AdaptCLIP: Adapting CLIP for Universal Visual Anomaly Detection

Bin-Bin Gao[1]   Yue Zhou [2,3]   Jiangtao Yan [1]   Yuezhi Cai [2]

Weixi Zhang [2]   Meng Wang [2]   Jun Liu [1]   Yong Liu [1]   Lei Wang [2]   Chengjie Wang[1,4]

[1]Tencent YouTu Lab   [2]Siemens AG   [3]Technical University of Munich   [4]Shanghai Jiao Tong University

Abstract

Universal visual anomaly detection aims to identify anomalies from novel or unseen vision domains without additional fine-tuning, which is critical in open scenarios. Recent studies have demonstrated that pre-trained vision-language models like CLIP exhibit strong generalization with just zero or a few normal images. However, existing methods struggle with designing prompt templates, complex token interactions, or requiring additional fine-tuning, resulting in limited flexibility. In this work, we present a simple yet effective method called AdaptCLIP based on two key insights. First, adaptive visual and textual representations should be learned alternately rather than jointly. Second, comparative learning between query and normal image prompt should incorporate both contextual and aligned residual features, rather than relying solely on residual features. AdaptCLIP treats CLIP models as a foundational service, adding only three simple adapters, visual adapter, textual adapter, and prompt-query adapter, at its input or output ends. AdaptCLIP supports zero-/few-shot generalization across domains and possesses a training-free manner on target domains once trained on a base dataset. Adapt-CLIP achieves state-of-the-art performance on 12 anomaly detection benchmarks from industrial and medical domains, significantly outperforming existing competitive methods. We will make the code and model of AdaptCLIP available at https://github.com/gaobb/AdaptCLIP.

1. Introduction

Universal visual anomaly detection (AD) aims to identify anomaly images and segment anomaly pixels from novel or unseen visual objects after learning a single model on a base or seen dataset. This is a more challenging task as it requires strong generalization when facing cross-domain datasets. Meanwhile, it is a more practical topic as people are more interested in fast adaptability in real-world scenarios, especially in low data regimes (i.e., few-shot and even zero-shot). For example, in medical image diagnosis and industrial visual quality inspection, it is difficult to

collect a large-scale dataset due to inherent scarcity and privacy protection. Recently, developing universal visual AD has attracted increasing attention because existing unsupervised ADs with either separated [8, 25, 34] or unified models [11, 48] perform poorly in unseen objects despite promising performance on seen objects.
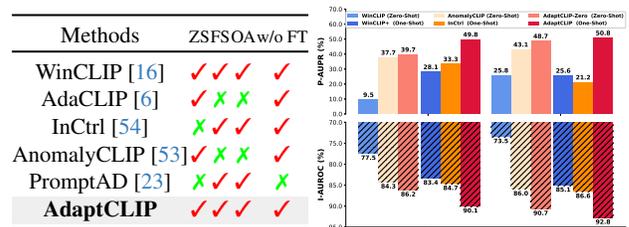

Methods	ZS	FS	OA	w/o FT
WinCLIP [16]	✓	✓	✓	✓
AdaCLIP [6]	✓	✗	✗	✓
InCtrl [54]	✗	✓	✓	✓
AnomalyCLIP [53]	✓	✗	✗	✓
PromptAD [23]	✗	✓	✓	✗
AdaptCLIP	✓	✓	✓	✓

Figure 1. Comparisons of state-of-the-arts and our **AdaptCLIP**. ✓ means satisfied and ✗ means not satisfied. Our method supports zero-/few-shot (ZS and FS) visual AD across different domains without fine-tuning (FT) on the target dataset. It only adds simple adapters at CLIP's input or output ends without complex token interactions, thus preserving CLIP's original ability (OA). The AdaptCLIP using only one normal image prompt achieves the best performance in image-level anomaly classification (I-AUROC) and pixel-level anomaly segmentation (P-AUPR) on 12 AD benchmarks from industrial and medical domains. Moreover, the zero-shot AdaptCLIP is also significantly better than existing zero-shot and even some one-shot approaches. The detailed results are reported in Tabs. 1 and 2. Best viewed in zoom.

To address this fragmentation, recent works have attempted to design universal models to recognize anomalies for unseen objects. They typically build on vision-language models (i.e., CLIP [29]) benefiting from strong generalization. WinCLIP [16] computes anomaly scores on dense patch windows. This brings large computational costs and memory burden, limiting high-resolution input or large pre-trained models. AnomalyCLIP [53] learns class-agnostic prompt embeddings to align patch-wise tokens thus avoiding dense window operations. It further refines vanilla CLIP by concatenating learnable tokens to intermediate layers of CLIP. AdaCLIP [6] further integrates visual knowledge into textual prompt embeddings. However, they may destroy inherent representations of CLIP. Therefore, *we want to explore whether we can achieve the same or even better AD*

performance while maintaining the original ability.

In contrast, humans perceive anomalies when an input significantly deviates from those normal patterns stored in our brains. There is evidence to support this point in neuroscience [31]. PatchCore [34] builds a memory bank storing normal features and PaDiM [7] learns a multivariate Gaussian distribution of normal features. At inference, anomalies are recognized by comparing input features with the memory bank or the learned distribution. However, these methods usually require a certain number of normal images and thus are limited in universal (i.e., open-world) scenarios. Two recent works, i.e., InCtrl [54] and PromptAD [23], have studied how to further improve performance with few-shot normal images. However, InCtrl only considers anomaly classification, while PromptAD needs to learn a new model for each class. Different from them, *we want to comprehensively explore a universal AD model, aiming to detect any anomalies in image-level and pixel-level from cross-domains without any training on target domains.*

Toward this end, we propose a simple but effective universal visual anomaly detection framework, called Adapt-CLIP. *The philosophy of AdaptCLIP is that "less and simpler could be better", and it contains three adapters designed by two key insights :* First, adaptive visual and textual representations should be learned alternately rather than jointly. Second, comparative learning between query and the corresponding normal image prompt should incorporate both contextual and aligned residual features, rather than relying solely on residual features. Our contributions are summarized as follows.

- We propose a simple but effective universal visual anomaly detection framework based on visual-language CLIP models, which is capable of detecting any visual anomalies at image- and pixel-level from cross-domain datasets without any training on target domains.
- We propose visual and textual adapters, and find that they should alternately learn adaptive visual and textual representation guided by the powerful vision-language representations from CLIP models.
- We propose a prompt-query adapter that aims to capture meta-perceptual capabilities between query image and the corresponding normal image prompt, based on their joint distribution of contextual features of the query and the aligned residual features between prompt and query.
- AdaptCLIP outperforms zero- and few-shot AD methods on 8 industrial and 4 medical benchmarks, as shown in Fig. 5. Meanwhile, AdaptCLIP possesses simpler adapters, fewer parameters, and competitive efficiency.

2. Related Works

Unsupervised ADs target to identify anomalies given sufficient normal training images. Most unsupervised

AD methods can be roughly grouped into three categories: embedding-, discrimination-, and reconstruction-based methods. Embedding-based methods, such as PaDiM [7], MDND [32], PatchCore [34], CS-Flow [35] and PyramidFlow [19], assume that offline features extracted from a pre-trained model preserve discriminative information and thus help to separate anomalies from normal samples. Discrimination-based methods, such as Cut-Paste [21], DRAEM [50], and SimpleNet [25], typically convert unsupervised AD to supervised ones by introducing pseudo (synthesized) anomaly samples. Reconstruction-based ADs, such as autoencoder [2, 12, 14, 40], generative adversarial networks [27, 45, 49] and reconstruction networks [24, 33, 51], assume that anomalous regions should not be able to be properly reconstructed and thus result in high reconstruction errors since they do not exist in normal training samples. The recent knowledge distillation [4, 8, 36, 42, 43] or feature reconstruction methods [11, 46, 48, 52] train a student or reconstruction network to match a fixed pre-trained teacher network and achieve a good balance between effectiveness and efficiency. However, all these methods are limited to recognizing anomalies of seen classes but often perform poorly on unseen classes. For a novel scenario, people have to collect sufficient normal images first and then retrain a model. This is inefficient and lacks the rapid adaptability required for practical applications.

Zero-Shot ADs have achieved impressive performance by utilizing large vision-language models, e.g., CLIP [29]. WinCLIP [16] designs two-class textual prompts and introduces multi-scale patch windows for accurate anomaly segmentation. It brings large computational costs and memory burden, limiting high-resolution input or large pre-trained models. AnomalyCLIP [53] learns class-agnostic prompt embeddings to align patch-wise tokens thus avoiding dense window operation. In addition, AnomalyCLIP refines vanilla CLIP representation by appending some learnable tokens to the middle layer of CLIP. Recently, Ada-CLIP [6] and VCP-CLIP [28] utilize similar ideas and further integrate visual knowledge into textual prompt embeddings. We argue that these additional operations make models more complex and may hurt the original capabilities of CLIP. Instead of visual-language models, ACR [20] and MuSc [22] perform zero-shot AD only requiring batch-level and full-shot testing images, but they may be limited in privacy protection scenarios. Different from these methods, we explore whether the same or even better AD performance is achieved while retaining the original ability of CLIP without any information on test data distribution.

Few-Shot ADs mainly pay attention to learning or using only a limited number of normal images, such as TDG [37], RegAD [15], GraphCore [44] and FastRecon [10]. Some works [9, 47] consider another few-shot setting where a
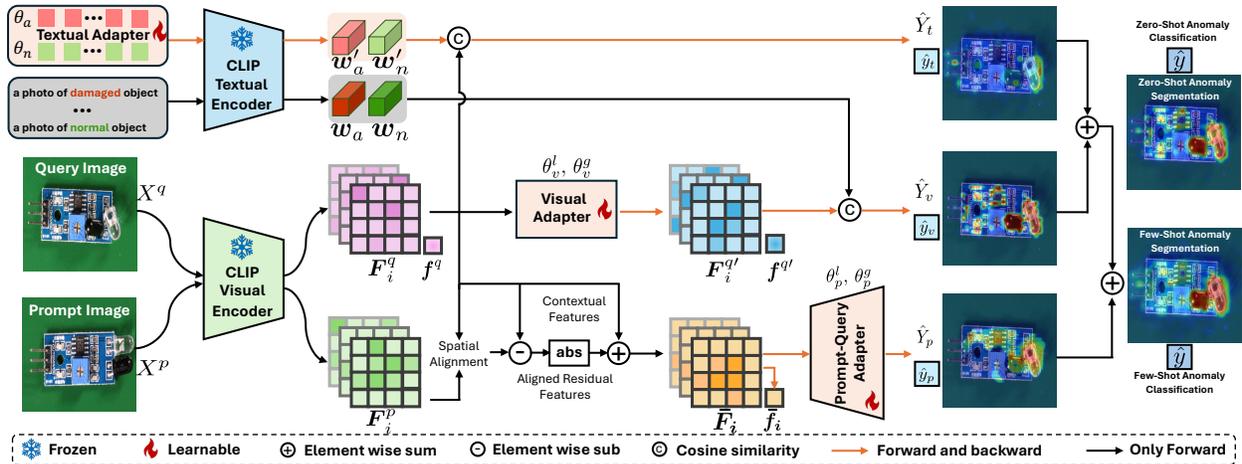
Figure 2. The framework of **AdaptCLIP**, which consists of three pluggable adapters, i.e., visual adapter, textual adapter, and prompt-query adapter. First, the first two adapters alternately learn visual and textual representations for zero-shot anomaly detection (Sec. 3.2). The prompt-query adapter further learns a comparison ability between query image and its corresponding normal prompt for few-shot anomaly detection (Sec. 3.3). Once trained, it can segment any anomalies providing only few-shot and even zero-shot normal image prompts.

limited number of samples is given from anomaly images. The performance of these methods lags behind unsupervised ADs. Recently, few-shot AD performance has been improved significantly by visual-language models. Win-CLIP+ [16] is the first work to apply CLIP models to few-shot AD, which stores normal tokens into a memory bank, then retrieves the nearest token for each query token using cosine similarity, and finally computes an anomaly map using the nearest distance. InCtrl [54] further integrates multi-level information, including patch-level residual maps and image-level residual features, and prior knowledge score using two-class textual prompts, to learn a holistic scoring function for anomaly classification. It does not consider pixel-level anomaly segmentation. PromptAD [23] introduces the concept of explicit anomaly margin, which mitigates the training challenge caused by the absence of anomaly training images. However, it requires re-training models when applied to target datasets. In contrast, we explore jointly optimizing anomaly classification and segmentation in a unified model, which can quickly adapt to novel scenarios only given few-shot normal image prompts, not involving additional re-training.

3. Methods

Problem Formulation: Our objective is to learn a universal AD model that detects any anomalies from diverse domains without any training on target dataset. Thus, a reasonable assumption is that there is a different distribution between training and testing sets. Formally, let $\mathcal{D}_{\text{base}} = \{X_i, Y_i, y_i\}_{i=1}^N$ be a training dataset, that consists of N normal and anomalous images, $X_i \in \mathcal{R}^{h \times w \times 3}$ is the i-th image, and $Y_i \in \mathcal{R}^{h \times w}$ and $y_i = \{0, 1\}$ is the corresponding anomaly mask and anomaly label, with $y_i = 0$ indicates normal and $y_i = 1$ signifies anomaly. The testing set $\mathcal{T}$

may consist of multiple different domains with various objects and anomaly types. Here, we denote the t-th novel domain as $\mathcal{D}_{\text{novel}}^t = \{X_i, Y_i, y_i\}_{i=1}^{N_t}$. Under a few-shot setting, a few normal images $\mathcal{P}_c = \{X_i\}_{i=1}^k$ are randomly drawn from each class of the target domain, where c is the class index and k is typically a small number, e.g., $k = \{1, 2, 4\}$. It is worth noting that $\mathcal{P}_c$ is only available during inference, and cannot be used in any way during training phase.

Overview: As illustrated in Fig. 2, the visual adapter adapts patch and image tokens with fixed two-class textual prompt embeddings. The textual adapter learns two-class prompt embeddings to align with the fixed patch and image tokens. The prompt-query adapter operates in a one-prompt meta-learning manner, leveraging the joint distribution of query context features and the aligned residual features between the prompt and query. In a zero-shot scenario, image-level anomaly score and pixel-level anomaly map can be obtained using textual and visual adapters (Sec. 3.2). In a few-shot scenario, anomaly score and map are derived by integrating predictions from zero-shot and prompt-query adapters (Sec. 3.3). Below we present them in detail.

3.1. Revisiting CLIP for Anomaly Detection

For a query image $X^q \in \mathcal{R}^{h \times w \times 3}$, we feed it to visual encoder $\mathcal{F}(\cdot)$ and obtain local patch tokens $\{F_i^q \in \mathcal{R}^d\}_{i=1}^{hw/p^2}$ and global image token $f^q \in \mathcal{R}^d$, where p is patch size. WinCLIP [16] introduces two-class prompts describing normal and abnormal states. For example, "a photo of a normal object" and "a photo of a damaged object". In practical application, one could design multiple textual descriptions for normal and abnormal states. Feeding these normal and abnormal descriptions to textual encoder $\mathcal{T}(\cdot)$, we can obtain the embeddings of normal $w_n \in \mathcal{R}^d$ and abnormal $w_a \in \mathcal{R}^d$. The pixel-level anomaly map is computed by

measuring the cosine similarities between all patch tokens and the textual embeddings, that is

$$\hat{Y} = \left[\frac{\exp(\langle \boldsymbol{w}_a, \boldsymbol{F}_i^q \rangle)}{\exp(\langle \boldsymbol{w}_a, \boldsymbol{F}_i^q \rangle) + \exp(\langle \boldsymbol{w}_n, \boldsymbol{F}_i^q \rangle)} \right], \quad (1)$$

where $\langle \cdot \rangle$ represents the cosine similarity, and $[\cdot]$ means that all patch-wise prediction scores are rearranged according to their spatial positions and interpolated to the original input resolution. Replacing $\boldsymbol{F}_i^q$ with $\boldsymbol{f}^q$ in Eq. 1, we can obtain an image-level anomaly score $\hat{y}$ for X^q, that is

$$\hat{y} = \frac{\exp(\langle \boldsymbol{w}_a, \boldsymbol{f}^q \rangle)}{\exp(\langle \boldsymbol{w}_a, \boldsymbol{f}^q \rangle) + \exp(\langle \boldsymbol{w}_n, \boldsymbol{f}^q \rangle)}. \quad (2)$$

3.2. AdaptCLIP with Alternating Learning

To adapt CLIP for universal visual anomaly detection, we design visual and textual adapters to alternately learn visual and textual representations. Specifically, the visual adapter learns adaptive visual tokens ($\boldsymbol{F}_i^{q\prime}$ and $\boldsymbol{f}^{q\prime}$) when fixing two-class static textual embeddings ($\boldsymbol{w}_a$ and $\boldsymbol{w}_n$), while the textual adapter learns two-class textual prompt embeddings ($\boldsymbol{w}_a'$ and $\boldsymbol{w}_n'$) when fixing visual tokens ($\boldsymbol{F}_i^q$ and $\boldsymbol{f}^q$).

Visual Adapter adapts vision tokens ($\boldsymbol{F}_i^q$ and $\boldsymbol{f}^q$) with fixed textual embeddings ($\boldsymbol{w}_a$ and $\boldsymbol{w}_n$). It consists of two branches, global and local, which transform global image token and local patch tokens, respectively. Architecturally, the global and local branches are implemented using a simple residual multi-layer perception (MLP), that is

$$\boldsymbol{F}_i^{q\prime} = \boldsymbol{F}_i^q + \text{MLP}(\boldsymbol{F}_i^q; \theta_v^l); \quad \boldsymbol{f}^{q\prime} = \boldsymbol{f}^q + \text{MLP}(\boldsymbol{f}^q; \theta_v^g), \quad (3)$$

where θ_v^l and θ_v^g are learnable parameters. Replacing $\boldsymbol{F}_i^q$ and $\boldsymbol{f}^q$ in Eqs. 1 and 2 with $\boldsymbol{F}_i^{q\prime}$ and $\boldsymbol{f}^{q\prime}$, we obtain pixel-level anomaly map $\hat{Y}_v$ and image-level anomaly score $\hat{y}_v$.

Textual Adapter aims to directly learn two-class prompts $\boldsymbol{\theta}_a, \boldsymbol{\theta}_n \in \mathcal{R}^{r \times d}$ without prompt templates, where $r > 0$ is the length of prompts. We feed them into the frozen textual encoder $\mathcal{T}(\cdot)$ of CLIP, and obtain the corresponding embeddings $\boldsymbol{w}_a'$ and $\boldsymbol{w}_n'$, that is

$$\boldsymbol{w}_a' = \mathcal{T}(\boldsymbol{\theta}_a), \boldsymbol{w}_n' = \mathcal{T}(\boldsymbol{\theta}_n). \quad (4)$$

Then, we replace the static $\boldsymbol{w}_a$ and $\boldsymbol{w}_n$ in Eqs. 1 and 2 with the learnable prompt embeddings $\boldsymbol{w}_a'$ and $\boldsymbol{w}_n'$ to derive local and global anomaly predictions, $\hat{Y}_t$ and $\hat{y}_t$.

Alternating Learning or Joint Learning? A possible question is whether we can learn visual and textual representations jointly. That is, in Eqs. 1 and 2, we simultaneously replace fixed textual embeddings and visual tokens with learnable prompt embeddings ($\boldsymbol{w}_a'$ and $\boldsymbol{w}_n'$) and adaptive visual tokens ($\boldsymbol{F}_i^{q\prime}$ and $\boldsymbol{f}^{q\prime}$). Indeed, this joint alignment mechanism is successful when a large-scale image-text dataset is available. However, we empirically find that it does not work well in the AD field, as shown in Tab. 5

(Lines 3 vs. 4). This is not surprising because the available training data scale is still relatively small and lacks fine-grained textual annotations. The joint learning easily overfits and leads to poor generalization on novel datasets. In contrast, the alternating learning helps us fully utilize the prior knowledge of the CLIP model and thus improve the cross-domain generalization.

3.3. AdaptCLIP with Comparative Learning

Compared to static or learnable textual prompts, using a normal image as a visual prompt is more intuitive. Therefore, we expect to learn a comparison ability between a query image X^q and its corresponding normal prompt X^p, which generalizes well to unseen objects. We find that applying multi-layer features yields better results. For simplicity, we use a single-layer feature in the following.

