



HAL
open science

Detecting Out-Of-Distribution Earth Observation Images with Diffusion Models

Georges Le Bellier, Nicolas Audebert

► **To cite this version:**

Georges Le Bellier, Nicolas Audebert. Detecting Out-Of-Distribution Earth Observation Images with Diffusion Models. EARTHVISION 2024 IEEE/CVF CVPR Workshop. Large Scale Computer Vision for Remote Sensing Imagery, Jun 2024, Seattle, United States. hal-04551408

HAL Id: hal-04551408

<https://hal.science/hal-04551408>

Submitted on 18 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Detecting Out-Of-Distribution Earth Observation Images with Diffusion Models

Georges Le Bellier¹

¹ Cnam, CEDRIC, EA4629
F-75141 Paris, France

georges.le-bellier@lecnam.net

Nicolas Audebert^{1,2}

² Univ. Gustave Eiffel, ENSG, IGN, LASTIG
F-94160 Saint-Mandé, France

nicolas.audebert@ign.fr

Abstract

Earth Observation imagery can capture rare and unusual events, such as disasters and major landscape changes, whose visual appearance contrasts with the usual observations. Deep models trained on common remote sensing data will output drastically different features for these out-of-distribution samples, compared to those closer to their training dataset. Detecting them could therefore help anticipate changes in the observations, either geographical or environmental. In this work, we show that the reconstruction error of diffusion models can effectively serve as unsupervised out-of-distribution detectors for remote sensing images, using them as a plausibility score. Moreover, we introduce ODEED, a novel reconstruction-based scorer using the probability-flow ODE of diffusion models. We validate it experimentally on SpaceNet 8 with various scenarios, such as classical OOD detection with geographical shift and near-OOD setups: pre/post-flood and non-flooded/flooded image recognition. We show that our ODEED scorer significantly outperforms other diffusion-based and discriminative baselines on the more challenging near-OOD scenarios of flood image detection, where OOD images are close to the distribution tail. We aim to pave the way towards better use of generative models for anomaly detection in remote sensing.

1. Introduction

In recent years, deep learning has grown to be a staple of image understanding in computer vision, Earth Observation included. Deep neural networks have been for several years now the state of the art for many tasks, from biomass estimation to land cover segmentation [64]. Despite their impressive capacity for generalization, these models are still trained on finite datasets. For predictive models, it is critical to be able to detect when new observations fall outside of this training set, as model performance can drop sharply. This task called *Out-Of-Distribution* (OOD) detection [58], is a specific case of anomaly detection and represents a major

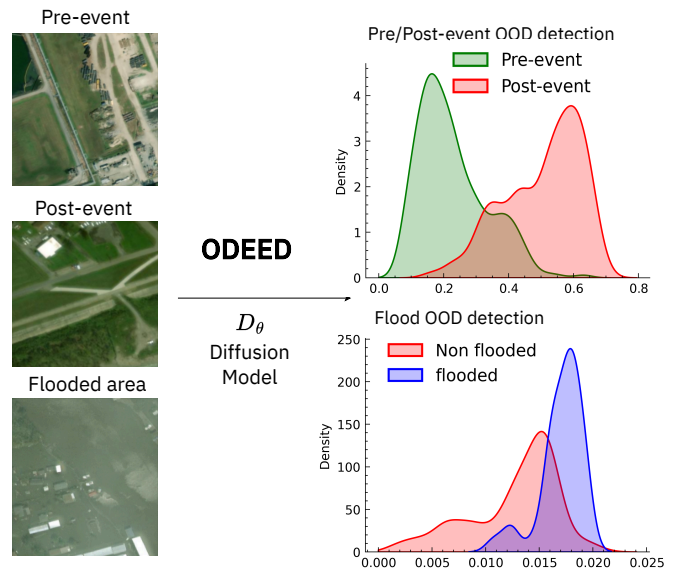


Figure 1. ODEED discriminates between pre and post-event images by making larger reconstruction errors on the latter w.r.t. the LPIPS metric (top). It can also isolate flooded areas from the other post-event images when evaluating the reconstruction similarity with the MSE (bottom).

challenge in improving the robustness of deep models.

Indeed, remote sensing imagery is facing a multitude of possible distribution shifts, from a change of acquisition sensor to a change of season to a change of geographical area. Being able to detect these OOD images is helpful to avoid considering degraded model predictions. In addition, disasters such as floods, forest fires, and storms are also unfrequent and catastrophic events, that are rarely observed in remote sensing datasets. While technically not always OOD, these events are “near” out-of-distribution observations as shown in Fig. 1, whose detection can be of great help for disaster management and surveillance. Identifying unusual observations can also help automate the curation of new datasets by eliminating outliers caused by cloud cover, sensor faults, or other artifacts.

Moreover, annotations are scarce in Earth Observation. The huge mass of available images remains largely unlabeled, due to the expert knowledge and time required to do so. For this reason, OOD detectors tailored to remote sensing data should be fully unsupervised, and be able to model the data distribution without labels. Generative models have been advanced as a way to model the distribution of the data and build likelihood estimators to detect OOD samples [60]. In particular, recent works in classical computer vision have established *diffusion models* as the new state-of-the-art to detect out-of-distribution images [17]. Diffusion models [51] are a family of unsupervised generative models that model the underlying distribution of training examples. They have outclassed previous models in image generation, such as Generative Adversarial Networks, due to their strong ability to model complex distributions without mode collapse. Promising new works have shown that diffusion models were applicable to Earth Observation imagery, *e.g.* for cloud removal [46] or image synthesis [29].

In this paper, we demonstrate the relevance of unconditional diffusion models for OOD detection in remote sensing. Several use cases are considered in our work: cloud detection, different geographical domains, pre- and post-disaster, and flooded/non-flooded images (cf. Fig. 1). We introduce a new ODEED specifically tailored to the latter use cases and show it outperforms all existing baselines.

2. Related Work

2.1. Diffusion models

Diffusion models are a powerful family of generative models that take inspiration from *score-matching* [22, 50, 54]. They were first derived in the discrete-time formulation as Markov chains [20, 49] before the introduction of continuous-time diffusion models building upon stochastic and ordinary differential equations [25, 52]. As generative models, it has been observed that denoising diffusion models produce high-quality samples both in unconditional [25, 52] and conditional image generation [12, 26, 43, 45] and outperform previous approaches, notably *Generative Adversarial Networks* (GANs) [12] as they do not exhibit mode collapse. Diffusion models proved their synthesis capabilities not only for images but also for video [7, 21], audio [31, 32, 44], text [3, 34], and protein design [56, 59]. In addition to the synthesis performances, diffusion models are now used for a variety of tasks, from image segmentation [1, 28] to inverse problems solving [27].

