

# Style-Aligned Image Composition for Robust Detection of Abnormal Cells in Cytopathology

Qiuyi Qi<sup>\*1</sup>

Xin Li<sup>\*2,3,5</sup>

Ming Kong<sup>1</sup>

Zikang Xu<sup>2,3</sup>

Bingdi Chen<sup>4,5</sup>

Qiang Zhu<sup>†1</sup>

S Kevin Zhou<sup>†2,3</sup>

QIQUYI@ZJU.EDU.CN

MERLEY@MAIL.USTC.EDU.CN

ZJUKONGMING@ZJU.EDU.CN

ZIKANGXU@MAIL.USTC.EDU.CN

INANOCHEN@TONGJI.EDU.CN

ZHUQ@ZJU.EDU.CN

SKEVINZHOU@USTC.EDU.CN

<sup>1</sup> Zhejiang University, Hangzhou 310058, China

<sup>2</sup> School of Biomedical Engineering, Division of Life Sciences and Medicine, University of Science and Technology of China (USTC), Hefei Anhui, 230026, China

<sup>3</sup> Center for Medical Imaging, Robotics, Analytic Computing & Learning (MIRACLE), Suzhou Institute for Advance Research, USTC, Suzhou Jiangsu, 215123, China

<sup>4</sup> The Institute for Biomedical Engineering & Nano Science, Tongji University School of Medicine, Shanghai, 200331, China

<sup>5</sup> Zhihui Medical Technology (Shanghai) Co., Ltd., Shanghai, 200333, China

**Editors:** Under Review for MIDL 2025

## Abstract

Challenges such as the lack of high-quality annotations, long-tailed data distributions, and inconsistent staining styles pose significant obstacles to training neural networks to detect abnormal cells in cytopathology robustly. This paper proposes a style-aligned image composition (SAIC) method that composes high-fidelity and style-preserved pathological images to enhance the effectiveness and robustness of detection models. Without additional training, SAIC first selects an appropriate candidate from the abnormal cell bank based on attribute guidance. Then, it employs a high-frequency feature reconstruction to achieve a style-aligned and high-fidelity composition of abnormal cells and pathological backgrounds. Finally, it introduces a large vision-language model to filter high-quality synthesis images. Experimental results demonstrate that incorporating SAIC-synthesized images effectively enhances the performance and robustness of abnormal cell detection for tail categories and styles, thereby improving overall detection performance. The comprehensive quality evaluation further confirms the generalizability and practicality of SAIC in clinical application scenarios. Our code will be released at <https://github.com/Joey-Qi/SAIC>.

**Keywords:** Cytopathological Diagnosis, Abnormal Cell Detection, Data Augmentation, Image Composition

## 1. Introduction

Due to its non-invasive, efficient, convenient, and cost-effective advantages, abnormal cell identification from cytopathological images has been widely applied in clinical diagnostics.

---

\* Equally contributed to this work.

† Corresponding author.

However, traditional cytopathological screening, such as the ThinPrep Cytology Test (TCT) relies on human experts to identify abnormal cells from gigapixel whole slide images (WSI), which is time-consuming, tedious, and heavily dependent on the subjective expertise of pathologists, posing challenges on scarce and imbalanced medical resources (Pan et al., 2024). Consequently, introducing deep learning methods to address abnormal cell detection in pathological images to improve efficiency and reduce the workload of physicians has become a prominent research focus (Yin et al., 2024).

Collecting training sets for abnormal cell detection models faces numerous issues. First, curating a large-scale and accurate dataset needs highly specialized expertise and substantial labor costs. Second, the imbalanced abnormal cell distribution exhibits a typical long-tailed characteristic, leading to performance and robustness concerns in tail categories. Lastly, pathological images acquired from different institutions or periods often exhibit significant differences in staining styles and image qualities, requiring robustness to handle the variation and diversity in real-world scenarios.

