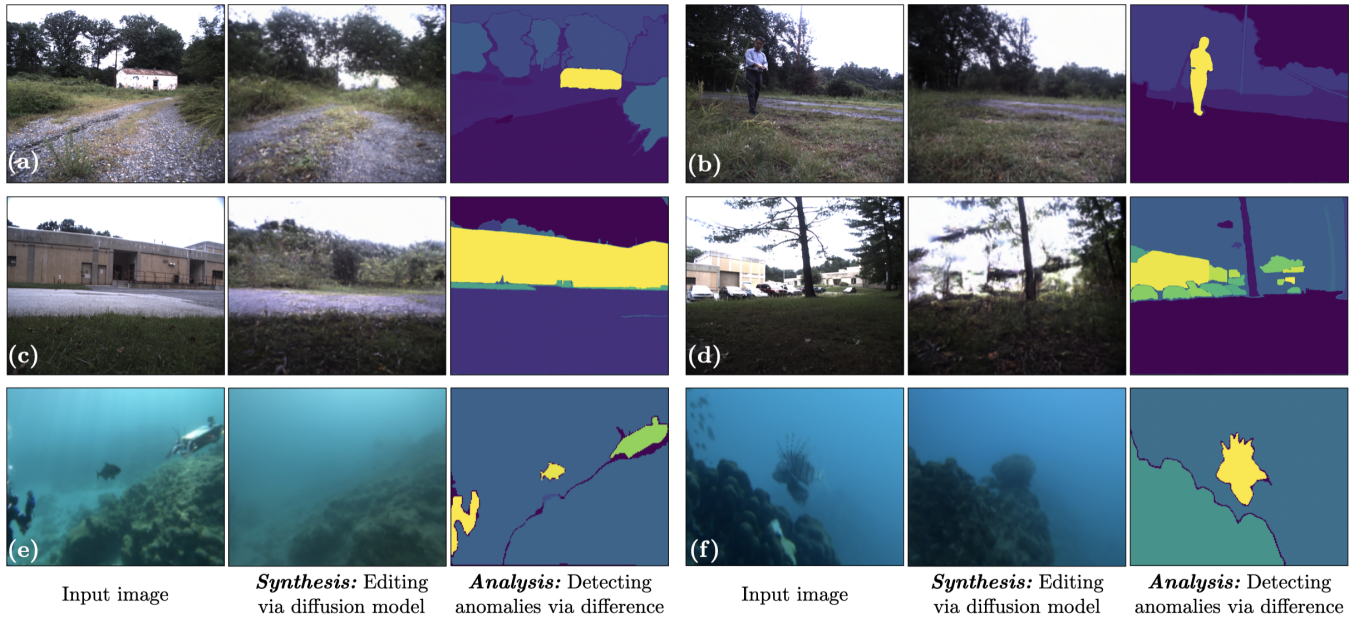


# Anomalies-by-Synthesis: Anomaly Detection using Generative Diffusion Models for Off-Road Navigation

Sunshine Jiang<sup>\*1</sup>, Siddharth Ancha<sup>\*1</sup>, Travis Manderson<sup>1</sup>, Laura Brandt<sup>1</sup>,  
Yilun Du<sup>1</sup>, Philip R. Osteen<sup>2</sup> and Nicholas Roy<sup>1</sup>

Website: <https://bit.ly/anomalies-by-synthesis><sup>†</sup>



**Fig. 1:** We present a method for anomaly detection in off-road images using an *analysis by synthesis* approach. **Synthesis:** A diffusion model trained on in-distribution data edits the input image to remove out-of-distribution segments. **Analysis:** Anomaly detection is framed as extracting the difference between input and edited images. The training dataset contains off-road images with: (a-d) natural vegetation, ground and sky, but no buildings, humans or vehicles, (e-f) underwater scenes with mostly the oceanbed and coral. The diffusion model makes interesting edits, such as (a) blending buildings into sky, (b) removing people, (c) growing moss over buildings, (e) removing fish, robots and divers, and (f) morphing camouflaged fish into rocks. Our method can detect small objects and multiple anomalies per image (d).

**Abstract**—In order to navigate safely and reliably in off-road and unstructured environments, robots must detect anomalies that are out-of-distribution (OOD) with respect to the training data. We present an *analysis-by-synthesis* approach for pixel-wise anomaly detection without making any assumptions about the nature of OOD data. Given an input image, we use a generative diffusion model to *synthesize* an edited image that removes anomalies while keeping the remaining image unchanged. Then, we formulate anomaly detection as *analyzing* which image segments were modified by the diffusion model. We propose a novel inference approach for guided diffusion by analyzing the ideal guidance gradient and deriving a principled approximation that bootstraps the diffusion model to predict guidance gradients. Our editing technique is purely test-time that can be integrated into existing workflows without the need for retraining or fine-tuning. Finally, we use a combination of vision-language

foundation models to compare pixels in a learned feature space and detect semantically meaningful edits, enabling accurate anomaly detection for off-road navigation.

## I. INTRODUCTION

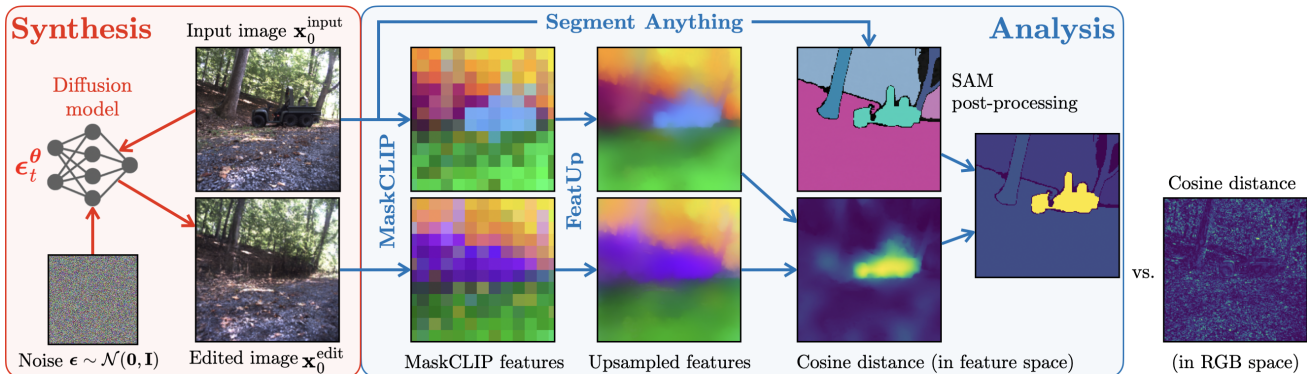
The need for reliable autonomous robotic navigation is increasing in unstructured, off-road environments like planetary exploration [5, 6], forests [7], deserts [8, 9] and underwater [10, 11]. However, the perception systems required for off-road navigation (e.g., semantic segmentation [12–14]) are often trained on relatively small datasets [15, 16] and deployed in environments where images often contain “out-of-distribution” (OOD) *anomalies* not well represented in the training data. Furthermore, lighting conditions, weather, and terrain can vary significantly between training and test environments. To navigate safely and reliably in unfamiliar environments, robots must identify such *anomalies* in order to anticipate potential perception failures. The aim of this work is to detect anomalous segments (if any) in a given

<sup>\*</sup>Denotes equal contribution. <sup>1</sup>MIT CSAIL, Cambridge, MA 02139, USA.

<sup>2</sup>DEVCOM Army Research Laboratory, Adelphi, MD 20783, USA.

