

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 BREAKING THE BIAS: QUANTIFYING THE ATTENTION OF INDUSTRIAL ANOMALY DETECTION

Anonymous authors

Paper under double-blind review

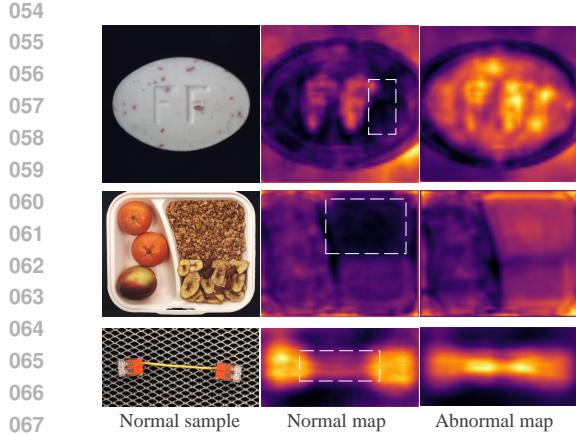
ABSTRACT

Industrial anomaly detection (IAD) predominantly utilizes unsupervised learning due to the scarcity and unpredictability of defect samples. A major challenge in unsupervised IAD methods is the inherent bias in normal samples, which causes models to focus on variable regions while overlooking potential defects in invariant areas. In this paper, we propose **Recalibrating Attention of Industrial Anomaly Detection** (RAAD), which decomposes and recalibrates the input data to highlight anomalies better. Additionally, Hierarchical Quantization Scoring (HQS) is introduced to refine the detection process by assigning quantization scores at multiple levels. These strategies work together to mitigate the bias toward normal samples and improve the accuracy of anomaly detection. We validate the effectiveness of RAAD on three IAD datasets: MVTec-AD, MVTec-LOCO, and VisA. The experimental results demonstrate that RAAD exhibits competitiveness in both detection and localization tasks, providing a robust solution for industrial anomaly detection. The source code will be released to promote further research and application.

1 INTRODUCTION

Industrial anomaly detection (IAD) is crucial for maintaining the quality and safety of manufacturing processes. Because of high annotation costs and the unpredictable nature of defects, unsupervised methods have become a practical solution for real-world anomaly detection, which is only trained on normal samples. However, traditional unsupervised methods face a fundamental challenge: during training, models tend to overfit the changing parts of normal samples while overlooking potential defects in unchanged regions. As shown in Figure 1, we visualize this challenge through anomaly heatmaps on MVTec-AD Bergmann et al. (2019a) and MVTec LOCO Bergmann et al. (2022) datasets. Within each group, from left to right, are the normal samples of the corresponding categories, the average anomaly heatmap for normal samples, and the average anomaly heatmap for anomalous samples. Brighter areas in the heatmaps indicate regions with a higher likelihood of receiving attention. The white boxes in the second column highlight how the model is misled by the inherent bias in normal samples. In other words, the attention maps derived from unsupervised training tend to highlight variable regions in normal samples, thereby neglecting invariant regions where subtle anomalies may reside. One might intuitively consider abandoning the attention mechanism. Nevertheless, ignoring attention maps entirely is not a viable solution, as they play a crucial role in anomaly detection. A key question arises: how to make the model allocate attention more reasonably?

A feasible solution is to solve this problem with two steps: first, directing the model’s attention toward the primary target, and then reallocating the attention for improved anomaly detection. The former can be achieved through model quantization, while the latter is accomplished via fine-tuning. During quantization, the reduction in parameter precision compels the model to prioritize learning and extracting the most critical information. Meanwhile, during the fine-tuning process, the model’s attention is recalibrated, enabling the redistribution of attention to better align with task-specific requirements. Building on this insight, we propose RAAD (**R**ecalibrating **A**ttention of **I**ndustrial **A**nomaly **D**etection), which firstly modifies attention maps with quantization and then fine-tuning them to recalibration. Meanwhile, we observe that convolutional neural networks are commonly used as backbone networks for extracting image features in industrial anomaly detection tasks, with each layer having a different impact on the model’s attention. To optimize the attention allocation process, we introduce **H**ierarchical **Q**uantization **S**coring (HQS), which



068
069
070
071
072
073
074
075
076
077

Figure 1: Visualization of heatmaps. These samples are from the MVTec-AD and MVTec LOCO datasets, which represent examples of industrial products, the average heatmap of normal samples, and the average heatmap for anomaly samples, respectively. It clearly shows the bias contained in the normal samples compared to the abnormal samples.

078
079
080
081
082
083
084
085

adaptively allocates bit-width according to each layer’s anomaly detection capability. In Figure 2(a), we visualize the anomaly heatmaps before and after model quantization. It can be observed that, compared to before quantization, the model’s attention is more spread across the main subject while ignoring the background. Subsequently, more precise anomaly detection was achieved after fine-tuning. Figure 2(b) illustrates the outputs of each convolutional layer in the teacher-student network, highlighting the layer-wise variation in focus across the image. This design leverages the distinct roles of network layers: shallow layers capture local details, while deeper layers extract global features, and it is most beneficial for enhancing model performance with fewer parameters.

086
087
088
089
090
091
092

Our main contributions are as follows: We break the inherent bias in attention allocation within unsupervised IAD, guiding models to better detect subtle anomalies in invariant regions. We proposed RAAD, which systematically refines attention maps, using quantization to reduce bias and recalibrating the attention map via fine-tuning to improve anomaly sensitivity. We introduce HQS, a module that dynamically allocates bit-widths based on each layer’s anomaly detection capability, optimizing the alignment between quantization and attention for enhanced efficiency and accuracy in IAD.

093

2 RELATED WORK

094

2.1 UNSUPERVISED INDUSTRIAL ANOMALY DETECTION.

095

096

097

Based on deep learning, visual detection has made significant achievements with the assistance of supervised learning, as cited in Kwon et al. (2019); Ruff et al. (2021). However, in real-world industrial scenarios, the scarcity of defect samples, the cost of annotation, and the lack of prior knowledge about defects may render supervised methods ineffective. In recent years, unsupervised anomaly detection (IAD) algorithms have been increasingly applied to industrial detection tasks, as referenced in Sydney et al. (2019); Xie et al. (2024); Tao et al. (2022). “Unsupervised” means that the training phase only includes normal images, without any defect samples. IAD refers to the task of differentiating defective images from the majority of non-defective images at the image level. Unsupervised IAD is mainly categorized into three types, i.e., the reconstruction-based methods, the synthesizing-based methods, and the embedding-based methods. Feature embedding-based methods have recently achieved state-of-the-art performance and can be specifically categorized into: teacher-student architecture Bergmann et al. (2020); Deng & Li (2022), normalizing flow Rezende &

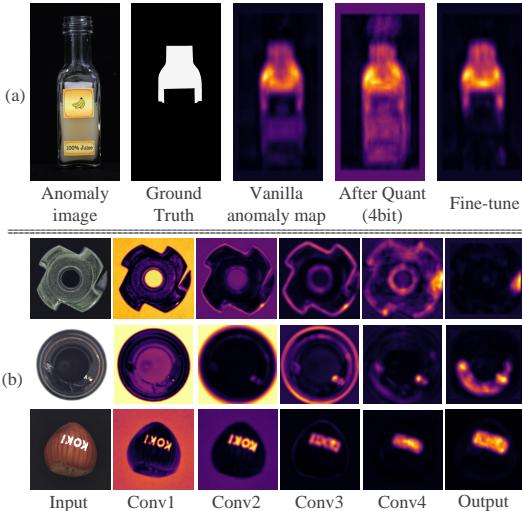


Figure 2: (a) Visualization of the attention maps at different stages of the model, from left to right, are the anomaly image, ground-truth, and predicted anomaly score. (b) The layer-wise attention outputs demonstrate the varying importance of each layer in anomaly detection.

108 Mohamed (2015); Rudolph et al. (2021), memory bank Roth et al. (2022); Cohen & Hoshen (2020),
 109 and one-class classification Sohn et al. (2020).

