

# Likelihood-based Out-of-Distribution Detection with Denoising Diffusion Probabilistic Models

Joseph Goodier  
jg852@bath.ac.uk  
Neill D. F. Campbell  
n.campbell@bath.ac.uk

Centre for Accountable, Responsible  
and Transparent Artificial Intelligence  
(ART-AI),  
University of Bath,  
Bath, UK

---

## Abstract

Out-of-Distribution detection between dataset pairs has been extensively explored with generative models. We show that likelihood-based Out-of-Distribution detection can be extended to diffusion models by leveraging the fact that they, like other likelihood-based generative models, are dramatically affected by the input sample complexity. Currently, all Out-of-Distribution detection methods with Diffusion Models are reconstruction-based. We propose a new likelihood ratio for Out-of-Distribution detection with Deep Denoising Diffusion Models, which we call the Complexity Corrected Likelihood Ratio. Our likelihood ratio is constructed using Evidence Lower-Bound evaluations from an individual model at various noising levels. We present results that are comparable to state-of-the-art Out-of-Distribution detection methods with generative models.

## 1 Introduction

Out-of-Distribution (OOD) detection is a sub-class of uncertainty estimation that is a critical topic in the field of machine learning (ML). With the increasing complexity and scale of modern ML models, it has become increasingly important to understand the limitations and potential failure modes of these models. OOD detection refers to the ability of a model to identify inputs that are significantly different from those it was trained on, which is often achieved by looking at the uncertainty of a sample under a model. By detecting OOD inputs, ML models can avoid making incorrect or misleading predictions, which is especially important in high-stakes applications such as healthcare, finance, and autonomous driving.

Generative models are particularly well suited for OOD detection because they are trained to model the distribution of the data, as opposed to the supervised learning approach that learns to map samples to labels. Likelihood estimations can be used to assess how likely a given sample is, under the learned model. The use of the likelihood for OOD detection has been applied to Variational AutoEncoders (VAEs) [14], flow-based models like GLOW [15] and auto-regressive-based models like PixelCNN++ [8].

Denoising Diffusion Probabilistic models (DDPMs) [10] provide a new addition to the landscape of generative models. DDPMs are generative models that add noise to an input




Figure 1: Plots displaying how the noising level of an in-distribution (ID) (FashionMNIST, **Blue**) sample and OOD (MNIST, **Red**) sample affects the pixel-wise loss and reconstruction fidelity. The samples are reconstructed (**Second & Fifth row**) in a single step, as opposed to iteratively. This plot demonstrates complexity bias as the more complex ID sample has a higher loss signal at lower noising levels than the less complex OOD sample (**Bottom row**), leading to a lower likelihood for the ID sample than the OOD sample. This plot also demonstrates that low noising levels, which are most corrupted by image complexity, are where low-level image features emerge.

sample and remove it using a noise prediction network. The amount of noise added is learned during training, allowing pure noise inputs to be denoised and samples to be generated based on the learned distribution. All current methods that use DDPMs for OOD detection do so with reconstruction-based methods. Similar to VAEs, DDPMs are trained using the Evidence Lower-Bound (ELBO) objective as an estimate of the marginal log-likelihood of the model. It follows that DDPMs could also be used for likelihood-based OOD detection. **In this paper** we present :

1. Evidence that input sample complexity dramatically affects the ELBO contributions from low noising levels in DDPMs, as is seen with other generative models.
2. A likelihood-based OOD detection method using DDPMs. We use a likelihood ratio that is calculated using ELBO evaluations from low noise levels over the total ELBO from all noise levels. We define it to be the **Complexity Corrected Likelihood Ratio (CCLR)**.

