# Towards Real-Time Unsupervised Visual Anomaly Detection in Industrial Settings

Miguel Campos-Romero[1,2][0009−0002−8576−5629], Manuel Carranza-García[1][0000−0002−4729−8604], and José C. Riquelme[1][0000−0002−8243−2186]

[1] Departamento de Lenguajes y Sistemas Informáticos. University of Sevilla, Av. Reina Mercedes, s/n, 41012 Sevilla, Spain
{migcamrom1, mcarranzag, riquelme}@us.es
[2] TKH AI Technology, Amstelplein 52, Amsterdam, 1096 BC, The Netherlands
m.campos@tkhaitechnology.com

**PhD thesis information:**
Start date: 10/2022.
Supervisors: Manuel Carranza García, José Cristóbal Riquelme Santos.
Department: Lenguajes y Sistemas Informáticos, Universidad de Sevilla.
Expected defense date: 09/2027.
Mode: Part-time Industrial PhD.

**Abstract.** Visual anomaly detection is a key task in automated industrial inspection, yet its practical deployment is often hindered by the scarcity and extreme heterogeneity of anomalous samples. While unsupervised methods trained exclusively on nominal data have emerged as a viable alternative, existing approaches frequently struggle to balance low-latency inference with the ability to detect both structural and logical inconsistencies. This thesis proposes a progressive research framework to bridge this gap. Our first contribution, InAI, leverages normalizing flows for high-speed structural defect detection, establishing a new state of the art on the MVTec AD benchmark while significantly reducing latency compared to embedding-based methods. Building on this, we developed MuDeNet, a dual-branch teacher-student architecture that integrates global contextual reasoning to identify logical anomalies, which are violations of semantic constraints that appear visually normal in isolation. MuDeNet demonstrates high competitive performance on the MVTec LOCO benchmark, demonstrating that structural and logical capabilities can be unified without compromising real-time constraints. Our ongoing research addresses advanced industrial challenges, such as high-variability environments and non-centered defects, as introduced in the MVTec AD 2 dataset. By synthesizing the architectural strengths of InAI and MuDeNet, this work aims to provide a robust, end-to-end framework for reliable, high-throughput anomaly detection in large-scale manufacturing.

**Keywords:** Deep Learning · Computer Vision · Unsupervised Learning · Anomaly Detection · Novelty Detection

# 1 Introduction

Visual anomaly detection aims to identify image patterns that deviate from normality. This task is especially relevant in industrial inspection, medical imaging [1], surveillance [3], and autonomous systems [14], where reliable defect or event detection is essential for safety, quality, and process control. In manufacturing, automated inspection enables scalable and repeatable quality control while reducing the cost and variability of manual inspection.

Existing approaches include reconstruction-based models [5], generative methods [2, 17], embedding-based techniques [12, 21], transformers [9, 19], teacher-student frameworks [6, 13], and distribution-based methods [22].

Despite this progress, current methods still struggle to satisfy industrial requirements. Reconstruction methods may introduce false positives; patch-based and memory-based approaches can be computationally expensive or limited in their contextual reasoning; and large vision-language models are often unsuitable for strict real-time deployment. A central open problem is therefore to detect both structural anomalies, which correspond to local visual defects, and logical anomalies, which arise from contextual or semantic inconsistencies, without compromising inference speed [8, 19].

This thesis addresses this gap through a progressive research strategy for real-time unsupervised industrial anomaly detection. The research explores novel frameworks designed to unify structural and logical defect detection within a high-throughput paradigm. By balancing accuracy with strict efficiency constraints, this work aims to provide a robust solution for next-generation automated inspection systems.

The following sections present the objectives, methodology, relevance, obtained results, and future research directions.