110 The most typical methods are the memory bank and teacher-student architecture. Memory bank
 111 methods embed normal features into a compressed space. Anomalous features are distant from
 112 the normal clusters within the embedding space. Regarding the teacher-student architecture, the
 113 teacher is a pre-trained and frozen CNN, and the student network is trained to mimic the teacher's
 114 output on training images. Since the student has not seen any anomalous images during training,
 115 it is generally unable to predict the teacher's output on these images, thereby achieving anomaly
 116 detection. Uninformed Students Bergmann et al. (2020) first introduced a new framework for anomaly
 117 detection known as the teacher-student anomaly detection framework. Reverse Distillation (RD)
 118 Deng & Li (2022) proposed a method where the student decoder learns to recover features from
 119 the compact embeddings of the teacher encoder. The GCCB Zhang et al. (2024) method employs a
 120 dual-student knowledge distillation framework, enhancing the ability to detect structural and logical
 121 anomalies. However, methods based on feature embeddings rely on the size of the memory bank or
 122 the capability of the teacher network. This reliance can lead to excessive memory usage, resulting in
 123 slower inference times, or may limit the model's generalization ability.

124 Reconstruction-based methods Haselmann et al. (2018); Ristea et al. (2022); Zavrtanik et al. (2021b)
 125 span from autoencoders Bergmann et al. (2019b); Zavrtanik et al. (2021a); Chen et al. (2023) and
 126 generative adversarial networks Yan et al. (2022); Duan et al. (2023) to Transformers You et al. (2022);
 127 Yao et al. (2023) and diffusion models Lu et al. (2023); Zhang et al. (2023). Among them, autoencoder
 128 methods rely on accurately reconstructing normal images and inaccurately reconstructing anomalous
 129 ones, detecting anomalies by comparing the reconstruction with the input image. Reconstruction-
 130 based methods are more likely to capture information from the entire image Liu et al. (2020).

131 2.2 QUANTIZED NEURAL NETWORKS.

132 Quantization aims to compress models by reducing the bit precision used to represent parameters
 133 and/or activations Cai et al. (2020). Existing neural network quantization algorithms can be divided
 134 into two categories based on their training strategy: post-training quantization (PTQ) and quantization-
 135 aware training (QAT). PTQ Nagel et al. (2019) refers to quantizing the model after training, without
 136 any fine-tuning or retraining, thus allowing for quick quantization but at the cost of reduced accuracy.
 137 In contrast, QAT Gong et al. (2019); Dong et al. (2020) adopts an online quantization strategy. This
 138 type of method utilizes the whole training dataset during the quantization process. As a result, it has
 139 higher accuracy but limited efficiency.

140 Recently, several studies have explored the integration of quantization techniques into anomaly
 141 detection tasks Sharmila & Nagapadma (2023). For example, Cho (2024); Jena et al. (2024) have
 142 even applied Post-Training Quantization for On-Device Anomaly Detection, striking a balance
 143 between computational efficiency and detection accuracy. However, it is important to clarify that
 144 model quantization is the method in this paper, not the goal.

145 Our method leverages the precision reduction characteristics of PTQ to achieve dimensionality
 146 reduction in weight precision, while also designing a mixed-precision quantization method specifically
 147 tailored for industrial anomaly detection.

148 3 METHOD

149 In this section, we will introduce our approach from three aspects: 1. Model architecture, which
 150 explains the details of the teacher-student model and the autoencoder we use. 2. Model quantization.
 151 How to evaluate the anomaly detection capability of each layer through the feature of the teacher-
 152 student model and determine the bit-width configuration. 3. Model training and fine-tuning methods.
 153 The specific operations for model initialization and fine-tuning.

154 3.1 MODEL ARCHITECTURE

155 Our model consists of a teacher-student model and an autoencoder, as illustrated in Figure 3. The
 156 RAAD process is divided into three steps: 1. Model Initialization: The model is trained on a dataset
 157 containing only normal images. During training, only the weights of the student model and the

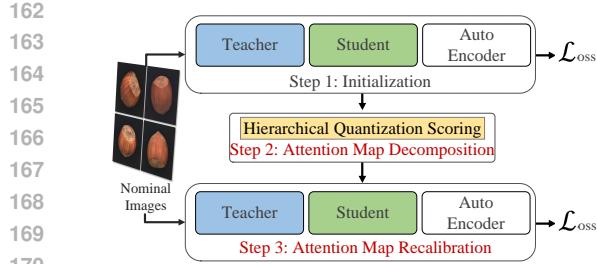


Figure 3: Pipeline of RAAD. Our architecture consists of three components: the teacher-student model and the autoencoder. During training and fine-tuning, we only use normal images. The process is divided into three steps: 1. Initial training of the model, 2. Decomposition of attention map in hierarchical quantitative scoring, detailed in Figure 4. 3. Fine-tuning of attention recalibration.

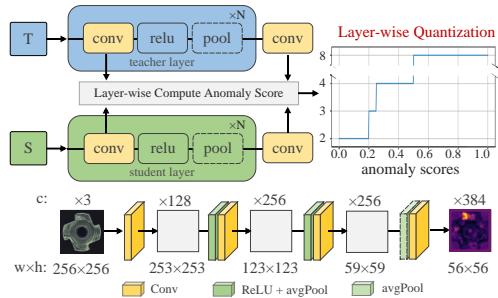


Figure 4: Hierarchical Quantization Scoring (HQS) Module. The teacher and student models are aligned layer by layer, with the anomaly scores calculated using the outputs of their respective convolutional layers. These scores are then converted into quantization bit-widths through a piecewise function. Below are the details of the teacher-student network (PDN).

autoencoder are updated. 2. Hierarchical Quantization Scoring: We first evaluate the anomaly detection capability of each layer in the network. Then, post-training quantization of the model layer by layer. 3. Attention Map Recalibration: Similar to the first step, we fine-tune the student model and the autoencoder. Before training, we employ a pre-trained WideResNet-101 (WRN-101) Zagoruyko & Komodakis (2016) on ImageNet to initialize the teacher model. By minimizing the Mean Squared Error (MSE) between the teacher model and the pre-trained network features. The loss function is as follows:

$$L_{pre} = \frac{1}{cwh} \|(\Psi(I)_{c,w,h} - \mu)(\sigma) - T_{c,w,h}(I)\|_F^2, \quad (1)$$

where I represents an image from the ImageNet, Ψ is a feature extractor composed of the second and third layers of the pre-trained WRN-101 network, and $T(\cdot)$ refers to the teacher model. μ and σ are the mean and standard of the $\Psi(I)$. $\|\cdot\|_F$ denotes the Frobenius norm (i.e., the square root of the sum of the squares of all elements).

We utilize the Patch Description Network (PDN) Batzner et al. (2024) as both the teacher and student model's feature extraction network. Unlike recent anomaly detection methods that commonly employ pre-trained CNN networks, such as DenseNet-201 Li et al. (2023); Huang et al. (2017) and WideResNet-101 Esser et al. (2019); Zagoruyko & Komodakis (2016), the PDN consists of only four convolutional layers. It is fully convolutional and can be applied to images of variable sizes. As the network depth and parameters are reduced, the running time and memory requirements of the model are correspondingly reduced, achieving an inference speed of 113 FPS for a single image on an NVIDIA RTX3090 GPU.

3.2 ATTENTION MAP DECOMPOSITION

The industrial anomaly detection model is trained using a block-level reconstruction strategy Li et al. (2021), followed by layer-by-layer quantization. The quantization formula is as follows:

$$r = S(q - Z), \quad q = \text{clip}(\text{round}(\frac{r}{S} + Z), 0, 255), \quad (2)$$

where r is the pre-quantized floating-point number, with the range (r_{\min}, r_{\max}) . q represents the quantized fixed-point number, with the range (q_{\min}, q_{\max}) . S is the scaling factor, given by:

$$\min_{\mathbf{w}} \mathbb{E} [\Delta \mathbf{z}^{(\ell),T} \mathbf{H}^{(\mathbf{z}^{(\ell)})} \Delta \mathbf{z}^{(\ell)}] = \min_{\mathbf{w}} \mathbb{E} \left[\Delta \mathbf{z}^{(\ell),T} \text{diag} \left(\left(\frac{\partial L}{\partial \mathbf{z}_1^{(\ell)}} \right)^2, \dots, \left(\frac{\partial L}{\partial \mathbf{z}_a^{(\ell)}} \right)^2 \right) \Delta \mathbf{z}^{(\ell)} \right], \quad (3)$$

216 $S = \frac{r_{\max} - r_{\min}}{q_{\max} - q_{\min}}$, and Z is the
 217 zero-point, given by: $Z = \text{clip}(\text{round}(-\frac{r_{\min}}{S}), 0, 255)$, where
 218 clip is a clipping function that
 219 constrains values within the range of 0
 220 to 255.
 221

