arXiv:2603.15300v1 [cs.CV] 16 Mar 2026

GATE-AD: Graph Attention Network Encoding For Few-Shot Industrial Visual Anomaly Detection

Aggelos Psiris[1], Yannis Panagakis[2,4], Maria Vakalopoulou[3,4], and Georgios Th. Papadopoulos[1,4]

[1] Harokopio University of Athens, Athens, Greece
[2] National and Kapodistrian University of Athens, Athens, Greece
[3] CentraleSupelec, University Paris Saclay, Paris, France
[4] Archimedes, Athena Research Center, Athens, Greece
{aggelospsiris,g.th.papadopoulos}@hua.gr
yannisp@di.uoa.gr
maria.vakalopoulou@centralesupelec.fr

Abstract. Few-Shot Industrial Visual Anomaly Detection (FS-IVAD) comprises a critical task in modern manufacturing settings, where automated product inspection systems need to identify rare defects using only a handful of normal/defect-free training samples. In this context, the current study introduces a novel reconstruction-based approach termed GATE-AD. In particular, the proposed framework relies on the employment of a masked, representation-aligned Graph Attention Network (GAT) encoding scheme to learn robust appearance patterns of normal samples. By leveraging dense, patch-level, visual feature tokens as graph nodes, the model employs stacked self-attentional layers to adaptively encode complex, irregular, non-Euclidean, local relations. The graph is enhanced with a representation alignment component grounded on a learnable, latent space, where high reconstruction residual areas (i.e., defects) are assessed using a Scaled Cosine Error (SCE) objective function. Extensive comparative evaluation on the MVTec AD, VisA, and MPDD industrial defect detection benchmarks demonstrates that GATE-AD achieves state-of-the-art performance across the 1- to 8-shot settings, combining the highest detection accuracy (increase up to 1.8% in image AUROC in the 8-shot case in MPDD) with the lowest per-image inference latency (at least 25.05% faster), compared to the best-performing literature methods. In order to facilitate reproducibility and further research, the source code of GATE-AD is available at https://github.com/gthpapadopoulos/GATE-AD.

Keywords: Few-shot industrial visual anomaly detection · graph attention networks · representation alignment

1 Introduction

Industrial Visual Anomaly Detection (IVAD) comprises a cornerstone of modern quality assurance pipelines, enabling the automated identification of de-

fects (such as scratches, dents, and structural failures) in manufactured products [18,19]. In practice, production lines change frequently and defective samples are rare by nature, making the collection of large labeled datasets prohibitively expensive. This motivates the **Few-Shot Industrial Visual Anomaly Detection (FS-IVAD)** setting, where a model must detect and localize anomalies using only a small number, typically 1 to 8, of defect-free reference images per product category, with no abnormal samples available during training [5, 12].

The core challenge in FS-IVAD is learning a compact, robust representation of *normality* from very limited data, such that deviations from it can be reliably identified during inference/test time. This is challenging due to the following main compounding factors: a) **Data scarcity**, where normality must be captured from as few as a single image, b) **Subtle variance**, where genuine structural failures must be distinguished from natural appearance variation, and c) **Domain shift**, where the learned appearance patterns of normality must generalize across unseen product categories [6, 7, 12].

Existing approaches fall broadly into three main categories: a) *Memory bank methods* [4, 6, 23, 30] retrieve anomaly scores by comparing test patches against a stored bank of extracted normal features, but with inference cost and memory that scale linearly with the number of reference shots. b) *Vision-language methods* [3, 12, 15, 20] leverage vision-text representations (e.g., CLIP-based ones) to align visual patches with textual descriptions of normality, but suffer from a fundamental granularity gap: Visual-Language Models (VLMs) are trained on image-level supervision, while industrial defects are fine-grained and local. c) *Reconstruction-based methods* [7, 16, 26] train a model to reconstruct normal patterns and score anomalies by residual error, but risk the identity mapping problem, where the model reconstructs anomalous regions equally well, and tend to overfit to the limited available normal samples.

The current work fundamentally argues that the key limitation shared across all FS-IVAD methodological categories is the failure to explicitly model **local relational structure** among image patches. Normal appearance is not merely a distribution over individual patch features; it is a pattern of contextual consistency between spatially adjacent regions. A scratch, crack, or contamination disrupts this local consistency in ways that patch-level feature matching or global reconstruction cannot always capture.

In order to address the above key challenge, a novel approach to reconstruction-based FS-IVAD is proposed in this work, termed **GATE-AD**. Specifically, the GATE-AD framework is grounded on the employment of a masked, representation-aligned Graph Attention Network (GAT) encoding scheme to learn robust normality appearance patterns from as few as a single normal sample per supported object/product category. By breaking down images into patches and leveraging dense, local, visual feature tokens (using powerful frozen pretrained encoders, namely DINOv2/DINOv3) as graph nodes, the model employs stacked self-attentional layers to adaptively encode complex, local-level relations. During inference time, high reconstruction residual areas (i.e., defects), compared to the

original/input visual features, are assessed using a Scaled Cosine Error (SCE) objective function. More precisely, the key novel contributions of GATE-AD are:

- GATE-AD employs a GAT encoder architecture, which is composed of stacked masked self-attentional layers. In contrast to other commonly used encoder types employed by similar reconstruction-based FS-IVAD techniques [7, 16, 31], the GAT's self-attentional layers allow the graph nodes to adaptively attend over their neighborhood's features, enabling the encoding of complex, irregular, non-Euclidean, patch-level relations with dynamic importance.
- In order to robustly handle the intrinsic feature over-smoothing problem, the graph is enhanced with a representation alignment component that is based on a learnable, latent space, where high reconstruction residual areas (i.e., defects) are assessed using a Scaled Cosine Error (SCE) objective function.
- GATE-AD is highly computationally efficient, since GATs inherently do not involve costly matrix operations.
- Extensive comparative evaluation on the MVTec AD, VisA, and MPDD industrial defect detection benchmarks demonstrates that GATE-AD achieves state-of-the-art performance across the 1- to 8-shot settings, combining the highest detection accuracy (increase up to 1.8% in image AUROC in the 8-shot case in MPDD) with the lowest per-image inference latency (at least 25.05% faster), compared to the best-performing literature methods.

In order to facilitate reproducibility and further research, the source code of GATE-AD is available at https://github.com/gthpapadopoulos/GATE-AD.

The remainder of the paper is organized as follows: Section 2 outlines the relevant FS-IVAD literature. Section 3 describes the proposed GATE-AD approach. Section 4 provides detailed experimental evaluation results, as well as comparisons with relevant literature methods. Finally, Section 5 concludes the paper and discusses future research directions.

2 Related Work

Visual feature embeddings and memory banks. Feature embedding-based methods rely on mapping an input image to a high-dimensional feature space (using powerful encoder architectures), where anomalous samples correspond to low-density areas or lie far from normal ones. In particular, PatchCore [23] employs a maximally representative memory bank of nominal patch-level features, coupled with a greedy coreset subsampling methodology. VisionAD [28] increases the support data diversity through simple data augmentation. In a similar fashion, AnomalyDINO [6] employs high-quality feature representations extracted by DINOv2. DFM [29] realizes a joint optimization of the feature extractor and the matching head. Additional approaches employ different graph-based representations [8, 25, 30]. Despite widespread adoption, common limitations of this category of methods are their inability to model intra-class variance, the linear scaling of the computational cost with respect to the memory bank size, and the increased sensitivity to object spatial alignment.

Vision-language model reasoning and adaptation. Vision-Language Models (VLMs) support linguistic descriptions that can often encapsulate high-level, information-rich priors for modeling abnormalities, which visual-only approaches need to learn from scratch and limited training data. In this context, WinCLIP+ [12] relies on the combination of state words and prompt templates with window/patch/image-level features. FADE [17] adapts the CLIP model to extract multi-scale, image patch embeddings that are better aligned with language. PromptAD [15] employs a dual-branch vision-language decoding network for both normality and abnormality information modeling. In a similar way, Lv et al. [20] learn a class-shared prompt generator. Moreover, InCTRL [33] learns an in-context residual learning model. AnomalyR1 [2] autonomously processes image and domain knowledge inputs. Despite 'open-vocabulary' detection capabilities, VLM-based approaches suffer from the inherent granularity gap (i.e., VLMs are typically trained using image-level information, while industrial abnormalities refer to fine-grained local defects), lack sufficient domain-specific knowledge, exhibit high sensitivity to prompt engineering, and are susceptible to noisy support sets.

Reconstruction-based frameworks and prototype refinement. Reconstruction-based approaches assume that a model trained to reconstruct normal patterns will result into high residual errors in case of abnormalities. In this context, FastRecon [7] employs a regression scheme with distribution regularization. Similarly, FastRef [26] relies on characteristic transfer from query features to prototypes and anomaly suppression through prototype alignment. Moreover, One-to-Normal [16] performs a personalized one-to-normal transformation of query images. FoundAD [32] learns a non-linear projection operator on the natural image manifold. Despite intuitiveness, key limitations include the risk of reconstructing both normal and abnormal samples equally well, the loss of reconstruction detail being (often) larger than the defects themselves, tendency to overfit to the available normal samples, and reduced ability to handle high intra-class variance.

3 Proposed approach

3.1 Overview

The proposed GATE-AD framework relies on the employment of a masked, representation-aligned GAT encoding scheme. As a reconstruction-based FS-IVAD approach [7, 26], its fundamental consideration comprises the learning of robust normality appearance patterns and, during inference time, high reconstruction residual areas (compared to the original/input visual features) are assessed as abnormalities/defects. The general architecture of GATE-AD is illustrated in Fig. 1, while its individual components are detailed in the followings.

3.2 Image patchification and feature extraction

Given an input image $I_{mg} \in \mathbb{R}^{H \times W \times 3}$, it is initially resized to a resolution equal to β pixel (smaller image dimension). Subsequently, a dense local feature

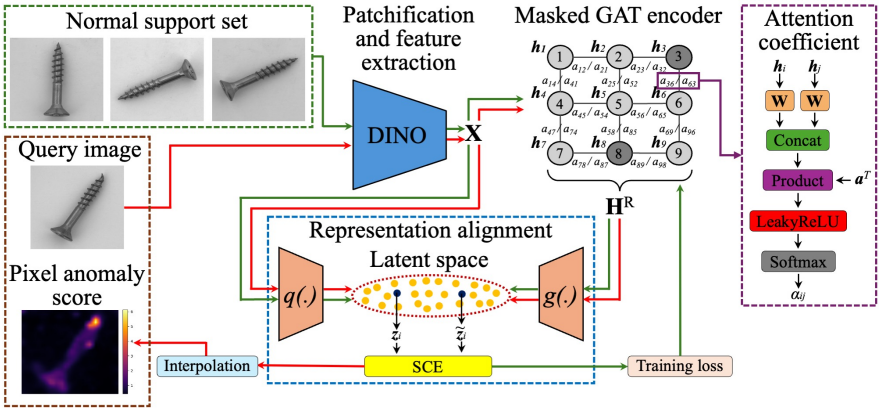

Fig. 1: General architecture of the proposed masked, representation-aligned GAT encoding scheme for FS-IVAD.

token $\mathbf{x}_i \in \mathbb{R}^D$, $i \in [1, N]$, is extracted for each 16x16/14x14 pixel image patch (standard patch size for DINOv3/DINOv2), using a frozen Visual Transformer (ViT) backbone (DINOv3 [24]/DINOv2 [22]). In order to extract a more robust and discriminative $\mathbf{x}_i$ representation, the common practice of estimating the patch token average from the last ζ transformer layers is adopted.

3.3 Masked graph attention network encoder

Encoder type selection. As described above, the fundamental building block of any reconstruction-based FS-IVAD approach comprises the use of an encoder module for learning the normality apperance pattern distribution. In the current work, a GAT encoder architecture is employed for that purpose, which is also suitable for processing the image patch-level representation $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ (Section 3.2). GATs [27] constitute a particular type of Graph Neural Network (GNN) architectures, which leverage a set of stacked masked self-attention layers and exhibit the following key advantageous characteristics: a) Anisotropic aggregation (selective importance), which allows each graph node to assign different weights to its neighbors, b) Increased inductive learning (generalization) capabilities, where the GAT attention mechanism (that is shared across all edges and depends only on node features) enables the model to efficiently handle previously unseen data without retraining, c) Improved robustness to noisy data, where the employed attention mechanism again allows to effectively minimize the influence of irrelevant neighbors, and d) Increased computational efficiency, where GATs are inherently computationally efficient, since they do not require costly matrix operations. In this respect, GATs enable the encoding of complex, irregular, non-Euclidean, local relations with dynamic importance.

GATs compare favorably with other commonly used encoder types employed by similar reconstruction-based FS-IVAD techniques [7,16,31]: a) Convolutional

Neural Networks (CNNs), which are restricted to Euclidean grids and model poor global context, b) ViTs, which may ignore fine-grained local relations and require large amounts of training data, c) Standard Autoencoders (AE), which tend to miss relational context between features, d) Generative Adversarial Networks (GANs), which exhibit training instability and are prone to mode collapse, and e) Diffusion Models (DMs), which are computational heavy and need massive training datasets.

