

AnomalyVFM – Transforming Vision Foundation Models into Zero-Shot Anomaly Detectors

Matic Fučka*

Vitjan Zavrtanik*

Danijel Skočaj

University of Ljubljana, Faculty of Computer and Information Science, Slovenia

{matic.fucka, vitjan.zavrtanik, danijel.skocaj}@fri.uni-lj.si

Abstract

Zero-shot anomaly detection aims to detect and localise abnormal regions in the image without access to any in-domain training images. While recent approaches leverage vision–language models (VLMs), such as CLIP, to transfer high-level concept knowledge, methods based on purely vision foundation models (VFMs), like DINOv2, have lagged behind in performance. We argue that this gap stems from two practical issues: (i) limited diversity in existing auxiliary anomaly detection datasets and (ii) overly shallow VFM adaptation strategies. To address both challenges, we propose AnomalyVFM, a general and effective framework that turns any pretrained VFM into a strong zero-shot anomaly detector. Our approach combines a robust three-stage synthetic dataset generation scheme with a parameter-efficient adaptation mechanism, utilising low-rank feature adapters and a confidence-weighted pixel loss. Together, these components enable modern VFMs to substantially outperform current state-of-the-art methods. More specifically, with RADIO as a backbone, AnomalyVFM achieves an average image-level AUROC of 94.1% across 9 diverse datasets, surpassing previous methods by significant 3.3 percentage points. [Project Page](#)

1. Introduction

Visual anomaly detection aims to identify abnormal regions at test time while training only on anomaly-free images. This represents a foundational task within manufacturing [4, 69, 89], medical imaging [21, 28, 60] and road obstacle detection [14, 66, 67]. In industrial inspection, it is typically assumed [18, 59] that many normal images are available during training. However, practical deployments often require detecting anomalies on arbitrary object classes without any or very few images. This extremely challenging setting has motivated recent interest in few-shot and zero-shot anomaly detection. Few-shot methods [39, 64, 86] re-




Figure 1. Vision–language models excel in zero-shot anomaly detection thanks to their high-level concept knowledge, but purely visual foundation models hold untapped potential. AnomalyVFM unlocks this potential by addressing the two practical limitations that hinder VFM underperformance: suboptimal training sets and suboptimal fine-tuning procedures.

quire a handful of normal images of the object class, while zero-shot methods [9, 53, 84] must generalise to unseen object classes with no in-domain images at all.

State-of-the-art zero-shot approaches [9, 53, 84] use vision–language models (VLMs) such as CLIP [54]. They typically use auxiliary anomaly detection datasets to train the model to output text embeddings that encode generic notions of normality and abnormality. Pretraining with image-text supervision introduces valuable high-level concept knowledge, which facilitates generalisation across different object categories. By contrast, pure vision foundation models (VFMs) such as DINOv2 [49] encode strong visual representations but have so far trailed behind VLM-based methods when used as a basis for zero-shot anomaly detection. This gap raises the question: Can VFMs, which are arguably better suited to the fundamentally visual nature of

*Equal contribution.

anomaly detection, be transformed into competitive zero-shot detectors?

We argue that two practical limitations explain why VFMs have underperformed in prior zero-shot work. First, existing auxiliary anomaly datasets [4, 89] lack sufficient diversity and coverage of realistic defects, which are required for training VFMs. This is not a problem for VLMs due to their high-level concept knowledge. When the model must generalise to arbitrary object classes, limited dataset diversity prevents learning broadly applicable cues. Second, most prior VFM adaptations [8, 9, 49] fine-tune only a small output head with simple pixel-wise losses, leaving the model’s internal visual representations essentially unchanged. This makes it challenging for the model to accurately learn the features necessary to distinguish between normal and abnormal appearances across various objects.

To address both points, we propose AnomalyVFM, a practical framework that transforms any modern VFM into a robust zero-shot anomaly detector. AnomalyVFM has two core components. First, a three-stage synthetic dataset generator that uses modern image generation models (e.g., FLUX [36]) to (i) create diverse anomaly-free object images, (ii) synthesise a wide variety of local defects by inpainting at sampled locations, and (iii) filter generated samples using a feature-based verification step to ensure the presence and relevance of defects. This creates a large and diverse auxiliary training set containing many object/background combinations. Second, we introduce a parameter-efficient adaptation mechanism tailored for VFMs: low-rank feature adapters are injected throughout the (transformer) backbone, coupled with a lightweight decoder and a confidence-weighted pixel loss that down-weights ambiguous supervision. The adapters enable the VFM to evolve its internal visual representations (not just the final head) with minimal additional parameters. The decoder converts these adapted features into pixel-level anomaly scores, and the confidence-weighted loss limits the impact of noisy gradients from imperfect synthetic labels. Crucially, AnomalyVFM is model-agnostic and practical: it can be applied to any pretrained VFM with a transformer backbone (as shown in Figure 1).

In summary, our contributions are:

- As our main contribution, we introduce AnomalyVFM, an effective framework that transforms any pretrained VFM into a competitive zero-shot anomaly detector using a parameter-efficient adaptation scheme and a synthetically generated dataset containing diverse data.
- As our secondary contribution, we propose a scalable scheme for generating synthetic anomaly detection datasets. We design a three-stage synthesis process that leverages modern generative models to produce diverse object instances, realistic local defects, and automatic feature-based verification to ensure data quality. This

yields data that is better suited for finetuning VFMs in comparison to existing datasets.

We validate our contributions by evaluating the proposed approach on nine standard industrial anomaly detection benchmarks, surpassing the previous best zero-shot anomaly detection methods by a significant 3.3 percentage points (p. p.) in image-level AUROC and 0.9 p. p. in pixel-level AUROC. Additionally, we demonstrate that AnomalyVFM also generalises well on medical anomaly detection benchmarks, even though it was not finetuned for this purpose. Finally, we demonstrate the versatility of AnomalyVFM by finetuning it on a few normal samples. With this, it is able to match the performance of state-of-the-art models in the few-shot regime.

2. Related Work

Anomaly detection can be categorized into several paradigms. Most commonly they are divided in reconstructive [3, 79], discriminative [18, 44, 78, 80] and embedding-based methods [13, 59]. Reconstructive approaches [2, 50] are trained to reconstruct anomaly-free images. Since the learnt models never see anomalous examples, the assumption is that such examples will be poorly reconstructed, making them detectable by reconstruction error. Discriminative methods [19, 57] are trained with synthetic anomalies under the assumption that this will generalise well to actual anomalies. Embedding-based methods [15, 85] fit a simple normality model, such as a coreset, on top of the features extracted from a pretrained encoder.

Training from generated data is common in solving or evaluating text-based tasks [6, 51, 71, 77]. It has also started being adapted in computer vision in areas such as video generation [83] and dynamic scene reconstruction [10], but has not yet seen widespread usage. For training data generation, powerful diffusion models [58] or flow-matching [17] approaches are commonly used [31, 35, 42]. In [42], a diffusion model is utilised to generate samples according to a set of labels for out-of-distribution object detection. The method focuses on outlier generation, encompassing different object classes, but overlooks near-in-distribution cases.

Anomaly synthesis has started to receive more attention in the past few years. The field was started by DRÆM [78], which synthesised anomalies by cropping and pasting parts of images from an external dataset [1]. Some approaches improved upon this later by refining the way the external images are augmented or moving the synthetic anomaly generation to the latent space [44, 56, 81]. Some of the later approaches improved upon the realism of the generated samples by using modern generative models, such as generative adversarial networks [16, 82] or diffusion models [24, 76]. However, all of these methods typically require substantial amounts of normal and/or abnormal data. Furthermore, these methods can only generate samples simi-

lar to the training set, i.e., seen anomalies, but fail to generate unseen anomalies. This makes them unsuitable for zero-shot anomaly detection. Different from previous approaches, our generation approach does not require any samples, neither normal nor abnormal.

Zero-shot anomaly detection methods detect anomalies at inference while never seeing an instance of the observed object during training. Most recent zero-shot anomaly detection methods [9, 11, 26, 84, 87] focus on utilising the general object appearance knowledge embedded in vision-language [54] models. A minority of methods have utilised Vision (only) Foundation Models for the task of zero-shot anomaly detection. SAA [8] uses the GroundingDINO [43] and SAM [34] with handcrafted anomaly prompts to directly segment anomalies. In [37], the method models the distribution of object appearance within the batch to detect anomalous samples. However, assumptions about the contents of a specific batch are often violated in real-world scenarios. Some methods [9, 49] have also investigated tuning VFMs directly on an auxiliary dataset, but achieved suboptimal results.

In this paper, we demonstrate that it is possible to achieve state-of-the-art performance using a pretrained VFM fine-tuned correctly on a diverse enough dataset.