# 2 Objectives

This thesis aims to advance unsupervised anomaly detection for real-time manufacturing inspection through a cumulative research strategy. Each stage addresses a more demanding scenario while preserving the capabilities achieved in previous phases. The research is organized around three main objectives:

1. **High-throughput structural anomaly detection.** The first objective addresses the detection of structural anomalies, such as scratches, cracks, holes, deformations, or contamination, using MVTec AD as the main benchmark [4]. To this end, we proposed InAI [10], an unsupervised method that models nominal feature distributions through normalizing flows. By mapping defect-free features to a Gaussian latent space, anomalous regions can be identified as deviations from the learned distribution. Its multiscale and multisemantic formulation improves robustness while avoiding nearest-neighbor retrieval, enabling high-speed inference compared with embedding-based methods such as PatchCore [21].
2. **Unified detection of structural and logical inconsistencies.** The second objective extends anomaly detection to logical inconsistencies, which cannot be identified from local appearance alone but require contextual reasoning. This stage focuses on MVTec LOCO [8], a benchmark that includes both structural and logical anomalies. To address this setting, we proposed MuDeNet [11], a lightweight dual-branch teacher-student architecture that combines local structural modeling with global contextual reasoning. This design enables robust detection and localization of both anomaly types under real-time constraints.

3. **Robustness to high-variability industrial environments.** The third objective, currently ongoing, addresses the more realistic scenarios introduced by MVTec AD 2 [15], including higher nominal variability, smaller and non-centered defects, border anomalies, and varying lighting conditions. This phase aims to synthesize the strengths of InAI and MuDeNet into a unified framework that preserves real-time performance while improving robustness in challenging industrial conditions.

## 3 Methodology and Work Plan

This section outlines the structured approach used to address each research objective. Given the cumulative nature of this thesis, a consistent methodology is applied to each benchmark (MVTec AD, LOCO, and AD 2) phase, as visualized in Fig. 1. This ensures that every contribution, from structural to logical anomaly detection, is built upon a rigorous foundation of baseline reproduction and iterative validation. The research workflow for each objective is organized as follows:

1. **State-of-the-art review:** Exploration and study of existing methods evaluated on the target benchmark.
2. **Baseline establishment:** Reproduction of reported scores to select a representative population for comparative study.
3. **Critical analysis:** Identification of the main strengths and weaknesses of current state-of-the-art approaches.
4. **Iterative prototyping:** Design and preliminary validation of the initial methodological proposals.
5. **Methodological refinement:** Development and iterative testing until a robust final architecture is obtained.
6. **Cross-dataset validation:** Extending the evaluation of the final proposal to other state-of-the-art datasets to ensure generalization.
7. **Dissemination:** Writing and publication of a research manuscript.

While each project follows the same sequential workflow, the execution has become increasingly efficient. The first stage (InAI) required an exhaustive study of normalizing flows and density estimation, establishing a strict reproducibility filter that favored open-source, high-impact published research. Building on this foundation, the development of MuDeNet was significantly streamlined. Prior experience with regression models and comparative analysis allowed for a faster transition to the architectural design phase. This cumulative expertise is now being applied to the third objective (MVTec AD 2), which is currently in development. By leveraging the technical and methodological proficiency gained from previous iterations, we have been able to navigate the challenges of this recent benchmark rapidly, ensuring a more agile path toward final validation and publication.

## 4 Relevance

Industrial visual inspection is a key component of large-scale manufacturing, as it contributes to quality control, reduces production costs, and enables the detection of defective products. Compared to manual inspection, automated systems offer greater consistency, scalability, and throughput while avoiding the limitations of human fatigue

