

## Denoising diffusion models for out-of-distribution detection

Mark S. Graham  
King's College London  
mark.graham@kcl.ac.uk

Walter H.L. Pinaya  
King's College London  
walter.diaz-sanz@kcl.ac.uk

Petru-Daniel Tudosiu  
King's College London  
petru.tudosiu@kcl.ac.uk

Parashkev Nachev  
University College London  
p.nachev@ucl.ac.uk

Sebastien Ourselin  
King's College London  
sebastien.ourselin@kcl.ac.uk

M. Jorge Cardoso  
King's College London  
m.jorge.cardoso@kcl.ac.uk

### Abstract

*Out-of-distribution detection is crucial to the safe deployment of machine learning systems. Currently, unsupervised out-of-distribution detection is dominated by generative-based approaches that make use of estimates of the likelihood or other measurements from a generative model. Reconstruction-based methods offer an alternative approach, in which a measure of reconstruction error is used to determine if a sample is out-of-distribution. However, reconstruction-based approaches are less favoured, as they require careful tuning of the model's information bottleneck - such as the size of the latent dimension - to produce good results. In this work, we exploit the view of denoising diffusion probabilistic models (DDPM) as denoising autoencoders where the bottleneck is controlled externally, by means of the amount of noise applied. We propose to use DDPMs to reconstruct an input that has been noised to a range of noise levels, and use the resulting multi-dimensional reconstruction error to classify out-of-distribution inputs. We validate our approach both on standard computer-vision datasets and on higher dimension medical datasets. Our approach outperforms not only reconstruction-based methods, but also state-of-the-art generative-based approaches. Code is available at <https://github.com/marksgraham/ddpm-ood>.*

### 1. Introduction

Out-of-distribution (OOD) detection plays a crucial role in the safe deployment of machine learning systems, ensuring that downstream models are only run on data sampled from the distribution they were trained on. OOD detection models can be broadly divided into unsupervised models, which only require in-distribution data for training, and supervised models, which require additional information such as classification labels or sample OOD data. Unsupervised

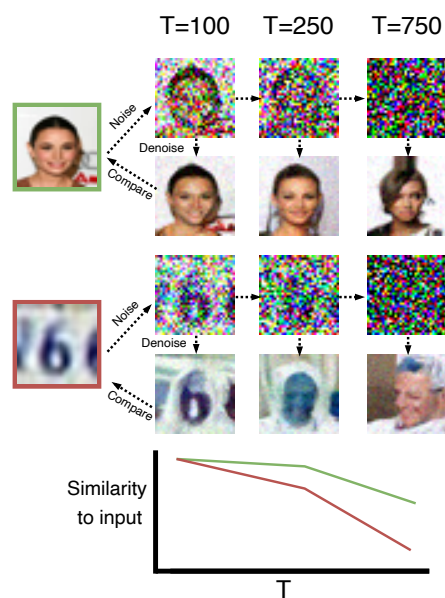


Figure 1. Reconstruction-based OOD detection, with the example of a model trained on CelebA. An in-distribution image from CelebA and an OOD image from SVHN are noised to various levels, reconstructed using the DDPM, and compared to the input. The similarity between inputs and reconstructions is plotted below.

models are appealing as they make no assumptions about the form OOD data will take or the type of downstream task (e.g. classification, segmentation) that will be performed.

The current dominant approach in unsupervised OOD detection is the use of the likelihood or other metrics from a generative model trained on the in-distribution data. However, it has been shown these models can exhibit egregious failures, such as a model trained on CIFAR10 assigning higher likelihoods to samples from the SVHN dataset than samples from CIFAR10 itself [3, 9, 24]. A number of methods have been proposed to address these shortcom-

ings [3, 23, 24, 30, 35]. These models have shown better performance in empirical benchmarks, but recent theoretical work suggests that all these methods will remain vulnerable against at least some OOD data [45].

Reconstruction-based methods offer an alternative approach to unsupervised OOD detection. They involve training a model to reconstruct in-distribution data and using the size of the reconstruction error to detect OOD inputs. However, compared to likelihood-based methods, reconstruction-based methods have received less attention in the literature [33, 44]. A likely reason is that these methods rely on an information bottleneck, such as a latent space that is smaller than the input, to effectively reconstruct in-distribution data but not OOD data. In practice, it is challenging to tune this bottleneck: too small and even the in-distribution data is poorly reconstructed, too large and even OOD data is successfully reconstructed. This need for tuning is undesirable and is likely a key reason these methods have typically been overlooked in favour of generative models for OOD detection.

Diffusion denoising probabilistic models (DDPM) [10, 36] present the information bottleneck issue in an interesting new light. These models are trained to incrementally remove noise from noised inputs. When an input is fully-noised, no information from the input itself is retained, and a DDPM will produce a new sample. However, when applying a DDPM to a partially-noised input, some information from the input is retained, and the denoising process is conditioned on that noisy input; that is, the model attempts to reconstruct the input. The amount of noise applied can be viewed as a variable information bottleneck. The bottleneck is not a property of the trained model itself, such as the size of the latent space in an autoencoder, but rather something we can control externally during model inference, meaning a single trained model can handle many different bottleneck levels. While this interpretation of DDPMs as autoencoders with an externally-controlled bottleneck has been previously discussed [5], to our knowledge, no work has attempted to use this property to perform reconstruction-based OOD detection with DDPMs.

