

MediCLIP: Adapting CLIP for Few-shot Medical Image Anomaly Detection

Ximiao Zhang¹, Min Xu^{1*}, Dehui Qiu¹, Ruixin Yan², Ning Lang², and
Xiuzhuang Zhou^{3*}

¹ College of Information and Engineering, Capital Normal University, Beijing, China

² Department of Radiology, Peking University Third Hospital, Beijing, China

³ School of Artificial Intelligence, Beijing University of Posts and
Telecommunications, Beijing, China

Abstract. In the field of medical decision-making, precise anomaly detection in medical imaging plays a pivotal role in aiding clinicians. However, previous work is reliant on large-scale datasets for training anomaly detection models, which increases the development cost. This paper first focuses on the task of medical image anomaly detection in the few-shot setting, which is critically significant for the medical field where data collection and annotation are both very expensive. We propose an innovative approach, MediCLIP, which adapts the CLIP model to few-shot medical image anomaly detection through self-supervised fine-tuning. Although CLIP, as a vision-language model, demonstrates outstanding zero-/few-shot performance on various downstream tasks, it still falls short in the anomaly detection of medical images. To address this, we design a series of medical image anomaly synthesis tasks to simulate common disease patterns in medical imaging, transferring the powerful generalization capabilities of CLIP to the task of medical image anomaly detection. When only few-shot normal medical images are provided, MediCLIP achieves state-of-the-art performance in anomaly detection and location compared to other methods. Extensive experiments on three distinct medical anomaly detection tasks have demonstrated the superiority of our approach. The code is available at <https://github.com/cnulab/MediCLIP>.

Keywords: Anomaly detection · Medical image · Few-shot learning.

1 Introduction

The medical image anomaly detection task involves distinguishing between normal and anomalous images with various lesion areas, and identifying the locations of these lesion areas. This offers significant insights for medical professionals in making clinical decisions, substantially reducing the risk associated with decision-making and enhancing the efficiency of medical professionals. Due to the scarcity of lesion images, research in this field is focused on unsupervised anomaly

* Corresponding author. Email: xumin@cnu.edu.cn, xiuzhuang.zhou@bupt.edu.cn.

detection settings, where models are trained solely on normal images, and any subtle deviation from the normal pattern is identified as an anomaly. Recently, significant progress has been made in medical image anomaly detection [23, 24, 4, 25], achieving comparable or even better anomaly detection performance than professional clinicians in certain detection tasks [5]. However, these methods utilize thousands or more medical images to train their anomaly detection models, increasing the developmental cost and hindering rapid deployment. In this paper, we focus on a more challenging task, namely, the few-shot medical image anomaly detection task. For each medical anomaly detection task, we provide only a few-shot normal images for model training, without any available anomaly images and pixel-level labels. This setup significantly reduces the development cost of medical anomaly detection models and avoids the expensive process of medical data collection and annotation.

Parallel to this, CLIP [19] demonstrates powerful generalization capabilities by mapping natural images and raw texts into a unified representational space, which can be easily applied to a variety of downstream tasks. MedCLIP [22] extends the capabilities of CLIP [19] for medical image classification by training on unpaired medical images and text. Despite this advancement, the method falls short in identifying and localizing lesion areas. Recent studies [28, 8, 7, 10] have applied CLIP [19] for zero-/few-shot anomaly detection with impressive performance. However, these methods require an auxiliary dataset with real anomaly images and corresponding pixel-level anomaly labels for model training, which is often challenging to obtain in the medical field. To address this limitation, we propose MediCLIP, which employs synthetic anomaly images for model training instead of relying on an auxiliary dataset. These synthetic images, crafted through a variety of carefully designed anomaly synthesis tasks, effectively emulate diverse abnormal patterns observed in medical imaging. In our approach, we leverage learnable prompts [27] to avoid the complex design of artificial text prompts, ensuring that the text embeddings can effectively generalize within medical images. Additionally, we use adapters to align the intermediate layer features from the CLIP [19] vision encoder with learnable text features, enabling multi-scale lesion localization capabilities.