Encoder network architecture. GAT networks are constructed through the successive stacking of their fundamental building block layer, the so called graph attentional layer [27]. Overall, the designed GAT encoder receives as input the extracted patch-level representation $\mathbf{X}$, where each feature token $\mathbf{x}_i \in \mathbb{R}^D$ is associated with (the location-aware) node $i \in [1, N]$ in the graph. Subsequently, the graph estimates updated feature representations (embeddings) $\mathbf{h}_i^r$ ($r \in [1, R]$) for each node i at layer r, by iteratively aggregating information from its local neighborhood (adjacent nodes). The latter is performed through the repetitive execution of the following steps: a) Neighborhood aggregation (message passing), where each node i collects information from the feature vectors of all its immediate neighbors $\mathcal{N}(i)$; in the current work, an 8-neighbor node connectivity is considered after experimentation, b) Feature update (combination), where once the neighborhood message is collected for node i, the model combines it with the node's own/current representation to produce a new/updated one, and c) (Layer-wise) iterative propagation, where the previous two steps are performed R times in a repetitive/layer-wise manner; at iteration r, each node vector $\mathbf{h}_i^r$ incorporates a feature summary of its r-hop neighborhood [21] (that defines the extent of the 'message passing' reach).

Graph attentional layer. The input to an individual attention layer is a set of (location-aware) node features $\mathbf{H} = \{\mathbf{h}_i\}_{i=1}^N$, $\mathbf{h}_i \in \mathbb{R}^F$, where F is the cardinality of the node feature vector. In case of the first layer in the network, $\mathbf{h}_i \equiv \mathbf{x}_i$ holds. The layer produces as output a new set of node features (of potentially different cardinality F') $\mathbf{H}' = \{\mathbf{h}_i'\}_{i=1}^N$, $\mathbf{h}_i' \in \mathbb{R}^{F'}$. In order to dynamically encode the relations between neighboring node features, an attention-based message passing mechanism is employed. In particular, a shared learnable linear transformation in the form of a weight matrix $\mathbf{W} \in \mathbb{R}^{F' \times F}$ is applied to every node, while a corresponding shared self-attentional mechanism $\alpha : \mathbb{R}^{F'} \times \mathbb{R}^{F'} \rightarrow \mathbb{R}$ is applied to neighboring nodes, resulting in the computation of the following attention coefficients:

$$e_{ij} = \alpha(\mathbf{W}\mathbf{h}_i, \mathbf{W}\mathbf{h}_j), \; j \in \mathcal{N}(i), \tag{1}$$

which indicate the importance of node j's features to node i. Among the various possible options for the attention mechanism, the common practice of implementing α as a single-layer feedforward neural network, parametrized by a weight vector $\mathbf{a} \in \mathbb{R}^{2F'}$, and applying the LeakyReLU nonlinearity is adopted in

the current work, as follows:

$$\alpha_{ij} = \frac{\exp\big(\text{LeakyReLU}\big(\mathbf{a}^T[\mathbf{W}\mathbf{h}_i \| \ \mathbf{W}\mathbf{h}_j]\big)\big)}{\sum_{m\in\mathcal{N}(i)} \exp\big(\text{LeakyReLU}\big(\mathbf{a}^T[\mathbf{W}\mathbf{h}_i \| \ \mathbf{W}\mathbf{h}_m]\big)\big)}, \tag{2}$$

where $\|$ denotes vector concatenation. Then, the normalized attention coefficients α_{ij} are used to compute a linear combination of the features corresponding to them, followed by a non-linearity operator $\sigma(.)$, so as to estimate the final output features $\mathbf{h}'_i$ for every node i, according to the following equation:

$$\mathbf{h}'_i = \sigma\big(\sum_{j\in\mathcal{N}(i)} \alpha_{ij}\mathbf{W}\mathbf{h}_j\big), \tag{3}$$

where for $\sigma(.)$ the typical choice of the ELU activation function is adopted.

Representation alignment. GATs (as a particular type of GNNs) suffer from the so called over-smoothing problem, where adding more network layers results into node representations becoming increasingly similar. In order to mitigate this issue, among various methodologies proposed in the literature (e.g., identity-aware GATs, structural identity encoding, etc.), a graph Representation Alignment (RA) [14] component is integrated in this work. In particular, the fundamental goal is to ensure that the produced representation of the last GAT layer $\mathbf{H}^R \in \mathbb{R}^F$ (where F is the cardinality of $\mathbf{H}^R$) is consistent/comparable with the input feature tokens' one $\mathbf{X}$. In other words, the designed RA scheme aims to force the GAT model to map similar structural or semantic properties in corresponding areas of the $\mathbf{H}^R$ and $\mathbf{X}$ spaces. In order to efficiently implement the latter (and also to handle the different cardinality values of $\mathbf{H}^R$ and $\mathbf{X}$), a common latent space $\mathbf{Z} \in \mathbb{R}^f$ ($f \leq F$) is defined, where both $\mathbf{H}^R$ and $\mathbf{X}$ are projected through a pair of learnable transformations. Following common practice, two lightweight projection heads, namely $g(.)$ (materialized as a Multi-Layer Perceptron (MLP) network (Linear→ReLu→Linear)) and $q(.)$ (single linear layer), are introduced to perform the mapping, as follows:

$$\mathbf{z}_i = q(\mathbf{x}_i), \ \tilde{\mathbf{z}}_i = g(\mathbf{h}_i^R), \quad \mathbf{z}_i, \tilde{\mathbf{z}}_i \in \mathbb{R}^f \ , \tag{4}$$

where the non-linearity operator is incorporated in $g(.)$ for handling the irregular, non-Euclidean relations in the developed graph. Among the various available options (e.g., cross-entropy, Mean Square Error (MSE), etc.) the Scaled Cosine Error (SCE) [10] objective function is adopted for assessing feature dissimilarity in $\mathbf{Z}$ and used for training. The latter aims to efficiently handle sensitivity (i.e., impact of feature dimensionality and vector norms) and selectivity (i.e., down-weighting easy samples' contribution) aspects, while it is formalized as follows:

$$\mathcal{L}_{SCE} = \frac{1}{N}\sum_i \left(1 - \frac{\mathbf{z}_i^T\tilde{\mathbf{z}}_i}{\|\mathbf{z}_i^T\| \, \|\tilde{\mathbf{z}}_i\|}\right)^\gamma, \ \gamma \geq 1, \tag{5}$$

which is averaged over all N nodes, while scaling factor γ comprises an adjustable hyper-parameter.

Feature masking. Similarly to any type of GNNs, (pure) GATs exhibit the tendency to memorize a specific global structure in the training data. The latter results into ignoring the actual links between data points. In order to overcome this limitation, graph Self-Supervised Learning (SSL) is employed for enforcing the model to learn how to aggregate local information. In particular, the node feature masking technique of GraphMAE [10] is adopted, where a percentage μ of the graph nodes is selected, based on a uniform random sampling strategy, and their original features $\mathbf{h}_i^r$ are replaced with a learnable mask token vector. In this way, each node relies on its neighboring ones to enhance/recover its features.

Graph regularization. In order to mitigate the inherent susceptibility of GATs to data over-fitting occurrences (that are also more likely to be observed in the low-data regime of FS-IVAD), graph regularization techniques are also employed. Among the various available options (e.g., residual connections, batch normalization, etc.), an attention dropout operator is employed, which is applied to the attention coefficients α_{ij} [27]. A percentage δ of the graph attention coefficients is selected, based on a uniform random sampling strategy.

3.4 Inference and anomaly scoring

Patch anomaly score. During inference time, a query/unseen image I_{mg} is provided as input to the backbone network and a dense local feature token $\mathbf{x}_i$ is extracted for each image patch. The extracted features $\mathbf{X}$ are then introduced to the learned GAT graph, which is trained using normal-only images from a support set. Subsequently, through its inference/forward-pass, the graph model outputs the feature representation of its last layer $\mathbf{H}^R$ and for each node i a patch anomaly score (reconstruction inconsistency measure) s_i is estimated, based on the defined SCE objective function (Eq. (5)).

Image anomaly score. For estimating an image-level anomaly score, the top-k s_i values (i.e., most anomalous image patches) are estimated, by selecting a fixed ratio ϵ of the highest s_i scores, and a simple averaging operator is applied.

Pixel anomaly score. In order to obtain pixel-level anomaly scores, the location-aware patch scores $\{s_i\}$ are positioned on a corresponding rectangular grid, forming a low-resolution/coarse anomaly map. Subsequently, light Gaussian smoothing is applied to the generated map for noise reduction (removal of spurious high-frequency fluctuations) [6], followed by bilinear interpolation up-sampling to the original image resolution.

4 Experimental results

4.1 Setup

Datasets. For the experimental evaluation, three widely used public benchmarks for high-resolution FS-IVAD are employed, namely MVTec AD [1], VisA [34],

and MPDD [13]. MVTec AD ($3,629$ defect-free training and $1,725$ test images) includes a broad spectrum of defect types, across 5 texture classes and 10 object categories. VisA comprises $10,821$ images ($9,621$ normal and $1,200$ anomalous), spanning 12 object categories in 3 domains. MPDD (888 training and 458 test images) focuses on real-world manufacturing conditions, containing a total of 6 object classes.

Experimental protocol and performance metrics. The proposed approach, as well as all relevant literature methods, are evaluated under the k/few-shot experimental setting, where k normal samples are provided per supported class for training purposes and, subsequently, evaluation is performed in the respective test set. Both image classification and segmentation performance are evaluated in the current study. Regarding the former (image classification), metrics 'Area Under the Receiver Operating Characteristic' (AUROC), which relies on plotting the true positive rate (sensitivity) against the false positive one (1-specificity), and mean Average Precision (mAP), which summarizes the precision-recall curve, are used. Concerning the latter (image segmentation), apart from AUROC, the Per-Region Overlap (PRO) metric is also employed, which focuses on the individual connected components of a given defect. All experiments are repeated three times (using fixed seed numbers), while the mean and standard deviation values are reported for each metric (using the $mean \pm_{std}$ notation).

Default GATE-AD configuration. The default GATE-AD implementation that is used for computing the detection results reported in the remaining of this section considers input image resolution $\beta = 512$ ($\beta = 488$) for DINOv3 (DINOv2), while the DINOv2 ViT-L/14, DINOv3 ViT-B/16 and DINOv3 ViT-L/16 backbone is used for MVTec AD, VisA, and MPDD, respectively. Additionally, the patch token average is computed from the last $\zeta = 8$ backbone transformer layers. The GAT encoder comprises $R = 3$ attentional layers, all of them with cardinality F equal to 256. The cardinality of latent representation Z is set $f = 256$. Moreover, the scaling factor in the SCE objective function is set $\gamma = 2$, the node masking ratio is defined $\mu = 0.2$, and the dropout rate is set $\delta = 0.3$. A ratio $\epsilon = 0.025$ in MVTec AD and $\epsilon = 0.01$ in VisA is used for defining the top-k most anomalous image patches. For MPDD, image-level scoring is performed via max pooling (instead of averaging), according to the relevant experimental protocol [7,30].

Hyper-parameter tuning. Regarding the various hyper-parameters of the proposed method, a thorough ablation study is performed on the MVTec AD and VisA benchmarks (the detailed results are provided in Appendix A). The main observations are the following: a) Input image resolution: β equal to 512 and 768 exhibit comparable performance (for a DINOv3 backbone); 512 selected for reduced computational cost. b) Backbone: None of the various distillation sizes of DINOv3/DINOv2 evaluated is universally optimal. c) Number of last transformer layers: $\zeta = 8$ leads to relatively increased performance, compared to other lower/higher values. d) Number of attentional layers: $R = 2, 3$ results into similar performance; $R = 3$ being slightly better. e) Cardinality of GAT layers: $F = 256$

is sufficient for reaching increased recognition performance, while maintaining low inference latency. f) Cardinality of latent representation: Setting $f = 256$ or higher leads to negligible variance in I-AUROC. $f = 256$ selected for achieving a low computational cost. g) Objective function: SCE is shown clearly advantageous over MSE and cosine error. $\gamma = 2$ in SCE leads to the most consistent results. h) Masking ratio: $\mu \in [0.1, 0.6]$ results into negligible performance variations; $\mu = 0.2$ leads to the best results. i) Dropout rate: $\delta \in [0.1, 0.4]$ results into the highest performance; $\delta = 0.3$ being the optimal value. j) Anomalous patches ratio: $\epsilon = 0.025$ in MVTec AD and $\epsilon = 0.01$ in VisA leads to the best overall performance. Significantly greater/lower values of ϵ in MVTec AD and greater ones in VisA lead to small performance decrease.

Implementation details. Regarding the implementation details of the proposed method, input images were resized to a resolution of $\beta = 512$ ($\beta = 488$) for DINOv3 (DINOv2). Image augmentation/rotation was only applied to images containing defect categories that exhibit pose variance, similarly to AnomalyDINO [6]. The proposed method was implemented using the PyTorch Geometric library for developing GNN networks. Training was performed using the Adam optimizer, a learning rate of 3×10^{-4} for up to $2,000$ epochs, and early stopping. All experiments were conducted on a single NVIDIA GeForce RTX 3080 Ti GPU with 12 GB VRAM.