## 1.1 Likelihood-based OOD detection

Bishop *et al.* [2] proposed the use of model likelihood for OOD detection, where the author suggested it could be viewed as a probability under the model. Here, a sample would be assigned as OOD given a one-tailed test using the likelihood calculated by the model. In an ideal setup, a model would assign a high likelihood to in-distribution (ID) samples and a low likelihood to OOD samples. Recently, a number of contemporaneous publications challenged this assumption [5, 9, 22]. Specifically relating to generative models, Nalisnick

*et al.* [22] and Choi *et al.* [5] claim that generative models, when trained on an ID dataset and treat another dataset as OOD, similar to the above paradigm, assign higher likelihood scores to some OOD data. In light of this result, the community has released a number of studies that investigate and seek to redeem the likelihood for OOD detection.

Ren *et al.* [27] claim that the reason these likelihood-based methods fail is that the likelihood is dramatically affected by population-level background statistics that are not relevant to the specific details of the ID data. To remedy this, the authors propose a likelihood ratio score between the likelihood of a sample under the model and a background model that is trained on random perturbations of the input data. In a similar vein, Serra *et al.* [32] claim that the corruption of the likelihood score is related to the complexity of the test sample. The authors propose a likelihood ratio, based on an offset to the likelihood that is calculated with a loss-less compression of the sample. Building on the likelihood ratio, Xiao *et al.* [36] propose another likelihood-based OOD score with a single model, called Likelihood Regret. The score is the difference in the likelihood of a sample under a model trained on the ID data, and the likelihood of the sample under the same model when it is over-fit to the sample. The greater the change in likelihood, the greater the likelihood regret score.

Havtorn *et al.* [8] claims that the low-level features learned by VAEs generalise well across datasets and dominate the likelihood scores. They propose using a hierarchical VAE, where features at different levels in the input data are learned at the different levels in latent representations. Ran *et al.* [26] propose a method using noise contrastive priors that they call Improved Noise Contrastive Priors (INCP), where they leverage synthetic OOD data created from ID samples to train a VAE. This contrastive prior can then be used to calculate an ELBO ratio, which is used as an OOD detection metric, with very high levels of success. Zhang *et al.* [38] demonstrate that OOD generalisation ability depends on non-local features, which are defined in opposition to local or low-level features. The authors estimate a non-local likelihood score by calculating a likelihood ratio of two auto-regressive models; one trained on local features and the other trained on global features. Other methods involve using the Watanabe-Akaike Information Criterion (WAIC) for an ensemble of models as an OOD detection score [6], whilst Morningstar *et al.* [21] use Density-of-States to assess the typicality of model statistics as an OOD detection score.

The common thread across these contributions is that the likelihood of a sample under a generative model is corrupted by low-level or local features, that generalise well between datasets. These features contribute in an outsized amount to the complexity of an image, where we define image complexity as *a measure of the degrees of freedom required to represent the images in the dataset*. The community has shown that what’s actually important for OOD detection are the high-level, non-local semantics of an image that are unique to the dataset. The likelihood in generative models captures both of these, and the low-level image complexity needs to be corrected for in some way to produce a robust, likelihood-based OOD detection method with generative models.

## 1.2 Deep Denoising Diffusion Models

DDPMs [10, 33] have emerged as a new class of generative models that have proven to be powerful and scaleable in many contexts. They are similar to normalising flows in that they map the input data to latent variables of the same dimension. They also resemble VAEs, as they estimate the likelihood through the ELBO objective [25]. Unlike VAEs, instead of a learned encoder, the noise is added to the data in a defined schedule. The learning is done in the inverse process by a noise prediction network, that sequentially “denoises” the data.

This inverse process is then used to produce samples from pure noise. We refer the reader to Yang *et al.* [37] for a thorough review of various other diffusion model implementations and applications.

### 1.3 Diffusion-based OOD Detection

A number of OOD detection methods that use DDPMs have been proposed. However, these methods all rely on reconstruction-based scores. The intuition of a reconstruction-based approach is that generative models should be able to faithfully reconstruct ID samples, and unfaithfully reconstruct OOD samples. The difference between the original sample and its reconstruction can be used as an OOD score.

