

# A Less Biased Evaluation of Out-of-distribution Sample Detectors

Alireza Shafaei  
<http://www.shafaei.ca>

Mark Schmidt  
<http://cs.ubc.ca/~schmidtm>

James J. Little  
<http://cs.ubc.ca/~little>

Computer Science Department  
University of British Columbia  
Vancouver, Canada

---

## Abstract

In the real world, a learning system could receive an input that is unlike anything it has seen during training. Unfortunately, out-of-distribution samples can lead to unpredictable behaviour. We need to know whether any given input belongs to the population distribution of the training/evaluation data to prevent unpredictable behaviour in deployed systems. A recent surge of interest in this problem has led to the development of sophisticated techniques in the deep learning literature. However, due to the absence of a standard problem definition or an exhaustive evaluation, it is not evident if we can rely on these methods. What makes this problem different from a typical supervised learning setting is that the distribution of outliers used in training may not be the same as the distribution of outliers encountered in the application. Classical approaches that learn inliers vs. outliers with only two datasets can yield optimistic results. We introduce OD-test, a three-dataset evaluation scheme as a more reliable strategy to assess progress on this problem. We present an exhaustive evaluation of a broad set of methods from related areas on image classification tasks. Contrary to the existing results, we show that for realistic applications of high-dimensional images the previous techniques have low accuracy and are not reliable in practice.

## 1 Introduction

If we present a natural image of an unknown class to the currently popular deep neural network that is trained to discriminate ImageNet [50] classes, we will get a prediction with a high (softmax) probability of an arbitrary class (see Fig. 1). With English speaking phone assistants, if we talk in another language, it will generate an English sentence that most often is not even remotely similar to what we have said. The silent failure of these systems is due to an implicit assumption: the input to the ImageNet classifier *will be* from the same ImageNet distribution, and the user *will be* speaking in English. However, in practice, any automation pipeline that involves a deep neural network will have a critical challenge:

*Can we trust the output of a neural network for a particular input?*



Figure 1: The predictions of several popular networks [12, 15, 16, 21, 55] that are trained on ImageNet on unseen data. The red predictions are entirely wrong, the green predictions are justifiable, and the orange predictions are less justifiable. The middle image is Gaussian noise. We show that thresholding the output probability is not a reliable defence.

One solution is to add a `None` class to the models to account for the absence of other classes. The first challenge is defining the `None` class. Would `None` mean all other image vectors or only vectors of natural images? Another non-trivial challenge is capturing the diversity of `None` class with a finite sample set to use for training. The third and the most prohibitive problem is that we have to significantly increase the complexity of our models to capture the diversity of the `None` class. Despite these challenges, we might be able to achieve reasonable results on low dimensional problems [14], but as the input dimension grows, the problem becomes more severe. Our argument rests on the assumption that learning a reliable `None` class is not a viable strategy for practical applications.

In supervised learning we typically assume the samples are independently and identically distributed (IID) – we expect the train and test examples to be drawn from a fixed population distribution. However, this condition cannot be easily enforced in deployed applications. When the IID assumption is not satisfied, the empirical error can no longer predict the performance. Without any assumptions, the outputs on out-of-distribution (OOD) samples can be arbitrarily bad. The OOD samples violate the *identically distributed* (ID) assumption. Thus, as long as we rely on the empirical error alone to train and evaluate deep neural networks, the first condition for making a reliable prediction (with bounded error) on any input would be whether the ID assumption is satisfied. When we *a priori* expect to change the underlying distribution, the traditional applicable frameworks are transfer learning, multitask learning, and zero-shot learning. In our setting, we only wish to *detect* out-of-distribution samples.

In real life deployment of products that use complex machinery such as deep neural networks (DNNs), we would have very little control over the input. In the absence of extrapolation guarantees, when the IID assumption is violated, the behaviour of the pipeline may be unpredictable. A reliable pipeline would first determine whether it can process a given sample, then it would use the prediction of the target neural network. Successful detection of such violations could also be used in active learning, unsupervised learning, learning with noisy data, or simply be a condition to invoking transfer learning strategies.

There has been a recent surge of interest in specifically OOD sample detection with deep neural networks [2, 13, 14, 22, 26, 52]. However, the commonly applied evaluation mechanisms are susceptible to overly optimistic results and cannot provide a conclusive evidence of reliability (we demonstrate this). There is an imminent need to define a standard problem and properly define a benchmark that is both realistic and principled to compare the previous and future approaches reliably.