4.2 Comparative evaluation

Baseline methods and reported performance. The proposed method is comparatively evaluated against key baseline literature works, covering all FS-IVAD methodological categories (Section 2), namely: a) Memory bank methods: SPADE [4], PatchCore [23], GraphCore [30], and AnomalyDINO [6], b) Vision-language methods: WinCLIP+ [12], PromptAD [15], IIPAD [20], and APRIL-GAN [3], and c) Reconstruction-based methods: FastRecon [7], FoundAD [32], and RegAD [11]. Table 1 illustrates the performance of the above approaches under the 1/2/4/8-shot FS-IVAD setting on the MVTec AD, VisA, and MPDD benchmarks. For all baseline methods, results are copied from the original publication works, except for AnomalyDINO in MPDD (where its publicly available implementation was used for experimentation). It needs to be highlighted that the normal sample selection process of GATE-AD is identical to the one of AnomalyDINO in all datasets (standardized sample selection process (Seed: 0)), while for the remaining methods such information is not available. Moreover, indicative 1-shot defect detection results from the application of GATE-AD and representative baseline methods (covering all literature methodological categories) to the MVTec AD, VisA, and MPDD benchmarks are depicted in Fig. 2.

Discussion. From the results presented in Table 1, the following key observations can be made: a) GATE-AD consistently achieves **state-of-the-art performance** in the wide majority of cases for both image classification and pixel segmentation tasks in all benchmarks (leading to an increase up to 1.8% in image

Table 1: FS-IVAD detection results on the MVTec AD, VisA, and MPDD benchmarks. Methods with an asterisk (*) use the standardized sample selection process (Seed: 0).

Method	MVTec AD				VisA				MPDD			
	Image		Pixel		Image		Pixel		Image		Pixel	
	AUROC	mAP	AUROC	PRO	AUROC	mAP	AUROC	PRO	AUROC	mAP	AUROC	PRO
1-shot												
SPADE [4] (arXiv'20)	82.9 ±2.6	91.7 ±1.2	92.0 ±0.3	85.7 ±0.7	80.7 ±5.0	82.3 ±4.3	96.2 ±0.4	85.7 ±1.1	-	-	-	-
PatchCore [23] (CVPR'22)	83.4 ±3.0	92.2 ±1.5	92.0 ±1.0	79.7 ±2.0	79.9 ±2.9	82.8 ±2.3	95.4 ±0.6	80.5 ±2.5	-	-	-	-
GraphCore [30] (ICLR'23)	89.9	-	95.6	-	-	-	-	-	**84.7**	-	95.2	-
AnomalyDINO* [6] (WACV'25)	96.6 ±0.4	98.2 ±0.2	**96.8 ±0.1**	92.7 ±0.1	87.4 ±1.2	89.0 ±1.0	97.8 ±0.1	92.5 ±0.5	73.9 ±1.5	75.2 ±2.3	95.5 ±0.2	87.4 ±0.6
WinCLIP+ [12] (CVPR'23)	93.1 ±2.0	96.5 ±0.9	95.2 ±0.5	87.1 ±1.2	83.8 ±4.0	85.1 ±4.0	96.4 ±0.4	85.1 ±2.1	-	-	-	-
APRIL-GAN [3] (CVPRW'23)	92.0 ±0.3	95.8 ±0.2	95.1 ±0.1	90.6 ±0.2	91.2 ±0.8	93.3 ±0.8	96.0 ±0.0	90.0 ±0.1	-	-	-	-
PromptAD [15] (WACV'24)	94.6 ±1.7	97.1 ±1.0	95.9 ±0.5	87.9 ±1.0	86.9 ±2.3	88.4 ±2.6	96.7 ±0.4	85.1 ±2.5	-	-	-	-
IIPAD [20] (ICLR'25)	94.2	97.2	96.4	89.8	85.4	87.5	96.9	87.3	-	-	-	-
FoundAD [32] (ICLR'26)	96.1	97.9	**96.8**	92.8	92.6	92.0	**99.7**	**98.0**	-	-	-	-
GATE-AD*	**97.7 ±0.2**	**98.7 ±0.1**	96.3 ±0.1	**93.0 ±0.1**	**93.5 ±0.2**	**94.3 ±0.3**	97.7 ±0.1	93.2 ±0.3	84.2 ±0.9	**86.2 ±0.8**	**97.5 ±0.1**	**93.4 ±0.2**
2-shot												
SPADE [4] (arXiv'20)	81.0 ±2.0	90.6 ±0.8	91.2 ±0.4	83.9 ±0.7	79.5 ±4.0	82.0 ±3.3	95.6 ±0.4	84.1 ±1.6	-	-	-	-
PatchCore [23] (CVPR'22)	86.3 ±3.3	93.8 ±1.7	93.3 ±0.6	82.3 ±1.3	81.6 ±4.0	84.8 ±3.2	96.1 ±0.5	82.6 ±2.3	-	-	-	-
GraphCore [30] (ICLR'23)	91.9	-	96.9	-	-	-	-	-	**85.4**	-	95.4	-
AnomalyDINO* [6] (WACV'25)	96.9 ±0.7	98.2 ±0.5	**97.0 ±0.2**	93.1 ±0.2	89.7 ±1.3	90.7 ±0.8	98.0 ±0.1	93.4 ±0.6	76.7 ±2.5	77.3 ±2.9	95.8 ±0.2	88.9 ±0.9
WinCLIP+ [12] (CVPR'23)	94.4 ±1.3	97.0 ±0.7	96.0 ±0.3	88.4 ±0.9	84.6 ±2.4	85.8 ±2.7	96.8 ±0.3	86.2 ±1.4	-	-	-	-
APRIL-GAN [3] (CVPRW'23)	92.4 ±0.3	96.0 ±0.2	95.5 ±0.0	91.3 ±0.1	92.2 ±0.3	94.2 ±0.3	96.2 ±0.0	90.1 ±0.1	-	-	-	-
PromptAD [15] (WACV'24)	95.7 ±1.5	97.9 ±0.7	96.2 ±0.3	88.5 ±0.8	88.3 ±2.0	**96.1 ±0.6**	97.1 ±0.3	85.8 ±2.1	-	-	-	-
IIPAD [20] (ICLR'25)	95.7	97.9	96.7	90.3	86.7	88.6	97.2	87.9	-	-	-	-
RegAD [11] (ECCV'22)	85.7	-	94.6	-	-	-	-	-	63.4	-	93.2	-
FastRecon [7] (ICCV'23)	91.0	-	95.9	-	-	-	-	-	73.7	-	97.0	-
FoundAD [32] (ICLR'26)	96.9	98.3	**97.0**	93.2	**93.8**	93.3	**99.7**	**98.2**	-	-	-	-
GATE-AD*	**98.0 ±0.2**	**98.8 ±0.2**	96.8 ±0.1	**93.4 ±0.1**	92.8 ±0.3	93.6 ±0.3	97.9 ±0.1	93.5 ±0.3	84.3 ±0.5	**85.5 ±0.9**	**97.5 ±0.1**	**93.4 ±0.1**
4-shot												
SPADE [4] (arXiv'20)	84.8 ±2.5	92.5 ±1.2	92.7 ±0.3	87.0 ±0.5	81.7 ±3.4	83.4 ±2.7	96.6 ±0.3	87.3 ±0.8	-	-	-	-
PatchCore [23] (CVPR'22)	88.8 ±2.6	94.5 ±1.5	94.3 ±0.5	84.3 ±1.6	85.3 ±2.1	87.5 ±2.1	96.8 ±0.3	84.9 ±1.4	-	-	-	-
GraphCore [30] (ICLR'23)	92.9	-	**97.4**	-	-	-	-	-	**85.7**	-	95.7	-
AnomalyDINO* [6] (WACV'25)	97.7 ±0.2	98.7 ±0.1	97.2 ±0.1	93.4 ±0.1	92.6 ±0.9	92.9 ±0.7	98.2 ±0.0	94.1 ±0.1	77.3 ±2.4	79.1 ±3.4	96.3 ±0.3	89.8 ±1.1
WinCLIP+ [12] (CVPR'23)	95.2 ±1.3	97.3 ±0.6	96.2 ±0.3	89.0 ±0.8	87.3 ±1.8	88.8 ±1.8	97.2 ±0.2	87.6 ±0.9	-	-	-	-
APRIL-GAN [3] (CVPRW'23)	92.8 ±0.2	96.3 ±0.1	95.9 ±0.0	91.8 ±0.1	92.6 ±0.4	94.5 ±0.3	96.2 ±0.0	90.2 ±0.1	-	-	-	-
PromptAD [15] (WACV'24)	96.6 ±0.9	98.5 ±0.5	96.5 ±0.2	90.5 ±0.7	89.1 ±1.7	90.8 ±1.3	97.4 ±0.3	86.2 ±1.7	-	-	-	-
IIPAD [20] (ICLR'25)	96.1	98.1	97.0	91.2	88.3	89.6	97.4	88.3	-	-	-	-
RegAD [11] (ECCV'22)	88.2	-	95.8	-	-	-	-	-	68.3	-	93.9	-
FastRecon [7] (ICCV'23)	94.2	-	97.0	-	-	-	-	-	79.9	-	**97.6**	-
FoundAD [32] (ICLR'26)	97.1	98.6	97.2	**93.6**	**94.4**	94.0	**99.7**	**98.4**	-	-	-	-
GATE-AD*	**97.9 ±0.2**	**98.8 ±0.2**	96.9 ±0.1	93.5 ±0.1	**94.4 ±0.2**	**94.6 ±0.3**	98.3 ±0.1	94.5 ±0.2	84.1 ±0.9	**84.9 ±0.8**	**97.6 ±0.1**	**93.7 ±0.1**
8-shot												
GraphCore [30] (ICLR'23)	95.9	-	**97.8**	-	-	-	-	-	86.0	-	95.9	-
AnomalyDINO* [6] (WACV'25)	**98.2 ±0.2**	**99.1 ±0.1**	97.4 ±0.1	**93.8 ±0.1**	93.8 ±0.3	94.3 ±0.4	**98.4 ±0.0**	**94.8 ±0.2**	80.5 ±3.0	82.8 ±3.9	96.8 ±0.2	90.4 ±0.5
RegAD [11] (ECCV'22)	91.2	-	96.7	-	-	-	-	-	71.9	-	95.1	-
FastRecon [7] (ICCV'23)	95.2	-	97.3	-	-	-	-	-	82.5	-	**97.9**	-
GATE-AD*	**98.2 ±0.1**	99.0 ±0.1	97.0 ±0.1	93.7 ±0.1	**95.2 ±0.2**	**95.3 ±0.3**	98.3 ±0.1	94.6 ±0.3	**87.8 ±0.5**	**89.7 ±0.4**	97.8 ±0.0	**94.4 ±0.0**

AUROC in the 8-shot case in MPDD, compared to the best performing baseline method). This demonstrates the significant advantage of the introduced GAT encoder in dynamically modeling complex, irregular, non-Euclidean, local/patch-level relations. b) GATE-AD demonstrates a clear **superiority in very low-shot counts** (i.e., 1/2-shots), while this detection prevalence is maintained as the number of shots increases (e.g., 4/8-shots) with decreased performance gap from other methods though. This is mainly due to the strong normality appearance patterns learned by the GAT encoder, even in the presence of very low training data. c) GATE-AD exhibits **high detection stability** across all shot counts and datasets, compared to other baseline methods whose performance fluctuates significantly. The latter demonstrates the robustness of GATE-AD, regardless of the particular normal support samples selected for training. d) GATE-AD shows **low variance in detection performance** across all settings, where GATE-AD's standard deviation values are significantly lower than other baseline methods. The latter again indicates the increased resilience of GATE-AD with respect to the normal support set formation. e) GATE-AD demonstrates **increased ef-**

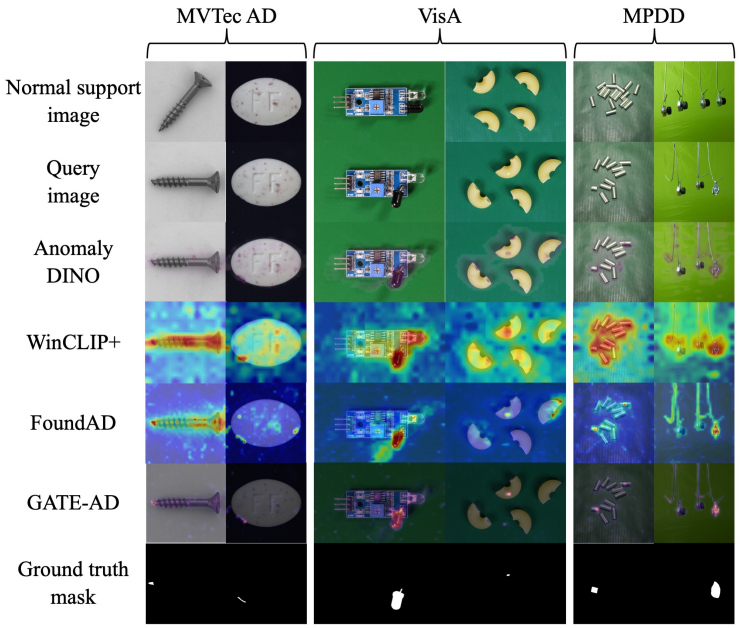

Fig. 2: Indicative 1-shot defect detection results. GATE-AD localizes more accurately and robustly defects, across multiple pose, placement, and abnormality type settings.

ficiency especially in complex, real-world manufacturing settings, as can be seen by the performance achieved in the MPDD dataset. More detailed object-oriented performance and indicative detection results for GATE-AD and other baseline methods are provided in Appendix B.

