is seen twice during training. If the number of OE samples is less than the OE batch size of 128, we sample with replacement. Note that applying data augmentation introduces some variety to the OE set, even in the extreme case of having only  $2^0 = 1$  sample.

#### 3.1. Results on the CIFAR-10 Benchmark

The results on CIFAR-10 are shown in Table 1. We observe that the unsupervised OE methods GT+ and HSC yield a comparable detection performance. Interestingly, the supervised Focal and BCE methods also show state-of-the-art performance, with BCE attaining the best mean AUC overall. To understand the informativeness of OE for unsupervised and supervised approaches, we compare the performance of HSC and BCE while varying the OE set size. The results in Figure 2(a) demonstrate that surprisingly few OE samples already yield a very competitive detection performance.

Table 1. Mean AUC detection performance in % (over 10 classes and 10 seeds) on the CIFAR-10 one vs. rest benchmark using 80MTI as OE (\* indicates results from the literature).

Unsupervised		Unsupervised OE			Supervised OE		
DSVDD*	GT+*	GT+*	DSAD	HSC	Focal*	Focal	BCE
64.8	90.1	95.6	94.5	95.9	87.3	95.8	<b>96.1</b>

#### 3.2. Results on the ImageNet Benchmark

The results on ImageNet are given in Table 2. GT+\* and Focal\* are the results from Hendrycks et al. (2019b), where this benchmark was introduced. Deep SAD, HSC, Focal, and BCE all outperform the current state of the art (GT+\*) by a surprisingly wide margin. We are unsure as to why the Focal\* results from Hendrycks et al. (2019b) are so poor since their experimental code is not public. We found perfor-

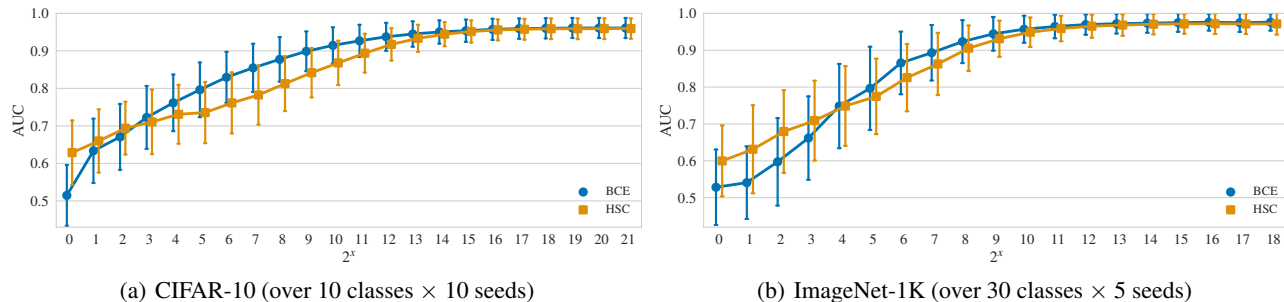


Figure 2. Mean AUC detection performance in % on the CIFAR-10 and ImageNet-1K one vs. rest benchmarks when varying the number of 80MTI and ImageNet-22K OE samples respectively.

mance to be insensitive to the choice of  $\gamma$  in the Focal loss (see Appendix E). Not balancing the number of nominal and OE samples in each batch may be one explanation. We also compare the performance of HSC and BCE while varying the OE set size. The results are in Figure 2(b). Again we see that there is a transition from HSC to BCE performing best, which can be understood as a transition from unsupervised OE to supervised OE. Remarkably, classification beats previous methods with only 64 OE samples.

Table 2. Mean AUC detection performance in % (over 30 classes and 10 seeds) on the ImageNet-1K one vs. rest benchmark using ImageNet-22K (with the 1K classes removed) as OE (\* indicates results from the literature).

Unsupervised OE			Supervised OE		
GT+*	DSAD	HSC	Focal*	Focal	BCE
85.7	96.7	97.3	56.1	97.5	<b>97.7</b>

In addition to studying the effect of OE set size, we also evaluate the effect of OE data diversity on detection performance. For this, we vary the number of anomaly classes that comprise the OE set. As expected, we find that performance overall increases with OE data diversity, but interestingly drawing OE samples from just one class (which is not present as anomaly at test time!) yields surprisingly good performance. We provide the full results in Appendix B.

### 3.3. Removing Multiscale Information

To investigate the hypothesis that the exceptional informativeness of OE samples is due to the multiscale structure of natural images, we perform an experiment in which we remove small scale features from the OE data. For this we compare the performance of HSC and BCE on the ImageNet one vs. rest task while increasingly blurring the OE samples with a Gaussian filter. The blurring gradually removes the small scale (high frequency) features from the OE data.

In Figure 3, we see that detection performance quickly drops with even a small amount of blurring. With sufficient blurring, the unsupervised OE method (HSC) performs better. We provide results for MNIST and CIFAR-10 in Appendix

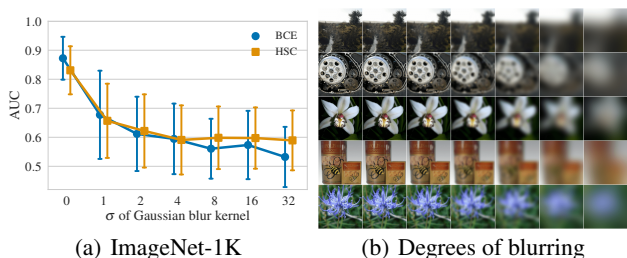


Figure 3. Mean AUC detection performance in % (over 30 classes and 5 seeds) on ImageNet-1K for  $2^6 = 64$  OE samples when increasingly blurring the OE images with a Gaussian kernel (see (a)). A visual example of the Gaussian blurring is shown in (b). The sharp decline in AUC suggests the exceptional informativeness of OE for images is due to their multiscale nature.