We validate the effectiveness of MediCLIP on three distinct medical imaging datasets: the chest X-ray image dataset CheXpert [11], the brain MRI image dataset BrainMRI [6], and the breast ultrasound image dataset BUSI [1]. Under the few-shot learning setting, MediCLIP demonstrates approximately a 10% improvement in performance compared to existing state-of-the-art methods. Notably, on the most challenging CheXpert [11] dataset, MediCLIP achieved 94% of the performance level of the advanced full-shot medical image anomaly detection method, SQUID [25], while using less than 1% of the training images. Additionally, we have also discovered that MediCLIP has an impressive zero-shot generalization capability. The anomaly detection ability of MediCLIP acquired by training in a medical anomaly detection task can be transferred to other medical anomaly detection tasks. This indicates that MediCLIP has the potential to serve as a unified model for medical anomaly detection. Through extensive




Fig. 1. The overall pipeline of our proposed MediCLIP framework.

experimentation, we have demonstrated that MediCLIP has both low cost and high performance, showcasing its capabilities in assisting professional doctors with real medical decision-making.

2 Method

In this section, we first present the problem definition of few-shot anomaly detection. Then we introduce the learnable prompts and adapters in MediCLIP. Finally, we describe our medical image anomaly synthesis strategy. The overall pipeline of MediCLIP is illustrated in Fig. 1.

2.1 Problem Definition

We provide a formal definition for the few-shot anomaly detection task. According to the few-shot learning setting, for an n -way and k -shot episode, a support set $D = \bigcup_{i=1}^n T_i$ contains normal images from n tasks, where T_i includes k normal images from the i th task. During the test phase, query images from n tasks are provided, and the model is trained to identify whether these query images are anomalous and to indicate the location of the anomalous regions. Due to the significant differences of various medical anomaly detection tasks, and in order to meet the actual needs of the medical system, in this paper, we set $n = 1$, which means one model for one task.

2.2 Prompt Learning

In CLIP [19], commonly used prompts like ‘A photo of a [CLS]’ mainly describe the overall semantic information of the nature image, and they struggle to capture the subtle details in medical imaging. Therefore, we use learnable prompts [27] instead of manually designed prompts for anomaly detection, ensuring that the text embeddings can generalize effectively in medical images while avoiding the complexity of prompt engineering. Specifically, we adopt the

following prompt format:

$$p = [V_1][V_2]...[V_M][CLS] \quad (1)$$

where $[V]$ and $[CLS]$ represent the learnable word embeddings and the non-learnable class embeddings in the prompt template, and M denotes the count of the learnable tokens. For class tokens, we use medically relevant terms such as `[healthy]` and `[normal]` for normal cases, and `[disease]` for anomalies. The complete list of class tokens is provided in the Appendix. Therefore, for normal and anomalous cases, we have $\mathcal{P}_n = \{p_{n_1}, p_{n_2}, \dots, p_{n_I}\}$ and $\mathcal{P}_a = \{p_{a_1}, p_{a_2}, \dots, p_{a_E}\}$, where $\mathcal{P}_n$ and $\mathcal{P}_a$ denote the prompt sets for normal and anomalous cases, and I and E represent the count of class tokens contained within $\mathcal{P}_n$ and $\mathcal{P}_a$. We define the text encoder of CLIP as $F(\cdot)$, where $F(\cdot)$ maps the prompt p to its feature representation $F(p) \in R^C$, where C signifies the dimension of prompt feature. The mean feature representations f_n and f_a for the prompt sets $\mathcal{P}_n$ and $\mathcal{P}_a$ are calculated respectively by $f_n = \frac{1}{I} \sum_{i=1}^I F(p_{n_i})$ and $f_a = \frac{1}{E} \sum_{i=1}^E F(p_{a_i})$.