Our contributions are as follows:

1. We introduce OD-test: a more realistic formulation of the OOD detection task.
2. We present a new benchmark to evaluate the existing techniques for OOD detection.
3. We provide analysis of several previously proposed methods under OD-test.
4. We release a PyTorch [28] package to replicate all the results (<https://github.com/ashafaei/OD-test>).

We demonstrate that the performance of current techniques quickly approaches the random prediction baseline as we make a transition to realistic high-dimensional images.

## 2 Related Work

The violation of the ID assumption is not the only way to wreak havoc on deep learning pipelines. Adversarial example [66] attacks are crafted signals in otherwise innocent-looking images that fool the neural networks into misclassification. While the OOD detection is a model-independent problem, adversarial images exploit inductive bias in the model families. We limit our attention to the OOD detection problem.

**The Uncertainty View** An uncertainty measure could be directly applied to reject OOD samples as we would expect the uncertainty to be high on such inputs. The MC-Dropout [7] approach is a feasible uncertainty estimation method for a variety of applications [4, 8, 10]. Lakshminarayanan et al. [22] show an ensemble of five neural networks (DeepEnsemble) trained with an adversarial-sample-augmented loss is sufficient to provide a measure of predictive uncertainty. We evaluate DeepEnsemble and MC-Dropout.

**The Abstention View** If we abstain from prediction at the cost of a penalty, we end up with the abstention view. A reject function makes the abstention choice and it could be a threshold on magnitude of the prediction [11] or chosen from a reject-hypothesis set [4, 5]. The reject function is chosen simultaneously with the predictive function through an abstention-augmented loss. The reliability of such reject function is contingent on evaluation on a fixed distribution. If we encounter an OOD sample, we do not know for sure if it would be rejected. Our formulation of the OOD detection task is similar to the abstention view with key differences that we will discuss in Section 3. We show that the prior work on OOD detection can be reduced to an abstract problem of choosing a reject function from a specific space.

**The Anomaly View** *Density estimation*-based techniques assume that low measure samples are outliers. These approaches tend to work well mostly within low-dimensional or well-defined distributions. PixelCNN++ [31] is an auto-regressive model with a tractable likelihood that could be used within a density estimation scheme. Note that density estimation is *not equivalent* to the binary OOD detection: a perfect density estimator can solve the outlier detection problem, but a perfect outlier detector does not necessarily have the information needed to solve the density estimation problem. *Proximity*-based methods use a distance measure and the train data to flag anomalies. A simple strategy is based on the K-nearest neighbours of a given input – we call this K-NNSVM. Clustering methods reject points that do not conform to any of the identified clusters. The one-class SVM [33] with a radial basis function learns a conservative region that covers the train data. Goldstein and Uchida [9] show that the proximity-based approaches are empirically the most effective outlier detectors over a range of datasets. *Reconstruction*-based methods learn to reconstruct the train data, then try to reconstruct each given input. The samples that cannot be reconstructed well are then flagged as anomalies. We use an autoencoder with a reconstruction threshold to test this idea (AETHreshold).

**The Novelty View** Open-set recognition and novelty detection study the detection of anomalies at a *semantic* level. These methods are typically concerned with recognition of unseen classes, *e.g.*, new objects in the scene. This is a special case of OOD detection where the OOD samples explicitly differ by the semantic content. However, the notion of OOD is more granular: an unseen viewpoint of a specific object violates the ID assumption, but it does not necessarily constitute a novelty. The notion of novelty is often underspecified in practice and results are limited to particular assumptions and problem definitions. Bendale and Boulton [2] present OpenMax, a replacement for the softmax layer that detects unknown classes through evaluation against a representative neural activation of each class.

**Deep Learning Literature** Guo et al. [10] observed that neural networks tend to be overconfident in predictions. They show that temperature scaling is an effective calibration strategy for neural networks. Hendrycks and Gimpel [13] show that it is possible to detect OOD samples by thresholding the softmax probabilities. More recently, Liang et al. [26] presented ODIN, a method based on temperature rescaling and input perturbation to detect OOD samples. Further extensions rely on statistics of hidden representations [25, 29], construct classifier ensembles with subsets of data [37], or perform regression on word embeddings [32].

There are also ideas that rely on GANs [6, 19, 24, 32] to detect anomalies or novelty in the data. To the best of our knowledge, training GANs [10] is an active area of research, and it is not apparent what design decisions would be appropriate to implement these ideas in practice. We are therefore unable to evaluate these ideas fairly at this time.