C, where we observe a similar decrease in performance for CIFAR-10. For MNIST, however, HSC outperforms BCE at any degree of blurring and retains a good performance. This suggests that there are two regimes in deep image AD: (i) on high-dimensional multiscale images, OE samples are highly informative and standard as well as hypersphere classification perform similarly well, and (ii) on low-dimensional images (e.g., MNIST), OE samples are also informative, but hypersphere performs better than standard classification, which reflects the classic intuition that anomaly detection requires a compact model.

## 4. Conclusion

We have shown that deep AD on images displays a phenomenon which is quite different from what is expected in classic AD. Compared to classic AD, a few example outliers are exceptionally informative on common image AD benchmarks. Furthermore, we have shown that this phenomenon is tied to the multiscale nature of natural images. Finally, please note that we do not claim that a supervised OE approach is *the* solution to AD in general, or that there is no utility for unsupervised OE. However, our results suggest that it may be time for the community to move to more challenging benchmarks (e.g., MVTec-AD (Bergmann et al., 2019)) to gauge the significance of deep AD works.

## References

- Abati, D., Porrello, A., Calderara, S., and Cucchiara, R. Latent space autoregression for novelty detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 481–490, 2019.
- Akçay, S., Atapour-Abarghouei, A., and Breckon, T. P. GANomaly: Semi-supervised anomaly detection via adversarial training. In *Asian Conference on Computer Vision*, pp. 622–637, 2018.
- Bergman, L. and Hoshen, Y. Classification-based anomaly detection for general data. In *International Conference on Learning Representations*, 2020.
- Bergmann, P., Fauser, M., Sattlegger, D., and Steger, C. MVTec AD—A comprehensive real-world dataset for unsupervised anomaly detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9592–9600, 2019.
- Chandola, V., Banerjee, A., and Kumar, V. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, 2009.
- Chapelle, O., Schölkopf, B., and Zien, A. *Semi-Supervised Learning*. The MIT Press, Cambridge, Massachusetts, 2006.
- Charbonnier, P., Blanc-Féraud, L., Aubert, G., and Barlaud, M. Deterministic edge-preserving regularization in computed imaging. *IEEE Transactions on Image Processing*, 6(2):298–311, 1997.
- Chen, J., Sathe, S., Aggarwal, C. C., and Turaga, D. S. Outlier detection with autoencoder ensembles. In *SIAM International Conference on Data Mining*, pp. 90–98, 2017.
- Cohen, G., Afshar, S., Tapson, J., and Van Schaik, A. EMNIST: Extending MNIST to handwritten letters. In *International Joint Conference on Neural Networks*, pp. 2921–2926, 2017.
- Deecke, L., Vandermeulen, R. A., Ruff, L., Mandt, S., and Kloft, M. Image anomaly detection with generative adversarial networks. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pp. 3–17, 2018.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, pp. 4171–4186, 2019.
- Emmott, A. F., Das, S., Dietterich, T., Fern, A., and Wong, W.-K. Systematic construction of anomaly detection benchmarks from real data. In *KDD 2013 Workshop on Outlier Detection and Description*, pp. 16–21, 2013.
- Erfani, S., Baktashmotlagh, M., Rajasegarar, S., Karunasekera, S., and Leckie, C. R1SVM: A randomised nonlinear approach to large-scale anomaly detection. In *AAAI Conference on Artificial Intelligence*, pp. 432–438, 2015.
- Erfani, S. M., Rajasegarar, S., Karunasekera, S., and Leckie, C. High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning. *Pattern Recognition*, 58:121–134, 2016.
- Gidaris, S., Singh, P., and Komodakis, N. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- Golan, I. and El-Yaniv, R. Deep anomaly detection using geometric transformations. In *Advances in Neural Information Processing Systems*, pp. 9758–9769, 2018.
- Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. MIT press, 2016.
- Görnitz, N., Kloft, M., Rieck, K., and Brefeld, U. Toward supervised anomaly detection. *Journal of Artificial Intelligence Research*, 46:235–262, 2013.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons, 2005.
- Hawkins, S., He, H., Williams, G., and Baxter, R. Outlier detection using replicator neural networks. In *International Conference on Data Warehousing and Knowledge Discovery*, volume 2454, pp. 170–180, 2002.
- Hendrycks, D., Mazeika, M., and Dietterich, T. G. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2019a.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in Neural Information Processing Systems*, pp. 15637–15648, 2019b.
- Huang, C., Cao, J., Ye, F., Li, M., Zhang, Y., and Lu, C. Inverse-transform autoencoder for anomaly detection. *arXiv preprint arXiv:1911.10676*, 2019.
- Huang, F. J. and LeCun, Y. Large-scale learning with SVM and convolutional nets for generic object categorization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 284–291, 2006.