Spatial Alignment: A simple way is to directly measure their difference by the absolute value of their residual feature, that is $|\boldsymbol{F}_i^q - \boldsymbol{F}_i^p|$, where $\boldsymbol{F}_i^q$ and $\boldsymbol{F}_i^p$ are the patch token of X^q and X^p, respectively. It may fail if the query and prompt images are not aligned in pixel space (e.g., due to rotation and translation). Therefore, we have to align query and prompt tokens for effective comparison. For any query token $\boldsymbol{F}_i^q$, we search the nearest one among all normal tokens $\{\boldsymbol{F}_j^p\}_{j=1}^{hw/p^2}$ using euclidean distance, that is

$$\boldsymbol{F}_i^{p\prime} = \boldsymbol{F}_k^p, k = \arg\min_j \|\boldsymbol{F}_i^q - \boldsymbol{F}_j^p\|_2. \quad (5)$$

Then, we take $\boldsymbol{F}_i^{p\prime}$ as aligned prompt token of $\boldsymbol{F}_i^q$. Now, we can derive the aligned residual feature, i.e., $|\boldsymbol{F}_i^q - \boldsymbol{F}_i^{p\prime}|$.

Joint contextual and aligned residual feature: The aligned residual feature highlights differences or anomaly regions well. However, it may lose contextual information or introduce noise. Intuitively, the contextual information is critical to identify anomalies. Therefore, we aggregate the original query tokens and the aligned residual features by an element-wise sum,

$$\bar{\boldsymbol{F}}_i = \boldsymbol{F}_i^q + |\boldsymbol{F}_i^q - \boldsymbol{F}_i^{p\prime}|. \quad (6)$$

Prompt-Query Adapter: The ultimate goal is to achieve pixel-level anomaly segmentation and image-level anomaly classification. Therefore, we propose a lightweight segmentation head $\mathcal{G}(\cdot; \theta_p^l)$ to learn anomaly segmentation based on the joint feature $\bar{F}$, that is

$$\hat{Y}_p = \mathcal{G}(\bar{F}; \theta_p^l), \quad (7)$$

where θ_p^l is its parameters. Specifically, the segmentation head consists of several transposed convolution blocks following a 1×1 convolution layer. Here, each transposed convolution block upsamples input feature by $2\times$, and it is composed of a 3×3 convolution, a BatchNorm, a ReLU, and a 2×2 deconvolution.

4

Meanwhile, we need to obtain a global image-level prediction. First, we perform average-pooling and max-pooling on the joint feature $\bar{F}$ along the spatial dimension and then take their weighted average as the global image representation. Then, a simple MLP is used to map the global feature to an image-level prediction score, that is

$$\hat{y}_p = \text{MLP}\big((\text{AvgPool}(\bar{F}) + \text{MaxPool}(\bar{F}))/2; \theta_p^g\big), \quad (8)$$

where θ_p^g is the parameter.

```
"""
image_encoder and textual_encoder: frozen CLIP model
visual_adapter: learned visual adapter
theta: learned two-class prompt embeddings
prompt_query_adapter: learned prompt-query adapter
Xq, Xp: query image and the normal image prompt
T:  fixed two-class textual prompts
inference_style: 'zero-shot' or 'few-shot' (default)
"""
# extract fixed visual and textual representations
Fq = image_encoder(Xq), W = textual_encoder(T)
# adapt visual tokens and textual embeddings, Eqs. 3 and 4
Fq_a = visual_adapter(Fq), W_a = textual_encoder(theta)

# AdaptCLIP via alternating learning for zero-shot AD
# compute anomaly map with softmax, Eq. 1
Yv = F.softmax(F.cosine_similarity(W, Fq_a))
Yt = F.softmax(F.cosine_similarity(W_a, Fq))
Y = (Yv + Yt)/2

# AdaptCLIP via comparative learning for few-shot AD
if inference_style == 'few-shot':
    Fp = image_encoder(Xp)
    # spatial alignment, Eq. 5
    Fp_a = Fp[torch.min((Fq - Fp)**2)[1], dim = -1)]
    # joint contextual and aligned residual feature, Eq. 6
    Fq_ = Fq + torch.abs(Fq - Fp_a)
    # perform anomaly segmentation, Eq. 7
    Yp = prompt_query_adapter(Fq_)
    Y = (Yv + Yt + Yp)/3
```

Figure 3. PyTorch pseudocode for the inference of AdaptCLIP.

3.4. Training and Inference

During training, we use cross-entropy loss for global image anomaly classification, and Focal and Dice losses for local patch anomaly segmentation, which is exactly the same as AnomalyCLIP [53]. For zero-shot inference, we average the predictions from visual and textual adapters. For few-shot inference, we fuse (i.e., average) all results from three adapters, i.e., prompt-query, visual and textual adapters, as the final predictions of AdaptCLIP. Fig. 3 shows Py-Torch pseudocode for the pixel-level inference of Adapt-CLIP, where $W = [\boldsymbol{w}_a^T; \boldsymbol{w}_n^T]$, $W_a = [\boldsymbol{w}_a'^T; \boldsymbol{w}_n'^T]$, Fq and Fp means the fixed $\boldsymbol{F}_i^q$ and $\boldsymbol{F}_i^p$, and Fq_a and Fp_a refers the adaptive $\boldsymbol{F}^{q'}$ and the aligned $\boldsymbol{F}^{p'}$. Here, we omit the image-level inference since it can be easily obtained by replacing local patch tokens with a global image token.

4. Experiments

4.1. Experimental Setup

Datasets: We comprehensively evaluate AdaptCLIP on multiple datasets from industrial and medical domains.

For industrial domain, we use MVTec [3], VisA [55], BTAD [26], MVTec3D [5], DTD [1], KSDD [38], MPDD [17], and large-scale Real-IAD [41]. In medical domain, we utilize brain tumor detection datasets, Br35H [13] and COVID-19 [30], as well as gastrointestinal polyp datasets, Kvasir [18] and Endo [39]. A detailed introduction to these datasets can be found in `Appendix`.

Evaluation Metrics: Following previous works, we use AUROC for image-level anomaly classification and AUPR for pixel-level anomaly segmentation in our main paper. Here, we emphasize that AUPR is better for anomaly segmentation, where the imbalance issue is very extreme between normal and anomaly pixels [55]. In `Appendix`, we also provide detailed comparisons using all metrics, including AUROC, AUPR, and F1_{max}.

Training and Testing Protocol: Following Anomaly-CLIP [53], we train AdaptCLIP using the testing data from MVTec and evaluate zero-/few-shot performance on other datasets. As for the evaluation of MVTec, we train Adapt-CLIP using the testing data of VisA. For fair comparison, all models are trained and evaluated using the same protocol.

Competing Methods: We compare our AdaptCLIP with diverse state-of-the-art zero-/few-shot AD methods including zero-shot WinCLIP [16], AnomalyCLIP [53], Ada-CLIP [6], and few-shot WinCLIP+ [16], InCtrl [54] and AnomalyCLIP+. Here, AnomalyCLIP+ is a strong baseline we build on AnomalyCLIP [53] by adding patch-level feature associations like WinCLIP+. More implementation details about AdaptCLIP and competing methods can be found in `Appendix`.

4.2. Comparisons with Zero-/Few-Shot Methods

Tabs. 1 and 2 present comparisons of AdaptCLIP to competing zero-/few-shot methods in image-level anomaly classification and pixel-level anomaly segmentation, respectively, on 8 real-world industrial and 4 medical AD datasets. Note that we only use image-level metrics to evaluate Br35H and Covid due to the lack of pixel-level annotations, and only report the results for Kvasir and Endo using pixel-level metrics since normal images are not included in these two datasets. Below we analyze these results in detail.

Generalization on Industrial Domain: Generally, Adapt-CLIP significantly outperforms all competing models on almost all industrial datasets across three few-shot settings, 1-shot, 2-shot and 4-shot. The performance of all methods generally gets better with more image prompts. Specifically, InCtrl [54] surpasses WinCLIP [16] due to additional fine-tuning on a base training dataset. Anomaly-CLIP [53] further achieves better generalization, which verifies the importance of learning object-agnostic prompts. AdaptCLIP exhibits superior performance, outperforming AnomalyCLIP [53] by a large margin (about 10%+ in pixel AUPR and 2%+ in image AUROC), particularly on chal-

Table 1. **Image-level anomaly classification comparisons with AUROC metric on industrial and medical domains.** The best and second-best results are highlighted in red and blue, respectively. The superscript[†] indicates that the results are our re-implementation with the same training and testing protocol as AnomalyCLIP and our AdaptCLIP. Note that the results are averaged over all categories on each dataset and the full results of each category are presented in `Appendix`, the same below.

Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Br35H	Covid	AVG
0	WinCLIP [16]	90.4	75.5	68.2	69.4	95.1	92.9	61.5	67.0	77.5	80.5	66.4	73.5
	AdaCLIP[†] [6]	90.7	81.7	89.9	76.2	92.7	96.6	64.0	73.3	83.1	96.7	69.4	83.0
	AnomalyCLIP [53]	91.6	82.0	88.3	73.9	93.9	97.8	77.5	69.5	84.3	94.2	77.7	86.0
	AdaptCLIP-Zero	93.5	84.8	91.0	78.6	96.0	98.1	73.6	74.2	86.2	94.8	86.5	90.7
1	WinCLIP+ [16]	93.6±0.4	80.0±2.4	84.4±1.5	74.1±0.4	97.9±0.2	93.8±0.4	69.3±2.9	74.7±0.2	83.4	80.1±2.1	90.1±3.6	85.1
	InCtrl [54]	91.3±0.4	83.2±2.4	88.5±0.4	75.3±1.3	97.9±0.3	92.0±0.9	73.0±2.7	76.6±0.0	84.7	83.9±6.4	89.2±5.3	86.6
	AnomalyCLIP+ [53]	95.2±0.2	86.1±0.7	88.5±0.8	76.7±2.1	98.0±0.2	97.5±0.3	83.4±2.6	78.2±0.0	88.0	90.8±5.1	87.3±2.6	89.1
	AdaptCLIP	94.5±0.5	90.5±1.2	93.4±0.0	81.7±1.5	98.0±0.0	96.9±0.3	83.8±2.2	81.8±0.3	90.1	93.7±2.4	91.8±2.5	92.8
2	WinCLIP+ [16]	94.5±1.0	82.7±1.0	85.8±1.8	74.3±0.3	98.1±0.2	93.8±0.2	69.3±2.3	76.1±0.1	84.3	81.6±0.6	91.8±2.5	86.7
	InCtrl [54]	91.8±0.9	86.3±1.4	86.2±2.0	75.4±0.5	98.3±0.2	91.6±0.9	74.2±1.8	78.5±0.0	85.3	86.1±1.7	89.7±5.1	87.9
	AnomalyCLIP+ [53]	95.4±0.1	87.8±0.5	89.2±1.1	78.3±1.3	98.2±0.1	97.9±0.2	83.4±1.5	78.3±0.0	88.6	91.5±4.0	89.3±2.7	90.4
	AdaptCLIP	95.7±0.6	92.2±0.8	93.4±0.2	82.9±1.1	98.3±0.0	97.2±0.0	84.4±0.7	82.9±0.2	90.8	94.0±1.7	94.9±0.9	94.5
4	WinCLIP+ [16]	95.3±0.1	84.3±0.6	87.8±0.8	75.7±0.3	98.2±0.0	94.0±0.2	71.2±1.6	77.0±0.0	85.4	82.3±0.4	92.9±2.1	87.6
	InCtrl [54]	93.1±0.7	87.8±0.2	67.5±2.4	78.1±1.1	97.7±0.1	91.6±0.9	78.6±2.3	81.8±0.0	84.5	89.1±1.2	91.4±4.1	90.3
	AnomalyCLIP+ [53]	96.1±0.1	88.8±0.5	90.5±1.2	79.2±1.3	98.4±0.1	97.8±0.1	86.3±1.8	78.4±0.0	89.4	91.1±4.4	91.4±3.0	91.3
	AdaptCLIP	96.6±0.3	93.1±0.2	93.3±0.3	84.2±0.6	98.5±0.1	97.0±0.2	86.8±1.1	83.9±0.2	91.7	93.7±2.0	95.8±0.9	94.8

Table 2. **Pixel-level anomaly segmentation comparisons with AUPR metric on industrial and medical domains.**

Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Kvasir	Endo	AVG
0	WinCLIP [16]	18.2	5.4	12.9	5.3	9.8	7.1	14.1	3.3	9.5	27.8	23.8	25.8
	AdaCLIP[†] [6]	39.1	31.0	42.9	37.5	75.2	48.2	25.9	30.5	41.3	36.6	43.7	40.1
	AnomalyCLIP [53]	34.5	21.3	45.5	30.5	62.6	51.9	28.9	26.7	37.7	39.6	46.6	43.1
	AdaptCLIP-Zero	38.3	26.1	41.8	31.4	68.7	58.3	25.3	28.2	39.7	45.3	52.0	48.7
1	WinCLIP+ [16]	38.3±0.8	15.8±0.2	41.3±2.6	18.4±1.1	47.8±0.9	19.2±0.3	29.8±2.0	13.9±0.2	28.1	27.6±2.9	23.6±0.1	25.6
	InCtrl [54]	47.8±1.1	17.7±0.6	44.1±1.4	18.7±0.5	64.3±0.5	26.7±0.7	27.9±2.2	19.1±0.0	33.3	22.1±1.7	20.3±3.7	21.2
	AnomalyCLIP+ [53]	40.8±0.1	24.8±0.9	41.3±1.1	30.6±1.1	67.4±0.4	47.5±0.5	34.2±0.8	27.9±0.0	39.3	46.9±3.9	47.8±4.9	47.4
	AdaptCLIP	53.7±0.9	38.9±0.3	60.6±1.0	40.7±0.6	76.9±0.1	57.8±1.2	33.5±2.5	36.6±0.1	49.8	49.2±4.7	52.4±4.7	50.8
2	WinCLIP+ [16]	39.5±0.6	17.2±0.8	42.8±1.3	19.1±0.8	48.2±0.9	19.0±0.5	30.7±1.1	14.8±0.1	28.9	29.1±0.2	27.6±2.3	28.4
	InCtrl [54]	49.2±0.7	18.5±0.2	44.2±0.8	20.3±0.6	64.4±0.4	26.4±2.5	29.2±1.3	20.1±0.0	34.0	24.9±1.9	24.5±7.5	24.7
	AnomalyCLIP+ [53]	41.5±0.1	26.2±0.7	41.9±0.6	32.4±1.5	68.1±0.2	47.6±0.4	35.3±1.1	28.1±0.0	40.1	47.3±2.9	49.6±4.8	48.5
	AdaptCLIP	55.1±0.5	40.7±0.6	61.0±0.6	42.3±1.1	77.4±0.2	57.5±1.1	35.0±0.7	37.8±0.1	50.9	49.0±4.1	53.1±4.2	51.1
4	WinCLIP+ [16]	41.2±0.9	18.1±1.3	44.0±0.4	19.9±0.6	49.3±0.1	19.1±0.7	32.0±0.2	15.4±0.2	29.9	29.6±0.8	27.7±0.5	28.7
	InCtrl [54]	50.9±0.3	19.2±0.6	44.0±0.2	22.2±1.2	64.9±0.3	26.0±1.4	31.4±0.8	21.0±0.0	35.0	24.7±1.6	22.3±1.0	23.5
	AnomalyCLIP+ [53]	42.4±0.0	27.5±1.1	45.8±3.0	33.4±1.3	68.5±0.2	46.4±0.7	36.8±1.0	28.2±0.0	41.1	45.9±1.5	49.2±3.4	47.6
	AdaptCLIP	57.2±0.8	41.8±0.6	62.3±0.3	44.5±0.3	78.2±0.2	56.4±1.4	37.4±1.1	39.1±0.3	52.1	47.5±2.7	52.2±3.1	49.9

lenging and large-scale datasets like VisA and Real-IAD. This reveals the power of comparative learning based on the joint contextual and aligned residual features for universal anomaly detection. Under zero-shot setting, AdaptCLIP-Zero significantly outperforms SOTA AdaCLIP [6] on anomaly classification, although it shows a slight weakness in industrial anomaly segmentation. However, AdaptCLIP is simpler, requires fewer learnable parameters (0.6M vs. 10.7M in Tab. 3), and generalizes better from the industrial to the medical domain. In addition, our one-shot Adapt-CLIP easily outperforms zero-shot AdaCLIP [6] if only one normal image prompt is available.

Generalization on Medical Domain: Our AdaptCLIP performs strongly on medical AD regardless of zero-shot or few-shot settings when applying the same model trained on an industrial dataset (i.e., MVTec). Surprisingly, it significantly outperforms SOTA AdaCLIP on image anomaly classification (i.e., 6.3% in AUROC) and pixel anomaly segmentation (i.e., 8.6% in AUPR). Notably, our approach still works even when replacing normal image prompts with anomaly images. This is meaningful for some special datasets that don't contain any normal images, such as

Table 3. Complexity and efficiency comparisons.

Shots	Methods	CLIP Models	Input Size	# Params (M)	Inf.Time (ms)
0	WinCLIP [16]	ViT-B-16+240	240×240	208.4 + 0.0	201.3
		ViT-B-16+240	512×512	208.4 + 0.0	3912.6
	AdaCLIP [6]	ViT-L/14@336px	518×518	428.8 + 10.7	212.0
	AnomalyCLIP [53]	ViT-L/14@336px	518×518	427.9 + 5.6	154.9
	AdaptCLIP-Zero	ViT-B-16+240	512×512	208.4 + 0.4	49.9
		ViT-L/14@336px	518×518	427.9 + 0.6	162.2
1	WinCLIP+ [16]	ViT-B-16+240	240×240	208.4 + 0.0	339.5
		ViT-B-16+240	512×512	208.4 + 0.0	7434.9
	InCtrl [54]	ViT-B-16+240	240×240	208.4 + 0.3	337.0
	AnomalyCLIP+ [53]	ViT-L/14@336px	518×518	427.9 + 5.6	158.6
	AdaptCLIP	ViT-B-16+240	512×512	208.4 + 1.4	54.0
		ViT-L/14@336px	518×518	427.9 + 1.8	168.2

Kvasir and Endo. Here, this success is mainly due to the proposed spatial alignment mechanism, as well as a strong prior assumption that anomaly pixels are mostly sparse.

Efficiency Comparison: We measure complexity and efficiency by the number of parameters and the forward inference time, as shown in Tab. 3. The evaluation is performed on one V100 GPU with batch size 32. The number of parameters of AdaCLIP and AnomalyCLIP is 17 times and 9 times that of our AdaptCLIP, respectively. Compared to SOTA, AdaptCLIP achieves competitive inference time yet better AD performance. When extending from zero-shot to one-shot, AnomalyCLIP+ and our AdaptCLIP require almost no additional inference time, unlike earlier WinCLIP.

Qualitative Results: Fig. 4 shows some selected visualizations from industrial and medical testing images using AdaptCLIP. Generally, few-shot normal image prompts help AdaptCLIP segment anomalies more accurately and produce fewer false positives than in a zero-shot manner.

4.3. Comparisons with Many-Shot Methods

In Tab. 4, we compare few-shot AdaptCLIP with many-shot and full-shot unified AD models. It can be seen that AdaptCLIP is better than the early many-shot methods, RegAD [15], and comparable to the latest PromptAD [23]. It is worth noting that PromptAD [23] requires re-training with few-shot normal images while our method remains training-free on target domains. Furthermore, our method outperforms full-shot methods, such as SimpleNet [25] and UniAD [48], and is also competitive with the latest OneNIP [11]. In short, our method has shown excellent performance, especially in the open-world scenario for universal anomaly detection, although there is still some small gap compared to state-of-the-art full-shot methods.

4.4. Ablation Studies

To demonstrate the effectiveness of the proposed three adapters in AdaptCLIP, TA: Texual Adapter, VA: Visual Adapter, and PQA: Prompt-Query Adapter, and two main insights, alternating learning, and comparative learning based on the joint contextual and aligned residual feature, we conduct experiments on MVTec and VisA, and report results in Tab. 5.

Table 4. Comparisons of image-level anomaly classification and pixel-level anomaly segmentation (using AUROC/AUPR metric, and the same as below) with many-shot and full-shot methods.

Methods	Shots	MVTec	VisA	BTAD
AdaptCLIP	1	94.5 / 53.7	90.5 / 38.9	93.4 / 60.6
	4	96.6 / 57.2	93.1 / 41.8	93.3 / 62.3
RegAD [15]	8	91.2 / 51.1	79.7 / 28.6	90.7 / 40.5
PromptAD [23]	4	96.6 / 52.9	89.1 / 31.5	-
SimpleNet [25]	full	78.2 / 24.8	89.2 / 33.1	90.3 / 36.2
UniAD [48]	full	96.5 / 44.7	90.8 / 33.6	92.2 / 50.9
OneNIP [11]	full	97.9 / 63.7	92.5 / 43.3	92.6 / 56.8

Table 5. Ablation studies about different components.