In this work, we apply DDPMs to perform reconstruction-based OOD detection. We measure the quality of a model’s reconstructions of an input noised to a range of different levels and propose to use this set of reconstruction error metrics to determine whether an image is OOD.

## 2. Related Work

Methods for OOD detection can be broadly categorised into supervised, requiring some additional labels or OOD data, or unsupervised, requiring only in-distribution data. This overview of related work focuses on unsupervised methods.

### 2.1. Generative based

A conceptually appealing approach to OOD detection involves fitting a generative model  $p(\mathbf{x}; \theta)$  to a data distribution  $\mathbf{x}$  and evaluating the likelihood of unseen samples under this model. The assumption is that OOD samples will be assigned a lower likelihood than in-distribution samples and can be identified using a simple threshold on this value [2]. It has since been demonstrated that OOD samples will not necessarily be assigned lower likelihoods than in-distribution data; for example, different families of generative models trained on CIFAR10 all assign higher likelihoods to images from SVHN [3, 9, 24].

Subsequent studies have sought to address this shortcoming in generative models. One approach suggests the failures may be due to likelihood estimation errors and proposes using the Watanabe-Akaike Information Criterion (WAIC) across the likelihood estimates of an ensemble of models to identify OOD samples [3]. Other studies postulate that likelihood estimates are affected by population-level background statistics and propose using a likelihood ratio to remove this effect, either obtaining the denominator from an additional trained model [30] or a measure of image complexity [35]. A closely related approach notes that these methods do not perform well for Variational Autoencoders (VAE) and seeks to develop a specific likelihood ratio for this class of model [43]. Another related approach notes that lower-level model features dominate the likelihood and proposes the use of a hierarchical VAE and the requirement that a sample is in-distribution across all levels of the hierarchy [8]. Another strand of work proposes flagging samples as OOD if their likelihoods do not lie within the typical set of a model - put simply, a sample may be considered OOD not only if its likelihood is lower than that of in-distribution data, but also if it is higher [25]. Follow-up work proposes assessing the typicality of multiple summary statistics from the model, not just the likelihood, to assess whether samples are OOD [23].

### 2.2. Reconstruction based

Reconstruction-based methods represent another paradigm for OOD detection. They involve training a model  $R$  to reconstruct an input,  $\hat{\mathbf{x}} = R(\mathbf{x})$ . The intuition is that if  $R$  contains an information bottleneck, such as a latent space of lower dimension than the input size, it will only be capable of faithfully reconstructing inputs from the distribution it was trained on, and will poorly reconstruct OOD inputs. This can be measured using some similarity measure between the input and reconstruction,  $S(\mathbf{x}, \hat{\mathbf{x}})$ , and flagging inputs that show low similarity.

Several works have highlighted practical issues in the use of reconstruction-based methods [4, 21, 28, 48]. If the information bottleneck is too small, a model cannot faithfully reconstruct even in-distribution samples. If the bottleneck is too large, the model is able to learn the identity function,

allowing OOD samples to be reconstructed with low error. The result is that it is often necessary to perform a dataset-specific tuning process to produce a model that performs well, limiting the utility of such models in practice.

Some work has sought to address these issues. It has been suggested to use the Mahalanobis distance [22] in an autoencoder’s feature space as an OOD metric [4]. Other work has sought to reduce the ability of an autoencoder to reconstruct OOD samples by introducing a memory module and forcing the decoder to decode directly from the memory, encouraging any reconstructions to look more similar to in-distribution data [7]. However, none of these reconstruction-based works address the fundamental issue of bottleneck selection and require some tuning of the information bottleneck for the specific in-distribution/OOD dataset pairing being considered.

Perhaps the work most closely related to ours is AnoD-DPM [40]. Whilst focusing on a slightly different problem of detecting localised anomalies, the work also proposes using reconstructions from a DDPM trained on in-distribution data. However, the authors propose to reconstruct from a single  $t$ -value, or noise bottleneck, with the choice of  $t$  tuned to the dataset being considered. We seek to address this shortcoming in our work with DDPM-based OOD detection, which reconstructions from a range of noise values, obviating the need for any dataset-specific tuning. Concurrent to this work, a complementary approach that involves corrupting inputs and reconstructing with DDPMs was also developed [20].

### 3. Method

To enable reconstruction-based OOD detection that is not dependent on a fixed information bottleneck, we propose making use of a trained DDPM [10] to reconstruct images. During training, samples  $\mathbf{x}_0$  are degraded according to a fixed process with Gaussian noise according to a timestep  $t$  and a noise variance schedule  $\beta_t$  to produce noised samples  $\mathbf{x}_t$ , such that

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

where  $0 \leq t \leq T$  and we define  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ . The schedule  $\beta_t$  is designed to increase with  $t$  and have the property that the fully noised  $\mathbf{x}_T$  is close to an isotropic Gaussian,  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ; i.e.  $\mathbf{x}_T$  contains no information about  $\mathbf{x}_0$ . We train a single network to iteratively reverse the diffusion process by estimating the parameters of the denoising step given  $\mathbf{x}_t$  and  $t$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}|\mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)) \quad (2)$$

Given this trained model and a test input  $\mathbf{x}_0$ , we can sample a set of  $\mathbf{x}_t$  for a range of values of  $t$  and estimate their reconstructions,  $\hat{\mathbf{x}}_{0,t} = p_\theta(\mathbf{x}_0|\mathbf{x}_t)$ . Measuring the similarity between each reconstruction and input  $S(\hat{\mathbf{x}}_{0,t}, \mathbf{x}_t)$  provides a

range of similarity scores that can be used to decide whether  $\mathbf{x}_0$  is in-distribution.