Correspondence email: [sancha@mit.edu](mailto:sancha@mit.edu).

<sup>†</sup>While this paper is fully self-contained, our project website contains (1) a 3-min overview video, (2) qualitative visualizations, (3) an appendix with detailed proofs, (4) Google Colab notebook, and (5) code, for easy access.



**Fig. 2:** Our proposed pipeline for pixel-wise anomaly detection. *Left to right:* In the synthesis step, a trained diffusion model edits a given input image to remove anomaly segments without modifying other parts of the image. In this case, the model blends the OOD vehicle into dirt in the background. The analysis step extracts anomalies by comparing the pair of images in the CLIP [1] feature space. First, MaskCLIP [2] computes low-resolution CLIP features for each image, which are upsampled using FeatUp [3]. In this figure, features are visualized via a t-SNE projection to three dimensions. Cosine distances between pixel features in the two images produce a raw anomaly map that highlights anomaly objects. In contrast, comparing the images directly in RGB space (extreme right) is noisy and unable to isolate OOD segments. Finally, SAM [4] processes the input image to generate segments; these are used to refine and clean the anomaly map.

RGB image that are out-of-distribution from the training data.

Prior work on anomaly detection has focused on *discriminative* models that directly map pixels to anomaly scores. These models classify features at the pixel- or segment-level as in-distribution or as anomalies [17–31]. On the other hand, recent advances in generative AI have led to the development of diffusion models [32, 33] that can model complex distributions and generate high-quality realistic images [34, 35]. In this work, we present an *analysis-by-synthesis* [36, 37] approach to anomaly detection (hence the title “anomalies-by-synthesis”). We use a *generative* diffusion model to “edit” the input image and remove anomalies while keeping the remaining input unchanged. The new, *synthesized* image represents what the scene would have looked like had it not contained any anomalies. Then, we frame the problem of anomaly detection as *analyzing* which parts of the image were edited by the generative model to remove anomalies. In other words, the difference between the input and synthesized images produces anomaly detections.

Our method has several advantages. First, we **do not make any assumptions about the nature of anomalies** the model can expect at test time. Our approach **does not require any OOD data during training**. We use diffusion guidance [34, 38–40] to edit the input image. This is a post-hoc test-time procedure that **does not require re-training or finetuning the diffusion model**. The diffusion model is only assumed to be trained on the standard objective of fitting to the training data. Therefore, our method can be applied to pre-trained or off-the-shelf diffusion models whose architectures and hyperparameters are carefully tuned for high-quality image synthesis.

We propose a **novel diffusion guidance approach** to ensure that the synthesized image is both similar to the input image and has a high probability under the training distribution. We **theoretically analyze the ideal guidance gradient**, which is intractable to compute. We then **propose a principled and tractable approximation** that, unlike prior methods, reuses the learned diffusion score function to compute the guidance

gradient. Our diffusion guidance approach modifies pixels corresponding to anomaly regions while keeping the in-distribution parts of the image as close to the input as possible.

Unfortunately, extracting anomaly segments from the edited image is not as straightforward as computing pixel-wise intensity differences with the input image. The diffusion model tends to slightly alter pixels throughout the entire image, including regions that are in-distribution. While the differences are subtle to the human eye, pixel-wise intensity differences are sensitive to even minor changes between images. This sensitivity produces extremely noisy anomaly masks riddled with false positives. We mitigate this issue using a **combination of foundation vision models**: MaskCLIP [2], FeatUp [3] and SAM [4]. By comparing images in CLIP [1] feature space, we are able to **detect semantically meaningful edits while being robust to subtle and inconsequential pixel intensity changes**.

We quantitatively validate the effectiveness of our approach on two public off-road land navigation datasets: RUGD [15] and RELLIS [16], and qualitatively evaluate on an underwater navigation dataset [11].

## II. PROBLEM FORMULATION: ANOMALY DETECTION AS POST-HOC ANALYSIS-BY-SYNTHESIS

**Pixel-wise anomaly detection:** We consider the problem of pixel-wise anomaly detection in RGB images. We denote an RGB image of dimensions  $H \times W$  by  $\mathbf{x}_0 \in [0, 1]^{3HW}$ , with a zero-subscript<sup>1</sup>. In our problem setting, we are given a set of  $N$  training images  $\mathcal{D}_{\text{train}} = \{\mathbf{x}_0^{(n)}\}_{n=1}^N \stackrel{i.i.d.}{\sim} q(\mathbf{x}_0)$  assumed to be sampled *i.i.d.* from a training distribution<sup>2</sup>  $q(\mathbf{x}_0)$ . We refer to  $q$  as the “training distribution” or the “in-distribution”. At test time, we are given an input image  $\mathbf{x}_0^{\text{input}}$ . A subset of its pixels belong to objects that are *out-of-distribution* (OOD).

<sup>1</sup>The zero-subscript in  $\mathbf{x}_0$  is introduced for consistency with notation used for diffusion models [32] (see Sec. III). It corresponds to the noise level  $t = 0$  implying that no noise has been added to the image. We slightly modify conventional notation to consistently specify timesteps in the subscript and model parameters in the superscript when applicable e.g.  $\mathbf{x}_t, p_t^\theta, \epsilon_t^\theta, \mu_t^\theta$ .

<sup>2</sup>We denote the training distribution by  $q(\cdot)$  following Ho et al. [32].

The task is to detect these anomalous pixels in the image. During training, the anomaly detector is not allowed to make any assumptions about the nature of OOD examples it can expect at test time. Anomalies are *implicitly* defined by how unlikely the respective object segments are under the training distribution  $q(\mathbf{x}_0)$ .

**Analysis-by-synthesis formulation:** We formulate the problem of pixel-wise anomaly detection in three stages.

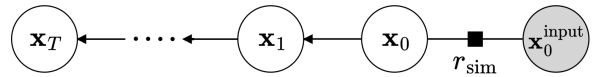
(i) **Training:** We train a whole-image generative model  $p^\theta(\mathbf{x}_0)$ , parameterized by weights  $\theta \in \mathbb{R}^W$ , to fit the training set  $\mathcal{D}_{\text{train}}$  and learn  $q(\mathbf{x}_0)$ . Learning  $p^\theta(\mathbf{x}_0) \approx q(\mathbf{x}_0)$  is the only objective in this stage; training is not specialized for anomaly detection. This de-coupling of the problem of *learning*  $p^\theta \approx q$  from the task of *inferring* anomaly segments given a learned model  $p^\theta$  allows us to repurpose pre-trained generative models for anomaly detection.