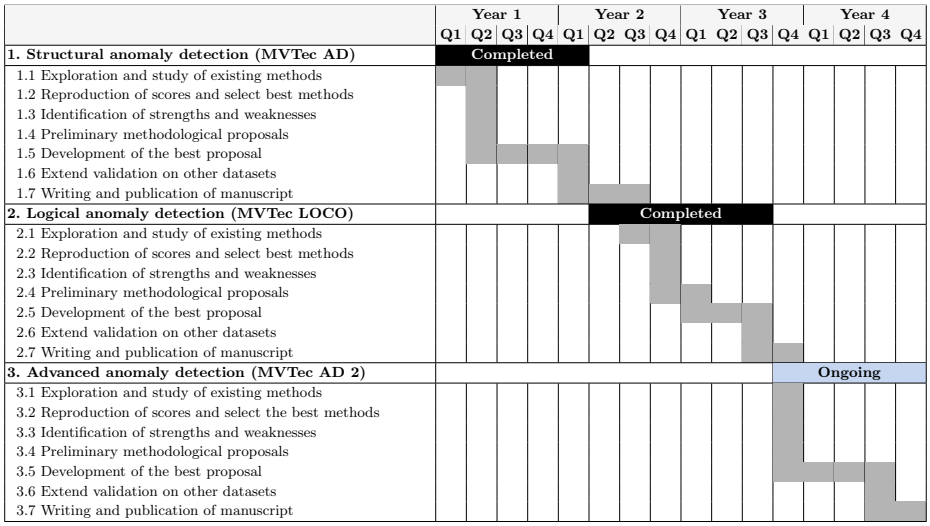



**Fig. 1.** Chronogram of the research plan (2022–2026). The project is structured into four annual cycles, with each objective following a systematic workflow across quarterly intervals.

and subjectivity in decision-making. However, the transition from successful laboratory experiments to reliable production-line deployment remains a significant challenge. This research is uniquely positioned to bridge this gap, as it is conducted within an Industrial Doctoral Program in collaboration with TKH Group.

Despite the practical importance of this field, unsupervised anomaly detection remains a complex problem. In industrial settings, anomalous samples are scarce and difficult to collect, which limits the applicability of supervised learning approaches. This necessitates the development of methods that can learn exclusively from defect-free nominal data. Additionally, industrial operations impose requirements that are not always addressed in existing literature, such as guaranteed real-time inference, high robustness across varying conditions, and the critical need for a near-zero false-positive rate or total recall.

A central contribution of this work is the validation of the "academic-to-industrial" transfer process. Our research demonstrates that high performance on static datasets, which often focus on maximizing metrics like AU-ROC, does not inherently translate to operational success. Ensuring an algorithm is production-ready requires addressing factors typically absent from benchmark datasets, such as meeting the strictly required false-positive rates or standardizing image acquisition for training and inference across diverse variables. This involves controlling lighting conditions, camera settings, lens selection, distance to the target, color quality and gamut, background, speed, and even mechanical vibrations in the machinery that could induce motion blur during image acquisition. While these aspects are rarely the focus of academic settings, they are essential for the viability of any model in a real-world environment.

Therefore, the relevance of this work lies in its dual perspective: beyond proposing novel architectures, it develops the end-to-end pipelines necessary to ensure that

academically sound models remain robust under the rigorous constraints of real-world industrial use.

## 5 Results

The first two phases of this thesis have successfully fulfilled their respective research objectives. Each proposed method was validated against the corrresponding benchmarks and compared with state-of-the-art (SOTA) approaches, demonstrating significant improvements in both detection accuracy and computational efficiency.

### 5.1 First stage

The initial objective led to the development of *InAI*, an unsupervised anomaly detection method based on normalizing flows, multiscale feature analysis, and multisemantic learning. This work resulted in an open-access publication in the Engineering Applications of Artificial Intelligence journal [10]. As illustrated in Figure 2, InAI employs an encoder-decoder architecture that maps defect-free features to a Gaussian latent space. Anomalies are subsequently identified as regions that deviate from this learned distribution.

InAI was validated primarily on the MVTec AD benchmark, achieving a SOTA Area Under the Receiver Operating Characteristic (AU-ROC) curve of 99.9%. Notably, the model demonstrated high data efficiency, outperforming competitors even when trained on only 20% of the nominal data. Further validation on MVTec LOCO and the Surface Defect Magnetic (SDMT) [16] dataset confirmed its robustness. Specifically, on MVTec LOCO, it achieved an AU-ROC of 87.0%, improving the previous SOTA by 3.6 percentage points.