In remote sensing, the application of diffusion models to Earth Observation is relatively new. RSDiff [47] trains a cascade of a low-resolution diffusion model that feeds into a text-conditioned super-resolution diffusion model that synthesizes remote sensing images based on textual descriptions. SatDM [4] uses a conditional DDPM, conditioned

on building footprints, to generate new labeled samples to train deep networks for building extractions. Espinosa and Crowley [13] extends this idea to a DDPM conditioned on semantic maps using ControlNet [61] to generate new synthetic aerial imagery of Scotland. While generating fake remote sensing imagery has limited practical applications, diffusion models have been adapted to solve classical EO tasks, such as change detection and cloud removal [46] and super-resolution [19, 35, 57]. In this work, we will leverage the modeling ability of diffusion models not to generate new data, but to estimate the *plausibility* of EO images and identify unusual observations, *i.e.* find images that are out of the distribution of “usual” acquisitions.

2.2. Out-of-distribution detection

Out-of-distribution (OOD) detection is a special case of anomaly detection. Early works focused on detecting OOD samples that fell outside the training dataset of machine learning models, to prevent failure in the predictions. The first approaches required a dataset of “unusual” samples in order to train OOD detectors in a supervised manner [33, 39]. These approaches restrict the definition of the OOD samples since we are not considering all possible OOD outside the distribution of interest [8, 55]. Recent approaches are *post-hoc*, *i.e.* they rely on features extracted from pre-trained models. In this way, we avoid retraining neural networks for the specific task of OOD detection and benefit from existing pretrained features. In general, the goal is to assign a *score* to an observation. If it is greater than a threshold, then the image is considered in-distribution, and OOD otherwise.

Numerous OOD scorers assume that there is a discriminative model, *e.g.* a classifier or a segmenter, that has been trained on some dataset. *Post-hoc* scorers try to discriminate between out-of-distribution and in-distribution samples based only on the predictions of a classifier or segmentation model. They assess the model certainty on its prediction and samples whose predictions have high uncertainty scores are considered OOD [36, 38]. However, there has recently been a surge in fully unsupervised OOD scorers, *i.e.* scorers that do not depend on a discriminative model. These scorers can therefore be used even on unlabeled datasets. For example, some post-hoc scorers have been introduced to work on the features learnt by self-supervised models [40, 48].

Similarly, reconstruction-based OOD scorers rely on generative models to reproduce an image given a corrupted version of it. The corruption can be *e.g.* downsampling for *Autoencoders* (AE) and *Variational Autoencoders* (VAE) [11, 63, 65] or noise for diffusion models [15, 18, 37]. In principle, the generative model has learned the training distribution and should have lower reconstruction errors on in-distribution samples (ID) than on OOD samples. The reconstruction error can be evaluated with different similarity metrics depending on the nature of the samples. For

images, common metrics are the Mean-Squared Error (MSE) and the Learned Perceptual Image Patch Similarity (LPIPS), the latter being more aligned with human perception [62]. As diffusion models proved robust reconstruction faculty, several OOD scorers employ them in the context of images denoising [15, 18] or inpainting [37]. In addition, as diffusion models have a variational interpretation [30], Goodier and Campbell [17] proposed an OOD scorer based on diffusion models’ ELBO, linking the reconstruction error to the statistical likelihood.

In remote sensing, OOD detection has been a somewhat niche topic. To the best of our knowledge, the only works on detecting OOD remote sensing images are those from Gawlikowski et al. [16] and Coca et al. [10]. Gawlikowski et al. [16] introduced an OOD scorer based on a Dirichlet Prior Network linked to a classification model trained on remote sensing images to detect a shift in the classes, in sensor characteristics, or in geographical areas that would reduce classification accuracy. Coca et al. [10] later adapted this work to a self-supervised model to detect satellite images containing burned areas. As these images were rare in the training dataset, they were able to characterize them as OOD. We will show that this approach is a promising avenue to detect unusual events in remote sensing imagery.

3. Methods

In this work, we investigate diffusion models for OOD detection and focus on three diffusion-based scorers:

- **Diffusion loss scorer** based on the time-truncated diffusion losses which average reconstruction errors;
- **One-step denoising scorer** focusing on the denoising performances at fixed timestep;
- **ODEED (ODE Encoding Decoding) scorer**, a new OOD scorer leveraging PF-ODE trajectory accuracy as a way to discriminate between in and out-of-distribution samples.

3.1. Background on Diffusion Models

Continuous-time diffusion models generalize discrete-time diffusion models to the infinite timesteps case. They do not rely on Markov chains but on *Stochastic Differential Equations* (SDEs) instead, both for the diffusion and the denoising processes. Let us consider a diffusion process $\{x_t\}_{t \in [0, T]}$, with fixed boundary conditions $x_0 \sim p_0(x) = p_{\text{data}}(x)$ and $x_T \sim p_T(x) = p_{\text{prior}}(x)$, and $p_t(x_t)$ denote the marginal density at time t . This diffusion process is a solution to the following SDE:

$$dx_t = f(x_t, t)dt + g(t)dw_t \quad (1)$$

where $f : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is the *drift*, $g : [0, T] \rightarrow \mathbb{R}_+^*$ is the *diffusion* coefficient and w_t is the standard Wiener process (*i.e.* Brownian motion). This diffusion process corrupts the original image x_0 into a Gaussian noise x_T . We know

from [2] that the forward SDE in Eq. (1) admits a reverse SDE mapping p_T to p_0 :

$$dx_t = [f(x_t, t) - g(t)^2 \nabla_x \log p_t(x_t)] dt + g(t)d\bar{w}_t \quad (2)$$

Moreover, these diffusion processes share the property of admitting a *deterministic* process preserving the same marginal distributions p_t for all $t \in [0, T]$. It is described by the *Probability Flow Ordinary Differential Equation* (PF-ODE):

$$dx_t = \left[f(x_t, t) - \frac{g(t)^2}{2} \nabla_x \log p_t(x_t) \right] dt \quad (3)$$

Then, the ODE allows encoding samples from the data distribution into the prior distribution by integrating Eq. (3) from $t = 0$ to $t = T$ and vice-versa.