(ii) **Synthesis:** Given an input image  $\mathbf{x}_0^{\text{input}}$  at test time, we wish to detect anomaly segments by synthesizing an edited image using the learned model  $p^\theta \approx q$ . We assume a similarity metric  $r_{\text{sim}} : \mathbb{R}^{3HW} \times \mathbb{R}^{3HW} \rightarrow \mathbb{R}_{\geq 0}$  between two images that quantifies how similar they are. In this work, we use the Gaussian kernel:  $r_{\text{sim}}(\mathbf{x}'_0, \mathbf{x}''_0) = \exp(-\lambda \|\mathbf{x}'_0 - \mathbf{x}''_0\|_2^2)$ , although any non-negative similarity metric can be used. We define the editing process, i.e., *synthesis*, as sampling from the following distribution:

$$\mathbf{x}_0^{\text{edit}} \sim q(\mathbf{x}_0^{\text{edit}} | \mathbf{x}_0^{\text{input}}) \propto \underbrace{q(\mathbf{x}_0^{\text{edit}})}_{\text{likelihood under training distribution}} \underbrace{r_{\text{sim}}(\mathbf{x}_0^{\text{edit}}, \mathbf{x}_0^{\text{input}})}_{\text{similarity to input image}} \quad (1)$$

The distribution in Eq. 1 is a product of two terms. The first term  $q(\cdot)$ , which is approximated by the learned generative model  $p^\theta(\cdot)$ , ensures that the edited image  $\mathbf{x}_0^{\text{edit}}$  is likely under the training distribution. The second term ensures that  $\mathbf{x}_0^{\text{edit}}$  is similar to the input image. Sampling from this conditional distribution synthesizes a new image that effectively “edits”  $\mathbf{x}_0^{\text{input}}$  by removing anomalies but keeping pixels in the in-distribution regions as close as possible to the input image.  $\mathbf{x}_0^{\text{edit}}$  can also be interpreted as a “projection” of  $\mathbf{x}_0^{\text{input}}$  onto the manifold of the training distribution  $p^\theta(\mathbf{x}_0) \approx q(\mathbf{x}_0)$ . Sampling from Eq. 1 is a challenging inference problem:  $q(\mathbf{x}_0)$  is a complex multimodal distribution over natural images, and multiplying by  $r_{\text{sim}}(\cdot)$  makes the distribution unnormalized with an intractable normalizing constant. In this work, we develop a diffusion guidance approach [34, 38, 39] to sample from this distribution. The synthesis procedure is “post-hoc” in the sense that the generative model is not explicitly trained to perform the editing task; editing is result of an inference procedure performed at test-time.

(iii) **Analysis:** We formulate detecting anomalies in  $\mathbf{x}_0^{\text{input}}$  as computing the difference between  $\mathbf{x}_0^{\text{input}}$  and  $\mathbf{x}_0^{\text{edit}}$ . As described in Sec. I, RGB intensities in  $[0, 1]^3$  may not be a suitable space for this comparison. We define the *analysis* task as finding an appropriate  $C$ -dimensional feature space  $f : [0, 1]^{H \times W \times 3} \rightarrow \mathbb{R}^{H \times W \times C}$  so that the anomaly score for the  $i$ -th pixel can be computed as  $\|f(\mathbf{x}_0^{\text{edit}})_i - f(\mathbf{x}_0^{\text{input}})_i\|$ . We use the CLIP [1] feature space for comparison, via a combination of MaskCLIP [2], FeatUp [3] and SAM [4].



**Fig. 3:** Probabilistic graphical model for the *conditional* forward diffusion process. The target variable we wish to sample is  $\mathbf{x}_0$ . Directed edges correspond to the standard forward diffusion process. The unnormalized factor  $r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$  conditions  $\mathbf{x}_0$  to be similar to the (fixed) input image.

### III. BACKGROUND ON DIFFUSION MODELS

Diffusion models [32, 38, 39] are recently developed, state-of-the-art generative models effective at learning and sampling from complex, multi-modal image distributions. They have been primarily used in the computer vision community for photorealistic image generation and creative applications [35, 41], and in the robot learning community to learn distributions over action trajectories [42, 43]. In this work, we use diffusion models for anomaly detection in images. Recall that we denote (noiseless) image vectors by  $\mathbf{x}_0 \in [0, 1]^{3HW}$ . We train a denoising diffusion probabilistic model (DDPM) [32]  $p^\theta(\mathbf{x}_0)$  to fit the training distribution  $q(\mathbf{x}_0)$  and sample from it. To do so, DDPM introduces  $T$  latent variables  $\mathbf{x}_1, \dots, \mathbf{x}_T$  by the *forward diffusion process*  $q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$ . At each timestep  $t$ , forward diffusion progressively adds a small amount of *i.i.d.* Gaussian noise with variance  $\beta_t > 0$  to  $\mathbf{x}_{t-1}$ . With an appropriate noise schedule  $\beta_{1:T}$ , the final latent variable  $\mathbf{x}_T$  is approximately normally distributed i.e.  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

Because forward diffusion is Gaussian, latent variables from intermediate timesteps  $\mathbf{x}_t$  can be directly sampled from  $\mathbf{x}_0$  as  $\mathbf{x}_t(\mathbf{x}_0, \epsilon) = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \epsilon$  where  $\alpha_t = \prod_{s=1}^t (1 - \beta_s)$  [32, 38] and  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  is a noise vector sampled from the standard normal distribution. DDPM trains a neural network  $\epsilon_t^\theta(\mathbf{x}_t)$  with weights  $\theta$  to predict  $\epsilon$  that generated  $\mathbf{x}_t$ .  $\epsilon_t^\theta$  is trained to minimize the loss  $\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|\epsilon - \epsilon_t^\theta(\mathbf{x}_t(\mathbf{x}_0, \epsilon))\|_2^2$ . This network allows DDPM to perform *reverse diffusion*. First,  $\mathbf{x}_T \sim p^\theta(\mathbf{x}_T) = \mathcal{N}(\mathbf{0}, \mathbf{I})$  is sampled from the standard normal. Then,  $\mathbf{x}_T$  is progressively denoised to  $\mathbf{x}_0$ . In the limit of small  $\beta_t$ , the reversal of forward diffusion  $q(\mathbf{x}_t | \mathbf{x}_{t+1})$  becomes Gaussian [38] and can be approximated as  $q(\mathbf{x}_t | \mathbf{x}_{t+1}) \approx p^\theta(\mathbf{x}_t | \mathbf{x}_{t+1}) = \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_{t+1}^\theta(\mathbf{x}_{t+1}), \tilde{\beta}_{t+1})$ , where  $\boldsymbol{\mu}_{t+1}^\theta$  is a reparametrization of  $\epsilon_{t+1}^\theta$  (see line 5 of Alg. 1). Finally, the reverse diffusion process produces a sample from  $p^\theta(\mathbf{x}_0) \approx q(\mathbf{x}_0)$ . We refer to Ho et al. [32] for more details on DDPM.

### IV. GUIDANCE GRADIENTS FOR CONDITIONAL DIFFUSION

Given an input image  $\mathbf{x}_0^{\text{input}}$ , we are interested in sampling from  $q(\mathbf{x}_0 | \mathbf{x}_0^{\text{input}}) \propto q(\mathbf{x}_0) r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$  as noted in Eq. 1. To do so, we first define a *conditional forward diffusion process* [34] as shown in Fig. 3. It can be denoted by  $q(\mathbf{x}_{0:T} | \mathbf{x}_0^{\text{input}}) = q(\mathbf{x}_0 | \mathbf{x}_0^{\text{input}}) q(\mathbf{x}_{1:T} | \mathbf{x}_0)$ , where  $\mathbf{x}_0$  is hypothetically sampled from our desired conditional distribution  $q(\mathbf{x}_0 | \mathbf{x}_0^{\text{input}})$ , followed by the standard forward diffusion process  $q(\mathbf{x}_{1:T} | \mathbf{x}_0)$ . A compelling property of diffusion models is that  $\epsilon_t^\theta$  can be used to reverse a *conditional* forward diffusion processes at test-time without having to re-train  $\epsilon_t^\theta$ . Specifically, *classifier guidance* [34, 38–40] can