2.3 Adapting CLIP for Anomaly Detection

The vanilla CLIP model [19] is designed for zero-/few-shot image-text classification and is not directly applicable to anomaly detection and localization. In this section, we introduce adapters to the vanilla CLIP model [19], adapting it for few-shot anomaly detection through a small number of learnable parameters. Specifically, for a medical image $X \in R^{H \times W \times 3}$ from the support set D , we randomly sample an anomaly region mask $Y \in R^{H \times W}$ and generate the corresponding synthetic anomaly image $\hat{X} = \Psi(X, Y)$ through a series of anomaly synthesis tasks, where $\Psi(\cdot)$ represents the anomaly synthesis function. The vision encoder of CLIP [19], denoted as, $G_j(\cdot)$, which extracts the j th layer feature of the image, represented as $G_j(\hat{X}) \in R^{H_j \times W_j \times C_j}$, where H_j , W_j , and C_j respectively represent the height, width, and number of channels of the j th layer's features. Consequently, for image $\hat{X}$, we obtain a multi-scale visual feature set $\{G_1(\hat{X}), G_2(\hat{X}), \dots, G_J(\hat{X})\}$. For each intermediate layer feature $G_j(\hat{X})$, we use an adapter $\phi_j(\cdot)$ to map it to the same number of channels as the prompt features f_n and f_a , resulting in $g_j = \phi_j(G_j(\hat{X}))$, where $g_j \in R^{H_j \times W_j \times C}$. We then calculate the similarity of g_j with f_n and f_a at the spatial location (h, w) :

$$S_{n_j}^{(h,w)} = \frac{\exp(\langle g_j^{(h,w)}, f_n \rangle / \tau)}{\sum_{f \in \{f_n, f_a\}} \exp(\langle g_j^{(h,w)}, f \rangle / \tau)}, S_{a_j}^{(h,w)} = \frac{\exp(\langle g_j^{(h,w)}, f_a \rangle / \tau)}{\sum_{f \in \{f_n, f_a\}} \exp(\langle g_j^{(h,w)}, f \rangle / \tau)} \quad (2)$$

where $\langle \cdot, \cdot \rangle$ represents the cosine similarity, τ is the temperature parameter, and $h < H_j$, $w < W_j$, $h, w \in \mathbb{N}$. For each visual feature g_j projected by adapter, we can obtain similarity matrices $S_{n_j}, S_{a_j} \in R^{H_j \times W_j}$. Then, we perform an aggregation operation on the multi-scale similarity matrix sets $\{S_{n_1}, S_{n_2}, \dots, S_{n_J}\}$ and $\{S_{a_1}, S_{a_2}, \dots, S_{a_J}\}$. We upsample them to the same spatial resolution of




Fig. 2. Synthetic anomaly image examples from diverse tasks.

$H \times W$, and compute their average to obtain $S_n, S_a \in R^{H \times W}$. To optimize the parameters in learnable prompts and adapters, we define the loss function as:

$$\mathcal{L} = \mathcal{L}_{focal}([S_n, S_a], Y) + \mathcal{L}_{dice}(S_a, Y) \quad (3)$$

where $\mathcal{L}_{focal}(\cdot, \cdot)$ and $\mathcal{L}_{dice}(\cdot, \cdot)$ respectively represent Focal loss [15] and Dice loss [14], and $[\cdot, \cdot]$ denotes concatenation along the channel. During the inference stage, we remove anomaly synthesis tasks, using the query image as $\hat{X}$. We employ S_a as the pixel-level anomaly map, and the maximum value in S_a is used as the image-level anomaly score.

2.4 Multi-task Anomaly Synthesis

Multi-task anomaly synthesis generates the anomaly image $\hat{X}$ based on a given source image X and a target mask Y . This paper introduces three types of anomaly synthesis tasks: CutPaste, GaussIntensityChange, and Source, corresponding to image blending, intensity variation, and deformation, respectively. Some examples are presented in Fig. 2.