222 By using the Fisher information
 223 matrix, we measure the inter-block
 224 relationships within layers. Any
 225 second-order error difference is
 226 transformed into blockwise output,
 227 activating the Fisher information
 228 matrix’s diagonal elements to
 229 match each element’s variance. The
 230 weighted objective is:

231
 232 where $\Delta z^{(l)}$ is the feature variation in the l -th layer, which is the difference between quantized and
 233 pre-quantized values. $\mathbb{E}[\cdot]$ represents the expected value. $H^{(z^{(l)})}$ is the Hessian matrix of features in
 234 the l -th layer.
 235

3.3 HIERARCHICAL QUANTIZATION SCORING

236 To further enhance the redistribution of attention in anomaly detection, we propose a Hierarchical
 237 Quantization Scoring (HQS) mechanism, which utilizes mixed-precision techniques to adaptively
 238 adjust the bit-width for each layer. As illustrated in Figure 4, both the teacher and student models
 239 are processed through the HQS module, where corresponding layers of the same depth are aligned,
 240 and their outputs after convolution are used to compute an anomaly score s ($s \in (0, 1)$). This score
 241 measures the alignment of attention between the two models, guiding the reallocation of attention to
 242 defect-prone regions.
 243

244 The bit-width b for the ℓ -th layer ($1 \leq \ell \leq N$) is determined as a function of its anomaly score.
 245 Layers with higher scores, which contribute more to accurate anomaly localization, are assigned
 246 greater precision, allowing deeper, semantically rich layers to focus on subtle anomalies. In contrast,
 247 shallower layers, which often capture redundant or noisy information, are more heavily compressed.
 248 This dynamic bit-width allocation leads to a more effective redistribution of attention, ensuring
 249 that the model emphasizes critical regions in the attention maps. The bit width b for the ℓ -th layer
 250 ($1 \leq \ell \leq N$) is defined as:
 251

$$b^{(\ell)} = \phi(\text{anomaly_score}) = \phi \left((c^{(\ell)} w^{(\ell)} h^{(\ell)})^{-1} \sum_c \left\| T_c^{(\ell)}(i) - S_c^{(\ell)}(i) \right\|_F^2 \right), \quad (4)$$

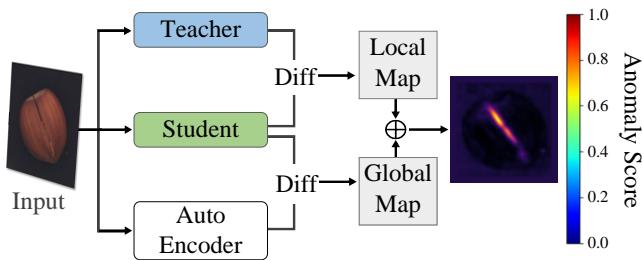
252 where i represents the output of the previous layer, when $\ell = 1$, $i = \text{conv}^{(1)}(I)$, with I being the
 253 image input to the model. $T(i), S(i) \in \mathbb{R}^{c \times w \times h}$, where c, w, h are the number of channels, width,
 254 and height of the output features of the ℓ -th layer, respectively. $\phi(\cdot)$ is a piecewise function that
 255 determines the bit width based on the hierarchical quantization score, and the detailed explanation of
 256 the piecewise function is in the appendix. We chose 2, 3, 4, and 8 bits for mixed precision because
 257 they are most common in practical deployment.
 258

3.4 ATTENTION MAP RECALIBRATION

261 During the training process, the teacher model, student model, and autoencoder are paired with each
 262 other to generate three losses: $L_{t-s}, L_{ae-s}, L_{t-ae}$. Formally, we apply the teacher T , student S , and
 263 autoencoder A to the training image I , with $T(I), S(I), A(I) \in \mathbb{R}^{C \times W \times H}$, and the loss expression
 264 for L_{t-s} is:

$$L_{t-s} = (CWH)^{-1} \sum_c \|T(I)_c - S(I)_c\|_F^2, \quad (5)$$

265 The expressions for L_{ae-s} and L_{t-ae} are similar to L_{t-s} , differing in that $T(I)_c - S(I)_c$ is replaced
 266 with $A(I)_c - S(I)_c$ and $T(I)_c - A(I)_c$, respectively. Note that to confine L_{t-s} to the most relevant



267 Figure 5: The inference process of the models. “Diff” refers
 268 to computing the element-wise squared difference between
 269 two collections of output feature maps and computing its
 270 average across feature maps. To obtain pixel anomaly scores,
 271 the anomaly maps are resized to match the input image using
 272 bilinear interpolation.

270
 271 Table 1: Comparison results with baseline and alternative quantification methods on MVTec-AD and
 272 MVTec-LOCO. $\cdot / \cdot / \cdot$ denotes Anomaly Image-level AU-ROC%/Image-level AP%/PRO%. Grey
 273 indicates the average value. \dagger denotes the unofficial implementation of EfficientAD. The best results
 274 are highlighted in **bold**.

275 Dataset	276 Baseline†	277 LSQ	278 OMPQ	279 RAAD
280 Weight/Activation	281 32/32	282 8/8	283 8/8	284 $\leq 8 / \leq 8$
285 MVTec-AD	286 96.98/97.44/91.38	287 97.21/96.58/86.19	288 98.77/97.97/92.17	289 98.90/97.83/92.92
290 bottle	291 100.0/100.0/94.58	292 99.92/98.32/88.82	293 100.0/100.0/93.92	294 100.0/100.0/93.97
295 cable	296 95.16/ 100.0/94.58	297 95.25/86.34/87.39	298 96.49/90.72/86.80	299 97.71/94.51/88.75
300 capsule	301 94.41/94.32/96.00	302 85.32/87.23/79.17	303 96.33/ 96.40/96.78	304 97.40/94.74/96.70
305 carpet	306 97.43 /98.91/91.11	307 98.17/98.51/90.38	308 98.71/98.84/91.09	309 98.79/98.85/91.99
310 grid	311 99.08/100.0/88.84	312 100.0/100.0/88.85	313 100.0/100.0/88.75	314 99.83/100.0/ 91.00
315 hazelnut	316 99.50/ 100.0/91.40	317 99.14/99.18/83.88	318 99.50/ 100.0/92.44	319 99.78/100.0/92.44
320 leather	321 86.68/93.23/97.09	322 98.30/97.20/97.87	323 99.79/98.92/98.22	324 98.30/93.88/97.87
325 metal_nut	326 98.43/98.76/91.86	327 97.99/98.24/89.09	328 98.77/ 98.91/93.28	329 98.82/98.91/93.66
330 pill	331 96.78/ 98.34/95.87	332 92.03/95.33/84.14	333 97.84/97.18/97.41	334 98.00/97.18/97.44
335 screw	336 93.72/93.44/89.87	337 96.67/96.55/90.96	338 98.48/95.08/ 94.97	339 98.56/96.64/94.45
340 tile	341 100.0/100.0/88.42	342 100.0/100.0/88.42	343 100.0/100.0/88.42	344 100.0/100.0/89.00
345 toothbrush	346 100.0/100.0/94.47	347 100.0/100.0/58.72	348 100.0/100.0/94.47	349 100.0/100.0/94.47
350 transistor	351 99.54/ 100.0/85.43	352 99.29/ 100.0/85.43	353 100.0/100.0/87.31	354 100.0/100.0/87.31
355 wood	356 98.77/97.63/87.41	357 99.03/97.69/87.57	358 98.68/ 98.36/86.65	359 98.50/ 98.36/87.00
360 zipper	361 95.24/94.43/91.67	362 97.05/94.18/92.15	363 96.95/ 95.16/92.05	364 97.84/94.44/97.84
365 MVTec-LOCO	366 84.09/ 78.51/83.32	367 86.26/82.34/81.16	368 89.60/ 88.59/86.19	369 89.75/87.85/86.76
370 breakfast_box	371 77.13/ 66.98/65.21	372 80.54/66.93/65.31	373 80.91/ 91.79/72.00	374 80.94/85.81/72.31
375 juice_bottle	376 96.41/ 96.79/97.28	377 99.62/98.72/98.29	378 97.86/97.41/97.96	379 98.21/97.01/98.07
380 pushpins	381 78.35/ 68.12/88.23	382 78.97/85.03/83.47	383 95.46/90.59/90.86	384 95.85/91.18/91.09
385 screw_bag	386 71.57/ 67.80/76.13	387 74.65/66.25/64.25	388 76.02/67.86/75.41	389 75.95/66.36/77.64
390 splicing_connectors	391 97.03/ 78.51/94.53	392 97.56/94.76/94.50	393 97.75/95.29/ 94.74	394 97.83/98.89/94.72

300
 301 parts of the image, the value of 10% is used for backpropagation in each of the three dimensions of
 302 the mean squared error D , where $D_{c,w,h} = \left(T(I)_{c,w,h} - S(I)_{c,w,h} \right)^2$.