---

**Algorithm 1:** Similarity-guided diffusion to edit input image and remove anomalies

---

```

1: procedure GUIDEDIFFUSION ( $\epsilon_t^\theta$ ,  $\mathbf{x}_0^{\text{input}}$ ,  $r_{\text{sim}}$ )
2:    $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
3:   for  $t = T - 1, \dots, 0$  do
4:      $\triangleright$  Standard reverse diffusion step to compute un-conditional mean of  $\mathbf{x}_t$  given  $\mathbf{x}_{t+1}$ 
5:      $\bar{\mathbf{x}}_t := \mu_{t+1}^\theta(\mathbf{x}_{t+1}) = \mathbb{E}_q[\mathbf{x}_t | \mathbf{x}_{t+1}] \approx \frac{1}{\sqrt{\alpha_{t+1}}} \left( \mathbf{x}_{t+1} - \frac{\beta_{t+1}}{\sqrt{1-\alpha_{t+1}}} \epsilon_{t+1}^\theta(\mathbf{x}_{t+1}) \right)$ 
6:      $\triangleright$  Computing the guidance gradient for conditioning on  $r_{\text{sim}}(\cdot)$ 
7a:     $\mu_0^\theta(\mathbf{x}_t) := \mathbb{E}_q[\mathbf{x}_0 | \mathbf{x}_t] \approx \frac{1}{\sqrt{\alpha_t}} \mathbf{x}_t - \frac{\sqrt{1-\alpha_t}}{\sqrt{\alpha_t}} \epsilon_t^\theta(\mathbf{x}_t)$ 
7b:     $\mathbf{g}_t(\mathbf{x}_t) := \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mathbf{x}_t, \mathbf{x}_0^{\text{input}})$ 
8a:     $\mathbf{g}_t^\theta(\mathbf{x}_t) := \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mu_0^\theta(\mathbf{x}_t), \mathbf{x}_0^{\text{input}})$ 
8b:     $\mathbf{g}_t(\mathbf{x}_t) := \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mathbf{x}_t, \mathbf{x}_0^{\text{input}})$ 
9:      $\triangleright$  Standard guided sampling step
10:     $\mathbf{x}_t \sim \bar{\mathbf{x}}_t + \sigma^2 \mathbf{g}_t^\theta(\bar{\mathbf{x}}_t) + \sigma \epsilon$  where  $\sigma^2 = \tilde{\beta}_{t+1}$ ,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
11:  return  $\mathbf{x}_0^{\text{edit}} := \mathbf{x}_0$ 

```

**Algorithm 1:** *Inputs:* (1)  $\epsilon_t^\theta$ : denoising diffusion model fit to training image distribution  $q(\mathbf{x}_0)$ . (2)  $\mathbf{x}_0^{\text{input}}$ : input image at test-time potentially containing anomaly segments. (3)  $r_{\text{sim}}$ : similarity metric between two images. We use  $r_{\text{sim}}(\mathbf{x}, \mathbf{y}) = \exp(-\lambda \|\mathbf{x} - \mathbf{y}\|_2^2)$ . *Output:*  $\mathbf{x}_0^{\text{edit}}$ : edited version of  $\mathbf{x}_0^{\text{input}}$  sampled from  $q(\mathbf{x}_0) r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$  that removes anomalies. While standard reverse diffusion (lines 5, 10) is designed to sample from  $q(\mathbf{x}_0)$ , the image similarity metric  $r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$  guides the generated image  $\mathbf{x}_0$  to be close to  $\mathbf{x}_0^{\text{input}}$  using guided diffusion [34, 38]. This is achieved by including the log-gradient of  $r_{\text{sim}}$  (line 8a). A key question during intermediate timesteps  $t$  is what pair of images should the similarity metric compare? A naïve approach [38] (line 8b) simply uses the intermediate sample  $\mathbf{x}_t$  and  $\mathbf{x}_0^{\text{input}}$ ; they however correspond to different noise levels and should not be compared directly. Another baseline attempts to bring  $\mathbf{x}_0^{\text{input}}$  to timestep  $t$  by expected forward diffusion (lines 7c, 8c). Instead, we propose using the trained diffusion model to estimate the expected noise-free version  $\mu_0^\theta(\mathbf{x}_t)$  of  $\mathbf{x}_t$  to compare with the input image  $\mathbf{x}_0^{\text{input}}$  at  $t = 0$ . While this requires a backward pass through the diffusion model (line 8a), we show that our approach is mathematically principled and empirically leads to better performance.

be used in the special case where one wishes to sample from  $q(\mathbf{x}_0 | y) \propto q(\mathbf{x}_0) q(y | \mathbf{x}_0)$  i.e., the conditioning factor takes the form of a classifier  $q(y | \mathbf{x}_0)$ . Classifier guidance adds a gradient step  $\mathbf{g}_t(\mathbf{x}_t) = \nabla_{\mathbf{x}_t} \log p(y | \mathbf{x}_t)$  at each reverse diffusion step, shown in line 8 (a, b, c) of Alg. 1. Intuitively, the guidance gradient encourages the conditioning factor of the resulting  $\mathbf{x}_t$  to increase, while balancing with the original objective of sampling from  $q(\mathbf{x}_0)$ .

**Mismatched-timesteps baseline:** In our case, the conditioning factor is not a classifier but is of the form  $r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$ . A naïve application of classifier guidance, as used in Sohl-Dickstein et al. [38], would be to choose the guidance gradient as  $\mathbf{g}_t(\mathbf{x}_t) = \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mathbf{x}_t, \mathbf{x}_0^{\text{input}})$ ; see line 8b of Alg. 1. However, this formulation compares a noiseless image  $\mathbf{x}_0^{\text{input}}$  with a noisy image  $\mathbf{x}_t$  corresponding to the noise level at timestep  $t$ . The mismatch between the timesteps of the two images being compared leads to poor performance.

**Forward timestep matching baseline:** A heuristic approach to match the timesteps of the images is to add noise to the input image. We call this the “forward timestep matching” baseline (lines 7c, 8c in Alg. 1). In this baseline,  $\mathbf{x}_0^{\text{input}}$  is transported to timestep  $t$  via forward diffusion. The noised image  $\mathbf{x}_t^{\text{input}}$  is then compared with  $\mathbf{x}_t$ . We now derive a more principled method for guided conditional diffusion, and show that the correct way to match image timesteps is in the reverse direction.

#### V. GENERALIZED SIMILARITY-CONDITIONED GUIDED DIFFUSION AND A PRINCIPLED APPROXIMATION

Conventional classifier guidance is restricted to using a classifier as the conditioning factor. Furthermore, it also requires training a classifier  $p(y | \mathbf{x}_t)$  for each intermediate

latent state [34, 39]. Instead, we extend classifier guidance to the more general setting where the conditioner can be any non-negative function:

**Theorem 1:** When a diffusion model  $\epsilon_t^\theta$  is trained to sample from  $q(\mathbf{x}_{0:T})$ , the conditional distribution  $q(\mathbf{x}_{0:T} | \mathbf{x}_0^{\text{input}}) \propto q(\mathbf{x}_{0:T}) r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})$  can be sampled by using the idealized guidance gradient  $\mathbf{g}_t^*(\mathbf{x}_t)$  during reverse diffusion:

$$\begin{aligned} \mathbf{g}_t^*(\mathbf{x}_t) &= \nabla_{\mathbf{x}_t} \log \int_{\mathbf{x}_0} q(\mathbf{x}_0 | \mathbf{x}_t) r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}}) d\mathbf{x}_0 \\ &= \nabla_{\mathbf{x}_t} \log \mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_t)} [r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})] \end{aligned} \quad (2)$$