Both the backward SDE Eq. (2) and the PF-ODE Eq. (3) equations include the *score function* $\nabla_x \log p_t$ that has no known close form in practical cases. In this paper, we are embracing the diffusion framework proposed by [25], in which the score function derives from a perfect denoiser D :

$$\nabla_x \log p_t(x) = \frac{1}{t^2} (D(x, t) - x) \quad (4)$$

We approximate this denoiser with a neural network D_θ optimized to reduce the reconstruction loss at every time t :

$$\mathbb{E}_{x_0 \sim p_0} \mathbb{E}_{\epsilon \sim \mathcal{N}(0, t^2 I)} [\lambda(t) \|D_\theta(x_0 + \epsilon) - x_0\|_2^2] \quad (5)$$

3.2. Detecting OOD with diffusion models

In addition to the generation abilities of continuous-time diffusion models, they can be used in the context of out-of-distribution detection. We leverage three scorers using diffusion models, based on their reconstruction performances.

First, we want to evaluate the effectiveness of the loss function as an OOD scorer. The intuition behind this is that diffusion models should generalize to in-distribution samples unseen during train but fail on OOD samples, resulting in a greater loss. Since in our case, both in and out-of-distribution are satellite images, we argue that the differences in reconstruction can be observed at small times. Indeed, diffusion models generate images by first placing the overall structure (large times) and then fill-in the details (small times) [6, 42]. Thus, we compute the reconstruction error Eq. (5) for several times below a given threshold t_0 . In other words, for a given image, we sample N noise corrupted versions of it with increasing times t_1, \dots, t_N , then we estimate the N reconstructions with the learned diffusion model D_θ and compute the weighted average loss:

$$S_{t_0}^{\text{loss}}(x_0) = -\mathbb{E}_{t, \epsilon} [\lambda(t) \|D_\theta(x_0 + \epsilon) - x_0\|_2^2] \quad (6)$$

where $t \sim \mathcal{U}(0, t_0)$ and $\epsilon \sim \mathcal{N}(0, t_0^2 I)$. We experiment with two different weighting functions $\lambda(t)$: the training weighting function [25] and a linear weighting as in [17].

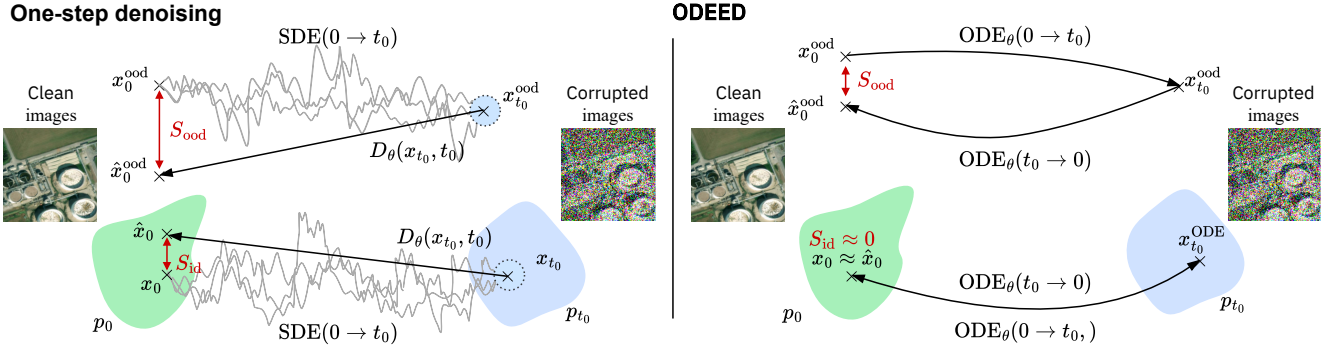


Figure 2. **Illustration of the one-step denoising and ODEED scorers.** (Left) The one-step denoiser samples multiple corrupted versions of the original image $x_{t_0} \sim p_{t_0}(x|x_0)$ thanks to the forward SDE and then evaluates similarity scores on the one-step reconstructions made with the diffusion model D_θ . (Right) The ODEED scorer encodes the initial image into a unique latent x_{t_0} with the PF-ODE estimated with D_θ and then decodes the latent. For the true PF-ODE, in-distribution samples’ reconstruction is perfect.

Second, as the learned diffusion model D_θ is a denoiser conditioned on time, we leverage its reconstruction performances in *one-step denoising*, as shown in Fig. 2 (left). For several values to t_0 , we investigate the ability of the model to reconstruct an initial image x_0 from a noised version $x_t \sim p_t(x_t|x_0)$. To do so, we sample several noises and average a reconstruction error across the samples, considering a single denoising step:

$$S_{t_0}^{\text{denoiser}}(x_0) = -\mathbb{E}_{\epsilon \sim \mathcal{N}(0, t_0^2 I)} [d(D_\theta(x_0 + \epsilon) - x_0)] \quad (7)$$

where $d(\cdot, \cdot)$ is a similarity measure, *e.g.* MSE or LPIPS. In practice, we observe that the denoiser D_θ reconstructs well from noised versions of x_0 for small values of t_0 .

3.3. ODEED scorer for OOD detection

Due to the stochastic nature of the diffusion process using the forward SDE Eq. (1), proceeding to several integrations of the forward SDE from an initial image x_0 to time t_0 produces multiple $x_{t_0} \sim p_{t_0}(x_{t_0}|x_0)$. It is also true that reversing the diffusion process from a latent x_T with the backward SDE Eq. (2) generates several denoised images. This means that stochastic samplers lead to various reconstructions that can be more or less distant from the initial image, depending on the level of stochasticity used, introducing estimation errors in the reconstruction scores.

In contrast, integrating the PF-ODE Eq. (3) from 0 to t_0 encodes the initial image x_0 into a *unique* noisy latent x_{t_0} . Conversely, integrating it from t_0 to 0 allows us to get back to the exact initial image. In other words, the trajectory is fully deterministic. This property is fundamental in the context of reconstruction-based scorers, as diffusion models can propose exact reconstructions of the initial images.