- Huber, P. J. and Ronchetti, E. M. *Robust Statistics*. John Wiley & Sons, 2nd edition, 2009.
- Ioffe, S. and Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning*, volume 37, pp. 448–456, 2015.
- Kim, K. H., Shim, S., Lim, Y., Jeon, J., Choi, J., Kim, B., and Yoon, A. S. RaPP: Novelty detection with reconstruction along projection pathway. In *International Conference on Learning Representations*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pp. 3581–3589, 2014.
- Kriegel, H.-P., Schubert, M., and Zimek, A. Angle-based outlier detection in high-dimensional data. In *International Conference on Knowledge Discovery & Data Mining*, pp. 444–452, 2008.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *International Conference on Computer Vision*, pp. 2980–2988, 2017.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pp. 3111–3119, 2013.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., and Lakshminarayanan, B. Do deep generative models know what they don’t know? In *International Conference on Learning Representations*, 2019.
- Nguyen, D. T., Lou, Z., Klar, M., and Brox, T. Anomaly detection with multiple-hypotheses predictions. In *International Conference on Machine Learning*, volume 97, pp. 4800–4809, 2019.
- Odena, A. Semi-supervised learning with generative adversarial networks. In *ICML 2016 Workshop on Data Efficient Machine Learning*, 2016.
- Oliver, A., Odena, A., Raffel, C., Cubuk, E. D., and Goodfellow, I. J. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, volume 31, pp. 3235–3246, 2018.
- Olshausen, B. A. and Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- Pang, G., Shen, C., Cao, L., and Hengel, A. V. D. Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2), 2021. doi: 10.1145/3439950.
- Perera, P., Nallapati, R., and Xiang, B. OCGAN: One-class novelty detection using GANs with constrained latent representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2898–2906, 2019.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *North American Chapter of the Association for Computational Linguistics*, pp. 2227–2237, 2018.
- Polonik, W. Measuring mass concentrations and estimating density contour clusters—an excess mass approach. *The Annals of Statistics*, 23(3):855–881, 1995.
- Rasmus, A., Berglund, M., Honkala, M., Valpola, H., and Raiko, T. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pp. 3546–3554, 2015.
- Ruff, L., Vandermeulen, R. A., Görnitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., and Kloft, M. Deep one-class classification. In *International Conference on Machine Learning*, volume 80, pp. 4390–4399, 2018.
- Ruff, L., Vandermeulen, R. A., Görnitz, N., Binder, A., Müller, E., Müller, K.-R., and Kloft, M. Deep semi-supervised anomaly detection. In *International Conference on Learning Representations*, 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. doi: 10.1109/JPROC.2021.3052449.
- Sakurada, M. and Yairi, T. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *2nd Workshop on Machine Learning for Sensory Data Analysis (MLSDA 2014)*, pp. 4–11, 2014.
- 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, 2017.
- Schlegl, T., Seeböck, P., Waldstein, S. M., Langs, G., and Schmidt-Erfurth, U. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical Image Analysis*, 54:30–44, 2019.
- Schölkopf, B. and Smola, A. J. *Learning with Kernels*. MIT press, 2002.

- Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., and Williamson, R. C. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7): 1443–1471, 2001.
- Steinwart, I., Hush, D., and Scovel, C. A classification framework for anomaly detection. *Journal of Machine Learning Research*, 6(Feb):211–232, 2005.
- Tax, D. M. J. *One-Class Classification*. PhD thesis, Delft University of Technology, 2001.
- Torralba, A., Fergus, R., and Freeman, W. T. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008.
- Tsybakov, A. B. On nonparametric estimation of density level sets. *The Annals of Statistics*, 25(3):948–969, 1997.
- Vert, R. and Vert, J.-P. Consistency and convergence rates of one-class SVMs and related algorithms. *Journal of Machine Learning Research*, 7(May):817–854, 2006.
- Wang, J., Sun, S., and Yu, Y. Multivariate triangular quantile maps for novelty detection. In *Advances in Neural Information Processing Systems*, pp. 5061–5072, 2019a.
- Wang, S., Zeng, Y., Liu, X., Zhu, E., Yin, J., Xu, C., and Kloft, M. Effective end-to-end unsupervised outlier detection via inlier priority of discriminative network. In *Advances in Neural Information Processing Systems*, pp. 5960–5973, 2019b.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference*, 2016.
- Zeiler, M. D. and Fergus, R. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pp. 818–833, 2014.
- Zenati, H., Romain, M., Foo, C.-S., Lecouat, B., and Chandrasekhar, V. Adversarially learned anomaly detection. In *IEEE International Conference on Data Mining*, pp. 727–736, 2018.
- Zhou, C. and Paffenroth, R. C. Anomaly detection with robust deep autoencoders. In *International Conference on Knowledge Discovery & Data Mining*, pp. 665–674, 2017.

# Appendix

## A. Methods Background

While there exist many shallow methods for AD, it has been observed that these methods perform poorly for high-dimensional data (Huang & LeCun, 2006; Kriegel et al., 2008; Erfani et al., 2015; 2016). Deep approaches have been proposed to fill this gap. The most common approaches to deep AD employ autoencoders trained on nominal data, where samples not reconstructed well are deemed anomalous (Hawkins et al., 2002; Sakurada & Yairi, 2014; Chen et al., 2017; Zhou & Paffenroth, 2017; Huang et al., 2019; Nguyen et al., 2019; Kim et al., 2020). Deep generative models have also been used to detect anomalies via a variety of methods (Schlegl et al., 2017; Deecke et al., 2018; Zenati et al., 2018; Schlegl et al., 2019), yet their effectiveness has been called into question (Nalisnick et al., 2019).

Another recent avenue of research on deep AD uses self-supervision on images (Gidaris et al., 2018; Golan & El-Yaniv, 2018; Wang et al., 2019b; Hendrycks et al., 2019b). In Golan & El-Yaniv (2018), the authors use a composition of image transformations—including identity, rotations, flips, and translations—to create a self-supervised classification task. Every training sample is transformed using each of these transformations and a label is assigned to every transformed sample corresponding to the applied transformation. This creates a multi-class classification task for predicting image transformations. A network is then trained on this data to predict the applied transformation. For a test sample these transformations and network outputs are utilized to determine an anomaly score.