See proof in App. I. Intuitively, this gradient guides  $\mathbf{x}_t$  to increase the expected similarity of its denoised version  $\mathbf{x}_0 \sim q(\mathbf{x}_0 | \mathbf{x}_t)$  with  $\mathbf{x}_0^{\text{input}}$ . The timesteps of the two images being compared now match, and correspond to  $t = 0$ , i.e., zero noise. This makes sense, as we want the eventually denoised image (not the intermediate noisy images) to match the input image. Unfortunately, the guidance gradient is intractable to compute due to the high-dimensional expectation over  $q(\mathbf{x}_0 | \mathbf{x}_t)$ . We propose approximating  $\mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_t)} [r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})]$  by a point estimate of  $q(\mathbf{x}_0 | \mathbf{x}_t)$ , namely, its expected value:

$$\begin{aligned} \mathbf{g}_t^*(\mathbf{x}_t) &= \nabla_{\mathbf{x}_t} \log \mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_t)} [r_{\text{sim}}(\mathbf{x}_0, \mathbf{x}_0^{\text{input}})] \\ &\approx \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_t)} [\mathbf{x}_0], \mathbf{x}_0^{\text{input}}) \\ &\approx \nabla_{\mathbf{x}_t} \log r_{\text{sim}}(\mu_0^\theta(\mathbf{x}_t), \mathbf{x}_0^{\text{input}}) = \mathbf{g}_t^\theta(\mathbf{x}_t) \end{aligned} \quad (3)$$

The point estimate  $\mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_t)} [\mathbf{x}_0]$  can be computed analytically by the diffusion model, which we denote by  $\mu_0^\theta(\mathbf{x}_t)$ . The expression of  $\mu_0^\theta(\mathbf{x}_t)$  in terms of  $\epsilon_t^\theta(\mathbf{x}_t)$  is shown in line 7a of Alg. 1, and derived in detail in App. II. Intuitively, since  $\epsilon_t^\theta(\mathbf{x}_t)$  is trained to predict the expected noise  $\epsilon$  that

produced  $\mathbf{x}_t$  from  $\mathbf{x}_0$ , the expected value of  $\mathbf{x}_0$  given  $\mathbf{x}_t$  can be computed by subtracting (a scaled version of)  $\epsilon_t^\theta(\mathbf{x}_t)$  from  $\mathbf{x}_t$ . Therefore, our proposed guidance gradient transforms  $\mathbf{x}_t$  in the *reverse* direction to  $t = 0$ , matching  $\mathbf{x}_0^{\text{input}}$ .

Another way to see why our approximation is principled is to examine its value when  $t \approx T$ . In forward diffusion,  $\mathbf{x}_0$  and  $\mathbf{x}_T$  are independently distributed [38, 39]. This implies that at the beginning of reverse diffusion, the ideal guidance gradient  $\mathbf{g}_T^*(\mathbf{x}_T) = \mathbf{0}$  for all  $\mathbf{x}_T$  (since the expectation in Eq. 2 is constant with respect to  $\mathbf{x}_T$ , therefore the gradient is zero). This property holds for our approximation  $\mathbf{g}_T^\theta(\mathbf{x}_T)$ . Assuming that the diffusion model is well-trained,  $\mu_0^\theta(\mathbf{x}_T) \approx \mathbb{E}_{q(\mathbf{x}_0 | \mathbf{x}_T)}[\mathbf{x}_0]$  is a constant function of  $\mathbf{x}_T$  due to the independence of  $\mathbf{x}_0$  and  $\mathbf{x}_T$ . Therefore  $\mathbf{g}_T^\theta(\mathbf{x}_T) = \nabla_{\mathbf{x}_T} \log r_{\text{sim}}(\mu_0^\theta(\mathbf{x}_T), \mathbf{x}_0^{\text{input}}) \approx \mathbf{0}$ . However, this property is violated by both baseline guidance gradients  $\mathbf{g}_T(\mathbf{x}_T)$  (lines 8b, 8c of Alg. 1), which, in general, can be non-zero.

As far as we are aware, ours is the first work to *leverage the learned diffusion model*  $\epsilon_t^\theta(\mathbf{x}_t)$  to compute the guidance gradient  $\mathbf{g}_t^\theta(\mathbf{x}_t)$ . Furthermore, our method backprops through the network  $\epsilon_t^\theta(\mathbf{x}_t)$  at test time in order to compute  $\mathbf{g}_t^\theta(\mathbf{x}_t)$ . In contrast, the baseline guidance gradients  $\mathbf{g}_t(\mathbf{x}_t)$  are purely analytical and do not use the diffusion model. We show in Sec. VII that utilizing the diffusion model to compute the guidance gradient improves anomaly detection performance.

## VI. EXTRACTING ANOMALY SEGMENTS FROM DIFFUSION-EDITED IMAGE USING VLP MODELS

The diffusion model allows us to synthesize an edited image  $\mathbf{x}_0^{\text{edit}}$  by replacing the anomalous regions of  $\mathbf{x}_0^{\text{input}}$  with content similar to the training distribution. Now, we wish to *analyze* the difference between the two images to detect the modified anomaly segments. Our pipeline is shown in Fig. 2. We found that directly comparing pixel intensities between the two images produces very noisy anomaly masks (shown in the extreme right of Fig. 2). Although the most salient edits made by the diffusion model are the removal of anomaly regions, the model also makes subtle changes across the entire image. Instead, we propose comparing the two images in a feature space that is invariant to subtle intensity changes. We wish to capture higher-level *semantic* changes, such as the replacement of the brown vehicle with mud in Fig. 2, even though they have similar RGB values.

Our pipeline leverages pre-trained vision-language models. First, we use MaskCLIP [2] to compute CLIP features [1] for each image. Because MaskCLIP has a large receptive field, it downsamples an input image of size  $224 \times 224$  to  $14 \times 14$ . Then, we use FeatUp [3] to upsample the CLIP features back to the full image resolution. FeatUp is designed to upsample feature embeddings to align them with object boundaries [3]. Finally, we compute an anomaly score for each pixel as the cosine distance between corresponding CLIP feature vectors. While this distance-based mask is able to highlight anomaly segments well, we further refine it using the SegmentAnything Model (SAM) [4]. SAM produces a set of accurate, open-world segments per image; we simply assign to each SAM segment the average cosine distance score averaged across

all pixels in that segment. This approach combines SAM’s fine segmentation capability with our coarser difference-based anomaly identification.

## VII. EXPERIMENTS

### A. Dataset and experimental design

We wish to identify pixels belonging to OOD anomaly objects. From the full semantic label set  $\mathbb{C} = \{1, \dots, C\}$ , we define a subset  $\mathbb{C}_{\text{ID}} \subset \mathbb{C}$  as the set of in-distribution classes, and the remaining subset  $\mathbb{C}_{\text{OOD}} = \mathbb{C} \setminus \mathbb{C}_{\text{ID}}$  as out-of-distribution classes. We split the overall dataset such that pixels in the training dataset  $\mathcal{D}_{\text{train}}$  belong exclusively to  $\mathbb{C}_{\text{ID}}$ , whereas test images contain pixels drawn from both  $\mathbb{C}_{\text{ID}}$  and  $\mathbb{C}_{\text{OOD}}$ . Importantly, we assume that the learner has no access to any amount of real or synthetic OOD data at training time, and we do not make any assumptions about the nature of OOD data that the model can expect at test time.

We quantitatively validate our approach using two off-road land navigation datasets: the RUGD [15] dataset and the RELIS-3D [16] dataset. Both are real-world datasets containing camera images collected in off-road environments using mobile robot platforms with manually labeled pixel-wise class annotations. They contain 7,453 and 6,235 labeled images, respectively. We split the semantic categories into in-distribution labels that contains mostly natural features and vegetation such as dirt, grass, sky etc., whereas classes like vehicle, building, person are defined as out-of-distribution. See App. III for a full list of in-distribution and out-of-distribution classes for each dataset.

