

MVREC: A General Few-shot Defect Classification Model Using Multi-View Region-Context

Shuai Lyu^{1,2}, Rongchen Zhang^{1,2}, Zeqi Ma^{1,2}
Fangjian Liao^{1,2}, Dongmei Mo^{1,2}, Waikung Wong^{1,2*}

¹ The Hong Kong Polytechnic University, Hong Kong SAR, China

² Laboratory for Artificial Intelligence in Design, Hong Kong SAR, China

{shuai.lyu, rongchen.zhang, mzeqi.ma, fangjian.liao, dongmei.mo}@connect.polyu.hk
calvin.wong@polyu.edu.hk

Abstract

Few-shot defect multi-classification (FSDMC) is an emerging trend in quality control within industrial manufacturing. However, current FSDMC research often lacks generalizability due to its focus on specific datasets. Additionally, defect classification heavily relies on contextual information within images, and existing methods fall short of effectively extracting this information. To address these challenges, we propose a general FSDMC framework called MVREC, which offers two primary advantages: (1) MVREC extracts general features for defect instances by incorporating the pre-trained AlphaCLIP model. (2) It utilizes a region-context framework to enhance defect features by leveraging mask region input and multi-view context augmentation. Furthermore, Few-shot Zip-Adapter(-F) classifiers within the model are introduced to cache the visual features of the support set and perform few-shot classification. We also introduce MVTEC-FS, a new FSDMC benchmark based on MVTEC AD, which includes 1228 defect images with instance-level mask annotations and 46 defect types. Extensive experiments conducted on MVTEC-FS and four additional datasets demonstrate its effectiveness in general defect classification and its ability to incorporate contextual information to improve classification performance.

Code — <https://github.com/ShuaiLYU/MVREC>

Introduction

Defect detection and classification (Jha and Babiceanu 2023) is a critical challenge in industrial manufacturing, as it involves identifying and categorizing defects within workpieces. However, in practical applications, the diversity of defect types and the low frequency of defect occurrences make it a particularly difficult task.

While Few-shot Learning (FSL) (Wang et al. 2020) has gained traction in general vision tasks like mini-Imagenet, its application to defect multi-classification (FSDMC) remains challenging. This disparity is evident in the limited availability of dedicated datasets and research focusing on FSDMC. Although Contrastive Vision-Language Pre-training (CLIP) (Radford et al. 2021) has demonstrated remarkable success in learning visual features from large-scale image-text pairs and adapting to downstream tasks

*Corresponding author.




Figure 1: Comparison of two different Classification models.

with few-shot learning, this type of application is nearly absent in the context of FSDMC. This is primarily due to the significant domain gap between general vision tasks and FSDMC. Secondly, defects inherently differ from normal surface areas, necessitating more contextual information for effective detection and classification. However, common classification models often involve cropping the defect region, resizing it to the model input size, and feeding it into a network, as shown in Figure 1 (a). This pre-treatment fails to retain important contextual information, such as the surrounding background and the size of the defect. The most popular multi-category datasets (Bergmann et al. 2019, 2022; Kim et al. 2020) with different product images are typically designed for anomaly detection rather than defect classification. Although the field of few-shot defect multi-classification has attracted considerable research attention (Cao et al. 2023; Zhan, Zhou, and Xu 2022; Zhao et al. 2023; Zhou et al. 2023; Liu et al. 2023; Xiao et al. 2022), the datasets used, such as the NEU-DET Dataset and the MTD Dataset, are limited by their focus on a single product category. There is a notable scarcity of multi-category datasets specifically proposed for FSDMC.

To mitigate these issues, we propose a general few-shot defect classification model using a multi-view region-context approach, called MVREC. Specifically, our approach begins by generating region-context visual features for the defect instance using the AlphaCLIP model (Sun et al. 2024), a transformer-based model that takes a defect image and its mask context as input to generate visual features from the masked region. By incorporating the mask region context, the network can perceive both the defect foreground region and its surrounding background, generating

target-specific features while maintaining input consistency. Furthermore, we propose a multi-view augmentation technique to generate multi-view features for a defect, maximizing the utility of few-shot samples and enhancing generalization ability. The multi-view region-context (MVREC) features can be extracted from the multi-view patches and masks of the defect instance, thereby enhancing the region-context features. Moreover, we propose two few-shot classifiers: the training-free Zip-Adapter, which predicts directly without training, and the fine-tuning Zip-Adapter-F, which adapts the MVREC features for better performance. Zip-Adapter and Zip-Adapter-F share the same structure, consisting of a Zero-initialized Projection (ZIP) module and a Scale-Dot-Product Attention (SDPA) module. Specifically, they store visual features and corresponding class labels from the support set images as key-value pairs. The SDPA module then calculates the visual feature similarity between the query defect instance and the support defects, outputting the classification logits through the weighted sum of the encoded labels from the support set. The ZIP module serves as an identity mapping and feature adapter, respectively, for Zip-Adapter and Zip-Adapter-F. Furthermore, we propose MVTec-FS, based on MVTec AD, to create a multi-category dataset suitable for the FSDMC task. This dataset features a diverse array of defect types and a balanced distribution. MVTec-FS includes 15 categories of product surface images and approximately 46 types of defects, making it a promising new benchmark in this field.

We tested MVREC across MVTec-FS and four other public defect datasets with classification annotations. Our results demonstrate superior performance in few-shot defect classification, outperforming existing models. In summary, our contributions can be outlined as follows:

- (1) We employ AlphaCLIP to extract general features from each defect instance, enhancing model generalizability, and design a new region-context-based defect classification framework that fully incorporates contextual information for more accurate defect classification.
- (2) We introduce multi-view context augmentation and Zip-Adapter(-F) classifiers for few-shot classification.
- (3) We reconstruct the popular MVTec AD dataset into a new FSDMC benchmark named MVTec-FS.
- (4) We conducted extensive experiments on multiple defect datasets, demonstrating the effectiveness of MVREC.