To our knowledge, one of the best performing AD method on image data is the self-supervised approach from Hendrycks et al. (2019b) which extends Golan & El-Yaniv (2018)’s method by using three classification heads to predict a combination of three types of transformations. They train their network on transformed nominal data as was done in Golan & El-Yaniv (2018). On a test sample the network’s certainty (how close to 1) on predicting correct transformations is used as an anomaly score, with certainty being a signifier that a sample is not anomalous. Essentially this assumes the network predictions to be less concentrated on the correct output for unfamiliar looking data. In that paper, the authors validate their method on the CIFAR-10 and ImageNet one vs. rest one-class learning benchmarks.

### A.1. Auxiliary Data and the State of the Art for Deep AD on Images

Many deep learning methods have been proposed to incorporate the large amount of unorganized data that is now easily accessible on the web. A common way to use this data is

via unsupervised or self-supervised learning. In the realm of NLP, word2vec (Mikolov et al., 2013) and more recent language models such as ELMo (Peters et al., 2018) or BERT (Devlin et al., 2019) are now standard and responsible for significant improvements on various NLP tasks. For image tasks, using an auxiliary dataset for pretraining has been found to be effective (Zeiler & Fergus, 2014). Moreover many deep semi-supervised methods have been introduced to enhance classification performance via incorporating unlabeled data into training (Kingma et al., 2014; Rasmus et al., 2015; Odena, 2016; Oliver et al., 2018).

The use of a large unstructured corpus of image data to improve deep AD was first proposed in Hendrycks et al. (2019a), where they call the general use of such data *outlier exposure* (OE). In Hendrycks et al. (2019b) the authors use OE to further improve existing self-supervised classification methods. They do this by training the aforementioned self-supervised methods to predict the uniform distribution for all transforms on OE samples, while leaving training on nominal samples unchanged. To our knowledge the AD method with OE presented in Hendrycks et al. (2019b) is the current state of the art on the CIFAR-10 and ImageNet image anomaly detection benchmarks, outperforming previous unsupervised AD methods with or without OE.

### A.2. Anomaly Detection as Binary Classification

Traditionally AD is understood as the problem of estimating the support (or level sets of the support) of the nominal data-generating distribution. This is also known as *density level set estimation* (Polonik, 1995; Tsybakov, 1997; Ruff et al., 2021). The motivation for density level set estimation is the common assumption that nominal data is concentrated whereas anomalies are not concentrated (Schölkopf & Smola, 2002). Steinwart et al. (2005) remark that the problem of density level set estimation can be interpreted as binary classification between the nominal and an anomalous distribution. Many of the classic AD methods (e.g., KDE or OC-SVM) implicitly assume the anomalies to follow a uniform, i.e. they make an uninformed prior assumption on the anomalous distribution (Steinwart et al., 2005). These methods, as well as a binary classifier trained to discriminate between nominal samples and uniform noise, are asymptotically consistent density level set estimators (Steinwart et al., 2005; Vert & Vert, 2006). Obviously it is better to directly estimate the level set rather than introducing the auxiliary task of classifying against uniform noise. Such a classification approach is particularly ineffective and inefficient in high dimensions since it would require massive amounts of noise samples to properly fill the sample space. As demonstrated through various experiments, however, we find that this intuition does not seem to extend to a deep approach to image anomaly detection when the anomalous examples are natural images.



## B. Diversity of the Outlier Exposure Data

Here we evaluate how data diversity influences detection performance for unsupervised and supervised OE, again comparing HSC to BCE. For this purpose, instead of 80MTI, we now use CIFAR-100 as OE varying the number of anomaly classes available for the CIFAR-10 benchmark. We further evaluate the methods on the MNIST one vs. rest benchmark where EMNIST-Letters is used as the OE dataset. For both experiments, the OE data is varied by choosing  $k$  classes at random for each random seed and using the union of these classes as the OE dataset.

The results are presented in Figure 4. As expected, the performance increases with the diversity of the OE dataset. Interestingly, drawing OE samples from just  $k = 1$  class, i.e. binary classification between the nominal class and a single OE class (which is not present as an anomaly class at test time!) already yields good detection performance on the CIFAR-10 benchmark. For example training a standard classification network to discern between automobiles and beavers performs competitively as an automobile anomaly detector, even when no beavers are present as anomalies during test time.

For the MNIST experiment we see that HSC outperforms BCE for any number of classes. We hypothesize that this is due to the lack of multiscale spatial structure in the MNIST and EMNIST datasets. This intuition is consistent with the classic understanding of AD mentioned in the introduction of the main paper (Section 1).




Figure 4. Mean AUC detection performance in % (over 10 classes with 10 seeds per class) on the MNIST with EMNIST-Letters OE (a) and CIFAR-10 with CIFAR-100 OE (b) one vs. rest benchmarks when varying the number of  $k$  classes that comprise the OE dataset.

## C. Removing Multiscale Information on MNIST and CIFAR-10

Here we present the experimental results on removing multiscale information (see Section 3.3) for MNIST and CIFAR-10. The results are given in Figures 5 and 6 for MNIST and CIFAR-10 respectively. For CIFAR-10, we can observe a similar decrease in performance as observed for ImageNet (see Figure 3). For MNIST, however, HSC outperforms

BCE at any degree of blurring and retains a good performance.