Based on this observation, we introduce **ODEED** a novel reconstruction-based OOD scorer based on the PF-ODE trajectories (Fig. 2 right). To do so, we approximate the true

score function in Eq. (3) with the diffusion model D_θ following Eq. (4). Let $\text{ODE}_\theta(\cdot, t_1 \rightarrow t_2)$ denote the operation of integrating the D_θ -approximated ODE from t_1 to t_2 with a numerical solver. Our scorer operates in two phases:

1. *Encoding*: we encode the clean image x_0 into a unique noisy latent $x_{t_0}^{\text{ODE}}$ by integrating the ODE from the starting time $t = 0$ to the intermediate time $t = t_0$.

$$x_{t_0}^{\text{ODE}} = \text{ODE}_\theta(x_0, 0 \rightarrow t_0) \quad (8)$$

2. *Decoding*: we solve the ODE with decreasing times from t_0 to 0, *i.e.* denoising the latent to a clean image that is the reconstruction \hat{x}_0^{ODE} of the original sample.

$$\hat{x}_0^{\text{ODE}} = \text{ODE}_\theta(x_{t_0}^{\text{ODE}}, t_0 \rightarrow 0) \quad (9)$$

Then, we define the OOD scorer from the similarity between the reconstructed and the original images.

$$S_{t_0}^{\text{ODE}}(x_0) = -d(x_0, \hat{x}_0^{\text{ODE}}) \quad (10)$$

Once again, this difference measures how well the diffusion model can reconstruct the clean observation x_0 . Compared to previous scorers, it is entirely *deterministic* and therefore should better represent whether x_0 belongs to the modeled distribution, without being influenced by the stochastic trajectories.

4. Experiments

4.1. Toy Problem: Cloud Cover Detection

To test the relevance of our approaches, we experiment with the simple task of discriminating between cloud-free images and cloudy images. Cloud detection can be achieved with well-designed features and deep learning techniques. Nevertheless, it is an easy use case on which we want to test the reconstruction-based approaches with diffusion models.



Figure 3. The three scenarios derived from the SpaceNet 8 dataset. The first two setups focus on the impact of floodings while the third one centers around geographical domain shift.

Method	AUC \uparrow	FPR95% \downarrow	R^2 \uparrow
Image mean	88.3	41.6	58.4
Diffusion MSE	<u>85.1</u>	29.1	<u>49.4</u>
Diffusion LPIPS	84.3	<u>29.2</u>	37.6

Table 1. Results for OOD detection on the Sentinel-2 Cloud Mask Catalogue where OOD samples are images with cloud coverage above 10%. Diffusion models-based approaches are competitive with the baseline and show good results for both AUC and FPR95%

Experimental setup For this task, we employ the pre-trained discrete-time diffusion model from [46], trained on 8000 64×64 pixels RGB cloud-free images from the Sentinel-2 Cloud Mask Catalogue dataset [14]. Cloud-free images are in-distribution samples and images with cloud coverage above 10% are OOD. Since the model was only trained on cloud-free images, we expect it to fail to reconstruct images with clouds. As the model is discrete time, we first consider only the one-step denoising scorers. We also use the mean pixel intensity as a naive scorer baseline.

4.2. SpaceNet 8

We then extend our experiments to more challenging applications by considering the SpaceNet 8 dataset [23].

Dataset SpaceNet 8 is a dataset of paired pre- and post-flooding images from East Louisiana and Germany. It is composed of 3-band RGB images from Maxar satellites. It includes segmentation masks for both the roads and the buildings while distinguishing between flooded and non-flooded objects. We downsample the original images by a factor of 2 to produce 256×256 patches and apply overlap on the training set. We obtain 5864 pre and post-event pairs for the Germany subset and 17660 pairs for Louisiana.

Experimental setup Using the images and annotations, we derive three different scenarios of OOD detection on

SpaceNet 8, each using a different definition of OOD. This allows us to test the abilities of different OOD scorers to detect various types of “out-of-distribution”, shown in Fig. 3:

1. **Pre-flood vs Post-flood images:** in this scenario, the OOD detector aims at discriminating between images prior to the flood (in-distribution samples) and images after the flood (out-of-distribution) within the same geographical domain (Germany or East-Louisiana).
2. **Flooded vs non-flooded areas images:** in this scenario, the OOD detector should discriminate between images with no flood (in-distribution) and images with a visible flood (out-of-distribution). In-distribution images include post-event images that do not have a visible flood. The labeling is done based on the segmentation masks: an image is considered OOD if it contains pixels belonging to the “flooded building” or “flooded road” class in the ground truth.¹
3. **Domain OOD:** in this case, we want to detect a geographical shift, *i.e.* discriminating images taken in East Louisiana from images taken over Germany. We consider only pre-event images in this setting.

We want to highlight that in all scenarios, the diffusion models are trained on *pre-event images only* and without any annotations, nor post-event images.

Baselines. We compare the diffusion models and ODEED with several well-established OOD detection techniques.

As a reconstruction-based baseline, we train an autoencoder on pre-event images for each geographical domain. We use either the MSE or the LPIPS distance on the reconstruction in image space as a score. We also consider the Mahalanobis distance D_M in the latent space Eq. (11) and estimate the covariance matrix Σ and the mean μ on the train set [5, 11]. The final score is then a linear combination of the latent space and image space distances:

$$S^M = - \left[\alpha D_M(E(x)) + \beta \|x - \hat{x}\|_2 \right] \quad (11)$$

$$\text{where } D_M(z) = \sqrt{(z - \mu)\Sigma^{-1}(z - \mu)} \quad (12)$$

For exhaustiveness, we also compare our models to OOD scorers-based discriminative models, even though these models require to be trained on an annotated dataset in the first place. We train a segmentation model f_θ on SpaceNet 8 for each pre-event domain to predict the classes “background”, “roads” and “buildings”. We then extract OOD scores from this model using standard approaches from the literature:

- Maximum Class Probability (MCP) [24, 41] simply consider that if the model is confident about a prediction then

¹Note that this proxy is perfectible and has some failure cases (*i.e.* flooded areas with no roads or buildings are considered “non-flooded”), which might inflate the number of false positives. FPR results for this setup should therefore be interpreted as a higher bound.