Graham *et al.* [7] use diffusion models to sequentially reconstruct samples for multiple noising levels. They use a combination of MSE and the LPIPS metric, which uses the distance between the deep features of the sample and reconstruction [39], as an OOD metric. Liu *et al.* [19] employ diffusion models for inpainting occluded samples. They also use the LPIPS metric, where they calculate the median distance between the deep features of the inpainted reconstruction and the original, unoccluded sample as an OOD score. In a related approach, Liu *et al.* [18] use a combination of a discriminator and DDPMs as an asymmetric interpolation method for OOD detection.

As previously stated, similar to VAEs, DDPMs are trained using the ELBO objective. It follows that DDPMs could also be used for likelihood-based OOD detection. These methods, when applied to other generative models, have historically demonstrated strong performance in OOD detection (See Table. 2) and have been explored more extensively than reconstruction-based methods. We posit that DDPMs, by extension, should also be able to exhibit a strong performance with likelihood-based OOD detection, given the current understanding of the considerations required to make likelihood-based OOD detection robust (See Fig. 1).

## 2 Methodology

### 2.1 Model

The model implementation used in this work is based on the discrete-time DDPM first introduced by Ho *et al.* [10]. Discrete-time DDPMs work by contaminating a sample,  $\mathbf{x}_0$ , with a distribution of noise defined by a time variable  $t$ , where  $0 \leq t \leq T$ . At  $T$ , the sample has been degraded to isotropic noise. The noise distribution added at each  $t$  is defined by a noising schedule,  $\beta_t$ , that produces noised samples  $\mathbf{x}_t$ . We employ the cosine noising schedule introduced by Nichol *et al.* [24],

$$\bar{\alpha}_t = \frac{f(t)}{f(0)}, f(t) = \cos\left(\frac{t/T + s}{1+s} * \frac{\pi}{2}\right)^2, \quad (1)$$

where  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=0}^t \alpha_s$  and  $s > 0$  is an offset to ensure that  $\beta_t$  is not too small at  $t = 0$ . This schedule was chosen because Nichol *et al.* [24] showed that a cosine schedule was shown to distribute the information that emerges from DDPMs when sampling more smoothly across  $t$ -values than a linear schedule. The noising process can then be defined as,

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t | \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}). \quad (2)$$

For the reverse process, a network learns the noise distribution added for each timestep, which is then removed from the sample. The reverse process can be defined as,

$$p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}|\mu_{\theta}(x_t, t), \Sigma_{\theta}(x_t, t)), \quad (3)$$

where both  $\mu_{\theta}(x_t, t)$  and  $\Sigma_{\theta}(x_t, t)$  are learned by the network, which is parameterised by  $\theta$ . During the learning process, a DDPM optimises the ELBO objective, which is the variational bound on the marginal log-likelihood [10]. In practice, the  $\Sigma_{\theta}(x_t, t)$  term in Eq. 3 is fixed as a function of  $t$  i.e. the variance is now  $\Sigma(t)$ ; which then allows the model to predict solely the noise term for each timestep. It follows that the ELBO reduces down to a sum of the L2 distance between the predicted and actual noise for each timestep,

$$\mathcal{L}_{\theta} = \frac{1}{T} \sum_{t=0}^T \left[ \|\hat{\epsilon}_{\theta,t} - \epsilon\|^2 \right], \quad (4)$$

where  $\epsilon$  is the sample of isotropic Gaussian noise, and  $\hat{\epsilon}_{\theta,t}$  is the approximation of the noise from the network. For a full derivation of the ELBO for a DDPM and how it reduces to Eq. 4, please see Prince, Ch. 18 [25].

## 2.2 Complexity Bias in DDPMS

It is well-established that generative models can assign a higher likelihood to some OOD samples than to ID samples [6, 8, 22]. Serra *et al.* [5] relate this to the idea that “generative models exhibit a strong bias towards the complexity of the corresponding inputs”. They provide evidence that there is a negative correlation between the image complexity of the input samples and the log-likelihood scores. The qualitative behaviour of log-likelihood is the same as the likelihood i.e. OOD samples should have a low log-likelihood under the model and vice versa for ID samples.