Furthermore, InAI established a new efficiency benchmark, with a 1.9 ms latency, 506 images/s throughput, and 362 MB memory usage. Compared to its leading competitor, InAI is approximately 10.7 times faster (reducing latency from 20.3 ms) and offers a nearly 12-fold increase in processing capacity, rising from 42 to 506 images/s.

### 5.2 Second stage

Building on these foundations, the second objective led to *MuDeNet*, a lightweight dual-branch teacher-student architecture designed to detect both structural and logical anomalies. While structural anomalies are local, logical anomalies require global contextual reasoning. This objective resulted in an open-access publication in the Information Fusion journal [11]. As shown in Figure 3, MuDeNet addresses this duality by fusing responses from two specialized modules during inference.

MuDeNet was first optimized on MVTec LOCO to ensure its dual-branch design effectively captured both anomaly types. The validation scope was then expanded to include MVTec AD and the ViSA [27] dataset, the latter providing a more diverse set of 12 industrial categories compared to the more specific SDMT.

On MVTec LOCO, MuDeNet significantly advanced the SOTA, surpassing the closest competitors by 5.8 and 8.0 percentage points in AU-ROC for anomaly detection and segmentation, respectively. Across MVTec AD, MVTec LOCO, and ViSA, the model achieved the best average performance, outperforming the leading competitor by 1.2 and 4.0 percentage points.

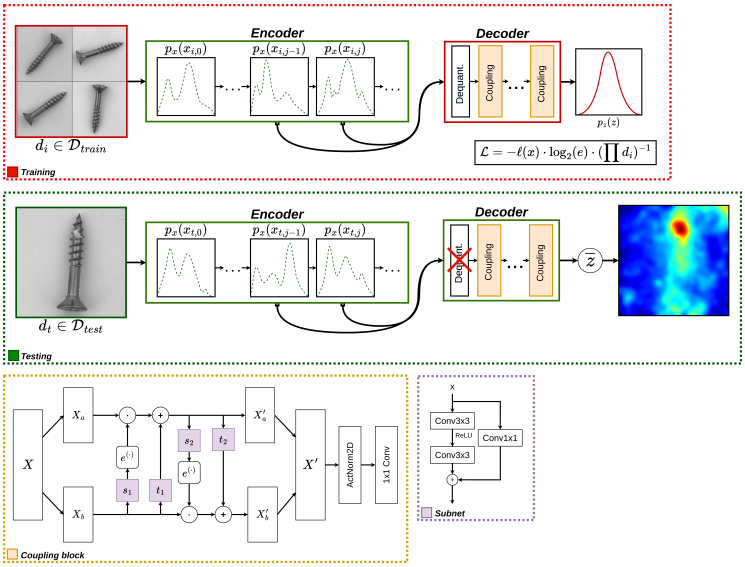



**Fig. 2.** Overview of InAI [10]. The method extracts semantic feature maps from the input images and models the distribution of defect-free features in a Gaussian latent space using normalizing flows. During inference, image regions that cannot be mapped to this learned distribution are identified as anomalies.

Furthermore, MuDeNet optimized the efficiency established in the previous stage, reducing latency to 1.6 ms and memory usage to 38.2 MB. By processing approximately 150 additional images per second compared to InAI, it demonstrates that superior logical detection can be achieved with a reduced computational footprint.

### 5.3 Performance analysis

Tables 1 and 2 summarize the quantitative results obtained across image, pixel, and region levels, comparing the proposed methods against several state-of-the-art competitors. Pixel-level AU-ROC is influenced by anomaly size, whereas region-level AU-ROC gives equal importance to each anomalous region, making it more suitable when small defects are critical.

Regarding image-level anomaly classification, InAI achieves the best results on MVTec AD, the main benchmark for structural anomaly detection. Subsequently, MuDeNet extended this performance to the more complex MVTec LOCO benchmark, becoming the reference for logical anomaly detection while matching InAI's performance on ViSA. This progression demonstrates how the research has successfully transitioned from mastering structural defect identification to addressing global semantic inconsistencies, ultimately improving the aggregate performance across all datasets.