No.	Methods	Shots	TA	VA	PQA	MVTec	VisA
0		0	✗	✗	✗	91.1 / 33.0	82.1 / 18.0
1	baselines	0	✓	✗	✗	92.2 / 31.4	82.9 / 19.7
2		0	✗	✓	✗	90.5 / 39.4	81.0 / 22.1
3	joint	0	✓	✓	✗	89.3 / 36.2	81.6 / 21.5
4	alternating	0	✓	✓	✗	93.5 / 38.3	84.8 / 26.1
5	w/o context	1	✗	✗	✓	62.6 / 7.0	85.3 / 28.7
6	w context	1	✗	✗	✓	88.1 / 50.2	88.9 / 38.1
7	AdaptCLIP	1	✓	✓	✓	94.2 / 52.5	92.0 / 38.8

Simple but effective baselines. The first baseline is the naive CLIP (Line 0), and it is simple and effective for zero-shot anomaly detection only using two-class textual prompts. However, it is still weak in pixel-level anomaly segmentation. The individual textual adapter and visual adapter are two addtional baselines. Specifically, the textual adapter can be seen as an extreme simplification of AnomalyCLIP [53], removing the textual prompt template and textual prompt tuning. The simple textual adapter performs better than the original AnomalyCLIP and naive CLIP in anomaly classification, although it is slightly inferior in anomaly segmentation (Lines 0 vs. 1). The visual adapter learns adaptive local patch tokens and global image tokens to align textual representations from CLIP in both patch and image levels. This significantly improves pixel-level anomaly segmentation (Lines 0 vs. 2).

Alternating learning is better than joint learning. We explore the impact of alternating learning and joint learning strategies on AdaptCLIP's performance. Alternating learning adapts visual or textual representations independently, while joint learning optimizes both representations simultaneously. As shown (Lines 3 vs. 4) in Tab. 5, the alternating learning strategy significantly enhances the performance of AdaptCLIP compared to joint learning. Alternating learning not only fully leverages the strong prior guidance of CLIP's visual and textual representations but also mitigates the risk of over-fitting due to fine-tuning on a small training dataset. Additionally, we observe that the visual adapter
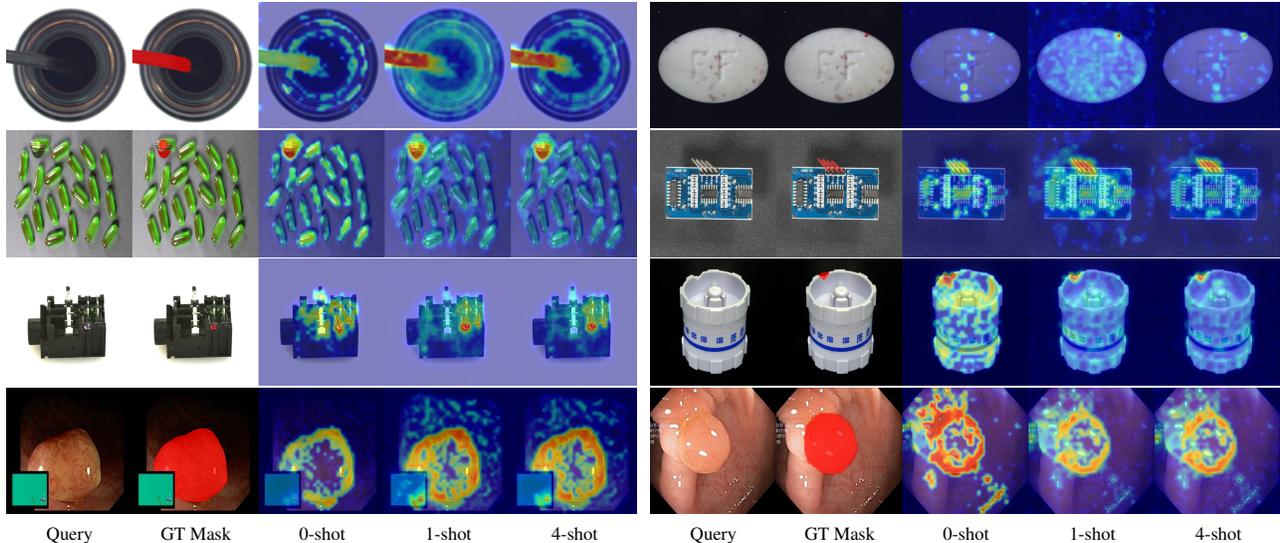
7

Figure 4. Qualitative comparisons of our AdaptCLIP with different prompt numbers on MVTec, VisA, Real-IAD, Kvasir and Endo. More qualitative results of AdaptCLIP can be found in `Appendix`. Best viewed in color and zoom.

alone excels in anomaly segmentation (Line 2), whereas the textual adapter alone performs better in anomaly classification (Line 1). By integrating the alternating learning into visual and textual adapters, AdaptCLIP generally achieves superior anomaly detection performance (Line 4).

The joint of contextual information and aligned residual features performs better than residual features alone. The aligned residual feature captures the distinctions between anomalous features and their corresponding normal counterparts. It effectively eliminates features related to individual objects and may improve generalization. However, we realize that isolated residual features may lose contextual information about visual objects, resulting in degraded model performance or even training failure (Line 5). Therefore, we propose a joint feature learning based on both contextual and aligned residual features, which further significantly boosts the model's performance (Lines 6 vs. 5). This means contextual information is equally important for anomaly identification. Notably, the optimal performance for AdaptCLIP is achieved when all proposed components are integrated (Line 7).

Effects on pre-trained CLIP models. We report zero- and one-shot results of AdaptCLIP using different CLIP models in Tab. 6. It can be seen that a larger pre-trained model always brings better performance, especially in image-level classification. Furthermore, our method equipped with a lightweight model (ViT-B-16+240) makes it possible to achieve competitive anomaly segmentation performance.

5. Conclusion

In this paper, we introduce a universal anomaly detection task, which focuses on generalizing anomaly detection models across domains, such as industrial and medical, and

Table 6. Ablation studies about different pre-trained CLIP models.

CLIP Models	Input Size	Shots	MVTec	VisA
ViT-B-16+240	512×512	0	83.9 / 38.3	75.4 / 19.5
ViT-L/14@336px	518×518	0	93.5 / 38.3	84.8 / 26.1
ViT-B-16+240	512×512	1	92.4 / 52.3	85.2 / 30.3
ViT-L/14@336px	518×518	1	94.2 / 52.5	92.0 / 38.8

in open scenarios, such as zero- or few-shot settings. Once the universal anomaly detection model is trained, it does not need any fine-tuning on the target dataset. Compared with single zero-shot or few-shot AD models, the universal anomaly detection model is more flexible, supporting zero-/few-shot inference via fixed or learnable textual prompts and a few normal image prompts, while providing both image-level and pixel-level anomaly predictions. We propose a universal anomaly detection framework, AdaptCLIP, which alternately learns adaptive visual representations and text prompt embeddings, as well as jointly learns comparisons based on the contextual information of query image and the aligned residual features between the query and the prompt. Extensive experiments on 8 standard industrial and 4 medical datasets show that AdaptCLIP significantly outperforms current competitive models in multiple settings.

Limitation: AdaptCLIP achieves good AD performance only given zero-/few-shot normal image prompts. However, it could cause the model to confuse normal and abnormal instances and finally result in a decreased performance when we provide anomaly images as normal image prompts. Fortunately, normal images are generally relatively easy to obtain in practical applications. In addition, it may work using abnormal images as visual prompts because most of the pixels may be normal even in anomaly images.

References

[1] Toshimichi Aota, Lloyd Teh Tzer Tong, and Takayuki Okatani. Zero-shot versus many-shot: Unsupervised texture anomaly detection. In *WACV*, 2023. 5

[2] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. Generalized denoising auto-encoders as generative models. In *NeurIPS*, 2013. 2

[3] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTec-AD: A comprehensive real-world dataset for unsupervised anomaly detection. In *CVPR*, 2019. 5

[4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed Students: Student-teacher anomaly detection with discriminative latent embeddings. In *CVPR*, 2020. 2

[5] Paul Bergmann., Xin Jin., David Sattlegger., and Carsten Steger. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. In *VISAPP*, 2022. 5

[6] Yunkang Cao, Jiangning Zhang, Luca Frittoli, Yuqi Cheng, Weiming Shen, and Giacomo Boracchi. Adaclip: Adapting clip with hybrid learnable prompts for zero-shot anomaly detection. In *ECCV*, 2024. 1, 2, 5, 6, 7, 4

[7] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. PaDiM: a patch distribution modeling framework for anomaly detection and localization. In *ICPR*, 2021. 2

[8] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In *CVPR*, 2022. 1, 2

[9] Choubo Ding, Guansong Pang, and Chunhua Shen. Catching both gray and black swans: Open-set supervised anomaly detection. In *CVPR*, 2022. 2

[10] Zheng Fang, Xiaoyang Wang, Haocheng Li, Jiejie Liu, Qiugui Hu, and Jimin Xiao. FastRecon: Few-shot industrial anomaly detection via fast feature reconstruction. In *ICCV*, 2023. 2

[11] Bin-Bin Gao. Learning to detect multi-class anomalies with just one normal image prompt. In *ECCV*, 2024. 1, 2, 7

[12] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *ICCV*, 2019. 2

[13] Ahmed Hamada. Br35h: Brain tumor detection 2020, 2020. 5

[14] Jinlei Hou, Yingying Zhang, Qiaoyong Zhong, Di Xie, Shiliang Pu, and Hong Zhou. Divide-and-Assemble: Learning block-wise memory for unsupervised anomaly detection. In *ICCV*, 2021. 2

[15] Chaoqin Huang, Haoyan Guan, Aofan Jiang, Ya Zhang, Michael Spratling, and Yan-Feng Wang. Registration based few-shot anomaly detection. In *ECCV*, 2022. 2, 7

[16] Jongheon Jeong, Yang Zou, Taewan Kim, Dongqing Zhang, Avinash Ravichandran, and Onkar Dabeer. WinCLIP: Zero-/few-shot anomaly classification and segmentation. In *CVPR*, 2023. 1, 2, 3, 5, 6, 7, 4

[17] Stepan Jezek, Martin Jonak, Radim Burget, Pavel Dvorak, and Milos Skotak. Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions. In *ICUMT*, 2021. 5

[18] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *International Conference on Multimedia Modeling*, pages 451–462, 2020. 5

[19] Jiarui Lei, Xiaobo Hu, Yue Wang, and Dong Liu. Pyramid-Flow: High-resolution defect contrastive localization using pyramid normalizing flow. In *CVPR*, 2023. 2

[20] Aodong Li, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan Mandt. Zero-shot anomaly detection via batch normalization. In *NeurIPS*, 2023. 2

[21] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. CutPaste: Self-supervised learning for anomaly detection and localization. In *CVPR*, 2021. 2

[22] Xurui Li, Ziming Huang, Feng Xue, and Yu Zhou. MuSc: Zero-shot industrial anomaly classification and segmentation with mutual scoring of the unlabeled images. In *ICLR*, 2024. 2

[23] Xiaofan Li, Zhizhong Zhang, Xin Tan, Chengwei Chen, Yanyun Qu, Yuan Xie, and Lizhuang Ma. PromptAD: Learning prompts with only normal samples for few-shot anomaly detection. In *CVPR*, 2024. 1, 2, 3, 7

[24] Wenrui Liu, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Diversity-measurable anomaly detection. In *CVPR*, 2023. 2

[25] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. SimpleNet: A simple network for image anomaly detection and localization. In *CVPR*, 2023. 1, 2, 7

[26] Pankaj Mishra, Riccardo Verk, Daniele Fornasier, Claudio Piciarelli, and Gian Luca Foresti. VT-ADL: A vision transformer network for image anomaly detection and localization. In *ISIE*, 2021. 5

[27] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. OC-GAN: One-class novelty detection using GANs with constrained latent representations. In *CVPR*, 2019. 2

[28] Zhen Qu, Xian Tao, Mukesh Prasad, Fei Shen, Zhengtao Zhang, Xinyi Gong, and Guiguang Ding. Vcp-clip: A visual context prompting model for zero-shot anomaly segmentation. In *ECCV*, 2024. 2

[29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2

[30] Tawsifur Rahman, Amith Khandakar, Yazan Qiblawey, Anas Tahir, Serkan Kiranyaz, Saad Bin Abul Kashem, Mohammad Tariqul Islam, Somaya Al Maadeed, Susu M Zughaier, Muhammad Salman Khan, et al. Exploring the effect of image enhancement techniques on covid-19 detection using chest x-ray images. *Computers in biology and medicine*, 2021. 5

[31] Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 1999. 2

[32] Oliver Rippel, Patrick Mertens, and Dorit Merhof. Modeling the distribution of normal data in pretrained deep features for anomaly detection. In *ICPR*, 2021. 2

[33] Nicolae-Cătălin Ristea, Neelu Madan, Radu Tudor Ionescu, Kamal Nasrollahi, Fahad Shahbaz Khan, Thomas B Moeslund, and Mubarak Shah. Self-supervised predictive convolutional attentive block for anomaly detection. In *CVPR*, 2022. 2

[34] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *CVPR*, 2022. 1, 2

[35] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Fully convolutional cross-scale-flows for image-based defect detection. In *WACV*, 2022. 2

[36] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *CVPR*, 2021. 2

[37] Shelly Sheynin, Sagie Benaim, and Lior Wolf. A hierarchical transformation-discriminating generative model for few shot anomaly detection. In *ICCV*, 2021. 2

[38] Domen Tabernik, Samo Šela, Jure Skvarč, and Danijel Skočaj. Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing*, 31(3):759–776, 2020. 5

[39] David Vázquez, Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Antonio M López, Adriana Romero, Michal Drozdzal, Aaron Courville, et al. A benchmark for endoluminal scene segmentation of colonoscopy images. *Journal of healthcare engineering*, 2017, 2017. 5

[40] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008. 2

[41] Chengjie Wang, Wenbing Zhu, Bin-Bin Gao, Zhenye Gan, Jiangning Zhang, Zhihao Gu, Shuguang Qian, Mingang Chen, and Lizhuang Ma. Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection. In *CVPR*, 2024. 5

[42] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. *BMVC*, 2021. 2

[43] Shenzhi Wang, Liwei Wu, Lei Cui, and Yujun Shen. Glancing at the patch: Anomaly localization with global and local feature comparison. In *CVPR*, 2021. 2

[44] Guoyang Xie, Jingbao Wang, Jiaqi Liu, Feng Zheng, and Yaochu Jin. Pushing the limits of fewshot anomaly detection in industry vision: Graphcore. In *ICLR*, 2023. 2

[45] Xudong Yan, Huaidong Zhang, Xuemiao Xu, Xiaowei Hu, and Pheng-Ann Heng. Learning semantic context from normal samples for unsupervised anomaly detection. In *AAAI*, 2021. 2

[46] Xincheng Yao, Ruoqi Li, Zefeng Qian, Yan Luo, and Chongyang Zhang. Focus the Discrepancy: Intra-and inter-correlation learning for image anomaly detection. In *ICCV*, 2023. 2

[47] Xincheng Yao, Ruoqi Li, Jing Zhang, Jun Sun, and Chongyang Zhang. Explicit boundary guided semi-push-pull contrastive learning for supervised anomaly detection. In *CVPR*, 2023. 2

[48] Zhiyuan You, Lei Cui, Yujun Shen, Kai Yang, Xin Lu, Yu Zheng, and Xinyi Le. A unified model for multi-class anomaly detection. In *NeurIPS*, 2022. 1, 2, 7

[49] Muhammad Zaigham Zaheer, Jin-ha Lee, Marcella Astrid, and Seung-Ik Lee. Old is Gold: Redefining the adversarially learned one-class classifier training paradigm. In *CVPR*, 2020. 2

[50] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. DRAEM: A discriminatively trained reconstruction embedding for surface anomaly detection. In *ICCV*, 2021. 2

[51] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *PR*, 112, 2021. 2

[52] Ying Zhao. OmniAL: A unified cnn framework for unsupervised anomaly localization. In *CVPR*, 2023. 2

[53] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. AnomalyCLIP: Object-agnostic prompt learning for zero-shot anomaly detection. In *ICLR*, 2024. 1, 2, 5, 6, 7, 4

[54] Jiawen Zhu and Guansong Pang. Toward generalist anomaly detection via in-context residual learning with few-shot sample prompts. In *CVPR*, 2024. 1, 2, 3, 5, 6, 7, 4

[55] Yang Zou, Jongheon Jeong, Latha Pemula, Dongqing Zhang, and Onkar Dabeer. Spot-the-difference self-supervised pretraining for anomaly detection and segmentation. In *ECCV*, 2022. 5

AdaptCLIP: Adapting CLIP for Universal Visual Anomaly Detection

Supplementary Material

6. Dataset Details

To validate the effectiveness of our method, we conduct comprehensive experiments on 12 public anomaly detection datasets covering two domains, industrial and medical, and three modalities, including photography, radiology, and endoscopy. We only use two test datasets for model pre-training and generalization evaluation on other test datasets, and their relevant information is reported in Tab. 7. Specifically, we train models using the test data from MVTec and evaluate zero-/few-shot performance on other datasets. As for the evaluation of MVTec, we train models using VisA's test data.

It should be noted that Real-IAD is the largest industrial anomaly detection dataset consisting of diverse categories (30 objects) and large-scale images (150k) among the utilized datasets. For the medical domain, we cannot find publicly available 2D medical AD datasets that include both image- and pixel-level annotations simultaneously. Therefore, we only report image-level classification performance on Br35H and Covid, while providing pixel-level anomaly segmentation performance on Kvasir and Endo.

In addition, we note that MPDD and all four medical datasets are pose-agnostic, and KSDD may contain noise. We empirically find that the performance of few-shot AD methods may be limited and sometimes may be worse than zero-shot methods on these datasets. We believe this is a shortcoming of all few-shot normal image prompt-based methods.

7. Implementation Details

We utilize the pre-trained CLIP (ViT-L/14@336) as the default CLIP model and extract local patch tokens from layers $\{6, 12, 18, 24\}$ and global image token from the last layer $\{24\}$. All images are resized to a resolution of 518×518 for training and testing. Regarding the visual and textual adapters, we only use features from the last layer (i.e., 24) of the CLIP visual encoder, while for the prompt-query adapter, we use features from all 4 layers. The visual adapter is a two-layer MLP, whose hidden layer dimension is 1/4 of the input layer, and the output layer dimension remains the same as the input layer. The length r of learnable textual prompt embeddings is set to 12 in the textual adapter. For the prompt-query adapter, the dimension of the first hidden layer is set to 128, and then the dimension of the next layer is halved until the last layer is set to 2 in both the lightweight segmentation head and the global MLP. We train models for 15 epochs with a learning rate of 0.001. All experiments are conducted using PyTorch with a single NVIDIA V100 GPU.

7.1. Competing Methods

For fair comparison, we compare state-of-the-art zero-shot methods, such as WinCLIP [16], AnomalyCLIP [53], and AdaCLIP [6], and few-shot methods, such as WinCLIP+ [16], InCtrl [54], AnomalyCLIP+, and PromptAD [23], with our AdaptCLIP using the same training protocol and few-shot normal image prompts. It is worth noting that the original InCtrl [54] only supports image-level few-shot AD, and we have appropriately extended it to allow pixel-level few-shot AD. In addition, AnomalyCLIP+ is an extension of AnomalyCLIP introducing feature association in WinCLIP+. In Tab. 8, we qualitatively analyze these methods in terms of **capability**, including zero-shot, few-shot, image-level anomaly classification, pixel-level anomaly segmentation, unified or separated models and original CLIP ability, and **complexity**, including pre-training, post-finetuning on target datasets, sliding windows, class names, learnable-prompts and few-shot normal image prompts. We summarize them in detail as follows.

WinCLIP [16] is the first zero-shot anomaly detection method based on a vision-language model, i.e., CLIP. WinCLIP designs two-class textual prompts and introduces multi-scale patch windows for accurate anomaly segmentation. However, it brings large computational costs and memory burden, limiting high-resolution input or large pre-trained models. Note that no official implementation of WinCLIP is available, our results are based on an unofficial implementation.

