Generative Defect Synthesis for Enhancing Industrial Anomaly Detection

Avinash Kumar Sharma and Tushar Shinde
Indian Institute of Technology Madras Zanzibar
shinde@iitmz.ac.in

Abstract

Anomaly detection in industrial manufacturing is challenged by the scarcity of defective samples and class imbalance. To address this, we propose a generative defect synthesis framework that leverages generative models to create realistic and diverse anomalies for training visual recognition systems. Our method preserves structural and textural consistency with normal samples while generating high-fidelity defects, enabling improved anomaly detection in low-data regimes. We validate our approach on the MVTec AD benchmark and demonstrate significant improvements in both image-level and pixel-level AUC scores. Furthermore, we explore model compression through 8-bit quantization, achieving up to 4× model size reduction with minimal performance degradation. A detailed case study in multiple category of objects and textures highlights the efficacy of synthetic defects and the optimal combination of real and synthetic data. This work bridges generative image modeling and industrial visual recognition, offering a scalable and efficient solution for real-world anomaly detection with reduced data annotation effort and enhanced deployability on edge devices.

1. Introduction and Related Work

Automated visual inspection (AVI) plays a vital role in industrial quality control, where accurate detection of surface anomalies is essential [5, 10]. While deep learning has enabled substantial progress in this area, its reliance on large-scale annotated datasets remains a significant limitation. In industrial domains, defective samples are rare, diverse, and expensive to label, prompting interest in generative methods to synthesize realistic anomalies.

Recent self-supervised and unsupervised methods [4, 12, 15] have shown promise in detecting anomalies without explicit supervision [3, 7, 14]. However, their generalization is often constrained by the scarcity and homogeneity of available data. Synthetic data generation via generative adversarial networks (GANs) [1], physics-based simulation [11], and advanced augmentations [6], provides a way to enrich training sets with diverse defect types. Still, the challenge remains: how can synthetic defects best emulate real-world conditions to boost detection performance?

The MVTec AD [3] benchmark provides a standard for evaluating industrial anomaly detection, with real-world variations across textures and objects. Feature-based models like PaDiM [7] and PatchCore [14] detect anomalies by modeling normal patch-level distributions. Recent work such as DRÆM [19] highlights the potential of synthetic anomalies in improving generalization. Moreover, practical deployment demands efficiency. Model compression techniques such as quantization and pruning [9, 18] are essential for resource-constrained settings like edge devices, ensuring low-latency inference without significant performance degradation.

In this paper, we propose a generative defect synthesis framework that: 1) Synthesizes structurally consistent and photorealistic anomalies to augment training data for deep anomaly detection. 2) Investigates the trade-off between real and synthetic data and the effect of realism on generalization. 3) Incorporates model quantization to enable lightweight deployment without compromising accuracy. Extensive experiments on the MVTec AD and synthetic datasets show that our method improves detection performance using fewer real defects, and scales effectively across object and texture categories.

The rest of the paper is organized as follows: Section 2 analyzes the MVTec Synthetic dataset and highlights core challenges. Section 3 details our proposed method and compression strategies. Section 4 presents the experimental setup and evaluation results. Finally, Section 5 concludes with key findings and future directions.

2. Dataset Characteristics and Analysis

The MVTec AD dataset [3] is a standard benchmark for unsupervised industrial anomaly detection, comprising 15 categories (10 objects and 5 textures) with 3,629 normal training images and 1,258 defective test samples across 48 defect types. Images reflect real-world manufacturing defects such as scratches, dents, and misalignments, and vary in resolution and structure across categories.

Table 1. Summary of MVTec AD Dataset and Synthetic Dataset Categories with Updated Defective Samples

Category	Type	# Defect Types	MVTec AD Dataset # Normal Samples	MVTec AD Dataset # Defective Samples	MVTec AD Dataset Imbalance Ratio	Synthetic MVTec AD Dataset # Normal Samples	Synthetic MVTec AD Dataset # Defective Samples	Synthetic MVTec AD Dataset Imbalance Ratio
Carpet	Texture	5	308	89	3.46	280+56 = 336	500	0.67
Grid	Texture	5	285	57	5.00	264+53 = 317	500	0.63
Leather	Texture	5	277	92	3.01	245+49 = 294	500	0.59
Tile	Texture	5	263	84	3.13	230+46 = 276	500	0.55
Wood	Texture	5	266	60	4.43	247+50 = 297	500	0.59
Bottle	Object	3	229	63	3.64	209+42 = 251	300	0.84
Cable	Object	8	282	92	3.07	224+45 = 269	800	0.34
Capsule	Object	5	242	109	2.22	219+44 = 263	500	0.53
Hazelnut	Object	4	431	70	6.16	391+79 = 470	400	1.18
Metal Nut	Object	4	242	93	2.60	220+44 = 264	400	0.66
Pill	Object	7	293	141	2.08	267+54 = 321	700	0.46
Screw	Object	5	361	119	3.03	320+64 = 384	500	0.77
Toothbrush	Object	1	72	30	2.40	60+12 = 72	100	0.72
Transistor	Object	4	273	40	6.83	213+43 = 256	400	0.64
Zipper	Object	7	272	119	2.29	240+48 = 288	600	0.48