A similar trend is observed in Table 2 at the pixel and region levels. While InAI already surpassed existing baselines, MuDeNet further enhanced localization robustness, especially within the logical anomaly domain. Since the MVTec LOCO benchmark does not support pixel-level AU-ROC, we report only its region-level performance.

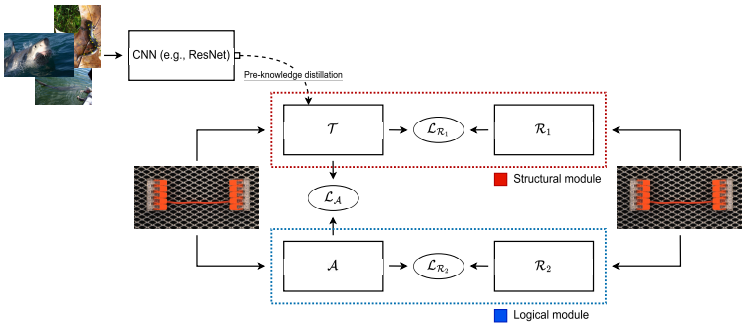



**Fig. 3.** Overview of MuDeNet [11]. The framework consists of two complementary teacher-student modules: a structural module ($\mathcal{T}$ and $\mathcal{R}_1$) for local defects, and a logical module ($\mathcal{A}$ and $\mathcal{R}_2$) for global anomalies. During inference, discrepancies between teacher and student outputs are fused across feature levels to obtain the final anomaly prediction.

**Table 1.** Image-level anomaly classification AU-ROC.

| Method | MVTec LOCO | MVTec AD | VisA | Average |
|---|---|---|---|---|
| GCAD [8] | 83.4 | 88.9 | 83.7 | 85.3 |
| SimpleNet [18] | 77.5 | 98.2 | 87.9 | 87.9 |
| TFA [19] | 81.6 | 98.7 | 91.4 | 90.5 |
| Patchcore [21] | 80.2 | 98.9 | 94.2 | 91.1 |
| DADF [25] | 83.7 | 98.3 | 94.1 | 92.0 |
| SFRAD [26] | 82.6 | 99.1 | 95.4 | 92.3 |
| AST [23] | 83.3 | 98.9 | 94.9 | 92.3 |
| SLSG [24] | 86.8 | 96.8 | 93.7 | 92.4 |
| REB [20] | 86.8 | 99.0 | 94.5 | 93.4 |
| InAI (ours) [10] | 87.5 | **99.6** | **96.6** | **94.6** |
| MuDeNet (ours) [11] | **93.3** | 99.5 | **96.6** | **96.4** |

**Table 2.** Anomaly localization AU-ROC. Values are reported as (pixel-level, region-level) tuples. For MVTec LOCO, only the region-level AU-ROC is provided because pixel-level evaluation cannot properly handle logical anomalies.

| Method | MVTec LOCO | MVTec AD | VisA | Average |
|---|---|---|---|---|
| SimpleNet [18] | 28.4 | (97.8, 89.6) | (91.8, 68.9) | (72.7, 62.3) |
| Patchcore [21] | 33.7 | (98.1, 92.7) | (94.6, 79.7) | (75.5, 68.7) |
| AST [23] | 38.7 | (94.2, 81.2) | (96.7, 81.5) | (76.6, 67.1) |
| SLSG [24] | 55.7 | (97.1, 89.5) | (94.8, 76.6) | (82.4, 73.9) |
| REB [20] | 63.4 | (98.9, 94.3) | (96.5, 85.7) | (86.2, 81.1) |
| GCAD [8] | 67.7 | (97.5, 90.9) | (97.6, 83.7) | (87.6, 83.5) |
| DADF [25] | 67.3 | (98.4, 94.1) | (98.5, 86.3) | (88.0, 82.5) |
| TFA [19] | 67.0 | (98.7, 94.5) | (98.8, 89.8) | (88.1, 83.8) |
| SFRAD [26] | 66.6 | (98.9, 93.5) | (98.8, 92.0) | (88.1, 84.0) |
| InAI (ours) [10] | 67.5 | (98.4, 94.6) | (98.8, 89.9) | **(88.2, 84.0)** |
| MuDeNet (ours) [11] | **75.7** | (**98.6**, **95.0**) | (**99.0**, **93.3**) | **(91.1, 88.0)** |