When training a vanilla DDPM, you are noising the sample to some  $t$ -value, predicting the noise that was added to the sample and taking a difference between actual noise and predicted noise as your loss. For low  $t$ -values, the noise added to a sample is very subtle (See Fig. 1). Therefore, for less complex images, it is easier to predict the noise and therefore the loss should be lower, even if these low-complexity samples are OOD. It follows then that the ELBO contributions from low  $t$ -values should also be affected by the complexity of the image statistics by an outsized amount through this process. In practice, when you are optimising the ELBO, you are minimising the upper bound on the negative marginal log-likelihood. Therefore, a *lower* loss implies a *higher* log-likelihood estimate.

It follows that an OOD sample with lower complexity can have a higher marginal log-likelihood estimate than an ID sample with higher complexity. This effect can be seen in DDPMS in Fig. 2, where we compare the average loss contributions across all  $t$ -values for RGB and single-channel dataset pairs of differing complexity. For the RGB dataset pairs (**Left column**) the lower loss scores for the less visually complex SVHN, when both ID and OOD, imply a higher marginal log-likelihood estimate and vice versa for CIFAR10. However, the difference is less when SVHN is the OOD class. For the single-channel image pairs (**Right column**), this difference is very prominent in the MNIST/FashionMNIST where the low-complexity MNIST is also the ID dataset. For the FashionMNIST/MNIST dataset pair, for low  $t$ -values, the lower complexity MNIST dataset has lower loss values. But as the  $t$ -values increase, the ID class displays lower loss values and, therefore, a higher marginal

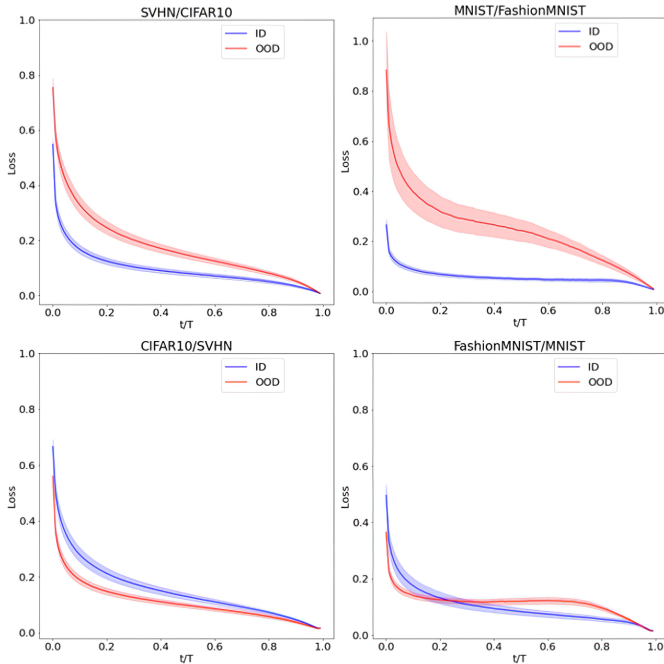


Figure 2: A plot showing the average loss as a function of  $t/T$  for dataset pairs. The top two plots demonstrate the expected behaviour that the OOD datasets have a higher average loss and therefore a lower log-likelihood under the model for SVHN/CIFAR10 (**Top Left**) and MNIST/FashionMNIST (**Top Right**). However, when the more complex dataset is ID, this expected behaviour no longer holds and there is a higher loss for ID dataset across  $t/T$  for CIFAR10/SVHN (**Bottom Left**) and for low  $t/T$  for FashionMNIST/MNIST (**Bottom Right**), and therefore a lower log-likelihood under the model.

log-likelihood estimate. This demonstrates that image complexity corrupts the ELBO contributions from low  $t$ -values, and needs to be corrected for to produce a robust likelihood-based OOD detection method using DDPMS.