3. Dataset Generation Scheme

To address the issue of data diversity, a collection of realistic images of objects, both with and without anomalies, is necessary. Additionally, each anomalous image should be accompanied by a pixel-level annotation. To do this, a three-stage generation scheme is proposed. (i) First, the initial image of the object is generated. (ii) Then, a realistic defect is inpainted on top of the object. (iii) Ultimately, the anomaly segmentation map is generated by subtracting the features of the normal image from those of the anomalous image, and based on this, poorly generated images are filtered. Each of these steps will be described in detail below. Some examples of generated samples are shown in Figure 2.

Anomaly-free Image Generation To generate the initial (anomaly-free) image I (Figure 3, I), an image generation model G is prompted with an anomaly-free text prompt p :

$$I = G(p). \quad (1)$$

As the image generation model G , the flow-matching-based FLUX model [36] is used in all experiments unless stated otherwise. Anomaly-free text prompt p is constructed as follows:

A close-up photo of [Object] for industrial visual inspection. Top-down view. Centered. [Texture] background.

The [Object] and [Texture] are replaced by an object or background class from a list of 100 objects and 50 backgrounds generated by an LLM (in our case GPT-4o [25]).

Anomalous Image Generation The anomalous image I_a is generated using the anomaly-free image I . To do so, the rough anomaly location R (in our case, a rectangle) has to be determined. As the first step, the foreground object mask M_{fg} must be extracted. In our case, this is done using a pretrained salient object segmentation network IS-Net [52] (Figure 3, M_{fg}). To generate R , the location of the anomaly (x, y) is first sampled on the foreground object, i.e., as a random positive pixel in M_{fg} . The initial location serves as the centre of the anomaly rectangle R (Figure 3, R). The width and height are uniformly sampled according to the desired anomaly width (w_{min}, w_{max}) and height (h_{min}, h_{max}) parameters:

$$w \sim \mathcal{U}(w_{min}, w_{max}), h \sim \mathcal{U}(h_{min}, h_{max}). \quad (2)$$

The anomaly is then generated by prompting the model G with an anomalous prompt p_a , while restricting the generation to the region R and maintaining I in other regions, i.e. inpainting. To generate an anomalous version of the generated image I , the prompt p_a additionally contains anomalous descriptions:

A close-up photo of a [Anomaly]
[Object] for industrial visual
inspection. Top-down view. Centered.
[Texture] background.

The [Anomaly] tag is replaced with a description of an anomaly, such as cracked, damaged, smudged, rotten. A list of [Anomaly] descriptions for each [Object] is generated by an LLM (again GPT-4o [25]). This ensures the [Anomaly] is relevant for the object. The [Object], [Anomaly] and [Texture] lists are listed in the Supplementary material.

No inpainting-specific models are used; instead, the characteristics of the iterative generation process (diffusion or flow matching) are used, using the RePaint approach [45]. The iteratively generated image I_a (Figure 3, I_a) contains the object generated in I with an anomaly in region R whose visual appearance follows the prompt in p_a . However, accurate prompt adherence is not a solved problem in image generation [17], so some generated images may not contain anomalies at all. To address this, a filtering process is proposed which removes the vast majority of examples where anomalies are not generated in I_a .

Dataset filtering To filter out examples with poor adherence to p_a , a comparison between the anomaly-free I and the corresponding anomalous I_a is performed. First, DINOv2 [49] features are extracted from I and I_a , obtaining




Figure 2. Examples of generated anomaly-free images I , anomalous images I_a and corresponding masks M .




Figure 3. Dataset generation pipeline. The image I is generated using a text-conditioned image generation model. Then, the foreground mask M_{fg} is extracted and an anomalous region R is sampled from it. Then, the anomalous image I_a is generated by inpainting an anomaly inside R . Finally, features are extracted from I and I_a , and then compared and thresholded to obtain M .

f and f_a , respectively. The extracted features are then compared using cosine distance, obtaining a distance map M_d . The maximum value of M_d is obtained as the distance score D . The mask is binarised according to a threshold T to obtain the final mask M . The generated sample is accepted if the distance D exceeds a set threshold T . The idea behind this filtering process is that if the anomaly generation step fails to adhere to p_a , an anomaly-free object region will be generated in region R instead of the anomaly. In this case, the inpainted region will be closer to the anomaly-free object appearance distribution, so the distance between f and f_a should be smaller, enabling the detection of failed examples.

Generated triplets containing an anomaly-free image I , an anomalous example I_a , and the corresponding anomaly mask M can be seen in Figure 2.

4. AnomalyVFM

Recent attempts [9] at using VFMs for zero-shot anomaly detection have only appended a simple MLP on top of the method to generate an anomaly mask, disregarding the adaptation of internal values, the design of the decoder and the losses used to train them. To improve upon this, an effective and parameter-efficient finetuning technique is proposed. First, we improve upon the decoder and inject feature adaptation modules within the VFM to enable adaptation of internal layers. In addition, we propose a confidence-

weighted loss to mitigate potential ambiguity caused by inaccurate labels. The adaptation network and the confidence-weighted loss will be described in detail below. The architecture of AnomalyVFM is shown in Figure 4.

Feature Adaptation Module and Decoder The input image I is input into the pretrained backbone F . Each transformer block b of F is integrated with a Feature Adaptation Module. More specifically, we integrate aLoRA [23] block into the attention mechanism [65] by injecting it into the query, value, and output projection layers, as shown in Figure 4. If not stated otherwise, the rank of LoRA is equal to 64. The extracted features f from the final block of the backbone F are reshaped to f_r . Then, f_r is input into a small convolutional decoder that upsamples the features. The decoder is composed of two sequential upsampling blocks, which are constructed as a convolutional layer, a GroupNorm layer [74], a ReLU activation function, and a bilinear upsampling operation. A final Convolutional Layer is used to output both the output anomaly segmentation map M_o and the confidence map c . The [CLS] tokens of the backbone are input into a simple linear layer, which predicts an image-level anomaly score A_o .

Confidence-weighted loss The Feature Adaptation Modules, Anomaly Decoder and Anomaly Score Predictor are trained jointly. For the image-level loss $\mathcal{L}_{img}$, Focal Loss [41] is used, while the base loss for segmentation $\mathcal{L}_{base}$ is a combination of Focal loss and $\mathcal{L}_1$ loss, following recent anomaly detection methods [75]:

$$\mathcal{L}_{base} = \mathcal{L}_1(M_o, M_{GT}) + \beta * \mathcal{L}_{focal}(M_o, M_{GT}), \quad (3)$$

where β is equal to 5. Additionally, to better handle the noisy segmentation masks that occur during data generation and any ambiguities in the ground truth masks, we weight the loss with the confidence output from the anomaly decoder, similar to 3D reconstruction methods [33, 70]. More specifically, the segmentation loss is defined as follows:

$$\mathcal{L}_{seg} = \mathcal{L}_{base}(M_o, M_{GT}) * C - \alpha \log(C), \quad (4)$$

where C is defined as $C = 1 + \exp(c)$, where c is the confidence map predicted by the decoder, and α is equal to 0.1.




Figure 4. Architecture of AnomalyVFM. All additions to the base VFM are colored in blue.

Table 1. Generalisation across different VFMs. Improvement over the baseline is shown in green. SD stands for Synthetic dataset, and FA stands for Feature Adaptors. The average results across 9 industrial datasets are reported.

Model	Additions		Image-level		Pixel-level	
	SD	FA	AUROC	F_1 -Max	AUROC	F_1 -Max
DINOv2 [49]	✓		83.0	78.9	80.4	23.9
			86.4 $\uparrow$ 3.4	79.3 $\uparrow$ 0.4	93.2 $\uparrow$ 12.8	41.2 $\uparrow$ 17.3
		✓	89.3 $\uparrow$ 6.3	82.7 $\uparrow$ 3.7	82.5 $\uparrow$ 2.1	27.1 $\uparrow$ 3.2
DINOv3 [61]	✓		90.2 $\uparrow$ 7.2	83.4 $\uparrow$ 4.5	93.4 $\uparrow$ 13.0	41.7 $\uparrow$ 17.8
		✓	85.3	80.1	88.0	32.5
		✓	89.0 $\uparrow$ 3.7	82.0 $\uparrow$ 1.9	95.0 $\uparrow$ 7.0	44.9 $\uparrow$ 12.4
RADIO [55]	✓		90.5 $\uparrow$ 5.2	83.9 $\uparrow$ 3.8	90.2 $\uparrow$ 2.2	39.3 $\uparrow$ 7.3
		✓	91.5 $\uparrow$ 6.2	84.7 $\uparrow$ 4.6	95.3 $\uparrow$ 7.2	44.6 $\uparrow$ 12.1
		✓	89.1	83.9	84.9	30.8
	✓		92.1 $\uparrow$ 3.0	87.2 $\uparrow$ 3.3	95.9 $\uparrow$ 11.0	43.2 $\uparrow$ 12.6
		✓	93.0 $\uparrow$ 3.9	88.9 $\uparrow$ 5.0	95.2 $\uparrow$ 10.3	42.8 $\uparrow$ 12.0
		✓	94.1 $\uparrow$ 5.0	87.6 $\uparrow$ 3.7	96.9 $\uparrow$ 12.0	44.3 $\uparrow$ 13.5

The full loss is the sum $\mathcal{L} = \mathcal{L}_{seg} + \mathcal{L}_{img}$.