The advantage of such a method over other reconstruction-based methods is that the information bottleneck is no longer a property of the network itself, such as the latent-space dimension in an autoencoder, but an externally chosen factor, the amount of noise applied to the input. This allows for reconstructions from a wide number of information bottlenecks, obviating the need for the pre-selection of the appropriate bottleneck for a given dataset through the choice of model architecture.

#### 3.1. The diffusion model

We use the DDPM model from [31], which is a time-conditioned UNet [32]. While it is possible to optimise the variational bound on the negative log-likelihood, it has been found that a simplified training objective works well in practice, and we make use of it in this work. In this simplified scheme, the variance is fixed to time-dependent constants,  $\Sigma_\theta(\mathbf{x}_t, t) = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \mathbf{I}$  and the network directly predicts the added noise  $\epsilon$  at step  $t$  [10]:

$$L_{\text{simple}}(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t)\|^2] \quad (3)$$

#### 3.2. Multiple reconstructions

Our method involves reconstructing an input  $\mathbf{x}_t$  for multiple values of  $t$ . In the DDPM sampling scheme,  $t$  steps are required to obtain the reconstruction  $\hat{\mathbf{x}}_{0,t}$ , with each step requiring an evaluation of the model. For a typical value  $T = 1000$  with reconstructions from 100 starting points equally spaced along the T-chain,  $T_{\text{start}} = [10, 20, 30, \dots, 990, 1000]$  we would need 50500 model evaluations to obtain all the reconstructions; equivalent to the computation required to obtain 50.5 samples from the model.

To expedite the process of obtaining multiple reconstructions, we make use of recent advances in fast sampling from diffusion models. In particular, we employ the PLMS sampler [18], which has been shown to substantially reduce the number of sampling steps whilst maintaining or even improving sample quality. If we select  $T = 1000$  but choose to use just 100 sampling steps, we are still able to perform 100 reconstructions. However, these reconstructions can be done with just 5050 model evaluations, equivalent to the compute required to obtain 5.05 samples from the model with a DPPM sampler, a  $10\times$  speed-up.

#### 3.3. Evaluating similarity

It is common to use the mean-squared error (MSE) between the input and reconstruction to evaluate similarity. In this work, we also choose to use the LPIPS metric, which uses the distance between the deep features of a network (in this case, Alexnet [15]) from two inputs as a measure of their

perceptual similarity. LPIPS has been shown to correlate well with human evaluations of image similarity [47]. Using both MSE and LPIPS gives a total of  $2N$  similarity measurements per input for the  $N$  reconstructions performed. We convert each measurement into a Z-score using the measurements from a validation set for each reconstruction and metric (MSE or LPIPS) separately. We average these  $2N$  Z-scores to produce an OOD score for each input.

## 4. Experiments and Results

### 4.1. Experimental details

**Datasets.** We evaluate our method both on a number of common computer vision benchmarks and, recognising that performance on these benchmarks don't necessarily reflect performance in the real world, on a set of higher-dimension medical imaging datasets. For the computer vision benchmarks we use four in-distribution datasets: Fashion-MNIST [41], CIFAR10 [16], CelebA [19], and SVHN [26]. For the grayscale FashionMNIST, we use MNIST [17] as an OOD dataset. For the other colour datasets, we use all other colour datasets as OOD datasets. CelebA images were resized to 32x32 to match the dimension of the other colour datasets. We also use vertically- and horizontally-flipped versions of each in-distribution dataset as further OOD datasets, giving a total of 15 pairs of in vs out-of-distribution datasets. For the medical evaluation we used images from the MedNIST dataset, consisting of six classes of different organs and modalities: Hand X-ray, Abdomen CT, Chest X-ray, Chest CT, Breast MRI, and Head CT, with 10,000 images per class. We trained models on each class and evaluated against all other classes. We performed evaluations at the dataset's native dimension of  $64 \times 64$ , and also repeated the experiments on data upsampled to  $128 \times 128$ .

**Baselines.** We benchmark against both generative- and reconstruction-based approaches. We used the Glow architecture [13] as the backbone for all generative approaches, following [23], and explored a number of methods in the literature to perform OOD using the trained model:

1. A threshold on the likelihood  $p(\mathbf{x}; \theta)$  as a simple baseline [2].
2. The Watanabe-Akaike Information Criterion (WAIC), obtained by training an ensemble of 5 models and calculating the WAIC score as  $\mathbb{E}_\theta[\log p(\mathbf{x}|\theta_n)] - \text{Var}_\theta[\log p(\mathbf{x}|\theta_n)]$  [3].
3. The single-sample typicality test where the typicality score for a sample  $\mathbf{x}$  is given by  $|\mathbb{H}[p(\mathbf{x}; \theta)] - \log p(\mathbf{x}; \theta)|$ , where  $\mathbb{H}[p(\mathbf{x}; \theta)]$  is calculated as an average over the training set [24].
4. Density-of-States Estimation (DoSE), to our knowledge the current state-of-the-art in unsupervised OOD detec-

tion. DoSE uses a density estimator to evaluate several features obtained from a model - for Glow models, the features are the model likelihood and its two constituent parts, the log-probability of the latent variable  $Z$ , and the log of the determinant between the input and  $Z$ . We used principal component analysis (PCA) to learn a whitening transform and trained a one-class support vector machine (SVM) on the transformed features from the validation set and used it to score new samples [23]. As detailed in their paper, both the PCA and SVM used the default implementations in scikit-learn [27].