	Method	Pre-flood/Post-flood		Non-flooded/flooded		Domain OOD	
		Germany	Louisiana	Germany	Louisiana	Germany	Louisiana
Discriminative Segmentation	MPC	52.6/98.9	41.1/98.4	67.4/80.0	61.0/94.3	69.8/93.4	47.7/95.5
	Neg-Entropy	59.6/97.7	42.8/95.9	64.3/80/0	60.7/97.1	76.3/92.2	48.8/94.3
	DeepKNN (k=5)	51.9/80.7	76.4/38.9	70.1/86.7	54.3/100	93.6/28.7	80.9/38.6
	Energy Logits	67.9/88.6	67.9/88.6	56.4/80.0	60.9/97.1	84.7/70.9	50.4/92.0
Generative Diffusion Loss	Training	50.4/84.1	52.3/88.5	59.7/100	71.0/71.4	67.0/86.0	55.9/98.8
	Linear	49.2/85.2	53.2/88.5	70.0/100	75.7/71.4	60.4/83.1	57.0/98.0
Reconstruction based							
Autoencoder	MSE	28.4/96.6	26.3/95.1	57.9/100	68.5/85.7	84.9/79.9	28.3/96.6
	LPIPS	21.1/97.7	27.5/96.3	55.6/100	69.5/77.1	75.2/54.9	41.4/100
	Mahalanobis	48.9/95.1	30.8/97.1	51.4/94.9	72.5/90.6	49.6/95.0	52.0/94.8
1-step denoising	MSE	28.8/81.8	42.2/73.3	68.5/100	73.4/77.1	86.3/20.9	33.5/97.7
	LPIPS	<u>74.5/59.1</u>	<u>90.9/35.7</u>	60.5/100	<u>76.8/65.7</u>	79.6/52.5	82.1/85.2
ODEED (Ours)	MSE	65.6/76.1	69.0/80.3	83.6/33.3	86.9/42.9	41.2/97.1	60.9/95.4
	LPIPS	87.9/20.5	94.5/24.6	<u>75.3/73.3</u>	64.1/85.7	54.3/97.5	68.3/70.4

Table 2. Results for OOD detection on the Germany and Louisiana subsets of SpaceNet 8 (AUC \uparrow / FPR95% \downarrow). ODEED in combination with the LPIPS metric yields the strongest results on the pre/post flood scenario for both domains. Paired with MSE, ODEED delivers the highest AUC and lowest FPR95% for the non-flooded/flooded scenario.

the maximum predicted class probability will be high.

$$S^{\text{MCP}}(x) = \max_k P(k = y|x) = \max_k f_\theta(x) \quad (13)$$

- Negative Entropy: if the model is unconfident about its prediction then the probability mass will be spread across all classes, resulting in a high entropy:

$$S^{\text{neg-H}}(x) = -H(p) = -\sum_{k=0}^{K-1} P(k = y|x) \log P(k = y|x) \quad (14)$$

- Energy-Logits (EL) [36] computes the *free-energy* on the logits distribution. We denote g_θ the sub-network of f_θ predicting the logits.

$$S^{\text{EL}}(x) = -E(x; g_\theta) = T \log \sum_{k=0}^{K-1} e^{g_\theta^{(k)}(x)/T} \quad (15)$$

- DeepKNN [53], which applies a k-nearest neighbors algorithm to the features extracted from the last layer.

Diffusion models Our diffusion models are trained based on Karras et al. [25], using the unconditional U-Net backbone from [52] with [64, 128, 128, 256, 256] channels in the downsampling blocks and the reverse for the upsampling blocks. Time information is sinusoidally embedded on 64 channels and then fed to a linear layer and SiLU activation. Models are trained for 40000 steps with a batch size of 80 (approximately 48 hours), a learning rate of $2e^{-4}$, and AdamW optimizer. To improve training stability, we use

gradient clipping and exponential moving average. We sample using the second-order Heun solver [25] with $L = 20$ discretization steps except in the case of small noise levels ($t_0 < 1$), where we integrate with 5 steps, as too many sampling steps hinder sampling quality. t_0 hyperparameter is tuned on a held-out validation set.

Autoencoders For a fair comparison, we derive the AE architecture from the U-Net backbone of the diffusion model by removing the time embedders and skip connections. We train the autoencoders for 20000 steps with a batch size of 80, a learning rate of $2e^{-4}$, and AdamW optimizer. We set $h = 512$ for the latent space dimension as it achieved the best reconstruction results on the validation set. Coefficients α and β of the Mahalanobis distance Eq. (11) are tuned on the validation set to match the reciprocal of the standard deviation of their respective distance [11].

Segmentation models We train DeepLabV3 segmentation models [9] for a maximum of 1400 steps with a batch size of 32 and use early stopping. We set the learning rate to $2e^{-4}$ and optimize the model’s parameters with AdamW.

5. Results

We evaluate our OOD detectors using two standard binary classification metrics, by varying the detection threshold: AUC (Area Under Curve) and FPR95% (False Positive Rate at Recall 0.95). For the latter, note that ID samples are considered as positives and OOD samples are negatives, *i.e.* FPR95% denote how many OOD samples are wrongly considered as in-distribution for a 95% recall of the ID samples.

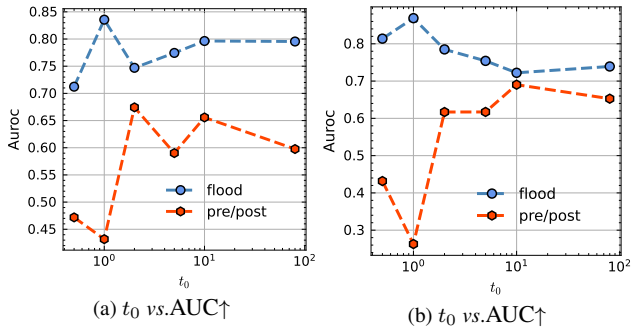


Figure 4. Impact of t_0 on AUC for ODEED + MSE metric on the geographical domains (a) Germany and (b) Louisiana.

5.1. Preliminary results: cloud cover

We report in Tab. 1 OOD detection results on the cloud cover dataset. We observe that the diffusion model’s reconstruction error is indeed higher on cloudy images, as shown by a high AUC using the diffusion + MSE score. This was expected as the model is trained only on cloud-free images. Diffusion-based methods produce competitive results in terms of AUC with respect to the naive baseline and achieve the best results in terms of FPR95% with a gain of 12.5 points for the diffusion MSE score against the naive baseline.