303 The total loss is the weighted summation of the three:

$$304 \quad Loss = \lambda_{t-s} L_{t-s} + \lambda_{ae-s} L_{ae-s} + \lambda_{t-ae} L_{t-ae}. \quad (6)$$

305 As illustrated in Figure 5, the inference process after training involves the teacher-student outputs
 306 a local anomaly map, while the autoencoder-student outputs a global anomaly map. These two
 307 anomaly maps are averaged to calculate a composite anomaly map, with its maximum value used
 308 as the image-level anomaly score, where the 2D anomaly score map $M \in R^{W \times H}$ is given by
 309 $M_{w,h} = C^{-1} \sum_c D_{c,w,h}$, which is the cross-channel average of D , assigning an anomaly score to
 310 each feature vector.

314 4 EXPERIMENT

315 In this section, we demonstrate the effectiveness of RAAD, by comparing the impact of different
 316 quantization methods on the performance of the model and comparing our proposed method with
 317 other advanced IAD methods. Moreover, we provide additional ablation studies.

320 4.1 DATASETS AND EVALUATION METRIC

321 **MVTec-AD** Bergmann et al. (2019a) dataset is a widely recognized anomaly detection benchmark
 322 that encompasses a diverse dataset of 5,354 high-resolution images from various domains. The
 323 data is divided into training and testing sets, with the training set containing 3,629 anomaly-free

Table 2: Comparison with some state-of-the-art methods. \therefore denotes image-level AUROC%, pixel-level AUROC%. Mean Anomaly Image-level AU-ROC% on MVTec-LOCO, and VisA.

Dataset	MVTec	LOCO	VisA
PatchCore	99.2/98.1	80.3	95.1
SimpleNet	99.6 /98.1	77.6	96.1
EfficientAD	96.9/97.1	84.0	95.3
RealNet	99.6 /99.0	-	96.3
RAAD	98.9/97.9	89.7	96.7

Table 3: Ablation studies on our RAAD. “Quant”: The model utilizes post-training quantization, where the weights and activations are quantized to 8-bit precision, followed by fine-tuning. “HQ5”: Using Layer-wise mixed precision Quantization.

Quant	HQS	I-AUC	PRO
✓		96.98	91.38
✓	✓	98.77	92.17
		98.90	92.92

images, ensuring a focus on normal samples. On the other hand, the test set consists of 1,725 images, providing a mix of both normal and abnormal samples for comprehensive evaluation. To aid in the anomaly localization evaluation, pixel-level annotations are provided. **MVTec-LOCO** Bergmann et al. (2022) dataset includes both structural and logical anomalies. It contains 3644 images from five different categories inspired by real-world industrial inspection scenarios. Structural anomalies appear as scratches, dents, or contaminations in the manufactured products. Logical anomalies violate underlying constraints, e.g., a permissible object being present in an invalid location or a required object not being present at all. **VisA** dataset Zou et al. (2022) proposes multi-instance IAD, comprising 10,821 high-resolution images, including 9,621 normal images and 1,200 anomaly images. This dataset is organized into 12 unique object classes. These 12 object classes can be further categorized into three distinct object types: Complex Structures, Multiple Instances, and Single Instances.

Evaluation Metric. We evaluated the performance at both the image-level and the pixel-level, using the Area Under the Receiver Operating Characteristic curve (AUROC) as the primary metric for quantifying image-level (I-AUC) and pixel-level (P-AUC) performance. To ensure a more equitable treatment of anomaly regions of varying sizes, we employed the Per-Region-Overlap (PRO) metric for anomaly segmentation.

4.2 IMPLEMENTATION DETAILS

We pre-train the teacher model using the pre-trained WideResnet101 Zagoruyko & Komodakis (2016) on the ImageNet dataset. Both the teacher and student models use the small version of the Patch Description Network (PDN), with the student model’s output feature dimension twice that of the teacher model, and the autoencoder encodes and decodes the complete image through a bottleneck of 64 latent dimensions. Our hyperparameter settings are as follows. λ_{t-s} , λ_{t-ae} , λ_{ae-s} are all set to 1. During both training and fine-tuning, the teacher model is frozen. The Adam optimizer is used with a learning rate of 0.00001 for the student model and the autoencoder. Experiments are conducted by default on an NVIDIA Geforce GTX 3090Ti with 24 GB of RAM. We train our model for 70k iterations, with a maximum of 60k iterations for fine-tuning training. However, our experiments show that the model often achieves the best performance before reaching the full 60,000 iterations. We compared our quantization method with HQS Esser et al. (2019) and OMPQ Ma et al. (2023). We fix the weight and activation of the first and last layer at 8 bits, following previous works, where the search space is 2/3/4/8 bits. Note that during comparative experiments, we disabled OMPQ’s mixed precision method.

4.3 MAIN RESULTS

We choose the architecture of EfficientAD Batzner et al. (2024) as our benchmark for evaluating our method. EfficientAD is an unsupervised IAD method utilizing a lightweight feature extractor, achieving both low error rates and high computational efficiency. We believe that the trade-off between accuracy and speed is the direction for future IAD development. It is worth noting that EfficientAD has not publicly released its code, so we reproduced their method and refer to it as baseline[†] in Table 1.

378
 379 Figure 6: We conducted evalua-
 380 tions across three datasets, and use
 381 “TRUE”/“FALSE” evaluation metrics.
 382 “TRUE”: Proportion of correct pre-
 383 dictions in the predicted mask M_{pred}
 384 relative to the ground truth (GT). “FALSE”:
 385 Proportion of erroneous predictions in
 386 M_{pred} relative to itself.

Dataset	Baseline	RAAD
MVTec	87.8 / 35.1	89.2 / 33.4
LOCO	63.1 / 55.8	66.0 / 52.9
VisA	93.3 / 75.3	93.6 / 75.2

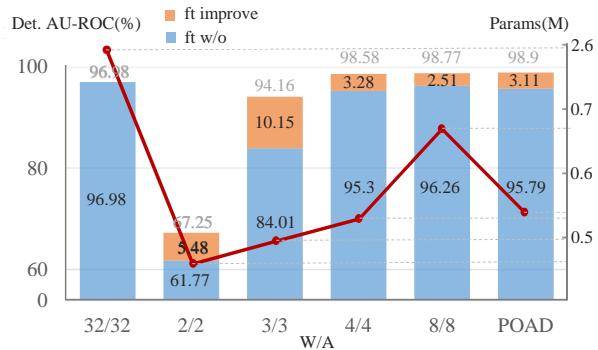
397 To demonstrate the improvement of RAAD on the model’s anomaly detection performance, we
 398 conducted extensive experiments on 32 datasets across three IAD datasets. We applied various
 399 post-training quantization methods to the original framework for comparison, with all quantization
 400 methods performing maximum quantization, i.e., both weights and activations are quantized to
 401 8 bits, ensuring the bit-width is greater than or equal to that used by RAAD quantization. As
 402 shown in Table 1, RAAD achieved average I-AUC scores of 98.8, 89.75, and 96.13 on MVTec-AD,
 403 MVTec-LOCO, and VisA, respectively. “LSQ” and “OMPQ” denote the results when our method’s
 404 quantization is replaced with the quantization methods from Esser et al. (2019) and Ma et al. (2023),
 405 respectively, both of which are representative post-training quantization methods. RAAD consistently
 406 outperforms baselines, emphasizing its effectiveness in both image-level anomaly detection and
 407 pixel-level anomaly localization. This demonstrates the framework’s adaptability to diverse anomaly
 408 characteristics.