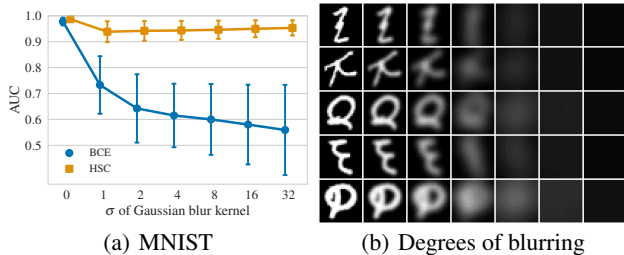


Figure 5. Mean AUC detection performance in % (over 10 classes and 10 seeds) on the MNIST one vs. rest benchmark when increasingly blurring the EMNIST OE images with a Gaussian kernel (see (a)). A visual example of the Gaussian blurring for some EMNIST OE samples is shown in (b). We see that HSC clearly outperforms BCE on MNIST, a dataset which has essentially no multiscale information.

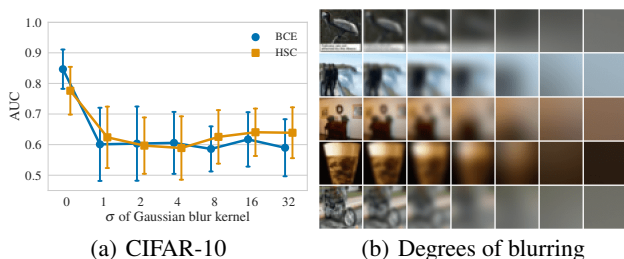


Figure 6. Mean AUC detection performance in % (over 10 classes and 10 seeds) on the CIFAR-10 one vs. rest benchmark for  $2^7 = 128$  OE samples when increasingly blurring the 80MTI OE images with a Gaussian kernel (see (a)). A visual example of the Gaussian blurring for some 80MTI OE samples is shown in (b). The rapid decrease in AUC on CIFAR-10 again suggests that the informativeness of OE on images is due to the multiscale structure of images.

## D. Hypersphere Classifier Sensitivity Analysis

Here we show results for the Hypersphere Classifier (HSC) we introduced in Section 3 when varying the radial function  $l(z) = \exp(-h(z))$ . For this, we run the CIFAR-10 one vs. rest benchmark with 80MTI OE experiment as presented in Table 1 in the main paper for different functions  $h : \mathbb{R}^r \rightarrow [0, \infty), z \mapsto h(z)$ . We also alter training to be with or without data augmentation in these experiments. The results are presented in Table 3. We see that data augmentation leads to an improvement in performance even in this case where we have the full 80MTI dataset as OE. HSC shows the overall best performance with data augmentation and using the robust Pseudo-Huber loss  $h(z) = \sqrt{\|z\|^2 + 1} - 1$ .

Table 3. Mean AUC detection performance in % (over 10 seeds) on the CIFAR-10 one vs. rest benchmark using 80MTI as OE for different choices of  $h(z)$  in the radial function  $l$  of the HSC.

Data augment.	$\ z\ _1$	$\ z\ _2$	$\ z\ _2^2$	$\sqrt{\ z\ _2^2 + 1} - 1$
w/o	90.6	92.3	89.1	91.8
w/	92.5	94.1	94.5	96.1

## E. Focal Loss With Varying $\gamma$

Here we include results showing how mean AUC detection performance changes with  $\gamma$  on the Focal loss. Since we balance every batch to contain 128 nominal and 128 OE samples during training, we set the weighting factor  $\alpha$  to be  $\alpha = 0.5$  (Lin et al., 2017). Again note that  $\gamma = 0$  corresponds to standard binary cross entropy. Figure 7 shows that mean AUC performance is insensitive to the choice of  $\gamma$  on the CIFAR-10 and ImageNet-1K one vs. rest benchmarks.




Figure 7. Focal loss detection performance in mean AUC in % when varying  $\gamma$  on the CIFAR-10 with 80MTI OE (a) and ImageNet-1K with ImageNet-22K OE (b) one vs. rest benchmarks.

## F. Network Architectures and Optimization

We always use the same underlying network  $\phi_\theta$  in each experimental setting for our HSC, Deep SAD, Focal, and BCE implementations to control architectural effects. For Focal and BCE, the output of the network  $\phi_\theta$  is followed by a linear layer with sigmoid activation. For the experiments on MNIST and CIFAR-10 we use standard LeNet-style networks having two and three convolutional layers followed by two fully connected layers respectively. We use batch normalization (Ioffe & Szegedy, 2015) and (leaky) ReLU activations in these networks. For our experiments on ImageNet we use the same WideResNet (Zagoruyko & Komodakis, 2016) as Hendrycks et al. (2019b), which has ResNet-18 as its architectural backbone. We use Adam (Kingma & Ba, 2015) for optimization and balance every batch to contain 128 nominal and 128 OE samples during training. For data augmentation, we use standard color jitter, random cropping, horizontal flipping, and Gaussian pixel

noise. We provide further dataset-specific details below.

### F.1. MNIST and CIFAR-10

On MNIST and CIFAR-10, we use LeNet-style networks having two and three convolutional layers and two fully connected layers respectively. Each convolutional layer is followed by batch normalization, a leaky ReLU activation, and max-pooling. The first fully connected layer is followed by batch normalization, and a leaky ReLU activation, while the last layer is just a linear transformation. The number of kernels in the convolutional layers are, from first to last: 16-32 (MNIST), and 32-64-128 (CIFAR-10). The fully connected layers have 64-32 (MNIST), and 512-256 (CIFAR-10) units respectively. We use Adam (Kingma & Ba, 2015) for optimization and balance every batch to contain 128 nominal and 128 OE samples during training. We train for 150 (MNIST) and 200 (CIFAR-10) epochs starting with a learning rate of  $\eta = 0.001$  and have learning rate milestones at 50, 100 (MNIST), and 100, 150 (CIFAR-10) epochs. The learning rate is reduced by a factor of 10 at every milestone.