For the reconstruction-based approaches, we used:

1. A threshold on the MSE from an autoencoder-based (AE) reconstruction  $\|\mathbf{x} - \hat{\mathbf{x}}\|_2$ .
2. A threshold on the Mahalanobis score, given by  $\alpha D_M(E(\mathbf{x})) + \beta \|\mathbf{x} - \hat{\mathbf{x}}\|_2$  where  $D_M(\mathbf{x})$  is the Mahalanobis distance from the latent space of  $\mathbf{x}$  to the latent space of the training set and  $\alpha$  and  $\beta$  are constants set to the reciprocal of the standard deviation of  $D_M(E(\mathbf{x}))$  and  $\|\mathbf{x} - \hat{\mathbf{x}}\|_2$  respectively, as evaluated on the validation set [4].
3. The MSE from reconstructions using a memory-augmented autoencoder (MemAE), which augments the standard AE with a memory bank and seeks to decode samples from atomic elements in this bank [7].
4. The MSE from a reconstruction from  $t = 250$  for a DDPM, following AnoDDPM [40]. As the method was intended for localised anomaly detection we modify it to make it more suitable for full-image OOD by training on Gaussian rather than simplex noise, as we found simplex noise hindered OOD performance. We refer to this as AnoDDPM-Mod.

**Implementation details.** For our method, we used the DDPM model as described in [31]<sup>1</sup>. We used a 3-layer UNet with [256, 512, 784] channels, with two residual blocks per layer and a single-headed attention block after each residual block in layers 2,3 of the downsampling branch and layer 1 of the upsampling branch. The timestep was sinusoidally embedded and passed through a two-layer MLP with a Swish activation function [29] to create a 1024-dim embedding. We used  $T = 1000$  during training and a linear noise schedule with  $\beta_t$  varying between 0.0015 and 0.0195. All models were trained for 300 epochs using the Adam optimiser [12] with a learning rate of  $2.5e^{-5}$ . At test time, we used the PLMS sampler with 100 timesteps and reconstructed from each of these 100 steps as starting points to produce 100 reconstructions per input. Our method was implemented

<sup>1</sup><https://github.com/CompVis/latent-diffusion/>



	FashionMNIST			CIFAR10				CelebA				SVHN			Rank	
	MNIST	VFlip	HFlip	SVHN	CelebA	VFlip	HFlip	CIFAR10	SVHN	VFlip	HFlip	CIFAR10	CelebA	VFlip		HFlip
<b>Generative-based</b>																
Likelihood [2]	8.5	55.5	51.1	6.1	52.3	50.5	50.0	67.4	5.9	57.9	50.1	99.0	<b>99.9</b>	50.2	50.3	5.3
WAIC [3]	8.8	55.4	51.1	6.0	52.5	50.4	50.0	67.0	5.69	57.8	50.1	99.0	<b>99.9</b>	50.2	50.3	5.7
Typicality [24]	81.1	51.2	49.6	88.6	40.3	50.0	50.0	66.5	88.8	51.2	50.0	97.1	99.8	50.0	49.9	6.6
Density-of-States [23]	<b>98.1</b>	67.1	55.3	96.4	54.6	51.0	50.1	84.9	99.4	69.6	49.9	98.9	<b>99.9</b>	50.3	50.9	3.5
<b>Reconstruction-based</b>																
AutoEncoder	75.0	59.0	50.4	3.2	<b>75.3</b>	50.1	49.9	42.1	2.7	53.2	49.9	<b>99.4</b>	<b>99.9</b>	49.9	49.9	6.5
AutoEncoder Mahlabonis [4]	94.9	79.5	63.0	4.46	71.8	50.9	50.0	64.2	9.6	71.6	50.1	99.3	99.8	50.6	50.6	3.9
MemAE [7]	56.9	59.0	48.7	4.21	69.4	50.3	49.9	51.5	5.8	56.4	49.9	98.6	99.5	49.8	49.7	7.3
AnoDDPM-Mod [40]	91.8	81.0	64.2	37.8	60.2	54.2	50.5	80.2	67.3	78.1	49.4	90.4	94.2	50.2	52.7	4.2
DDPM (ours)	97.4	<b>88.6</b>	<b>65.1</b>	<b>97.9</b>	68.5	<b>63.2</b>	<b>50.5</b>	<b>99.0</b>	<b>100</b>	<b>93.3</b>	<b>50.3</b>	99.0	99.6	<b>58.2</b>	<b>61.6</b>	<b>1.9</b>

Table 1. Results on computer vision datasets. AUC score for each comparison. Bold text indicates the highest value per column. The overall rank is calculated as the average across the ranks for each of the 15 comparisons (lower is better).

in PyTorch and is available at <https://github.com/marksgraham/ddpm-ood>.