WinCLIP+ [16] combines language- and visual-guided predictions for better anomaly classification and segmentation. The language-guided prediction is the same as in WinCLIP. For visual-guided prediction, it first simply stores multi-scale features from few-shot normal images into a memory bank, and then measures the anomaly score using the distance or similarity between each query feature and the nearest feature from the memory bank. The final anomaly score is derived by averaging these two scores.

AnomalyCLIP [53] learns object-agnostic text prompts that capture generic normality and abnormality in an image regardless of its foreground objects. However, AnomalyCLIP requires fine-tuning on an auxiliary domain dataset including normal and anomaly images. AnomalyCLIP is a zero-shot anomaly detection method and it is capable of recognizing any anomalies. We use the official model to report performance for anomaly classification and segmentation.

AnomalyCLIP+ is one of our baselines. Note that the original AnomlayCLIP only supports zero-shot anomaly

Table 7. Key statistics of industrial and medical datasets with different attributes. ✓ means satisfied and ✗ means not satified.

Domain	Dataset	Modality	Category	# Classes	Pose-Agnostic	Anomaly Anotations		Train	Test	
						Image-Level	Pixel-Level	# Normal	# Normal	# Anomaly
Industrial	MVTec	Photography	Obj & Texture	15	✗	✓	✓	3,629	467	1,258
	VisA	Photography	Obj	12	✗	✓	✓	8,659	962	1,200
	BTAD	Photography	Obj & Texture	3	✗	✓	✓	1,799	451	290
	MVTec3D	Photography	Obj	10	✗	✓	✓	2,656	249	948
	DTD	Photography	Texture	12	✗	✓	✓	1,200	357	947
	KSDD	Photography	Texture	1	✗	✓	✓	857	286	54
	MPDD	Photography	Obj	6	✓	✓	✓	888	176	282
	Real-IAD	Photography	Obj	30	✗	✓	✓	36,465	63,256	51,329
Medical	Br35H	Radiology (MRI)	Brain	1	✓	✓	✗	0	1500	1500
	Covid	Radiology (X-ray)	Chest	1	✓	✓	✗	0	1,341	219
	Kvasir	Endoscopy	Gastrointestinal tract	1	✓	✗	✓	0	0	1,000
	Endo	Endoscopy	Gastrointestinal tract	1	✓	✗	✓	0	0	200

Table 8. Comprehensive comparisons of state-of-the-art zero-/few-shot AD methods and our AdaptCLIP in terms of capabilities and complexity. ✓ means satisfied and ✗ means not satified.

Methods	Capability						Complexity					
	zero-shot	few-shot	image-cls.	pixel-seg.	unified	ori-ability	pre-training	post-finetuning	sliding-wins	class-names	learnable-prompts	image-prompts
WinCLIP/WinCLIP+ [16]	✓	✓	✓	✓	✓	✓	✗	✗	✓	✓	✗	✓
AdaCLIP [6]	✓	✗	✓	✓	✓	✗	✓	✗	✗	✓	✓	✗
InCtrl [54]	✗	✓	✓	✗	✗	✓	✓	✗	✗	✓	✗	✓
AnomalyCLIP [53]	✓	✗	✓	✓	✓	✗	✓	✗	✗	✗	✓	✗
PromptAD [23]	✗	✓	✓	✓	✗	✓	✓	✓	✗	✓	✓	✓
AdaptCLIP	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓

detection. To achieve few-shot anomaly detection, we introduce feature association based on a memory mechanism, similar to WinCLIP+, to AnomalyCLIP. The final anomaly results are also the average of zero-shot predictions guided by learnable textual prompts and few-shot predictions guided by few-shot normal image prompts.

InCtrl [54] integrates multi-level information, including patch-level residual maps, image-level residual features, and prior knowledge score using two-class textual prompts, to learn a holistic scoring function for anomaly classification. However, it does not consider pixel-level anomaly segmentation. In this paper, we simply use the patch-level residual map as the pixel-level anomaly prediction, which is essentially similar to WinCLIP+. It is worth noting that InCtrl provides multiple models for different shot numbers. We use its official models for extensive evaluation. In addition, since it does not provide a 1-shot model, we use the 2-shot model to evaluate 1-shot performance.

AdaCLIP [6] further integrates visual knowledge from query images into textual prompt embeddings for enhancing the interaction of visual features and textual prompt embeddings. Different from AnomalyCLIP, AdaCLIP trains base models using more auxiliary datasets, including industrial and medical, which is not conducive to cross-domain evaluation. For a fair comparison, we retrain AdaCLIP models using the same training protocol as AnomalyCLIP, and conduct comprehensive evaluations on multiple datasets for zero-shot anomaly detection.

PromptAD [23] introduces an explicit anomaly margin to mitigate the training challenge caused by the absence of anomaly training images. Instead of using a unified paradigm (i.e., one model for all classes) in AnomalyCLIP and AdaCLIP, PromptAD uses a separate paradigm (i.e. one model for one class). Therefore, it needs to re-train a model with few-shot normal images when applied to each class of the target datasets. In addition, image-level anomaly classification and pixel-level segmentation models also need to be trained separately. In this paper, we only compare PromptAD with our method on MVTec and VisA because it involves fine-tuning for each class of target datasets.

8. Compelete Experimental Results

In our main paper, we compare state-of-the-art methods with our AdaptCLIP using AUROC for image-level anomaly classification and AUPR for pixel-level anomaly segmentation. Here, we provide more comprehensive comparisons, including image-level anomaly classification in AUPR and $F1_{max}$ in Tabs. 9 and 10, and pixel-level anomaly segmentation in AUROC and $F1_{max}$ in Tabs. 11 and 12, respectively. To more intuitively show the performance trends between zero-shot and few-shot methods on different datasets, we show comprehensive comparisons using all three metrics (AUROC, AUPR and $F1_{max}$) for image-level anomaly classification and pixel-level anomaly segmentation, as shown in Fig. 5. In addition, we only report the averaged results of all categories for each dataset in our main paper. Here, we also provide more detailed reports in Tabs. 13, 14, 15, 16, 17, 18, 19, 20, 21, and 22 for each

category on MVTec, VisA, MVTec3d, DTD, MPDD, Real-IAD, BTAD, KSDD, Br35H, Covid, Kvasir and Endo, respectively.

9. More Visualizations

In our main paper, we only visualize some selected examples from MVTec, VisA, Real-IAD, Kvasir and Endo to compare zero-shot and few-shot AdaptCLIP. Here, we show more visualizations for all 91 categories from 8 industrial and 2 medical datasets, as shown in Figs. 6, 7, 8, 9, 10, 11, 12, 13, 14 and 15.

Image-Level Classification

Pixel-Level Segmentation

(a) AUROC. (b) AUPR. (c) F1$_{\max}$.
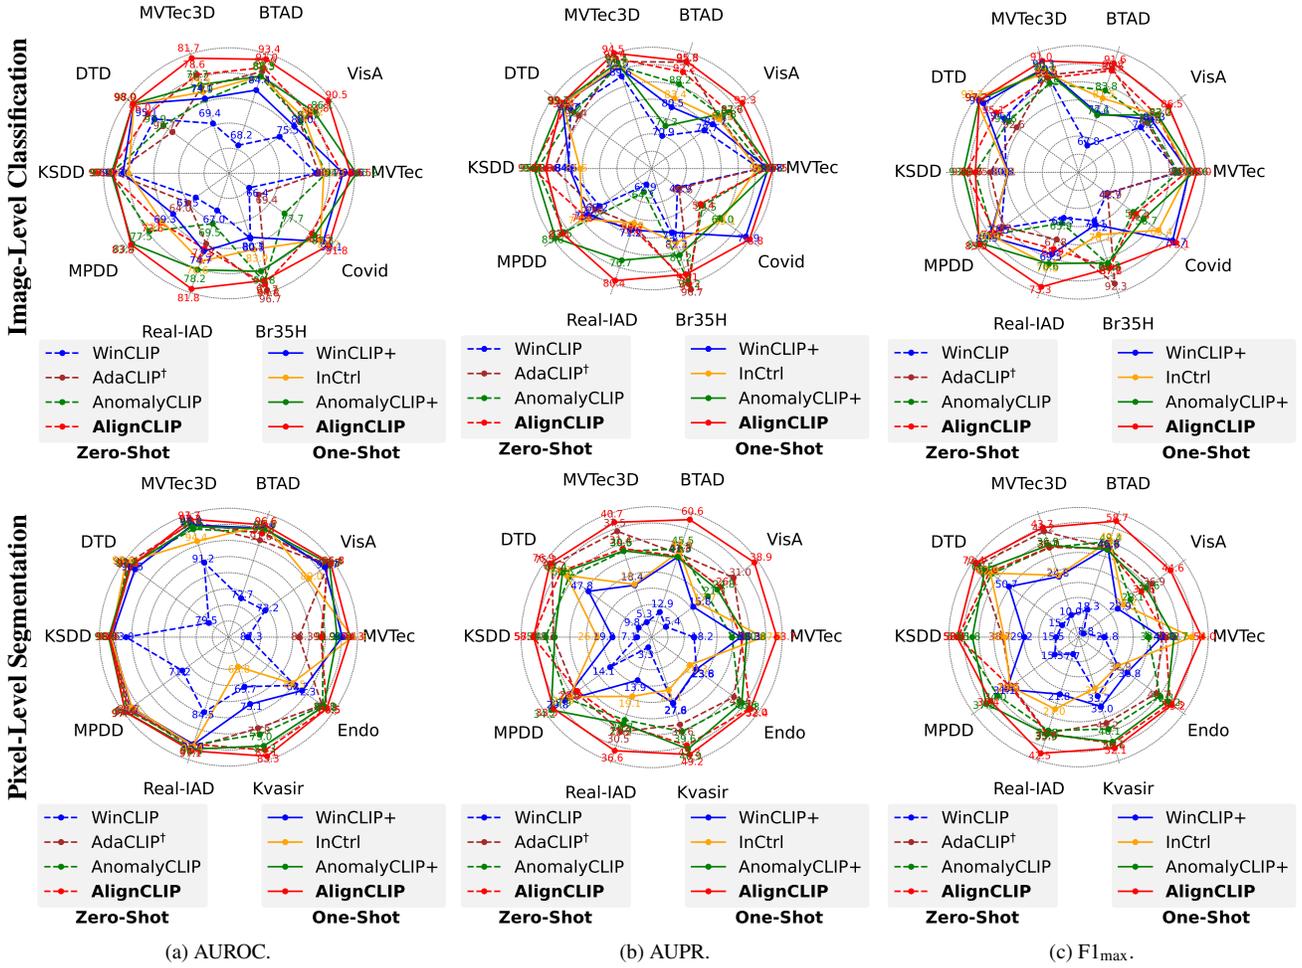
Figure 5. Image-level anomaly classification and pixel-level anomaly segmentation comparisons of state-of-the-art zero-/one-shot methods and our **AdaptCLIP** with all three metrics, AUROC, AUPR and F1$_{\max}$. The one-shot AdaptCLIP utilizes a training-free manner on target domains and achieves more accurate anomaly classification and segmentation on 8 industrial and 4 medical benchmarks.

Table 9. Image-level anomaly classification comparisons with **AUPR** metric on industrial and medical domains.

Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Br35H	Covid	AVG
0-shot	WinCLIP [16]	95.6	78.7	70.9	89.6	97.7	84.9	69.2	62.9	81.2	82.2	42.9	62.6
	AdaCLIP† [6]	95.6	84.3	95.5	92.7	96.4	89.3	70.8	70.2	86.9	96.7	43.6	70.2
	AnomalyCLIP [53]	96.4	85.3	88.2	91.8	97.2	94.2	82.5	64.3	87.5	94.2	55.5	74.9
	AdaptCLIP	96.7	87.6	92.2	93.4	98.4	95.7	74.5	70.8	88.8	95.1	54.4	74.8
1-shot	WinCLIP+ [16]	96.8±0.2	81.7±1.5	80.5±3.3	91.5±0.2	99.0±0.1	84.6±0.9	73.5±2.2	71.2±0.3	84.9	79.4±1.9	76.9±6.4	78.2
	InCtrl [54]	95.2±0.3	84.1±1.5	83.4±9.3	91.9±0.7	98.9±0.2	81.5±2.2	75.7±0.8	69.9±0.0	85.1	82.6±8.7	65.1±5.6	73.9
	AnomalyCLIP+ [53]	97.2±0.1	87.7±1.1	74.2±1.5	92.4±0.9	99.2±0.2	95.2±0.1	85.6±2.8	76.7±0.0	88.5	86.2±7.8	64.0±4.9	75.1
	AdaptCLIP	97.5±0.1	92.3±0.9	95.8±0.9	94.5±0.5	99.1±0.0	91.8±0.2	83.1±3.7	80.4±0.2	91.8	92.1±3.1	78.8±4.6	85.5
2-shot	WinCLIP+ [16]	97.3±0.5	84.0±0.7	82.5±3.2	91.6±0.1	99.1±0.1	84.5±0.6	73.6±2.0	72.7±0.1	85.7	80.5±0.4	80.3±3.1	80.4
	InCtrl [54]	95.5±0.7	86.8±1.7	81.6±8.0	91.8±0.2	99.1±0.3	81.0±2.5	75.9±0.7	71.7±0.0	85.4	84.8±2.9	66.2±4.4	75.5
	AnomalyCLIP+ [53]	97.3±0.1	89.1±0.7	75.4±1.5	92.9±0.4	99.3±0.1	95.6±0.2	85.4±2.4	76.9±0.0	89.0	86.8±6.2	66.7±4.5	76.8
	AdaptCLIP	97.9±0.2	93.6±0.6	95.9±0.1	94.8±0.4	99.2±0.0	92.4±0.4	84.7±2.1	81.5±0.1	92.5	92.2±2.5	83.6±1.9	87.9
4-shot	WinCLIP+ [16]	97.7±0.0	85.5±0.9	88.1±1.4	92.2±0.1	99.2±0.1	84.9±0.5	75.3±0.1	73.6±0.1	87.1	81.0±0.6	81.1±3.2	81.1
	InCtrl [54]	96.3±0.5	88.0±0.3	80.9±1.7	92.8±0.6	98.3±0.4	84.6±1.7	79.5±2.7	75.6±0.0	87.0	88.6±1.0	66.9±3.4	77.8
	AnomalyCLIP+ [53]	97.8±0.0	90.1±0.7	77.5±3.1	93.3±0.4	99.4±0.1	95.0±0.2	88.0±2.2	77.1±0.0	89.8	86.3±6.5	70.7±4.7	78.5
	AdaptCLIP	98.4±0.2	94.3±0.2	96.4±0.1	95.3±0.2	99.3±0.0	91.7±0.9	87.7±2.4	82.6±0.0	93.2	91.8±2.9	85.8±1.2	88.8

4

Figure 6. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **MVTec**.

Query | GT Mask | 0-shot | 1-shot | 4-shot | Query | GT Mask | 0-shot | 1-shot | 4-shot

Table 10. Image-level anomaly classification comparisons with $F1_{max}$ metric on industrial and medical domains.

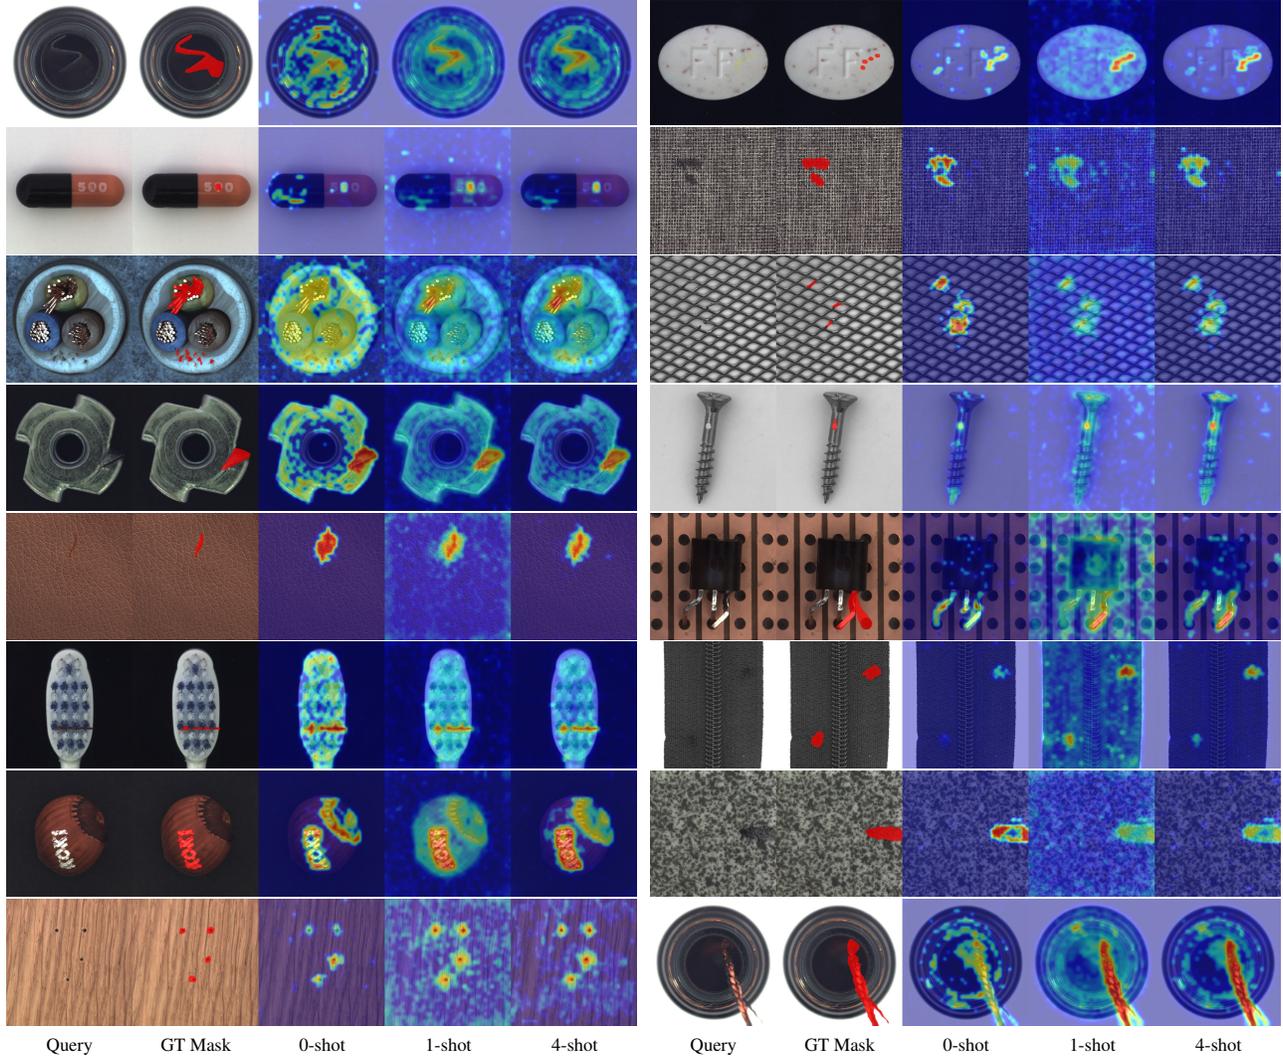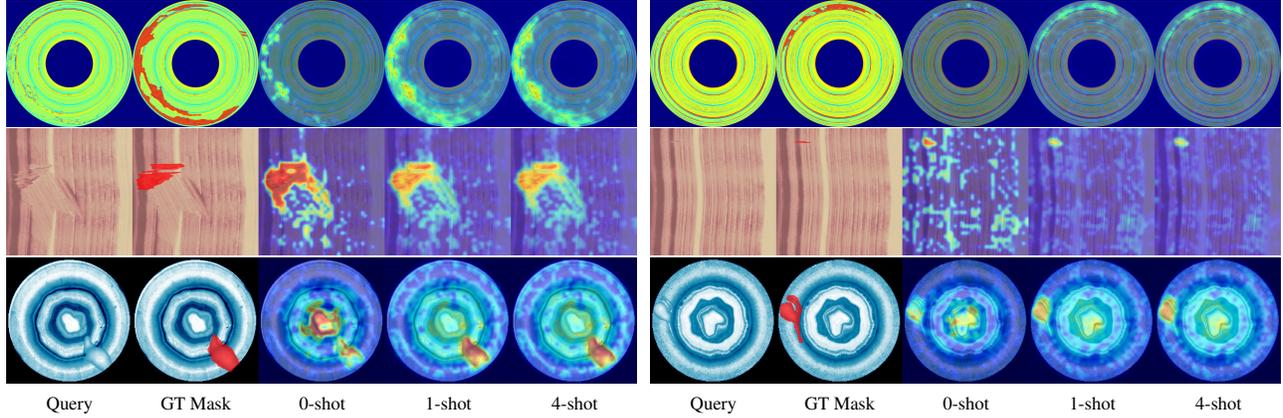| Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Br35H	Covid	AVG
0-shot	WinCLIP [16]	92.7	78.2	67.8	89.7	94.1	80.8	77.5	65.3	80.8	74.1	42.7	58.4
	AdaCLIP† [6]	92.4	80.0	90.2	89.6	92.6	85.4	76.7	67.8	84.3	92.3	42.9	67.6
	AnomalyCLIP [53]	92.7	80.4	83.8	88.8	93.6	89.7	80.4	65.9	84.4	86.8	54.2	70.5
	AdaptCLIP	93.7	83.0	89.5	89.3	95.1	92.3	79.7	68.9	86.4	87.7	55.4	71.6
1-shot	WinCLIP+ [16]	94.0±0.4	81.3±0.8	77.1±3.5	90.3±0.2	96.7±0.3	80.1±1.6	81.5±0.9	69.5±0.1	83.8	75.2±1.8	71.7±5.5	73.5
	InCtrl [54]	93.9±0.4	83.1±1.4	81.5±5.6	89.8±0.3	97.7±0.1	80.3±1.2	80.9±2.2	70.8±0.0	84.8	78.1±2.7	65.4±3.0	71.8
	AnomalyCLIP+ [53]	94.6±0.2	83.2±0.6	76.6±0.8	90.1±0.5	97.2±0.4	93.0±1.3	84.7±1.3	70.6±0.0	86.3	86.3±4.9	58.7±6.1	72.5
	AdaptCLIP	95.0±0.0	86.5±1.0	91.6±1.0	91.0±0.3	97.2±0.1	89.6±0.5	85.3±1.3	73.3±0.2	88.7	87.8±3.2	73.1±6.6	80.5
2-shot	WinCLIP+ [16]	94.5±0.4	82.3±1.1	78.7±2.9	90.3±0.2	97.0±0.1	80.3±2.0	81.0±1.2	70.3±0.1	84.3	76.4±0.3	75.1±2.0	75.8
	InCtrl [54]	94.2±0.2	84.3±1.3	81.4±5.8	90.0±0.2	97.9±0.1	80.5±0.2	81.4±1.7	72.1±0.0	85.2	78.4±1.2	65.8±1.5	72.1
	AnomalyCLIP+ [53]	94.9±0.1	84.5±0.5	77.2±0.5	90.2±0.4	97.7±0.1	93.6±1.0	84.8±0.5	70.7±0.0	86.7	86.9±4.1	61.7±5.8	74.3
	AdaptCLIP	95.4±0.1	88.0±0.7	91.7±0.5	91.2±0.2	97.7±0.1	89.4±0.9	85.7±0.7	74.2±0.1	89.2	88.3±2.0	78.7±2.4	83.5
4-shot	WinCLIP+ [16]	94.8±0.1	82.8±0.6	83.8±0.3	90.4±0.1	97.1±0.1	81.0±2.2	82.0±0.4	70.8±0.0	85.3	76.7±0.2	75.8±2.3	76.3
	InCtrl [54]	94.7±0.2	85.1±0.2	87.5±0.1	90.5±0.3	98.1±0.1	78.7±1.1	84.5±0.5	74.2±0.0	86.7	82.7±1.1	67.0±4.3	74.9
	AnomalyCLIP+ [53]	95.5±0.1	85.2±0.3	78.9±1.1	90.0±0.3	97.9±0.2	92.5±0.9	85.2±1.0	70.8±0.0	87.0	86.6±4.3	66.0±5.4	76.3
	AdaptCLIP	96.0±0.0	88.5±0.2	92.1±0.4	91.3±0.3	97.8±0.1	88.6±1.7	87.8±0.0	75.2±0.1	89.7	88.3±2.3	81.2±1.1	84.8