To evaluate the calibration of the scores, we perform a linear regression of the OOD score with the cloud coverage percentage. The naive baseline (mean pixel intensity) achieves the highest R^2 coefficient (more clouds implies more bright pixels). However, diffusion model scorers also correlate well with the cloud coverage, indicating that more cloudy images tend to be classified as “more OOD”. These first results show that diffusion-based scorers discriminate between in and out-of-distribution samples on the simple task of cloudy image detection. This suggests their effectiveness for use cases where no naive baselines exist.

5.2. Main results

We report in Tab. 2 the OOD scorer’s performances on the Germany and Louisiana subsets of SpaceNet 8. In the near-OOD setting (pre-flood/post-flood and non-flooded/flooded), AE exhibits poor OOD detection abilities, with AUC well under the 0.5 random baseline. Discriminative OOD detectors fare better but still misclassify a lot of OOD images as in-distribution, as shown by the high FPR95. In contrast, the best AUC and FPR95% results are achieved with our ODEED method for the pre/post-event and the non-flooded/flooded scenarios. On the pre/post-event OOD setup, the combination of ODEED + LPIPS outperforms the best baseline by a margin of 13.4 points on the Germany dataset and 3.6 pts on the Louisiana dataset in terms of AUC. We notice a gain of -38.6pts of FPR95% on the pre/post-event

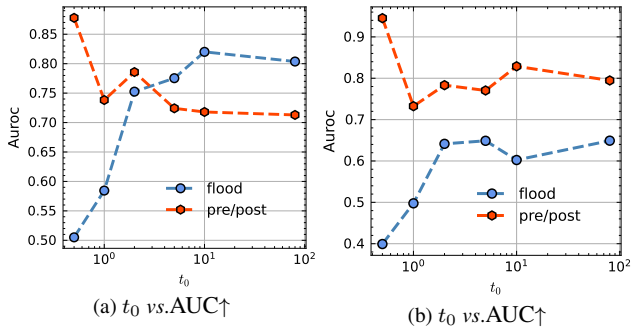


Figure 5. Impact of t_0 on AUC for ODEED + LPIPS metric on (a) Germany (b) Louisiana.

Germany experiment when using the ODEED + LPIPS and a gain of -11.1 for the same setup on the Louisiana dataset. Interestingly, while best results are obtained with the pair ODEED + LPIPS for the pre/post-event setup on both areas, the MSE brings the ODE-encoding to top performances on the non-flooded/flooded scenarios with 83.6 AUC and 33.3 FPR95% on Germany and 86.9 AUC and 42.9 FPR95% on Louisiana. We hypothesize that this difference is linked to the difference between ID and OOD samples in the two scenarios: in the first scenario the satellite images have different global aspects, whereas floods are locally spatialized elements of the same global context.

For all setups, reconstruction-based methods based on diffusion models outperform the autoencoder-based ones.

The largest performance gap reaches 39.0 (87.9 vs. 48.9) and 63.7 pts (94.5 vs. 30.8) in AUC on the pre/post-flood scenario for respectively Germany and Louisiana. The narrowest gap is set on the Germany/Louisiana domain OOD case where autoencoder-based methods are on par with diffusion models (respectively 85.6 and 86.3 in top AUC).

Yet, the ODEED scorer performs poorly on the domain change detection. We suspect that the model generalizes enough to reconstruct images from other domains for small noise levels. This idea is reinforced by the decent performances of the one-step denoisers in this scenario: best AUC (82.1) with LPIPS on the Louisiana dataset and second best on the Germany dataset (86.3) with MSE, while achieving the best FPR95% (20.9). DeepKNN yields top results in domain OOD on the Germany dataset and second best on the Louisiana domain. This could indicate that segmentation models learn features that are more domain-specific than the reconstruction models, and therefore generalize worse when switching geographical areas. Finally, Tab. 2 shows that diffusion model losses, independently of the weighting function used, are poor likelihood estimators in our settings. The various dynamics of the reconstruction errors may explain the poor OOD detection performance of diffusion losses which average the latter over a large time range.

Domain Model	Pre-flood/Post-flood				Non-flooded/flooded				
	Germany		Louisiana		Germany		Louisiana		
	D_θ G	D_θ L	D_θ L	D_θ G	D_θ G	D_θ L	D_θ L	D_θ G	
1-step denoising	MSE	28.8/81.8	26.2/100	42.2/73.3	23.2/98.0	68.5/100	65.6/100	73.4/77.1	71.9/74.3
	LPIPS	74.5/59.1	76.8/51.1	90.9/35.7	87.5/31.6	60.5/100	64.5/100	76.8/65.7	77.4/71.4
ODEED	MSE	65.6/76.1	61.0/70.5	69.0/80.3	80.8/30.7	83.6/33.3	70.0/60.0	86.9/42.9	76.0/51.4
	LPIPS	87.9/20.5	84.0/18.2	94.5/24.6	94.6/06.1	75.3/73.3	45.8/80.0	64.1/85.7	60.8/60.0

Table 3. Results obtained in OOD detection (AUC \uparrow / FPR95% \downarrow) using, on a given geographical domain, a model trained on another domain. We note that the differences in performance are slight, both in AUC and in FPR95. A minor difference can be seen in the case of ODE encoding in the pre/post scenario, where the high AUC over Louisiana is preserved (94.5 vs. 94.6 when using LPIPS).

Effect of t_0 on OOD detection We now evaluate the impact of the corruption time t_0 used in the ODEED method, as it can be tuned on a validation set. We observed experimentally that its effect on the OOD detection performances depends on the task and on the similarity measure considered. For example, Fig. 4 highlights that using the MSE as the score results in higher AUC scores for small t_0 in the flood OOD scenario, while it reaches the highest performance for medium intermediate times for the pre/post OOD setup. This might indicate that smaller t_0 should be preferred for more localised anomalies. On the contrary, Fig. 5 shows that the LPIPS ability to discriminate between the pre and post-event images is higher for early t_0 whereas giving the best results on non-flooded/flooded images for intermediate noise levels. We attribute this to the perceptual effect of LPIPS, where localised changes result in small LPIPS variations, and therefore need a higher t_0 to isolate flood events.