One mainstream approach to addressing these issues is to introduce data augmentation. Traditional methods like affine transformations or noise addition increase the diversity of geometric and local perturbations. However, they fail to effectively address the lack of diversity in long-tailed categories and staining styles. With the advancement of GANs (Goodfellow et al., 2020) and diffusion models (Ho et al., 2020), generation-based data augmentation has gained significant attention in pathological image analysis. For example, (Hou et al., 2019) proposed a GAN-based hybrid synthesis pipeline for generating pathological images using predefined rules and textures. (Shen et al., 2024) explored applying parameter-efficient fine-tuning (PEFT) techniques to customize diffusion models for synthesizing cervical cytopathological images. Despite their effectiveness, these methods require fine-tuning, and thus remain constrained by data scale, distribution, and style biases. Another mainstream strategy is to produce augmented data by composing existing foregrounds and backgrounds. For example, Paint-by-Example (Yang et al., 2023) and ObjectStitch (Song et al., 2023) use CLIP (Radford et al., 2021) image encoder to convert the foreground image as an embedding for guidance, thus painting a semantic consistency object on the background image. However, these methods are not specially designed for pathological diagnosis and exhibit deficiencies in style consistency and fidelity.

This paper proposes a novel training-free **Style Aligned Image Composition (SAIC)** framework, which seamlessly “injects” abnormal cells into specified locations of pathological images while ensuring high fidelity and consistency in categories, types, areas, and staining styles between the foreground and background. Specifically, SAIC consists of three steps: (1) **Attribute-based selection**: Using prior knowledge of category, type, and area to select candidate abnormal cells from an existing cell bank; (2) **Style-aligned composition**: Performing online staining style alignment by reconstructing high-frequency details between the abnormal cell candidates and the style reference image to ensure style consistency in the synthesized area; (3) **LVLm-based filtration**: Leveraging a large visual-language model (LVLm) to filter high-quality samples from synthesized pathological images. Experimental results demonstrate that SAIC achieves high-fidelity, style-preserving data augmentation, effectively enhances the detection of tail categories and rare styles of abnormal cells, and improves overall cell detection performance.

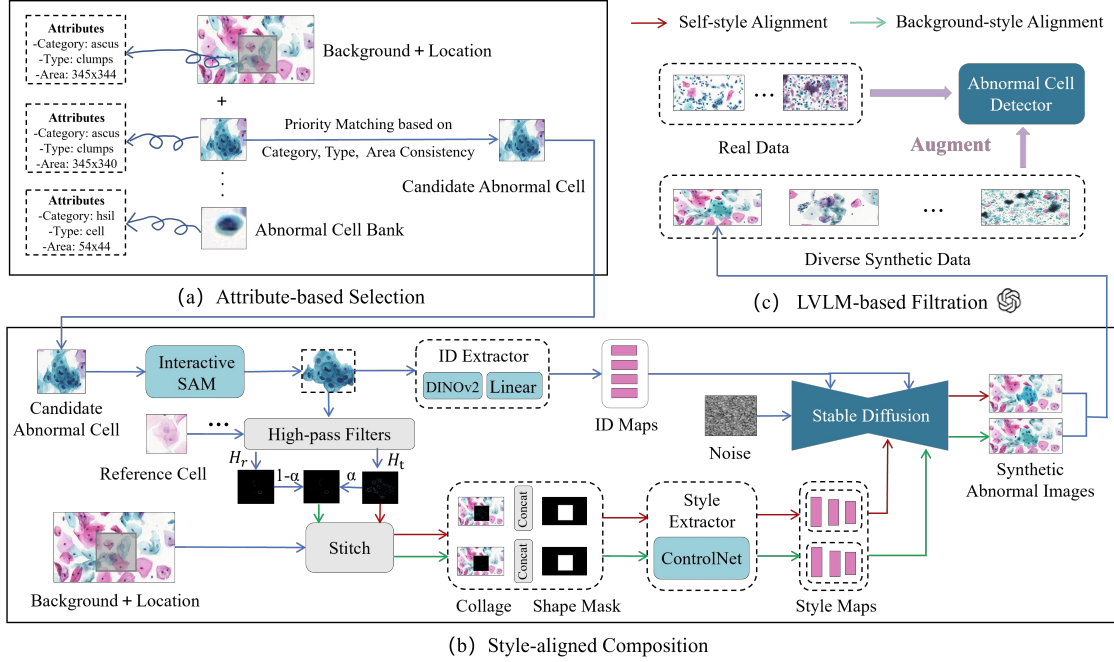


Figure 1: Overall pipeline of SAIC.