2.1. Challenges in Real-World Anomaly Detection

The distribution of normal and defective samples across categories is presented in Table 1. The imbalance ratio (IR), defined as:

$$IR = \frac{N_{\text{normal}}}{N_{\text{defective}}} \tag{1}$$

varies significantly, with some categories exhibiting severe class imbalance, impacting model generalization. Three key limitations hinder generalization in practical deployment:

- **Class Imbalance:** Many categories, such as *transistor* (IR = 6.83) and *hazelnut* (IR = 6.16), exhibit severe imbalance, reducing the model's exposure to defective patterns and biasing learning toward normal features.
- **Defect Diversity:** Defect variety is inconsistent. For example, *toothbrush* contains only a single defect type, limiting generalization, whereas *zipper* includes up to seven, providing richer anomaly patterns.
- **Limited Realism:** Controlled lighting, clean backgrounds, and clearly visible defects simplify the MVTec AD dataset, contrasting with real-world settings where defects are subtle (e.g., micro-cracks), and lighting, occlusion, and noise add complexity.

2.2. Motivation for Synthetic Data

To address these challenges, we propose a generative framework that synthesizes diverse, high-fidelity anomalies with controllable attributes (location, size, texture, contrast). Our contributions include:

- **Balancing class distributions** by generating synthetic anomalies for under-represented categories.
- **Enhancing defect diversity** using GANs guided by structured masks and texture priors.
- **Simulating defect subtlety** via controlled perturbations that mimic realistic wear or micro-damage.

This synthetic augmentation not only enriches training data but also enables systematic study of how anomaly type, severity, and realism influence detection. Our goal is to bridge the domain gap between synthetic and real defects, enhancing robustness of deep anomaly detectors in industrial contexts.

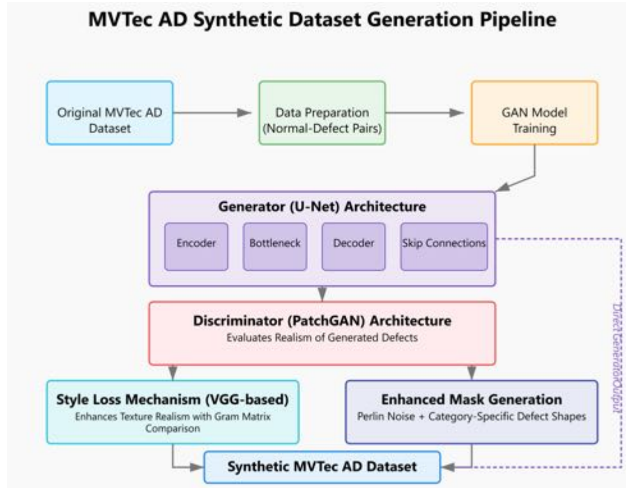

Figure 1. Synthetic defect generation pipeline. Normal images are combined with category-specific masks, producing synthetic defect images via a GAN-based model.

3. Methodology

We propose a modular deep learning framework to generate high-fidelity synthetic defect datasets tailored for industrial anomaly detection. Our method leverages conditional GANs guided by defect masks and incorporates controllable parameters, enabling realistic and diverse anomaly generation with fine-grained spatial precision.

Overview. As illustrated in Figure 1, our pipeline begins with defect-free samples from the MVTec AD dataset. Each image is paired with a synthetically generated defect mask, forming triplets of *(normal image, mask, synthetic defect)*. A U-Net-based generator learns to localize and synthesize defects guided by these masks, while a PatchGAN discriminator enforces local texture realism.