Cross domain performances The poor results of ODEED on geographical domain OOD lead us to question the generalization of diffusion models to new areas. Tab. 3 shows that OOD detection performance remains close, even when switching the models’ domains. Naturally, reconstruction errors are higher on the domain for which the diffusion model has not been trained, *i.e.* D_θ Germany provides lower quality reconstruction on Louisiana images than D_θ Louisiana (and vice-versa). Nevertheless, we show in Fig. 6 that the distribution of the scores is shifted both for in and out-of-distribution samples. This means the detection threshold has also to be shifted, however, detection abilities are preserved, as illustrated by the similar AUC and FPR95% for the same-domain and cross-domain setups. This is promising for the detection of anomalies even under non-stationary data distributions.

6. Conclusion

In this work, we evaluated the effectiveness of diffusion models to detect out-of-distribution Earth Observation images. We introduced the ODEED scorer that leverages the deterministic reconstruction capabilities of continuous-time diffusion models. We evaluated these approaches for 1) cloud

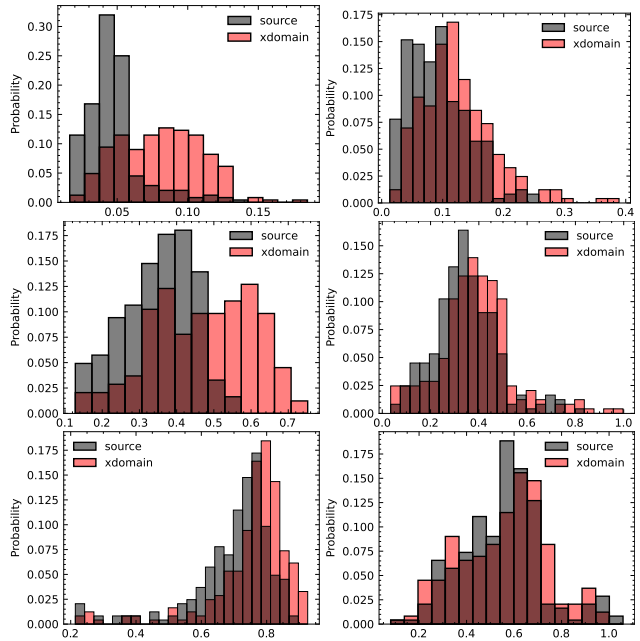


Figure 6. LPIPS shift on pre/post-event Louisiana OOD in function of t_0 , from 0.05 (top), 0.5 (middle), to 5.0 (bottom), for ID (left) and OOD (right) samples. On the Louisiana dataset, LPIPS increases when using D_θ Germany (xdomain) instead of D_θ Louisiana (source). As t_0 increases, the shift is less severe.

detection and 2) a challenging ensemble of OOD detection tasks on the Space-Net 8 dataset. We demonstrated that our ODEED scorer significantly outperforms baselines in the more challenging flood-related scenarios, showing the interest of diffusion models to detect “near-out-of-distribution” remote sensing imagery, such as images of floods. These findings open the door to new approaches to detect rare events from unlabeled EO data using generative modeling.

Acknowledgements This research was conducted as part of the research project MAGE (ANR-22-CE23-0010) funded by the *Agence Nationale de la Recherche*. This work was granted access to the HPC resources of IDRIS under the allocation AD011014327 made by GENCI.

References

- [1] Tomer Amit, Tal Shaharbany, Eliya Nachmani, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv e-prints*, pages arXiv–2112, 2021. [2](#)
- [2] Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. [3](#)
- [3] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021. [2](#)
- [4] Orkhan Baghirli, Hamid Askarov, Imran Ibrahimli, Ismat Bakhishov, and Nabi Nabiyeu. Satdm: Synthesizing realistic satellite image with semantic layout conditioning using diffusion models. *arXiv preprint arXiv:2309.16812*, 2023. [2](#)
- [5] Wele Gedara Chaminda Bandara, Nithin Gopalakrishnan Nair, and Vishal M Patel. Ddpm-cd: Remote sensing change detection using denoising diffusion probabilistic models. *arXiv preprint arXiv:2206.11892*, 2022. [5](#)
- [6] Giulio Biroli, Tony Bonnaire, Valentin de Bortoli, and Marc Mézard. Dynamical regimes of diffusion models. *arXiv preprint arXiv:2402.18491*, 2024. [3](#)
- [7] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [8] Bertrand Charpentier, Daniel Zügner, and Stephan Günnemann. Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts. *Advances in Neural Information Processing Systems*, 33:1356–1367, 2020. [2](#)
- [9] L-C Chen. Rethinking atrous convolution for semantic image segmentation. *Computing Research Repository*, 2017. [6](#)
- [10] Mihai Coca, Iulia Coca Neagoe, and Mihai Datcu. Hybrid DNN-Dirichlet Anomaly Detection and Ranking: Case of Burned Areas Discovery. 60:1–16. [3](#)
- [11] Taylor Denouden, Rick Salay, Krzysztof Czarnecki, Vahdat Abdelzad, Buu Phan, and Sachin Vernekar. Improving reconstruction autoencoder out-of-distribution detection with mahalanobis distance. *CoRR*, abs/1812.02765, 2018. [2](#), [5](#), [6](#)
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [2](#)
- [13] Miguel Espinosa and Elliot J. Crowley. Generate your own scotland: Satellite image generation conditioned on maps. *NeurIPS 2023 Workshop on Diffusion Models*, 2023. [2](#)
- [14] Alistair Francis, John Mrziglod, Panagiotis Sidiropoulos, and Jan-Peter Muller. Sentinel-2 cloud mask catalogue, 2020. [5](#)
- [15] Ruiyuan Gao, Chenchen Zhao, Lanqing Hong, and Qiang Xu. Diffguard: Semantic mismatch-guided out-of-distribution detection using pre-trained diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1579–1589, 2023. [2](#), [3](#)
- [16] Jakob Gawlikowski, Sudipan Saha, Anna Kruspe, and Xiao Xiang Zhu. Out-of-distribution detection in satellite image classification. In *RobustML Workshop at ICLR 2021*, pages 1–5. ICRL, 2021. [3](#)
- [17] Joseph Goodier and Neill DF Campbell. Likelihood-based out-of-distribution detection with denoising diffusion probabilistic models. *BMVC*, 2023. [2](#), [3](#)
- [18] Mark S. Graham, Walter H.L. Pinaya, Petru-Daniel Tudosiu, Parashkev Nachev, Sebastien Ourselin, and Jorge Cardoso. Denoising diffusion models for out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2948–2957, 2023. [2](#), [3](#)
- [19] Lintao Han, Yuchen Zhao, Hengyi Lv, Yisa Zhang, Hailong Liu, Guoling Bi, and Qing Han. Enhancing remote sensing image super-resolution with efficient hybrid conditional diffusion model. *Remote Sensing*, 15(13), 2023. [2](#)
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [2](#)
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv:2204.03458*, 2022. [2](#)
- [22] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005. [2](#)
- [23] Ronny Hänsch, Jacob Arndt, Dalton Lunga, Matthew Gibb, Tyler Pedelose, Arnold Boedihardjo, Desiree Petrie, and Todd M. Bacastow. SpaceNet 8 - The Detection of Flooded Roads and Buildings. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1471–1479. [5](#)
- [24] Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. To trust or not to trust a classifier. *Advances in neural information processing systems*, 31, 2018. [5](#)
- [25] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. [2](#), [3](#), [6](#)
- [26] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [27] Bahjat Kawar, Roy Ganz, and Michael Elad. Enhancing diffusion-based image synthesis with robust classifier guidance. *Transactions on Machine Learning Research*, 2022. [2](#)
- [28] Aliasghar Khani, Saeid Asgari, Aditya Sanghi, Ali Mahdavi Amiri, and Ghassan Hamarneh. Slime: Segment like me. In *The Twelfth International Conference on Learning Representations*, 2023. [2](#)
- [29] Samar Khanna, Patrick Liu, Linqi Zhou, Chenlin Meng, Robin Rombach, Marshall Burke, David B. Lobell, and Stefano Ermon. Diffusionsat: A generative foundation model for satellite imagery. In *The Twelfth International Conference on Learning Representations*, 2024. [2](#)
- [30] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. [3](#)