We implemented and trained all baselines to enable comparisons across the full set of dataset pairings considered, using the author’s implementations when available. We trained Glow models<sup>2</sup> using the architectural details outlined in [23, 25]: we used 3 blocks (except for the FashionMNIST model, which used 2), each with 8 layers. Each layer used an activation norm, an inverted  $1 \times 1$  convolution and affine coupling with 400 hidden channels per layer. Each model was trained for 100 epochs with a batch size of 64, using the Adamx optimiser with a learning rate of  $5e^{-4}$ , weight decay of  $5e^{-5}$  and a 10 epoch warmup.

We trained<sup>3</sup> three-layer AEs with [128, 128, 256] features in the encoder and [128, 128, 3] in the decoder for colour datasets, and [32, 16, 8] in the encoder and [16, 32, 1] in the decoder for grayscale as in [7]. Each layer is followed by a batch normalisation [11] and a leaky ReLU activation, with upsampling implemented using transposed convolutions. The MemAE contained an additional memory module with sparse shrinkage. Following [7], we used a memory size of 100 for the grayscale datasets and 500 for colour datasets, with a shrinkage threshold of  $2.5e^{-3}$ . Models were trained for 100 epochs with Adam with a learning rate of  $1e^{-4}$  and a batch size of 10.

## 4.2. Results for computer vision datasets

Results are presented in Table 1, reported as AUC scores. The DDPM has the highest average rank across all experiments, with the state-of-the-art DoSE ranked second-highest. In a direct comparison, the DDPM outperforms DoSE on

13/15 dataset pairings. The increases in performance afforded by the DDPM are sometimes substantial, most notably when the OOD dataset is a vertically- or horizontally-flipped version of the in-distribution dataset. For example, on FashionMNIST vs VFlip, DDPM improves performance over DoSE from an AUC of 67.1 to 88.6 and on HFlip from 55.3 to 65.1. The DDPM is the only method able to perform substantially better than chance on certain pairings: for example, no other method scores higher than 50.9 on either SVHN vs VFlip or HFlip, whilst DDPM scores 58.2 and 61.6, respectively.

Fig. 2 shows some reconstructions from the model trained on the SVHN dataset. Reconstructions from all four models are included in Supplementary. Reconstructions of the in-distribution SVHN input still retain similarity to the input up until noising to  $t = 500 - 600$ , whilst the OOD reconstructions start to look dissimilar to their inputs after noising to  $t = 100$  and bear almost no resemblance by  $t = 400$ . The plot also shows that for  $t \gtrsim 700$  the noised images retain very little information from the input and the model outputs begin to resemble unconditioned samples more than they do reconstructions. This suggests that reconstructions from higher  $t$  contribute little to the OOD signal and can potentially be discarded, though we view it as an advantage of our method that it performs well across dataset pairings without any post-hoc need for selecting the range of values to reconstruct from. We explore this further in Sec. 4.5.

As widely reported [3, 24, 30], using the likelihood alone performs very poorly on a number of dataset pairings, such as FashionMNIST vs MNIST and CIFAR10 vs SVHN. Typicality substantially improves performance on the three dataset pairings that simple likelihood performs the worst on, but offers little improvements on other dataset pairings. In agreement with [23] we found that WAIC provides little advantage

<sup>2</sup><https://github.com/y0ast/Glow-PyTorch>

<sup>3</sup><https://github.com/donggong1/memae-anomaly-detection>

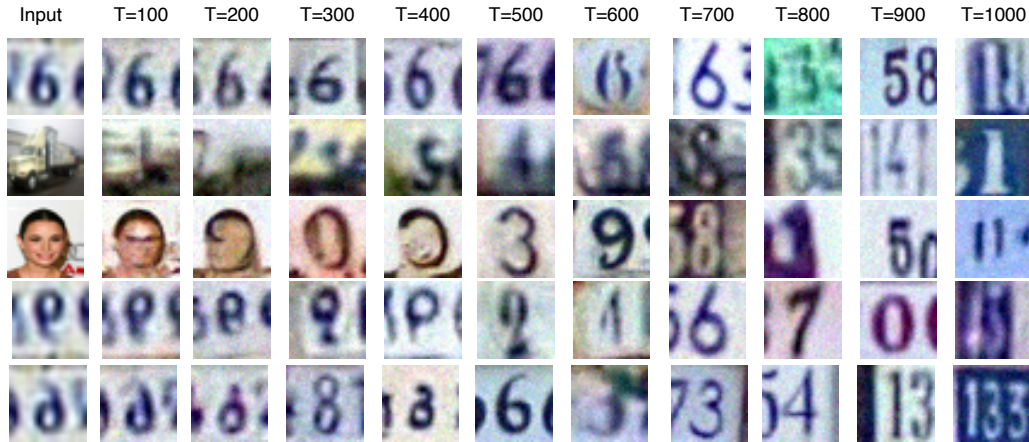


Figure 2. Example reconstructions from a model trained on SVHN for ten different  $t$ -values spaced equally across the chain. Plot shows an in-distribution input (top row) and OOD inputs from CIFAR10, CelebA, SVHN VFlip, and SVHN HFlip (rows 2-5).