Related work

Defect Detection and Classification

Various models, including object detection, segmentation, and classification, have been applied to defect detection. In recent years, the MVTec AD (Bergmann et al. 2019) dataset has been widely studied for anomaly detection tasks (Lyu, Mo, and Wong 2024; Pang et al. 2022). These models learn from normal samples to identify anomalies. However, defect classification, which involves identifying specific defect types, is more challenging due to the rarity and diversity of defects and the limited availability of relevant datasets. FSDMC models, such as CAO (Cao et al. 2023), Fabric (Zhan, Zhou, and Xu 2022), and FANet (Zhao et al.

2023), have been proposed. However, these methods are often dataset-specific and require complex training processes, including meta-learning and metric learning. Additionally, using a subset of defect types as base classes to train a base model and then evaluating novel classes is common, but this approach may not be practical for real-world applications.

Clip-based Few-shot Classification

The most common few-shot methods include meta-learning and metric learning. Meta-learning methods learn a model that can quickly adapt to new tasks with minimal training data. Metric learning methods learn a distance metric that can effectively measure the similarity between samples. Recently, large language models and multi-modal pre-training models, such as GPT (Ouyang et al. 2022; Zhu et al. 2023) and CLIP (Radford et al. 2021), have emerged as powerful tools. Related research (Gu et al. 2024; Jeong et al. 2023) has been applied to defect detection tasks, showing impressive performance. Numerous adapter-based methods have been proposed to adapt the CLIP model to specific tasks with few samples, such as CLIP-Adapter (Gao et al. 2024), Tip-Adapter (Zhang et al. 2022), CoOp (Zhou et al. 2022), and SuS-X (Udandarao, Gupta, and Albanie 2023), but most of these methods are designed to jointly learn image and text features for general vision tasks.

Region-Context-based Models

Traditional classification networks typically evolve by cropping the defect region and resizing it, without explicitly utilizing the region context. When defects vary in size, as shown in the example in Figure 1, cropping and resizing the defect region may result in the loss of crucial contextual information. To address this issue, region-context models incorporate region context as a prompt to predict target information, thereby preserving the contextual information of the target. For example, the SAM network uses prompts in the form of points and bounding boxes. To enable CLIP to focus on specific regions within the entire image, various methods (Zhong et al. 2022; Zhou, Loy, and Dai 2022; Shtedritski, Rupprecht, and Vedaldi 2023; Sun et al. 2024) have been explored. AlphaCLIP is an innovative enhancement of the CLIP model, designed to improve its ability to focus attention on specific regions (Sun et al. 2024). This architecture enables AlphaCLIP to provide precise control over the emphasis of image content.

MVREC

In this chapter, we introduce the MVREC feature extraction module and present the training-free Zip-Adapter classifier, along with its fine-tuning version, Zip-Adapter-F.

Multi-View Region-Context Feature Extraction

We first introduce the multi-view region-context (MVREC) feature extraction process for defect instances. To effectively capture the visual representation of defects and explicitly mine contextual information, we employ the pre-trained AlphaCLIP model to extract visual features from images using their mask prompts. Given the small data

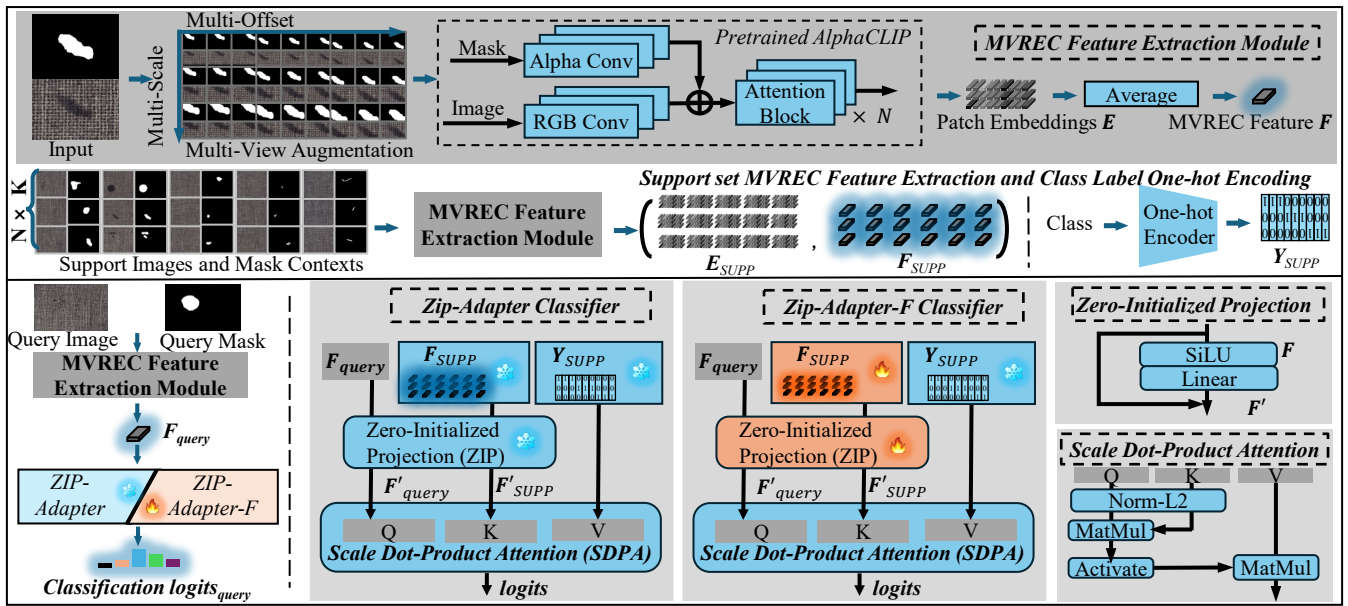


Figure 2: The framework of MVREC. First, the MVREC feature extraction is introduced. Given an N -way K -shot task, the MVREC features for the support set are collected. Then the inference process with Zip-Adapter or Zip-Adapter-F is shown.