CutPaste [13, 21, 2] randomly selects an image patch and pastes it to target location. We employ the enhanced version from [21, 2] that utilizes Poisson image editing [17] for seamless image blending. The CutPaste task can simulate misplacement-like anomalies in medical imaging, such as fractures.

GaussIntensityChange simulates density variations in medical imaging by altering pixel intensities. Specifically, we sample noise σ from a standard Gaussian distribution $\sigma \sim \mathcal{N}(0, \mathbf{I})$, and then binarize σ based on a threshold value of 0. Subsequently, the binarized noise is processed through a Gaussian filter to obtain $\hat{\sigma}$. The final anomaly image is defined as:

$$\hat{X} = X \odot (1 - Y) + (X + \gamma \hat{\sigma}) \odot Y \quad (4)$$

where $\odot$ denotes the element-wise multiplication operation, and γ is the intensity variation factor, when $\gamma > 0$, intensity increases, and when $\gamma < 0$, intensity decreases. The task is particularly effective in mimicking the density variations seen in tumors or cysts within radiographic images.

Source [2] pushes all points in the mask Y away from a center c . We define the radius of Y in a certain direction as r . For the source position l in this direction, we compute the corresponding target position $\hat{l}$:

$$\hat{l} = c + r \frac{l - c}{\|l - c\|_2} \left(\frac{\|l - c\|_2}{r} \right)^\alpha, \quad c \in \mathbb{L}_Y, \forall l \in \mathbb{L}_Y \quad (5)$$

Table 1. Performance comparison in anomaly detection: MediCLIP versus other methods under few-shot settings, using Image-AUROC (%) as the evaluation metric.

Datasets	k	PatchCore	SimpleNet	SQUID	AnoDDPM	CutPaste	PatchSVDD	RegAD	MediCLIP
CheXpert	4	60.0 $\pm$ 1.8	63.5 $\pm$ 2.4	66.0 $\pm$ 1.9	55.8 $\pm$ 1.3	63.6 $\pm$ 2.8	57.5 $\pm$ 1.1	59.8 $\pm$ 1.4	70.1$\pm$1.6
	8	59.9 $\pm$ 2.2	63.8 $\pm$ 0.7	66.6 $\pm$ 3.9	57.3 $\pm$ 3.5	65.0 $\pm$ 2.9	61.3 $\pm$ 1.5	62.0 $\pm$ 3.0	70.8$\pm$1.5
	16	64.6 $\pm$ 1.6	65.6 $\pm$ 3.7	67.6 $\pm$ 4.7	59.7 $\pm$ 2.8	65.2 $\pm$ 3.2	62.7 $\pm$ 2.3	60.2 $\pm$ 1.7	72.5$\pm$1.7
	32	64.3 $\pm$ 1.2	67.6 $\pm$ 1.1	69.4 $\pm$ 2.2	66.9 $\pm$ 3.2	71.4 $\pm$ 2.0	61.7 $\pm$ 1.4	65.6 $\pm$ 1.4	73.3$\pm$1.3
BrainMRI	4	70.8 $\pm$ 3.5	77.0 $\pm$ 3.0	69.5 $\pm$ 4.4	67.2 $\pm$ 1.8	72.3 $\pm$ 1.8	78.4 $\pm$ 1.4	67.5 $\pm$ 1.0	94.1$\pm$0.7
	8	75.9 $\pm$ 2.9	79.3 $\pm$ 2.9	75.8 $\pm$ 1.9	73.0 $\pm$ 3.6	77.0 $\pm$ 4.9	79.5 $\pm$ 0.8	75.1 $\pm$ 2.8	94.9$\pm$0.4
	16	79.5 $\pm$ 1.9	81.7 $\pm$ 4.4	77.2 $\pm$ 0.6	77.2 $\pm$ 2.5	80.2 $\pm$ 2.7	78.3 $\pm$ 3.0	82.4 $\pm$ 1.6	94.8$\pm$0.5
	32	81.7 $\pm$ 0.2	82.6 $\pm$ 0.9	79.3 $\pm$ 1.7	80.0 $\pm$ 2.1	79.0 $\pm$ 4.3	78.8 $\pm$ 1.4	83.0 $\pm$ 2.4	95.4$\pm$0.4
BUSI	4	81.7 $\pm$ 1.3	76.1 $\pm$ 2.5	55.4 $\pm$ 2.2	69.3 $\pm$ 1.3	75.7 $\pm$ 1.4	69.2 $\pm$ 1.1	74.7 $\pm$ 5.0	88.5$\pm$0.6
	8	82.6 $\pm$ 2.1	80.4 $\pm$ 4.8	58.7 $\pm$ 1.6	72.6 $\pm$ 4.6	78.2 $\pm$ 2.1	70.6 $\pm$ 2.4	75.8 $\pm$ 2.7	89.0$\pm$0.3
	16	87.6 $\pm$ 1.0	84.1 $\pm$ 2.7	64.8 $\pm$ 4.5	74.4 $\pm$ 2.9	80.0 $\pm$ 0.9	73.7 $\pm$ 6.0	76.5 $\pm$ 1.2	90.8$\pm$0.4
	32	88.9 $\pm$ 0.1	88.0 $\pm$ 0.3	67.8 $\pm$ 0.4	76.2 $\pm$ 4.0	80.2 $\pm$ 1.7	77.9 $\pm$ 1.6	78.6 $\pm$ 1.2	91.0$\pm$0.3