5

Figure 7. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **BTAD**.

Table 11. Pixel-level anomaly segmentation comparisons with **AUROC** metric on industrial and medical domains.

Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Kvasir	Endo	AVG
0-shot	WinCLIP [16]	82.3	73.2	72.7	91.2	79.5	93.0	71.2	84.5	81.0	69.7	68.2	69.0
	AdaCLIP† [6]	88.3	95.7	91.6	97.1	98.3	97.6	95.5	96.1	95.0	77.8	83.8	80.8
	AnomalyCLIP [53]	91.1	95.5	94.2	96.2	97.9	98.1	96.5	95.1	95.6	79.0	84.2	81.6
	AdaptCLIP	90.9	95.7	93.8	97.2	97.7	98.1	95.9	94.9	95.5	82.1	86.5	84.3
1-shot	WinCLIP+ [16]	93.4±0.2	94.7±0.1	95.6±0.2	96.8±0.2	96.5±0.1	97.6±0.1	94.7±1.3	95.0±0.0	95.5	73.1±2.6	72.3±0.3	72.7
	InCtrl [54]	94.6±0.2	89.0±0.2	96.6±0.1	94.4±0.2	98.6±0.1	97.8±0.2	94.4±1.0	95.4±0.0	95.1	65.8±1.3	66.5±6.2	66.2
	AnomalyCLIP+ [53]	92.8±0.0	96.4±0.1	95.3±0.3	96.6±0.1	97.6±0.1	98.6±0.1	97.4±0.2	96.5±0.0	96.4	81.3±0.4	84.8±1.5	83.1
	AdaptCLIP	94.3±0.1	96.8±0.0	96.6±0.2	97.7±0.0	97.4±0.0	98.2±0.1	97.4±0.2	97.1±0.0	96.9	83.3±0.5	86.5±1.2	84.9
2-shot	WinCLIP+ [16]	93.8±0.1	95.1±0.1	95.7±0.1	96.9±0.2	96.6±0.1	97.6±0.1	94.2±0.2	95.3±0.0	95.7	74.6±0.5	75.5±1.6	75.1
	InCtrl [54]	95.2±0.2	89.8±0.2	96.7±0.1	94.7±0.1	98.7±0.1	97.6±0.6	94.3±0.2	96.0±0.0	95.4	70.0±0.9	70.2±6.1	70.1
	AnomalyCLIP+ [53]	92.9±0.1	96.7±0.1	95.5±0.2	96.8±0.1	97.7±0.1	98.6±0.1	97.5±0.1	96.6±0.0	96.5	81.7±1.0	85.1±1.4	83.4
	AdaptCLIP	94.5±0.0	97.1±0.0	96.7±0.1	97.8±0.0	97.6±0.0	98.1±0.1	97.7±0.0	97.3±0.0	97.1	83.5±0.8	86.6±1.1	85.1
4-shot	WinCLIP+ [16]	94.2±0.2	95.1±0.2	95.9±0.1	97.0±0.1	96.8±0.1	97.5±0.2	94.7±0.3	95.5±0.0	95.8	75.0±0.8	75.3±0.3	75.2
	InCtrl [54]	95.8±0.2	90.2±0.2	96.8±0.0	95.2±0.1	98.7±0.1	97.5±0.3	95.1±0.5	96.4±0.0	95.7	70.6±1.2	69.4±1.3	70.0
	AnomalyCLIP+ [53]	93.2±0.1	96.9±0.1	95.7±0.1	97.0±0.1	97.8±0.0	98.6±0.1	97.8±0.1	96.7±0.0	96.7	81.3±0.4	84.5±0.9	82.9
	AdaptCLIP	94.8±0.1	97.3±0.0	96.8±0.0	98.0±0.0	97.8±0.1	98.0±0.1	97.9±0.1	97.4±0.0	97.3	83.0±0.5	85.8±0.7	84.4

Table 12. Pixel-level anomaly segmentation comparisons with **F1max** metric on industrial and medical domains.

Shots	Methods	Industrial									Medical		
		MVTec	VisA	BTAD	MVTec3D	DTD	KSDD	MPDD	Real-IAD	AVG	Kvasir	Endo	AVG
0-shot	WinCLIP [16]	24.8	8.8	18.3	10.0	15.7	15.5	15.3	7.7	14.5	35.7	32.9	34.3
	AdaCLIP† [6]	42.1	36.9	46.6	42.2	70.7	38.7	29.5	35.4	42.8	44.2	48.7	46.5
	AnomalyCLIP [53]	38.0	28.1	49.4	36.0	62.2	56.5	34.0	34.5	42.3	46.1	50.3	48.2
	AdaptCLIP	43.2	32.6	46.3	36.1	63.6	58.1	29.9	35.6	43.2	50.6	54.5	52.6
1-shot	WinCLIP+ [16]	42.0±0.8	22.9±0.2	46.8±2.3	24.8±0.7	50.7±0.7	29.2±0.2	31.1±2.2	21.8±0.2	33.7	39.0±1.9	36.8±0.3	37.9
	InCtrl [54]	51.0±1.6	25.2±0.4	49.9±1.3	25.3±0.7	62.9±0.5	38.0±0.6	29.9±1.7	27.0±0.0	38.7	33.2±1.1	33.0±3.7	33.1
	AnomalyCLIP+ [53]	45.7±0.2	34.6±0.6	45.7±1.4	36.9±1.1	66.5±0.5	51.8±0.3	37.8±0.9	35.9±0.0	44.4	50.1±1.1	53.2±2.3	51.7
	AdaptCLIP	54.0±0.7	44.6±0.4	58.7±0.7	43.7±0.7	72.4±0.2	58.4±0.5	36.4±2.4	42.5±0.1	51.3	52.1±0.7	55.2±2.0	53.7
2-shot	WinCLIP+ [16]	43.2±0.7	24.3±0.8	48.1±1.1	25.5±0.5	51.0±0.7	29.6±0.1	32.0±1.3	23.0±0.0	34.6	40.2±0.6	39.1±0.9	39.7
	InCtrl [54]	52.3±0.8	26.3±0.5	49.9±0.8	27.0±0.5	63.0±0.4	37.5±2.1	31.0±1.2	28.0±0.0	39.4	36.4±0.4	35.2±4.2	35.8
	AnomalyCLIP+ [53]	46.3±0.1	36.0±0.5	46.2±1.3	38.5±1.4	66.9±0.3	51.9±0.5	38.9±0.8	36.0±0.0	45.1	50.6±1.9	54.1±2.1	52.4
	AdaptCLIP	55.0±0.3	46.1±0.4	58.8±0.4	45.4±1.1	72.8±0.3	58.1±0.6	38.2±0.2	43.4±0.2	52.2	52.6±1.4	55.8±1.7	54.2
4-shot	WinCLIP+ [16]	44.6±0.6	25.3±1.2	49.5±0.7	26.3±0.5	51.7±0.2	29.9±0.8	33.5±0.4	23.9±0.3	35.6	40.7±0.8	39.2±0.2	40.0
	InCtrl [54]	54.1±0.4	27.2±0.6	50.3±0.7	28.9±1.0	63.1±0.2	36.7±1.0	33.6±1.9	29.0±0.0	40.4	36.9±0.8	34.6±1.4	35.8
	AnomalyCLIP+ [53]	47.1±0.1	37.5±0.9	49.6±1.2	39.4±1.0	67.2±0.1	51.2±0.6	40.1±1.1	36.2±0.0	46.0	50.1±0.8	53.8±1.6	52.0
	AdaptCLIP	56.8±0.7	47.2±0.5	59.6±0.2	47.6±0.2	73.5±0.1	57.2±0.7	40.7±1.1	44.6±0.3	53.4	52.1±0.8	55.6±1.2	53.9

Table 13. Image-level anomaly classification and pixel-level anomaly segmentation results on **MVTec** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	carpet	100.0	99.9	100.0	99.2	68.2	67.2
	bottle	90.7	97.1	90.8	92.4	60.3	56.4
	hazelnut	94.8	97.0	92.4	97.6	51.0	53.7
	leather	100.0	99.9	99.5	99.1	27.9	43.7
	cable	83.6	89.9	82.5	76.6	15.6	24.7
	capsule	95.2	99.0	95.0	95.8	35.4	40.2
	grid	98.9	99.5	97.3	97.4	27.5	39.7
	pill	86.8	97.0	93.9	89.9	33.1	35.2
	transistor	88.7	87.1	81.1	69.3	15.6	19.5
	metal_nut	93.7	98.3	94.7	76.5	27.6	34.5
	screw	85.9	93.8	90.3	97.9	28.8	34.5
	toothbrush	86.4	92.5	95.1	87.7	13.5	22.9
	zipper	99.3	99.8	98.3	91.8	44.4	47.7
	tile	99.6	99.7	98.2	95.8	68.2	67.8
	wood	98.7	99.5	96.7	96.9	57.5	60.9
	mean	93.5	96.7	93.7	90.9	38.3	43.2
1-shot	carpet	100.0±0.0	100.0±0.0	100.0±0.0	99.2±0.0	71.3±0.2	68.4±0.1
	bottle	99.1±0.1	99.7±0.0	97.7±0.6	95.6±0.0	75.1±0.3	71.1±0.4
	hazelnut	99.9±0.1	99.9±0.0	99.3±0.6	98.9±0.0	74.7±0.5	70.3±0.4
	leather	100.0±0.0	100.0±0.0	100.0±0.0	99.6±0.0	60.2±0.6	60.8±0.2
	cable	88.7±2.0	92.9±1.3	85.9±0.8	88.3±0.7	40.7±0.9	42.7±1.1
	capsule	90.0±5.2	97.7±1.3	95.5±1.2	97.3±0.1	40.4±6.0	43.7±3.9
	grid	92.2±1.0	97.4±0.3	90.2±0.6	96.5±0.3	35.9±0.3	43.6±0.4
	pill	92.6±0.5	98.1±0.2	95.2±0.3	93.3±0.1	55.3±0.3	55.3±0.3
	transistor	94.8±1.6	93.3±1.7	89.0±1.8	74.6±0.4	25.7±1.6	26.9±1.4
	metal_nut	99.2±0.4	99.8±0.1	98.7±0.2	89.0±0.3	64.9±1.3	62.2±0.9
	screw	74.0±3.8	88.2±2.0	87.4±0.1	97.5±0.1	19.5±3.8	25.1±3.9
	toothbrush	90.7±1.2	96.0±0.2	92.6±3.0	97.6±0.2	47.8±1.0	49.9±0.5
	zipper	98.1±0.3	99.5±0.1	97.9±0.3	93.5±0.1	49.9±0.5	52.1±0.3
	tile	99.5±0.1	99.8±0.0	97.5±0.3	96.1±0.0	72.3±0.1	68.7±0.0
	wood	99.7±0.1	99.9±0.0	98.7±0.8	97.1±0.2	72.3±0.6	68.6±0.3
	mean	94.5±0.5	97.5±0.1	95.0±0.0	94.3±0.1	53.7±0.9	54.0±0.7
2-shot	carpet	100.0±0.0	100.0±0.0	100.0±0.0	99.3±0.0	71.3±0.4	68.5±0.1
	bottle	99.2±0.2	99.7±0.0	98.1±0.4	95.8±0.0	75.9±0.4	72.0±0.3
	hazelnut	99.8±0.2	99.9±0.1	98.8±0.9	98.9±0.0	74.0±0.9	70.1±0.3
	leather	100.0±0.0	100.0±0.0	100.0±0.0	99.6±0.0	60.5±0.6	60.8±0.1
	cable	90.5±1.3	94.2±0.9	88.2±1.2	88.6±0.0	41.9±0.8	44.1±1.0
	capsule	94.3±4.9	98.7±1.2	96.3±1.1	97.5±0.2	45.3±5.9	47.0±4.1
	grid	93.2±2.9	97.7±1.0	91.2±2.5	97.4±0.4	37.8±2.2	44.5±1.0
	pill	93.3±0.2	98.2±0.1	95.5±0.3	93.3±0.2	55.4±1.2	55.6±1.2
	transistor	95.1±0.5	93.3±0.7	87.5±0.7	75.0±0.3	26.0±1.0	26.8±0.9
	metal_nut	99.4±0.5	99.9±0.1	98.9±0.4	89.5±0.2	66.8±1.2	64.3±1.3
	screw	79.2±2.7	91.1±1.4	87.8±0.6	97.9±0.1	22.8±2.3	28.1±2.2
	toothbrush	93.5±4.7	97.2±2.0	94.7±4.0	98.1±0.5	53.4±6.2	54.2±4.8
	zipper	98.2±0.0	99.5±0.0	97.9±0.3	93.7±0.0	50.2±0.3	52.0±0.3
	tile	99.4±0.0	99.8±0.0	97.5±0.3	96.2±0.0	72.4±0.1	68.8±0.1
	wood	99.7±0.1	99.9±0.0	98.9±0.4	97.3±0.1	72.8±0.6	68.8±0.2
	mean	95.7±0.6	97.9±0.2	95.4±0.1	94.5±0.0	55.1±0.5	55.0±0.3
4-shot	carpet	100.0±0.0	100.0±0.0	99.6±0.3	99.2±0.0	71.4±0.3	68.4±0.2
	bottle	99.5±0.2	99.8±0.0	98.2±0.3	95.9±0.0	76.4±0.4	72.6±0.4
	hazelnut	99.9±0.0	100.0±0.0	99.5±0.3	98.9±0.0	74.8±0.9	70.9±0.6
	leather	100.0±0.0	100.0±0.0	100.0±0.0	99.6±0.0	60.9±0.6	60.9±0.3
	cable	92.1±0.7	95.3±0.4	88.5±0.5	88.8±0.4	44.0±0.6	45.9±1.0
	capsule	97.5±0.4	99.5±0.1	97.1±0.5	97.6±0.1	49.4±1.0	49.8±1.0
	grid	95.0±2.2	98.3±0.7	93.0±2.1	97.5±0.4	40.0±2.6	44.9±1.3
	pill	94.0±0.1	98.3±0.1	95.9±0.1	93.4±0.0	56.0±0.4	56.7±0.6
	transistor	95.5±1.3	94.0±2.0	89.4±0.9	76.1±0.2	27.1±0.4	27.9±0.5
	metal_nut	99.5±0.5	99.9±0.1	98.9±0.4	90.2±0.2	69.9±1.1	67.9±1.4
	screw	82.6±3.5	92.8±2.0	89.2±1.6	98.2±0.1	30.1±2.5	34.2±2.1
	toothbrush	96.6±4.3	98.4±2.0	96.9±2.5	98.8±0.4	62.5±7.9	62.6±7.1
	zipper	98.3±0.2	99.5±0.0	98.0±0.2	93.8±0.0	50.1±0.2	51.9±0.1
	tile	99.3±0.1	99.7±0.0	97.4±0.2	96.3±0.0	72.6±0.1	68.9±0.0
	wood	99.8±0.0	99.9±0.0	99.2±0.0	97.3±0.2	73.1±0.5	68.9±0.3
	mean	96.6±0.3	98.4±0.2	96.0±0.0	94.8±0.1	57.2±0.8	56.8±0.7