At inference, the image I is passed through the model, which directly returns both the output anomaly segmentation mask M_o and the image-level anomaly score A_o .

5. Experiments

5.1. Datasets

We evaluate AnomalyVFM on 9 industrial and 9 medical anomaly detection datasets as standard in other zero-shot methods [9, 15]. For industrial anomaly detection, MVTec AD [4], VisA [89], BTAD [48], MPDD [27], Real-IAD [69], KSDD [62], KSDD2 [7], DAGM [72], and DTD-Synthetic [1] were used, and for medical anomaly detection, HeadCT [60], BrainMRI [30], BR35H [21], ISIC [12], ClinicDB [5], ColonDB [63], Kvasir [28], Endo [22] and TN3K [20] were used. The evaluation metrics follow AdaCLIP [9], where the AUROC and F1-max are used for image-level anomaly detection, and the pixel-wise AUROC and the pixel-wise F1-max are used for anomaly localisation. We compare AnomalyVFM to recent state-of-the-art

approaches that are trained on auxiliary data. More specifically, when evaluated on the MVTec AD [4] dataset, the recent zero-shot AD methods are trained on the VisA [89] test set and are trained on the MVTec AD test set when evaluated on other datasets. In contrast, AnomalyVFM is trained solely on automatically generated data.

5.2. Implementation Details

In the data generation pipeline, the FLUX [36] conditional image generation model is used. The (w_{min}, w_{max}) and (h_{min}, h_{max}) are set to $(50, 350)$ in all experiments, when generating images of dimension 1024×1024 . The filtering threshold, T , is set to 0.3 in all experiments. A synthetic dataset of 10,000 images was generated for all experiments, unless stated otherwise. More details about the generated dataset are provided in the Supplementary Material.

Zero-shot anomaly detection training is performed on generated data. AnomalyVFM is trained for 500 iterations with a batch size of 32 using the AdamW optimiser and a learning rate of 10^{-4} . The RADIOv2.5 [55] ViT-L with a patch-size of 16 is used as the backbone for most experiments. Since RADIO has been trained on multiple resolutions, the input images are resized to 768×768 for training and evaluation. The confidence parameter α in Equation 4 is set to 0.1 in all experiments.

5.3. Generalisation of the proposed framework

First, we evaluate our contribution across a set of diverse VFMs to verify our claims. More specifically, we use DINOv2 [49], DINOv3 [61] and RADIO [55]. We evaluate each VFM in 4 different settings to verify our contribution. We modulate two settings: the training dataset and the adaptation strategy. For the dataset, we either follow the standard practice [9, 53, 84] of training on the test set of MVTec AD (the model is trained on the test set of VisA when evaluated on MVTec AD) or the dataset generated with the proposed synthetic generation procedure. For the adaptation strategy,

Table 2. Comparisons of zero-shot anomaly detection methods on industrial inspection datasets. The best performance is colored in red and the second best in blue.

Metric	Dataset	SAA [8] ToC'25	WinCLIP [26] CVPR'23	AnomalyCLIP [84] ICLR'24	AdaCLIP [9] ECCV'24	AACLIP [47] CVPR'25	Bayes-PFL [53] CVPR'25	FAPrompt [88] ICCV'25	AnomalyVFM
Image-level (AUROC, max-F1)	MVTec AD	(63.5, 87.4)	(91.8, 92.9)	(91.6, 92.7)	(89.2, 90.6)	(90.5, 90.4)	(92.3, 93.1)	(91.1, 92.2)	(94.9, 94.1)
	VisA	(67.1, 75.9)	(78.1, 80.7)	(82.0, 80.4)	(85.8, 83.1)	(84.6, 78.8)	(87.0, 84.1)	(82.8, 81.3)	(93.6, 90.1)
	BTAD	(59.0, 89.7)	(68.2, 67.8)	(88.2, 83.8)	(88.6, 88.2)	(94.8, 93.7)	(93.2, 91.9)	(90.7, 88.1)	(96.0, 91.0)
	MPDD	(42.7, 73.9)	(61.4, 77.5)	(77.5, 80.4)	(76.0, 82.5)	(75.1, 79.8)	(81.2, 83.5)	(76.6, 80.4)	(85.5, 87.8)
	RealIAD	(51.4, 64.6)	(74.7, 69.8)	(78.7, 80.0)	(79.2, 73.5)	(81.3, 76.4)	(85.2, 78.7)	(81.6, 75.2)	(88.0, 81.6)
	KSDD	(68.6, 37.6)	(93.3, 79.0)	(84.5, 71.1)	(97.1, 90.7)	(69.3, 57.1)	(88.2, 56.0)	(81.3, 71.1)	(92.5, 69.7)
	KSDD2	(91.6, 67.0)	(94.2, 71.5)	(94.1, 80.0)	(95.9, 86.7)	(95.9, 84.4)	(97.3, 87.6)	(95.6, 84.8)	(97.1, 79.2)
	DAGM	(87.1, 88.8)	(91.8, 87.6)	(97.7, 90.1)	(99.1, 97.5)	(93.2, 79.4)	(97.7, 95.7)	(97.3, 89.3)	(99.6, 95.8)
	DTD	(94.4, 93.5)	(95.1, 94.1)	(93.9, 93.6)	(95.5, 94.7)	(90.4, 92.8)	(95.1, 95.1)	(95.9, 94.7)	(99.4, 99.0)
Average		(69.5, 75.4)	(83.2, 80.1)	(87.6, 83.6)	(89.6, 87.5)	(86.1, 81.4)	(90.8, 85.1)	(88.1, 84.1)	(94.1, 87.6)
Pixel-level (AUROC, max-F1)	MVTec AD	(75.5, 38.1)	(88.7, 43.4)	(91.1, 39.1)	(88.7, 43.4)	(91.4, 46.4)	(91.8, 49.0)	(90.8, 39.3)	(92.7, 45.2)
	VisA	(76.5, 31.6)	(95.5, 37.7)	(95.5, 28.3)	(95.5, 37.7)	(94.8, 30.2)	(95.6, 34.3)	(95.6, 27.6)	(96.2, 31.2)
	BTAD	(65.8, 14.8)	(92.1, 51.7)	(94.2, 49.7)	(92.1, 51.7)	(97.3, 55.1)	(93.9, 52.0)	(95.8, 52.6)	(92.3, 49.7)
	MPDD	(81.7, 18.9)	(96.1, 34.9)	(96.5, 34.2)	(96.1, 32.8)	(96.7, 30.0)	(97.8, 35.0)	(95.5, 31.9)	(97.0, 38.1)
	RealIAD	(73.5, 4.5)	(87.2, 10.8)	(96.3, 39.0)	(97.2, 43.0)	(96.2, 40.2)	(97.2, 41.2)	(96.2, 38.3)	(96.4, 40.4)
	KSDD	(78.8, 6.6)	(97.7, 54.5)	(90.6, 42.5)	(97.7, 54.5)	(87.1, 28.0)	(96.5, 6.6)	(93.1, 47.2)	(99.0, 10.1)
	KSDD2	(79.9, 63.4)	(94.4, 23.9)	(98.5, 59.8)	(98.5, 67.0)	(99.5, 63.4)	(97.0, 62.0)	(99.1, 60.4)	(99.3, 55.9)
	DAGM	(91.5, 57.5)	(91.5, 57.5)	(95.6, 58.9)	(95.5, 57.5)	(96.2, 53.3)	(95.9, 49.8)	(98.6, 60.2)	(99.4, 61.3)
	DTD	(97.9, 71.6)	(97.9, 71.6)	(97.9, 62.2)	(97.9, 71.6)	(95.8, 59.6)	(98.4, 65.2)	(98.1, 61.9)	(99.4, 66.5)
Average		(80.1, 34.1)	(93.5, 42.9)	(95.1, 46.0)	(95.0, 51.0)	(95.0, 45.1)	(96.0, 43.9)	(95.9, 46.6)	(96.9, 44.3)

Table 3. Comparisons of zero-shot anomaly detection methods on medical datasets. † - AdaCLIP is also trained with auxiliary medical datasets. Other methods are not.