We summarize the contributions of this paper as follows: (1) We propose an image composition-based data augmentation architecture for cytopathological abnormal cell detection. (2) Compared to existing generation-based data augmentation methods, the proposed Style Aligned Image Composition (SAIC) framework offers advantages such as being training-free, style-aligned, and high-fidelity. (3) Extensive experimental results demonstrate that SAIC-synthesized images effectively improve abnormal cell detection performance, particularly for tail categories and rare staining styles, while with high-fidelity.

## 2. Method

Figure 1 illustrates the overall pipeline of SAIC. First, **Attribute-based Selection** selects a candidate abnormal cell from an abnormal cell bank based on category, type, and area consistency. Then, **Style-aligned Composition** extracts ID maps and style maps, parallelly producing two synthetic images with self- and background-style alignment. Finally, **LVLM-based Filtration** picks the image with higher fidelity as the final output.

### 2.1. Attribute-based Selection

We first create an abnormal cell bank  $\mathcal{C}_{Bank}$  that contains various abnormal cells with region and attribute annotations, which can be easily acquired from the labeled training set. Given a background of cytopathological image  $B$  and the target region  $L$  with the original cell  $c_{orig}$ , we select a candidate abnormal cell  $c_{cand}$  for composition from the cell bank  $\mathcal{C}_{Bank}$  according to the category  $\hat{m}$ , type  $\hat{t}$  (cell/clumps) and area size  $\hat{a}$  of  $c_{orig}$ :

$$\begin{aligned}
 c_{cand} &= \arg \min_{c \in \mathcal{C}_{Bank}} |a_c - \hat{a}| \\
 s.t. \ m_c &= \hat{m} \text{ and } t_c = \hat{t}
 \end{aligned} \tag{1}$$

where  $m_c$ ,  $t_c$ , and  $a_c$  represent the abnormal cell  $c$ 's category, type, and area, respectively.

## 2.2. Style-aligned Composition

As the candidate abnormal cell is selected, we must compose it with the background while preserving identity and style consistency. To this end, we design a three-step process:

**Identity map extraction.** For identity preservation, we extract discriminative identity features from the abnormal cell as an ID map. First, we segment the cell region with an interactive SAM (Kirillov et al., 2023) and center-align it. Then extract the visual feature with a DINOv2 encoder (Oquab et al., 2023), followed up with a visual-to-text linear projection to get the ID maps  $\mathcal{F}_{ID}$ .

**Style map extraction.** For style preservation, we extract the style map with high-frequency information to incorporate style-aware detail guidance. Specifically, first, we select a style reference cell  $c_{ref}$  from the abnormal cell bank  $\mathcal{C}_{Bank}$  to provide high-frequency information akin to the background's staining style. The selection process is denoted as:

$$\begin{aligned}
 S(c_{orig}, c) &= \frac{\text{DINOv2}(c_{orig}) \cdot \text{DINOv2}(c)}{\|\text{DINOv2}(c_{orig})\| \|\text{DINOv2}(c)\|}, \\
 C_{ref} &= \arg \max_{c \in \mathcal{C}_{bank}} S(c_{orig}, c)
 \end{aligned} \tag{2}$$

Subsequently, we extract high-frequency maps  $H_t$  and  $H_r$  for the candidate abnormal cell and the reference cell, respectively by High-pass Filters (Kanopoulos et al., 1988). Considering the style preservation of both the candidate cell and background, we parallelly extract two kinds of high-frequency maps for self- and background-style alignment:

$$H_n = \begin{cases} H_t, & \text{self-style alignment.} \\ \alpha \cdot H_t + (1 - \alpha) \cdot H_r, & \text{background-style alignment} \end{cases} \tag{3}$$

In practice, the reconstruction coefficient  $\alpha$  is empirically set to 0.1 to keep the most important textual information of the candidate cell.

Finally, we stitch both kinds of high-frequency maps into the background and concatenate them with the shape mask of the composition location as inputs of a ControlNet (Zhang et al., 2023) to generate hierarchical style maps  $\mathcal{F}_{style}$ .