All the previous studies primarily focus on low-dimensional MNIST [23], SVHN [27], and CIFAR [20] datasets. We evaluate several previously proposed solutions in controlled experiments on datasets with varying complexity. We show that, in such low-dimensional spaces, simple anomaly detection methods work as well, thus stressing that a more comprehensive evaluation is necessary for assessments of the current and future work.

### 3 OD-test: A Less Biased Evaluation of Outlier Detectors

Let us define the source distribution  $\mathcal{D}_s$  to be our input distribution. The objective is to decide whether a given sample belongs to  $\mathcal{D}_s$ . We define a reject function  $r : \mathcal{X} \rightarrow \{0, 1\}$  that makes this binary decision. Note that this decision can be made independently from the ultimate prediction task. While in the abstention view we reject the samples that the predictive function is likely to mislabel, here we reject the samples that do not belong to the source distribution  $\mathcal{D}_s$ , hence decoupling the reject function and the predictive function.

If the reject function flags an input, then the sample does not belong to the source distribution; thus, the output of the pipeline may not be reliable. On the other hand, if the function accepts an input, we can continue the pipeline with the ID assumption. This form of rejection is more general than the previous work. In addition to the previous methods, we can study new approaches that operate in the input space directly (*e.g.* K-NNSVM, AThreshold).

The  $r$  function is a binary classifier; the classes are in-distribution vs. out-of-distribution. To learn the classifier, the standard approach is to adopt the supervised learning assumptions and use an outlier dataset  $\mathcal{D}_v$ . In that scenario, what would happen if an outlier that we encounter is not represented by  $\mathcal{D}_v$ ? We end up with the original problem again. The traditional supervised outlier detection may overestimate our ability to detect outliers. A high accuracy in this scenario may not yield an accurate model for practice in many settings where the outlier may not look like samples from  $\mathcal{D}_v$ . The actual OOD samples are beyond

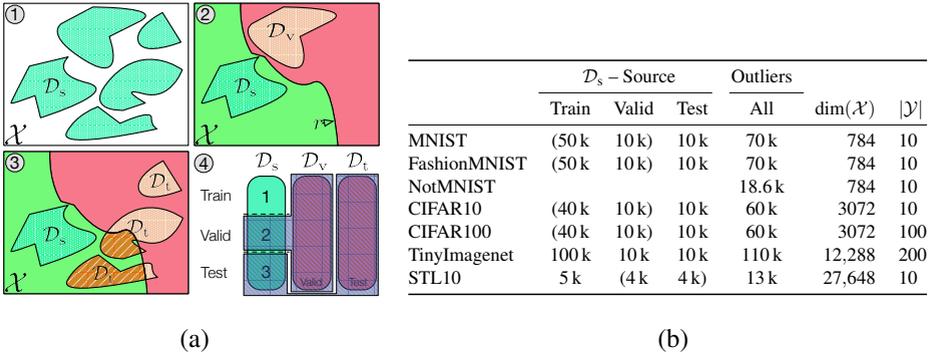


Figure 2: The OD-test. a.1) A set of distributions visualized as shapes within  $\mathcal{X}$ . The source distribution  $\mathcal{D}_s$  is identified in the image; everything else is an outlier. a.2) We pick one validation distribution  $\mathcal{D}_v$  and learn a binary reject function  $r$  that partitions the input space  $\mathcal{X}$  based on  $\mathcal{D}_s$  and  $\mathcal{D}_v$  only. a.3) We evaluate  $r$  on other distributions (as  $\mathcal{D}_t$ ) and measure the accuracy. a.4) The dataset splits for each step. b) A summary of the datasets. The datasets are split as indicated by the parentheses.

our direct reach and our models can easily overfit in distinguishing  $\mathcal{D}_s$  from  $\mathcal{D}_v$  (we verify this empirically, see Fig. 3). We present a less optimistic evaluation framework that prevents scoring high through overfitting.