Metric	Dataset	SAA [8] ToC'25	WinCLIP [26] CVPR'23	AnomalyCLIP [84] ICLR'24	AdaCLIP† [9] ECCV'24	AACLIP [47] CVPR'25	Bayes-PFL [53] CVPR'25	FAPrompt [88] ICCV'25	AnomalyVFM
Image-level (AUROC, max-F1)	HeadCT	(46.8, 68.0)	(84.1, 79.8)	(93.0, 88.4)	(91.4, 85.2)	(96.9, 93.1)	(92.6, 86.3)	(93.0, 88.2)	(94.8, 90.5)
	BrainMRI	(34.4, 76.7)	(89.9, 86.9)	(90.0, 86.5)	(94.8, 91.2)	(80.2, 91.5)	(95.2, 94.4)	(95.5, 93.2)	(92.9, 92.5)
	BR35H	(33.2, 67.3)	(81.6, 74.4)	(94.2, 86.8)	(97.7, 92.4)	(95.4, 90.2)	(97.0, 93.2)	(96.6, 90.3)	(94.4, 90.2)
	Average	(38.1, 70.7)	(85.2, 80.4)	(92.4, 87.2)	(94.6, 89.6)	(90.8, 91.6)	(94.9, 91.3)	(95.0, 90.6)	(94.0, 91.1)
Pixel-level (AUROC, max-F1)	ISIC	(83.8, 74.2)	(83.3, 64.1)	(89.4, 71.6)	(89.3, 71.4)	(94.6, 80.4)	(92.3, 76.8)	(90.1, 72.0)	(90.8, 74.4)
	ClinicDB	(66.2, 29.1)	(74.3, 30.7)	(82.9, 42.4)	(84.4, 58.2)	(89.6, 54.1)	(89.6, 51.7)	(83.2, 43.4)	(92.0, 57.6)
	ColonDB	(71.8, 31.5)	(61.2, 19.6)	(81.9, 37.5)	(90.4, 58.2)	(84.1, 38.1)	(82.1, 39.2)	(84.1, 38.8)	(85.6, 42.8)
	Kvasir	(86.2, 65.9)	(38.6, 27.0)	(79.0, 46.2)	(95.0, 77.1)	(87.3, 57.3)	(85.3, 54.8)	(81.6, 48.8)	(90.6, 63.2)
	Endo	(79.4, 51.6)	(43.7, 25.3)	(84.2, 50.3)	(96.6, 80.1)	(90.2, 59.7)	(89.2, 57.9)	(86.4, 52.7)	(92.2, 64.4)
	TN3K	(66.8, 32.6)	(67.2, 30.0)	(81.4, 47.8)	(77.2, 41.9)	(80.5, 43.0)	(85.4, 42.4)	(84.4, 49.2)	(89.0, 55.6)
Average		(75.7, 47.5)	(61.4, 32.8)	(83.1, 49.3)	(88.8, 64.5)	(87.7, 55.4)	(87.3, 53.8)	(85.0, 50.8)	(90.0, 59.7)

we either train a simple decoder and leave the internal representations unchanged or we employ the proposed adaptation strategy. To demonstrate generalisation, we evaluated our model on nine industrial datasets described in Section 5.1. The results can be seen in Table 1. All Chosen VFM achieve a significant improvement in performance in both detection and localisation, showcasing the generality and the strength of the contribution. On average, the image-level AUROC is improved by 6.1 p. p. and the pixel-level AUROC is improved by 10.7 p. p. This is also visually represented in Figure 1, where it can be seen that all of the VFMs match the performance of current methods.

5.4. Comparison to zero-shot methods

Quantitative results In Table 2, the comparison of AnomalyVFM to the state-of-the-art zero-shot anomaly detection methods on industrial datasets is shown. Anoma-

lyVFM outperforms the state-of-the-art considerably in both anomaly detection and anomaly localisation. More specifically, it outperforms the next best method (Bayes-PFL) in terms of image-level AUROC by a substantial 3.3 percentage points (p. p.). AnomalyVFM also slightly improves the results in terms of image-level F1-Max. Additionally, AnomalyVFM significantly improves results on the widely used MVTec AD, VisA, and Real IAD, achieving results close to those of full-shot methods, reiterating the contribution of our model.

In terms of anomaly localisation, AnomalyVFM also improves upon previous methods in terms of pixel-level AUROC and achieves competitive results in terms of pixel-level F1-Max. More specifically, AnomalyVFM improves previous methods by 0.9 p. p. in terms of pixel-level AUROC. In terms of pixel-level F1-Max, it trails behind AdaCLIP [9], which achieves lower scores in terms of pixel-




Figure 5. Qualitative comparison of the anomaly segmentation masks produced by AnomalyVFM and two other best-performing methods. In the first row, the image is shown. In the next three rows, the anomaly segmentations produced by Bayes-PFL [53], AdaCLIP [9] and AnomalyVFM are depicted, and in the last row, the ground truth mask is depicted.

level AUROC.

The results for zero-shot anomaly detection and localisation in the medical domain are presented in Table 3. AnomalyVFM achieves competitive results in terms of detection and improves previous methods in terms of localisation scores. In terms of pixel-level AUROC, AnomalyVFM improves previous methods by 1.2 p. p. More importantly, the results demonstrate the generalisation of AnomalyVFM to the medical domain, despite not being finetuned on any medical data.

Qualitative results Qualitative examples can be seen in Figure 5. AnomalyVFM produces sharper anomaly masks in comparison to Bayes-PFL and is able to localise anomalies even in cases where AdaCLIP fails. Additionally, it is able to detect both small defects (Columns 4 and 5) and larger defects (Columns 2, 3 and 10). Additionally, it successfully detects medical defects (Columns 9 and 12).

5.5. Comparison to few-shot methods

To verify the effectiveness of AnomalyVFM as a backbone, we have fine-tuned the zero-shot model for an additional 50 iterations using a few normal samples. We have chosen MVTEC AD [4] and VisA [89] as our evaluation datasets due to their widespread use. The results can be seen in Table 4. AnomalyVFM achieves the highest image-level AUROC in all settings on MVTEC AD and in the 1-shot setting on VisA.

Remarkably, despite being designed for the zero-shot regime, AnomalyVFM matches or even surpasses the performance of recent few-shot methods such as INP-Former [46], without any architecture changes and with minimal fine-tuning. These results highlight the robustness and

Table 4. Comparison to few-shot methods on MVTEC AD and VisA benchmarks. Results are in image-level AUROC. The best results are marked in bold.

Method	MVTEC AD			VisA		
	1-shot	2-shot	4-shot	1-shot	2-shot	4-shot
PatchCore [59]CVPR’22	83.4	86.3	88.8	79.9	81.6	85.3
WinCLIP+ [26]CVPR’23	93.1	94.4	95.2	83.8	84.6	87.3
PromptAD [38]CVPR’24	94.6	95.7	96.6	86.9	88.3	89.1
INP-Former [46]CVPR’25	96.1	97.0	97.6	91.4	94.6	96.4
AnomalyVFM	97.2	97.9	98.2	93.8	94.2	94.5

transferability of AnomalyVFM, underscoring its potential as a strong and versatile backbone for future anomaly detection research.

6. Ablation study

Ablation experiments validating the individual contributions of AnomalyVFM are performed on 9 industrial datasets presented in the Section 5.1. Results are shown in Table 5. Additional experiments are presented in the Supplementary Material.

Image Generation Model FLUX [36] is used as the default image generation model in our experiments. To verify the importance of this choice, we replaced it with two recent generative models: QWEN-Image [73] and WAN [68]. This leads to a very slight decrease in performance: 0.1 p. p. and 0.4 p. p. in image-level AUROC, respectively and 0.5 p. p. and 2.1 p. p. in pixel-level AUROC, respectively. This shows that our generation pipeline is robust to this choice.

Dataset Filtering To measure the importance of verifying

Table 5. Ablation of the anomaly detection method components.

Group	Condition	Image-level		Pixel-level	
		AUROC	F_1 -Max	AUROC	F_1 -Max
<i>Image Generation</i>	FLUX [36] → QWEN-Image [73]	-0.1	+0.1	-0.5	-0.2
	FLUX [36] → WAN [68]	-0.4	-0.8	-2.1	-1.0
	No Filtering	-3.8	-2.7	-14.6	-18.7
	No Foreground Selection	-1.4	-0.8	-5.8	-12.3
<i>Module Ablation</i>	No Confidence Loss	-0.6	-0.1	-2.0	-6.3
	LoRA [23] → AdaLN [40]	-0.7	-1.0	-2.9	-1.9
	LoRA [23] → VPT [29]	-1.0	+1.2	-0.5	+4.0
<i>AnomalyVFM</i>		94.1	87.6	96.9	44.3

that the generated images actually do contain anomalies, we have omitted the dataset filtering step. This leads to a decrease in image-level AUROC for a significant 3.8 p. p. and a decrease in pixel-level AUROC for 14.6 p. p. This showcases both the problems with current image generation models and the necessity of having clean data.