Table 2. Performance comparison in anomaly localization: MediCLIP versus other methods on the BUSI dataset [1], using Pixel-AUROC (%) as the evaluation metric.

k	PatchCore	SimpleNet	SQUID	AnoDDPM	PatchSVDD	RegAD	MediCLIP
4	78.5 $\pm$ 2.6	73.5 $\pm$ 4.9	56.8 $\pm$ 5.1	68.3 $\pm$ 1.9	74.4 $\pm$ 0.9	71.9 $\pm$ 3.1	88.0$\pm$0.3
8	78.1 $\pm$ 1.3	77.6 $\pm$ 3.4	56.8 $\pm$ 1.2	74.7 $\pm$ 4.3	73.6 $\pm$ 3.4	76.1 $\pm$ 2.6	88.6$\pm$1.0
16	80.2 $\pm$ 0.8	78.2 $\pm$ 1.5	65.5 $\pm$ 2.8	79.5 $\pm$ 1.4	78.3 $\pm$ 2.2	74.6 $\pm$ 1.0	89.9$\pm$0.9
32	80.3 $\pm$ 0.1	79.8 $\pm$ 0.4	68.4 $\pm$ 2.2	79.1 $\pm$ 1.3	77.6 $\pm$ 3.4	76.3 $\pm$ 0.4	90.1$\pm$1.0

where α controls the intensity of repulsion from the center, with a larger α resulting in stronger repulsion. Finally, the pixel at the source location is replaced with the pixel at the target location, resulting in $\hat{x}_l = x_i$. The Source task simulates proliferative anomalies found in medical imaging, such as cardiomegaly.


3 Experiments

3.1 Experimental Setup

Datasets. We conducted extensive experiments on three datasets, including the Stanford CheXpert dataset [11], the BrainMRI dataset [6], and the BUSI dataset [1]. The CheXpert dataset [11] comprises chest X-ray images from clinical patients at Stanford Hospital, encompassing 12 different diseases (such as Cardiomegaly, Edema, Pleural Effusion, etc.). The BrainMRI dataset [1] contains 2D human brain MRI images, including both normal and tumor-affected cases. The BUSI dataset [1] consists of breast ultrasound images from female patients aged 25 to 75. These images are classified into normal, benign, and malignant categories, with pixel-level annotations marking disease locations. We treated all disease images as anomalies. For each dataset, we provided $k = \{4, 8, 16, 32\}$ normal images for model training. During the testing phase, the CheXpert [11], BrainMRI [6], and BUSI [1] datasets contain 250, 65, and 101 normal images, respectively, along with 250, 155, and 647 anomaly images.