Mask-Guided Defect Synthesis. The generator G receives a normal image x and defect mask m to produce a defect-laden output $G(x, m)$. The final synthesized image is constructed as:

$$y = x \odot (1 - m) + G(x, m) \odot m, \tag{2}$$

ensuring that only masked regions are modified. The U-Net architecture with skip connections enables preservation of fine spatial structures in unmasked areas.

Discriminator and Style Loss. A PatchGAN-based discriminator D evaluates local image patches, encouraging the synthesis of texture-consistent anomalies. To further enhance visual fidelity, we introduce a style loss:

$$L_{\text{style}} = \sum_l \lambda_l \| G_l(F_l(y)) - G_l(F_l(\hat{y})) \|_F^2, \tag{3}$$

where $F_l(\cdot)$ denotes feature activations and $G_l(\cdot)$ the corresponding Gram matrices [8]. This encourages the synthesized image $\hat{y}$ to match the style of real defect textures.

Category-Specific Mask Generation. Defect masks are generated using category-aware Perlin noise [13] modulated by geometry and boundary proximity:

$$m(x, y) = \begin{cases} 1, & \text{if } s(x, y) = 1 \land P(x, y) > \tau_1 \\ 1, & \text{if } s(x, y) = 0 \land P(x, y) > \tau_2 \land d(x, y) < \epsilon \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

This generates morphologically plausible masks that respect object boundaries and real-world defect variability.

Controllable Defect Attributes. Our system supports parametric control over defect attributes, location, size, intensity, and texture, allowing the generation of diverse samples tailored for specific inspection tasks. This facilitates targeted training and robust evaluation of anomaly detection models [16].

Loss Functions. The generator is optimized with a composite loss:

$$L_G = \lambda_{\text{adv}} L_{\text{adv}} + \lambda_{\text{pixel}} L_{\text{pixel}} + \lambda_{\text{style}} L_{\text{style}}, \tag{5}$$

where L_{adv} promotes realism, L_{pixel} preserves structural consistency, and L_{style} enforces perceptual similarity. The discriminator is trained using a standard GAN loss.

Model Compression via Quantization. To reduce inference cost, we apply post-training quantization by mapping 32-bit weights to b-bit integers [17].

4. Experiments and Results

4.1. Experimental Setup

All experiments were conducted on the Kaggle platform using NVIDIA Tesla P100 and G4 GPUs. The implementation was in PyTorch to facilitate reproducibility.

Architecture Details. Our adversarial framework consists of a U-Net-based generator with skip connections and a five-layer PatchGAN discriminator, both employing spectral normalization and LeakyReLU activations. A VGG-16-based perceptual loss module enhances style fidelity by extracting features from layers `relu1_2` to `relu4_3`.

Training Protocol. Models were trained using the Adam optimizer ($\beta_1 = 0.5$, $\beta_2 = 0.999$) for 30 epochs with a batch size of 4. The learning rate was set to 2×10^{-4} and linearly decayed post 20 epochs. A 3-epoch warm-up and gradient clipping (threshold 1.0) were employed for stability. Early stopping was triggered after three non-improving validation epochs.

Loss Function. The total objective combines adversarial, pixel-wise, and perceptual components:

$$\mathcal{L}_{total} = \lambda_{adv}\mathcal{L}_{GAN} + \lambda_{pixel}\mathcal{L}_{L1} + \lambda_{style}\mathcal{L}_{VGG}, \tag{6}$$

where $\lambda_{adv} = 1$, $\lambda_{pixel} = 100$, and $\lambda_{style} = 100$ ensure a balance between realism and content preservation.

Table 2. Benchmark comparison between Anomalib and our baseline implementation on the MVTec AD and Synthetic MVTec AD datasets. The difference between our implementation and Anomalib is shown in parentheses.