**Conditioned Composition.** Injecting the ID map  $\mathcal{F}_{ID}$  and style map  $\mathcal{F}_{style}$  into a Stable Diffusion model (Rombach et al., 2022) to guide the composition. Specifically,  $\mathcal{F}_{ID}$  are integrated via cross-attention at each UNet layer for identity preservation, while  $\mathcal{F}_{style}$  are concatenated with decoder features at each resolution for style preservation. The latent representation of the synthesized image is generated as follows:

$$\mathbf{z}_t = \alpha_t \hat{\mathbf{x}}_\theta(\epsilon, \mathcal{F}_{ID}, \mathcal{F}_{style}) + \sigma_t \epsilon, \tag{4}$$

where  $\alpha_t$  and  $\sigma_t$  are denoising hyperparameters, which stay aligned with the setting of Stable Diffusion.

All the aforementioned models, including the interactive SAM segmenter, DINOv2-based encoder and its Linear layer, ControlNet, and Stable Diffusion, directly apply the parameters pre-trained on general datasets and do not need specific domain fine-tuning.

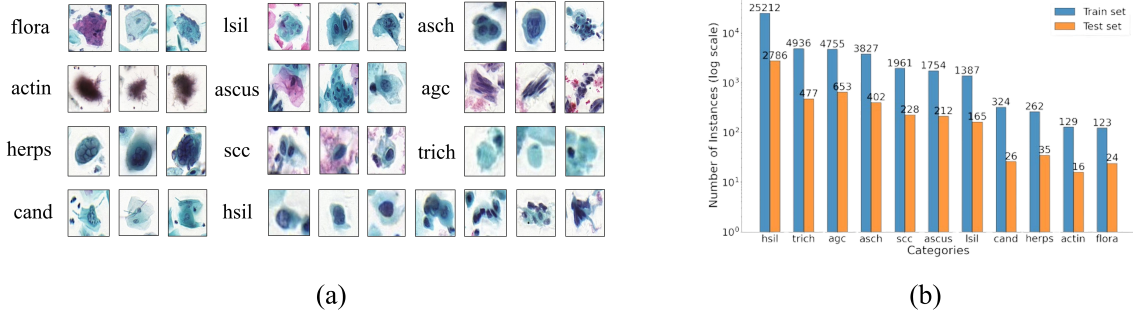


Figure 2: Overview of Comparison detector Database. Figures (a) and (b) respectively demonstrate examples of each category and their distribution in the dataset.

### 2.3. LVLM-based Filtration

As we produce two synthetic images with self- and background-style aligned in parallel (See Appendix A for examples), we need to keep the more harmonized one as the final output to raise the augmented data quality. We introduce an LVLM (GPT-4) to make the choice. The detailed prompt setting of filtration and the LVLM’s filtration ratio of the two styled images are shown in Appendix B.

## 3. Experiment

### 3.1. Datasets and Settings

**Dataset.** We conduct experiments on the Comparison Detector Database (Liang et al., 2018), the largest public dataset for cervical cancer cell detection. This dataset comprises 7,410 cervical microscopic images with 50,447 abnormal cells across 11 categories, as illustrated in Figure 2. We define the categories with fewer than 500 cells as tail categories, while the remainder are non-tail categories. The abnormal cell bank  $\mathcal{C}_{Bank}$  is composed of 824 randomly selected abnormal cells, where each category includes 68-90 samples.

**Implementation details.** We evaluate data augmentation methods using two object detectors: YOLOv8 (Varghese and Sambath, 2024) and Faster R-CNN (Ren, 2015). Models are trained using the SGD optimizer (Ruder, 2016) with an initial learning rate of 0.01, a momentum of 0.937, a weight decay of 0.0005, and a batch size of 8, for 150 epochs.

**Evaluation metrics.** We use mAP to evaluate augmentation effectiveness. Specifically,  $AP_{50}$  (a strong indicator of good localization and classification scores) is calculated with an IoU threshold of 0.5 for each category, and averaged for  $mAP_{50}$ . Additionally, we use FID (Heusel et al., 2017) to assess the overall realism of synthetic images, and DINOv2 Score to evaluate the foreground fidelity.

**Methods for comparison.** We compare three types of data augmentation methods. (1) Copy & Paste: A simple method by copying abnormal cells from the abnormal cell bank, resizing and pasting them onto the specified locations in the background images. (2) Generation-based method: We adopt GLIGEN (Li et al., 2023), an advanced diffusion model for image inpainting, following the setting of (Shen et al., 2024), we fine-tuned on

Table 1: Performance comparison of different methods (**Best results**, second best results).