volume characteristic of few-shot learning tasks, we utilize multi-view context augmentation to generate multi-view patches of defect images, thereby expanding the available dataset for subsequent processing. Specifically, we employ two context augmentation methods to achieve this. The first method, multi-scale augmentation, involves cropping Num_{scale} patches at different scales from the defect patches and their corresponding masks, centered on the defect. The second method involves offsetting the center of the defect to generate Num_{offset} defect patches with different offsets at each scale. By applying these two augmentation methods, a total of $V = \text{Num}_{scale} \times \text{Num}_{offset}$ patches can be obtained for each defect instance. The AlphaCLIP model extracts MVREC patch embeddings $E \in \mathbb{R}^{V \times C}$ from the multi-view patches, where C is the number of feature channels. These embeddings are then averaged to produce a single MVREC feature $F \in \mathbb{R}^C$.

Support set MVREC Feature Extraction and Class Label Encoding

For N -way K -shot classification tasks, the MVREC patch embeddings $E_{SUPP} \in \mathbb{R}^{NK \times V \times C}$ and MVREC features $F_{SUPP} \in \mathbb{R}^{NK \times C}$ are first extracted for the support set. Then, the one-hot encoded class labels $Y_{SUPP} \in \mathbb{R}^{NK \times N}$ are extracted. The MVREC features F_{SUPP} and Y_{SUPP} are used to build the cached key-value pairs for the FSDMC task. Additionally, the MVREC patch embeddings E_{SUPP} and the one-hot encoded labels Y_{SUPP} are used as training data to fine-tune the Zip-Adapter.

Training-free Zip-Adapter Classifier

In this section, we introduce the method of utilizing MVREC visual features to construct the zero-initialized pro-

jection classifier (Zip-Adapter) for FSDMC tasks. The Zip-Adapter classifier consists of a zero-initialized projection (ZIP) module and a scaled dot-product attention (SDPA) module. It stores the MVREC features F_{SUPP} along with the encoded labels Y_{SUPP} of the support set samples. The ZIP module includes a single linear layer, a residual connection, and a SiLU activation function, with the linear layer initialized to zeros. The output of the ZIP module is the adapted feature F' , generated as follows:

$$F' = \text{SiLU}(\text{Linear}(F)) + F, \quad (1)$$

here, F represents the MVREC feature for either the support sample or the query sample. In the Zip-Adapter, the ZIP module is designed to serve as an identity transformation by initializing the linear layer with zeros and using a residual connection. The SDPA module is a scaled dot-product attention mechanism, which calculates the visual similarity between the query defect instance and the support set. It then produces the classification logits by performing a weighted sum of the support encoded labels Y_{SUPP} . The SDPA module is defined as follows:

$$\text{logits}_{query} = Y_{SUPP} \cdot \psi(\text{Sim}(F'_{query}, F'_{SUPP})), \quad (2)$$

where the ψ is the activation function (Zhang et al. 2022) for modulating the cosine similarity:

$$\psi(x) = \exp(-\beta(1-x)), \quad (3)$$

β controls the sharpness of the curve. And Sim is the cosine similarity function. The output of the SDPA module, logits_{query} , represents the classification logits of the query defect instance. The class with the highest logit value is identified as the predicted class.

Training Zip-Adapter-F Classifier

Zip-Adapter-F enhances the visual features for better performance by fine-tuning the Zip-Adapter classifier, making

both the ZIP module and the cached visual features of the support set learnable. Our Zip-Adapter-F combines a cache-based mechanism with an adapter-based mechanism, using the Zip-Adapter as the base model. The fine-tuning process involves two training objectives: (1) optimizing the cross-entropy (CE) loss between the predicted logits $logits_{query}$ and the labels Y_{query} . The CE loss \mathcal{L}_{CE} is defined as:

$$\mathcal{L}_{CE}(logits_{query}, Y_{query}) = - \sum_i y_i \log(p_i), \quad (4)$$

where y_i and p_i represents the label and predicted probability distribution for class i .

The second part uses the triplet loss to optimize the intra-class compactness and inter-class separability of the adapted feature $f_{adapted}$ of the ZIP module within a batch. The triplet loss $\mathcal{L}_{triplet}$ is defined as:

$$\mathcal{L}_{triplet}(F'_{query}) = \max(d(F'_{anchor}, F'_{positive}) - d(F'_{anchor}, F'_{negative}) + \alpha, 0), \quad (5)$$

where F'_{anchor} , $F'_{positive}$, and $F'_{negative}$ are the embeddings (feature vectors) of an anchor sample, a positive sample (same class as an anchor), and a negative sample (different class from anchor), within a batch, respectively. $d(\cdot, \cdot)$ is a distance function used to measure the similarity between embeddings. α is a margin hyperparameter that specifies the minimum difference between the distances of positive and negative pairs required for the loss to be zero. The overall loss function for finetuning Zip-Adapter-F is:

$$\mathcal{L}_{Zip-Adapter-F} = \mathcal{L}_{CE} + \lambda \cdot \mathcal{L}_{triplet}, \quad (6)$$

where λ is a hyperparameter that balances the importance of the two parts in the overall loss. After Zip-Adapter-F is trained, it can be used to classify query defect instances similar to the Zip-Adapter classifier.