- [31] Yuma Koizumi, Heiga Zen, Kohei Yatabe, Nanxin Chen, and Michiel Bacchiani. Specgrad: Diffusion probabilistic model based neural vocoder with adaptive noise spectral shaping. 2022. [2](#)
- [32] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. [2](#)
- [33] Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. Training confidence-calibrated classifiers for detecting out-of-distribution samples. In *International Conference on Learning Representations*, 2018. [2](#)
- [34] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343, 2022. [2](#)
- [35] Jinzhe Liu, Zhiqiang Yuan, Zhaoying Pan, Yiqun Fu, Li Liu, and Bin Lu. Diffusion model with detail complement for super-resolution of remote sensing. *Remote Sensing*, 14(19), 2022. [2](#)
- [36] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in neural information processing systems*, 33:21464–21475, 2020. [2](#), [6](#)
- [37] Zhenzhen Liu, Jin Peng Zhou, Yufan Wang, and Kilian Q. Weinberger. Unsupervised out-of-distribution detection with diffusion inpainting. In *Proceedings of the 40th International Conference on Machine Learning*. JMLR.org, 2023. [2](#), [3](#)
- [38] David Macêdo, Tsang Ing Ren, Cleber Zanchettin, Adriano LI Oliveira, and Teresa Ludermitr. Entropic out-of-distribution detection. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021. [2](#)
- [39] Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. *Advances in neural information processing systems*, 31, 2018. [2](#)
- [40] Sina Mohseni, Mandar Pitale, J. B. S. Yadawa, and Zhangyang Wang. Self-Supervised Learning for Generalizable Out-of-Distribution Detection. 34(04):5216–5223. [2](#)
- [41] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. [5](#)
- [42] Gabriel Raya and Luca Ambrogioni. Spontaneous symmetry breaking in generative diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. [3](#)
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. 2021. [2](#)
- [44] Simon Rouard and Gaëtan Hadjeres. Crash: Raw audio score-based generative modeling for controllable high-resolution drum sound synthesis. *arXiv preprint arXiv:2106.07431*, 2021. [2](#)
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [2](#)
- [46] Fulvio Sanguigni, Mikolaj Czerkawski, Lorenzo Papa, Irene Amerini, and Bertrand Le Saux. Diffusion models for earth observation use-cases: from cloud removal to urban change detection. *ArXiv*, abs/2402.06684, 2023. [2](#), [5](#)
- [47] Ahmad Sebaq and Mohamed ElHelw. Rsdiff: Remote sensing image generation from text using diffusion model. *arXiv preprint arXiv:2309.02455*, 2023. [2](#)
- [48] Vikash Sehwal, Mung Chiang, and Prateek Mittal. SSD: A Unified Framework for Self-Supervised Outlier Detection. [2](#)
- [49] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265, Lille, France, 2015. PMLR. [2](#)
- [50] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [51] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution, 2020. [2](#)
- [52] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020. [2](#), [6](#)
- [53] Yiyou Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. Out-of-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*, pages 20827–20840. PMLR, 2022. [6](#)
- [54] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. [2](#)
- [55] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4921–4930, 2022. [2](#)
- [56] Kevin E Wu, Kevin K Yang, Rianne van den Berg, Sarah Alamdari, James Y Zou, Alex X Lu, and Ava P Amini. Protein structure generation via folding diffusion. *Nature Communications*, 15(1):1059, 2024. [2](#)
- [57] Yi Xiao, Qiangqiang Yuan, Kui Jiang, Jiang He, Xianyu Jin, and Liangpei Zhang. Ediffsr: An efficient diffusion probabilistic model for remote sensing image super-resolution. *IEEE Transactions on Geoscience and Remote Sensing*, 2023. [2](#)
- [58] Jingkan Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized Out-of-Distribution Detection: A Survey. [1](#)
- [59] Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023. [2](#)
- [60] Lily Zhang, Mark Goldstein, and Rajesh Ranganath. Understanding Failures in Out-of-Distribution Detection with Deep Generative Models. In *Proceedings of the 38th International Conference on Machine Learning*, pages 12427–12436. PMLR. [2](#)
- [61] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In

Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 3836–3847, 2023. 2

- [62] 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. 3
- [63] Yibo Zhou. Rethinking reconstruction autoencoder-based out-of-distribution detection. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7369–7377, 2022. 2
- [64] X. X. Zhu, D. Tuia, L. Mou, G. S. Xia, L. Zhang, F. Xu, and F. Fraundorfer. Deep Learning in Remote Sensing: A Comprehensive Review and List of Resources. 5(4):8–36. 1
- [65] 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