Detector	Method	mAP <sub>50</sub>	Tail				Non-tail							
			flora	actin	herps	cand	lsil	ascus	scc	asch	agc	trich	hsil	
YOLOv8	Baseline	51.7	68.9	64.3	83.6	<u>54.7</u>	<u>43.8</u>	28.0	21.1	<b>19.8</b>	<u>66.3</u>	<b>68.4</b>	50.1	
	Copy & Paste	<u>52.1</u>	72.5	<u>68.9</u>	<b>86.7</b>	48.3	43.7	28.5	<b>22.9</b>	18.2	66.2	66.7	<u>50.4</u>	
	GLIGEN (CVPR 2023)	51.8	<u>77.6</u>	65.6	79.9	<b>54.8</b>	41.2	<u>30.9</u>	20.3	17.6	66.2	66.7	49.2	
	Paint-by-Example (CVPR 2023)	48.7	64.6	62.9	79.9	44.3	40.9	<u>30.9</u>	20.4	15.2	63.2	65.0	48.5	
	ObjectStitch (CVPR 2023)	50.2	69.9	66.0	84.8	47.0	40.4	30.0	17.6	17.3	64.3	65.2	49.3	
	<b>SAIC (Ours)</b>	<b>54.8</b>	<b>84.4</b>	<b>77.5</b>	<u>85.8</u>	51.7	<b>44.3</b>	<b>31.5</b>	<u>22.8</u>	<u>19.6</u>	<b>67.6</b>	<u>66.9</u>	<b>50.7</b>	
Faster R-CNN	Baseline	59.4	72.8	78.9	83.5	68.8	<b>60.7</b>	43.0	31.8	<b>23.7</b>	<u>69.6</u>	<u>67.9</u>	52.6	
	Copy & Paste	59.5	70.5	77.6	<u>83.6</u>	73.5	<u>59.9</u>	<u>44.7</u>	<u>35.0</u>	22.8	68.3	66.8	51.3	
	GLIGEN (CVPR 2023)	59.5	76.6	77.4	82.0	<b>79.5</b>	57.3	42.5	32.9	19.8	69.1	66.1	50.9	
	Paint-by-Example (CVPR 2023)	59.1	<u>81.4</u>	74.7	77.2	72.3	58.2	43.8	31.0	22.7	67.6	<b>68.4</b>	<u>52.7</u>	
	ObjectStitch (CVPR 2023)	<u>59.7</u>	77.1	<u>82.5</u>	82.9	68.2	55.6	44.1	34.8	21.9	<b>70.4</b>	65.8	<b>52.9</b>	
	<b>SAIC (Ours)</b>	<b>61.9</b>	<b>83.7</b>	<b>85.6</b>	<b>85.9</b>	<u>76.3</u>	59.4	<b>44.8</b>	<b>36.9</b>	<u>22.9</u>	68.4	65.2	52.3	

the Comparison detector Database for 50k iterations. (3) Composition-based methods: We include Paint-by-Example (Yang et al., 2023) and ObjectStitch (Song et al., 2023), two mainstream models that support the same input format as ours and require no additional parameter fine-tuning.

### 3.2. Validation of Data Augmentation Effectiveness

We validate the effectiveness of our SAIC through its improvement in detector performance and complement to staining styles in the training set.

**Improvement in detector performance.** Under identical experimental conditions of adding 5,696 synthetic images into the initial training set, we compared the performance improvements achieved by different augmentation methods. As shown in Table 1, SAIC yields the best average performance improvement across both detection models, enhancing YOLOv8 by 3.1 points and Faster R-CNN by 2.5 points. This method is especially effective for tail categories. For example, flora achieves improvements of 15.5 points for YOLOv8 and 10.9 points for Faster R-CNN; and actinomyces (actin) gain 13.2 points for YOLOv8 and 6.7 points for Faster R-CNN.

**Complement to staining styles.** For the four tail categories, we use color histograms to roughly represent their staining styles and perform t-SNE analysis on style distributions of their training sets, augmented data, and test sets. As shown in Figure 3, SAIC effectively complements the staining styles distribution in the training set, thereby enhancing the detector’s robustness to staining variations. More investigations of data augmentation effectiveness are shown in Appendix C.