### F.2. ImageNet

On ImageNet, we use exactly the same WideResNet (Zagoruyko & Komodakis, 2016) as was used in Hendrycks et al. (2019b), which has a ResNet-18 as architectural backbone. We use Adam (Kingma & Ba, 2015) for optimization and balance every batch to contain 128 nominal and 128 OE samples during training. We train for 150 epochs starting with a learning rate of  $\eta = 0.001$  and milestones at epochs 100 and 125. The learning rate is reduced by a factor of 10 at every milestone.

## G. Results on Individual Classes

For the CIFAR-10 one vs.-rest benchmark experiments from Section 3.1, we report the results for all individual classes and methods in Table 4. We additionally report the results with standard deviations for our implementations in Table 4. For the ImageNet-1K one vs.-rest benchmark experiments from Section 3.2, we present the performance over all individual classes with standard deviations in Table 6. For the experiments on varying the number of OE samples, we include plots for all individual classes in Figure 8 for CIFAR-10 and in Figures 9 and 10 for ImageNet-1K respectively. Lastly, for the experiments on varying the diversity of OE data on MNIST with EMNIST-Letters OE, and CIFAR-10 with CIFAR-100 OE, we added the plots for all individual classes as well in Figures 11 and 12.

Complete code also available at: <https://github.com/lukasruff/Classification-AD>

Table 4. Mean AUC detection performance in % (over 10 seeds) for all individual classes and methods on the CIFAR-10 one vs. rest benchmark with 80MTI OE from Section 3.1.

Class	SVDD*	Unsupervised				Unsupervised OE			Supervised OE		
		DSVDD*	GT*	IT*	GT+*	GT+*	DSAD	HSC	Focal*	Focal	BCE
Airplane	65.6	61.7	74.7	78.5	77.5	90.4	94.2	96.3	87.6	95.9	<b>96.4</b>
Automobile	40.9	65.9	95.7	89.8	96.9	<b>99.3</b>	98.1	98.7	93.9	98.7	98.8
Bird	65.3	50.8	78.1	86.1	87.3	<b>93.7</b>	89.8	92.7	78.6	92.3	93.0
Cat	50.1	59.1	72.4	77.4	80.9	88.1	87.4	89.8	79.9	88.8	<b>90.0</b>
Deer	75.2	60.9	87.8	90.5	92.7	<b>97.4</b>	95.0	96.6	81.7	96.6	97.1
Dog	51.2	65.7	87.8	84.5	90.2	<b>94.3</b>	93.0	94.2	85.6	94.1	94.2
Frog	71.8	67.7	83.4	89.2	90.9	97.1	96.9	97.9	93.3	97.8	<b>98.0</b>
Horse	51.2	67.3	95.5	92.9	96.5	<b>98.8</b>	96.8	97.6	87.9	97.6	97.6
Ship	67.9	75.9	93.3	92.0	95.2	<b>98.7</b>	97.1	98.2	92.6	98.0	98.1
Truck	48.5	73.1	91.3	85.5	93.3	<b>98.5</b>	96.2	97.4	92.1	97.5	97.7
Mean AUC	58.8	64.8	86.0	86.6	90.1	95.6	94.5	95.9	87.3	95.8	<b>96.1</b>

Table 5. Mean AUC detection performance in % (over 10 seeds) with standard deviations for all individual classes for our implementations on the CIFAR-10 one vs. rest benchmark with 80MTI OE from Section 3.1.


Class	Unsupervised OE		Supervised OE	
	DSAD	HSC	Focal	BCE
Airplane	94.2 ± 0.34	96.3 ± 0.13	95.9 ± 0.11	96.4 ± 0.17
Automobile	98.1 ± 0.19	98.7 ± 0.07	98.7 ± 0.09	98.8 ± 0.06
Bird	89.8 ± 0.54	92.7 ± 0.27	92.3 ± 0.32	93.0 ± 0.14
Cat	87.4 ± 0.38	89.8 ± 0.27	88.8 ± 0.33	90.0 ± 0.27
Deer	95.0 ± 0.22	96.6 ± 0.17	96.6 ± 0.10	97.1 ± 0.10
Dog	93.0 ± 0.30	94.2 ± 0.13	94.1 ± 0.21	94.2 ± 0.12
Frog	96.9 ± 0.22	97.9 ± 0.08	97.8 ± 0.07	98.0 ± 0.09
Horse	96.8 ± 0.15	97.6 ± 0.10	97.6 ± 0.16	97.6 ± 0.09
Ship	97.1 ± 0.21	98.2 ± 0.09	98.0 ± 0.11	98.1 ± 0.08
Truck	96.2 ± 0.22	97.4 ± 0.13	97.5 ± 0.12	97.7 ± 0.16
Mean AUC	94.5 ± 3.30	95.9 ± 2.68	95.8 ± 2.97	96.1 ± 2.71

Table 6. Mean AUC detection performance in % (over 10 seeds) for all individual classes for our implementations of the ImageNet-1K one vs. rest benchmark with ImageNet-22K OE from Section 3.2. Note that for GT+\* and Focal\*, as reported in Table 2 in the main paper, Hendrycks et al. (2019b) do not provide results on a per class basis.