	64 × 64					128 × 128					Rank		
	Abdomen	Breast	Chest	CXR	Hand	Head	Abdomen	Breast	Chest	CXR		Hand	Head
<b>Generative-based</b>													
Likelihood [2]	92.9	98.7	97.1	54.9	73.6	96.0	93.5	99.5	99.5	90.1	82.2	93.	5.8
Typicality [24]	94.2	98.7	96.0	70.3	82.2	90.1	93.9	99.6	99.5	83.4	75.9	86.5	6.2
Density-of-States [23]	<b>99.9</b>	99.8	99.2	76.5	88.9	<b>99.9</b>	96.8	99.9	<b>99.7</b>	98.5	79.5	96.2	3.2
<b>Reconstruction-based</b>													
AutoEncoder	99.0	99.4	95.6	96.9	89.5	90.0	99.2	99.7	97.8	94.7	91.5	90.4	4.8
AutoEncoder Mahlabonis [4]	99.6	99.7	95.3	99.0	95.6	93.6	<b>99.9</b>	99.8	99.5	99.2	95.1	94.1	3.3
MemAE [7]	75.8	99.7	85.5	93.0	87.8	60.9	97.7	98.0	96.0	82.2	85.7	58.1	6.6
AnoDDPM-Mod [40]	96.4	88.4	96.7	99.0	98.5	87.7	98.5	98.9	95.0	99.4	99.1	94.1	4.6
DDPM (ours)	99.7	<b>100</b>	<b>99.3</b>	<b>100</b>	<b>99.9</b>	99.7	99.8	<b>100</b>	99.5	<b>100</b>	<b>99.9</b>	<b>99.9</b>	<b>1.5</b>

Table 2. Results on medical datasets, at both image size  $64 \times 64$  and  $128 \times 128$ . Values for each dataset represent the AUC when considering all other datasets at the same image dimension as OOD datasets.

over using just the likelihood, and sometimes degrades performance.

All three AE-based reconstruction methods have inconsistent performance. They perform well on SVHN vs CIFAR10/CelebA but fail on other pairings, such as CIFAR10 vs SVHN. Inspection of the results reveals this class of method tends to perform better when the in-distribution datasets have simpler textures than the OOD datasets, as this causes higher OOD reconstruction error - this explains the good performance on SVHN vs CIFAR10/CelebA. Conversely, when in-distribution datasets are more complicated, the models can reconstruct simpler OOD datasets with low error and have poor OOD performance.

### 4.3. Results on medical datasets

Results are shown in Table 2, with some sample reconstructions in Supplementary. We excluded WAIC from this

comparison as Table 1 shows no benefit over simple likelihood whilst it requires training an ensemble of five models. The advantage of the DDPM over the other methods is more pronounced in these comparisons than in the computer vision datasets: it is the highest ranked method and also performs the most consistently, with 99.3 its lowest performance for any dataset pairing. A possible explanation for this is that at higher image sizes the noise level at which the model transitions from reconstructing to sampling is higher, so in-distribution data is successfully reconstructed in a greater number of the total reconstructions, resulting in a ‘cleaner’ in-distribution signal. We consider it an advantage of the method that it performs well on  $32 \times 32$  images where the reconstructions at higher noise values don’t contribute to the OOD signal, whilst also being capable of using this information at higher image sizes - this shows the method does not require any dataset specific tuning.

#### 4.4. Bottleneck size

To explore the importance of bottleneck size for reconstruction-based models, we considered the effect of changing the bottleneck on OOD performance. For the AE models we tried decreasing the size of the latent space, and both increasing and decreasing the memory size for MemAE models. For AnoDDPM-Mod we changed the  $t$  value the image was noised to before reconstruction. Results are shown in Supplementary. They show that for every method there was no single setting that worked well for all dataset pairings. For example, a 4-layer MemAE with a memory size of 2000 achieved 91.5 on FashionMNIST vs MNIST, substantially higher than the 3-layer model with recommended memory size, but this architecture’s performance on SVHN vs CIFAR10 dropped to 89.3, making it the poorest-performing of any model (not just MemAE models) on this pairing. These results highlight the need for undesirable dataset-specific tuning for standard reconstruction-based OOD detection, and underscore the value of our method that automatically includes results from a range of noise bottlenecks, avoiding the need for fine-tuning.

#### 4.5. Performance and number of reconstructions

We evaluated the performance of the DDPM model as the number of reconstructions performed was reduced on the computer vision datasets. Results are shown in Fig. 3 with a full breakdown by dataset pairing included in Supplementary. Reconstruction starting points were uniformly sub-sampled across the range of possible starting points; e.g. to reduce from 100 to 25 reconstructions we selected every fourth starting point. The results show that we can reduce the number of reconstructions from 100 to 25 with a mean drop in AUC across experiments of 0.58 (with AUC scores reported in the range of 0-100). Performing 25 reconstructions requires 1225 model evaluations, only slightly more than required to obtain a single sample from the model using the DDPM sampler over 1000 steps. Dropping even further to just 13 reconstructions gives a mean drop in AUC of 1.49, still outperforming the state-of-the-art DoSE method on 13/15 dataset pairings, but requiring just 637 evaluations, i.e. slightly over half those required to obtain a single sample from the model. Reducing the number of reconstructions below this point begins to affect performance substantially.