### B. Evaluation metrics

We evaluate our approach using the following metrics. The results of our evaluation are presented in Table I. (i) **AUC-PR**: We compute the area under the precision-recall curve between the per-pixel anomaly score, and the ground-truth classification of the pixel as in-distribution. (ii)  **$F_1^*$  score**: We borrow this metric from existing anomaly detection benchmarks for structured urban-driving scenes [49]. The  $F_1^*$ -score summarizes true positive, false positive and false negative detections averaged over different detection thresholds, and normalized by the size of ground-truth segments to prevent large objects dominating the metric. Not all classes in  $\mathbb{C}_{\text{ID}}$  are equally in-distribution; some classes (e.g. grass) occur more frequently than others (e.g. rock-bed). Therefore, we weight the classes in  $\mathbb{C}_{\text{ID}}$  by their frequency in the training dataset when computing AUC-PR and  $F_1^*$  scores.

### C. Baselines and ablations

We compare against multiple baselines in Table I. **SSIM** (row 1): We use the structural similarity index (SSIM) [44] as an anomaly score that compares images using pixel intensities (and not a learned feature space). Baselines 2-4 represent the conventional paradigm of directly predicting anomaly scores from pixels. We fit a **GMM** (row 2) with 20 components on pixel-level MaskCLIP [2] features and uses the negative log-likelihood as the anomaly score. **Nearest-neighbor search** (row 3): is akin to memorizing the training dataset. We collect

			RUGD dataset [15]		RELLIS dataset [16]	
			AUC-PR ( $\uparrow$ )	$F_1^*$ score ( $\uparrow$ )	AUC-PR ( $\uparrow$ )	$F_1^*$ score ( $\uparrow$ )
Baselines	1	SSIM [44]	0.293	0.410	0.278	0.506
	2	GMM w/ MaskCLIP [2] + SAM	0.554	0.587	0.549	0.534
	3	Nearest neighbor search w/ MaskCLIP [45] + SAM	0.705	0.629	0.536	0.498
	4	Normalizing flow [17] w/ MaskCLIP [2] + SAM	0.616	0.596	0.541	0.545
	5	Guided diffusion without timestep matching [38]	0.629	0.591	0.462	0.522
	6	Guided diffusion with forward timestep matching	0.697	0.578	0.462	0.514
Ablations	7	Ours with ResNet [46]	0.424	0.524	0.255	0.484
	8	Ours without SAM [4]	0.645	0.737	0.396	0.546
	9	Ours with RGB features	0.302	0.591	0.144	0.482
	10	Ours with DINOv2 [45]	0.685	0.598	0.416	0.518
	11	Ours with ViT [47]	0.692	0.571	0.255	0.493
	12	Ours with CLIP [1]	0.694	0.585	0.499	0.532
	13	Ours with DINOv1 [48]	0.665	0.622	0.416	0.521
Ours	14	Ours without FeatUp	<b>0.724</b>	<b>0.858</b>	0.475	<b>0.549</b>
	15	Ours (Reverse matching + MaskCLIP + SAM) with FeatUp	0.709	0.599	<b>0.568</b>	0.540

TABLE I: Anomaly detection accuracy on RUGD [15] dataset and the RELLIS-3D [16] dataset.  Our methods (in green) are compared against  baselines (in red), and  ablations (in blue) removing or changing one component of our method at a time. See text for details.

50,000 MaskCLIP pixel feature vectors sampled from training images. For a given feature vector of pixel  $i$  at test time, we output the distance of the nearest neighbor in the collected set. This is treated as the anomaly score. **Normalizing flow** [17, 20] (row 4): We train a normalizing flow on MaskCLIP pixel features from the training set similar to the GMM baseline. We follow Ancha et al. [17] and use a GMM as a stronger base distribution. Negative log-likelihoods of the normalizing flows, which can be computed exactly [50], are used as anomaly scores. **Diffusion guidance baselines**: We compare against vanilla guided diffusion with no timestep matching [38] (line 8b of Alg. 1), and forward timestep matching (lines 7c, 8c of Alg. 1). We also perform ablation experiments in rows 7-14, by either removing one of our contributions at a time, or changing the VLP feature space in the analysis component. We find that our method outperforms all baselines and ablations, on both datasets and evaluation metrics. On RUGD, we find that not using FeatUp [3] can improve quantitative performance. We hypothesize that this might be because FeatUp can average upsampled pixel features near segment boundaries, affecting the resulting cosine distances. However, FeatUp can be useful for certain applications where high-resolution is crucial when detecting anomalies. We use FeatUp when generating all visualizations in this paper.

We also apply our method to an under-water navigation dataset [11] where the training distribution contains images of the ocean bed, corals and plants. We are unable to perform a quantitative evaluation due to the lack of ground truth semantic segmentation labels. We qualitatively evaluate our method on anomalous images that contain fish, divers and robots in Fig. 1. We observe that the diffusion model makes interesting edits to the input image. It either removes anomalies altogether, or blends them into the background. Both types of edits are detected by our analysis pipeline. See our [project website](#) and appendix for more qualitative results.

## VIII. RELATED WORK

Prior work on anomaly detection [51] has predominantly focused on *discriminative* models that directly map pixels

to anomaly scores, classifying features at the pixel- or segment- level as in-distribution or anomalies [17–31] using principal component analysis (PCA) [52–54], random feature projection [55–57], novelty functions [58, 59], feature-space comparison of neural embeddings [60–63] and more recently, evidential uncertainty estimation [19, 20, 64] for off-road navigation [17, 65–67]. Other works that take an analysis-by-synthesis [36, 37] approach use *generative* models such as autoencoders [68, 69] and GANs [70–74]. However, these approaches can suffer from high false-positive rates [51, 68]. More recently, a small set of works have employed diffusion models for anomaly detection [75–78]. However, these works focus on specialized domains like medical imaging [75, 76] and industrial inspection [77, 78]; to the best of our knowledge, we are the first to use this approach for natural images in off-road navigation. Furthermore, Zhang et al. [77] train on synthetic anomaly examples, whereas our method does not require real or synthetic OOD examples during training.

## IX. CONCLUSIONS

In this work, we presented an *analysis-by-synthesis* approach for pixel-wise anomaly detection in off-road images. Given an input image, we used a diffusion model to *synthesize* an edited image that removes anomalies while keeping the remaining image unchanged. We then formulated anomaly detection as *analyzing* which image segments were modified by the diffusion model. We proposed a novel inference approach for guided diffusion by theoretically analyzing the ideal guidance gradient and deriving a principled approximation. Unlike prior methods, this approach bootstraps the diffusion model to predict edits to the image. Our editing technique is purely test-time and be integrated into existing workflows without re-training or finetuning. Finally, we presented a combination of vision-language foundation models to compare pixels in a more semantically meaningful feature space in order to identify segments that were modified by the diffusion model. We hope this work paves the way towards generative approaches for accurate and interpretable anomaly detection for off-road navigation.