### 2.3 Likelihood Ratio for DDPMS

From Eq. 4, we define a decomposed ELBO for DDPMS as,

$$\mathcal{L}_\theta^{<k} = \frac{1}{k} \sum_{t=0}^k w(t) [\|\hat{\epsilon}_{\theta,t} - \epsilon\|_1], \quad (5)$$

where  $k$  is a noising threshold defined by  $0 < k < T$ .  $\mathcal{L}_\theta^{<k}$  defines the ELBO for samples noised up to the threshold  $k$ . For a well-chosen value of  $k$ ,  $\mathcal{L}_\theta^{<k}$  should capture low  $t$ -value loss contributions that are affected by sample complexity.

During training and inference, we use an L1 rather than an L2 loss. The main reason for this change relates to the findings of Saharia *et al.* [28] that the use of L1 reduces sample diversity, which in turn reduces the probability of hallucinations in diffusion models. This is important for OOD detection, as a tightly learned distribution over the training data is critical for the ability of the method to identify OOD samples. The L1 loss has also been shown to lead to increased training stability for DDPMS [4, 29].

We also include a linear importance weight  $w(t) = 0.5 + t/T$ . Various DDPM implementations in the literature have been shown to be special cases of weighting the loss, and have been shown to improve the performance of DDPMS in terms of image generation quality [12]. Kingma *et al.* [13] show that if a weighting term is monotonic with  $t$ , then “the weighted loss corresponds to maximizing the ELBO of noise-perturbed data”. Our coefficient  $w(t)$  increases the importance of the L1 contributions added at higher  $t$ -values, to reduce the influence on the ELBO of the complexity bias at low  $t$ -value contributions. This term also has the property that,  $\frac{1}{T} \int_{t=0}^T w(t) dt = 1$ .

The  $\mathcal{L}_\theta^{<k}$  can be interpreted as the background model that captures background statistics from the original likelihood ratio for OOD work by Ren *et al.* [17], but for DDPMS and using only one model. It’s worth noting here that if  $k = T$ , then  $\mathcal{L}_\theta^{<T}$  becomes the full DDPM ELBO term  $\mathcal{L}_\theta$ . The only difference is the importance weights  $w(t)$ . If  $k$  is well chosen, it is possible to construct a likelihood ratio that corrects for the influence of complexity bias that affects the ELBO score as an OOD estimate. We define this ratio to be the CCLR,

$$CCLR_{k/T} = \mathcal{L}_\theta^{<k} - \mathcal{L}_\theta^{<T}. \quad (6)$$

The ordering is due to the fact that  $\mathcal{L}_\theta^{<k}$  and  $\mathcal{L}_\theta^{<T}$  are ELBO estimates, which implies they are bounds on the *negative* marginal log-likelihood. Therefore, they are multiplied by  $-1$  so that the likelihood ratio has the property that an OOD sample should have a low likelihood ratio under this model. CCLR is also a difference in the log-likelihood space, which implies a ratio in the likelihood space. See Ablation in Appendix A for justification of algorithmic implementation choices.

## 2.4 OOD Detection Algorithm

In order to optimise CCLR to perform OOD detection with a DDPM, we have to take a number of model-specific steps to enable an inference estimate. These steps are to ensure that both  $\mathcal{L}_\theta^{<k}$  and  $\mathcal{L}_\theta^{<T}$  are calculated from the same number of sample estimates. Firstly, each sample is expanded and passed through the model as a batch of size  $N_{bs}$ . This is to ensure that we get sufficient coverage of  $t$ -values to marginalise over, as the batch size of the samples,  $\mathbf{xs}$ , must match the batch size of the  $t$ -values,  $\mathbf{ts}$ , when passed through the DDPM. However, the  $\mathcal{L}_\theta^{<k}$  term is calculated from  $T/k$  times fewer samples than the full  $\mathcal{L}_\theta^{<T}$  term. To account for this, we pass the batch through the model  $n$  number of times, where  $n = T/k$ , combine the L1 losses across the  $t$  dimension and then calculate the  $\mathcal{L}_\theta^{<k}$  term. This corrects for the difference of samples across  $t$  with samples across noise. The full algorithm for OOD detection is as follows:

# 3 Experiments & Results

## 3.1 Model Implementation

The backbone DDPM code is based on the discrete-time DDPM from the library `denoising-diffusion-pytorch` [14]. All models were instantiated with  $T=1000$  timesteps, trained for 100 epochs and with a learning rate of  $2.0 \times 10^{-5}$ . A cosine noising schedule was used for all trials, following Nichol *et al.* [24]. We use the L1 norm as the training loss for all objectives. The noise prediction network was instantiated with 64 feature maps for the initial block. There were 4 blocks, each with multipliers for the feature maps of (1,2,4,8), which result in map resolutions of  $32 \times 32$ ,  $16 \times 16$ ,  $8 \times 8$  and  $4 \times 4$  respectively.

**Algorithm 1:** Complexity Corrected Likelihood Ratio (CCLR)**INPUT:** Data sample  $x_0$ , noising threshold  $k$ , inference batch-size  $N_{bs}$ , diffusion model  $\theta$ **OUTPUT:**  $CCLR_{k/T}$ 

```

0:  $\mathbf{ts} = \text{linspace}(0, T, N_{bs})$ 
0:  $\mathbf{w} = \text{linspace}(0.5, 1.5, N_{bs})$ 
0:  $\mathbf{xs} = \text{expand}(N_{bs}, \mathbf{x}_0)$ 
0:  $\mathcal{L}_\theta^{<T} = \text{mean}(\mathbf{w} * \text{diffuse}(\mathbf{xs}, \mathbf{ts}))$ 
0:  $n = T/k$ 
0:  $\mathbf{ls} = []$ 
1: for  $\_$  in  $\text{range}(n)$  do
1:    $\mathbf{ls.append}(\text{diffuse}(\mathbf{xs}, \mathbf{ts}))$ 
2: end for
2:  $\mathcal{L}_\theta^{<k} = \text{mean}(\mathbf{w} * \text{logsumexp}(\mathbf{ls})[:k])$ 
2:  $CCLR_{k/T} = \mathcal{L}_\theta^{<k} - \mathcal{L}_\theta^{<T}$ 

```

## 3.2 Experiments

For single-channel dataset pairs, we train on FashionMNIST (Xiao *et al.* [5]) and treat the test set as ID. The OOD test set will be MNIST (Lecun *et al.* [17]). For RGB dataset pairs, we train on CIFAR10 (Krizhevsky [16]) and treat the test set as ID. The OOD test set will be SVHN (Netzer *et al.* [23]). All images are resized to  $32 \times 32$ . For all dataset pairs, we randomly sample 1000 images from the ID test set and the OOD test set and assess the method’s ability to distinguish whether a sample is ID or OOD, on a sample-by-sample basis. We expand each sample to an inference batch size of  $N_{bs} = 100$ . We use the Area-Under-Receiver-Operating Characteristics (AUROC) as the baseline metric to evaluate the performance of OOD detection methods [6, 8, 22, 27, 32, 36, 38].

## 3.3 Results

Dataset Pair	AUROC scores					
	$\mathcal{L}_\theta^{<T}$	$CCLR_{1/2}$	$CCLR_{1/3}$	$CCLR_{1/5}$	$CCLR_{1/10}$	$CCLR_{1/20}$
FashionMNIST/MNIST	0.784	0.985	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>	<b>1.000</b>
CIFAR10/SVHN	0.941	<b>0.964</b>	0.958	0.955	0.960	0.947

Table 1: AUROC scores for FashionMNIST/MNIST and CIFAR10/SVHN dataset pairs. Results are reported for the whole ELBO objective as an OOD score, as well as likelihood ratios for various  $k/T$ -values.