4.3 Ablation study

Baseline methods and reported performance. In order to assess the importance of the main architectural blocks of GATE-AD, a comprehensive ablation study is performed, involving the evaluation of the following variants: a) 'GATE-AD' with default configuration (as detailed in Section 4.1), b) 'GATE-AD (GCN)', where the GAT layers are replaced with conventional Graph Convolutional Network (GCN) ones, c) 'Transformer Encoder (TE)', where each GAT layer is replaced by a global self-attention mechanism (i.e., a transformer block), d) 'GATE-AD without RA (GATE-AD w/o RA)', where the GAT's RA component (Section 3.3) is discarded and the SCE objective (Eq. (5)) is directly estimated between the GAT's last layer feature representation $\mathbf{H}^R$ and the input feature tokens $\mathbf{X}$, e) 'GATE-AD Autoencoder (GATE-AD AE)', where the GAT's RA is again discarded and an MLP graph decoder is integrated for formulating a graph AE; the SCE objective is now estimated between the AE output representation and the input feature token one, f) 'Masked Autoencoder (MAE)', which comprises

Table 2: FS-IVAD ablation study on the MVTec AD, VisA, and MPDD benchmarks.

	MVTec AD		VisA		MPDD	
Shot–Method	I-AUROC	P-AUROC	I-AUROC	P-AUROC	I-AUROC	P-AUROC
1-shot						
GATE-AD	**97.7** ±0.2	**96.3** ±0.1	**93.5** ±0.2	**97.7** ±0.1	**84.3** ±1.1	**97.4** ±0.2
GATE-AD (GCN)	97.3 ±0.2	95.9 ±0.2	92.7 ±0.4	97.4 ±0.1	82.1 ±0.9	95.7 ±0.1
TE	97.1 ±0.2	95.7 ±0.1	93.4 ±0.2	97.5 ±0.1	82.4 ±0.9	96.3 ±0.1
GATE-AD w/o RA	94.1 ±0.5	95.1 ±0.1	80.5 ±0.3	96.0 ±0.2	55.7 ±1.1	94.4 ±0.4
GATE-AD AE	93.2 ±1.0	95.3 ±0.1	80.8 ±0.5	96.2 ±0.1	59.2 ±2.5	94.2 ±0.3
MAE	66.6 ±5.0	62.0 ±10.1	55.0 ±1.6	71.2 ±0.6	45.0 ±2.7	87.2 ±6.2
2-shot						
GATE-AD	**98.0** ±0.2	**96.8** ±0.1	**92.8** ±0.3	**97.9** ±0.1	**84.5** ±0.7	**97.4** ±0.2
GATE-AD (GCN)	97.6 ±0.2	96.4 ±0.2	**92.8** ±0.5	97.6 ±0.1	82.9 ±0.8	96.5 ±0.1
TE	97.1 ±1.0	96.5 ±0.2	92.1 ±1.0	97.5 ±0.2	83.4 ±1.0	97.0 ±0.1
GATE-AD w/o RA	94.8 ±0.2	95.7 ±0.1	82.5 ±1.4	96.4 ±0.2	58.2 ±1.0	94.6 ±0.5
GATE-AD AE	94.2 ±0.3	95.1 ±0.1	82.3 ±1.2	96.5 ±0.2	57.8 ±1.1	94.4 ±0.3
MAE	70.1 ±1.2	65.2 ±6.2	64.2 ±0.1	72.2 ±0.9	44.5 ±4.8	84.9 ±3.4
4-shot						
GATE-AD	**97.9** ±0.2	**96.9** ±0.1	**94.4** ±0.2	**98.3** ±0.1	**84.5** ±0.7	**97.5** ±0.2
GATE-AD (GCN)	97.7 ±0.2	96.7 ±0.1	94.0 ±0.2	98.1 ±0.1	83.3 ±0.8	97.3 ±0.1
TE	97.4 ±0.6	96.8 ±0.2	94.1 ±0.6	98.2 ±0.1	83.6 ±0.4	**97.5** ±0.1
GATE-AD w/o RA	95.6 ±0.4	96.2 ±0.2	87.9 ±1.3	96.8 ±0.3	58.6 ±3.2	95.1 ±0.4
GATE-AD AE	94.9 ±0.4	95.7 ±0.1	87.8 ±0.3	97.0 ±0.1	59.8 ±3.4	95.0 ±0.2
MAE	70.0 ±1.2	60.0 ±6.2	64.2 ±0.1	72.4 ±0.2	44.5 ±4.8	84.9 ±3.4
8-shot						
GATE-AD	**98.2** ±0.1	**97.0** ±0.1	**95.2** ±0.2	**98.3** ±0.1	**87.3** ±0.8	**97.8** ±0.1
GATE-AD (GCN)	98.1 ±0.2	96.9 ±0.2	95.1 ±0.2	**98.3** ±0.1	87.2 ±0.7	97.7 ±0.1
TE	97.7 ±0.3	96.9 ±0.2	94.5 ±0.5	**98.3** ±0.1	87.1 ±0.4	97.7 ±0.1
GATE-AD w/o RA	95.8 ±0.6	95.5 ±0.2	90.5 ±0.3	97.1 ±0.3	64.5 ±1.3	95.5 ±0.3
GATE-AD AE	95.9 ±0.5	95.6 ±0.1	89.7 ±0.4	97.2 ±0.2	65.2 ±0.9	95.4 ±0.2
MAE	72.2 ±0.1	66.2 ±0.1	62.3 ±0.1	70.8 ±1.2	46.4 ±2.0	77.3 ±10.7

the conventional MAE SSL method [9], aiming at directly reconstructing the original input feature tokens. Table 2 illustrates the obtained experimental results; due to page length limitations, only the Image AUROC (I-AUROC) and Pixel AUROC (P-AUROC) metrics are provided.

Discussion. From the results presented in Table 2, the following key observations can be made (in correspondence with the defined experiments/variants): a) The default GATE-AD configuration **outperforms all other variants** in all cases, demonstrating the efficiency of the developed GAT encoder with RA in the FS-IVAD setting. b) The drop in performance (especially in the 1/2-shot setting) introduced by the GCN layers (GATE-AD (GCN)), demonstrates the **usefulness of the GAT's anisotropic aggregation mechanism**, which allows the model to assign varying importance to neighboring nodes. c) The decreased performance of TE indicates the **robustness of GAT layers in efficiently modeling fine-grained, local relations**, even in the case of the low-data FS-IVAD regime. d) The significant reduction in performance when the RA component is discarded (GATE-AD w/o RA) clearly highlights the **importance of the latent space mapping for accurately assessing reconstruction residuals**. e) The notable performance drop introduced by the integration of a decoder (GATE-AD AE) evidently illustrates the **effectiveness of normality modeling by means of a latent space**, rather than aiming at reconstructing the actual input feature tokens. f) Standard SSL techniques (MAE) exhibit

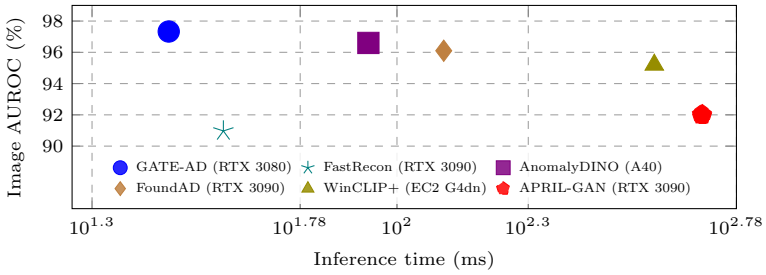

Fig. 3: FS-IVAD detection performance (image-level AUROC) versus per-image inference time (ms), under the 1-shot setting on the MVTec AD benchmark.

remarkably lower detection performance, which confirms that common **masked autoencoders lack the necessary relational context and inductive bias capabilities**, required to model normality from very few samples.

4.4 Time performance

Reported performance and discussion. In order to assess computational complexity aspects, Fig. 3 illustrates the achieved detection accuracy (image-level AUROC) versus the measured inference time (ms) per-image for GATE-AD and key baselines methods of the literature (covering all FS-IVAD methodological categories (Section 2)), under the 1-shot setting on the MVTec AD benchmark. It needs to be mentioned that the performance of GATE-AD is measured using an NVIDIA GeForce RTX 3080 GPU (Section 4.1), while for FastRecon [7], AnomalyDINO [6], FoundAD [32], WinCLIP+ [12], and APRIL-GAN [3] results are copied from the respective original publication works (also involving more high-performing GPUs). From the presented results, it can be seen that GATE-AD combines the highest detection accuracy with the lowest per-image inference time (25.05% faster than the most time-efficient baseline method, namely FastRecon). The latter demonstrates that the inherently computationally efficient nature of GATs (they do not require costly matrix operations) is elegantly coupled with their ability to model discriminant, local-level relations. A detailed ablation study regarding time performance (during both training and inference phases), concerning the main architectural blocks of GATE-AD (Section 4.3), is provided in Appendix C.

5 Conclusion

In this study, the GATE-AD approach to Few-Shot Industrial Visual Anomaly Detection (FS-IVAD) was introduced. By utilizing dense, local, visual feature tokens and incorporating a masked, representation-aligned Graph Attention Network (GAT) encoding scheme, the model effectively learns complex, irregular, non-Euclidean, patch-level relations, while robustly addressing the inherent

graph feature over-smooth-ing issue. Extensive comparative evaluation across the MVTec AD, VisA, and MPDD public benchmarks demonstrated that GATE-AD elegantly combines the highest detection accuracy with the lowest per image inference latency, compared to the best-performing literature methods. Future research directions include the integration with multi-modal foundation models for encoding more complex semantic relations, the optimization of the latent space construction to model more explicitly the normal samples intra-class variance, and the adaptation of the GAT encoder to continual learning scenarios.

References

1. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: The MVTec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. IJCV **129**(4), 1038–1059 (2021) 8
2. Chao, Y., Liu, J., Tang, J., Wu, G.: Anomalyr1: A GRPO-based end-to-end MLLM for industrial anomaly detection. arXiv preprint arXiv:2504.11914 (2025) 4
3. Chen, X., Han, Y., Zhang, J.: APRIL-GAN: A zero-/few-shot anomaly classification and segmentation method for CVPR 2023 VAND workshop challenge tracks 1&2. arXiv preprint arXiv:2305.17382 (2023) 2, 10, 11, 14
4. Cohen, N., Hoshen, Y.: Sub-image anomaly detection with deep pyramid correspondences. arXiv preprint arXiv:2005.02357 (2020) 2, 10, 11
5. Cui, Y., Liu, Z., Lian, S.: A survey on unsupervised anomaly detection algorithms for industrial images. IEEE Access **11**, 55297–55315 (2023). https://doi.org/10.1109/ACCESS.2023.3282993 2
6. Damm, S., Laszkiewicz, M., Lederer, J., Fischer, A.: Anomalydino: Boosting patch-based few-shot anomaly detection with DINOv2. In: IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV). pp. 1319–1329. IEEE (2025) 2, 3, 8, 10, 11, 14, 23, 26
7. Fang, Z., Wang, X., Li, H., Liu, J., Hu, Q., Xiao, J.: Fastrecon: Few-shot industrial anomaly detection via fast feature reconstruction. In: ICCV. pp. 17481–17490 (2023) 2, 3, 4, 5, 9, 10, 11, 14
8. Gu, Z., Zhu, B., Zhu, G., Chen, Y., Tang, M., Wang, J.: Univad: A training-free unified model for few-shot visual anomaly detection. In: CVPR. pp. 15194–15203 (2025) 3
9. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR. pp. 16000–16009 (2022) 13
10. Hou, Z., Liu, X., Cen, Y., Dong, Y., Yang, H., Wang, C., Tang, J.: Graphmae: Self-supervised masked graph autoencoders. In: Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min. pp. 594–604 (2022) 7, 8
11. Huang, C., Guan, H., Jiang, A., Zhang, Y., Spratling, M., Wang, Y.F.: Registration based few-shot anomaly detection. In: ECCV. pp. 303–319. Springer (2022) 10, 11
12. Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: Winclip: Zero-/few-shot anomaly classification and segmentation. In: CVPR. pp. 19606–19616 (2023) 2, 4, 10, 11, 14, 26
13. Jezek, S., Jonak, M., Burget, R., Dvorak, P., Skotak, M.: Deep learning-based defect detection of metal parts: Evaluating current methods in complex conditions. In: 13th Int. Congr. Ultra Modern Telecommun. Control Syst. Workshops (ICUMT). pp. 66–71. IEEE (2021) 9