Table 14. Image-level anomaly classification and pixel-level anomaly segmentation results on **VisA** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	candle	87.4	90.5	81.4	98.7	21.6	35.2
	capsules	93.9	96.8	91.6	94.2	33.4	44.1
	cashew	86.1	94.0	84.0	93.4	22.1	27.2
	chewinggum	96.7	98.7	96.4	99.5	82.7	77.8
	fryum	91.5	95.9	88.5	95.3	26.5	33.6
	macaroni1	83.3	85.1	75.6	98.1	16.4	23.0
	macaroni2	69.9	69.2	71.5	98.2	2.2	6.4
	pcb1	84.1	86.1	80.4	96.2	15.7	24.6
	pcb2	64.6	67.8	69.9	93.0	12.4	20.0
	pcb3	65.3	72.0	67.3	87.9	4.6	10.4
	pcb4	97.7	97.5	93.5	95.3	31.9	36.8
	pipe_fryum	96.7	98.0	95.6	98.4	44.0	52.4
	mean	84.8	87.6	83.0	95.7	26.1	32.6
1-shot	candle	92.9±0.9	94.0±0.6	87.9±1.1	98.9±0.0	23.7±1.0	37.6±1.4
	capsules	94.0±0.3	96.5±0.1	90.7±0.6	96.8±0.6	29.0±2.5	36.1±1.7
	cashew	96.2±0.5	98.3±0.2	93.2±0.7	92.8±0.7	46.7±1.1	53.8±0.7
	chewinggum	98.1±0.0	99.2±0.0	96.3±0.5	99.7±0.0	81.5±0.3	76.4±0.4
	fryum	96.3±0.6	98.4±0.3	93.2±1.2	96.7±0.0	41.2±1.4	45.8±0.9
	macaroni1	93.2±0.6	94.2±0.7	85.1±0.3	99.0±0.0	27.0±1.3	33.9±0.6
	macaroni2	74.7±2.5	79.3±1.2	71.0±1.8	98.5±0.2	14.2±0.7	25.6±0.2
	pcb1	83.6±12.6	84.6±9.4	80.4±7.7	98.3±0.0	57.9±2.3	58.5±1.3
	pcb2	82.4±0.9	85.2±0.8	76.1±1.4	94.7±0.2	18.9±0.6	28.9±0.7
	pcb3	79.5±3.5	81.8±3.0	74.3±1.5	91.8±0.0	28.8±0.9	34.6±2.1
	pcb4	95.6±3.6	95.9±2.4	90.7±5.2	95.8±0.1	34.6±3.1	38.6±2.4
	pipe_fryum	99.6±0.1	99.7±0.1	98.5±0.0	99.2±0.2	63.4±4.3	65.6±2.7
	mean	90.5±1.2	92.3±0.9	86.5±1.0	96.8±0.0	38.9±0.3	44.6±0.4
2-shot	candle	95.0±0.5	95.3±0.3	89.1±0.8	98.9±0.0	24.1±0.8	38.0±0.7
	capsules	94.5±1.0	96.7±0.4	91.9±0.8	97.4±0.3	32.9±1.7	39.0±1.0
	cashew	96.0±0.3	98.2±0.1	93.2±0.6	93.4±0.6	46.9±1.5	54.0±1.2
	chewinggum	98.5±0.0	99.3±0.0	96.7±0.6	99.6±0.0	81.8±0.4	76.6±0.2
	fryum	96.4±0.1	98.5±0.0	94.0±0.7	96.9±0.1	42.0±1.6	46.3±1.5
	macaroni1	91.8±0.8	93.3±0.9	85.1±1.0	99.0±0.0	26.1±1.0	32.9±1.7
	macaroni2	76.8±3.7	81.5±2.5	72.1±2.5	98.8±0.2	15.2±0.6	26.3±0.7
	pcb1	91.6±2.5	90.6±2.0	84.9±3.6	98.6±0.1	65.2±5.8	64.3±4.7
	pcb2	84.9±2.9	87.2±2.3	78.2±2.7	95.3±0.2	20.9±0.3	30.8±0.3
	pcb3	83.9±1.9	85.4±2.1	78.4±0.7	92.1±0.0	33.3±0.9	39.0±0.7
	pcb4	97.8±0.9	97.6±0.7	93.4±1.9	96.4±0.2	38.3±3.4	41.8±3.0
	pipe_fryum	99.6±0.1	99.7±0.0	98.0±0.4	99.2±0.1	61.8±0.3	64.5±0.5
	mean	92.2±0.8	93.6±0.6	88.0±0.7	97.1±0.0	40.7±0.6	46.1±0.4
4-shot	candle	96.0±0.1	95.8±0.2	90.3±0.3	98.9±0.0	24.0±0.6	37.8±0.4
	capsules	95.2±0.5	97.0±0.2	92.5±0.9	97.7±0.2	34.0±1.0	39.7±0.5
	cashew	95.9±0.1	98.2±0.0	93.6±0.2	93.6±0.5	46.4±1.3	54.3±1.0
	chewinggum	98.4±0.2	99.3±0.1	96.5±0.4	99.6±0.0	81.8±0.6	76.6±0.4
	fryum	96.9±0.3	98.7±0.1	94.8±0.0	96.9±0.1	42.4±1.0	46.3±1.0
	macaroni1	93.0±0.7	94.4±0.6	86.5±0.7	99.0±0.0	27.4±1.9	33.7±1.3
	macaroni2	79.6±2.3	83.9±1.5	72.6±2.3	99.1±0.1	16.3±0.4	27.4±0.1
	pcb1	92.0±2.4	91.0±1.8	84.9±2.5	98.6±0.1	67.5±4.7	66.2±3.8
	pcb2	86.5±1.5	88.3±1.2	79.0±2.5	95.6±0.0	23.0±0.2	32.7±0.2
	pcb3	84.8±1.2	86.6±0.7	78.2±1.5	92.3±0.0	36.3±0.6	42.7±0.3
	pcb4	98.8±0.3	98.4±0.4	95.5±0.6	96.8±0.1	40.2±3.2	43.4±3.0
	pipe_fryum	99.7±0.0	99.7±0.0	98.2±0.2	99.3±0.0	62.6±0.4	65.3±0.1
	mean	93.1±0.2	94.3±0.2	88.5±0.2	97.3±0.0	41.8±0.6	47.2±0.5

8

Table 15. Image-level anomaly classification and pixel-level anomaly segmentation results on **MVTec 3D** with zero-shot and few-shot AdaptCLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	bagel	93.4	98.4	92.7	99.6	73.1	67.2
	cable_gland	73.6	92.4	90.2	98.2	30.1	33.6
	carrot	78.7	94.7	91.3	98.1	24.5	31.4
	cookie	72.7	92.0	88.0	96.3	42.9	44.2
	tire	78.9	93.7	87.4	97.9	29.7	36.7
	rope	76.8	90.2	81.5	96.9	27.2	29.8
	foam	81.2	95.0	89.9	89.9	23.7	32.4
	dowel	74.0	91.1	90.1	96.6	5.8	11.9
	peach	86.5	96.4	91.9	99.1	27.5	37.4
	potato	70.8	90.0	89.8	99.4	29.9	36.8
	mean	78.6	93.4	89.3	97.2	31.4	36.1
1-shot	bagel	95.9±1.0	99.0±0.2	95.0±1.2	99.6±0.0	77.0±0.4	71.2±0.2
	carrot	85.3±2.1	96.5±0.6	92.1±0.8	99.2±0.0	40.2±1.1	44.3±0.7
	dowel	68.6±5.6	89.6±2.5	89.7±0.7	97.1±0.1	9.7±0.5	14.3±0.4
	potato	76.4±2.1	92.1±0.3	90.4±0.4	99.5±0.1	41.5±5.0	46.4±3.7
	rope	95.3±3.1	98.0±1.3	93.4±2.8	98.0±0.4	40.6±1.2	46.0±0.5
	cable_gland	70.5±4.2	91.5±1.6	90.2±0.4	97.7±0.1	22.0±4.3	28.0±5.9
	cookie	74.7±0.4	92.2±0.3	88.3±0.2	97.5±0.2	53.0±3.8	51.7±3.6
	foam	79.7±1.8	94.5±0.6	90.2±0.9	89.6±0.4	32.7±0.4	42.8±0.2
	peach	88.4±2.1	97.0±0.5	91.8±0.6	99.3±0.0	48.8±1.8	48.7±2.0
	tire	82.0±2.0	94.4±0.6	88.8±0.8	98.8±0.0	41.3±1.2	44.3±1.1
	mean	81.7±1.5	94.5±0.5	91.0±0.3	97.7±0.0	40.7±0.6	43.7±0.7
2-shot	bagel	96.9±0.7	99.3±0.2	95.7±1.0	99.7±0.0	76.8±0.7	71.3±0.5
	carrot	86.8±1.4	97.1±0.3	92.3±0.7	99.1±0.0	41.0±2.0	45.0±1.5
	dowel	69.6±6.2	89.9±2.8	89.4±0.4	97.3±0.4	10.4±2.3	15.0±2.9
	potato	77.0±0.5	92.0±0.2	90.3±0.2	99.6±0.0	43.3±2.1	48.3±0.6
	rope	96.0±2.5	98.3±1.0	93.6±2.0	98.1±0.2	41.3±0.8	47.0±0.4
	cable_gland	71.3±3.9	91.6±1.5	90.3±0.5	97.8±0.2	25.6±1.8	32.2±3.1
	cookie	77.5±0.8	93.3±0.2	88.5±0.3	97.7±0.1	56.8±0.7	55.6±0.4
	foam	80.6±1.6	94.7±0.5	91.0±0.5	90.1±0.3	33.2±0.3	42.9±0.0
	peach	90.1±1.2	97.5±0.3	92.2±0.4	99.4±0.1	52.4±6.1	51.3±4.8
	tire	83.3±2.4	94.9±0.8	89.2±0.8	98.9±0.0	42.7±1.0	45.6±0.8
	mean	82.9±1.1	94.8±0.4	91.2±0.2	97.8±0.0	42.3±1.1	45.4±1.1
4-shot	bagel	96.9±0.2	99.3±0.0	95.9±0.5	99.7±0.0	77.2±1.4	71.5±1.1
	carrot	88.3±0.3	97.4±0.2	92.7±0.3	99.0±0.0	43.4±0.1	47.9±0.5
	dowel	73.9±3.6	91.0±0.8	90.3±0.9	98.0±0.2	15.4±3.4	21.1±3.4
	potato	78.8±1.2	92.9±0.6	90.2±0.6	99.6±0.0	44.7±1.9	49.3±0.8
	rope	96.0±0.7	98.3±0.3	92.5±0.9	98.1±0.1	41.2±1.0	47.2±0.8
	cable_gland	75.6±1.1	93.2±0.4	90.7±0.6	98.4±0.2	33.2±2.4	38.5±1.3
	cookie	79.2±1.6	93.8±0.4	89.0±0.8	97.8±0.0	57.9±0.5	56.8±0.7
	foam	80.0±1.7	94.6±0.5	91.0±0.7	90.5±0.2	33.4±0.1	43.0±0.2
	peach	90.0±0.9	97.5±0.2	92.2±0.6	99.4±0.0	54.7±4.9	53.2±3.2
	tire	83.4±1.3	95.2±0.4	88.4±0.5	98.9±0.0	44.1±1.1	46.8±0.8
	mean	84.2±0.6	95.3±0.2	91.3±0.3	98.0±0.0	44.5±0.3	47.6±0.2

Table 16. Image-level anomaly classification and pixel-level anomaly segmentation results on **DTD** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	Woven_001	100.0	99.8	100.0	99.8	78.2	71.7
	Woven_127	94.8	95.4	88.4	94.5	54.0	54.6
	Woven_104	98.9	99.7	98.1	95.2	67.6	64.8
	Stratified_154	98.1	99.5	96.4	99.6	77.6	68.9
	Blotchy_099	98.3	99.4	99.4	99.5	79.7	69.4
	Woven_068	97.0	98.3	93.0	98.7	55.8	53.4
	Woven_125	100.0	99.9	100.0	99.5	73.9	66.4
	Marbled_078	99.1	99.7	98.1	99.2	75.4	67.5
	Perforated_037	93.7	98.4	93.9	93.5	62.3	61.7
	Mesh_114	85.6	94.2	84.7	93.7	58.0	58.1
	Fibrous_183	98.8	99.6	97.5	99.1	68.8	60.7
	Matted_069	87.3	96.7	91.2	99.5	73.6	65.8
	mean	96.0	98.4	95.1	97.7	68.7	63.6
1-shot	Woven_001	100.0±0.0	99.8±0.0	100.0±0.0	99.6±0.1	78.9±0.5	74.1±0.2
	Woven_127	97.4±1.0	98.2±0.6	95.6±1.3	93.2±0.5	57.6±1.9	59.8±1.9
	Woven_104	98.2±0.4	99.6±0.1	97.1±0.3	94.8±0.6	67.6±0.8	66.1±0.5
	Stratified_154	99.4±0.4	99.8±0.1	98.1±0.9	99.7±0.0	84.3±0.8	74.6±1.2
	Blotchy_099	100.0±0.0	99.9±0.0	99.8±0.3	99.9±0.0	92.1±0.1	85.8±0.0
	Woven_068	96.5±1.0	98.2±0.6	93.9±2.5	96.8±0.3	64.8±2.9	60.4±2.7
	Woven_125	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	84.6±0.1	76.8±0.0
	Marbled_078	100.0±0.0	99.9±0.0	100.0±0.0	99.6±0.0	87.1±0.5	80.4±0.6
	Perforated_037	95.1±1.0	98.8±0.2	94.1±1.2	93.0±0.5	67.4±0.9	65.1±0.5
	Mesh_114	90.0±0.4	96.2±0.2	89.4±0.3	92.8±0.1	58.8±0.5	59.5±0.5
	Fibrous_183	100.0±0.0	99.9±0.0	100.0±0.0	99.6±0.0	90.6±0.2	84.1±0.2
	Matted_069	99.4±0.1	99.8±0.0	98.3±0.8	99.5±0.1	88.9±0.4	82.6±0.1
	mean	98.0±0.0	99.1±0.0	97.2±0.1	97.4±0.0	76.9±0.1	72.4±0.2
2-shot	Woven_001	100.0±0.0	99.8±0.0	100.0±0.0	99.6±0.0	79.1±0.8	74.0±0.5
	Woven_127	97.8±0.4	98.3±0.3	96.6±0.3	94.1±0.7	61.7±1.5	63.4±1.1
	Woven_104	99.0±0.2	99.7±0.0	98.3±0.6	95.7±0.4	69.8±0.5	66.9±0.4
	Stratified_154	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	83.6±0.3	74.1±0.2
	Blotchy_099	100.0±0.0	99.9±0.0	99.8±0.3	99.9±0.0	92.1±0.0	85.8±0.0
	Woven_068	96.9±0.5	98.4±0.2	94.6±1.2	96.8±0.4	65.4±3.0	61.2±2.9
	Woven_125	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	84.4±0.1	76.7±0.2
	Marbled_078	100.0±0.0	99.9±0.0	100.0±0.0	99.6±0.0	87.2±0.5	80.6±0.4
	Perforated_037	96.1±0.4	99.0±0.1	95.2±0.3	92.9±0.1	66.9±0.7	64.7±0.1
	Mesh_114	90.1±0.6	96.2±0.2	88.4±0.4	93.3±0.5	58.3±0.6	59.0±0.4
	Fibrous_183	100.0±0.0	99.9±0.0	99.8±0.3	99.6±0.0	91.1±0.4	84.5±0.4
	Matted_069	99.7±0.0	99.8±0.0	99.2±0.3	99.6±0.1	89.5±0.3	82.9±0.1
	mean	98.3±0.0	99.2±0.0	97.7±0.1	97.6±0.0	77.4±0.2	72.8±0.3
4-shot	Woven_001	100.0±0.0	99.8±0.0	99.8±0.3	99.7±0.0	80.0±1.5	74.6±1.2
	Woven_127	98.4±0.2	98.7±0.2	96.8±0.5	94.7±1.2	62.6±2.5	63.8±1.7
	Woven_104	99.1±0.0	99.7±0.0	97.7±0.3	95.9±0.1	70.1±0.2	67.0±0.3
	Stratified_154	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	84.4±1.0	75.7±1.1
	Blotchy_099	100.0±0.0	99.9±0.0	100.0±0.0	99.9±0.0	92.1±0.0	85.7±0.0
	Woven_068	96.9±0.5	98.4±0.3	95.5±1.6	97.7±0.4	68.5±3.2	64.4±2.7
	Woven_125	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	84.9±0.2	77.4±0.3
	Marbled_078	100.0±0.0	99.9±0.0	100.0±0.0	99.6±0.0	87.2±0.3	80.5±0.3
	Perforated_037	96.5±0.2	99.1±0.0	95.3±0.3	93.5±0.1	68.3±0.3	65.4±0.3
	Mesh_114	90.8±1.3	96.5±0.4	89.1±0.2	94.2±0.3	59.4±0.8	60.0±0.5
	Fibrous_183	100.0±0.0	99.9±0.0	100.0±0.0	99.7±0.0	91.4±0.2	84.8±0.2
	Matted_069	99.9±0.1	99.9±0.0	99.8±0.3	99.7±0.0	89.8±0.2	82.9±0.2
	mean	98.5±0.1	99.3±0.0	97.8±0.1	97.8±0.1	78.2±0.2	73.5±0.1

Table 17. Image-level anomaly classification and pixel-level anomaly segmentation results on **MPDD** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	bracket_black	65.8	73.1	77.9	96.4	15.8	26.4
	bracket_brown	56.3	72.1	80.3	92.6	2.9	6.7
	bracket_white	61.4	54.7	71.4	99.5	4.0	9.1
	connector	71.4	53.3	61.9	96.1	14.4	20.5
	metal_plate	88.6	94.8	91.8	93.3	60.4	61.5
	tubes	97.8	98.9	94.8	97.8	54.6	55.1
	mean	73.6	74.5	79.7	95.9	25.3	29.9
1-shot	bracket_black	76.7±5.1	79.9±6.9	81.5±2.8	97.7±0.3	8.6±2.2	19.5±3.5
	bracket_brown	63.0±5.1	74.3±5.3	80.3±0.0	93.0±0.4	3.2±0.2	6.6±0.3
	bracket_white	82.6±5.7	78.7±9.8	82.8±0.8	99.4±0.3	9.7±5.0	17.4±6.7
	connector	83.5±7.6	67.2±13.0	71.9±5.6	96.7±0.3	19.8±9.5	24.5±6.9
	metal_plate	100.0±0.0	99.8±0.0	99.8±0.3	98.9±0.4	93.1±2.8	85.7±2.9
	tubes	96.9±1.4	98.6±0.7	95.4±2.3	98.8±0.1	66.5±2.0	64.8±2.0
	mean	83.8±2.2	83.1±3.7	85.3±1.3	97.4±0.2	33.5±2.5	36.4±2.4
2-shot	bracket_black	76.0±4.0	79.4±6.8	80.0±0.8	97.7±0.3	10.8±4.1	21.8±6.0
	bracket_brown	63.0±2.8	73.9±3.7	80.1±0.3	93.7±0.2	3.7±0.1	7.0±0.3
	bracket_white	79.4±3.9	76.2±8.3	78.6±2.6	99.4±0.3	8.7±4.3	17.9±7.5
	connector	89.8±3.2	79.4±6.3	79.0±1.9	97.0±0.2	22.6±7.5	27.8±4.2
	metal_plate	100.0±0.0	99.8±0.0	100.0±0.0	99.2±0.1	95.1±0.6	87.6±0.8
	tubes	98.2±0.7	99.1±0.3	96.6±0.9	98.9±0.3	69.0±3.1	66.8±2.9
	mean	84.4±0.7	84.7±2.1	85.7±0.7	97.7±0.0	35.0±0.7	38.2±0.2
4-shot	bracket_black	80.7±2.6	85.4±3.3	81.5±0.5	98.1±0.0	15.1±1.7	27.2±2.5
	bracket_brown	66.1±2.9	76.3±3.8	81.2±0.3	94.4±0.2	4.5±0.3	8.5±0.5
	bracket_white	83.3±4.8	79.6±8.9	82.6±1.0	99.5±0.3	9.2±4.3	19.1±7.1
	connector	93.0±0.9	86.4±3.0	85.3±1.0	97.4±0.2	29.7±4.5	32.8±2.5
	metal_plate	100.0±0.0	99.8±0.0	100.0±0.0	99.3±0.0	95.9±0.4	89.0±0.6
	tubes	97.7±0.8	99.0±0.3	96.0±1.3	98.9±0.1	70.2±1.8	67.7±2.2
	mean	86.8±1.1	87.7±2.4	87.8±0.0	97.9±0.1	37.4±1.1	40.7±1.1

| Query | GT Mask | 0-shot | 1-shot | 4-shot | Query | GT Mask | 0-shot | 1-shot | 4-shot |
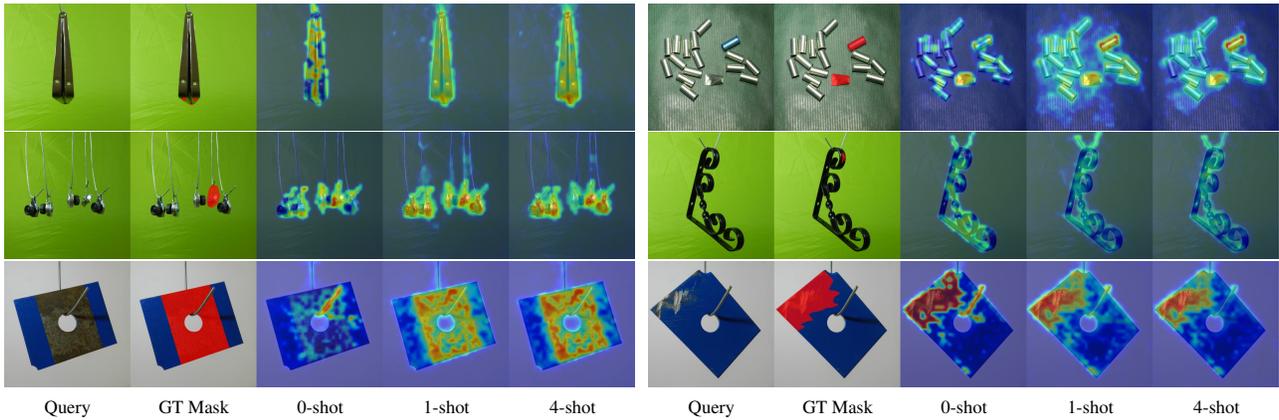
Figure 8. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **MPDD**.