411 In addition, we also compared RAAD with several competitive methods across multiple datasets
 412 using various evaluation metrics, as shown in Table 2. We compare RAAD with PatchCore Roth
 413 et al. (2022), GCAD Bergmann et al. (2022), and SimpleNet Liu et al. (2023). Besides EfficientAD,
 414 the results of the other methods are from “paper with code”. RAAD’s average I-AUC score across
 415 the three datasets is 95.12, which is 5.37, 7.95, 6.75, and 2.97 higher than the other methods,
 416 demonstrating the best overall anomaly detection performance, proving RAAD achieves powerful
 417 image-level detection and pixel-level anomaly localization.

418 4.4 EMPIRICAL STUDIES

420 **Effectiveness of different components.** We investigate the effectiveness of each component of
 421 RAAD in Table 3. We set the baseline as the performance of EfficientAD-S on MVTec AD, comparing
 422 the Mean I-AUC and Mean PRO. Utilizing post-training quantization leads to improvements of
 423 1.79% and 0.79%, respectively. Incorporating Hierarchical Quantization Scoring (HQS) results in
 424 improvements of 0.13% and 0.75% respectively, indicating that the introduction of HQS can further
 425 enhance the model performance. Quantization alone improves performance by mitigating attention
 426 bias, but HQS significantly optimizes attention distribution, emphasizing its critical role in refining
 427 anomaly sensitivity.

428 **Prove the quantization effect.** In order to better prove the ability of quantification to solve problems,
 429 we added two metrics. In Table 6, we evaluate three benchmark datasets (MVTec, VisA, and LOCO).
 430 The results of “Vanilla” (original EfficientAD) and “RAAD” (our method) are reported as all class
 431 averaged values. “TRUE”: Proportion of correct predictions in the predicted mask M_{pred} relative
 432 to the ground truth (GT), formulated as: $TRUE = \frac{M_{pred} \cap GT}{GT}$. A higher TRUE indicates better
 433 alignment between predictions and GT. “FALSE”: Proportion of erroneous predictions in M_{pred}
 434 relative to itself, formulated as: $FALSE = \frac{M_{pred} - M_{pred} \cap GT}{M_{pred}}$. A lower FALSE signifies higher
 435 precision in predictions. Analysis table, TRUE: In MVTec/VisA, Minimal differences before/after



394 Figure 7: Quantizing and fine-tuning models using varying
 395 bit-width. Constructing stacked bar charts with I-AUC
 396 corresponds to the left. Plotting line graphs with model
 397 parameters corresponding to the right.

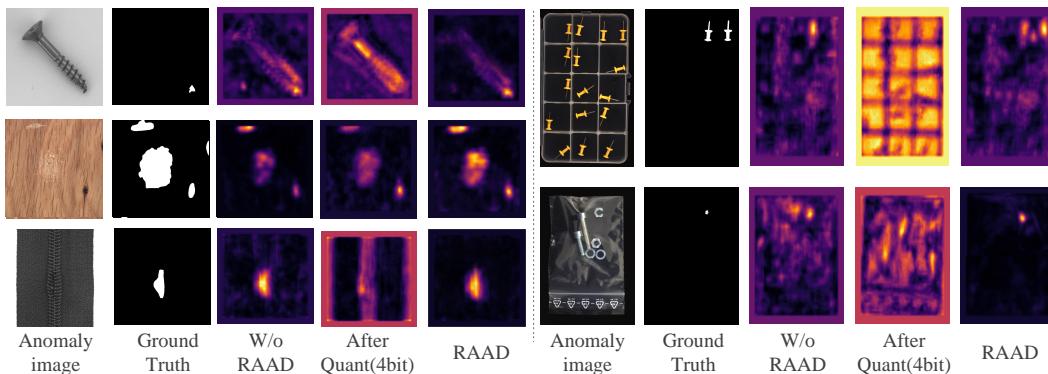


Figure 8: Qualitative results of RAAD on the MVTec-AD dataset Bergmann et al. (2019a) and MVTec-AD LOCO Bergmann et al. (2022). Within each group, from left to right, are the abnormal image, true value, baseline method prediction, quantized model prediction, RAAD prediction.

quantization. In LOCO, Significant TRUE improvement after quantization, attributed to weaker baseline performance allowing larger GT coverage. Similar trends were observed in underperforming MVTec/VisA classes. FALSE: In MVTec, increased FALSE post-quantization, indicating broader attention coverage. In VisA, Slight FALSE decrease due to poor baseline performance.

Quantization and Hierarchical Quantization Scoring (HQS). Figure 7 illustrates the effect of varying quantization bit-widths on model performance, including fine-tuning results for different configurations. Constructing stacked bar charts with I-AUC corresponds to the left vertical axis. Plotting line graphs with model parameters corresponding to the right vertical axis. It can be seen that wider bit-widths better preserve the original model performance, while lower bit-widths increase precision loss, posing challenges for fine-tuning to recover dimensionality. However, it is also observed that fine-tuning is crucial for improving the performance of the quantized model, especially when using lower bit-width quantization, which can lead to more significant improvement. Notably, after applying the HQS method, where the quantization width does not exceed 8 bits, the performance before fine-tuning is slightly lower than that of 8-bit quantization. However, after fine-tuning, it surpasses 8-bit quantization, demonstrating that a more suitable bit-width can better preserve key correspondences. At the same time, the line chart in Figure 7 visualizes the model parameters. Compared to other results with fixed bit widths, RAAD achieves higher accuracy with fewer parameters.

Qualitative Results Figure 8 presents the qualitative results of RAAD on the MVTec-AD and MVTec LOCO datasets. We have visualized the anomaly maps at different stages. Within each group, from left to right, are the anomaly image, ground-truth, predicted anomaly score from EfficientAD-S Batzner et al. (2024), predicted anomaly score from after model quantization (with the quantization bit set to 4-bit to highlight the differences), and the anomaly maps generated by RAAD. It is evident that the anomaly maps after quantization exhibit significant diffusion. The anomaly maps produced by RAAD are more accurate than the baseline, with lower anomaly probabilities in the normal regions.

5 CONCLUSION

Unsupervised IAD methods generally suffer from intrinsic bias in normal samples, which results in misaligned attention. This bias causes models to focus on variable regions while overlooking potential defects in invariant areas. In response, to this work, we propose RAAD (Recalibrating Attention of Industrial Anomaly Detection), a comprehensive framework that decomposes and recalibrates attention maps through a two-stage quantization process. By employing the Hierarchical Quantization Scoring (HQS) mechanism, RAAD optimally redistributes computational resources to enhance defect sensitivity. Qualitative and quantitative experiments show that our method can allocate the model attention properly, breaking the bias of unsupervised IAD, and achieving effective attention redistribution.

486

6 ETHICS STATEMENT

488 This work adheres to the ICLR Code of Ethics. In this study, no human subjects or animal experimen-
 489 tation was involved. All datasets used, including MVTec-AD Bergmann et al. (2019a), MVTec-LOCO
 490 Bergmann et al. (2022) and VisA dataset Zou et al. (2022), were sourced in compliance with relevant
 491 usage guidelines, ensuring no violation of privacy. We have taken care to avoid any biases or dis-
 492 criminatory outcomes in our research process. No personally identifiable information was used, and
 493 no experiments were conducted that could raise privacy or security concerns. We are committed to
 494 maintaining transparency and integrity throughout the research process.

496

7 REPRODUCIBILITY STATEMENT

498 We have made every effort to ensure that the results presented in this paper are reproducible. All code
 499 and datasets have been made publicly available in an anonymous repository to facilitate replication
 500 and verification. The experimental setup, including training steps, model configurations, and hardware
 501 details, is described in detail in the paper. We have also provided a full description of HQS, to assist
 502 others in reproducing our experiments.

504

REFERENCES

506 Kilian Batzner, Lars Heckler, and Rebecca König. Efficientad: Accurate visual anomaly detection at
 507 millisecond-level latencies. In *Proceedings of the IEEE/CVF Winter Conference on Applications*
 508 of *Computer Vision*, pp. 128–138, 2024.

510 Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtect ad — a comprehensive
 511 real-world dataset for unsupervised anomaly detection. In *2019 IEEE/CVF Conference on Com-
 512 puter Vision and Pattern Recognition (CVPR)*, Jun 2019a. doi: 10.1109/cvpr.2019.00982. URL
 513 <http://dx.doi.org/10.1109/cvpr.2019.00982>.

514 Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving
 515 unsupervised defect segmentation by applying structural similarity to autoencoders. In *Proceedings*
 516 of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics
 517 Theory and Applications, Jan 2019b. doi: 10.5220/0007364503720380. URL <http://dx.doi.org/10.5220/0007364503720380>.