Anomaly Location Importance In our pipeline, the anomaly region is selected by sampling a rectangle R on the foreground M_{fg} produced by an external model. To verify the importance of this step, we set M_{fg} to be equal to the whole image. This means that sometimes the generated images contain a defect in the background, so the model is trained to focus not only on the object but also on the background. This leads to a decrease of 1.4 p. p. in image-level AUROC and 5.8 p. p. in pixel-level AUROC. This highlights the importance of selecting the inpainting location intelligently.

Confidence Loss To show the importance of the introduced confidence loss, we have retrained the model without it. This has led to a slight decrease in both image-level and pixel-level AUROC (0.6 p. p. and 2.0 p. p. respectively). This reiterates the importance of this loss for optimal performance.

Adapter Architecture To even further show the generality of our framework, LoRA [23] adapters were exchanged with two other Parameter Efficient Techniques, AdaLN [40] and VPT [29]. This has led to a decrease in performance of 0.7 p. p. and 1.0 p. p. in image-level AUROC and 2.9 p. p. and 0.5 p. p. in pixel-level AUROC, respectively. All of these results are still significantly above SOTA and show the robustness of our framework to this choice. Additionally, it shows possible extensions to newer VFMs, which might have different architectures.

Inference Speed and Computational Complexity The inference speed can be seen in Table 6. The protocol from EfficientAD [2] was used to calculate them. AnomalyVFM is significantly faster than its main competitors. AnomalyVFM requires approximately 2 hours to train on a single A100 GPU and has 345.8 million parameters. Out of these 35.4 million are trainable.

Limitations At present, the main bottleneck lies in the image generation stage, which takes approximately one day on

Table 6. Results for average inference time of a single sample with NVIDIA A100 GPU. Inference times are reported in milliseconds.

Method	Bayes-PFL [53]	AdaCLIP [9]	<i>AnomalyVFM</i>
Inference [ms]	208.5	82.4	20.5

an A100 GPU, whereas model training requires only about two hours. Additional discussion of these limitations and related analyses can be found in the Supplementary.

7. Conclusion

We present AnomalyVFM, a practical and model-agnostic framework that transforms any pretrained Vision Foundation Model into a strong zero-shot anomaly detector. Unlike prior approaches that rely on high-level concept knowledge from vision-language models, AnomalyVFM leverages the rich visual representations of VFMs and enhances them through two key innovations. First, we introduced a three-stage synthetic dataset generator that produces diverse and realistic training samples, capturing a broad range of object categories and defect types. Second, we designed a parameter-efficient adaptation strategy that inserts low-rank adapters throughout the backbone and employs a confidence-weighted loss to refine the model’s representations with minimal parameters and robust supervision.

Together, these components allow VFMs to generalise to unseen object classes and outperform existing VLM-based methods in the zero-shot regime. More specifically, we achieve an average image-level AUROC of 94.1% across 9 diverse industrial datasets, improving upon previous methods by a significant 3.3. percentage points. Additionally, we demonstrate the effectiveness of AnomalyVFM as a potent backbone by finetuning it on a few normal samples without any bells and whistles. With this, AnomalyVFM achieves a performance comparable to SOTA in the few-shot regime.

Looking ahead, further efforts to improve defect realism and resulting labels are a good avenue for future research. Additionally, integrating depth data via monodepth foundational models, such as Marigold [32], could be used to enable zero-shot RGBD anomaly detection. Most importantly, the results also indicate that AnomalyVFM could be used as a backbone for future few-shot and full-shot models.

Acknowledgements This work was in part supported by the ARIS research projects MUXAD (J2-60055) and AI4Science (GC-0001), research programme P2-0214 and the supercomputing network SLING (ARNES, EuroHPC Vega).

References

- [1] Toshimichi Aota, Lloyd Teh Tzer Tong, and Takayuki Okatani. Zero-shot versus many-shot: Unsupervised texture anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5564–5572, 2023. [2](#), [5](#)
- [2] Kilian Batzner, Lars Heckler, and Rebecca König. EfficientAD: Accurate Visual Anomaly Detection at Millisecond-Level Latencies. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 128–138, 2024. [2](#), [8](#)
- [3] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving Unsupervised defect segmentation by applying structural similarity to autoencoders. *ArXiv*, abs/1807.02011, 2018. [2](#)
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTec AD—A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9592–9600, 2019. [1](#), [2](#), [5](#), [7](#)
- [5] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño. Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized medical imaging and graphics*, 43:99–111, 2015. [5](#)
- [6] Vadim Borisov, Kathrin Seßler, Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. Language models are realistic tabular data generators. *arXiv preprint arXiv:2210.06280*, 2022. [2](#)
- [7] Jakob Božič, Domen Tabernik, and Danijel Skočaj. Mixed supervision for surface-defect detection: From weakly to fully supervised learning. *Computers in Industry*, 129: 103459, 2021. [5](#)
- [8] Yunkang Cao, Xiaohao Xu, Chen Sun, Yuqi Cheng, Zongwei Du, Liang Gao, and Weiming Shen. Segment Any Anomaly without Training via Hybrid Prompt Regularization. *arXiv preprint arXiv:2305.10724*, 2023. [2](#), [3](#), [6](#)
- [9] Yunkang Cao, Jiangning Zhang, Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi. AdaCLIP: Adapting CLIP with hybrid learnable prompts for zero-shot anomaly detection. In *European Conference on Computer Vision*, pages 55–72. Springer, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [10] Weirong Chen, Ganlin Zhang, Felix Wimbauer, Rui Wang, Nikita Araslanov, Andrea Vedaldi, and Daniel Cremers. Back on track: Bundle adjustment for dynamic scene reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4951–4960, 2025. [2](#)
- [11] Xuhai Chen, Jiangning Zhang, Guanzhong Tian, Haoyang He, Wuhao Zhang, Yabiao Wang, Chengjie Wang, and Yong Liu. Clip-AD: A Language-Guided Staged Dual-Path Model for Zero-shot Anomaly Detection. In *International Joint Conference on Artificial Intelligence*, pages 17–33. Springer, 2024. [3](#)
- [12] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, pages 168–172. IEEE, 2018. [5](#)
- [13] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International conference on pattern recognition*, pages 475–489. Springer, 2021. [2](#)
- [14] Anja Delić, Matej Grcic, and Siniša Šegvić. Outlier detection by ensembling uncertainty with negative objectness. In *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25-28, 2024*. BMVA, 2024. [1](#)
- [15] Hanqiu Deng and Xingyu Li. Anomaly Detection via Reverse Distillation from One-Class Embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9737–9746, 2022. [2](#), [5](#)
- [16] Yuxuan Duan, Yan Hong, Li Niu, and Liqing Zhang. Few-shot defect image generation via defect-aware feature manipulation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(1):571–578, 2023. [2](#)
- [17] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In *Forty-first International Conference on Machine Learning*, 2024. [2](#), [3](#)
- [18] Matic Fučka, Vitjan Zavrtanik, and Danijel Skočaj. TransFusion—a Transparency-based Diffusion Model for Anomaly Detection. In *European conference on computer vision*, pages 91–108. Springer, 2025. [1](#), [2](#)
- [19] Matic Fučka, Vitjan Zavrtanik, and Danijel Skočaj. SALAD – Semantics-Aware Logical Anomaly Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. [2](#)
- [20] Haifan Gong, Guanqi Chen, Ranran Wang, Xiang Xie, Mingzhi Mao, Yizhou Yu, Fei Chen, and Guanbin Li. Multi-task learning for thyroid nodule segmentation with thyroid region prior. In *2021 IEEE 18th international symposium on biomedical imaging (ISBI)*, pages 257–261. IEEE, 2021. [5](#)
- [21] A. Hamada. Br35h: Brain tumor detection. <https://www.kaggle.com/datasets/ahmedhamada0/brain-tumor-detection>, 2020. Online; accessed 2020. [1](#), [5](#)
- [22] Steven A Hicks, Debesh Jha, Vajira Thambawita, Pål Halvorsen, Hugo L Hammer, and Michael A Riegler. The