14. Khoshraftar, S., An, A.: A survey on graph representation learning methods. ACM Transactions on Intelligent Systems and Technology **15**(1), 1–55 (2024) 7
15. Li, Y., Goodge, A., Liu, F., Foo, C.S.: Promptad: Zero-shot anomaly detection using text prompts. In: IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV). pp. 1093–1102 (2024) 2, 4, 10, 11
16. Li, Y., Zhang, S., Li, K., Lao, Q.: One-to-normal: Anomaly personalization for few-shot anomaly detection. NeurIPS **37**, 78371–78393 (2024) 2, 3, 4, 5
17. Li, Y., Ivanova, E., Bruveris, M.: FADE: Few-shot/zero-shot anomaly detection engine using large vision-language model. BMVC (2024) 4
18. Li, Z., Yan, Y., Wang, X., Ge, Y., Meng, L.: A survey of deep learning for industrial visual anomaly detection. Artif. Intell. Rev. **58**(9), 279 (2025) 2
19. Liu, J., Xie, G., Wang, J., Li, S., Wang, C., Zheng, F., Jin, Y.: Deep industrial image anomaly detection: A survey. Mach. Intell. Res. **21**(1), 104–135 (2024) 2
20. Lv, W., Su, Q., Xu, W.: One-for-all few-shot anomaly detection via instance-induced prompt learning. In: ICLR (2025), https://openreview.net/forum?id=Zzs3JwknAY 2, 4, 10, 11
21. Nikolentzos, G., Dasoulas, G., Vazirgiannis, M.: k-hop graph neural networks. Neural Networks **130**, 195–205 (2020) 6
22. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023) 5
23. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: CVPR. pp. 14318–14328 (2022) 2, 3, 10, 11
24. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. arXiv preprint arXiv:2508.10104 (2025) 5
25. Tao, F., Xie, G.S., Zhao, F., Shu, X.: Kernel-aware graph prompt learning for few-shot anomaly detection. In: AAAI. pp. 7347–7355. No. 7 (2025) 3
26. Tian, L., Li, Y., Dai, Y., Chen, W., Liu, X., Chen, B.: Fastref: Fast prototype refinement for few-shot industrial anomaly detection. arXiv preprint arXiv:2506.21398 (2025) 2, 4
27. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017) 5, 6, 8
28. Wang, Q., Guo, J., Gao, S., Wang, H., Xiong, L., Hu, J., Guo, H., Zhang, W.: Search is all you need for few-shot anomaly detection. arXiv preprint arXiv:2504.11895 (2025) 3
29. Wu, S., Wang, Y., Liu, X., Yang, Y., Wang, R., Guo, G., Doermann, D., Zhang, B.: DFM: Differentiable feature matching for anomaly detection. In: CVPR. pp. 15224–15233 (2025) 3
30. Xie, G., Wang, J., Liu, J., Zheng, F., Jin, Y.: Pushing the limits of fewshot anomaly detection in industry vision: GraphCore. In: ICLR (2023) 2, 3, 9, 10, 11
31. You, Z., Yang, K., Luo, W., Cui, L., Zheng, Y., Le, X.: Adtr: Anomaly detection transformer with feature reconstruction. In: Int. Conf. Neural Inform. Process. pp. 298–310. Springer (2022) 3, 5
32. Zhai, G., Zhou, Y., Deng, X., Heckler, L., Navab, N., Busam, B.: Foundation visual encoders are secretly few-shot anomaly detectors. The Fourteenth International Conference on Learning Representations (ICLR) (2026) 4, 10, 11, 14, 26
33. Zhu, J., Pang, G.: Toward generalist anomaly detection via in-context residual learning with few-shot sample prompts. In: CVPR. pp. 17826–17836 (2024) 4

34. Zou, Y., Jeong, J., Pemula, L., Zhang, D., Dabeer, O.: Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In: ECCV. pp. 392–408. Springer (2022) 8

Appendix

A Hyper-parameter tuning

This section provides a detailed analysis of the staged ablation study conducted to determine the most suitable and robust hyper-parameter configuration for the proposed GATE-AD framework. In order to ensure a stable and consistent experimentation protocol, the study follows a sequential tuning strategy. In particular, starting from the default GATE-AD implementation (Section 4.1), the effect of each individual of the following hyper-parameters on the overall performance of the proposed framework is evaluated on the MVTec AD and VisA benchmarks: a) Input image resolution, b) Backbone, c) Number of last transformer layers, d) Number of attentional layers, e) Cardinality of GAT layers, f) Cardinality of latent representation, g) Objective function, h) Masking ratio, i) Dropout rate, and j) Anomalous patches ratio. In each of the subsequent experiments, the value of only a single hyper-parameter is varied, while all remaining ones are maintained fixed to their default value.

A.1 Input image resolution

Fig. 4 illustrates how the input image resolution β (Section 3.2) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that resolutions of 512 and 768 pixels exhibit comparable performance. On the contrary, a resolution of 256 pixels results in noticeably lower I-AUROC and PRO scores across all shot counts. Additionally, while 768 pixels occasionally introduce slightly higher PRO scores in certain settings, the improvement over 512 pixels is marginal. To this end, a resolution of 512 pixels is selected for achieving an optimal balance between high detection accuracy and reduced computational cost.

A.2 Backbone

Fig. 5 illustrates how the backbone network (Section 3.2) affects the performance of GATE-AD across the MVTec AD, VisA, and MPDD benchmarks. From the presented results it can be seen that no single backbone or distillation size is proven to be universally optimal across all datasets. For MVTec AD, DINOv2 ViT-L/14 generally yields the best performance. For VisA, DINOv3 ViT-B/16 is proven to be the most effective. For MPDD, DINOv3 ViT-L/16 exhibits the highest accuracy.

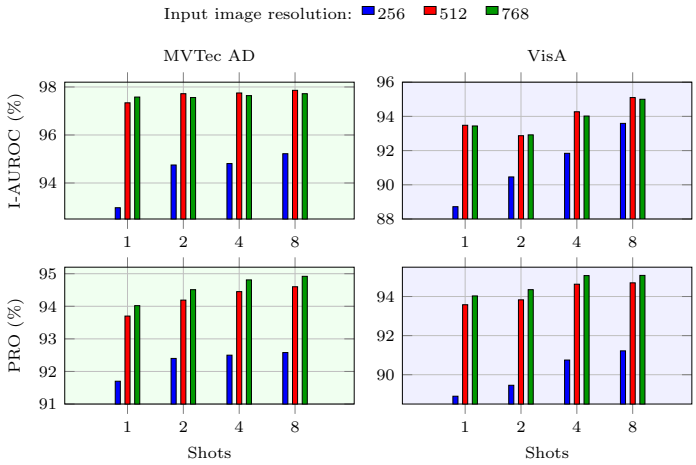

Fig. 4: Effect of the input image resolution on the MVTec AD and VisA benchmarks.

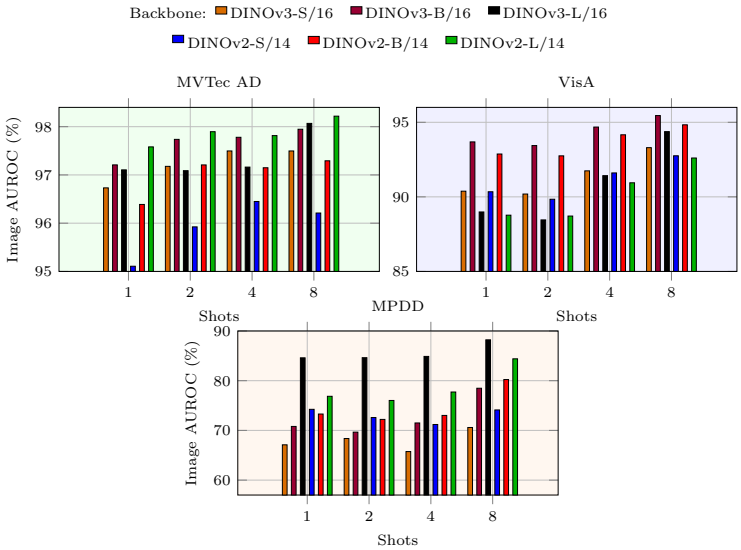

Fig. 5: Effect of the backbone on the MVTec AD, VisA, and MPDD benchmarks.

A.3 Number of last transformer layers

Fig. 6 illustrates how the number of the last transformer layers ζ (Section 3.2) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that increasing the number of layers generally improves Image AUROC, until it stabilizes or slightly declines

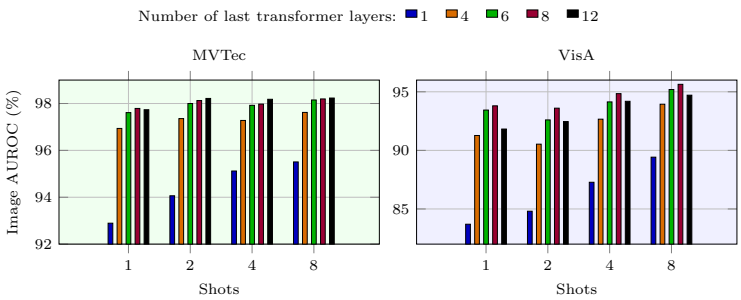

Fig. 6: Effect of the number of last transformer layers on the MVTec and VisA benchmarks.

when including too many layers. In particular, averaging patch token from the last $\zeta = 8$ layer leads to the highest overall performance.

A.4 Number of attentional layers

Fig. 7 illustrates how the number of the number of attentional layers R (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that setting $R = 2$ or $R = 3$ results in similar performance, though using 3 layers is slightly better in terms of achieving the highest I-AUROC and PRO scores. Further increasing the network depth ($R = 4$) does not necessarily yield better results and, in some cases, leads to slight performance decrease. Based on the above observations, $R = 3$ is used to ensure robust encoding of local relations.

A.5 Cardinality of GAT layers

Fig. 8 illustrates how the cardinality of the GAT layers F (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that $F = 256$ is sufficient for reaching high recognition performance. Further increasing the cardinality (beyond 256) results into negligible or inconsistent improvements in Image AUROC and PRO scores. $F = 256$ is selected, since it maintains high detection accuracy, while ensuring low inference latency and reduced computational cost.

A.6 Cardinality of latent representation

Fig. 9 illustrates how the cardinality of the latent representation f (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that setting $f = 256$ or higher leads to negligible variance in image AUROC. To this end, f is selected equal to 256 for achieving a low computational cost, without sacrificing recognition performance.

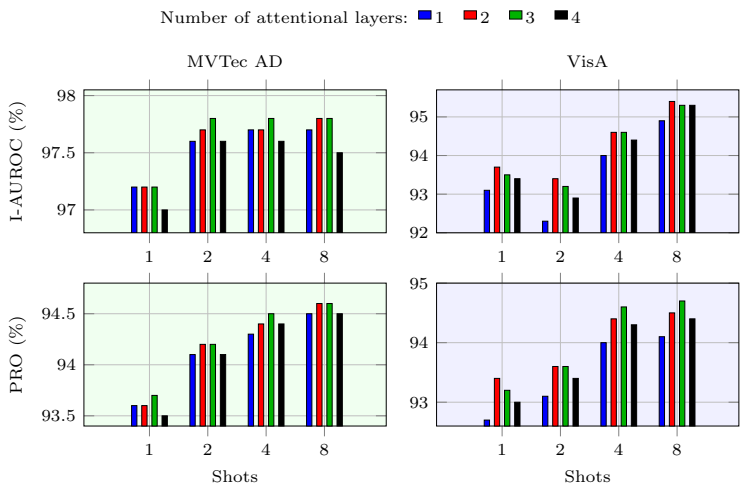

Fig. 7: Effect of the number of attentional layers on the MVTec AD and VisA benchmarks.

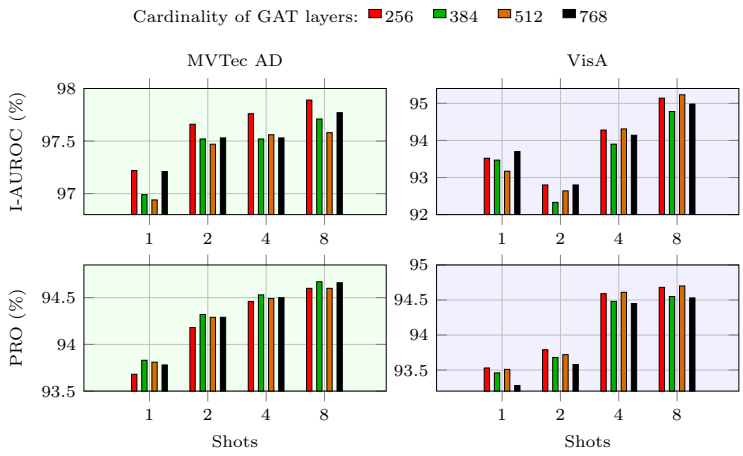

Fig. 8: Effect of the cardinality of the GAT layers on the MVTec AD and VisA benchmarks.

A.7 Objective function

Fig. 10 illustrates how the objective function (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that SCE is clearly advantageous over both MSE and cosine error. Concerning the SCE function, setting the scaling factor $\gamma = 2$ leads to the most consistent results.