Specifically, we introduce a third “target” distribution  $\mathcal{D}_t$  to measure whether a method can actually detect outliers that are not only outside of  $\mathcal{D}_s$  but that also might be outside of  $\mathcal{D}_v$ . The idea is to treat the problem as a binary classification between three different datasets. Similar to the supervised outlier detection, we begin by learning a reject function to distinguish  $\mathcal{D}_s$  from  $\mathcal{D}_v$ . For evaluation, we use a third unseen distribution  $\mathcal{D}_t$  instead of  $\mathcal{D}_v$  (see Fig. 1(a)).  $\mathcal{D}_t$  represents OOD examples that were not encountered during training – a more realistic evaluation setting for uncontrolled scenarios. Ideally,  $\mathcal{D}_v$  and  $\mathcal{D}_t$  should have no similarity. In practice, we recommend choosing a collection of datasets with no label overlap. In our experiments, we cycle through all choices of outliers ( $\mathcal{D}_v, \mathcal{D}_t$  from Tab. 1(b)) and average the results. The pseudocode of the evaluation procedure is outlined in Alg. 1.

### Algorithm 1: OD-test

```

input :  $\mathcal{D}_s = (\mathcal{D}_s^{\text{train}}, \mathcal{D}_s^{\text{valid}}, \mathcal{D}_s^{\text{test}})$  the source dataset.
input :  $\mathcal{D} = \{\mathcal{D}_i\}$  outlier set.
input :  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  the method under evaluation.
1 begin
2    $A \leftarrow \{\}$ 
3   /* Generate a reject-hypothesis class  $\mathcal{R}$  using  $\mathcal{D}_s^{\text{train}}$ . */
4    $\mathcal{R} \leftarrow \mathcal{M}(\mathcal{D}_s^{\text{train}})$ 
5   for  $\mathcal{D}_v \in \mathcal{D}$  do
6     /* Find the best binary classifier in  $\mathcal{R}$ . */
7      $r \leftarrow \text{train}(\mathcal{R}, \{\mathcal{D}_s^{\text{valid}} : 0, \mathcal{D}_v : 1\})$ 
8     for  $\mathcal{D}_t \in \mathcal{D} \setminus \{\mathcal{D}_v\}$  do
9       /* Evaluate accuracy of  $r$ . */
10       $\text{acc} \leftarrow \text{eval}(r, \{\mathcal{D}_s^{\text{test}} : 0, \mathcal{D}_t : 1\})$ 
11      add acc to  $A$ 
12 return  $\text{mean}(A)$ 

```

As an example, let us walk through the stages of evaluation for PbThreshold, the

method that thresholds the maximum probability of a discriminative neural network. We need a trained deep neural network and a threshold for the maximum probability that would reject the outliers. We first train a deep neural network on  $\mathcal{D}_s$  to discriminate the image classes (line 3). The reject class  $\mathcal{R}$  returned on line 3 will have a single free parameter  $\tau$ , the threshold to use on the underlying classifier. We pick the optimal threshold  $\tau$  in the next step. On line 5, we pick the best threshold  $\tau$  to discriminate between  $\mathcal{D}_s$  and  $\mathcal{D}_v$ . After finding the best threshold  $\tau$  on line 7, we evaluate the learned reject function on  $\mathcal{D}_s$  and  $\mathcal{D}_t$ . For unsupervised OOD detection methods the evaluation is a single loop over  $\mathcal{D}_t$  (see supplementary material).

Note that to score high on OD-test we *do not need to perform density estimation*. However, if we can perform density estimation well, we can score high on OD-test too. If the method  $\mathcal{M}$  successfully learns a density function for  $\mathcal{D}_s$ , we only would have to pick a single threshold to reject OOD samples. Methods that yield a confidence (or uncertainty) estimate can be used similarly. A binary classifier trained with the traditional supervised learning approaches is also not sufficient to score high on this benchmark since we change the second distribution during the test stage. In Section 4 we show how a traditional binary classifier would fall short. The methods that learn conservative boundaries around  $\mathcal{D}_s$  will have a higher chance of success. All the existing approaches can be implemented within this framework.

We equalize the binary classes and only measure accuracy. Furthermore, we require the *methods* to pick the optimal parameters such as the threshold. These choices simplify aggregation, analysis, and comparison of the results. We can meaningfully average over multiple experiments and robustly compare methods in a variety of conditions. Similar to the supervised learning regime, we can incorporate the prior knowledge on abundance and the risk associated with the OOD samples into the evaluation by modifying the implicit objective on line 5 and 7. We leave the proper choice of application dependent schemes to the practitioner and focus on assessing the discriminative power of the methods in the fixed 50/50 scenario *without* any prior knowledge. Through this constraint, we ensure that the methods that rely on the prior likelihood cannot perform better than random prediction.