520 Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-
 521 teacher anomaly detection with discriminative latent embeddings. In *2020 IEEE/CVF Conference*
 522 on *Computer Vision and Pattern Recognition (CVPR)*, Jun 2020. doi: 10.1109/cvpr42600.2020.
 523 00424. URL <http://dx.doi.org/10.1109/cvpr42600.2020.00424>.

524 Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. Beyond
 525 dents and scratches: Logical constraints in unsupervised anomaly detection and localization.
 526 *International Journal of Computer Vision*, 130(4):947–969, 2022.

528 Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, MichaelW. Mahoney, and Kurt Keutzer. Zeroq:
 529 A novel zero shot quantization framework. *Cornell University - arXiv*, Cornell University - arXiv,
 530 Jan 2020.

531 Ruitao Chen, Guoyang Xie, Jiaqi Liu, Jinbao Wang, Ziqi Luo, Jinfan Wang, and Feng Zheng. Easynet:
 532 An easy network for 3d industrial anomaly detection. In *Proceedings of the 31st ACM International*
 533 *Conference on Multimedia*, pp. 7038–7046, 2023.

535 Hyuntae Cho. Cnn-based autoencoder and post-training quantization for on-device anomaly detection
 536 of cartesian coordinate robots. In *2024 IEEE 14th Annual Computing and Communication*
 537 *Workshop and Conference (CCWC)*, pp. 0662–0666. IEEE, 2024.

538 Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences.
 539 *Cornell University - arXiv*, Cornell University - arXiv, May 2020.

540 Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding.
 541 In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2022.
 542 doi: 10.1109/cvpr52688.2022.00951. URL <http://dx.doi.org/10.1109/cvpr52688.2022.00951>.

544 Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W Mahoney, and Kurt Keutzer.
 545 Hawq-v2: Hessian aware trace-weighted quantization of neural networks. *Advances in neural*
 546 *information processing systems*, 33:18518–18529, 2020.

548 Yuxuan Duan, Yan Hong, Li Niu, and Liqing Zhang. Few-shot defect image generation via defect-
 549 aware feature manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 550 volume 37, pp. 571–578, 2023.

551 Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019.

554 Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazen Lin, Fengwei Yu, and
 555 Junjie Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks.
 556 In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4852–4861,
 557 2019.

558 Matthias Haselmann, Dieter P. Gruber, and Paul Tabatabai. Anomaly detection using deep learning
 559 based image completion. In *2018 17th IEEE International Conference on Machine Learning and*
 560 *Applications (ICMLA)*, Dec 2018. doi: 10.1109/icmla.2018.00201. URL <http://dx.doi.org/10.1109/icmla.2018.00201>.

563 Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected
 564 convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern*
 565 *recognition*, pp. 4700–4708, 2017.

566 Sushovan Jena, Arya Pulkit, Kajal Singh, Anoushka Banerjee, Sharad Joshi, Ananth Ganesh, Dinesh
 567 Singh, and Arnav Bhavsar. Unified anomaly detection methods on edge device using knowledge
 568 distillation and quantization. *arXiv preprint arXiv:2407.02968*, 2024.

570 Donghwoon Kwon, Hyunjoo Kim, Jinoh Kim, Sang C Suh, Ikkyun Kim, and Kuinam J Kim. A
 571 survey of deep learning-based network anomaly detection. *Cluster Computing*, 22:949–961, 2019.

573 Hanxi Li, Jianfei Hu, Bo Li, Hao Chen, Yongbin Zheng, and Chunhua Shen. Target before shooting:
 574 Accurate anomaly detection and localization under one millisecond via cascade patch retrieval.
 575 *arXiv preprint arXiv:2308.06748*, 2023.

576 Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fei Yu, Wei Wang, and Shi Gu.
 577 Brecq: Pushing the limit of post-training quantization by block reconstruction. *Cornell University*
 578 - *arXiv, Cornell University - arXiv*, Feb 2021.

579 Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyan Wu, Bir Bhanu, Richard J.
 580 Radke, and Octavia Camps. Towards visually explaining variational autoencoders. In *2020*
 581 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2020. doi: 10.
 582 1109/cvpr42600.2020.00867. URL <http://dx.doi.org/10.1109/cvpr42600.2020.00867>.

585 Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. Simplenet: A simple network for image
 586 anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer*
 587 *Vision and Pattern Recognition*, pp. 20402–20411, 2023.

588 Fanbin Lu, Xufeng Yao, Chi-Wing Fu, and Jiaya Jia. Removing anomalies as noises for industrial
 589 defect localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,
 590 pp. 16166–16175, 2023.

592 Yuexiao Ma, Taisong Jin, Xiawu Zheng, Yan Wang, Huixia Li, Yongjian Wu, Guannan Jiang, Wei
 593 Zhang, and Rongrong Ji. Ompq: Orthogonal mixed precision quantization. In *Proceedings of the*
AAAI conference on artificial intelligence, volume 37, pp. 9029–9037, 2023.

594 Markus Nagel, Mart Van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization
 595 through weight equalization and bias correction. In *2019 IEEE/CVF International Conference*
 596 *on Computer Vision (ICCV)*, Oct 2019. doi: 10.1109/iccv.2019.00141. URL <http://dx.doi.org/10.1109/iccv.2019.00141>.

598 Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *International Conference on Machine Learning, International Conference on Machine Learning*, Jul
 600 2015.

602 Nicolae-Catalin Ristea, Neelu Madan, Radu Tudor Ionescu, Kamal Nasrollahi, Fahad Shahbaz Khan,
 603 Thomas B. Moeslund, and Mubarak Shah. Self-supervised predictive convolutional attentive
 604 block for anomaly detection. In *2022 IEEE/CVF Conference on Computer Vision and Pattern*
 605 *Recognition (CVPR)*, Jun 2022. doi: 10.1109/cvpr52688.2022.01321. URL <http://dx.doi.org/10.1109/cvpr52688.2022.01321>.

607 Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler.
 608 Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF conference*
 609 *on computer vision and pattern recognition*, pp. 14318–14328, 2022.

611 Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but differnet: Semi-supervised
 612 defect detection with normalizing flows. In *2021 IEEE Winter Conference on Applications*
 613 *of Computer Vision (WACV)*, Jan 2021. doi: 10.1109/wacv48630.2021.00195. URL <http://dx.doi.org/10.1109/wacv48630.2021.00195>.

615 Lukas Ruff, Jacob R. Kauffmann, Robert A. Vandermeulen, Gregoire Montavon, Wojciech Samek,
 616 Marius Kloft, Thomas G. Dietterich, and Klaus-Robert Muller. A unifying review of deep and
 617 shallow anomaly detection. *Proceedings of the IEEE*, pp. 756–795, May 2021. doi: 10.1109/jproc.
 618 2021.3052449. URL <http://dx.doi.org/10.1109/jproc.2021.3052449>.

620 BS Sharmila and Rohini Nagapadma. Quantized autoencoder (qae) intrusion detection system for
 621 anomaly detection in resource-constrained iot devices using rt-iot2022 dataset. *Cybersecurity*, 6
 622 (1):41, 2023.

623 Kihyuk Sohn, Chun-Liang Li, Jinsung Yoon, Minho Jin, and Tomas Pfister. Learning and evaluating
 624 representations for deep one-class classification. *Learning, Learning*, Nov 2020.

626 Raghavendra Chalapathy (University of Sydney, Capital Markets Cooperative Research Centre (CM-
 627 CRC)), Sanjay Chawla (Qatar Computing Research Institute (QCRI), and HBKU) HBKU). Deep
 628 learning for anomaly detection: A survey. *arXiv preprint arXiv:2308.06748*, 2019.

630 Xian Tao, Xinyi Gong, Xin Zhang, Shaohua Yan, and Chandranath Adak. Deep learning for unsuper-
 631 vised anomaly localization in industrial images: A survey. *IEEE Transactions on Instrumentation*
 632 *and Measurement*, 71:1–21, 2022.

633 Chengjie Wang, Wenbing Zhu, Bin-Bin Gao, Zhenye Gan, Jiangning Zhang, Zhihao Gu, Shuguang
 634 Qian, Mingang Chen, and Lizhuang Ma. Real-iad: A real-world multi-view dataset for bench-
 635 marking versatile industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on*
 636 *Computer Vision and Pattern Recognition*, pp. 22883–22892, 2024.