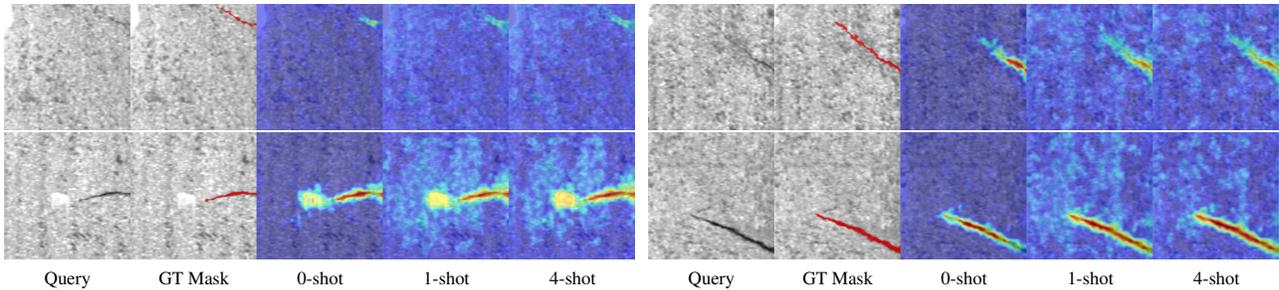| Query | GT Mask | 0-shot | 1-shot | 4-shot | Query | GT Mask | 0-shot | 1-shot | 4-shot |

Figure 9. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **KSDD**.

Table 18. Image-level anomaly classification and pixel-level anomaly segmentation results on **Real-IAD** with zero-shot and one-shot AdaptCLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	audiojack	63.5	48.2	53.4	94.9	11.3	20.9
	bottle_cap	66.2	63.1	61.2	98.0	26.6	36.8
	button_battery	70.6	78.0	73.4	98.2	50.8	54.2
	end_cap	66.1	70.9	74.0	95.2	10.0	18.2
	eraser	86.4	85.8	75.9	99.4	37.2	45.6
	fire_hood	80.0	70.6	64.8	99.0	37.1	42.8
	mint	66.4	65.7	65.0	93.6	26.2	38.6
	mounts	79.9	66.7	67.5	97.9	39.0	46.3
	pcb	61.9	73.0	75.5	94.4	7.4	13.5
	phone_battery	82.1	81.1	70.7	76.1	23.3	34.5
	plastic_nut	68.4	55.5	54.7	96.1	33.3	43.4
	plastic_plug	69.5	57.6	60.5	96.3	23.0	33.4
	porcelain_doll	77.9	60.7	67.5	99.1	42.9	44.9
	regulator	73.0	53.4	53.7	92.1	16.2	29.3
	rolled_strip_base	87.1	93.5	85.7	98.6	35.2	40.8
	sim_card_set	89.4	92.9	84.6	99.5	58.8	59.6
	switch	67.8	72.8	70.5	84.6	11.8	18.9
	tape	93.0	91.4	82.8	97.0	40.5	50.0
	terminalblock	76.8	81.8	71.5	96.8	36.5	44.2
	toothbrush	58.4	66.2	69.9	91.1	11.1	17.7
	toy	63.3	71.4	73.5	81.9	9.6	17.8
	toy_brick	69.3	61.0	61.6	97.6	22.7	29.8
	transistor1	71.3	75.5	75.5	91.8	22.7	32.3
	u_block	77.2	64.9	62.8	98.6	33.9	43.1
	usb	63.0	57.7	65.0	93.2	5.1	9.9
	usb_adaptor	72.8	62.2	63.6	97.1	15.3	25.3
	vcpill	80.2	75.6	67.7	98.0	48.4	53.0
	wooden_beads	82.6	79.7	73.1	98.4	41.9	46.2
	woodstick	75.6	55.9	56.2	97.5	46.6	52.2
	zipper	86.1	90.5	86.2	95.5	22.3	26.2
	mean	74.2	70.8	68.9	94.9	28.2	35.6
1-shot	audiojack	72.9±0.7	62.8±1.7	57.6±1.2	96.8±0.0	24.7±4.3	32.0±4.7
	bottle_cap	80.2±0.5	77.8±0.6	71.5±0.3	98.7±0.0	32.6±0.2	39.2±0.0
	button_battery	75.1±0.6	82.4±0.4	74.1±0.2	98.3±0.0	58.5±1.0	57.8±0.7
	end_cap	75.1±2.2	81.3±1.6	75.5±0.7	96.7±0.2	26.0±1.5	33.7±1.4
	eraser	88.2±1.1	86.2±1.2	75.2±1.6	99.5±0.0	50.6±0.1	50.9±0.1
	fire_hood	87.0±0.3	79.0±0.4	74.1±0.3	99.2±0.0	33.3±0.5	41.6±0.2
	mint	70.8±0.4	70.1±0.1	66.8±0.3	92.7±0.1	26.5±0.4	37.6±0.3
	mounts	85.7±0.4	72.0±0.7	73.8±0.6	98.4±0.0	37.0±0.4	43.1±0.6
	pcb	73.7±1.5	83.2±0.9	76.3±0.4	97.4±0.1	24.3±1.1	32.0±1.3
	phone_battery	88.2±0.9	88.0±1.1	78.7±1.5	99.2±0.0	59.5±1.2	62.9±0.7
	plastic_nut	79.0±0.1	72.7±0.6	64.0±0.3	97.6±0.0	41.9±0.2	44.9±0.3
	plastic_plug	82.0±0.5	78.2±0.6	68.2±1.0	97.2±0.3	26.9±0.8	37.7±0.2
	porcelain_doll	92.5±0.3	88.3±0.7	79.8±0.5	99.5±0.0	49.0±0.5	49.4±0.3
	regulator	73.0±0.8	58.3±1.0	53.8±1.1	96.0±0.2	10.5±1.2	19.3±1.3
	rolled_strip_base	90.5±0.5	95.0±0.4	88.0±0.4	99.4±0.0	34.1±0.4	42.9±1.0
	sim_card_set	95.8±0.2	96.6±0.1	90.6±0.3	99.7±0.0	58.4±0.6	61.8±0.1
	switch	83.4±0.7	86.9±0.6	76.4±0.6	93.8±0.0	34.1±0.9	43.2±0.9
	tape	94.1±0.3	92.9±0.3	84.7±0.6	98.9±0.0	45.5±0.2	51.3±0.1
	terminalblock	84.6±0.7	88.1±0.6	77.2±0.5	98.8±0.0	47.7±0.4	52.1±0.6
	toothbrush	75.2±1.0	81.3±0.9	71.6±0.2	93.6±0.3	23.3±1.3	31.7±1.3
	toy	69.4±2.7	76.2±2.1	75.4±0.9	84.1±0.4	12.7±0.8	19.6±1.1
	toy_brick	70.6±0.5	64.3±1.2	61.5±0.3	97.2±0.1	22.9±3.0	29.0±2.7
	transistor1	73.8±1.3	77.9±2.2	75.2±0.4	95.5±0.0	31.1±3.9	38.0±3.8
	u_block	83.8±0.4	78.5±0.8	68.2±0.8	99.0±0.0	49.7±1.3	55.9±1.0
	usb	79.0±0.5	79.3±0.6	70.5±0.7	96.6±0.1	23.9±0.6	34.8±0.5
	usb_adaptor	83.4±0.2	78.9±0.2	70.5±0.6	98.8±0.0	31.3±0.3	36.4±0.4
	vcpill	85.1±0.4	83.7±0.8	72.6±1.0	97.9±0.0	53.8±1.2	60.3±1.2
	wooden_beads	88.3±0.3	88.3±0.2	77.6±0.3	98.7±0.0	49.0±0.4	53.7±0.3
	woodstick	79.7±0.6	65.7±0.9	59.7±0.5	97.7±0.0	53.4±1.1	55.0±0.9
	zipper	95.0±1.0	96.9±0.6	90.8±1.3	97.4±0.2	26.4±0.9	27.0±0.8
	mean	81.8±0.3	80.4±0.2	73.3±0.2	97.1±0.0	36.6±0.1	42.5±0.1

Table 19. Image-level anomaly classification and pixel-level anomaly segmentation results on **Real-IAD** with 2-shot and 4-shot AdaptCLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
2-shot	audiojack	74.7±1.1	66.2±0.9	58.5±1.4	96.9±0.0	29.4±4.2	36.8±4.3
	bottle_cap	80.5±0.6	77.8±0.8	71.9±0.7	98.8±0.0	32.7±0.8	39.2±0.6
	button_battery	76.3±1.0	83.6±0.6	74.4±0.5	98.5±0.0	61.3±0.4	60.1±0.4
	end_cap	77.3±0.4	83.0±0.1	76.3±0.3	97.0±0.0	27.3±0.2	34.5±0.5
	eraser	88.9±1.0	87.1±0.9	76.2±1.3	99.5±0.0	49.9±0.0	50.5±0.0
	fire_hood	87.4±0.3	79.5±0.3	74.5±0.7	99.3±0.0	33.3±0.4	41.4±0.2
	mint	71.9±0.4	71.5±0.8	67.4±0.2	93.2±0.0	27.4±0.7	38.0±0.5
	mounts	86.5±0.1	72.5±0.2	74.5±0.3	98.5±0.0	37.1±0.4	43.3±0.3
	pcb	76.1±0.8	85.0±0.6	76.8±0.3	97.8±0.0	28.2±0.8	35.5±1.0
	phone_battery	89.5±0.1	89.2±0.0	79.6±0.1	99.3±0.0	61.9±0.4	63.4±0.1
	plastic_nut	79.9±0.1	74.0±0.3	65.3±0.2	97.8±0.0	42.0±0.1	44.8±0.2
	plastic_plug	82.6±0.3	78.3±0.4	67.7±0.5	97.4±0.1	28.1±0.1	37.8±0.2
	porcelain_doll	92.9±0.1	89.0±0.2	80.6±0.3	99.6±0.0	49.2±0.2	49.2±0.1
	regulator	74.1±0.8	60.6±1.9	55.7±0.5	96.7±0.1	15.7±2.9	25.2±4.0
	rolled_strip_base	92.4±1.1	95.8±0.8	89.7±0.8	99.6±0.0	37.1±2.5	44.8±2.6
	sim_card_set	96.2±0.2	96.9±0.1	90.8±0.3	99.7±0.0	58.3±0.3	61.6±0.0
	switch	85.9±0.6	89.2±0.6	78.4±0.9	93.9±0.0	36.6±0.4	45.9±0.2
	tape	94.5±0.2	93.4±0.2	85.8±0.4	99.0±0.0	46.0±0.3	51.2±0.1
	terminalblock	86.4±1.2	89.4±0.9	78.6±0.8	98.9±0.0	49.2±0.9	53.1±0.5
	toothbrush	76.2±0.6	82.1±0.4	72.0±0.5	93.6±0.4	23.5±1.9	32.3±2.0
	toy	68.7±3.7	75.7±2.8	74.9±1.4	84.2±0.2	11.4±1.0	17.8±1.1
	toy_brick	71.4±0.2	65.3±0.2	61.7±0.2	97.4±0.1	23.6±0.3	30.1±0.9
	transistor1	74.6±1.3	78.9±1.7	75.2±0.2	95.4±0.2	32.1±2.6	38.7±2.9
	u_block	84.3±0.3	79.2±0.4	68.7±0.4	99.1±0.0	49.6±0.6	55.9±0.5
	usb	81.4±0.3	82.1±0.5	72.7±0.7	96.9±0.0	26.3±0.9	37.1±0.7
	usb_adaptor	84.0±0.1	79.4±0.1	71.5±0.4	98.9±0.0	31.9±0.2	36.5±0.3
	vcpill	88.0±0.5	86.2±0.4	74.9±0.7	98.0±0.0	54.3±0.7	60.6±0.8
	wooden_beads	88.8±0.2	88.7±0.2	78.4±0.3	98.7±0.0	48.7±0.2	53.7±0.2
	woodstick	80.8±0.2	67.3±0.1	61.4±0.2	97.9±0.0	54.6±0.3	55.8±0.2
	zipper	96.2±0.4	97.6±0.3	92.5±0.7	97.7±0.0	27.3±0.2	27.8±0.3
	mean	82.9±0.2	81.5±0.1	74.2±0.1	97.3±0.0	37.8±0.1	43.4±0.2
4-shot	audiojack	75.7±0.6	67.6±1.6	59.4±0.7	97.1±0.1	33.7±3.1	41.2±3.2
	bottle_cap	81.6±0.2	79.1±0.4	73.3±0.3	98.8±0.0	33.2±0.3	39.3±0.2
	button_battery	76.8±0.3	83.9±0.4	74.6±0.4	98.6±0.0	62.5±0.8	60.8±0.7
	end_cap	79.6±0.5	84.6±0.4	77.7±0.3	97.3±0.0	28.2±0.6	34.9±0.4
	eraser	89.9±0.1	88.1±0.0	78.1±0.2	99.6±0.0	49.9±0.4	50.5±0.2
	fire_hood	87.6±0.2	79.9±0.3	74.9±0.3	99.3±0.0	33.4±0.5	41.4±0.4
	mint	72.9±0.1	73.1±0.5	67.7±0.2	93.5±0.0	28.6±0.5	38.3±0.3
	mounts	86.5±0.2	72.1±0.3	75.1±0.3	98.6±0.0	36.6±0.8	42.3±0.4
	pcb	76.5±1.1	85.4±0.6	76.7±0.5	98.1±0.0	31.4±0.9	38.1±1.0
	phone_battery	90.3±0.4	89.8±0.3	81.0±0.6	99.3±0.0	62.9±0.4	63.4±0.2
	plastic_nut	80.4±0.1	74.9±0.1	65.2±0.3	97.9±0.0	41.7±0.2	44.9±0.1
	plastic_plug	82.7±0.1	78.4±0.1	68.4±0.4	97.4±0.1	28.1±0.1	37.9±0.3
	porcelain_doll	92.9±0.2	88.7±0.3	80.6±0.3	99.6±0.0	49.2±0.3	49.0±0.3
	regulator	75.6±0.4	63.5±2.1	56.5±0.7	97.3±0.1	18.8±3.7	29.1±4.6
	rolled_strip_base	93.8±1.1	96.4±0.9	91.5±0.1	99.6±0.0	40.1±1.9	47.3±1.8
	sim_card_set	96.4±0.2	97.0±0.2	91.3±0.4	99.7±0.0	58.1±0.7	61.3±0.3
	switch	87.6±0.5	90.7±0.5	80.2±0.5	94.0±0.0	38.5±0.3	47.6±0.4
	tape	94.9±0.1	93.9±0.3	86.1±0.8	99.0±0.0	45.9±0.3	51.1±0.2
	terminalblock	89.1±0.2	91.4±0.1	81.2±0.5	99.1±0.0	50.2±0.7	53.7±0.4
	toothbrush	77.6±0.4	83.2±0.1	72.9±0.5	93.6±0.1	23.9±0.9	33.0±1.1
	toy	70.4±2.6	78.2±1.8	74.5±0.6	84.3±0.1	16.7±1.7	24.1±1.8
	toy_brick	72.5±0.4	66.7±1.2	62.6±0.4	97.5±0.0	26.4±2.3	31.9±1.7
	transistor1	78.9±1.1	84.1±0.9	76.5±0.5	95.5±0.1	37.2±0.7	43.6±0.5
	u_block	85.0±0.3	80.8±0.4	71.2±0.9	99.2±0.0	50.9±0.5	56.7±0.2
	usb	83.2±0.1	84.0±0.2	75.0±0.0	97.2±0.0	29.6±0.2	39.7±0.1
	usb_adaptor	84.8±0.2	79.9±0.2	72.4±0.5	98.9±0.0	32.0±0.4	36.4±0.4
	vcpill	88.5±0.4	86.9±0.5	76.0±0.7	98.0±0.0	54.8±0.6	61.2±0.4
	wooden_beads	88.9±0.0	88.7±0.1	78.4±0.1	98.7±0.0	48.4±0.1	53.5±0.0
	woodstick	81.0±0.3	67.8±0.5	61.5±0.9	97.9±0.0	55.3±0.2	56.4±0.1
	zipper	97.0±0.0	98.1±0.0	93.9±0.0	97.8±0.0	28.3±0.1	29.4±0.2
	mean	83.9±0.2	82.6±0.0	75.2±0.1	97.4±0.0	39.1±0.3	44.6±0.3

Table 20. Image-level anomaly classification and pixel-level anomaly segmentation results on **BTAD** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	01	96.1	98.6	95.8	93.3	51.1	54.0
	02	81.1	96.9	93.2	93.4	60.5	61.6
	03	95.9	81.0	79.5	94.7	13.8	23.4
	mean	91.0	92.2	89.5	93.8	41.8	46.3
1-shot	01	97.8±0.5	99.2±0.2	96.5±1.0	96.4±0.1	61.2±0.3	60.3±0.5
	02	84.1±0.3	97.5±0.0	93.7±0.0	96.3±0.1	71.8±0.4	65.2±0.3
	03	98.2±0.8	90.7±2.9	84.5±4.0	97.0±0.4	48.6±3.0	50.6±2.0
	mean	93.4±0.0	95.8±0.9	91.6±1.0	96.6±0.2	60.6±1.0	58.7±0.7
2-shot	01	97.9±0.2	99.1±0.0	96.1±0.5	96.5±0.0	61.4±0.1	60.3±0.2
	02	84.0±0.7	97.5±0.1	93.5±0.0	96.3±0.1	71.6±0.2	65.1±0.0
	03	98.5±0.1	91.2±0.5	85.5±0.9	97.3±0.2	49.8±1.7	51.0±1.1
	mean	93.4±0.2	95.9±0.1	91.7±0.5	96.7±0.1	61.0±0.6	58.8±0.4
4-shot	01	97.2±0.9	99.0±0.3	96.1±0.5	96.5±0.0	61.7±0.2	60.4±0.2
	02	83.8±0.3	97.4±0.0	93.6±0.2	96.3±0.0	71.8±0.1	65.3±0.1
	03	98.8±0.1	92.6±0.3	86.6±1.5	97.5±0.1	53.4±0.8	53.2±0.7
	mean	93.3±0.3	96.4±0.1	92.1±0.4	96.8±0.0	62.3±0.3	59.6±0.2

Table 21. Image-level anomaly classification and pixel-level anomaly segmentation results on **KSDD** with zero-shot and few-shot Adapt-CLIP.

Shot	Catergoies	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	electrical commutators	98.1	95.7	92.3	98.1	58.3	58.1
1-shot	electrical commutators	96.9±0.3	91.8±0.2	89.6±0.5	98.2±0.1	57.8±1.2	58.4±0.5
2-shot	electrical commutators	97.2±0.0	92.4±0.4	89.4±0.9	98.1±0.1	57.5±1.1	58.1±0.6
4-shot	electrical commutators	97.0±0.2	91.7±0.9	88.6±1.7	98.0±0.1	56.4±1.4	57.2±0.7

Table 22. Image-level anomaly classification and pixel-level anomaly segmentation results on four **medical datasets, Br35H, Covid, Kvasir and Endo** with zero-shot and few-shot AdaptCLIP.

Shot	Datasets	Anomaly Classification			Anomaly Segmentation		
		I-AUROC	I-AUPR	I-F1$_{max}$	P-AUROC	P-AUPR	P-F1$_{max}$
0-shot	Br35H	94.8	95.1	87.7	-	-	-
	Covid	86.5	54.4	55.4	-	-	-
	Kvasir	-	-	-	82.1	45.3	50.6
	Endo	-	-	-	86.5	52.0	54.5
1-shot	Br35H	93.7±2.4	92.1±3.1	87.8±3.2	-	-	-
	Covid	91.8±2.5	78.8±4.6	73.1±6.6	-	-	-
	Kvasir	-	-	-	83.3±0.5	49.2±3.8	52.1±0.7
	Endo	-	-	-	86.5±1.2	52.4±4.7	55.2±2.0
2-shot	Br35H	94.0±1.7	92.2±2.5	88.3±2.0	-	-	-
	Covid	94.9±0.9	83.6±1.9	78.7±2.4	-	-	-
	Kvasir	-	-	-	83.5±0.8	49.0±3.4	52.6±1.4
	Endo	-	-	-	86.6±1.1	53.1±4.2	55.8±1.7
4-shot	Br35H	93.7±2.0	91.8±2.9	88.3±2.3	-	-	-
	Covid	95.8±0.9	85.8±1.2	81.2±1.1	-	-	-
	Kvasir	-	-	-	83.0±0.5	47.5±2.2	52.1±0.8
	Endo	-	-	-	85.8±0.7	52.2±3.1	55.6±1.2

| Query | GT Mask | 0-shot | 1-shot | 4-shot | | Query | GT Mask | 0-shot | 1-shot | 4-shot |

Figure 10. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **VisA**.