In Table 1, we present AUROC results for dataset pairs for RGB and single-channel image dataset comparisons. For each dataset pair, we present AUROC scores using just the ELBO as an OOD score, and for CCLR scores using a range of  $k/T$  values. For the FashionMNIST/MNIST dataset pair, the ELBO objective for DDPMs is shown to be a poor estimator for OOD detection. The CCLR for various  $k/T$  values proves to be a very strong OOD detection score for single-channel dataset pairs. All  $k/T$ -values that were  $< 1/3$  lead to a 1.00 AUROC score, or the datasets were completely separated using this score up to 3sf (See Fig. 3). For CIFAR10/SVHN, the ELBO objective was a surprisingly good OOD estimator of 0.941, and is within 0.006 of the lowest CCLR score. The CCLR proved to be a strong estimator with peak AUROC scores of 0.964 at  $k/T = 1/2$  (See Fig. 3).

Our results for FashionMNIST/MNIST are on par with the state-of-the-art results of the Local Auto-Regressive Model [33] and INCP-VAE [26] and outperform all current DDPM-based approaches for this dataset pair (See Table 2). For CIFAR10/SVHN our results are






Figure 3: Violin plots displaying histograms of the CCLR score for both CIFAR10/SVHN (Left) and FashionMNIST/MNIST (Right) dataset pairs. For CIFAR10/SVHN violin plot, the CCLR scores were calculated using  $k/T = 1/2$ . For FashionMNIST/MNIST, CCLR scores were calculated using  $k/T = 1/5$ . Both of the selected  $k/T$ -values gave the highest AUROC scores from Table 1 for each respective dataset pair.

Method	Dataset Pair	
	FashionMNIST/MNIST	CIFAR10/SVHN
Likelihood-based		
WAIC (model ensemble) [9]	0.766	<b>1.000</b>
Likelihood Ratio [10]	0.997	0.912
Likelihood Regret [16]	0.988	0.875
BIVA with LLR [9]	0.984	0.891
MSMA KD Tree [10]	0.693	0.991
Probabilistic Auto-Encoder [9]	0.997	-
S using PCNN and FLIF [10]	0.967	0.929
S using Glow and FLIF [10]	0.998	0.950
DoSE using VAE [10]	0.998	-
DoSE using GLOW [10]	-	0.973
Local Auto-Regressive Model [16]	<b>1.000</b>	0.969
INCP-VAE [10]	<b>1.000</b>	<b>1.000</b>
CCLR using DDPM (Ours)	<b>1.000</b>	0.964
Reconstruction-based		
Mahalanobis* [9]	0.986	0.991
DDPM + SSIM [9]	0.974	0.979
Diffusion-Based Neighborhood [16]	-	0.950
LMD [16]	<b>0.992</b>	<b>0.992</b>

Table 2: AUROC scores for FashionMNIST/MNIST and CIFAR10/SVHN dataset pairs for OOD detection with generative models for both likelihood-based and reconstruction-based methods. \* implies knowledge of OOD labels.

competitive with the other likelihood-based methods. However, on this dataset pair, CCLR is outperformed by the reconstruction-based methods using DDPMS.

## 4 Conclusions & Future Work

In this work, we propose a likelihood ratio for OOD detection, CCLR, with DDPMS. We show that likelihood-based OOD detection can be extended to diffusion models by leveraging the fact that DDPMS, like other likelihood-based generative models, are affected by the input sample complexity. For future work, there have been improvements that have been proposed for DDPMS that could benefit our OOD detection method. These include architecture improvements for the noise prediction network [10], new image transforms [9, 60] and reformulation of the DDPM objective and noising schedule [17, 18]. Our work uses a vanilla DDPM to act as a proof of concept for likelihood-based OOD detection with DDPMS, and it is likely that performance gains could be made by implementing these improvements.

## 5 Acknowledgements

This work was supported by Innovative Physics Ltd, the UKRI Centre for Doctoral Training in Accountable, Responsible and Transparent AI (grant number EP/S023437/1) and the UKRI CAMERA project (EP/T022523/1). The authors would like to thank Jessica Nicholson for her help in editing this paper.