Figure 3: Some modified cases are displayed, in which there are multiple defect instances with different types.

MVTec-FS Dataset

Although few-shot defect multi-classification has garnered considerable research attention, datasets like NEU-DET (He et al. 2020) and MTD (Huang, Qiu, and Yuan 2020) are limited to a single product category. Recently, anomaly detection has also gained attention, with several multi-category datasets (Bergmann et al. 2019, 2022; Kim et al. 2020) being proposed. However, these datasets are not designed for defect classification. The MVTec AD (Bergmann et al. 2019) dataset, the most popular benchmark for anomaly detection, features 15 product categories (5 textiles and 10 objects), offering significant diversity and generalization. In its original configuration, the training set consists of normal images, while the testing set includes both normal and defect

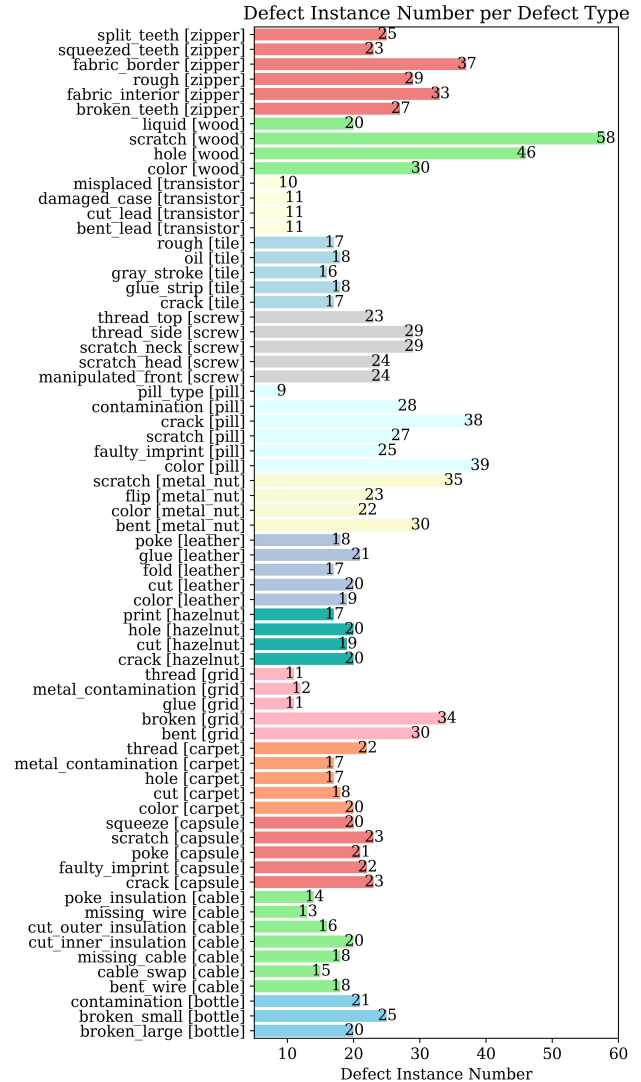


Figure 4: Details of MVTec-FS dataset.

images, with defect images labeled by masks. This dataset contains about 47 defect types (ranging from 8 to 26 images per type), making it suitable for FSDMC tasks. However, FSDMC tasks have rarely been studied on this dataset.

We selected 1,228 defective images from the MVTec AD and labeled them with instance-level masks, creating a new benchmark dataset named MVTec-FS. Since the toothbrush category contains only one defect type, it was excluded from MVTec-FS. The number of defect instances per defect type is presented, totaling 46 types with instances ranging from 9 to 58, as shown in Fig 4. The original annotations were at the image level and did not account for multiple defects within a single image. We used the connected component algorithm to convert image-level masks to instance-level masks, followed by necessary human corrections. Some examples are shown in Figure 3. For each defect type, 50% of the defects are used as the training set to sample the support set, while the remaining 50% constitute the testing set (query set).

Experiments

Experiments Setting

We conducted experiments on the MVTec-FS dataset and four other datasets to evaluate our MVREC using accuracy metrics. The few-shot setup is defined as N-way K-shot, where K is set to 1, 3, or 5. The support set is sampled from the training set, and the query set consists of all images in the testing set. The evaluation was conducted on the query set for each of the five sampled support sets, and the average classification accuracy was calculated to provide a more robust assessment. Ablation studies were performed on the MVTec-FS dataset to assess the effectiveness of the various components. For AlphaCLIP, we selected the ViT-L/14 (Dosovitskiy et al. 2020) backbone. The MVREC visual feature used three scales, representing the commonly used large, medium, and small settings. The number of offsets was set to 9, based on a grid layout similar to a tic-tac-toe board. We set β to 32 for the Zip-Adapter and 1 for the Zip-Adapter-F classifiers. When training the Zip-Adapter-F, we used the AdamW optimizer with a learning rate of 0.0001. The model was updated for 500 iterations, training on all MVREC features of the support set in each iteration. For the triplet loss item, the hyperparameters α and λ were set to 0.5 and 4, respectively.

In our experiment, we evaluated a variety of baseline classifiers based on the AlphaCLIP backbone, including: 1. **CLIP-ZeroShot** (Radford et al. 2021): This approach leverages the zero-shot capability of the CLIP model. We generated text embeddings for each class description and computed the similarity between the test image embeddings and these text embeddings, classifying them based on the highest similarity. 2. **CLIP-KNN**: This method uses the K-Nearest Neighbors algorithm with the CLIP features. The most similar K (K=1) support samples are retrieved for a query sample, and the class with the majority vote is selected as the prediction. 3. **CLIP-ProtoNet**: This approach builds a Prototypical Network (Snell, Swersky, and Zemel 2017) on top of the CLIP, where proxy features representing each class are calculated from the support set, and test images are classified based on their similarity to these class proxies. 4. **CLIP-Adapter** (Gao et al. 2024): This method involves adding adapter layers on top of the CLIP. These layers are trained for the new classification task, adjusting the image features accordingly. 5. **Tip-Adapter** (Zhang et al. 2022): This method constructs a key-value cache model using CLIP-extracted features from the few-shot data and performs recognition in a retrieval-based manner. Tip-Adapter-F treats the visual cache as learnable parameters and optimizes them to improve performance.