637 Guoyang Xie, Jinbao Wang, Jiaqi Liu, Jiayi Lyu, Yong Liu, Chengjie Wang, Feng Zheng, and Yaochu
 639 Jin. Im-iad: Industrial image anomaly detection benchmark in manufacturing. *IEEE Transactions*
 640 *on Cybernetics*, 54(5):2720–2733, 2024. doi: 10.1109/TCYB.2024.3357213.

641 Xudong Yan, Huaidong Zhang, Xuemiao Xu, Xiaowei Hu, and Pheng-Ann Heng. Learning semantic
 642 context from normal samples for unsupervised anomaly detection. *Proceedings of the AAAI*
 643 *Conference on Artificial Intelligence*, pp. 3110–3118, Sep 2022. doi: 10.1609/aaai.v35i4.16420.
 644 URL <http://dx.doi.org/10.1609/aaai.v35i4.16420>.

646 Xincheng Yao, Ruqi Li, Zefeng Qian, Yan Luo, and Chongyang Zhang. Focus the discrepancy:
 647 Intra-and inter-correlation learning for image anomaly detection. In *Proceedings of the IEEE/CVF*
 648 *International Conference on Computer Vision*, pp. 6803–6813, 2023.

648 Zhiyuan You, Lei Cui, Yujun Shen, Kai Yang, Xin Lu, Yu Zheng, and Xinyi Le. A unified model
 649 for multi-class anomaly detection. *Advances in Neural Information Processing Systems*, 35:
 650 4571–4584, 2022.

651 Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British
 652 Machine Vision Conference 2016*, Jan 2016. doi: 10.5244/c.30.87. URL <http://dx.doi.org/10.5244/c.30.87>.

653 Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Dr-AEm – a discriminatively trained re-
 654 construction embedding for surface anomaly detection. *arXiv: Computer Vision and Pattern
 655 Recognition, arXiv: Computer Vision and Pattern Recognition*, Aug 2021a.

656 Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly
 657 detection. *Pattern Recognition*, pp. 107706, Apr 2021b. doi: 10.1016/j.patcog.2020.107706. URL
 658 <http://dx.doi.org/10.1016/j.patcog.2020.107706>.

659 Jie Zhang, Masanori Suganuma, and Takayuki Okatani. Contextual affinity distillation for image
 660 anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of
 661 Computer Vision*, pp. 149–158, 2024.

662 Xinyi Zhang, Naiqi Li, Jiawei Li, Tao Dai, Yong Jiang, and Shu-Tao Xia. Unsupervised surface
 663 anomaly detection with diffusion probabilistic model. In *Proceedings of the IEEE/CVF Interna-
 664 tional Conference on Computer Vision*, pp. 6782–6791, 2023.

665 Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference
 666 self-supervised pre-training for anomaly detection and segmentation. In *European Conference on
 667 Computer Vision*, pp. 392–408. Springer, 2022.

668

A INFERENCE SPEED, GPU MEMORY, AND MODEL PARAMETERS

676 Table 4: Comparing different PTQ meth-
 677 ods and bit widths on the inference
 678 speed, GPU memory, and model param-
 679 eters.

Method	W/A	FPS	MB	Params
baseline	32/32	167	30.70	8.05
LSQ	2/2	147	30	0.50
	3/3	147	30	0.75
	4/4	147	30	1.00
	8/8	147	30	2.01
OMPQ	2/2	112	61	1.32
	3/3	112	61	1.20
	4/4	112	61	1.46
	8/8	112	61	1.74
RAAD	2/2	113	61	1.32
	3/3	113	61	1.20
	4/4	113	61	1.46
	8/8	113	61	1.74

Table 5: Comparison of existing advanced IAD methods and RAAD in inference time, memory usage, and model parameters.

Method	FPS	MEM(MB)	Params(M)
EfficientAD-S	167	30.7	8.05
PatchCore	25	267	68
RD4D	6	18	150
SimpleNet	4	8	72
CPR	125	3226	2.88
RAAD	113	61	1.46

Table 6: The advanced IAD method uses detection strate-
 gies and image-level detection accuracy I-AUC and pixel-
 level accuracy PRO.

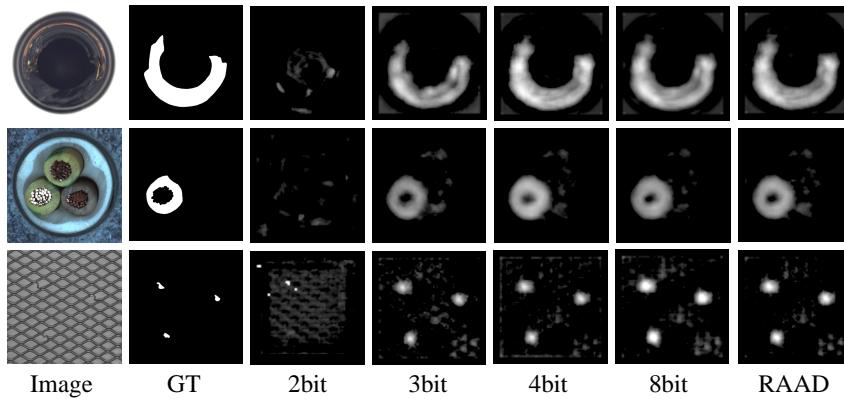
Method	strategy	AU-ROC	AU-PRO
EfficientAD-S	T-S+AutoE	96.98	91.38
PatchCore	Memory Bank	99	93.5
RD4D	T-S	98.5	93.9
SimpleNet	AD Synthesis	98.1	92.9
CPR	Memory Bank	99.7	97.8
RAAD	T-S+AutoE	98.9	92.92

697 In Table 4, the impact of different post-training quantization methods and bit widths on the Inference
 698 Speed, GPU Memory, and Model Parameters is demonstrated.

699
 700 In Table 5 and Table 6, a comparison of the existing state-of-the-art IAD methods with RAAD in
 701 terms of model size and inference time is presented, along with the detection strategies used by
 these methods. We believe that industrial defect methods suitable for edge devices should not only

702 have a low false negatives rate and false positives rate but also possess three key attributes: 1. Fast
 703 inference speed (high FPS), 2. Low GPU memory, and 3. Small model parameters. In Table 5, it
 704 can be observed that the memory bank strategy using feature embeddings achieves the best results in
 705 image-level detection but also incurs a certain memory overhead. Reconstruction-based strategies
 706 perform better at the pixel level.

707



721 Figure 9: Visualization of anomaly maps on MVTec AD with baseline and different PTQ methods.
 722

723 B ADDITIONAL EXPERIMENTS

724 B.1 COMPARISON WITH MODELS USING TRANSFORMER ARCHITECTURE

725 We compared the image-level AU-ROC metrics on the MVTec-AD dataset with methods using
 726 Transformer architectures. It can be observed that Transformer-based methods show significant per-
 727 formance variation across different categories, whereas RAAD is more balanced. This demonstrates
 728 that RAAD has stronger adaptability and generalization across different categories, and its overall
 729 score is superior to other methods.

730 We also tested our method based on the FOD approach. As shown in Table 7, FOD-Quant represents
 731 the quantized FOD model. It is important to clarify that our method is based on a CNN model,
 732 and the attention mentioned in the text is an analogy, not equivalent to the attention mechanism in
 733 Transformers.

734 B.2 USING DIFFERENT NETWORKS

735 In Table 8, the impact of replacing the RAAD feature extraction network PDN with Wide-ResNet of
 736 different depths on model accuracy, model size, and inference time is demonstrated. Since the depth
 737 of Wide-ResNet needs to satisfy $(\text{depth} - 4) \% 6 = 0$, we choose depths of 28, 40, 64, and 100.

738 B.3 FULL RESULT OF THE VISA.

739 We provide the complete results on the VisA datasets in Table 9. As shown in the Table, we evaluated
 740 the anomaly detection presented in the image-level AU-ROC, and pixel-level AP.

741 C QUALITATIVE RESULTS

742 We visualized the results on the MVTec-AD dataset, as shown in Figure 9, demonstrating the impact
 743 of our method using different quantization bit widths on anomaly maps. Moreover, in Figure 10, we
 744 compare the different PTQ methods on MVTec AD and MVTec LOCO-AD. The baseline results are
 745 obtained using EfficientAD-S Batzner et al. (2024).

756
757 Table 7: Compared to the Image-level AU-ROC metric on the MVTec-AD dataset using Transformer
758 architecture methods, “FOD-Quant”: Quantitative evaluation results of the FOD.