## References

- [1] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022.
- [2] Christopher M Bishop. Novelty detection and neural network validation. *IEE Proceedings-Vision, Image and Signal processing*, 141(4):217–222, 1994.
- [3] Vanessa Böhm and Uroš Seljak. Probabilistic auto-encoder. *arXiv preprint arXiv:2006.05479*, 2020.
- [4] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020.
- [5] Hyunsun Choi, Eric Jang, and Alexander A Alemi. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018.
- [6] Taylor Denouden, Rick Salay, Krzysztof Czarnecki, Vahdat Abdelzad, Buu Phan, and Sachin Vernekar. Improving reconstruction autoencoder out-of-distribution detection with mahalanobis distance. *arXiv preprint arXiv:1812.02765*, 2018.
- [7] Mark S Graham, Walter HL Pinaya, Petru-Daniel Tudosiu, Parashkev Nachev, Sebastien Ourselin, and M Jorge Cardoso. Denoising diffusion models for out-of-distribution detection. *arXiv preprint arXiv:2211.07740*, 2022.
- [8] Jakob D Havtorn, Jes Frellsen, Søren Hauberg, and Lars Maaløe. Hierarchical vaes know what they don’t know. In *International Conference on Machine Learning*, pages 4117–4128. PMLR, 2021.
- [9] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:1812.04606*, 2018.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [11] Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. *arXiv preprint arXiv:2301.11093*, 2023.
- [12] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021.
- [13] Diederik P Kingma and Ruiqi Gao. Understanding the diffusion objective as a weighted integral of elbos. *arXiv e-prints*, pages arXiv–2303, 2023.

- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [15] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- [16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [17] 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.
- [18] Luping Liu, Yi Ren, Xize Cheng, and Zhou Zhao. Out-of-distribution detection with diffusion-based neighborhood.
- [19] Zhenzhen Liu, Jin Peng Zhou, Yufan Wang, and Kilian Q Weinberger. Un-supervised out-of-distribution detection with diffusion inpainting. *arXiv preprint arXiv:2302.10326*, 2023.
- [20] Ahsan Mahmood, Junier Oliva, and Martin Styner. Multiscale score matching for out-of-distribution detection. *arXiv preprint arXiv:2010.13132*, 2020.
- [21] Warren Morningstar, Cusuh Ham, Andrew Gallagher, Balaji Lakshminarayanan, Alex Alemi, and Joshua Dillon. Density of states estimation for out of distribution detection. In *International Conference on Artificial Intelligence and Statistics*, pages 3232–3240. PMLR, 2021.
- [22] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don’t know? *arXiv preprint arXiv:1810.09136*, 2018.
- [23] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [24] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [25] Simon J.D. Prince. *Understanding Deep Learning*. MIT Press, 2023. URL <https://udlbook.github.io/udlbook/>.
- [26] Xuming Ran, Mingkun Xu, Lingrui Mei, Qi Xu, and Quanying Liu. Detecting out-of-distribution samples via variational auto-encoder with reliable uncertainty estimation. *Neural Networks*, 145:199–208, 2022.
- [27] Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. *Advances in neural information processing systems*, 32, 2019.
- [28] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.

- [29] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [30] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- [31] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017.
- [32] Joan Serrà, David Álvarez, Vicenç Gómez, Olga Slizovskaia, José F Núñez, and Jordi Luque. Input complexity and out-of-distribution detection with likelihood-based generative models. *arXiv preprint arXiv:1909.11480*, 2019.
- [33] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [34] Phil Wang. denoising-diffusion-pytorch. <https://github.com/lucidrains/denoising-diffusion-pytorch.git>, 2020.
- [35] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [36] Zhisheng Xiao, Qing Yan, and Yali Amit. Likelihood regret: An out-of-distribution detection score for variational auto-encoder. *Advances in neural information processing systems*, 33:20685–20696, 2020.
- [37] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022.
- [38] Mingtian Zhang, Andi Zhang, and Steven McDonagh. On the out-of-distribution generalization of probabilistic image modelling. *Advances in Neural Information Processing Systems*, 34:3811–3823, 2021.
- [39] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.