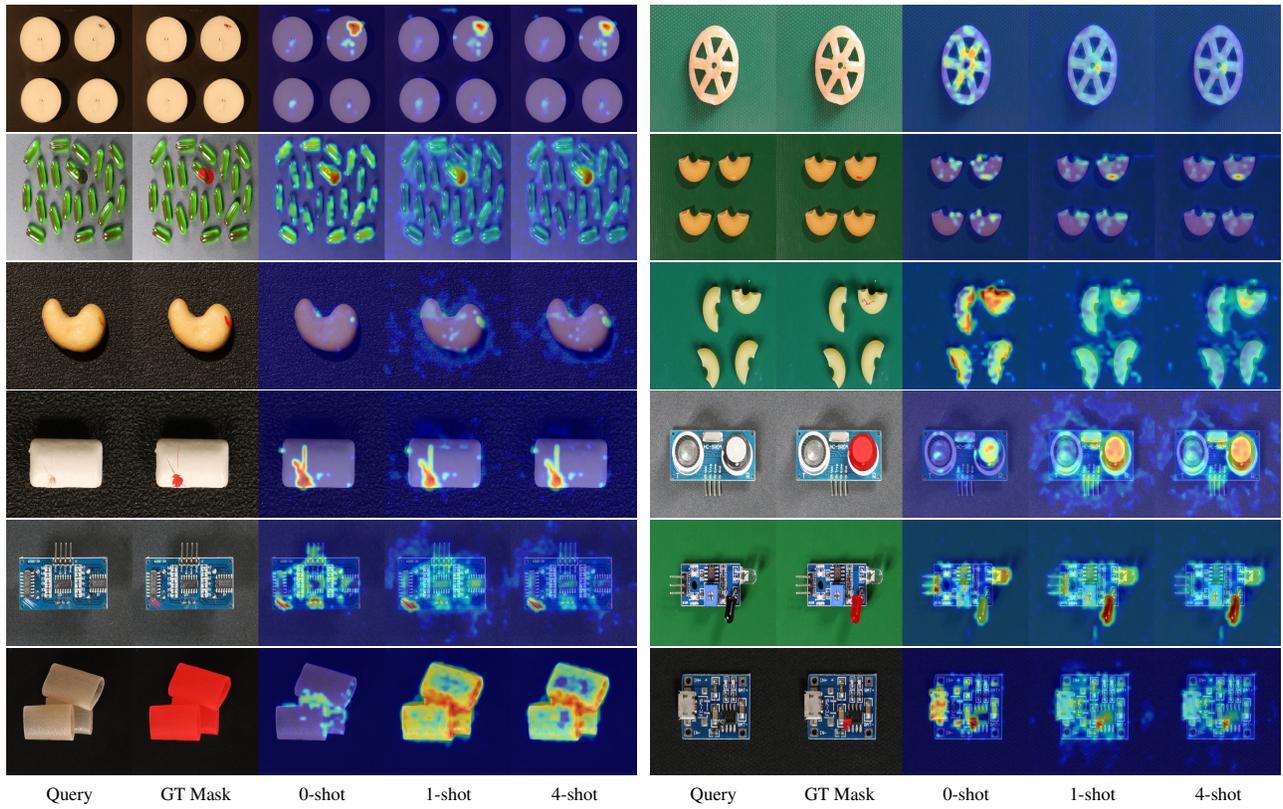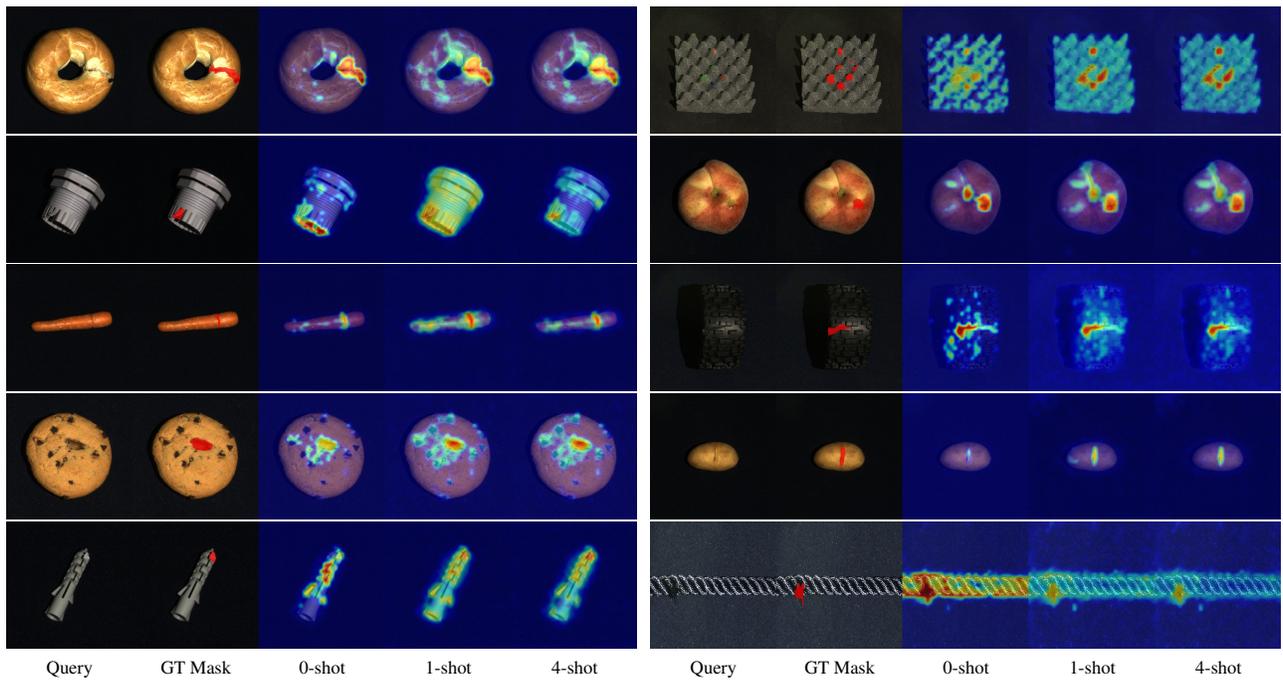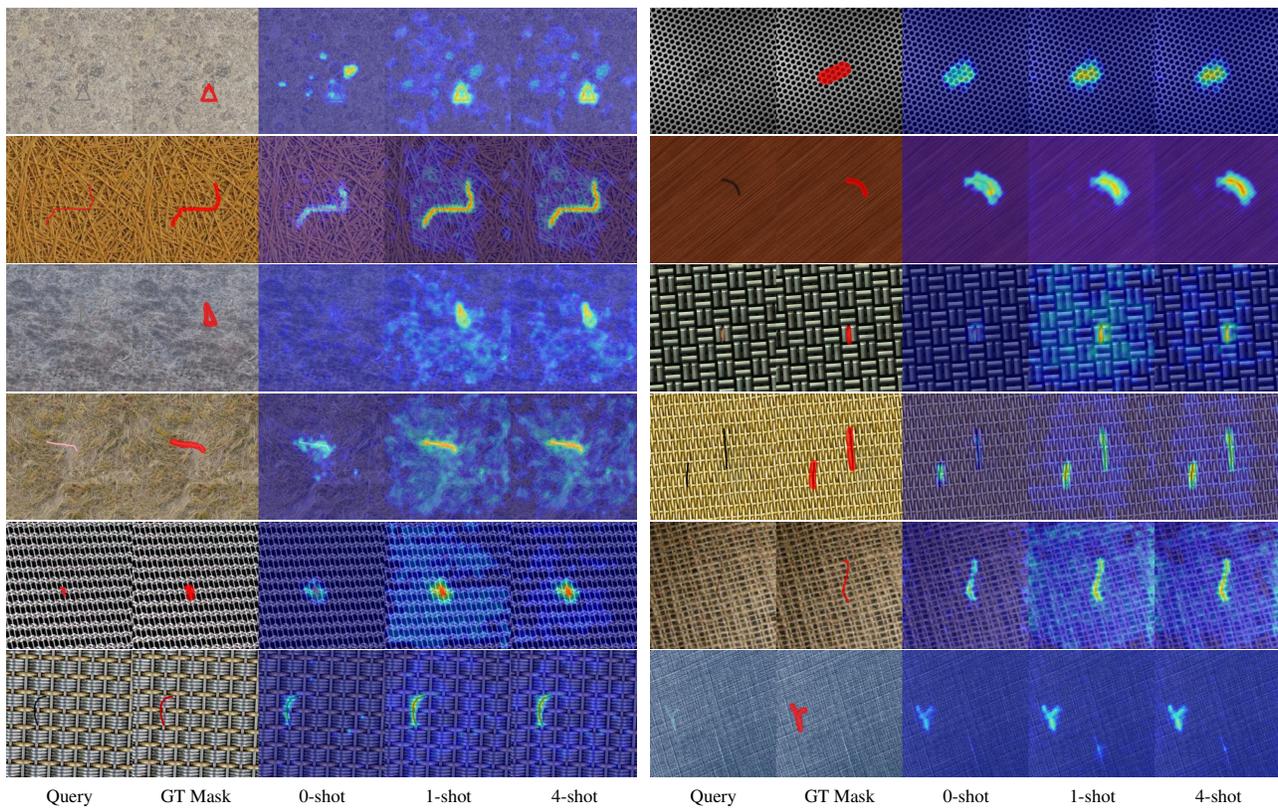| Query | GT Mask | 0-shot | 1-shot | 4-shot | | Query | GT Mask | 0-shot | 1-shot | 4-shot |

Figure 11. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **MVTec 3D**.

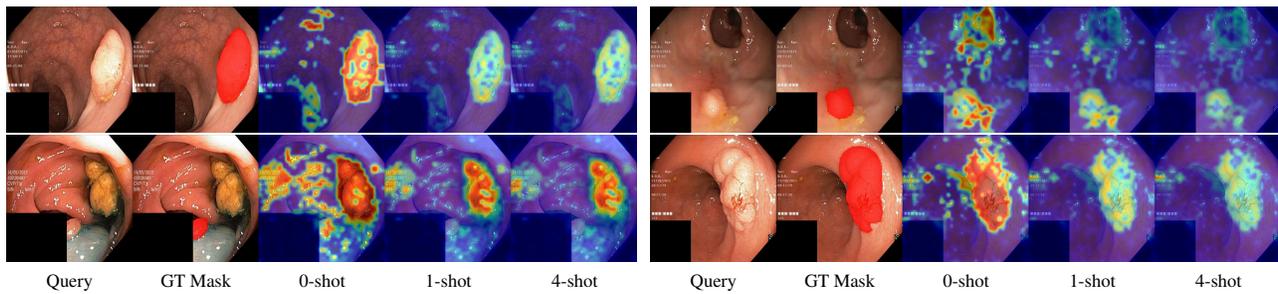Figure 12. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **DTD**.

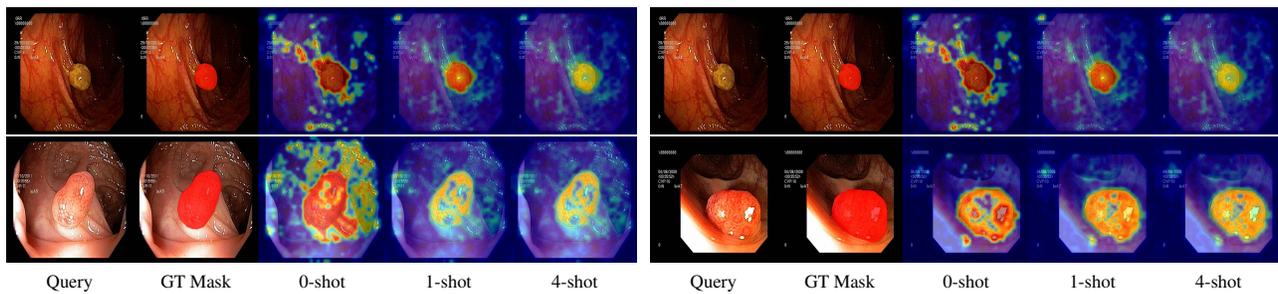Figure 13. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **Kvasir**.

Figure 14. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **Endo**.

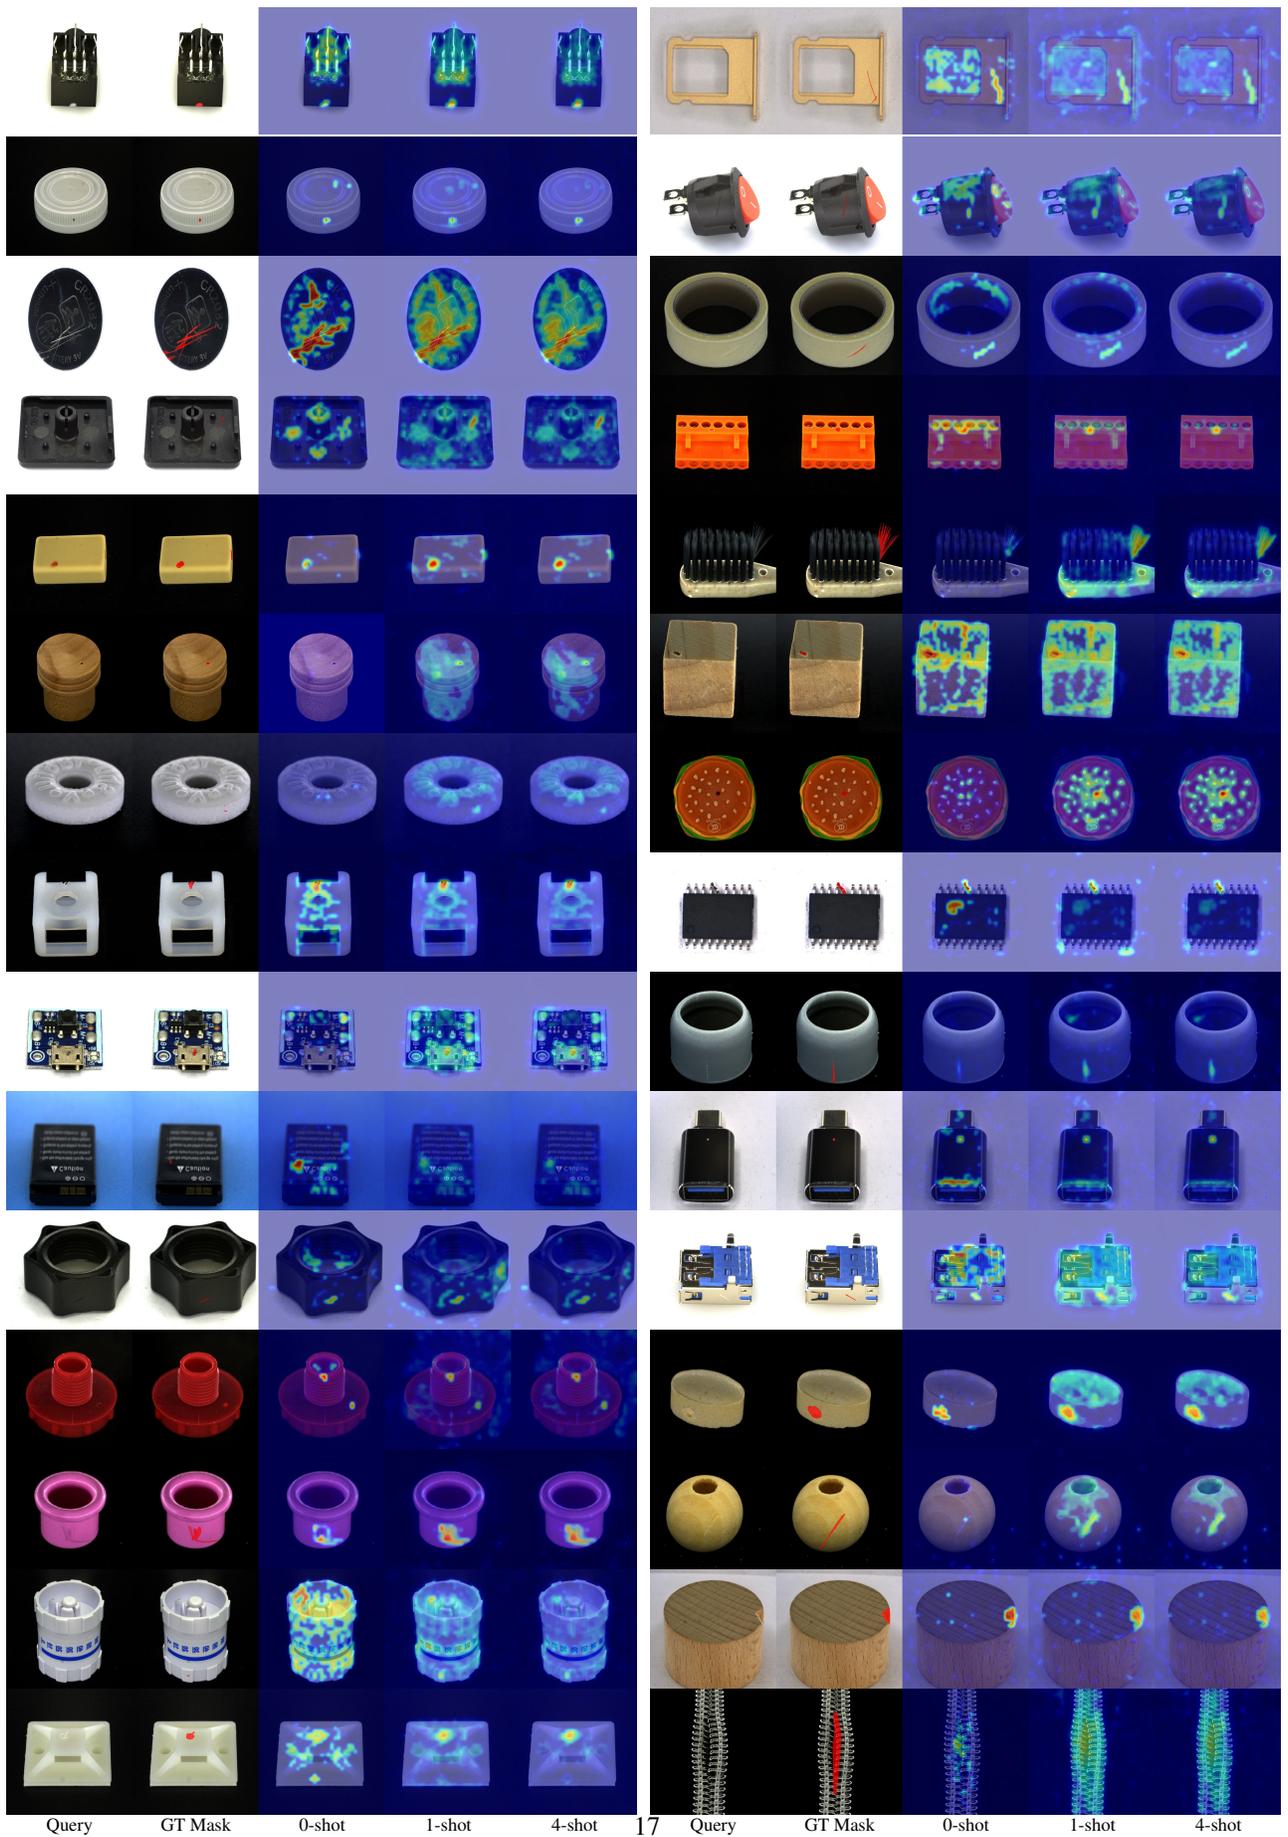| Query | GT Mask | 0-shot | 1-shot | 4-shot | | Query | GT Mask | 0-shot | 1-shot | 4-shot |

Figure 15. Qualitative comparisons of our AdaptCLIP with different prompt numbers on **Real-IAD**.