Category	Image-Level AUC			Pixel-Level AUC		
	Anomalib MVTec AD	Our Impl. (Diff.) MVTec AD	Our Impl. (Diff.) Synthetic	Anomalib MVTec AD	Our Impl. (Diff.) MVTec AD	Our Impl. (Diff.) Synthetic
bottle	0.9990	1.0000 (+0.0010)	1.0000 (+0.0010)	0.9850	0.9737 (-0.0113)	0.9661 (-0.0189)
cable	0.8780	0.9202 (+0.0422)	1.0000 (+0.1220)	0.9700	0.9529 (-0.0171)	0.9793 (+0.0093)
capsule	0.9270	0.9358 (+0.0088)	0.9957 (+0.0687)	0.9880	0.9791 (-0.0089)	0.9886 (+0.0006)
carpet	0.9950	0.9992 (+0.0042)	0.9997 (+0.0047)	0.9910	0.9844 (-0.0066)	0.9872 (-0.0038)
grid	0.9420	0.9582 (+0.0162)	0.9994 (+0.0574)	0.9700	0.9390 (-0.0310)	0.9832 (+0.0132)
hazelnut	0.9640	0.9586 (-0.0054)	1.0000 (+0.0360)	0.9850	0.9712 (-0.0138)	0.9753 (-0.0097)
leather	1.0000	1.0000 (+0.0000)	0.9993 (-0.0007)	0.9930	0.9858 (-0.0072)	0.9919 (-0.0011)
metal nut	0.9890	0.9936 (+0.0046)	1.0000 (+0.0110)	0.9820	0.9641 (-0.0179)	0.9551 (-0.0269)
pill	0.9390	0.9193 (-0.0197)	0.9986 (+0.0596)	0.9660	0.9510 (-0.0150)	0.9758 (+0.0098)
screw	0.8450	0.8203 (-0.0247)	0.9950 (+0.1500)	0.9880	0.9724 (-0.0156)	0.9955 (+0.0075)
tile	0.9740	0.9899 (+0.0159)	0.9987 (+0.0247)	0.9550	0.9027 (-0.0523)	0.9578 (+0.0028)
toothbrush	0.9470	1.0000 (+0.0530)	1.0000 (+0.0530)	0.9780	0.9806 (+0.0026)	0.9883 (+0.0103)
transistor	0.9750	0.9925 (+0.0175)	0.9989 (+0.0239)	0.9680	0.9513 (-0.0167)	0.9805 (+0.0125)
wood	0.9930	0.9939 (+0.0009)	1.0000 (+0.0070)	0.9570	0.9088 (-0.0482)	0.9658 (+0.0088)
zipper	0.9820	0.9422 (-0.0398)	0.9972 (+0.0152)	0.9800	0.9761 (-0.0039)	0.9872 (+0.0072)
avg	0.9500	0.9616 (+0.0116)	0.9988 (+0.0488)	0.9790	0.9596 (-0.0194)	0.9785 (-0.0005)

4.2. Results and Analysis

We evaluate the performance of our synthetic defect generation approach across the 15 categories of the MVTec AD dataset, comparing our method to Anomalib [2] and our baseline PaDiM [7] implementation.

4.2.1. Benchmark Results on MVTec AD Dataset

Table 2 presents a comparative analysis between Anomalib and our method. Performance is assessed using Area Under the Curve (AUC) at both the image and pixel levels.

Image-Level AUC. Our method shows an average gain of +0.0116 over Anomalib, with the most significant improvement in the *toothbrush* category (+0.0530). This indicates superior defect classification capability.

Pixel-Level AUC. Although our method remains competitive, the pixel-level AUC shows a slight decrease of -0.0194 on average, highlighting challenges in spatial anomaly localization. Notably, *tile* and *wood* categories exhibit larger declines, suggesting potential areas for improvement in segmentation precision. The disparity between image-level and pixel-level performance suggests that while our method excels in anomaly detection, further work is needed to improve spatial localization. Future research will explore enhanced feature extraction and thresholding strategies to address this.

4.2.2. Results on Synthetic MVTec AD Dataset

Table 2 compares our approach with Anomalib on the synthetic MVTec AD dataset. Our method shows a notable improvement in image-level AUC (+0.0488 on average), with significant gains in categories like *cable* (+0.1220) and *screw* (+0.1500), highlighting the efficacy of our synthetic defect generation in boosting classification performance.

Pixel-Level AUC. At the pixel level, performance remains comparable to Anomalib, with a negligible average difference of -0.0005. This indicates that our synthetic approach retains robust spatial localization, especially in categories such as *cable* and *screw*, though further refinement is needed for categories like *metal nut* and *bottle*. These re-

Table 3. Performance comparison on MVTec Synthetic Data of different quantization methods across multiple categories using evaluation metrics Image-level AUC (I-AUC), Image-level AUC (P-AUC), Compression Ratio (CR), and Huffman encoding-based CR (HCR).