Category		InTra (CVPR 2021)	UniAD (NeurIPS 2022)	FOD (ICCV 2023)	FOD (Quant)	RAAD
textures	Carpet	98.8	98.5	99.6	97.9	98.7
	Grid	100	96.5	99.6	99.4	99.8
	Leather	100	98.8	100	99.8	98.3
	Tile	98.2	91.8	100	99.9	100
	Wood	97.5	93.2	98.8	98.2	98.5
	average	98.9	95.7	99.6	99.1	99.0
object	Bottle	100	98.1	100	99.8	100
	Cable	70.3	97.3	98.4	98.5	97.7
	Capsule	86.5	98.5	95.45	95.4	97.4
	Hazelnut	95.7	98.1	100	100	99.7
	Metal_nut	96.9	94.8	99.9	98.1	98.8
	Pill	90.2	95	94.2	95.0	98.0
	Screw	95.7	98.3	94.7	90.7	98.5
	Toothbrush	100	98.4	95.0	93.3	100
	Transistor	95.7	97.9	99.75	99.5	100
	Zipper	99.4	96.8	94.0	98.7	97.8
	average	93.0	97.3	97.1	96.9	98.8
	average all category	95.0	96.8	97.9	92.9	98.9

781
782 Table 9: The full result in the VisA dataset. We report
783 the I-AUC and I-AP.

VisA	Baseline	RAAD
candle	91.72 / 81.90	92.92 / 87.50
capsules	85.15 / 76.03	85.13 / 79.46
cashew	86.15 / 76.03	98.14 / 95.05
chewinggum	98.68 / 96.08	99.64 / 98.99
fryum	97.72 / 95.88	98.61 / 96.91
macaroni1	96.54 / 87.85	98.48 / 86.84
macaroni2	89.26 / 78.26	92.05 / 82.57
pcb1	99.17 / 96.08	99.53 / 97.98
pcb2	98.94 / 95.92	99.28 / 96.08
pcb3	95.71 / 89.11	97.56 / 92.93
pcb4	99.35 / 94.34	99.60 / 96.15
pipe.fryum	99.54 / 98.00	99.72 / 99.00
average	94.73 / 88.79	96.72 / 92.05

784 Table 8: Using Wide-ResNet of different
785 depths on Inference FPS, GPU Memory,
786 model parameters. And evaluate
787 Detection AU-ROC and Segmentation
788 AU-PRO in the capsule category. The
789 units of Mem and params are MB and
790 M, respectively.

Depth	28	40	64	100
FPS	178	125	91	65
Mem	9.6	11	14	18
Params	2.5	2.8	3.6	4.8
I-AUC	89.1	98	92.3	80.9
PRO	85.1	91.3	87.8	83.9

D ALGORITHM

804 The following algorithm explains and proves the effectiveness of quantization and fine-tuning.

805 For example, once an iteration of fine-tuning:

$$\hat{W}_1 = W^+ - \eta \bigtriangledown L(W^+) = W^* + \Delta W - \eta H(W^*) \Delta W \Leftrightarrow \hat{W} - W^* = (I - \eta H(W^*)) \Delta W \quad (7)$$

806 Next, we explain why $(I - \eta H(W^*)) \Delta W$ is better than $H(W^*) \Delta W$. We decompose $H(W^*)$ into
807 its eigenvalues as: $H(W^*) \Delta W = U \Lambda U^T$. 1. When U is large, the model is currently at a sharp

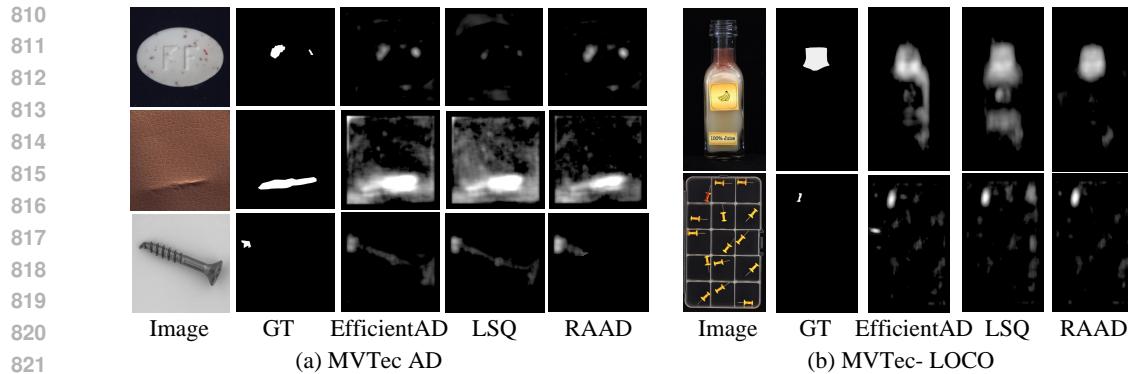


Figure 10: Visualization of anomaly maps on MVTec AD and MVTec-LOCO.

Algorithm 1 Proof of Regularization through Quantification

Input: the initial model parameters W ; after training W^* ; quantization model parameters W^+ ; learning step η ; training loss $L(\cdot)$; quantization $Quant(\cdot)$;

- 1: $W^* = \arg \min_W L(W)$, training convergence $\nabla L(W^*) = 0$;
- 2: $W^+ = Quant(W^*) = W^* + \Delta W$;
- 3: $\hat{W} = W^+ - \eta \nabla L(W^+) \Leftrightarrow \hat{W} - W^+ = -\eta \nabla L(W^+)$;
- 4: $\nabla L(W^+) \xrightarrow{2} \nabla L(W^* + \Delta W)$;
- 5: using Taylor expansion, $\nabla L(W^*) + H(W^*) \Delta W + R(\Delta W) \xrightarrow{1} H(W^*) \Delta W + R(\Delta W) \approx H(W^*) \Delta W$;
- 6: $\because H(W^*) \Delta W \neq 0 \therefore$ After quantization-fine tuning, the gradient is no longer 0.

minimum. $I - \eta H(W^*)$ facilitates escaping from this local minimum. After multiple iterations, the model will tend to approach regions with predominantly smaller Hessian eigenvalues, known as flat minima. 2. When U is small, indicating that the model is at a flat minimum, $I - \eta H(W^*) \approx I$, and the model performance remains stable.

It is well known that the test loss L_{test} is equal to the training loss L_{train} plus the generalization ability L_{gene} . When a model overfits the training data by learning specific details and noise, it loses generalization ability, resulting in a high L_{test} despite a low L_{train} . This occurs because the model has learned features specific to the training data that do not apply to new, unseen data. Quantization, as a regularization technique, enhances model generalization and achieves optimization of model performance during fine-tuning.

E FUTURE WORKS

Unlike model compression, our approach employs model quantization not with the aim of reducing model size, but rather to achieve dimensionality reduction at the precision level. Some industrial anomaly detection efforts currently utilize model quantization during deployment to reduce model size while maintaining certain accuracy levels. For instance, CPR Li et al. (2023) has achieved 1016 FPS using TensorRT on NVIDIA RTX4090 GPU.

In future work, we intend to propose a post-training quantization method specifically designed for the industrial anomaly detection domain, utilizing PyTorch. This method would be applied to post-training models to reduce model size while maintaining or even enhancing model accuracy.

Additionally, we aim to conduct experiments on the more challenging IAD dataset. Currently, most state-of-the-art methods have reached saturation (AUROC exceeding 99%) on mainstream datasets like MVTec, making it difficult to distinguish between methods, which leads to unsatisfactory performance in practical applications. A recently introduced dataset, Real-IAD Wang et al. (2024), has garnered attention due to its large scale, real-world context, and multi-view nature. In future work, we plan to experiment with this dataset.

864 **F LLM USAGE**
865866 Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript.
867 Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring
868 clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing,
869 grammar checking, and enhancing the overall flow of the text.870 It is important to note that the LLM was not involved in the ideation, research methodology, or
871 experimental design. All research concepts, ideas, and analyses were developed and conducted by
872 the authors. The contributions of the LLM were solely focused on improving the linguistic quality of
873 the paper, with no involvement in the scientific content or data analysis.874 The authors take full responsibility for the content of the manuscript, including any text generated or
875 polished by the LLM. We have ensured that the LLM-generated text adheres to ethical guidelines and
876 does not contribute to plagiarism or scientific misconduct.
877878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917