- endotect 2020 challenge: evaluation and comparison of classification, segmentation and inference time for endoscopy. In *International Conference on Pattern Recognition*, pages 263–274. Springer, 2021. 5
- [23] Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 4, 8
- [24] Teng Hu, Jiangning Zhang, Ran Yi, Yuzhen Du, Xu Chen, Liang Liu, Yabiao Wang, and Chengjie Wang. Anomalydiffusion: Few-shot anomaly image generation with diffusion model. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(8):8526–8534, 2024. 2
- [25] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. GPT-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 3
- [26] Jongheon Jeong, Yang Zou, Taewan Kim, Dongqing Zhang, Avinash Ravichandran, and Onkar Dabeer. WinCLIP: Zero-/Few-Shot Anomaly Classification and Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19606–19616, 2023. 3, 6, 7
- [27] Stepan Jezek, Martin Jonak, Radim Burget, Pavel Dvorak, and Milos Skotak. Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions. In *2021 13th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, pages 66–71, 2021. 5
- [28] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas De Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *International conference on multimedia modeling*, pages 451–462. Springer, 2019. 1, 5
- [29] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727. Springer, 2022. 8
- [30] Pranita Balaji Kanade and PP Gumaste. Brain tumor detection using mri images. *Brain*, 3(2):146–150, 2015. 5
- [31] Laurynas Karazija, Iro Laina, Andrea Vedaldi, and Christian Rupprecht. Diffusion Models for Open-Vocabulary Segmentation. In *European Conference on Computer Vision*, pages 299–317. Springer, 2024. 2
- [32] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9492–9502, 2024. 8
- [33] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018. 4
- [34] 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*, pages 4015–4026, 2023. 3
- [35] Orest Kupyn and Christian Rupprecht. Dataset Enhancement with Instance-Level Augmentations. In *European Conference on Computer Vision*, pages 384–402. Springer, 2024. 2
- [36] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2, 3, 5, 7, 8, 1
- [37] Aodong Li, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan Mandt. Zero-Shot Anomaly Detection via Batch Normalization. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [38] Xiaofan Li, Zhizhong Zhang, Xin Tan, Chengwei Chen, Yanyun Qu, Yuan Xie, and Lizhuang Ma. PromptAD: Learning Prompts with only Normal Samples for Few-Shot Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16838–16848, 2024. 7
- [39] Xiaofan Li, Zhizhong Zhang, Xin Tan, Chengwei Chen, Yanyun Qu, Yuan Xie, and Lizhuang Ma. PromptAD: Learning Prompts with only Normal Samples for Few-Shot Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16838–16848, 2024. 1
- [40] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & Shifting Your Features: A New Baseline for Efficient Model Tuning. *Advances in Neural Information Processing Systems*, 35:109–123, 2022. 8
- [41] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 4
- [42] Jiahui Liu, Xin Wen, Shizhen Zhao, Yingxian Chen, and Xiaojuan Qi. Can OOD Object Detectors Learn from Foundation Models? In *European Conference on Computer Vision*, pages 213–231. Springer, 2024. 2
- [43] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding DINO: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer, 2024. 3
- [44] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. SimpleNet: A Simple Network for Image Anomaly Detection and Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20402–20411, 2023. 2
- [45] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022. 3
- [46] Wei Luo, Yunkang Cao, Haiming Yao, Xiaotian Zhang, Jianan Lou, Yuqi Cheng, Weiming Shen, and Wenyong Yu. Exploring intrinsic normal prototypes within a single image for universal anomaly detection. In *Proceedings of the*

- Computer Vision and Pattern Recognition Conference*, pages 9974–9983, 2025. 7
- [47] Wenxin Ma, Xu Zhang, Qingsong Yao, Fenghe Tang, Chenxu Wu, Yingtai Li, Rui Yan, Zihang Jiang, and S Kevin Zhou. Aa-clip: Enhancing zero-shot anomaly detection via anomaly-aware clip. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 4744–4754, 2025. 6, 1
- [48] Pankaj Mishra, Riccardo Verk, Daniele Fornasier, Claudio Piciarelli, and Gian Luca Foresti. VT-ADL: A vision transformer network for image anomaly detection and localization. In *30th IEEE/IES International Symposium on Industrial Electronics (ISIE)*, 2021. 5
- [49] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*, 2024. 1, 2, 3, 5
- [50] Jonathan Pirnay and Keng Chai. Inpainting transformer for anomaly detection. In *International Conference on Image Analysis and Processing*, pages 394–406. Springer, 2022. 2
- [51] Anamaria-Roberta Preda, Christoph Mayr-Dorn, Atif Mashkoor, and Alexander Egyed. Supporting high-level to low-level requirements coverage reviewing with large language models. In *Proceedings of the 21st International Conference on Mining Software Repositories*, pages 242–253, 2024. 2
- [52] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly Accurate Dichotomous Image Segmentation. In *European Conference on Computer Vision*, pages 38–56. Springer, 2022. 3
- [53] Zhen Qu, Xian Tao, Xinyi Gong, ShiChen Qu, Qiyu Chen, Zhengtao Zhang, Xingang Wang, and Guiguang Ding. Bayesian Prompt Flow Learning for Zero-Shot Anomaly Detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 30398–30408, 2025. 1, 5, 6, 7, 8
- [54] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 3
- [55] Mike Ranzinger, Greg Heinrich, Jan Kautz, and Pavlo Molchanov. AM-RADIO: Agglomerative vision foundation model reduce all domains into one. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12490–12500, 2024. 5
- [56] Blaž Rolih, Matic Fučka, and Danijel Skočaj. SuperSimpleNet: Unifying Unsupervised and Supervised Learning for Fast and Reliable Surface Defect Detection. In *International Conference on Pattern Recognition*, 2024. 2
- [57] Blaž Rolih, Matic Fučka, and Danijel Skočaj. No Label Left Behind: A Unified Surface Defect Detection model for all Supervision Regimes. *Journal of Intelligent Manufacturing*, 2025. 2
- [58] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [59] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards Total Recall in Industrial Anomaly Detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2022. 1, 2, 7
- [60] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14902–14912, 2021. 1, 5
- [61] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 5
- [62] Domen Tabernik, Samo Šela, Jure Skvarč, and Danijel Skočaj. Segmentation-Based Deep-Learning Approach for Surface-Defect Detection. *Journal of Intelligent Manufacturing*, 2019. 5
- [63] Nima Tajbakhsh, Suryakanth R Gurudu, and Jianming Liang. Automated polyp detection in colonoscopy videos using shape and context information. *IEEE transactions on medical imaging*, 35(2):630–644, 2015. 5
- [64] Fenfang Tao, Guo-Sen Xie, Fang Zhao, and Xiangbo Shu. Kernel-aware graph prompt learning for few-shot anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7347–7355, 2025. 1
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [66] Tomáš Vojtř and Jiří Matas. Image-consistent detection of road anomalies as unpredictable patches. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5491–5500, 2023. 1
- [67] Tomáš Vojtř, Jan Šochman, and Jiří Matas. Pixood: Pixel-level out-of-distribution detection. In *European Conference on Computer Vision*, pages 93–109. Springer, 2024. 1
- [68] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 7, 8, 1
- [69] Chengjie Wang, Wenbing Zhu, Bin-Bin Gao, Zhenye Gan, Jiangning Zhang, Zhihao Gu, Shuguang Qian, Mingang Chen, and Lizhuang Ma. Real-IAD: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Com-*

- puter Vision and Pattern Recognition, pages 22883–22892, 2024. [1](#), [5](#)
- [70] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. DUST3R: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. [4](#)
- [71] Chenxi Whitehouse, Monojit Choudhury, and Alham Fikri Aji. LLM-powered data augmentation for enhanced cross-lingual performance. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. [2](#)
- [72] Matthias Wieler and Tobias Hahn. Weakly supervised learning for industrial optical inspection. In *DAGM symposium in*, page 11, 2007. [5](#)
- [73] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. [7](#), [8](#), [1](#)
- [74] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. [4](#)
- [75] Minghui Yang, Peng Wu, and Hui Feng. Memseg: A semi-supervised method for image surface defect detection using differences and commonalities. *Engineering Applications of Artificial Intelligence*, 119:105835, 2023. [4](#)
- [76] Shuai Yang, Zhifei Chen, Pengguang Chen, Xi Fang, Yixun Liang, Shu Liu, and Yingcong Chen. Defect spectrum: A granular look of large-scale defect datasets with rich semantics. In *Computer Vision – ECCV 2024*, pages 187–203, Cham, 2024. Springer Nature Switzerland. [2](#)
- [77] Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. GPT3Mix: Leveraging large-scale language models for text augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. [2](#)
- [78] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. DRÆM - a Discriminatively Trained Reconstruction Embedding for Surface Anomaly Detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8330–8339, 2021. [2](#)
- [79] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021. [2](#)
- [80] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. DSR—a dual subspace re-projection network for surface anomaly detection. In *European conference on computer vision*, pages 539–554. Springer, 2022. [2](#)
- [81] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Cheating depth: Enhancing 3d surface anomaly detection via depth simulation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2164–2172, 2024. [2](#)
- [82] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-gan: High-fidelity defect synthesis for automated defect inspection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2524–2534, 2021. [2](#)
- [83] Qi Zhao, Xingyu Ni, Ziyu Wang, Feng Cheng, Ziyang Yang, Lu Jiang, and Bohan Wang. Synthetic video enhances physical fidelity in video synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12135–12146, 2025. [2](#)
- [84] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. AnomalyCLIP: Object-agnostic Prompt Learning for Zero-shot Anomaly Detection. In *ICLR*, 2024. [1](#), [3](#), [5](#), [6](#)
- [85] Yixuan Zhou, Xing Xu, Jingkuan Song, Fumin Shen, and Heng Tao Shen. Msflow: Multiscale flow-based framework for unsupervised anomaly detection. *IEEE transactions on neural networks and learning systems*, 2024. [2](#)
- [86] Jiawen Zhu and Guansong Pang. Toward generalist anomaly detection via in-context residual learning with few-shot sample prompts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17826–17836, 2024. [1](#)
- [87] Jiaqi Zhu, Shaofeng Cai, Fang Deng, Beng Chin Ooi, and Junran Wu. Do LLMs Understand Visual Anomalies? Uncovering LLM’s Capabilities in Zero-shot Anomaly Detection. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 48–57, 2024. [3](#)
- [88] Jiawen Zhu, Yew-Soon Ong, Chunhua Shen, and Guansong Pang. Fine-grained abnormality prompt learning for zero-shot anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22241–22251, 2025. [6](#), [1](#)
- [89] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. SPot-the-Difference Self-supervised Pre-training for Anomaly Detection and Segmentation. In *European Conference on Computer Vision*, pages 392–408. Springer, 2022. [1](#), [2](#), [5](#), [7](#)