Metric	I-AUC	P-AUC	CR	HCR	I-AUC	P-AUC	CR	HCR	I-AUC	P-AUC	CR	HCR	I-AUC	P-AUC	CR	HCR	I-AUC	P-AUC	CR	HCR
Category	Bottle				Cable				Capsule				Carpet				Grid			
1-bit Q	0.6583	0.6035	32.0	32.0	0.6841	0.8111	32.0	32.0	0.6425	0.7119	32.0	32.0	0.4607	0.6291	32.0	32.0	0.6507	0.7275	32.0	32.0
2-bit Q	0.6470	0.7081	16.0	29.1	0.6604	0.8302	16.0	29.1	0.6276	0.7708	16.0	29.1	0.3780	0.5799	16.0	29.1	0.7852	0.7653	16.0	29.1
4-bit Q	0.4877	0.4563	8.0	25.1	0.5986	0.5856	8.0	25.1	0.8863	0.1628	8.0	25.1	0.7790	0.8026	8.0	25.1	0.4738	0.3054	8.0	25.1
8-bit Q	1.0000	0.9660	4.0	7.3	1.0000	0.9793	4.0	7.3	0.9944	0.9885	4.0	7.3	0.9996	0.9871	4.0	7.3	0.9994	0.9830	4.0	7.3
Baseline	1.0000	0.9662	1.0	1.0	1.0000	0.9793	1.0	1.0	0.9942	0.9886	1.0	1.0	0.9996	0.9872	1.0	1.0	0.9994	0.9832	1.0	1.0
Category	Hazelnut				Leather				Metal Nut				Pill				Screw			
1-bit Q	0.6221	0.4522	32.0	32.0	0.7329	0.7285	32.0	32.0	0.9125	0.2993	32.0	32.0	0.6798	0.2310	32.0	32.0	0.4502	0.6113	32.0	32.0
2-bit Q	0.6443	0.2889	16.0	29.1	0.7147	0.6975	16.0	29.1	0.8653	0.2744	16.0	29.1	0.6788	0.2240	16.0	29.1	0.4630	0.6107	16.0	29.1
4-bit Q	0.6565	0.4511	8.0	25.1	0.8552	0.9541	8.0	25.1	0.9444	0.4470	8.0	25.1	0.6131	0.4020	8.0	25.1	0.6484	0.2623	8.0	25.1
8-bit Q	1.0000	0.9752	4.0	7.3	0.9996	0.9919	4.0	7.3	1.0000	0.9557	4.0	7.3	0.9982	0.9757	4.0	7.3	0.9947	0.9955	4.0	7.3
Baseline	1.0000	0.9753	1.0	1.0	0.9997	0.9919	1.0	1.0	1.0000	0.9552	1.0	1.0	0.9980	0.9758	1.0	1.0	0.9950	0.9955	1.0	1.0
Category	Tile				Toothbrush				Transistor				Wood				Zipper			
1-bit Q	0.4935	0.5186	32.0	32.0	0.6208	0.2019	32.0	32.0	0.5107	0.5673	32.0	32.0	0.5305	0.5324	32.0	32.0	0.4517	0.6509	32.0	32.0
2-bit Q	0.5332	0.4688	16.0	29.1	0.6217	0.1894	16.0	29.1	0.4994	0.7387	16.0	29.1	0.5394	0.5481	16.0	29.1	0.4787	0.7617	16.0	29.1
4-bit Q	0.8222	0.8100	8.0	25.1	0.6450	0.2065	8.0	25.1	0.4715	0.5913	8.0	25.1	0.7564	0.7933	8.0	25.1	0.5235	0.6462	8.0	25.1
8-bit Q	0.9987	0.9580	4.0	7.3	1.0000	0.9883	4.0	7.3	0.9988	0.9810	4.0	7.3	1.0000	0.9651	4.0	7.3	0.9968	0.9871	4.0	7.3
Baseline	0.9984	0.9579	1.0	1.0	1.0000	0.9883	1.0	1.0	0.9989	0.9804	1.0	1.0	1.0000	0.9657	1.0	1.0	0.9969	0.9872	1.0	1.0

sults affirm the potential of synthetic defect generation to improve classification accuracy while maintaining spatial anomaly localization. Future work will focus on refining defect synthesis to enhance pixel-wise anomaly detection.

Visual Results. Figure 2 shows qualitative comparisons among normal images, real anomalies, and synthetically generated defects. Our synthetic anomalies accurately replicate real-world structural and texture variations. For instance, *bottle* defects capture subtle texture discontinuities, while *screw* defects exhibit realistic localized surface anomalies.