## REFERENCES

- [1] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. **Learning transferable visual models from natural language supervision**. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 8748–8763, 2021.
- [2] Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. **MaskCLIP: Masked self-distillation advances contrastive language-image pretraining**. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10995–11005, 2023.
- [3] Stephanie Fu, Mark Hamilton, Laura E. Brandt, Axel Feldmann, Zhoutong Zhang, and William T. Freeman. **FeatUp: A Model-Agnostic Framework for Features at Any Resolution**. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.
- [4] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. **Segment Anything**. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026, 2023.
- [5] John Bares, Martial Hebert, Takeo Kanade, Eric Krotkov, Tom Mitchell, Reid Simmons, and William Whittaker. **Ambler: An autonomous rover for planetary exploration**. *Computer*, 22(6):18–26, 1989.
- [6] Mauro Massari, Giovanni Giardini, Franco Bernelli-Zazzera, et al. **Autonomous navigation system for planetary exploration rover based on artificial potential fields**. In *Proceedings of Dynamics and Control of Systems and Structures in Space (DCSSS) 6th Conference*, pages 153–162, 2004.
- [7] Jonas Frey, David Hoeller, Shehryar Khattak, and Marco Hutter. **Locomotion policy guided traversability learning using volumetric representations of complex environments**. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5722–5729. IEEE, 2022.
- [8] Xiangyun Meng, Nathan Hatch, Alexander Lambert, Anqi Li, Nolan Wagener, Matthew Schmittle, JoonHo Lee, Wentao Yuan, Zoey Chen, Samuel Deng, et al. **TerrainNet: Visual Modeling of Complex Terrain for High-speed, Off-road Navigation**. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [9] Amirreza Shaban, Xiangyun Meng, JoonHo Lee, Byron Boots, and Dieter Fox. **Semantic terrain classification for off-road autonomous driving**. In *Conference on Robot Learning*, pages 619–629. PMLR, 2022.
- [10] Ryan M Eustice, Oscar Pizarro, and Hanumant Singh. **Visually augmented navigation for autonomous underwater vehicles**. *IEEE Journal of oceanic Engineering*, 33(2):103–122, 2008.
- [11] Travis Manderson, Juan Camilo Gamboa, Stefan Wapnick, Jean-François Tremblay, Florian Shkurti, Dave Meger, and Gregory Dudek. **Vision-Based Goal-Conditioned Policies for Underwater Navigation in the Presence of Obstacles**. In *Proceedings of Robotics: Science and Systems (RSS)*, July 2020.
- [12] Abhinav Valada, Gabriel L Oliveira, Thomas Brox, and Wolfram Burgard. **Deep multispectral semantic scene understanding of forested environments using multimodal fusion**. In *2016 International Symposium on Experimental Robotics (ISER)*, pages 465–477. Springer, 2017.
- [13] Abhinav Valada, Johan Vertens, Ankit Dhall, and Wolfram Burgard. **AdapNet: Adaptive semantic segmentation in adverse environmental conditions**. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4644–4651. IEEE, 2017.
- [14] Tianrui Guan, Divya Kothandaraman, Rohan Chandra, Adarsh Jagan Sathyamoorthy, Kasun Weerakoon, and Dinesh Manocha. **GA-Nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments**. *IEEE Robotics and Automation Letters*, 7(3):8138–8145, 2022.
- [15] Maggie Wigness, Sungmin Eum, John G Rogers, David Han, and Heesung Kwon. **A RUGD dataset for autonomous navigation and visual perception in unstructured outdoor environments**. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5000–5007. IEEE, 2019. URL <http://rugd.vision/>.
- [16] Peng Jiang, Philip Osteen, Maggie Wigness, and Srikanth Saripalli. **RELLIS-3D dataset: Data, benchmarks and analysis**. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1110–1116. IEEE, 2021. URL <http://www.unmannedlab.org/research/RELLIS-3D>.
- [17] Siddharth Ancha, Philip R. Osteen, and Nicholas Roy. **Deep Evidential Uncertainty Estimation for Semantic Segmentation under Out-Of-Distribution Obstacles**. In *IEEE International Conference on Robotics and Automation (ICRA)*, Yokohama, Japan, May 2024.
- [18] Dennis Thomas Ulmer, Christian Hardmeier, and Jes Frellesen. **Prior and Posterior Networks: A Survey on Evidential Deep Learning Methods For Uncertainty Estimation**. *Transactions on Machine Learning Research (TMLR)*, 2023. ISSN 2835-8856.
- [19] 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.
- [20] Bertrand Charpentier, Oliver Borchert, Daniel Zügner, Simon Geisler, and Stephan Günnemann. **Natural posterior network: Deep Bayesian predictive uncertainty for exponential family distributions**. In *International Conference on Learning Representations*, 2022.

- [21] Matej Grić, Petra Bevandić, Zoran Kalafatić, and Siniša Šegvić. **Dense out-of-distribution detection by robust learning on synthetic negative data.** *Sensors*, 24(4):1248, 2024.
- [22] Yuyuan Liu, Choubo Ding, Yu Tian, Guansong Pang, Vasileios Belagiannis, Ian Reid, and Gustavo Carneiro. **Residual pattern learning for pixel-wise out-of-distribution detection in semantic segmentation.** In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1151–1161, 2023.
- [23] Dan Hendrycks and Kevin Gimpel. **A baseline for detecting misclassified and out-of-distribution examples in neural networks.** *arXiv preprint arXiv:1610.02136*, 2016.
- [24] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. **Simple and scalable predictive uncertainty estimation using deep ensembles.** *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- [25] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. **A simple unified framework for detecting out-of-distribution samples and adversarial attacks.** *Advances in Neural Information Processing Systems*, 31, 2018.
- [26] Denis Gudovskiy, Tomoyuki Okuno, and Yohei Nakata. **Concurrent misclassification and out-of-distribution detection for semantic segmentation via energy-based normalizing flow.** In *Uncertainty in Artificial Intelligence*, pages 745–755. PMLR, 2023.
- [27] Yu Tian, Yuyuan Liu, Guansong Pang, Fengbei Liu, Yuanhong Chen, and Gustavo Carneiro. **Pixel-wise energy-biased abstention learning for anomaly segmentation on complex urban driving scenes.** In *European Conference on Computer Vision*, pages 246–263. Springer, 2022.
- [28] Petra Bevandić, Ivan Krešo, Marin Oršić, and Siniša Šegvić. **Dense open-set recognition based on training with noisy negative images.** *Image and vision computing*, 124:104490, 2022.
- [29] Robin Chan, Matthias Rottmann, and Hanno Gottschalk. **Entropy maximization and meta classification for out-of-distribution detection in semantic segmentation.** In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5128–5137, 2021.
- [30] Nazir Nayal, Misra Yavuz, Joao F Henriques, and Fatma Güney. **RbA: Segmenting unknown regions rejected by all.** In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 711–722, 2023.
- [31] Jan Ackermann, Christos Sakaridis, and Fisher Yu. **Maskomaly: Zero-shot mask anomaly segmentation.** *arXiv preprint arXiv:2305.16972*, 2023.
- [32] Jonathan Ho, Ajay Jain, and Pieter Abbeel. **Denoising diffusion probabilistic models.** *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020.
- [33] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. **Diffusion models in vision: A survey.** *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2023.
- [34] Prafulla Dhariwal and Alexander Nichol. **Diffusion models beat GANs on image synthesis.** *Advances in Neural Information Processing Systems (NeurIPS)*, 34: 8780–8794, 2021.
- [35] 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 (CVPR)*, pages 10684–10695, 2022.
- [36] Thomas G Bever and David Poeppel. **Analysis by synthesis: a (re-) emerging program of research for language and vision.** *Biolinguistics*, 4(2-3):174–200, 2010.
- [37] Alan Yuille and Daniel Kersten. **Vision as Bayesian inference: analysis by synthesis?** *Trends in cognitive sciences*, 10(7):301–308, 2006.
- [38] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. **Deep unsupervised learning using nonequilibrium thermodynamics.** In *International conference on machine learning (ICML)*, pages 2256–2265. PMLR, 2015.
- [39] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. **Score-Based Generative Modeling through Stochastic Differential Equations.** In *9th International Conference on Learning Representations (ICLR), 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- [40] Yilun Du, Conor Durkan, Robin Strudel, Joshua B Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, and Will Sussman Grathwohl. **Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and MCMC.** In *International Conference on Machine Learning (ICML)*, pages 8489–8510. PMLR, 2023.
- [41] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. **DALL-E: Zero-shot text-to-image generation.** In *International conference on machine learning (ICML)*, pages 8821–8831. PMLR, 2021.
- [42] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. **Diffusion policy: Visuomotor policy learning via action diffusion.** In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [43] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. **Octo: An open-source generalist robot policy.** In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024.
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. **Image quality assessment: from error visibility to structural similarity.** *IEEE Transactions on*