### 3.3. Quality Evaluation of Augmented Image

We evaluate the quality of augmented images synthesized by SAIC through qualitative comparisons, quantitative comparisons, and a user study.

**Qualitative comparisons.** As shown in Figure 4, Copy & Paste produces synthetic images with a low informational density as it does not introduce new information. The Generation-based method GLIGEN (Li et al., 2023) generates visually realistic images but struggles with tail category representation and fidelity to real cells due to diffusion models’ bias toward simpler, in-distribution samples. In contrast, our SAIC synthesizes images






Figure 3: t-SNE analysis of staining style complement on tail categories. The first/second value of legends indicates the detection state before/after augmentation, and T/F represents the true/false detection.

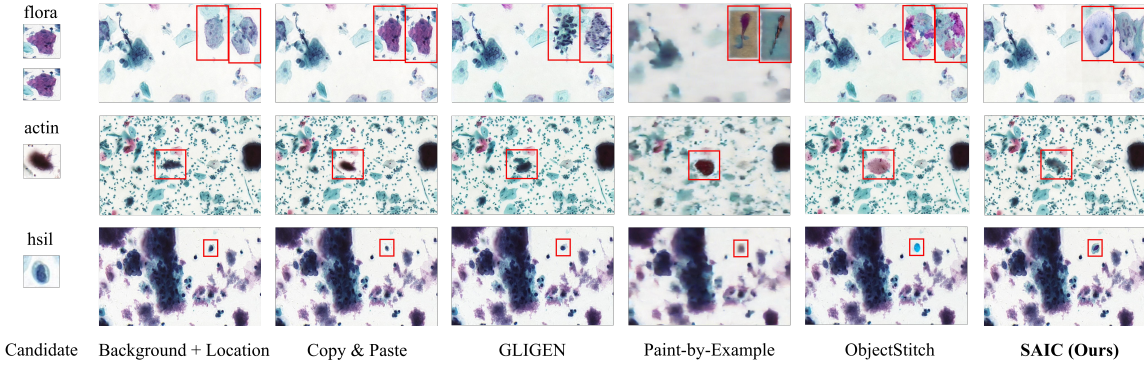


Figure 4: Qualitative comparisons across augmentation methods on flora, actin (top-2 tail categories) and hsil (top-1 non-tail category).

with high fidelity and rich informational density, outperforming other Composition-based methods. More examples are shown in Appendix D.

**Quantitative comparisons.** As shown in Table 2, compared to other composition-based methods, our SAIC significantly excels in the overall realism of synthetic images (indicated by FID), and achieves the highest foreground fidelity (indicated by DINOv2 Score). Although the fine-tuned GLIGEN (Li et al., 2023) offers marginally better overall realism, it falls short of preserving the fidelity of candidate cells used as the foreground, which limits its effectiveness for downstream tasks. In contrast, our SAIC provides a balanced performance, excelling in both aspects.

**User study.** We conduct a user study involving 8 experienced pathologists to evaluate the quality of SAIC-synthesized images. Each pathologist was required to assess 50 images (25 synthetic and 25 real) for their realism within a 30-minute timeframe. According to the mean and standard deviation results presented in Table 3, the average accuracy of distinguishing between real and synthetic images was 50%. And the judgment distributions of actual real and synthetic images are consistent. The results demonstrate that the fidelity of the synthetic images is sufficient to deceive human observers and reaffirm their high quality in supporting cytopathological diagnostics.

Table 2: Quantitative comparisons.

Framework	Method	FID ↓	DINOv2 Score ↑
Generation	GLIGEN	<b>9.1</b>	<u>81.6</u>
Composition	Paint-by-Example	191.0	74.3
Composition	ObjectStitch	95.6	79.2
Composition	<b>SAIC (Ours)</b>	<u>9.7</u>	<b>86.5</b>

Table 3: User study results.

	Pred. Real	Pred. Syn	Total
Real	14.875±3.295	10.125±3.295	25
Syn	14.875±2.976	10.125±2.976	25
Total	29.750±5.449	20.250±5.449	50

Table 4: Ablation study on impacts of various strategies.