AnomalyVFM – Transforming Vision Foundation Models into Zero-Shot Anomaly Detectors

Supplementary Material

In this Appendix, we provide extensive additional details and supporting information that extend beyond the scope of the main manuscript. The Appendix is organised as follows:

- **Limitations** in Section A.
- **Discussion about the Dataset Generation Phase** in Section B.
- **Results of competing methods when trained on the synthetic dataset and a discussion about them** in Section C.
- **Extended synthetic dataset details** in Section D.
- **Extended ablation studies** in Section E.
- **Additional qualitative results** in Section F.
- **Data generation data** in Section G.

A. Limitations

The main limitation currently is the time required to generate the synthetic dataset, which takes approximately one day on an A100 GPU, whereas model training requires only about two hours. While this represents a lot of time, it is a one-time investment, and the same dataset can be used for every VFM. With the improvements to the generation speed of current image generation models, we expect this time to drop even further. In Section E, we also conducted additional experiments, demonstrating that good performance can be achieved with fewer than 10,000 images, meaning the generation phase can be shorter if needed.

Additionally, while AnomalyVFM performs well on medical datasets, its performance could be further improved. In our preliminary attempts, the pretrained image generation models [36, 68, 73] failed to output realistic medical images suitable for zero-shot anomaly detection training. While this was not needed for industrial anomaly detection, fine-tuning the image generator on an auxiliary medical imaging dataset may enable the image-generation model to output data of suitable quality.

B. Discussion about dataset generation phase

While our synthetic dataset generation works well, it could be further improved. More specifically, the anomaly mask estimation and image filtering could be further improved. Although the dataset filtering is quite robust, some images without anomalies still pass through. Some examples of this can be seen in Figure 1. A trained AnomalyVFM could be used to further filter the data and improve the data quality even further. On top of that, the amount and the content of [Object] tags could be improved. Based on the experi-




Figure 1. Failure Cases in Image Generation Process

Table 1. Comparison of performance of competing methods when trained on the proposed synthetic dataset versus when using the default datasets. SD stands for Synthetic Dataset

Method	SD	Image-level		Pixel-level	
		AUROC	F1-Max	AUROC	F1-Max
AACLIP [47]		86.1	81.4	95.0	45.1
	✓	85.6 ↓ 0.5	81.1 ↓ 0.3	93.5 ↓ 1.5	44.2 ↓ 0.9
AnomalyCLIP [84]		87.6	83.6	95.1	46.0
	✓	87.5 ↓ 0.1	82.5 ↓ 1.1	95.3 ↑ 0.3	45.8 ↓ 0.2
FAPrompt [88]		88.1	84.1	95.9	46.6
	✓	88.5 ↑ 0.4	84.4 ↑ 0.3	96.2 ↑ 0.3	45.9 ↓ 0.7
AdaCLIP [9]		89.6	87.5	95.0	51.0
	✓	87.1 ↓ 2.5	84.9 ↓ 2.6	92.9 ↓ 2.1	47.9 ↓ 3.1
Bayes-PFL [53]		90.8	85.1	96.0	43.9
	✓	91.2 ↑ 0.4	85.6 ↑ 0.5	96.1 ↑ 0.1	43.8 ↓ 0.1

ment in Section E, we hypothesise that this would improve the performance even further. We have, however, left this for future work.

To ensure that **no data leakage** occurred during the generation phase, we manually reviewed the [Object] tags and excluded any tags that were included in the evaluation test sets. We have left [Anomaly] and [Texture] as they were generated, as these represent more general concepts.

Table 2. Dataset Statistics for the generated dataset

Dataset Statistic	Value
No. of images	10,000
No. of different objects	100
No. of different backgrounds	50
No. of different anomalies	204
No. of object background combinations	4,596
Avg. Anomalous Area	2.52%
Min. Anomalous Area	0.28%
Max. Anomalous Area	11.24%

C. Training Competing methods with the proposed synthetic dataset

To demonstrate that the diversity of the datasets is not problematic for VLM-based methods, we retrained them using the proposed synthetic dataset. The results can be seen in Table 1. While it does help for some methods, it does not significantly alter the results. This indicates that VLM methods do not suffer from the same problem of inadequate data diversity as VFMs.

D. Synthetic Dataset Details

Here, we provide details about the synthetic dataset generated for training our model. High-level statistics can be seen in Table 2. The generated dataset contains all of the possible objects and backgrounds. Additionally, it contains 204 different anomalies, significantly more than current datasets (e.g. MVTec AD [4] contains 73 different anomaly types). The generated anomalies are, in general, relatively small (on average, they account for 2.52% of the image). In contrast, in MVTec AD, they occupy 4.39% of the image. A more detailed visualisation of the anomaly area distribution is depicted in Figure 2.

E. Additional Ablation Studies

In this section, we present additional experiments that verify the design choices in AnomalyVFM. Most of the results are presented in Table 3

Filtering Threshold To verify the impact of the threshold T used during dataset filtering, we re-filtered the dataset using various values of T . On top of the performance metrics, we also measured the rejection rate (i.e., the percentage of images discarded). The results and the rejection rate can be seen in Figure 3. The results show that the image-level AUROC is quite robust to the set threshold, while the pixel-level AUROC is more reliant on a correct choice of a threshold. At the default setting, the rejection rate is approximately 30%, showcasing that prompt adherence is far from a solved problem in generative models.

Number of [Object] tags To verify the importance




Figure 2. Anomalous Area Distribution in the generated synthetic dataset.

of having a diverse dataset, we varied the number of [Object] tags during the synthetic data generation phase. The results can be seen in Figure 4. The results consistently rise with the number of [Object] tags. The performance with 20 [Object] tags is similar to the performance when AnomalyVFM is trained on MVTec AD [4], which has 15 different objects inside the dataset. We have not gone above 100 tags, as that is the list we initially generated with an LLM. In the future, we will increase this to see if the performance can be improved even further.


Number of images During all of our experiments, we used 10,000 generated images. To verify the importance of this, we have tried several different quantities: 100, 500, 1,000, and 10,000. The results are depicted in Figure 5. The performance increases steadily with each increment. We hypothesise that further scaling could improve performance even further. We have not done so to maintain a training set size similar to that of the related methods.

LoRA Rank To verify the robustness of the proposed method towards the rank of the LoRA adapters, we varied this parameter. More specifically, we decreased the rank to 32 and then increased it to 128. Decreasing it leads to a decrease of 0.3 p. p. in image-level AUROC and 0.3 p. p. in pixel-level AUROC. Increasing the LoRA rank leads to no differences in image-level metrics, while the pixel-level AUROC decreases for 0.3 p. p. This shows the robustness of the proposed method to this parameter.

LoRA Positions In the implementation, LoRA is added to query, value and projection layers inside the attention mechanism. This was done based on the insights from the open-

Table 3. Additional ablations of the anomaly detection method components.

Group	Condition	Image-level		Pixel-level	
		AUROC	F_1 -Max	AUROC	F_1 -Max
Module Ablation	ViT-L $\rightarrow$ ViT-B	-1.8	-1.9	-1.2	-2.5
	ViT-L $\rightarrow$ ViT-H	-0.6	-0.4	-0.8	-2.3
	LoRA Rank 64 $\rightarrow$ 32	-0.3	0.0	-0.3	+0.4
	LoRA Rank 64 $\rightarrow$ 128	0.0	+0.1	-0.3	-0.2
	LoRA Positions: QKV and Proj	-0.1	+0.2	-0.2	-0.2
	LoRA Positions: All Norm Layers	-0.1	-0.4	-0.4	-1.1
	LoRA Positions: All Linear Layers	-0.3	-1.3	-0.1	-0.5
AnomalyVFM	LoRA Positions: QV and Proj	94.1	87.6	96.9	44.3

Figure 3. Model performance and rejection rate in relation to filtering threshold T .

source community on how to efficiently adapt image generation models. To verify the importance of this choice, we performed experiments with more layouts. All of the layouts keep a similar performance, showcasing robustness to this choice. The largest dip in performance is observed when LoRA adaptors are added to all linear layers. We assume this is the case as the model cannot pass the information globally but rather only locally.