We extend the previous work by evaluating over a broader set of datasets with varying levels of complexity. The variation in complexity allows for a fine-grained evaluation of the techniques. Since OOD detection is closely related to the problem of density estimation, the dimensionality of the input image will be of vital importance in practical assessments. As the input dimensionality increases, we expect the task to become much more difficult. Therefore, to provide a more accurate picture of performance, it is crucial to evaluate the methods on high dimensional data. Table 1(b) summarizes the datasets that we use. We also evaluate with uniform, and Gaussian noise for outliers. We evaluate the following methods:

- **BinClass**. A traditional binary classifier that is directly trained on  $\mathcal{D}_s$  vs.  $\mathcal{D}_v$ .
- **PbThreshold**. A threshold on the softmax.
- **ScoreSVM**. An SVM [9] classifier on the logits.
- **LogisticSVM**. Similar to **ScoreSVM**, but the underlying classifier is trained with logistic loss.
- **ODIN**. A threshold on the scaled softmax outputs of the perturbed input.
- **K-NNSVM**. A linear SVM on the sorted Euclidean distance between the input and the k-nearest training samples. Note that a threshold on the average distance is a special case of **K-NNSVM**.
- **AEThreshold**. A threshold on the autoencoder (AE) reconstruction error of the given input. We train AEs with binary cross-entropy (BCE) and mean squared error (MSE).

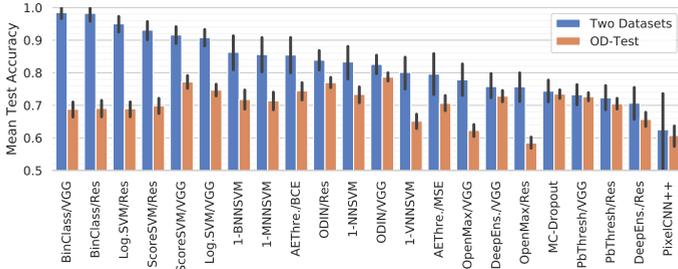


Figure 3: Evaluation with two datasets versus OD-test. Evaluating OOD detectors with only two distributions can be misleading for practical applications. The error bars are the 95% confidence level. The two-dataset evaluations are over all possible pairs of datasets ( $n = 46$ ), whereas the OD-test evaluations are over all possible triplets ( $n = 308$ ).

- **K-MNNSVM, K-BNNSVM.** A  $K$ -NNSVM applied on the learned hidden representations of an autoencoder with MSE or BCE.
- **K-VNNSVM.** Similar to the previous, except we use the learned representation of a variational autoencoder [18].
- **MC-Dropout.** A threshold on the entropy of average prediction of 7 evaluations per input.
- **DeepEnsemble.** Similar to MC-Dropout, except we average over the predictions of 5 networks that are trained independently with adversarial-augmented loss.
- **PixelCNN++.** A threshold on the log-likelihood of each input.
- **OpenMax.** Similar to ScoreSVM, but we use the calibrated output of the OpenMax module that also includes a probability for an unknown class.

We use two generic architectures: VGG-16 [65] and Resnet-50 [12] and reuse the same base classifier for all the methods to provide a fair and meaningful comparison. We train these architectures with cross-entropy loss (CE), and  $k$ -way logistic loss (KWL). CE loss enforces mutual exclusion in the predictions while KWL loss does not. We test these two loss functions to see if the exclusivity assumption of CE hurts the ability to predict OOD samples. CE loss cannot make a `None` prediction without an explicitly defined `None` class, but KWL loss can make `None` predictions through low activations of all the classes.

Note that our formulation of the problem separates the target task from the OOD sample detection. Thus, it is plausible to use, for OOD detection, a different predictive model from the actual predictive model of the target problem if there is an advantage. We tune the hyper-parameters of these methods following the best practices and the published guidelines in the respective articles. The implementation details, a discussion of evaluation cost, and the performance statistics of the above methods are in the supplementary material. The PyTorch implementation with the pre-trained models and all the numerical results is available on <https://github.com/ashafaei/OD-test>.