Class	Unsupervised OE		Supervised OE	
	DSAD	HSC	Focal	BCE
acorn	98.5 ± 0.28	98.8 ± 0.42	99.0 ± 0.15	99.0 ± 0.19
airliner	99.6 ± 0.16	99.8 ± 0.10	99.9 ± 0.02	99.8 ± 0.04
ambulance	99.0 ± 0.13	99.8 ± 0.13	99.2 ± 0.14	99.9 ± 0.07
american alligator	92.9 ± 1.06	98.0 ± 0.32	94.7 ± 0.67	98.2 ± 0.27
banjo	97.0 ± 0.51	98.2 ± 0.41	97.0 ± 0.33	98.7 ± 0.22
barn	98.5 ± 0.29	99.8 ± 0.05	98.7 ± 0.24	99.8 ± 0.08
bikini	96.5 ± 0.84	98.6 ± 0.57	97.2 ± 0.89	99.1 ± 0.30
digital clock	99.4 ± 0.33	96.8 ± 0.79	99.8 ± 0.03	97.2 ± 0.29
dragonfly	98.8 ± 0.28	98.4 ± 0.16	99.1 ± 0.21	98.3 ± 0.04
dumbbell	93.0 ± 0.53	91.6 ± 0.88	94.0 ± 0.04	92.6 ± 0.97
forklift	90.6 ± 1.43	99.1 ± 0.33	94.2 ± 0.90	99.5 ± 0.09
goblet	92.4 ± 1.05	93.8 ± 0.38	93.8 ± 0.27	94.7 ± 1.43
grand piano	99.7 ± 0.06	97.4 ± 0.37	99.9 ± 0.04	97.6 ± 0.34
hotdog	95.9 ± 2.01	98.5 ± 0.34	97.2 ± 0.05	98.8 ± 0.34
hourglass	96.3 ± 0.37	96.9 ± 0.26	97.5 ± 0.17	97.6 ± 0.48
manhole cover	98.5 ± 0.29	99.6 ± 0.34	99.2 ± 0.09	99.8 ± 0.01
mosque	98.6 ± 0.29	99.1 ± 0.26	98.9 ± 0.30	99.3 ± 0.15
nail	92.8 ± 0.80	94.0 ± 0.76	93.5 ± 0.32	94.5 ± 1.37
parking meter	98.5 ± 0.29	93.3 ± 1.64	99.3 ± 0.04	94.7 ± 0.76
pillow	99.3 ± 0.14	94.0 ± 0.47	99.2 ± 0.14	94.2 ± 0.42
revolver	98.2 ± 0.30	97.6 ± 0.25	98.6 ± 0.11	97.7 ± 0.68
rotary dial telephone	90.4 ± 1.99	97.7 ± 0.50	92.2 ± 0.33	98.3 ± 0.75
schooner	99.1 ± 0.18	99.2 ± 0.20	99.6 ± 0.02	99.1 ± 0.26
snowmobile	97.7 ± 0.86	99.0 ± 0.22	98.1 ± 0.15	99.1 ± 0.25
soccer ball	97.3 ± 1.70	92.9 ± 1.18	98.6 ± 0.13	93.6 ± 0.61
stingray	99.3 ± 0.20	99.1 ± 0.33	99.7 ± 0.04	99.2 ± 0.10
strawberry	97.7 ± 0.64	99.1 ± 0.20	99.1 ± 0.03	99.2 ± 0.22
tank	97.3 ± 0.51	98.6 ± 0.18	97.3 ± 0.47	98.9 ± 0.13
toaster	97.7 ± 0.56	92.2 ± 0.78	98.3 ± 0.05	92.2 ± 0.65
volcano	89.6 ± 0.44	99.5 ± 0.09	91.6 ± 0.90	99.4 ± 0.19
Mean AUC	96.7 ± 2.98	97.3 ± 2.53	97.5 ± 2.43	97.7 ± 2.34

## Rethinking Assumptions in Deep Anomaly Detection


Figure 8. Mean AUC detection performance in % (over 10 seeds) for all classes of the CIFAR-10 one vs. rest benchmark from Section 3.1 when varying the number of 80MTI OE samples. These plots correspond to Figure 2(a), but here we report the results for all individual classes.




(a) Class: airplane




(b) Class: automobile




(c) Class: bird




(d) Class: cat




(e) Class: deer




(f) Class: dog




(g) Class: frog



(h) Class: horse



(i) Class: ship



(j) Class: truck

Figure 9. Mean AUC detection performance in % (over 5 seeds) for all classes of the ImageNet-1K one vs. rest benchmark from Section 3.2 when varying the number of ImageNet-22K OE samples. These plots correspond to Figure 2(b), but here we report the results for all individual classes (from class 1 (acorn) to class 15 (hourglass)).




Figure 10. Mean AUC detection performance in % (over 5 seeds) for all classes of the ImageNet-1K one vs. rest benchmark from Section 3.2 when varying the number of ImageNet-22K OE samples. These plots correspond to Figure 2(b), but here we report the results for all individual classes (from class 16 (manhole cover) to class 30 (volcano)).




Figure 11. Mean AUC detection performance in % (over 10 seeds) for all MNIST classes from the experiment in Appendix B on varying the number of classes  $k$  of the EMNIST-Letters OE data. These plots correspond to Figure 4(a), but here we report the results for all individual classes.