Model Size RADIO has multiple model sizes. To verify the importance of this parameter, we exchanged it with a smaller (ViT-B) and larger (ViT-H) model. Using a smaller model leads to a decrease of 1.8 p. p. in image-level AUROC and 1.2 p. p. in pixel-level AUROC. A larger model leads to a decrease of 0.6 p. p. in image-level AUROC and 1.2 p. p. in pixel-level AUROC. This shows that ViT-L is the optimal choice. We also hypothesise that increasing the number of [Object] tags and the total number of images would make ViT-H more optimal.

F. Additional Qualitative Examples

In this section, we add additional qualitative examples of anomaly segmentations produced by AnomalyVFM. The examples can be seen in Figure 6. AnomalyVFM can detect anomalies across a wide range of objects.

G. Image Generation Data

To enable reproducibility and to ensure transparency, we provide the list of [Object], [Anomaly] and [Texture] used in the synthetic dataset generation. The lists of [Object] and [Anomaly] tags can be seen in Table 4 and Table 5. The list of [Texture] tags can be seen in Table 6.




Figure 4. Model performance in comparison to the number of [Object] tags.




Figure 5. Model performance in comparison to the number of images in the training set.




Figure 6. Qualitative examples of anomaly segmentation masks produced by AnomalyVFM. In the first row, the image is shown. In the next row, the anomaly segmentation produced by AnomalyVFM is depicted, and in the last row, the ground truth mask is depicted.

Table 4. [Object] and [Anomaly] data used for synthetic data generation. Here are listed objects from A to L.

[Object]	[Anomaly]
apple	[bruised, wrinkled, rotten, moldy, dented, discolored, soft spots]
apple slice	[oxidized, bruised, dried out]
asphalt	[cracked, pitted, faded, eroded, oil-stained, uneven]
ball	[deflated, scuffed, punctured, faded, cracked surface]
banana	[bruised, overripe, blackened, split peel, mushy, spotted]
battery	[leaking, corroded, dented, faded label]
belt	[cracked leather, frayed edges, worn holes, peeling finish]
bicycle	[flat tire, rusty chain, scratched frame, worn seat]
board game	[torn box, missing pieces, faded board, bent cards]
bread	[stale, moldy, crumbling, burnt, hardened, soggy]
brushed aluminum	[scratched, dented, stained, faded, oxidized, pitted]
butter	[rancid, melted, discolored, greasy residue, hardened]
car tire	[bald tread, cracked rubber, punctured, worn sidewall]
carbon fiber	[frayed, chipped, cracked, delaminated, scratched, discolored]
carrot	[softened, cracked, dehydrated, spotted, moldy, bent]
chair	[scratched wood, stained cushion, wobbly leg]
chalkboard	[scratched, smudged, cracked, chipped, stained, uneven]
cheese	[moldy, dried out, cracked, discolored, sweating, crumbly]
chocolate bar	[melted, bloomed, crumbled, discolored]
concrete	[cracked, pitted, stained, eroded, chipped, weathered]
cookies	[crumbled, stale, burnt, moldy]
cork	[cracked, crumbled, stained, dried out, warped, pitted]
corrugated metal	[dented, rusted, bent, scratched, corroded, pitted]
denim	[frayed, torn, stained, faded, pilled, worn]
doll	[torn clothing, missing eye, stained, frayed hair, loose limbs]
drill	[worn chuck, scratched casing, broken switch, dented battery]
egg	[cracked, leaking, discolored shell, dented, rotten, thin shell]
fabric	[frayed, torn, stained, faded, pilled, snagged]
fur	[matted, shedding, stained, torn, faded, dull]
garden hose	[cracked, leaking, kinked, faded]
garlic	[sprouted, dried out, moldy]
glasses	[scratched lenses, bent frame, loose arms, cloudy lenses]
gloves	[frayed fingers, stretched out, stained]
grape	[wrinkled, moldy, shriveled]
grill	[rusty grates, blackened residue, scratched body]
hammer	[rusty, chipped, bent, dented, scratched, loose head]
hat	[faded color, stretched, frayed edges]
headphones	[frayed cable, scratched ear cups, loose padding]
helmet	[scratched shell, cracked foam, loose straps]
hemp fabric	[frayed, torn, faded, pilled, stained, snagged]
jacket	[broken zipper, faded color, torn lining]
jeans	[worn knees, frayed hem, ripped pocket, faded]
key	[bent, worn teeth, rusty, scratched surface]
kite	[torn fabric, bent frame, frayed string, missing tail]
laminate	[scratched, chipped, peeled, bubbled, stained, warped]
lamp	[flickering, scratched base, broken switch]
laptop	[scratched casing, cracked hinge, faded keyboard keys]
lettuce	[wilting, yellowing, rotting]
light bulb	[burnt out, cracked, blackened, loose filament]
linen	[wrinkled, stained, faded, torn, frayed, pilled]

Table 5. [Object] and [Anomaly] data used for synthetic data generation. Here are listed objects from M to Z.

[Object]	[Anomaly]
mesh	[frayed, torn, snagged, discolored, brittle, stretched]
milk carton	[dented, leaking, stained, faded label, torn packaging]
mirror	[scratched, chipped edge, cloudy, stained surface]
onion	[sprouted, dried layers, rotting]
orange	[dried skin, moldy, bruised, discolored]
paintbrush	[frayed bristles, stiffened bristles, dried paint]
paper	[torn, wrinkled, stained, yellowed, brittle, moldy]
parquet flooring	[scratched, warped, faded, chipped, stained, dull]
phone	[cracked screen, scratched back, worn buttons]
plastic	[scratched, cracked, discolored, warped, brittle, faded]
pliers	[rusty, loose grip, scratched, chipped, stiff joint]
plywood	[warped, splintered, chipped, stained, delaminated, cracked]
potato	[sprouted, rotting, green spots, wrinkled, moldy, soft spots]
rake	[bent tines, rusty, loose handle]
rattan	[splintered, frayed, cracked, stained, brittle, discolored]
rubber floor	[cracked, brittle, discolored, stiffened, melted, torn]
saw	[rusty blade, dull teeth, chipped handle, bent blade]
scarf	[pilled fabric, snagged threads, stained]
screwdriver	[worn tip, rusty, scratched, bent, cracked handle]
shoes	[worn sole, scuffed leather, torn fabric, faded color]
shovel	[rusty blade, dented handle, worn grip]
smooth ceramic tile	[chipped, cracked, stained, crazed, dull, scratched]
smooth glass	[scratched, cracked, chipped, foggy, stained, shattered]
smooth metal	[rusted, scratched, dented, corroded, pitted, tarnished]
smooth wood plank	[cracked, splintered, warped, knotted, rotten, scratched, stained]
socks	[hole in toe, stretched elastic, faded]
stainless steel	[scratched, dented, stained, scuffed, fingerprinted, corroded]
stone tile	[chipped, cracked, eroded, stained, pitted, weathered]
strawberry	[moldy, bruised, shrinking]
synthetic fiber	[frayed, torn, stained, faded, pilled, melted]
table	[scratched surface, dented corner, water stains]
tape measure	[cracked casing, faded markings, stuck mechanism]
teddy bear	[ripped seam, matted fur, faded color, stained, missing stuffing]
tent	[torn fabric, bent poles, moldy spots]
tomato	[soft spots, cracked skin, moldy]
toy car	[scratched paint, missing wheel, cracked body, loose parts]
TV remote	[worn-out buttons, cracked case, faded labels]
velvet	[crushed, faded, stained, pilled, torn, frayed]
wallet	[worn edges, cracked leather, faded color, frayed stitching]
wallpaper	[peeled, torn, stained, faded, bubbled, wrinkled]
watch	[scratched face, broken strap, faded markings, cracked casing]
whiteboard	[scratched, stained, ghosting, cracked, faded, dented]
window	[scratched glass, cracked, foggy]
woven mat	[frayed, torn, faded, loose fibers, stained, worn]
wrench	[rusty, scratched, dented, worn edges, corroded]
yo-yo	[scratched, cracked, tangled string, chipped edge]

Table 6. [Texture] data used for synthetic dataset generation.

[Texture]
asphalt, bamboo, brick, brushed aluminum, canvas carbon fiber, ceramic, chalkboard, clouds, concrete cork, corrugated metal, denim, fabric, fleece foam, fur, glass, granite, grass gravel, hemp fabric, ice, laminate, linen marble, mesh, metal, mirror, painted wall paper, parquet flooring, pebbles, plastic, plywood rattan, rubber, sand, snow, stainless steel stone, synthetic fiber, tarpaulin, terrazzo, tile velvet, wallpaper, whiteboard, wire mesh, woven mat