Table 3. Zero-shot anomaly detection performance of MediCLIP with support set size $k = 16$.

Train Task	Test Task	Image-AUROC (%)
CheXpert	CheXpert	72.5 $\pm$ 1.7
	BrainMRI	92.1 $\pm$ 1.2
	BUSI	88.8 $\pm$ 1.0
BrainMRI	CheXpert	67.7 $\pm$ 0.8
	BrainMRI	94.8 $\pm$ 0.5
	BUSI	84.8 $\pm$ 1.1
BUSI	CheXpert	70.7 $\pm$ 0.8
	BrainMRI	91.3 $\pm$ 1.0
	BUSI	90.8 $\pm$ 0.4

**Fig. 3.** Qualitative results of MediCLIP. Each group contains the anomaly image and the predicted anomaly map.

Implementation Details. We used the ViT-L/14 [19] as the pre-trained CLIP model, retaining the features of the 12th, 18th, and 24th layers of the visual encoder for anomaly detection. We set the number of learnable tokens M to 8 and used the linear layer as our adapter, setting the temperature parameter τ in Eq. (2) to 0.07. For the anomaly synthesis tasks, we uniformly sampled the intensity variation factor γ in Eq. (4) within the range of $[-0.6, -0.4) \cup [0.4, 0.6)$, and α in Eq. (5) within $[\sqrt{2}, 4)$. For the CutPaste and GaussIntensityChange tasks, we utilized a Perlin noise generator [18] to capture various anomalous shapes and binarize them into the anomaly mask Y . For the Source task, since it needs to compute the radius of the anomalous region in any direction, we used ellipses or rectangles with random sizes and rotation angles as the anomaly mask. We employed three types of anomaly synthesis tasks with equal probability for anomaly synthesis. The input image size was 224×224 pixels. For the three datasets, we used a consistent hyperparameter setting. Each experiment was independently conducted three times, sampling different support sets for MediCLIP training.

Baselines and Metrics. To comprehensively evaluate the few-shot anomaly detection performance of MediCLIP, we established the following baselines for comparison: deep feature embedding-based methods PatchCore [20] and SimpleNet [16], reconstruction-based methods SQUID [25] and AnoDDPM [24], self-supervised learning-based method CutPaste [13], one-class classification method PatchSVDD [26], and few-shot anomaly detection method RegAD [9]. For the evaluation metrics, we used Image-AUROC to evaluate anomaly detection performance and Pixel-AUROC to evaluate anomaly localization performance.

3.2 Experimental Results

We compared the anomaly detection performance of MediCLIP with other methods on three datasets, as shown in Table 1. Under the few-shot setting, other methods struggle to achieve good generalization due to the scarcity of normal images. In contrast, through a series of carefully crafted anomaly synthesis tasks, MediCLIP effectively transfers the powerful generalization capability of the CLIP model to the task of medical anomaly detection. MediCLIP achieved a performance improvement of approximately 10% over other methods. Table 2 displays

Table 4. Ablation study results for learnable token and adapter in MediCLIP with support set size $k = 16$, using Image-AUROC (%) as the evaluation metric.

	Learnable Token	Adapter	CheXpert	BrainMRI	BUSI
1	-	-	42.5 $\pm$ 0.0	41.6 $\pm$ 0.0	52.6 $\pm$ 0.0
2	✓	-	57.7 $\pm$ 1.3	71.8 $\pm$ 1.3	80.3 $\pm$ 1.4
3	-	✓	71.4 $\pm$ 2.1	94.2 $\pm$ 0.9	90.3 $\pm$ 0.7
4	✓	✓	72.5$\pm$1.7	94.8$\pm$0.5	90.8$\pm$0.4

Table 5. Anomaly detection performance of MediCLIP under different training settings, using Image-AUROC (%) as the evaluation metric.