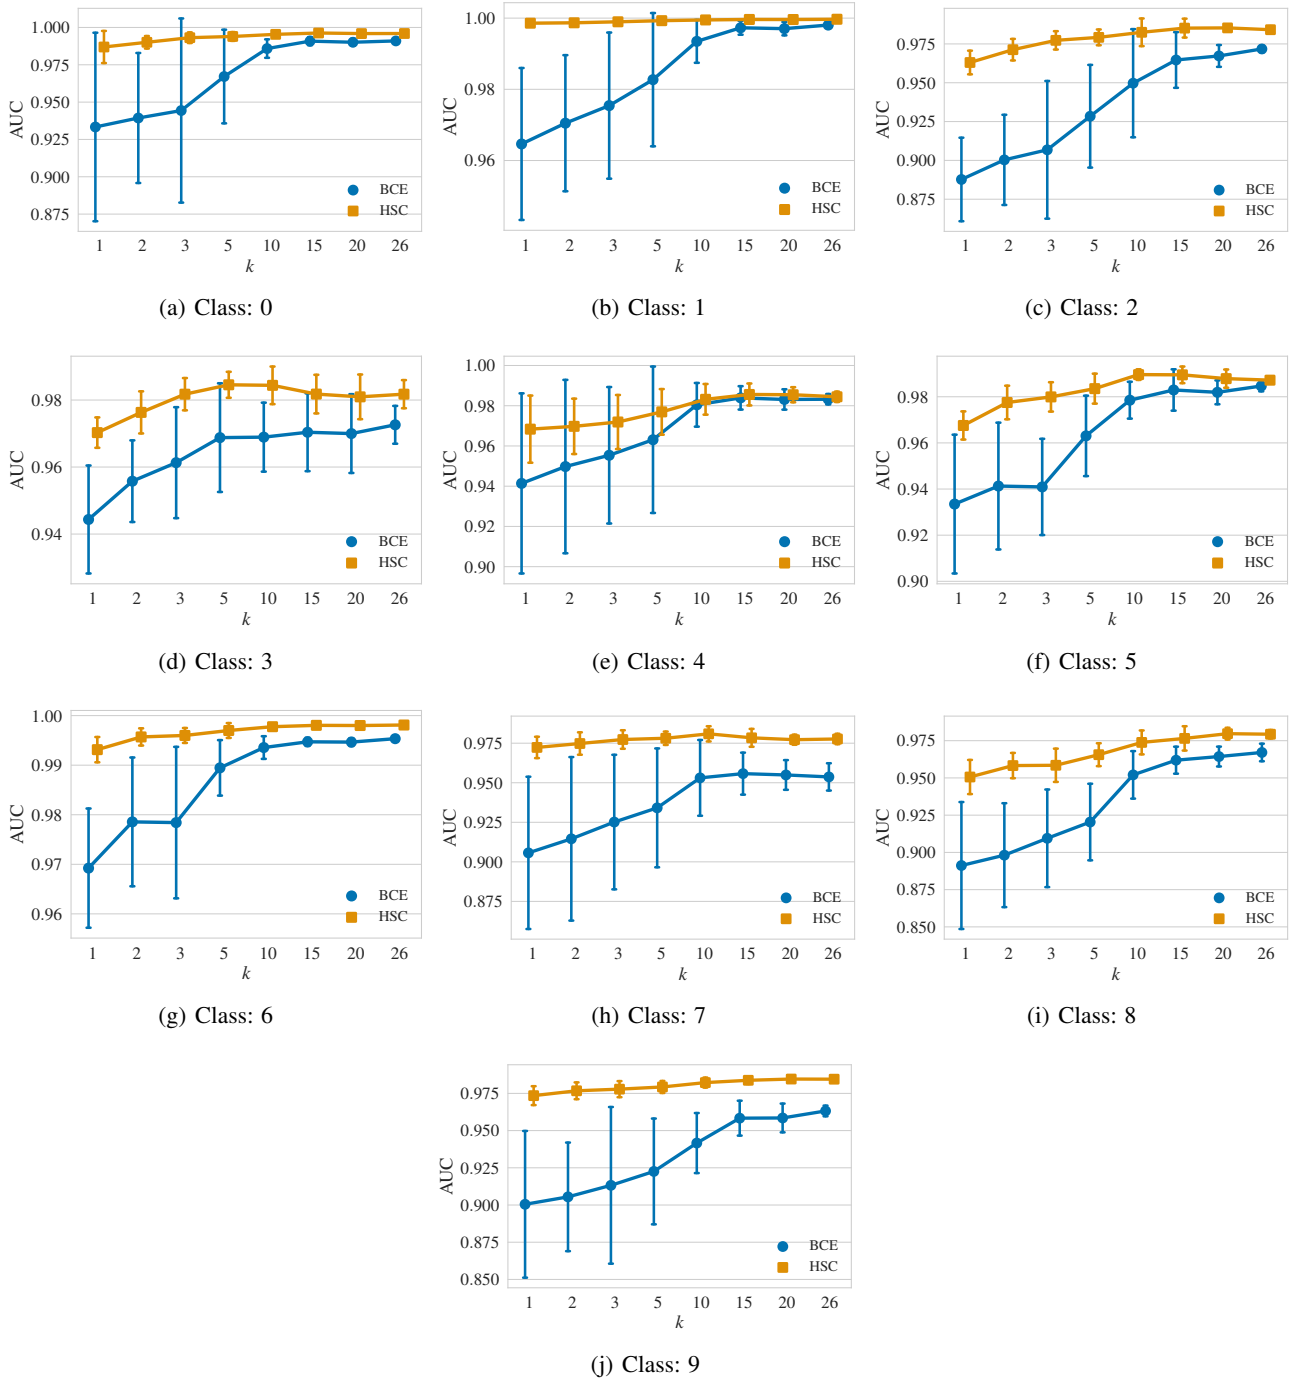







Figure 12. Mean AUC detection performance in % (over 10 seeds) for all CIFAR-10 classes from the experiment in Appendix B on varying the number of classes  $k$  of the CIFAR-100 OE data. These plots correspond to Figure 4(b), but here we report the results for all individual classes.




(a) Class: airplane




(b) Class: automobile




(c) Class: bird




(d) Class: cat




(e) Class: deer




(f) Class: dog




(g) Class: frog



(h) Class: horse



(i) Class: ship



(j) Class: truck