- Image Processing (TIP)*, 13(4):600–612, 2004.
- [45] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, et al. **DINOv2: Learning Robust Visual Features without Supervision**. *IEEE Transactions on Machine Learning Research (TMLR)*, 2023.
- [46] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. **Deep residual learning for image recognition**. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [47] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. **An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale**. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.
- [48] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. **Emerging Properties in Self-Supervised Vision Transformers**. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9650–9660, 2021.
- [49] Robin Chan, Krzysztof Lis, Svenja Uhlemeyer, Hermann Blum, Sina Honari, Roland Siegwart, Pascal Fua, Mathieu Salzmann, and Matthias Rottmann. **Segment-MeIfYouCan: A benchmark for anomaly segmentation**. In J. Vanschoren and S. Yeung, editors, *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS) Workshops Track on Datasets and Benchmarks*, volume 1, 2021.
- [50] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. **Normalizing flows for probabilistic modeling and inference**. *The Journal of Machine Learning Research (JMLR)*, 22(1):2617–2680, 2021.
- [51] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for anomaly detection: A review. *ACM Computing Surveys (CSUR)*, 54(2):1–38, 2021.
- [52] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):1–37, 2011.
- [53] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*, pages 583–588, 1997.
- [54] Chong Zhou and Randy C Paffenroth. Anomaly detection with robust deep autoencoders. In *Proceedings of the ACM Special-Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International conference*, pages 665–674, 2017.
- [55] Ping Li, Trevor J Hastie, and Kenneth W Church. Very sparse random projections. In *Proceedings of the ACM Special-Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International conference*, pages 287–296, 2006.
- [56] Guansong Pang, Longbing Cao, Ling Chen, and Huan Liu. Learning representations of ultrahigh-dimensional data for random distance-based outlier detection. In *Proceedings of the ACM Special-Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International conference*, pages 2041–2050, 2018.
- [57] Tomáš Pevný. Loda: Lightweight on-line detector of anomalies. *Machine Learning*, 102:275–304, 2016.
- [58] Marco AF Pimentel, David A Clifton, Lei Clifton, and Lionel Tarassenko. A review of novelty detection. *Signal Processing*, 99:215–249, 2014.
- [59] Charles Andrew Richter. *Autonomous navigation in unknown environments using machine learning*. PhD thesis, Massachusetts Institute of Technology, 2017.
- [60] Sarah M Erfani, Sutharshan Rajasegarar, Shanika Karunasekera, and Christopher Leckie. High-dimensional and large-scale anomaly detection using a linear one-class svm with deep learning. *Pattern Recognition*, 58:121–134, 2016.
- [61] Radu Tudor Ionescu, Fahad Shahbaz Khan, Mariana-Iuliana Georgescu, and Ling Shao. Object-centric auto-encoders and dummy anomalies for abnormal event detection in video. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7842–7851, 2019.
- [62] Dan Xu, Jingkuan Song, Yan Yan, E Ricci, Niculae Sebe, et al. Learning deep representations of appearance and motion for anomalous event detection. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 1–12, 2015.
- [63] Wenchao Yu, Wei Cheng, Charu C Aggarwal, Kai Zhang, Haifeng Chen, and Wei Wang. Netwalk: A flexible deep embedding approach for anomaly detection in dynamic networks. In *Proceedings of the ACM Special-Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International conference*, pages 2672–2681, 2018.
- [64] Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. **Deep evidential regression**. *Advances in Neural Information Processing Systems*, 33:14927–14937, 2020.
- [65] Xiaoyi Cai, Siddharth Ancha, Lakshay Sharma, Philip R Osteen, Bernadette Bucher, Stephen Phillips, Jiuguang Wang, Michael Everett, Nicholas Roy, and Jonathan P How. **EVORA: Deep Evidential Traversability Learning for Risk-Aware Off-Road Autonomy**. *IEEE Transactions on Robotics*, 40:3756–3777, 2024.
- [66] Xiaoyi Cai, James Queeney, Tong Xu, Aniket Datar, Chenhui Pan, Max Miller, Ashton Flather, Philip R Osteen, Nicholas Roy, Xuesu Xiao, et al. **PIETRA: Physics-Informed Evidential Learning for Traversing Out-of-Distribution Terrain**. *arXiv preprint arXiv:2409.03005*, 2024.
- [67] Kshitij Sirohi, Sajad Marvi, Daniel Büscher, and Wolfram Burgard. **Uncertainty-aware panoptic segmentation**. *IEEE Robotics and Automation Letters (RAL)*, 8(5):2629–

2636, 2023.

- [68] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021.
- [69] Xu Dan, Elisa Ricci, Yan Yan, Jingkuan Song, Nicu Sebe, et al. Learning deep representations of appearance and motion for anomalous event detection. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 8–1, 2015.
- [70] Takahiro Nakao, Shouhei Hanaoka, Yukihiro Nomura, Masaki Murata, Tomomi Takenaga, Soichiro Miki, Takeyuki Watadani, Takeharu Yoshikawa, Naoto Hayashi, and Osamu Abe. Unsupervised deep anomaly detection in chest radiographs. *Journal of Digital Imaging*, 34:418–427, 2021.
- [71] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *Proceedings of the International Conference on Information Processing in Medical Imaging (IPMI)*, pages 146–157, 2017.
- [72] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. fanogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical Image Analysis*, 54:30–44, 2019.
- [73] Samet Akçay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 622–637, 2019.
- [74] Samet Akçay, Amir Atapour-Abarghouei, and Toby P Breckon. Skip-ganomaly: Skip connected and adversarially trained encoder-decoder anomaly detection. In *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2019.
- [75] Julia Wolleb, Florentin Bieder, Robin Sandkühler, and Philippe C Cattin. **Diffusion models for medical anomaly detection**. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 35–45. Springer, 2022.
- [76] Walter HL Pinaya, Mark S Graham, Robert Gray, Pedro F Da Costa, Petru-Daniel Tudosiu, Paul Wright, Yee H Mah, Andrew D MacKinnon, James T Teo, Rolf Jager, et al. **Fast unsupervised brain anomaly detection and segmentation with diffusion models**. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 705–714. Springer, 2022.
- [77] Hui Zhang, Zheng Wang, Zuxuan Wu, and Yu-Gang Jiang. **DiffusionAD: Norm-guided one-step denoising diffusion for anomaly detection**. *arXiv preprint arXiv:2303.08730*, 2023.
- [78] Arian Mousakhan, Thomas Brox, and Jawad Tayyub. **Anomaly Detection with Conditioned Denoising Diffusion Models**. *arXiv preprint arXiv:2305.15956*, 2023.
- [79] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. **CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning**. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.