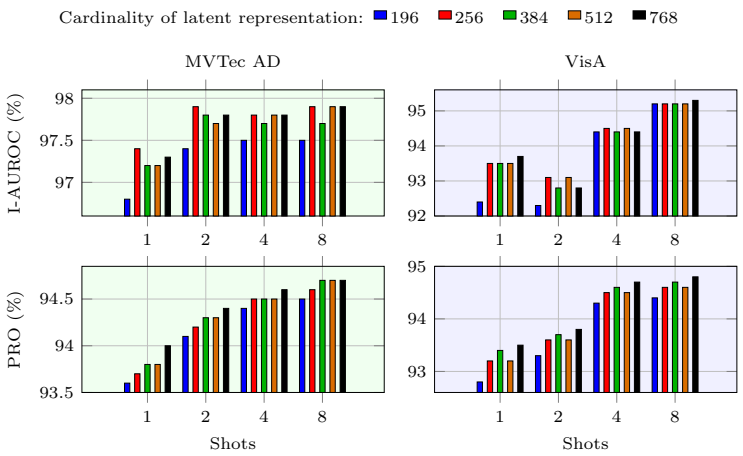

Fig. 9: Effect of the cardinality of the latent representation on the MVTec AD and VisA benchmarks.

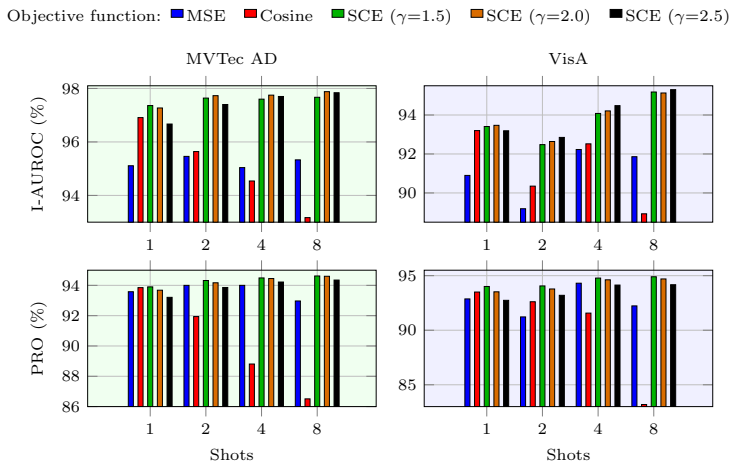

Fig. 10: Effect of the objective function on the MVTec AD and VisA benchmarks.

A.8 Masking ratio

Fig. 11 illustrates how the masking ratio μ (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that $\mu \in [0.1, 0.6]$ results into negligible performance variations, where selecting $\mu = 0.2$ leads to the best overall results.

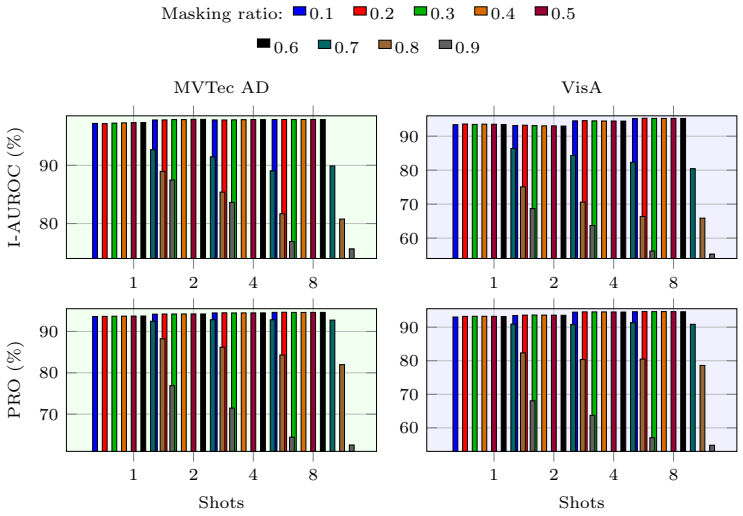

Fig. 11: Effect of the masking ratio on the MVTec and VisA benchmarks.

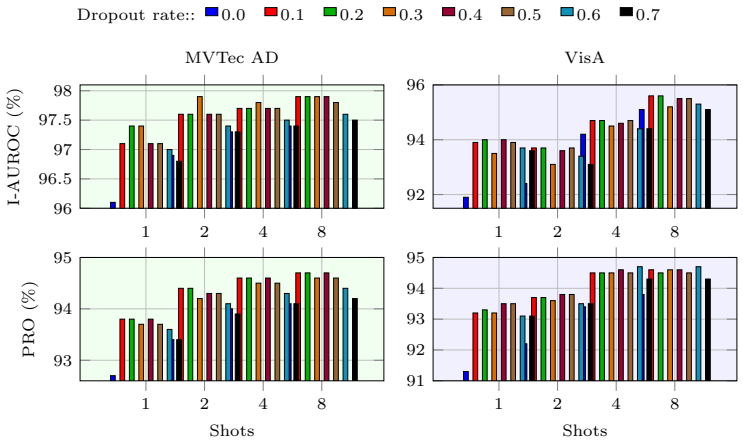

Fig. 12: Effect of the dropout rate ablation on the MVTec and VisA benchmarks.

A.9 Dropout rate

Fig. 12 illustrates how the dropout rate δ (Section 3.3) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that $\delta \in [0.1, 0.4]$ results in the highest detection performance, where $\delta = 0.3$ being the optimal value.

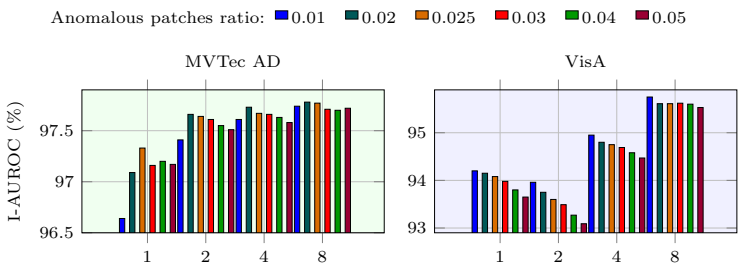

Fig. 13: Effect of the anomalous patches ratio on the MVTec AD and VisA benchmarks.

A.10 Anomalous patches ratio

Fig. 13 illustrates how the anomalous patches ratio ϵ (Section 3.4) affects the performance of GATE-AD across the MVTec AD and VisA benchmarks. From the presented results it can be seen that a ratio $\epsilon = 0.025$ in MVTec AD and $\epsilon = 0.01$ in VisA leads to the best overall performance. Significant deviations from these values (either much higher or lower in MVTec AD, or much higher in VisA) lead to a small decrease in detection performance.

B Object detection performance

This section provides a comprehensive, object-specific evaluation of the GATE-AD framework, complementary to the overall/aggregated defect detection results discussed in Section 4.2. In particular, detailed quantitative, qualitative and failure case results are discussed.

B.1 Quantitative evaluation results

Tables 3-5 illustrate the object-level detection performance of GATE-AD and AnomalyDINO [6] across the MVTec AD, VisA, and MPDD benchmarks. In order to ensure a fair and consistent comparison, both methods utilize the standardized sample selection process (Seed: 0). Key insights on a per dataset basis are provided in the following.

MVTec AD (Table 3):

- GATE-AD achieves high detection scores even in the 1-shot setting. accomplishing an 100.0% image AUROC for classes 'Bottle' and 'Leather'.
- GATE-AD performs robustly for challenging classes, like 'Cable' (94.7% image AUROC in the 1-shot setting) and 'Capsule' (93.1% image AUROC in the 1-shot case), significantly outperforming AnomalyDINO.
- GATE-AD maintains high localization accuracy in most cases, exhibiting up to 98.0% pixel PRO for class 'Leather' in the 1-shot setting.

Table 3: FS-IVAD object-level detection results for GATE-AD and AnomalyDINO on the MVTec AD benchmark. Both methods use the standardized sample selection process (Seed: 0).

Object	1-shot		2-shot		4-shot		8-shot	
	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO
Image AUROC (%)								
Bottle	**100.0** ±0.0	99.7 ±0.2	99.8 ±0.1	**99.9** ±0.1	**100.0** ±0.0	99.9 ±0.1	**100.0** ±0.0	**100.0** ±0.0
Cable	**94.7** ±0.1	92.7 ±0.8	**95.5** ±0.1	92.4 ±1.1	**94.3** ±0.1	93.8 ±0.9	95.0 ±0.4	**95.2** ±0.3
Capsule	**93.1** ±0.1	90.2 ±5.5	**94.3** ±0.7	89.2 ±7.9	95.0 ±0.9	**95.8** ±0.5	**95.8** ±0.8	95.6 ±0.5
Carpet	99.8 ±0.0	**100.0** ±0.0	99.9 ±0.0	**100.0** ±0.0	99.8 ±0.0	**100.0** ±0.0	99.9 ±0.0	**100.0** ±0.0
Grid	**99.6** ±0.0	99.1 ±0.2	**100.0** ±0.0	99.2 ±0.4	**100.0** ±0.0	99.5 ±0.3	**100.0** ±0.0	99.5 ±0.1
Hazelnut	**97.7** ±0.9	97.5 ±2.6	98.9 ±0.4	**99.6** ±0.5	99.6 ±0.3	**99.8** ±0.1	99.9 ±0.1	**100.0** ±0.0
Leather	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0
Metal nut	99.4 ±0.3	**99.9** ±0.1	99.2 ±0.3	**100.0** ±0.0	99.7 ±0.1	**100.0** ±0.0	99.7 ±0.0	**100.0** ±0.0
Pill	**96.3** ±1.0	93.7 ±0.9	**97.3** ±0.6	95.4 ±0.7	**96.5** ±0.5	96.0 ±0.2	**97.7** ±0.3	97.2 ±0.2
Screw	**94.6** ±0.2	93.2 ±0.3	**95.1** ±0.2	93.5 ±0.8	**94.5** ±0.0	92.7 ±2.3	**94.5** ±0.2	93.5 ±1.1
Tile	99.5 ±0.1	**100.0** ±0.0	99.7 ±0.0	**100.0** ±0.0	99.9 ±0.0	**100.0** ±0.0	**100.0** ±0.0	**100.0** ±0.0
Toothbrush	**98.1** ±0.4	97.4 ±0.5	97.5 ±0.0	**98.1** ±1.0	96.4 ±0.0	**97.5** ±0.6	96.7 ±0.4	**97.7** ±1.6
Transistor	**93.0** ±0.4	90.9 ±1.2	**93.1** ±0.7	89.4 ±4.6	**93.3** ±0.4	93.2 ±2.2	94.6 ±0.1	**96.2** ±1.3
Wood	**99.8** ±0.0	98.0 ±0.2	**100.0** ±0.0	98.0 ±0.1	**100.0** ±0.0	97.9 ±0.2	**100.0** ±0.0	98.3 ±0.4
Zipper	**99.8** ±0.2	97.4 ±0.9	**99.9** ±0.1	98.9 ±0.4	**99.8** ±0.1	99.0 ±0.4	**99.9** ±0.1	99.6 ±0.1
Mean	**97.7** ±0.2	96.6 ±0.4	**98.0** ±0.2	96.9 ±0.7	**97.9** ±0.2	97.7 ±0.2	**98.2** ±0.1	**98.2** ±0.2
Pixel PRO (%)								
Bottle	95.7 ±0.1	**95.9** ±0.3	95.7 ±0.1	**96.3** ±0.0	95.9 ±0.1	**96.7** ±0.1	96.0 ±0.1	**96.5** ±0.4
Cable	**89.8** ±0.4	89.4 ±0.4	**90.4** ±0.1	89.5 ±0.4	90.1 ±0.1	**90.4** ±0.3	**90.5** ±0.1	**90.5** ±0.2
Capsule	96.6 ±0.1	**97.1** ±0.3	96.7 ±0.1	**97.3** ±0.6	97.0 ±0.1	**97.9** ±0.1	97.1 ±0.1	**97.9** ±0.1
Carpet	96.7 ±0.0	**97.8** ±0.0	96.8 ±0.0	**97.9** ±0.0	96.7 ±0.1	**97.8** ±0.0	96.7 ±0.1	**97.8** ±0.0
Grid	96.5 ±0.2	**97.2** ±0.1	96.6 ±0.2	**97.2** ±0.1	96.6 ±0.1	**97.2** ±0.1	96.6 ±0.1	**97.2** ±0.0
Hazelnut	96.7 ±0.1	**97.4** ±0.4	97.1 ±0.0	**98.0** ±0.3	97.3 ±0.1	**98.0** ±0.1	97.3 ±0.1	**98.1** ±0.1
Leather	**98.0** ±0.2	97.9 ±0.1	**97.9** ±0.2	97.8 ±0.0	**97.9** ±0.1	97.6 ±0.1	**97.8** ±0.2	97.6 ±0.1
Metal nut	92.2 ±0.1	**94.2** ±0.0	92.7 ±0.2	**94.6** ±0.2	93.4 ±0.0	**95.3** ±0.1	93.8 ±0.2	**95.4** ±0.3
Pill	96.6 ±0.1	**97.3** ±0.1	96.9 ±0.1	**97.5** ±0.1	96.7 ±0.1	**97.6** ±0.1	96.9 ±0.1	**97.7** ±0.1
Screw	**96.1** ±0.1	93.4 ±0.4	**96.6** ±0.1	94.3 ±0.4	**96.6** ±0.1	94.5 ±0.9	**96.8** ±0.1	95.2 ±0.5
Tile	**89.7** ±0.4	88.0 ±0.2	**90.1** ±0.5	87.6 ±0.4	**90.0** ±0.4	87.1 ±0.4	**89.9** ±0.4	86.7 ±0.2
Toothbrush	**96.0** ±0.0	94.0 ±0.8	**95.9** ±0.1	94.7 ±0.3	**95.5** ±0.0	94.7 ±0.3	**95.7** ±0.0	95.1 ±0.8
Transistor	**68.2** ±0.3	67.3 ±2.1	**70.5** ±0.1	68.4 ±3.1	**70.9** ±0.1	70.6 ±1.4	71.4 ±0.1	**75.3** ±1.9
Wood	93.4 ±0.1	**94.7** ±0.1	93.6 ±0.0	**94.6** ±0.0	93.4 ±0.0	**94.6** ±0.1	93.9 ±0.1	**94.4** ±0.3
Zipper	**93.2** ±0.1	89.2 ±1.2	**94.2** ±0.4	90.2 ±0.4	**94.1** ±0.3	91.2 ±0.3	**94.8** ±0.3	91.1 ±0.4
Mean	**93.0** ±0.2	92.7 ±0.1	**93.4** ±0.1	93.1 ±0.2	**93.5** ±0.1	93.4 ±0.1	93.7 ±0.1	**93.8** ±0.1