Results on MVTec-FS Dataset

In Table 1, we present the classification accuracy (%) of various few-shot models evaluated on the MVTec-FS dataset under different few-shot learning configurations. To ensure a fair comparison, we report results across several few-shot settings, specifically with 0, 1, 3, and 5 shots. The accuracies for 14 product categories, as well as the average accuracy, are reported in separate columns. As shown in the table,

our MVREC demonstrates outstanding performance in all few-shot setups, regardless of whether Zip-Adapter or Zip-Adapter-F is used. The Zip-Adapter-F achieves the highest accuracy of 89.4% with 5 shots, which is 6.9% higher than the second-best, LinearProb. By comparing different product categories, it is evident that the Zip-Adapter-F achieves the highest accuracy in most categories, demonstrating its effectiveness in few-shot learning scenarios.

Ablation Study

Contributions of MVREC feature and Zip-Adapter(-F). First, we assess the effectiveness of the MVREC feature and Zip-Adapter(-F) classifiers by comparing their performance to other classifiers, regardless of MVREC usage. As shown in Table 2, the MVREC feature consistently boosts the performance of all classifiers across different few-shot settings, with the most significant gain of 11.6% in CLIP-Adapter with 1-shot, demonstrating its general effectiveness for few-shot defect classification. Zip-Adapter-F consistently outperforms most classifiers, regardless of MVREC use, highlighting its inherent strength. When combined with MVREC, Zip-Adapter-F achieves the best results across all few-shot settings, maximizing its potential and making it the most effective approach for FSDMC. Notably, Zip-Adapter and Tip-Adapter yield identical results before training, as they are mathematically equivalent at that stage.

Mask Region-Context. We investigate the impact of the mask region-context on model performance. As previously mentioned, the mask region-context helps the model focus on the defect instance without cropping the region based on defect size, which could otherwise result in the loss of contextual information. We evaluate two scenarios where the mask region-context is removed: (1) using a whole-foreground mask as the region-context input to AlphaCLIP, and (2) using CLIP without any mask region-context. The results, shown in Table 3, demonstrate that removing the mask context leads to a noticeable decrease in accuracy. We also consider the impact of different cropping styles. When cropping by defect size and using vanilla CLIP, the worst results are obtained, further emphasizing the importance of mask region-context. Cropping by fixed size and using AlphaCLIP as the feature extractor achieves the best performance, highlighting the effectiveness of MVREC.

Multi-View Context Augmentation. From the results in Table 3, we observe that different augmentation methods have varying impacts on classification accuracy. When single augmentations are used, multi-scale, multi-offset, and multi-rotation augmentations show significant improvements in both Zip-Adapter and Zip-Adapter-F. When double augmentations are applied, the combination of multi-scale and multi-offset yields the best results, indicating that these augmentations are complementary and can be combined to achieve better performance. Multi-scale augmentation allows the model to learn features at various resolutions, which is crucial for capturing both fine and coarse details in the images. Meanwhile, multi-offset augmentation helps the model learn robust features by shifting the image and mask context, thereby improving the model’s robustness.