Figure 4 presents 8 of the 32 anomaly scenarios selected from the three datasets. For each scenario, the input image and its corresponding ground-truth mask are displayed alongside the anomaly maps predicted by InAI, MuDeNet, and the eight competing approaches, where higher intensity regions indicate a higher probability of defects. The first two prediction columns correspond to our proposed frameworks: InAI and MuDeNet. These qualitative results underscore the complexity of industrial anomaly detection. The defects often occupy a small fraction of the image, yet most baseline methods generate a considerable number of false positives. However, the visualization highlights a clear progressive improvement between our two contributions. While InAI establishes a strong detection baseline, it exhibits a higher false-positive rate than MuDeNet, our most recent architecture. This refinement is particularly evident in rows 4, 5, and 6, where MuDeNet drastically suppresses these false positives compared to InAI while simultaneously outperforming all state-of-the-art competitors in localization accuracy.

Figure 5 illustrates the trade-off between inference latency and image-level anomaly classification performance. The horizontal axis represents latency in logarithmic scale, while the vertical axis reports the average AU-ROC across all datasets. This visualization highlights the iterative advancements made throughout this thesis: while InAI established a robust baseline by balancing high accuracy with low-latency execution, MuDeNet significantly optimized this trade-off.

## 6 Conclusions and future work

This paper presented the progress of an industrial thesis focused on real-time unsupervised anomaly detection for manufacturing inspection. The research followed a progressive strategy, moving from structural defect detection in standard industrial scenarios to more complex settings involving logical inconsistencies and realistic deployment constraints.

The results obtained thus far demonstrated that unsupervised frameworks can achieve superior performance without compromising industrial efficiency. Our first

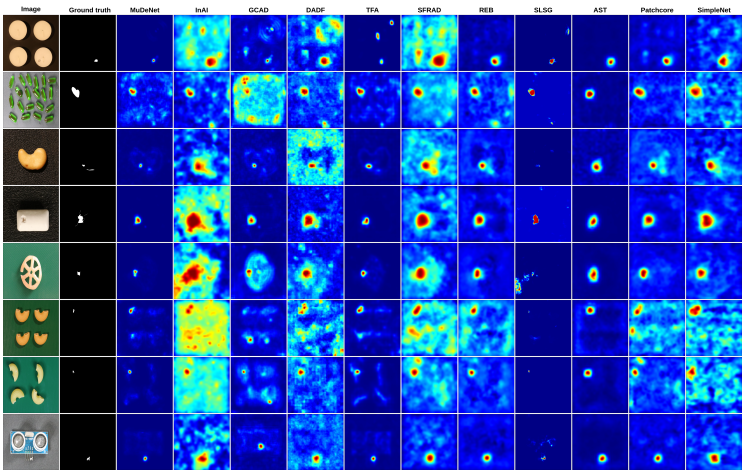



**Fig. 4.** Qualitative results. From left to right: input image, ground-truth mask, and the predicted anomaly maps, where higher intensity indicates a higher probability of a defect. Our proposed frameworks, InAI and MuDeNet, are displayed in the first two prediction columns for direct comparison against state-of-the-art competitors.