VisA (Table 4):

- GATE-AD exhibits high performance for most classes across all settings, e.g., 'Capsules' (99.3% image AUROC in the 1-shot case) and 'Chewinggum' (98.8% image AUROC in the 1-shot case), clearly surpassing the respective AnomalyDINO metrics.
- For complex electronic components (i.e., printed circuit boards), GATE-AD maintains high detection accuracy: 'PCB1' (98.0% image AUROC in the 1-shot setting), 'PCB2' (92.0% image AUROC in the 1-shot case), and 'PCB4' (88.3% image AUROC in the 1-shot setting).
- For food and organic shapes, GATE-AD exhibits increased localization performance, e.g., 'Capsules' (99.3% image AUROC in the 1-shot setting) and 'Chewing gum'(98.8% image AUROC in the 1-shot case).

MPDD (Table 5):

Table 4: FS-IVAD object-level detection results for GATE-AD and AnomalyDINO on the VisA benchmark. Both methods use the standardized sample selection process (Seed: 0).

Object	1-shot		2-shot		4-shot		8-shot	
	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO
Image AUROC (%)								
Candle	**97.8** ±0.4	87.9 ±0.3	**97.6** ±0.3	89.4 ±3.0	**97.5** ±0.3	91.3 ±2.9	**97.7** ±0.2	93.5 ±1.2
Capsules	**99.3** ±0.3	98.4 ±0.5	**99.5** ±0.2	98.9 ±0.1	**99.6** ±0.2	99.2 ±0.1	**99.6** ±0.2	99.2 ±0.1
Cashew	**95.5** ±0.4	86.1 ±3.6	**96.1** ±0.2	89.4 ±3.8	**98.7** ±0.2	94.5 ±0.7	**98.6** ±0.2	95.3 ±0.6
Chewing gum	**98.8** ±0.2	98.0 ±0.4	**99.1** ±0.1	98.6 ±0.4	**99.3** ±0.1	98.8 ±0.2	**99.4** ±0.1	98.8 ±0.2
Fryum	**97.4** ±0.4	94.8 ±0.5	**98.4** ±0.1	96.5 ±0.2	**98.7** ±0.2	97.0 ±0.1	**98.6** ±0.2	97.6 ±0.4
Macaroni1	**96.0** ±0.4	87.5 ±1.1	**95.4** ±0.5	87.5 ±0.9	**95.9** ±0.3	89.5 ±1.4	**96.9** ±0.3	90.1 ±1.7
Macaroni2	**75.9** ±1.4	62.2 ±4.3	**73.3** ±1.2	66.9 ±1.9	**72.6** ±1.0	70.0 ±1.7	**75.7** ±0.7	74.9 ±0.4
PCB1	**98.0** ±0.3	91.5 ±2.0	**97.9** ±0.4	91.2 ±2.7	**97.6** ±0.4	94.0 ±2.1	**97.6** ±0.3	95.5 ±0.5
PCB2	**92.0** ±0.5	84.8 ±1.2	**91.5** ±0.6	88.1 ±2.5	**91.4** ±0.7	91.1 ±1.7	92.1 ±0.3	**92.6** ±0.3
PCB3	**85.7** ±0.3	84.9 ±3.3	**91.1** ±0.2	89.4 ±3.8	94.0 ±0.2	**94.3** ±0.4	93.6 ±0.4	**95.6** ±0.2
PCB4	**88.3** ±0.8	79.9 ±13.7	78.2 ±2.0	**87.4.3-**	91.6 ±1.0	**96.2** ±2.6	96.2 ±0.3	**98.0** ±0.3
Pipe fryum	**96.9** ±0.6	92.7 ±2.7	**95.9** ±0.7	93.3 ±1.2	**96.1** ±0.5	94.6 ±1.9	**96.7** ±0.5	95.0 ±1.6
Mean	**93.5** ±0.2	87.4 ±1.2	**92.8** ±0.3	89.7 ±1.3	**94.4** ±0.2	92.6 ±0.9	**95.2** ±0.2	93.8 ±0.3
Pixel PRO (%)								
Candle	**97.0** ±0.2	96.8 ±0.4	**97.1** ±0.2	97.0 ±0.2	**97.3** ±0.2	97.2 ±0.1	**97.4** ±0.2	97.3 ±0.2
Capsules	**97.0** ±0.3	95.1 ±0.7	**97.1** ±0.3	95.5 ±0.2	**97.4** ±0.3	96.3 ±0.4	**97.7** ±0.3	96.7 ±0.2
Cashew	**96.4** ±0.2	96.1 ±0.9	**96.8** ±0.3	96.7 ±0.7	96.9 ±0.4	**97.4** ±0.5	96.9 ±0.3	**97.3** ±0.2
Chewing gum	**93.0** ±1.0	92.0 ±1.0	**93.2** ±0.9	92.9 ±0.3	**93.7** ±0.7	93.0 ±0.1	**93.9** ±0.5	93.1 ±0.3
Fryum	92.8 ±0.3	**93.2** ±0.2	93.8 ±0.4	**93.9** ±0.3	94.0 ±0.3	**94.5** ±0.4	93.9 ±0.3	**94.9** ±0.3
Macaroni1	97.1 ±0.3	**97.5** ±0.3	97.3 ±0.3	**97.9** ±0.3	97.6 ±0.3	**98.3** ±0.2	97.9 ±0.2	**98.6** ±0.2
Macaroni2	**96.1** ±0.3	92.0 ±0.7	**96.1** ±0.3	93.0 ±0.4	**96.4** ±0.3	93.9 ±0.8	**96.8** ±0.3	95.0 ±0.6
PCB1	**93.5** ±0.5	92.6 ±0.2	**94.2** ±0.4	92.5 ±0.5	**94.4** ±0.3	93.3 ±0.5	**94.2** ±0.3	93.9 ±0.2
PCB2	88.8 ±0.5	**89.9** ±0.2	88.9 ±0.4	**90.7** ±0.3	89.5 ±0.3	**91.4** ±0.2	89.6 ±0.3	**92.0** ±0.1
PCB3	84.8 ±0.2	**88.5** ±1.3	86.9 ±0.2	**90.8** ±0.5	89.8 ±0.3	**91.7** ±0.4	88.7 ±0.3	**93.1** ±0.3
PCB4	**84.1** ±0.5	78.5 ±6.8	**82.9** ±0.7	82.0 ±6.1	**88.9** ±0.5	84.1 ±1.5	**90.0** ±0.6	87.9 ±2.3
Pipe fryum	97.7 ±0.1	**98.0** ±0.0	97.8 ±0.1	**97.9** ±0.2	**97.9** ±0.1	97.8 ±0.1	**97.8** ±0.1	97.6 ±0.1
Mean	**93.2** ±0.3	92.5 ±0.5	**93.5** ±0.3	93.4 ±0.6	**94.5** ±0.2	94.1 ±0.1	94.6 ±0.3	**94.8** ±0.2

Table 5: FS-IVAD object-level detection results for GATE-AD and AnomalyDINO on the MPDD benchmark. Both methods use the standardized sample selection process (Seed: 0).

Object	1-shot		2-shot		4-shot		8-shot	
	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO	GATE-AD	AnomalyDINO
Image AUROC (%)								
Bracket black	**83.7** ±3.0	51.0 ±6.5	**83.4** ±2.0	59.4 ±8.0	**80.9** ±2.0	64.2 ±8.1	**85.7** ±0.6	69.5 ±2.7
Bracket brown	49.7 ±1.9	**52.0** ±4.6	51.7 ±1.3	**56.8** ±4.4	56.1 ±2.1	**58.6** ±1.2	57.0 ±2.0	**61.3** ±2.8
Bracket white	**80.8** ±3.3	70.8 ±2.3	**81.1** ±2.2	72.0 ±2.8	**81.0** ±1.9	65.7 ±16.7	**93.3** ±1.7	72.5 ±20.1
Connector	**93.8** ±1.0	80.6 ±4.1	**91.8** ±0.8	83.1 ±3.8	**89.4** ±1.0	83.7 ±1.7	**93.6** ±1.3	87.9 ±3.1
Metal plate	99.3 ±0.6	**100.0** ±0.0	99.7 ±0.3	**100.0** ±0.0	99.8 ±0.3	**100.0** ±0.0	99.5 ±0.5	**100.0** ±0.0
Tubes	**97.6** ±0.8	89.2 ±1.4	**98.0** ±0.6	88.8 ±1.5	**97.6** ±0.7	91.6 ±1.9	**97.5** ±0.7	91.6 ±2.1
Mean	**84.2** ±0.9	73.9 ±1.5	**84.3** ±0.5	76.7 ±2.5	**84.1** ±0.9	77.3 ±2.4	**87.8** ±0.5	80.5 ±3.0
Pixel PRO (%)								
Bracket black	**97.4** ±0.5	89.1 ±0.6	**97.2** ±0.5	90.3 ±0.2	**97.3** ±0.3	91.5 ±1.7	**98.3** ±0.0	93.4 ±0.9
Bracket brown	**85.8** ±0.3	80.6 ±1.2	**85.7** ±0.3	83.7 ±1.7	**86.6** ±0.3	84.9 ±1.7	**86.9** ±0.6	85.3 ±2.6
Bracket white	**95.4** ±0.2	95.4 ±0.7	95.2 ±0.3	**96.3** ±1.6	95.4 ±0.3	**96.3** ±1.6	**97.6** ±0.2	95.5 ±2.0
Connector	**90.2** ±0.5	84.4 ±2.4	**89.9** ±0.5	86.2 ±2.8	**90.2** ±0.4	88.0 ±2.8	**91.2** ±0.3	89.0 ±0.7
Metal plate	**92.8** ±0.2	89.0 ±0.3	**93.4** ±0.2	89.4 ±0.9	**93.8** ±0.1	89.7 ±0.6	**93.6** ±0.1	90.5 ±0.8
Tubes	**98.7** ±0.2	86.1 ±1.4	**98.7** ±0.2	87.6 ±1.4	**98.6** ±0.2	88.5 ±1.1	**98.6** ±0.2	88.8 ±0.9
Mean	**93.4** ±0.2	87.4 ±0.6	**93.4** ±0.1	88.9 ±0.9	**93.7** ±0.1	89.8 ±1.1	**94.4** ±0.0	90.4 ±0.5

– GATE-AD achieves high detection scores even in the 1-shot setting, accomplishing an 99.3% image AUROC for class 'Metal plate' and a corresponding 97.6% for category 'Tubes'.

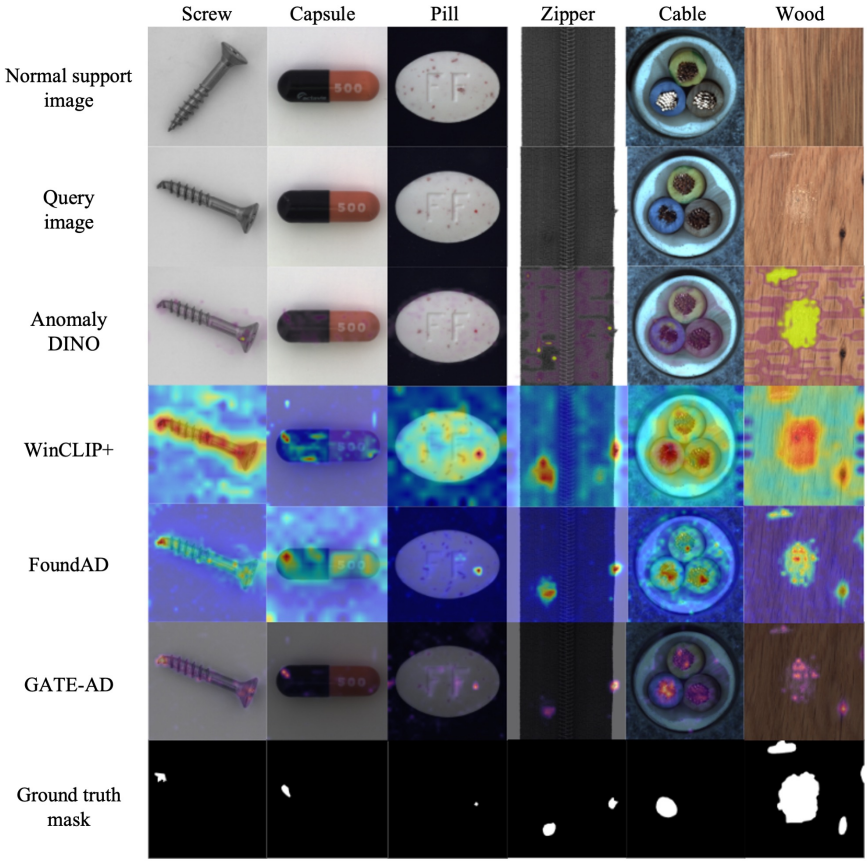

Fig. 14: Indicative 1-shot defect detection results on the MVTec AD benchmark.