## 4 Results and Discussion

We evaluate the methods in a controlled regime against the datasets in Tab. 1(b). We run over 10,000 experiments on all combinations of  $\mathcal{D}_s$ ,  $\mathcal{D}_v$ , and  $\mathcal{D}_t$ . First we analyze the aggregated

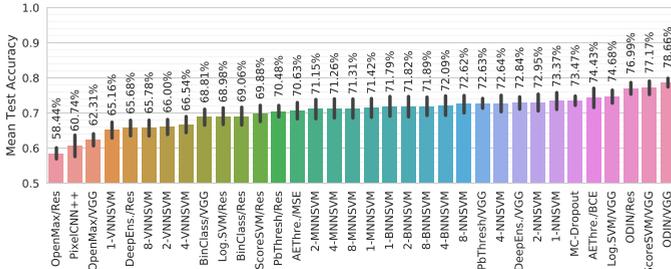


Figure 4: The average test accuracy of the OOD detection methods over 308 experiments/method with 95% confidence level. /VGG or /Res indicates the backing network architecture. #-NN./ is the number of nearest neighbours. A random prediction would have an accuracy of 0.5.

results, then we look at the breakdown of the accuracy per source dataset. Figure 3 and Figure 4 show the average accuracy. Each method under OD-test is tested over the same set of 308 experiments consisting of *all compatible triplet of datasets*. The two-dataset evaluations are averaged over 46 experiments (all possible pairs). See the project page for the list of experiments.

Figure 3 compares the mean test accuracy of methods within OD-test and the two-dataset setting. Methods that perform well in distinguishing two datasets fail when a third dataset is introduced. The gap in relative performance within each evaluation highlights the importance of having a more realistic assessment in practice. The general trend of the evaluation indicates that the methods that have a higher degree of freedom for OOD detection (the space characterized by  $\mathcal{R}$ ) are the most susceptible ones to overestimating accuracy within the traditional evaluation approach. This observation is consistent with the bias-variance tradeoff in learning theory.

BinClass, the direct binary classifier, consistently achieves a near-perfect accuracy on the train and validation of the same datasets (see Fig. 3), however, once tested on a third dataset that it has not seen before, the average accuracy drops to 68% for both VGG and Resnet. This demonstrates why we need to adopt a new evaluation scheme. The classifier overfits to the two distributions ( $\mathcal{D}_s$  and  $\mathcal{D}_v$ ), but it cannot distinguish a third distribution ( $\mathcal{D}_t$ ). Because of the diversity of OOD samples, we may always encounter new input that we have not seen before.

MC-Dropout and DeepEnsemble, the two uncertainty techniques, do not seem to provide a strong enough signal to distinguish the two classes compared to the simpler ScoreSVM. Interestingly, MC-Dropout has a higher accuracy than DeepEnsemble. Considering the training cost of DeepEnsemble, MC-Dropout is a more favourable choice.

VGG-backed and Resnet-backed methods significantly differ in accuracy. The gap indicates the sensitivity of the methods to the underlying networks. PbThreshold, ScoreSVM, and ODIN all prefer VGG over Resnet even though Resnet networks outperform the VGG variants in image classification. This means that the image classification accuracy may not be the only relevant factor in performance of these methods. ODIN is less sensitive to the underlying network. Furthermore, training the networks with KWL loss consistently reduces the accuracy of OOD detection methods on average. ScoreSVM/VGG and ScoreSVM/Res both outperform LogisticSVM/VGG, and LogisticSVM/Res respectively. Similarly, the autoencoders that were trained with BCE loss (AEThre./BCE) outperform the ones

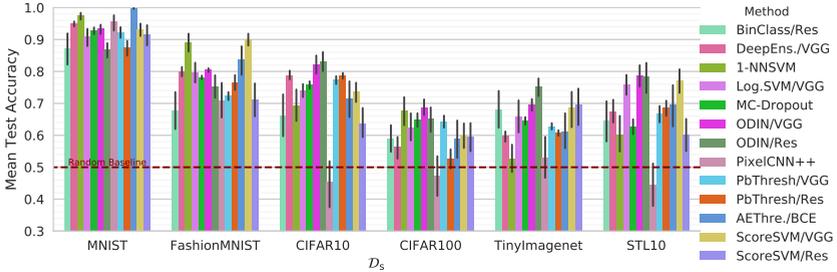


Figure 5: The test accuracy over 50 experiments/bar with 95% confidence level.

trained with MSE loss (AEThresh./MSE). Note that we are comparing identical architectures.

Within the nearest-neighbour methods,  $\#-(X)$  NNSVM, the number of the nearest neighbours does not significantly impact the accuracy on average. However, performing the nearest-neighbour in the input space directly outperforms nearest-neighbour in the latent representations of autoencoders (BNNSVM, and MNNSVM) and VAE (VNNSVM). Interestingly, 1-NNSVM has a higher accuracy than thresholding the probability (PbThresh) and DeepEnsemble on average within OD-test *and* two-dataset evaluation (Fig. 3). For  $\#-$ NNSVM, if the reference samples fit the GPU memory, a naive implementation could be faster than a forward pass on the neural networks of large datasets like TinyImagenet.