Attribute-based Selection Category	Area	Type	Style-aligned Composition	LVLM-based Filtration	mAP <sub>50</sub> ↑		FID ↓	DINOv2 Score ↑
					YOLOv8	Faster R-CNN		
✓					51.5	58.4	12.0	86.8
✓					53.2	60.3	11.4	86.4
✓	✓				53.6	60.5	10.6	<b>87.0</b>
✓	✓	✓			53.9	60.9	10.2	<u>86.9</u>
			✓		50.6	58.8	10.5	85.2
✓	✓	✓	✓		<u>54.1</u>	61.3	<b>9.5</b>	85.8
✓	✓	✓	✓	✓	<b>54.8</b>	<b>61.9</b>	<u>9.7</u>	86.5

### 3.4. Ablation Study

We conduct comprehensive ablation studies to validate the effectiveness of the various steps employed in SAIC. The comparison results presented in Table 4 confirm that our designed steps improve the composition quality and significantly enhance detection performance. Specifically, Attribute-based Selection, despite its simplicity, yields robust improvements. Incorporating Style-aligned Composition and LVLM-based Filtration further boosts performance. Qualitative results of the ablation study are provided in Appendix E.

### 3.5. Cross Domain Application

We explore the application of SAIC in a new cellular domain (Circulating Tumor Cells) without fine-tuning, demonstrating its strong cross-domain applicability and highlighting its potential to positively impact a variety of downstream applications. Some examples are shown in Appendix F.

## 4. Conclusion

This paper proposes **Style Aligned Image Composition (SAIC)**, a training-free data augmentation architecture for cytopathological abnormal cell detection, to address issues of limited, long-tailed distributions and biased staining styles in pathological image data. By introducing *Attribute-based Selection*, *Style-aligned Composition*, and *LVLM-based Filtration*, SAIC achieves high-fidelity and style-preserved data augmentation. Experimental results demonstrate that, compared to the existing data augmentation methods, SAIC-synthesized data more effectively enhances the performance and robustness of abnormal cell detection models for pathological images, showing notable advantages for tail categories. Moreover, SAIC exhibits outstanding fidelity and generalizability. This framework provides a universal data augmentation solution for cytopathological images and can potentially impact various cross-domain applications positively.



## Acknowledgments

This work is supported by the Natural Science Foundation of China under Grant 62271465, Suzhou Basic Research Program under Grant SYG202338, and Open Fund Project of Guangdong Academy of Medical Sciences, China (No. YKY-KF202206).

## References

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Le Hou, Ayush Agarwal, Dimitris Samaras, Tahsin M Kurc, Rajarsi R Gupta, and Joel H Saltz. Robust histopathology image analysis: To label or to synthesize? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8533–8542, 2019.
- Nick Kanopoulos, Nagesh Vasanthavada, and Robert L Baker. Design of an image edge detection filter using the sobel operator. *IEEE Journal of solid-state circuits*, 23(2):358–367, 1988.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22511–22521, 2023.
- Yixiong Liang, Zhihong Tang, Meng Yan, Jialin Chen, Qing Liu, and Yao Xiang. Comparison-based convolutional neural networks for cervical cell/clumps detection in the limited data scenario. *arXiv preprint arXiv:1810.05952*, 2018.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- Qiao Pan, Bin Yang, Dehua Chen, and Mei Wang. Enhancing model generalization of cervical fluid-based cell detection through causal feature extraction: A novel method. In *Asian Conference on Machine Learning*, pages 1055–1070. PMLR, 2024.

- 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Shaoqing Ren. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- Zhenrong Shen, Manman Fei, Xin Wang, Jiangdong Cai, Sheng Wang, Lichi Zhang, and Qian Wang. Two-stage cytopathological image synthesis for augmenting cervical abnormality screening. *arXiv preprint arXiv:2402.14707*, 2024.
- Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, and Daniel Aliaga. Objectstitch: Object compositing with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18310–18319, 2023.
- Rejin Varghese and M Sambath. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6. IEEE, 2024.
- Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18381–18391, 2023.
- Juanjuan Yin, Qian Zhang, Xinyi Xi, Menghao Liu, Wenjing Lu, and Huijuan Tu. Enhancing cervical cell detection through weakly supervised learning with local distillation mechanism. *IEEE Access*, 2024.
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.