- GATE-AD demonstrates consistent localization, achieving 98.7% pixel PRO in the 1-shot setting for category 'Tubes' and a corresponding 97.4% for class 'Bracket black'.
- Especially for the localization task, GATE-AD demonstrates higher pixel PRO rates for the vast majority of classes across all settings.

B.2 Qualitative evaluation results

Figs. 14-16 illustrate indicative 1-shot defect detection results from the application of GATE-AD, AnomalyDINO [6], WinCLIP+ [12], and FoundAD [32] across the MVTec AD, VisA, and MPDD benchmarks. In order to ensure a fair comparison, all baseline methods are evaluated using their default (open source) configuration, while the same single normal support image is employed. Overall, the presented visualizations reinforce the above quantitative analysis results

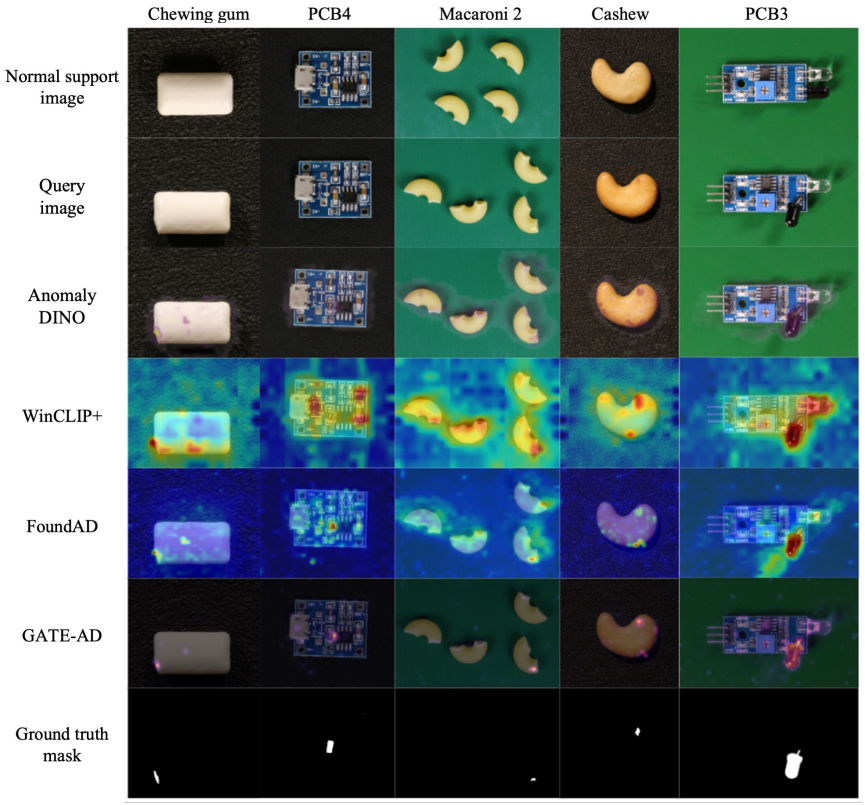

Fig. 15: Indicative 1-shot defect detection results on the VisA benchmark.

(Section B.1), demonstrating that GATE-AD localizes defects more accurately and robustly than other baseline methods, across multiple pose, placement, and abnormality type settings. Key insights on a per dataset basis are provided in the following.

MVTec AD (Fig. 14):

- GATE-AD demonstrates superior localization for both structural defects (e.g., in 'Screw' and 'Cable') and surface anomalies (e.g., in 'Pill' and 'Wood').
- While baseline methods (e.g., WinCLIP+) often produce diffuse or noisy heatmaps that cover large, non-defective areas, GATE-AD produces focused, high-intensity responses that align closely with the ground truth masks.
- GATE-AD effectively distinguishes between natural texture variations and genuine defects in categories such as 'Wood'.

VisA (Fig. 15):

- GATE-AD excels in scenes containing multiple objects or complex electronic layouts, such as 'PCB4' and 'PCB3'.

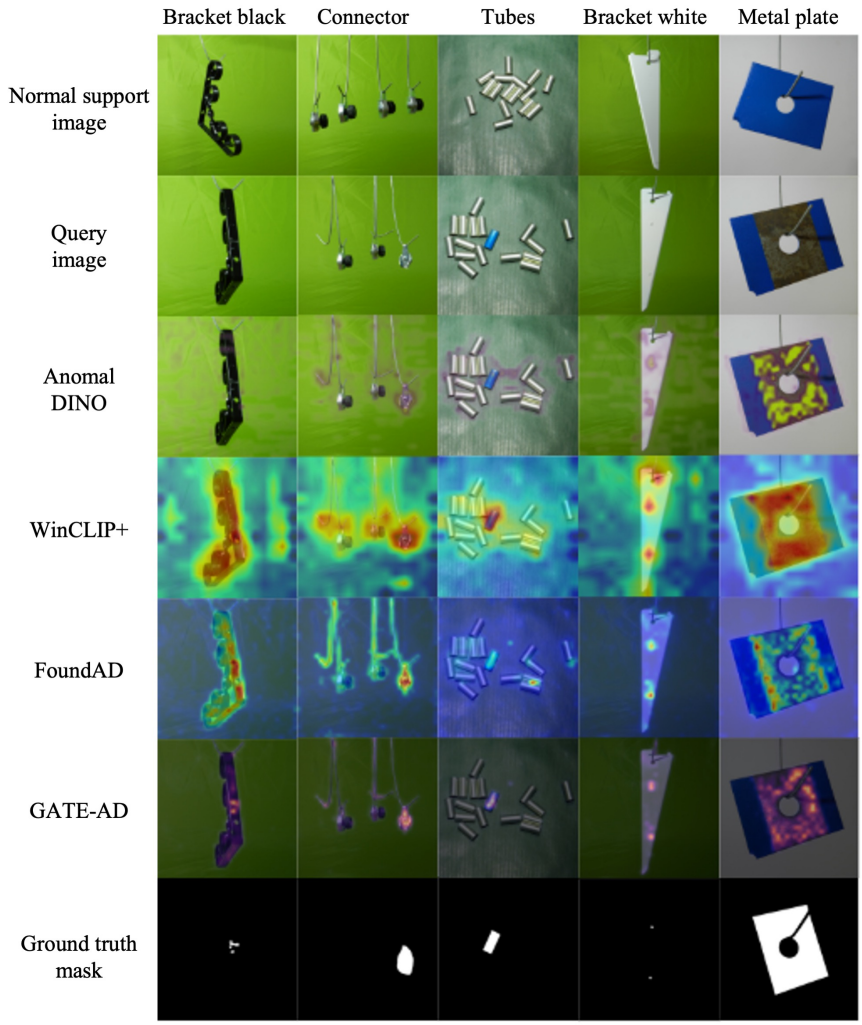

Fig. 16: Indicative 1-shot defect detection results on the MPDD benchmark.

- For organic or irregular shapes (e.g., 'Macaroni 2' and 'Cashew'), GATE-AD successfully detects fine-grained anomalies that other reconstruction-based methods (e.g., FoundAD) may fail to distinguish.
- GATE-AD generated anomaly maps exhibit minimal activation in background regions, in sharp contradistinction to most baseline methods.

MPDD (Fig. 16):

- GATE-AD demonstrates significant robustness in the case of varied poses and lighting conditions (e.g., 'Bracket black' and 'Tubes' classes).

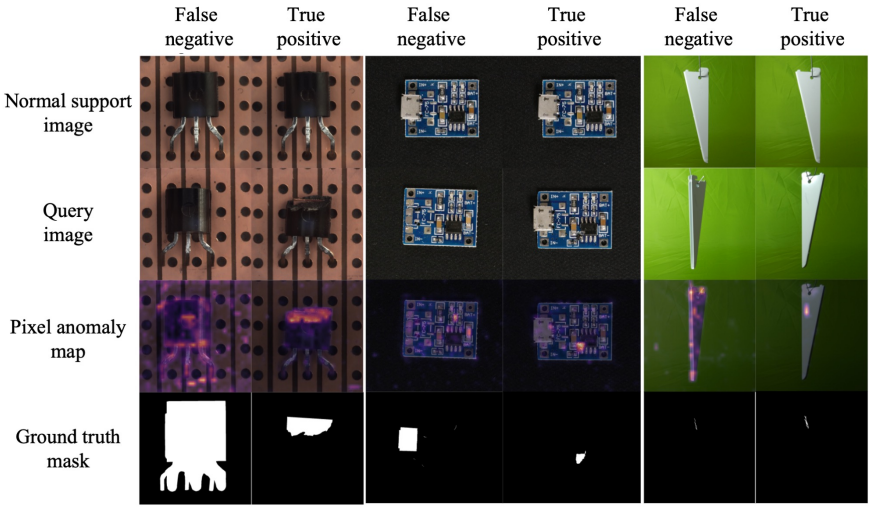

Fig. 17: Indicative examples of defect detection failure cases across the MVTec AD, VisA, and MPDD benchmarks.

- GATE-AD is also shown to accurately identify missing components or surface failures (e.g., 'Bracket black' and 'Connector' categories).
- Even when the query image has a different orientation than the single support image, GATE-AD detects local defect occurrences (e.g., 'Metal plate' class).

B.3 Failure cases

Fig. 17 illustrates indicative scenarios where GATE-AD may fail to detect or accurately localize anomalies across the MVTec AD, VisA, and MPDD benchmarks. Key insights extracted from the provided results are as follows:

- Primary source of failure comprises when anomalies arise from global or semantic changes, rather than local surface defects. This includes instances of rotated or misplaced components.
- If a query image contains a pose variation that was not present in the available (small) normal support set, GATE-AD may struggle to distinguish this natural variation from a genuine defect.
- Since GATE-AD is grounded on local-level reconstruction of patch features, it may be sensitive to distinct pixel-level mismatches, but less effective at capturing high-level structural anomalies that do not significantly disrupt local feature consistency.

Table 6: Per-image training and inference time (ms) on the MVTec AD, VisA, and MPDD benchmarks.

	MVTec AD		VisA		MPDD	
Method	Training	Inference	Training	Inference	Training	Inference
GATE-AD	**6.7** ±0.9	**0.03** ±0.00	**7.8** ±0.9	**0.04** ±0.00	4.8 ±0.6	**0.03** ±0.00
GATE-AD (GCN)	9.0 ±3.5	0.03 ±0.00	10.4 ±3.4	0.04 ±0.00	**4.4** ±0.3	0.03 ±0.00
TE	62.9 ±16.1	0.04 ±0.00	83.1 ±23.4	0.06 ±0.00	46.6 ±10.0	0.05 ±0.00
GATE-AD w/o RA	25.4 ±4.0	0.04 ±0.00	28.4 ±3.9	0.05 ±0.00	14.8 ±0.3	0.04 ±0.00
GATE-AD AE	16.1 ±6.5	0.04 ±0.00	19.0 ±6.7	0.05 ±0.00	18.6 ±1.3	0.04 ±0.00
MAE	52.4.4-	0.05 ±0.00	64.5 ±6.8	0.06 ±0.00	75.4 ±1.5	0.05 ±0.00

C Time performance

This section provides a detailed analysis of the computational complexity and time performance of the GATE-AD framework, focusing on both training and inference phases, across the MVTec AD, VisA, and MPDD benchmarks. In order to assess the importance and contribution of the main architectural blocks of GATE-AD, a comprehensive ablation study is performed, involving the evaluation of the same variants investigated in Section 4.3. In particular, Table 6 illustrates the average per-image training and inference time (ms), across all shot settings on the MVTec AD, VisA, and MPDD benchmarks. Key insights extracted from the provided results are as follows:

- Replacing GAT layers with conventional convolutional ones (GATE-AD (GCN)) results in slightly higher training time in most cases, while maintaining similar inference speed.
- Utilizing a global self-attention mechanism (TE) introduces an outstanding increase in both training and inference time, due to its higher computational overhead.
- Variants without the RA component (GATE-AD w/o RA) or those utilizing an MLP graph decoder (GATE-AD AE) generally exhibit significantly longer training and inference durations.
- The masked autoencoder (MAE) variant comprises one of the most computationally expensive ones.