PixelCNN++, the method that estimates the log-likelihood, has a surprisingly low accuracy on this problem on average. We suspect the auto-regressive nature of the model, specifically when coupled with the IID assumption, may be the reason for its failure. The network approximates the likelihood only in the vicinity of the training data. If we evaluate the model on points that are far from the training data, the estimates are not reliable anymore.

Figure 5 shows the average test accuracy across each source dataset  $\mathcal{D}_s$ . For the full figure, see supplementary material. Our quantification of performance shows that all the methods have a much lower accuracy on high-dimensional data than the low-dimensional data.

In low-dimensional datasets,  $K$ -NNSVM performs similarly or better than the other methods. In the high-dimensional case, however, the accuracy approaches the random baseline quickly. Interestingly,  $K$ -NNSVM performs better on STL10 ( $96 \times 96$ ) than TinyImagenet ( $64 \times 64$ ) which might be due to the higher diversity in TinyImagenet compared to STL10. Given the high accuracy of  $K$ -NNSVM in 784 dimensions, it might be feasible to learn an embedding for high-dimensional data, or learn a kernel function, to replace the original image space and enjoy a high accuracy. Going from CIFAR10 to CIFAR100, the dimensionality and the dataset size remains the same, the only changing factor is the diversity of the data and that seems to make the problem as difficult as the higher dimensional datasets. Except for  $K$ -NNSVM, the accuracy of other methods drops significantly in this transition.

AEThreshold has a near perfect accuracy on MNIST, however, the performance drops quickly on complex datasets. AEThreshold also outperforms the density estimation method PixelCNN++ on all datasets. We did not explore autoencoder architectures – more research on better architectures or reconstruction constraints for AEThreshold may potentially have a high pay-off. The top-performing method, ODIN, is influenced by the number of classes in the dataset. Similar to PbThreshold, ODIN depends on the maximum signal in the class predictions, therefore the increased number of classes would directly affect both of the methods. Furthermore, neither of them consistently prefers VGG over Resnet within all datasets. Overall, ODIN consistently outperforms others in high-dimensional settings, but all

the methods have a relatively low average accuracy in the 60%-78% range.

Overall, we can summarize the main observations as follows: (i) Outlier detection with two datasets yielded overly optimistic results with VGG and Resnet. (ii) The MC-Dropout and DeepEnsemble uncertainty methods were not reliable enough for OOD detection. (iii) A more accurate image classifier did not lead to a more accurate outlier detector on average. (iv) The nearest neighbour methods were competitive in low-dimensional settings, both in computational cost and accuracy. (v) The latent representations of vanilla (variational)-autoencoders were not useful for this task when combined with nearest neighbour methods. (vi) The state-of-the-art auto-regressive density estimation method had a surprisingly low accuracy, performing worse than the random prediction baseline in some settings. (vii) Although ODIN outperforms other methods in realistic high-dimensional settings, its average accuracy is still below 80%.

To perform supervised OOD sample detection in practice, we have to pick a method and choose a training outlier set  $\mathcal{D}_v$ . Assuming that  $\mathcal{D}_v$  may not represent the full spectrum of anomalies, we should pick methods that do not overfit to  $\mathcal{D}_v$ . OD-test tells us which methods are less likely to overfit to a chosen  $\mathcal{D}_v$  and therefore be more reliable in the face of an unseen OOD sample. Our results show that a two-dataset evaluation scheme can be too optimistic in identifying the best available method. OD-test is a more realistic evaluation of OOD sample detectors. In practice, the outlier set  $\mathcal{D}_v$  should contain the largest variety of anomalies that we can use, and the method should be the one that is more accurate and less likely to overfit.

## 5 Conclusion

By detecting OOD samples (the outliers), we can ensure the deep learning pipelines operate as expected. We assume that collecting a representative set of outliers for None prediction is not always practical. Therefore, we prefer methods that detect unseen outliers (the unknown unknowns) over methods that only detect previously seen outliers (the known unknowns). To assess how well methods cope with unknown unknowns we introduced a third dataset in the evaluation of outlier detectors. Furthermore, we averaged the accuracy across all combinations of outlier datasets to reduce the measurement bias. We called this evaluation strategy OD-test. OD-test is a new formulation of the problem that provides a more realistic assessment of the OOD detection methods. We also presented a new benchmark for OOD sample detection within image classification pipelines based on OD-test. We showed that the traditional supervised learning approach to OOD detection does not always yield reliable results – the previous assessments of the OOD detectors are too optimistic to be practical in many scenarios. We presented a comprehensive evaluation of a diverse set of approaches across a wide variety of datasets for the OOD detection problem. Furthermore, we showed that *none of the methods is suitable out-of-the-box for high-dimensional images*. We release the open-source PyTorch project with the pre-trained models to replicate the results of our study. We invite the community to tackle the outlined challenges in this work.