We further explored reducing the maximum value of  $t$  that reconstructions were performed from,  $\max_T$ , motivated by the observation in Fig. 2 that for the computer vision datasets the reconstructions from higher  $t$ -values do not seem heavily conditioned on the input and might not add to the OOD signal. Fig. 4 shows the mean drop in dataset performance, compared to the original model reported in Table 1, as  $\max_T$  was reduced from 1000, and as the number of reconstructions was reduced for a given value of  $\max_T$ . The results show that we can do better on the model-evaluations vs per-

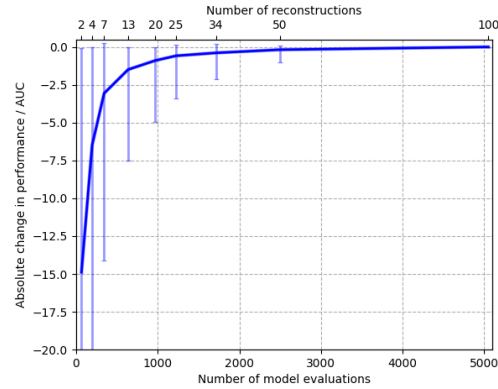


Figure 3. Results show the mean, minimum and maximum drop in absolute AUC as the number of reconstructions used to perform OOD detection is reduced, compared to results obtained using 100 reconstructions. The number of model evaluations is plotted on the lower x-axis, as this quantity directly correlates with total inference time. Full results are tabulated in the supplement.

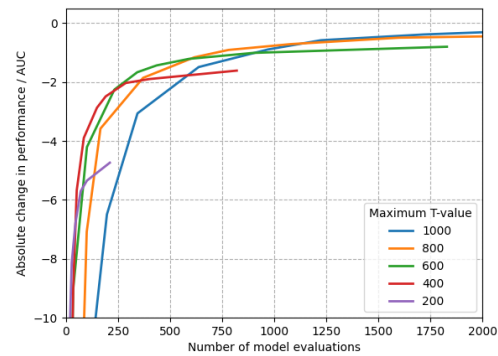


Figure 4. Results show the mean drop in absolute AUC, relative to the model with  $\max_T = 1000$  and 100 reconstructions, as both  $\max_T$  and the number of reconstructions performed for a given  $\max_T$  are reduced.

formance curve by reducing  $\max_T$ , meaning for a limited compute budget it is better to focus on reconstructions from lower  $t$ -values. However, this result is likely dependent on the size of the input. There is a relationship between the amount of noise applied to an image and the size of features that are detectable in that image [5], and our medical image experiments suggest that as input size increases, the maximum value of  $t$  that provides useful OOD signal will also increase. Reducing  $\max_T$  did not lead to mean performance better than the model with  $\max_T = 1000$  and 100 reconstructions, suggesting our approach is robust to the inclusion of non-informative signal and that there is no need for dataset-specific tuning of the bottlenecks. Detailed results are tabulated in Supplementary.

	FashionMNIST			CIFAR10			CelebA			SVHN					
	MNIST	VFlip	HFlip	SVHN	CelebA	VFlip	HFlip	CIFAR10	SVHN	VFlip	HFlip	CIFAR10	CelebA	VFlip	HFlip
Metric: MSE+LPIPS															
Classification method:															
GMM	95.1	82.3	60.8	96.8	55.8	53.1	49.8	98.3	100	88.9	50.1	97.6	98.6	52.3	55.1
SVM	95.6	82.5	60.7	96.4	56.8	52.8	49.9	97.5	100	87.6	50.0	97.8	99.2	52.6	55.1
Z-score average	97.4	88.6	65.1	97.9	68.5	63.2	50.5	99.0	100	93.3	50.3	99.0	99.6	58.2	61.6
Classification method: Z-score average															
Metric:															
MSE	92.4	83.6	63.8	40.2	70.3	57.2	50.5	86.3	77.5	86.5	50.2	97.3	99.7	58.8	60.8
LPIPS	95.4	81.5	61.6	98.5	50.5	58.6	50.3	98.3	100	87.6	50.2	92.3	87.5	53.4	56.0
MSE+LPIPS	97.4	88.6	65.1	97.9	68.5	63.2	50.5	99.0	100	93.3	50.3	99.0	99.6	58.2	61.6

Table 3. AUC score for variants of the DDPM method. In the top set of results, the similarity metrics used are fixed to MSE+LPIPS, and we try three classification methods: a GMM, SVM, and a Z-score average. In the bottom set, the classification method is fixed to Z-score averaging, and we use three different similarity metrics: MSE, LPIPS, and MSE+LPIPS.

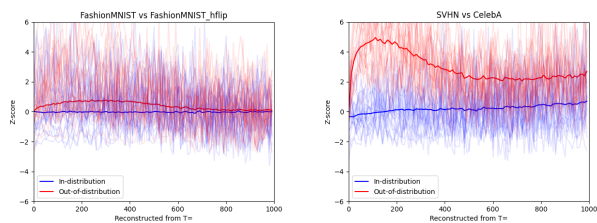


Figure 5. Plots of Z-scores of similarity for all 100 reconstructions of an image. Thin lines represent plots from individual inputs, thick lines show the means of these values. For FashionMNIST vs FashionMNIST HFlip (left), Z-score averaging performs much better than an SVM or GMM, whilst for SVHN vs CelebA (right), the difference in performance is small.