		CheXpert	BrainMRI	BUSI
Auxiliary Datasets	MVTec-AD	64.4 $\pm$ 0.7	88.2 $\pm$ 1.1	77.0 $\pm$ 0.9
	VisA	55.1 $\pm$ 1.0	85.2 $\pm$ 1.1	74.7 $\pm$ 2.1
Anomaly Synthesis ($k = 16$)	CutPaste	67.6 $\pm$ 1.3	94.8 $\pm$ 1.2	89.3 $\pm$ 1.8
	GaussIntensityChange	70.0 $\pm$ 1.8	94.0 $\pm$ 0.2	90.0 $\pm$ 0.6
	Source	67.3 $\pm$ 1.1	92.2 $\pm$ 1.0	87.9 $\pm$ 1.0
	GaussIntensityChange & CutPaste	71.8 $\pm$ 0.4	95.2$\pm$0.4	90.1 $\pm$ 0.8
	GaussIntensityChange & Source	71.1 $\pm$ 1.1	94.1 $\pm$ 0.5	90.4 $\pm$ 1.1
	All Synthesis Tasks	72.5$\pm$1.7	94.8 $\pm$ 0.5	90.8$\pm$0.4

the anomaly localization performance of MediCLIP and other methods on the BUSI dataset [1], indicating that MediCLIP is more effective in locating lesion areas. Fig. 3 provides a visualization of the anomaly localization results by MediCLIP, demonstrating its accurate identification of anomaly regions, thereby offering valuable references for real medical decision-making.

Table 3 demonstrates MediCLIP’s zero-shot anomaly detection performance on various datasets. The experimental results indicate that MediCLIP’s anomaly detection capability can be generalized across different tasks, demonstrating that the multi-task anomaly synthesis strategy can effectively simulate common anomaly patterns contained in various medical images, further evidencing its potential as a unified solution in medical anomaly detection.

3.3 Ablation Studies

Table 4 validates the effectiveness of the learnable token and adapter in MediCLIP. We replaced the learnable token with hard prompt templates provided in [12], such as ‘A photo of the [CLS]’. Additionally, we used average pooling along the channel instead of the adapter. The experimental results in Table 4 demonstrate the irreplaceability of the learnable token and adapter. Table 5 compares the anomaly detection performance of MediCLIP under different training settings. For the auxiliary datasets, we follow the settings in [28, 8, 7], using the industrial anomaly detection datasets MVTec-AD [3] and VisA [29] for model training, and assess the anomaly detection performance in medical images. Compared to auxiliary datasets, the medical anomaly synthesis task achieved better generalization performance. Furthermore, the combination of multiple tasks can further enhance the anomaly detection performance of MediCLIP.

4 Conclusion

This paper first comprehensively explores the task of few-shot medical image anomaly detection, a field with broad applications yet insufficiently studied. We introduce MediCLIP, an adaptation of the CLIP model tailored for medical anomaly detection, integrating multi-task anomaly synthesis. Characterized by its efficiency and superior performance, MediCLIP demonstrates effective detection of various diseases across diverse medical imaging types. Our comprehensive experiments on three medical datasets underscore MediCLIP’s capability in accurate anomaly detection and localization, establishing its promise as a unified and vital tool in medical diagnostics.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China under Grant 62177034 and Grant 61972046, the Beijing Natural Science Foundation under grant Z190020, and the Proof of Concept Program of Zhongguancun Science City and Peking University Third Hospital under Grant HDCXZHKC2022202.

Disclosure of Interests. The authors have no competing interests.