FS	Classifier	Carpet	Grid	Leather	Tile	Wood	Bottle	Cable	Capsule	Hazelnut	MetalNut	Pill	Screw	Transistor	Zipper	Average
0	CLIP-ZeroShot	25.0	51.0	43.2	42.9	75.7	34.4	21.1	20.8	62.9	36.0	14.8	23.3	33.3	39.0	37.4
1	CLIP-Adapter	66.4	42.9	65.5	70.5	64.0	51.9	61.8	44.5	66.3	59.6	54.3	72.7	63.8	77.3	61.5
	CLIP-ProtoNet	60.5	42.0	62.7	71.0	63.4	50.6	62.8	46.8	62.9	62.0	59.3	74.7	55.2	74.6	60.6
	CLIP-KNN	59.6	42.0	62.7	69.5	62.9	50.6	62.8	47.2	62.9	61.6	58.5	75.0	55.2	74.4	60.4
	CLIP-LinearProb	60.5	41.6	64.6	73.3	68.0	55.0	62.8	49.1	69.1	60.0	63.7	73.0	56.2	73.9	62.2
	Tip-Adapter	60.0	42.5	62.7	70.0	63.4	50.6	62.8	47.2	62.9	62.0	59.3	74.7	55.2	74.6	60.6
	Tip-Adapter-F	60.9	44.9	63.2	72.9	64.6	51.3	63.2	46.8	65.1	61.2	62.7	74.3	61.0	76.6	62.0
	Zip-A (MVREC)	73.6	49.8	79.5	94.3	69.4	58.8	80.4	57.7	77.1	71.2	66.9	68.0	77.1	79.0	71.6
Zip-A-F (MVREC)	78.6	50.6	82.3	96.2	71.4	58.1	77.5	60.4	77.1	71.6	72.4	67.3	89.5	79.3	73.7	
3	CLIP-Adapter	70.9	57.6	77.3	89.5	81.4	68.1	77.9	61.5	76.0	72.0	68.6	85.7	88.6	85.1	75.7
	CLIP-ProtoNet	71.4	57.6	78.2	86.7	82.3	70.0	80.0	63.0	77.7	71.2	69.1	87.0	85.7	83.4	76.0
	CLIP-KNN	71.4	57.6	78.2	86.7	82.3	70.0	80.0	63.0	77.7	71.2	69.1	87.0	85.7	83.4	76.0
	CLIP-LinearProb	74.1	60.0	82.3	91.0	83.1	71.3	82.1	63.4	78.3	73.6	75.6	86.7	87.6	83.4	78.0
	Tip-Adapter	65.5	55.9	77.3	86.2	75.7	68.1	79.0	60.4	77.7	73.2	67.4	86.3	84.8	74.9	73.7
	Tip-Adapter-F	72.7	60.8	80.5	89.1	82.3	71.9	81.8	61.5	78.3	73.2	73.3	87.0	87.6	85.4	77.5
	Zip-A (MVREC)	81.8	56.7	87.7	97.1	82.9	63.1	91.6	63.4	76.6	78.4	74.8	80.0	97.1	82.9	79.6
Zip-A-F (MVREC)	85.0	71.8	90.9	97.6	90.0	76.9	93.0	74.0	82.3	83.6	83.5	88.3	100.0	88.8	86.1	
5	CLIP-Adapter	75.5	64.9	87.7	91.0	86.9	67.5	83.9	69.8	78.9	78.8	77.3	90.7	93.3	87.6	81.0
	CLIP-ProtoNet	73.6	59.6	83.6	89.1	84.9	67.5	84.9	72.1	74.9	78.4	74.6	89.3	99.1	87.1	79.9
	CLIP-KNN	74.1	55.5	81.8	89.5	79.7	65.6	78.3	62.3	73.1	79.2	67.6	89.7	97.1	79.8	76.7
	CLIP-LinearProb	79.6	66.9	88.6	93.3	88.3	70.0	86.7	69.8	78.9	80.0	81.2	92.7	92.4	86.3	82.5
	Tip-Adapter	72.7	59.2	80.9	88.1	84.9	63.8	82.1	65.7	77.7	78.4	67.6	90.3	98.1	78.3	77.7
	Tip-Adapter-F	74.5	65.3	88.2	91.0	89.7	68.1	85.6	70.2	77.1	80.8	80.7	91.0	96.2	87.6	81.9
	Zip-A (MVREC)	84.5	60.0	89.6	97.6	89.4	61.9	93.0	75.1	84.0	82.8	75.1	88.0	99.1	83.2	83.1
Zip-A-F (MVREC)	85.9	80.8	92.7	97.6	96.6	77.5	93.0	81.1	88.6	91.2	84.7	92.0	100.0	90.0	89.4	

Table 1: Classification accuracy (%) on MVTec-FS of different models. AlphaCLIP is used to extract the visual feature for all classifiers. Zip-A and Zip-A-F stands for Zip-Adapter and Zip-Adapter-F for short. The best results are highlighted in bold.

Different Training Setting of Zip-Adapter(-F). As shown in Figure 5, the combination of trainable support features and a trainable ZIP module resulted in the highest accuracy across all few-shot setups.

Visualization

To better illustrate the function of MVREC, we used t-SNE (Van der Maaten and Hinton 2008) to visualize the support MVREC features in Zip-Adapter-F, as shown in Figure 5. Different colors represent 5 defect classes from the 5-shot leather images of the MVTec-FS. The changes in the distribution indicate that multi-view augmentation and fine-tuning help the model learn more discriminative features.




Figure 5: t-SNE projections of the MVREC features F_{SUPP} for support set. From left to right are 1) F_{SUPP} without multi-view augmentation. 2) F_{SUPP} 3) finetuned F_{SUPP} .

Comparison on Other Datasets

We evaluated MVREC on several public datasets, including: 1) **NEU-DET** (He et al. 2020), a metal surface defect dataset for detection model research; 2) **PCB Defect**

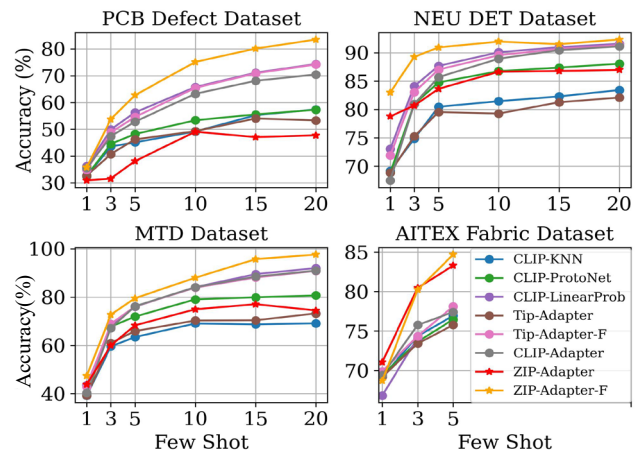


Figure 6: Classification accuracy (%) on other four datasets of different models with quantitative values.