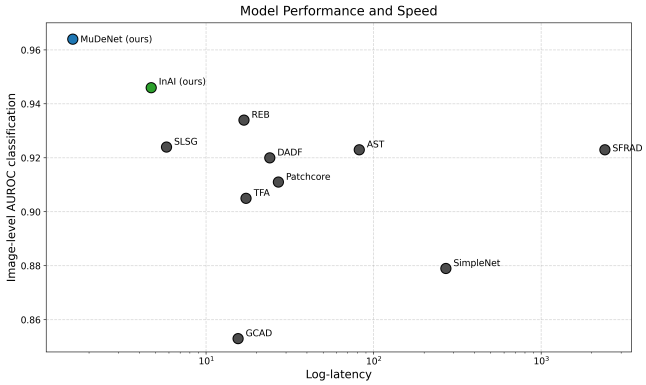



**Fig. 5.** Trade-off between anomaly detection performance and inference latency.

contribution, InAI, proved that structural anomalies can be effectively modeled as deviations in a Gaussian latent space using normalizing flows. Our subsequent work, MuDeNet, extended this capability through a dual-branch architecture, successfully incorporating the contextual reasoning required to identify logical anomalies within millisecond latencies.

Crucially, by conducting this research in collaboration with TKH Group, we validated these methods against rigorous industrial requirements, including near-zero false-positive rates and standardized image acquisition, ensuring the proposed pipelines remain robust under the constraints of high-speed production lines.

Future work will extend this research toward more realistic industrial inspection scenarios, particularly those involving three-dimensional data. Many manufacturing environments rely not only on 2D images but also on point clouds, which provide geometric information that can reveal depth variations, surface deformations, and subtle defects that may not be visible in RGB images alone. In this context, MVTec 3D-AD [7] represents a relevant benchmark, as it combines 3D anomaly scenarios with their corresponding 2D images, closely reflecting real industrial inspection settings where visual and geometric information are analyzed together. Extending anomaly detection to this domain could represent a significant contribution to both academia and industry.

Alternatively, other significant existing limitations could be addressed. Many influential works rely on large convolutional neural networks trained at scale on massive datasets of natural images, as they possess strong generalization capabilities and are powerful tools for image processing. However, they lack the anomaly context to extract optimal features for anomaly detection. Pre-trained transformers with DINOv2 offer robust representation capabilities, but they are computationally expensive, limiting their applicability in time-critical scenarios. Therefore, exploring alternatives to inject anomaly context to pre-trained models for image analysis would have a broad impact on both current and future methods.

## Acknowledgment

This research has been supported by Grant PID2023-146037OB-C21 funded by MICIU/AEI/10.13039/501100011033, and TKH Artificial Intelligence.

## References

1. Ahmad, I., Alqurashi, F.: Early cancer detection using deep learning and medical imaging: A survey. Critical Reviews in Oncology/Hematology p. 104528 (2024)
2. Akcay, S., Atapour-Abarghouei, A., Breckon, T.P.: GANomaly: Semi-supervised anomaly detection via adversarial training. In: Asian Conference on Computer Vision. pp. 622–637 (2019). https://doi.org/10.1007/978-3-030-20893-6_39
3. Anoopa, S., Salim, A.: Survey on anomaly detection in surveillance videos. Materials Today: Proceedings **58**, 162–167 (2022)
4. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: The MVTec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. International Journal of Computer Vision **129**(4), 1038–1059 (2021). https://doi.org/10.1007/s11263-020-01400-4
5. Bergmann, P., et al.: Improving unsupervised defect segmentation by applying structural similarity to autoencoders. In: Proceedings of the 14th International

Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications. pp. 372–380 (2019). https://doi.org/10.5220/0007364503720380
6. Bergmann, P., et al.: Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4183–4192 (2020)
7. Bergmann, P., et al.: The MVTEC 3D-AD dataset for unsupervised 3D anomaly detection and localization. arXiv preprint arXiv:2112.09045 (2021)
8. Bergmann, P., et al.: Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization. International Journal of Computer Vision **130**(4), 947–969 (2022). https://doi.org/10.1007/s11263-022-01578-9
9. Cai, X., Xiao, R., Zeng, Z., Gong, P., Ni, Y.: ITran: A novel transformer-based approach for industrial anomaly detection and localization. Engineering Applications of Artificial Intelligence **125**, 106677 (2023). https://doi.org/10.1016/j.engappai.2023.106677
10. Campos-Romero, M., Carranza-García, M., Riquelme, J.C.: Advancing unsupervised anomaly detection with normalizing flow and multi-scale ensemble learning. Engineering Applications of Artificial Intelligence **137**, 109088 (2024)
11. Campos-Romero, M., Carranza-García, M., Sips, R.J., Riquelme, J.C.: MuDeNet: A multi-patch descriptor network for anomaly modeling. Information Fusion p. 104214 (2026)
12. Defard, T., Setkov, A., Loesch, A., Audigier, R.: PaDiM: A patch distribution modeling framework for anomaly detection and localization. In: International Conference on Pattern Recognition. pp. 475–489 (2021). https://doi.org/10.1007/978-3-030-68799-1_35
13. Deng, H., Li, X.: Anomaly detection via reverse distillation from one-class embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9737–9746 (2022)
14. Fang, Y., et al.: Anomaly diagnosis of connected autonomous vehicles: A survey. Information Fusion **105**, 102223 (2024). https://doi.org/https://doi.org/10.1016/j.inffus.2024.102223, https://www.sciencedirect.com/science/article/pii/S1566253524000010
15. Heckler-Kram, L., et al.: The MVTec AD 2 dataset: Advanced scenarios for unsupervised anomaly detection. International Journal of Computer Vision **134**(4), 175 (2026)
16. Huang, Y., Qiu, C., Yuan, K.: Surface defect saliency of magnetic tile. The Visual Computer **36**(1), 85–96 (2020)
17. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: CutPaste: Self-supervised learning for anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9659–9669 (2021). https://doi.org/10.1109/CVPR46437.2021.00954
18. Liu, Z., Zhou, Y., Xu, Y., Wang, Z.: Simplenet: A simple network for image anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20402–20411 (2023). https://doi.org/10.1109/CVPR52729.2023.01954
19. Luo, W., Yao, H., Yu, W.: Template-based feature aggregation network for industrial anomaly detection. Engineering Applications of Artificial Intelligence **131**, 107810 (2024). https://doi.org/10.1016/j.engappai.2023.107810
20. Lyu, S., Mo, D., keung Wong, W.: Reb: Reducing biases in representation for industrial anomaly detection. Knowledge-Based Systems **290**, 111563 (2024). https://doi.org/https://doi.org/10.1016/j.knosys.2024.111563

21. Roth, K., et al.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14298–14308 (2022). https://doi.org/10.1109/CVPR52688.2022.01392
22. Rudolph, M., Wandt, B., Rosenhahn, B.: Same same but DifferNet: Semi-supervised defect detection with normalizing flows. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision. pp. 1906–1915 (2021). https://doi.org/10.1109/WACV48630.2021.00195
23. Rudolph, M., Wehrbein, T., Rosenhahn, B., Wandt, B.: Asymmetric student-teacher networks for industrial anomaly detection. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2592–2602 (2023). https://doi.org/10.1109/WACV56688.2023.00262
24. Yang, M., Liu, J., Yang, Z., Wu, Z.: Slsg: Industrial image anomaly detection with improved feature embeddings and one-class classification. Pattern Recognition p. 110862 (2024). https://doi.org/https://doi.org/10.1016/j.patcog.2024.110862
25. Yao, H., et al.: Dual-attention transformer and discriminative flow for industrial visual anomaly detection. IEEE Transactions on Automation Science and Engineering (2023). https://doi.org/10.1109/TASE.2023.3322156
26. Zhang, F., et al.: Low-shot unsupervised visual anomaly detection via sparse feature representation. IEEE Transactions on Neural Networks and Learning Systems (2024). https://doi.org/10.1109/TNNLS.2024.3420818
27. Zou, Y., et al.: Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In: European Conference on Computer Vision. pp. 392–408. Springer (2022)