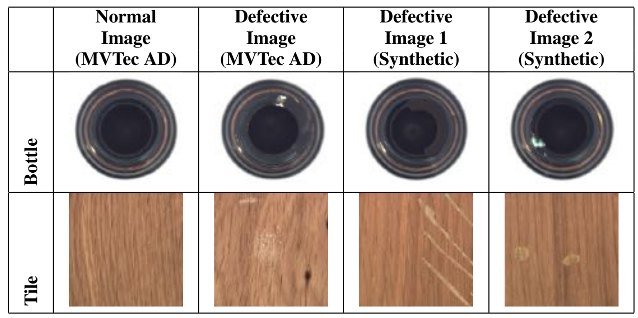

Figure 2. Visualization of normal and defective images from MVTec AD and synthetic dataset

4.3. Implications for Synthetic Data Generation

Our findings underscore the effectiveness of synthetic data in anomaly detection. As shown in Table 2, well-designed synthetic defects improve anomaly detection, with an average image-level AUC increase of +0.0488 over Anomalib. Categories with complex geometry and texture, such as *cable* and *screw*, benefit most, while simpler textures see more modest improvements. This highlights the importance of tailoring synthetic data generation to object complexity. In Table 2, we observe that categories with higher imbalance ratios benefit most from synthetic data augmentation. This aligns with real-world challenges where defect samples are scarce, reinforcing the value of synthetic data for addressing such imbalances.

4.4. Impact of Model Compression

Model quantization substantially reduces model size while striving to preserve performance. As shown in Table 3, 8-bit quantization maintains near-baseline accuracy across all categories, with minimal degradation in I-AUC and P-AUC. For instance, in the *bottle* category, 8-bit quantization achieves an I-AUC of 1.0000 and P-AUC of 0.9660, nearly identical to the baseline (I-AUC: 1.0000, P-AUC: 0.9662), while achieving a 4× compression. However, aggressive quantization (e.g., 1-bit and 2-bit) results in notable accuracy degradation. In the *capsule* category, 1-bit quantization yields an I-AUC of 0.6425 and P-AUC of 0.7119, illustrating the trade-off between model size and accuracy. The performance drop is more pronounced in complex categories like *screw* and *zipper*, where spatial and textural feature fidelity is critical.

5. Conclusion

We presented a synthetic data generation framework for industrial anomaly detection, combining a U-Net generator, PatchGAN discriminator, VGG-based style loss, and category-specific mask generation. Evaluated on 15 MVTec AD categories, our approach improved Image-Level AUC by +0.0488 on average, with notable gains in complex classes like *cable* and *screw*. Key insights include: (1) synthetic data is most impactful in data-scarce regimes, (2) domain-specific fidelity outweighs photorealism, and (3) customized generation strategies are essential across categories. We further demonstrated that 8-bit quantization reduces model size by 3-4× with minimal accuracy loss, supporting efficient edge deployment. Future work will extend to high-resolution, 3D-aware synthesis, physical modeling, and adaptation to unseen categories, with applications in fields like medical imaging and remote sensing.

References

[1] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 622–637. Springer, 2019. 1

[2] Samet Akcay, Dick Ameln, Ashwin Vaidya, Barath Lakshmanan, Nilesh Ahuja, and Utku Genc. Anomalib: A deep learning library for anomaly detection. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 1706–1710. IEEE, 2022. 3

[3] 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

[4] Raghavendra Chalapathy, Edward Toth, and Sanjay Chawla. Group anomaly detection using deep generative models. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part I 18*, pages 173–189. Springer, 2019. 1

[5] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. 1

[6] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 1

[7] 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. 1, 3

[8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 3

[9] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 1

[10] Rahima Khanam, Muhammad Hussain, Richard Hill, and Paul Allen. A comprehensive review of convolutional neural networks for defect detection in industrial applications. *IEEE Access*, 2024. 1

[11] Sergio Naval Marimont and Giacomo Tarroni. Anomaly detection through latent space restoration using vector quantized variational autoencoders. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1764–1767. IEEE, 2021. 1

[12] Seungju Park, Kyung Hwa Lee, Beomseok Ko, and Namkug Kim. Unsupervised anomaly detection with generative adversarial networks in mammography. *Scientific Reports*, 13 (1):2925, 2023. 1

[13] Ken Perlin. An image synthesizer. *ACM Siggraph Computer Graphics*, 19(3):287–296, 1985. 3

[14] 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

[15] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pages 4393–4402. PMLR, 2018. 1

[16] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019. 3

[17] Tushar Shinde. Adaptive quantization and pruning of deep neural networks via layer importance estimation. In *Workshop on Machine Learning and Compression, NeurIPS 2024*, 2024. 3

[18] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4820–4828, 2016. 1

[19] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8330–8339, 2021. 1