Dataset (Huang and Wei 2019), released by The Open Lab on Human-Robot Interaction of Peking University; 3) **Magnetic Tile Surface Defects (MTD)** (Huang, Qiu, and Yuan 2020), which contains 6 common magnetic tile defects; and 4) **AITEX Fabric Defect** (Silvestre-Blanes et al. 2019), a fabric defect dataset with 12 types of defects, from which seven defect types with at least 10 samples are selected. For each dataset, 50% of the data is used as the training set for sampling the support set, and the other 50% is used as the

FS	MVREC	CLIP-Adapter	CLIP-ProtoNet	CLIP-KNN	CLIP-LinearProb	Tip-Adapter	Tip-Adapter-F	Zip-Adapter	Zip-Adapter-F
1	✗	61.5	60.6	60.4	62.2	60.6	62.0	60.6	62.4
	✓ Gain	73.1 +11.6	71.6 +11.0	71.3 +10.9	71.8 +9.6	71.6 +11.0	73.0 +11	71.6 +11.1	73.7 +11.3
3	✗	75.7	76.0	71.0	78.0	73.7	77.5	73.7	77.9
	✓ Gain	85.7 +10.0	83.4 +7.6	78.9 +7.9	84.8 +6.8	79.6 +5.9	85.7 +8.2	79.6 +5.9	86.1 +8.2
5	✗	81.0	79.9	76.7	82.5	77.7	81.9	77.7	82.2
	✓ Gain	89.2 +8.2	86.1 +6.2	82.6 +5.9	88.1 +5.6	83.1 +5.4	89.2 +7.3	83.1 +5.4	89.4 +7.2

Table 2: Gains of MVREC Representation and Zip-Adapter(-F) Classifiers across different Few-Shot Settings.

Classifier	Crop Size	Feature Extractor	Few Shot Setup		
			1	3	5
Zip-Adapter	Defect	CLIP	57.7	65.5	68.5
	Fixed	CLIP	58.1	64.8	68.6
	Fixed	AlphaCLIP Wo.M	61.8	68.0	70.1
	Fixed	AlphaCLIP	71.6	79.6	83.1
Zip-Adapter-F	Defect	CLIP	61.4	77.0	82.0
	Fixed	CLIP	65.1	80.6	85.5
	Fixed	AlphaCLIP Wo.M	66.3	79.6	85.4
	Fixed	AlphaCLIP	73.7	86.1	89.4

Table 3: Classification accuracy (%) on MVTec-FS with different region-context handling methods. Wo.M stands for Without Mask. Cropping by fixed size and using AlphaCLIP with a masked region-context both preserves and effectively utilizes the context.

Classifier	Augmentation				Few Shot Setup		
	Scale	Rotate	Flip	Offset	1	3	5
Zip-Adapter	✗	✗	✗	✗	60.6	73.7	77.7
	✓	✗	✗	✗	68.1	78.3	81.4
	✗	✓	✗	✗	68.5	77.9	82.1
	✗	✗	✓	✗	55.3	64.8	68.2
	✗	✗	✗	✓	70.7	78.5	80.7
	✓	✓	✗	✗	70.5	78.4	82.0
	✓	✗	✓	✗	60.3	71.1	74.5
	✓	✗	✗	✓	71.6	79.6	83.1
Zip-Adapter-F	✗	✗	✗	✗	62.4	77.9	82.2
	✓	✗	✗	✗	69.3	82.9	86.4
	✗	✓	✗	✗	70.5	83.2	86.8
	✗	✗	✓	✗	54.7	72.6	78.3
	✗	✗	✗	✓	72.4	83.8	87.1
	✓	✓	✗	✗	73.1	84.3	88.4
	✓	✗	✓	✗	62.6	79.5	85.0
	✓	✗	✗	✓	73.7	86.1	89.4

Table 4: Classification accuracy (%) on MVTec-FS of different augmentation methods.

testing set (query set). In addition to 1, 3, and 5 shots, we also evaluated performance with 10, 15, and 20 shots on the NEU-DET, PCB Defect, and MTD datasets for a more comprehensive comparison. The results, in Figure 6, demonstrate that Zip-Adapter-F (MVREC) achieves the best performance on all datasets, with performance improving as the

Zip-Adapter-F Config		Few Shot Setup		
Trainable support feature	Trainable ZIP	1	3	5
✓	✗	72.67	85.55	89.39
✗	✓	73.45	85.89	89.32
✓	✓	73.74	86.12	89.41

Table 5: Classification accuracy (%) on MVTec-FS of different training settings of ZIFA-Adapter-F.

Dataset	Defect Type (Instance Number)	Annotations
NEU_DET	Crazing(689), Pitted_surface(432), Rolled_in_scale(628), Patches(881), Scratches(548), Inclusion(1011)	Bbox
PCB	Spurious_copper(503), Short(491), Spur(488), Mouse_bite(492), Missing_hole(497), Open_circuit(482)	Bbox
MTD	Break(108), Crack(69), Fray(37), Uneven(103), Blowhole(115)	Mask
AITEX	Broken_end(11), Broken_yarn(16), Cuts_elvage(12), Weftcrack(15), Fuzzyball(42), Nep(19), Broken_pick(65)	Mask

Table 6: Defect types and their counts in Other four datasets.

number of shots increases.

Discussion and Conclusion

This paper introduces MVREC, an instance-level few-shot classification approach that can be applied to various labeled formats, such as bounding boxes and masks. Extensive experiments on five datasets demonstrate that it is a versatile and effective approach for FSDMC.

Limitations. First, we have not yet explored the unified model that can be developed to handle different defect datasets after a single training session. Second, our study mainly focuses on using image features extracted by CLIP, without exploring the potential of CLIP’s text encoder for multi-model research. We hope this work inspires future research and the development of more advanced methods.

Acknowledgements

This research is supported by Laboratory for Artificial Intelligence in Design (Project Code: RP3-3) under InnoHK Research Clusters, Hong Kong SAR Government.