References

1. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. *Data in brief* **28**, 104863 (2020)
2. Baugh, M., Tan, J., Müller, J.P., Dombrowski, M., Batten, J., Kainz, B.: Many tasks make light work: Learning to localise medical anomalies from multiple synthetic tasks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 162–172. Springer (2023)
3. Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9592–9600 (2019)
4. Cai, Y., Chen, H., Yang, X., Zhou, Y., Cheng, K.T.: Dual-distribution discrepancy for anomaly detection in chest x-rays. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 584–593. Springer (2022)
5. Cao, K., Xia, Y., Yao, J., Han, X., Lambert, L., Zhang, T., Tang, W., Jin, G., Jiang, H., Fang, X., et al.: Large-scale pancreatic cancer detection via non-contrast ct and deep learning. *Nature Medicine* pp. 1–11 (2023)
6. Chakrabarty, N.: Brain mri images for brain tumor detection (2019), <https://www.kaggle.com/datasets/navoneel/brain-mri-images-for-brain-tumor-detection>
7. Chen, X., Han, Y., Zhang, J.: A zero-/few-shot anomaly classification and segmentation method for cvpr 2023 vand workshop challenge tracks 1&2: 1st place on zero-shot ad and 4th place on few-shot ad. *arXiv preprint arXiv:2305.17382* (2023)
8. Chen, X., Zhang, J., Tian, G., He, H., Zhang, W., Wang, Y., Wang, C., Wu, Y., Liu, Y.: Clip-ad: A language-guided staged dual-path model for zero-shot anomaly detection. *arXiv preprint arXiv:2311.00453* (2023)

9. Huang, C., Guan, H., Jiang, A., Zhang, Y., Spratling, M., Wang, Y.F.: Registration based few-shot anomaly detection. In: European Conference on Computer Vision. pp. 303–319. Springer (2022)
10. Huang, C., Jiang, A., Feng, J., Zhang, Y., Wang, X., Wang, Y.: Adapting visual-language models for generalizable anomaly detection in medical images. arXiv preprint arXiv:2403.12570 (2024)
11. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghighi, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 590–597 (2019)
12. Jeong, J., Zou, Y., Kim, T., Zhang, D., Ravichandran, A., Dabeer, O.: Winclip: Zero-/few-shot anomaly classification and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19606–19616 (2023)
13. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9664–9674 (2021)
14. Li, X., Sun, X., Meng, Y., Liang, J., Wu, F., Li, J.: Dice loss for data-imbalanced nlp tasks. arXiv preprint arXiv:1911.02855 (2019)
15. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
16. Liu, Z., Zhou, Y., Xu, Y., Wang, Z.: Simplesnet: A simple network for image anomaly detection and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20402–20411 (2023)
17. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. In: Seminal Graphics Papers: Pushing the Boundaries, Volume 2, pp. 577–582 (2023)
18. Perlin, K.: An image synthesizer. ACM Siggraph Computer Graphics **19**(3), 287–296 (1985)
19. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
20. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14318–14328 (2022)
21. Schlüter, H.M., Tan, J., Hou, B., Kainz, B.: Natural synthetic anomalies for self-supervised anomaly detection and localization. In: European Conference on Computer Vision. pp. 474–489. Springer (2022)
22. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 3876–3887 (2022)
23. Wolleb, J., Bieder, F., Sandkühler, R., Cattin, P.C.: Diffusion models for medical anomaly detection. In: International Conference on Medical image computing and computer-assisted intervention. pp. 35–45. Springer (2022)
24. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 650–656 (2022)

25. Xiang, T., Zhang, Y., Lu, Y., Yuille, A.L., Zhang, C., Cai, W., Zhou, Z.: Squid: Deep feature in-painting for unsupervised anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23890–23901 (2023)
26. Yi, J., Yoon, S.: Patch svdd: Patch-level svdd for anomaly detection and segmentation. In: Proceedings of the Asian conference on computer vision (2020)
27. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *International Journal of Computer Vision* **130**(9), 2337–2348 (2022)
28. Zhou, Q., Pang, G., Tian, Y., He, S., Chen, J.: Anomalyclip: Object-agnostic prompt learning for zero-shot anomaly detection. In: The Twelfth International Conference on Learning Representations (2024)
29. Zou, Y., Jeong, J., Pemula, L., Zhang, D., Dabeer, O.: Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In: European Conference on Computer Vision. pp. 392–408. Springer (2022)