#### 4.6. Variants of the proposed method

We investigated the effects of changing the method used to classify OOD samples. We first tried using a one-class SVM or a Gaussian Mixture Model (GMM) to score samples. Results in Table 3 show that, while all methods perform well, the GMM and SVM are consistently outperformed by the simple Z-score averaging method. Fig. 5 shows the Z-scores for a dataset pairing where Z-score averaging substantially outperforms an SVM or GMM and a pairing where the difference is small. These plots suggest that Z-score averaging has an advantage when the signals from in-distribution and OOD inputs are noisy and overlap substantially, suggesting that the GMM/SVM overfit to these noisy signals, impairing classification performance. Table 3 also shows that using MSE or LPIPS alone can be unreliable; both similarity measures have some dataset pairings they give poor performance on. Combining them is nearly always better than using either alone and provides robust performance above all dataset pairings; suggesting the two metrics are complementary.

#### 4.7. Limitations

A key limitation of the method is that it is more computationally expensive than competing methods. This might limit the contexts for which the method can be used, even with the potential efficiency enhancements discussed in Sec. 4.5. For example, the method may be suitable in medical imaging applications where it may be acceptable to flag an input as OOD within minutes, but less suitable in self-driving cars where predictions need to be made in real time. However, there is considerable scope for improving the efficiency of these models. The field of DDPMs is currently very active with substantial interest in fast sampling, and our method can directly benefit from advances in sampling speed [1, 6, 14, 34, 37–39, 42, 46]. Finally, while the steps along a single reconstruction chain cannot be parallelised, reconstructions from different starting points can be. This could reduce the reconstruction time down to the amount required to reconstruct the longest chain.

#### 5. Conclusion

In this work, we explored how DDPMs can be used to perform unsupervised OOD detection. We propose performing reconstructions of a number of inputs noised to different extents, addressing a drawback of standard reconstruction techniques that require the choice of a single, fixed bottleneck. We tested our method on both computer vision benchmarks and medical imaging datasets with a higher image size. These experiments show that our DDPM-based method outperforms not only reconstruction-based methods but also state-of-the-art generative approaches.

**Acknowledgements** MG, WS, PN, SO, and MJC are supported by a grant from the Wellcome Trust (WT213038/Z/18/Z).



## References

- [1] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In *International Conference on Learning Representations*, 2021. 8
- [2] Christopher M Bishop. Novelty detection and neural network validation. *IEE Proceedings-Vision, Image and Signal processing*, 141(4):217–222, 1994. 2, 4, 5, 6
- [3] Hyunsun Choi, Eric Jang, and Alexander A Alemi. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018. 1, 2, 4, 5
- [4] 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. 2, 3, 4, 5, 6
- [5] Sander Dieleman. Diffusion models are autoencoders, 2022. 2, 7
- [6] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Genie: Higher-order denoising diffusion solvers. *arXiv preprint arXiv:2210.05475*, 2022. 8
- [7] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1705–1714, 2019. 3, 4, 5, 6
- [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. 2
- [9] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2018. 1, 2
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 5
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [13] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 4
- [14] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021. 8
- [15] Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014. 3
- [16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [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. 4
- [18] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations*, 2021. 3
- [19] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 4
- [20] Zhenzhen Liu, Jin Peng Zhou, Yufan Wang, and Kilian Q Weinberger. Unsupervised out-of-distribution detection with diffusion inpainting. *arXiv preprint arXiv:2302.10326*, 2023. 3
- [21] Olga Lyudchik. Outlier detection using autoencoders. Technical report, 2016. 2
- [22] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. National Institute of Science of India, 1936. 3
- [23] 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. 2, 4, 5, 6
- [24] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don’t know? In *International Conference on Learning Representations*, 2018. 1, 2, 4, 5, 6
- [25] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, and Balaji Lakshminarayanan. Detecting out-of-distribution inputs to deep generative models using typicality. *arXiv preprint arXiv:1906.02994*, 2019. 2, 5
- [26] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 4
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. 4
- [28] Marco AF Pimentel, David A Clifton, Lei Clifton, and Lionel Tarassenko. A review of novelty detection. *Signal processing*, 99:215–249, 2014. 2
- [29] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017. 4
- [30] 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. 2, 5
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3, 4
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3
- [33] Mohammadreza Salehi, Hossein Mirzaei, Dan Hendrycks, Yixuan Li, Mohammad Hossein Rohban, and Mohammad Sabokrou. A unified survey on anomaly, novelty, open-set, and out-of-distribution detection: Solutions and future challenges. *arXiv preprint arXiv:2110.14051*, 2021. 2
- [34] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2021. 8
- [35] 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. In *International Conference on Learning Representations*, 2019. 2
- [36] 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. 2
- [37] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 8
- [38] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems*, 34:11287–11302, 2021. 8
- [39] Daniel Watson, William Chan, Jonathan Ho, and Mohammad Norouzi. Learning fast samplers for diffusion models by differentiating through sample quality. In *International Conference on Learning Representations*, 2021. 8
- [40] Julian Wyatt, Adam Leach, Sebastian M Schmon, and Chris G Willcocks. Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 650–656, 2022. 3, 4, 5, 6
- [41] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 4
- [42] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. In *International Conference on Learning Representations*, 2021. 8
- [43] 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. 2
- [44] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021. 2
- [45] Lily Zhang, Mark Goldstein, and Rajesh Ranganath. Understanding failures in out-of-distribution detection with deep generative models. In *International Conference on Machine Learning*, pages 12427–12436. PMLR, 2021. 2
- [46] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022. 8
- [47] 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. 4
- [48] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International conference on learning representations*, 2018. 2