References

- Bergmann, P.; Batzner, K.; Fauser, M.; Sattlegger, D.; and Steger, C. 2022. Beyond Dents and Scratches: Logical Constraints in Unsupervised Anomaly Detection and Localization. *Int. J. Comput. Vis.*, 130(4): 947–969.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD - A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, 9592–9600. Computer Vision Foundation / IEEE.
- Cao, Y.; Zhu, W.; Yang, J.; Fu, G.; Lin, D.; and Cao, Y. 2023. An effective industrial defect classification method under the few-shot setting via two-stream training. *Optics and Lasers in Engineering*, 161: 107294.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Gao, P.; Geng, S.; Zhang, R.; Ma, T.; Fang, R.; Zhang, Y.; Li, H.; and Qiao, Y. 2024. CLIP-Adapter: Better Vision-Language Models with Feature Adapters. *Int. J. Comput. Vis.*, 132(2): 581–595.
- Gu, Z.; Zhu, B.; Zhu, G.; Chen, Y.; Tang, M.; and Wang, J. 2024. AnomalyGPT: Detecting Industrial Anomalies Using Large Vision-Language Models. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada, 1932–1940*. AAAI Press.
- He, Y.; Song, K.; Meng, Q.; and Yan, Y. 2020. An End-to-End Steel Surface Defect Detection Approach via Fusing Multiple Hierarchical Features. *IEEE Transactions on Instrumentation and Measurement*, 69(4): 1493–1504.
- Huang, W.; and Wei, P. 2019. A PCB dataset for defects detection and classification. *arXiv preprint arXiv:1901.08204*.
- Huang, Y.; Qiu, C.; and Yuan, K. 2020. Surface defect saliency of magnetic tile. *The Visual Computer*, 36(1): 85–96.
- Jeong, J.; Zou, Y.; Kim, T.; Zhang, D.; Ravichandran, A.; and Dabeer, O. 2023. WinCLIP: Zero-/Few-Shot Anomaly Classification and Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 19606–19616. IEEE.
- Jha, S. B.; and Babiceanu, R. F. 2023. Deep CNN-based visual defect detection: Survey of current literature. *Computers in Industry*, 148: 103911.
- Kim, J.; Jeong, K.; Choi, H.; and Seo, K. 2020. GAN-Based Anomaly Detection In Imbalance Problems. In Bartoli, A.; and Fusiello, A., eds., *Computer Vision - ECCV 2020 Workshops - Glasgow, UK, August 23-28, 2020, Proceedings, Part VI*, volume 12540 of *Lecture Notes in Computer Science*, 128–145. Springer.
- Liu, Z.; Song, Y.; Tang, R.; Duan, G.; and Tan, J. 2023. Few-shot defect recognition of metal surfaces via attention-embedding and self-supervised learning. *Journal of Intelligent Manufacturing*, 34(8): 3507–3521.
- Lyu, S.; Mo, D.; and Wong, W. 2024. REB: Reducing biases in representation for industrial anomaly detection. *Knowledge-Based Systems*, 290: 111563.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Pang, G.; Shen, C.; Cao, L.; and van den Hengel, A. 2022. Deep Learning for Anomaly Detection: A Review. *ACM Comput. Surv.*, 54(2): 38:1–38:38.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Shtedritski, A.; Rupprecht, C.; and Vedaldi, A. 2023. What does CLIP know about a red circle? Visual prompt engineering for VLMs. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 11953–11963. IEEE.
- Silvestre-Blanes, J.; Albero-Albero, T.; Miralles, I.; Pérez-Llorens, R.; and Moreno, J. 2019. A public fabric database for defect detection methods and results. *Autex Research Journal*, 19(4): 363–374.
- Snell, J.; Swersky, K.; and Zemel, R. S. 2017. Prototypical Networks for Few-shot Learning. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 4077–4087.
- Sun, Z.; Fang, Y.; Wu, T.; Zhang, P.; Zang, Y.; Kong, S.; Xiong, Y.; Lin, D.; and Wang, J. 2024. Alpha-clip: A clip model focusing on wherever you want. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13019–13029.
- Udandarao, V.; Gupta, A.; and Albanie, S. 2023. SuS-X: Training-Free Name-Only Transfer of Vision-Language Models. In *IEEE/CVF International Conference on Com-*

puter Vision, ICCV 2023, Paris, France, October 1-6, 2023, 2725–2736. IEEE.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).

Wang, Y.; Yao, Q.; Kwok, J. T.; and Ni, L. M. 2020. Generalizing from a Few Examples: A Survey on Few-shot Learning. *ACM Comput. Surv.*, 53(3).

Xiao, W.; Song, K.; Liu, J.; and Yan, Y. 2022. Graph embedding and optimal transport for few-shot classification of metal surface defect. *IEEE Transactions on Instrumentation and Measurement*, 71: 1–10.

Zhan, Z.; Zhou, J.; and Xu, B. 2022. Fabric defect classification using prototypical network of few-shot learning algorithm. *Computers in Industry*, 138: 103628.

Zhang, R.; Zhang, W.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-Adapter: Training-Free Adaption of CLIP for Few-Shot Classification. In Avidan, S.; Brostow, G. J.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXV*, volume 13695 of *Lecture Notes in Computer Science*, 493–510. Springer.

Zhao, W.; Song, K.; Wang, Y.; Liang, S.; and Yan, Y. 2023. FaNet: Feature-aware network for few shot classification of strip steel surface defects. *Measurement*, 208: 112446.

Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; and Gao, J. 2022. RegionCLIP: Region-based Language-Image Pre-training. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 16772–16782. IEEE.

Zhou, C.; Liu, M.; Zhang, S.; Wei, P.; and Chen, B. 2023. Few-shot classification of screen defects with class-agnostic mask and context-based classifier. *IEEE Transactions on Instrumentation and Measurement*, 72: 1–16.

Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract Free Dense Labels from CLIP. In Avidan, S.; Brostow, G. J.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXVIII*, volume 13688 of *Lecture Notes in Computer Science*, 696–712. Springer.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to Prompt for Vision-Language Models. *Int. J. Comput. Vis.*, 130(9): 2337–2348.

Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. *CoRR*, abs/2304.10592.