**Acknowledgments** We would like to thank the reviewers for their helpful comments. We gratefully acknowledge the support of NVIDIA Corporation through the donation of the GPUs used for this research. This work was supported in part by the National Science and Engineering Research Council of Canada (NSERC), and the Canadian Institute for Advanced Research (CIFAR) Learning in Machines and Brains program.

## References

- [1] Peter L. Bartlett and Marten H. Wegkamp. Classification with a Reject Option using a Hinge Loss. *JMLR*, 2008.
- [2] Abhijit Bendale and Terrance E. Boult. Towards Open Set Deep Networks. In *CVPR*, 2016.
- [3] Corinna Cortes and Vladimir Vapnik. Support-vector Networks. *Machine learning*, 20(3):273–297, 1995.
- [4] Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with Rejection. *International Conference on Algorithmic Learning Theory (ALT)*, 2016.
- [5] Corinna Cortes, Giulia DeSalvo, Mehryar Mohri, and Scott Yang. On-line Learning with Abstention. *ArXiv e-prints*, 2017.
- [6] Lucas Deecke, Robert Vandermeulen, Lukas Ruff, Stephan Mandt, and Marius Kloft. Anomaly Detection with Generative Adversarial Networks, 2018.
- [7] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *ICML*, 2016.
- [8] Yarin Gal and Zoubin Ghahramani. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. In *NIPS*, 2016.
- [9] Markus Goldstein and Seiichi Uchida. A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data. *PLoS one*, 11(4), 2016.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *NIPS*, 2014.
- [11] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On Calibration of Modern Neural Networks. In *ICML*, 2017.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [13] Dan Hendrycks and Kevin Gimpel. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *ICLR*, 2017.
- [14] Dan Hendrycks, Mantas Mazeika, and Thomas G Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018.
- [15] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. In *CVPR*, 2017.
- [16] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level Accuracy with 50x Fewer Parameters and < 0.5 MB Model Size. *ArXiv e-prints*, 2016.
- [17] Alex Kendall and Yarin Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *NIPS*, 2017.

- [18] Diederik P Kingma and Max Welling. Auto-encoding Variational Bayes. In *ICLR*, 2014.
- [19] Mark Kliger and Shachar Fleishman. Novelty Detection with GAN. *ArXiv e-prints*, 2018.
- [20] Alex Krizhevsky and Geoffrey Hinton. Learning Multiple Layers of Features from Tiny Images. 2009.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. Imagenet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012.
- [22] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *NIPS*, 2017.
- [23] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [24] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training Confidence-calibrated Classifiers for Detecting Out-of-Distribution Samples. In *ICLR*, 2018.
- [25] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *NeurIPS*, 2018.
- [26] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. In *ICLR*, 2018.
- [27] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Ng. Reading Digits in Natural Images with Unsupervised Feature Learning. *NIPS Workshops*, 2011.
- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic Differentiation in PyTorch. *NIPS Workshops*, 2017.
- [29] Igor M. Quintanilha, Roberto de M. E. Filho, José Lezama, Mauricio Delbracio, and Leonardo O. Nunes. Detecting Out-Of-Distribution Samples Using Low-Order Deep Features Statistics, 2018.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.
- [31] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *ICLR*, 2017.
- [32] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker Discovery. In *International Conference on Information Processing in Medical Imaging*, 2017.

- [33] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support Vector Method for Novelty Detection. In *NIPS*, 2000.
- [34] Gabi Shalev, Yossi Adi, and Joseph Keshet. Out-of-distribution detection using multiple semantic label representations. In *NeurIPS*, 2018.
- [35] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-scale Image Recognition. *ArXiv e-prints*, 2014.
- [36] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing Properties of Neural Networks. In *ICLR*, 2014.
- [37] Apoorv Vyas, Nataraj Jammalamadaka, Xia Zhu, Dipankar Das, Bharat Kaul, and Theodore L Willke. Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In *ECCV*, 2018.