## Appendix A. Staining Style Alignment

Figure 5 shows the composition results with self- and background-style alignment. Through self-style alignment, synthetic cells can effectively retain their own staining styles. However, this may result in inconsistencies with the staining style of the background image. By applying background-style alignment, synthetic cells can integrate more harmoniously into the background.

Note that we do not simply apply background-style alignment. The main reason is that, given the limited size of the abnormal cell bank, the reference cells selected using the DINOv2 Score may still be inconsistent with the staining style of the background image. In such cases, applying background-style alignment could negatively impact the fidelity of synthetic images.

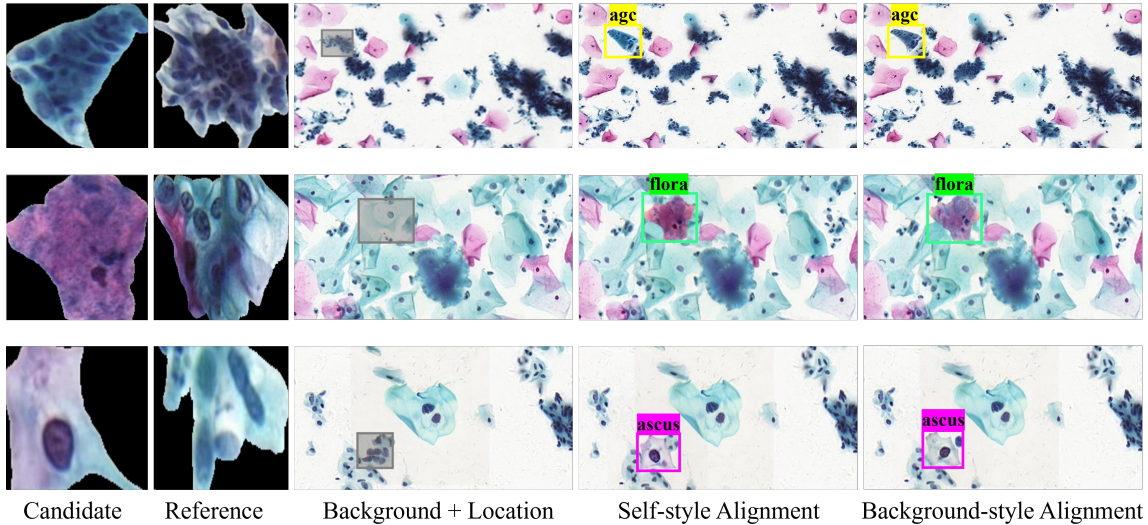


Figure 5: Demonstration of staining style alignment.

## Appendix B. LVLM-based Filtration

Figure 6 shows the detailed prompt setting of LVLM-based Filtration. Leveraging the perception of large vision-language models like GPT-4 for images and their understanding of text, we can design appropriate prompts to enable them to automatically filter the more harmonized one from two synthetic images and provide comprehensive reasons.

Note that in the subsequent experiments, through LVLM-based filtration, 64% of the synthetic images produced by our SAIC are derived from background-style alignment, while 36% are derived from self-style alignment.

## Appendix C. More Investigations of Data Augmentation Effectiveness

**Effectiveness of data augmentation across scaling ratios.** To validate the effectiveness of SAIC in relieving the practical puzzle of large-scale training dataset collection,




Figure 6: Demonstration of LVLM-based Filtration.




Figure 7: More investigations of data augmentation effectiveness.

we evaluate our framework’s performance with varying amounts of initial training data. Reduced training sets are sampled from the original data at proportions of 0.1 to 0.9, maintaining class distribution, and consistent data augmentation is applied to each reduced set. As shown in Figure 7 (a), our SAIC significantly enhances detector performance compared to the baseline. However, as the initial training data increases, the baseline accuracy and the improvement from augmentation gradually converge, likely due to the limited diversity of the data distribution.

**Effect of data augmentation across expanding ratios.** We also evaluate the impact of different degrees of data augmentation on the training set sampled from the original data at the proportion of 0.1. As shown in Figure 7 (b), the improvement on detector performance initially improves but declines as the degree of augmentation increases. This decline occurs because excessive synthetic data skews the training distribution, reducing alignment with real test data.

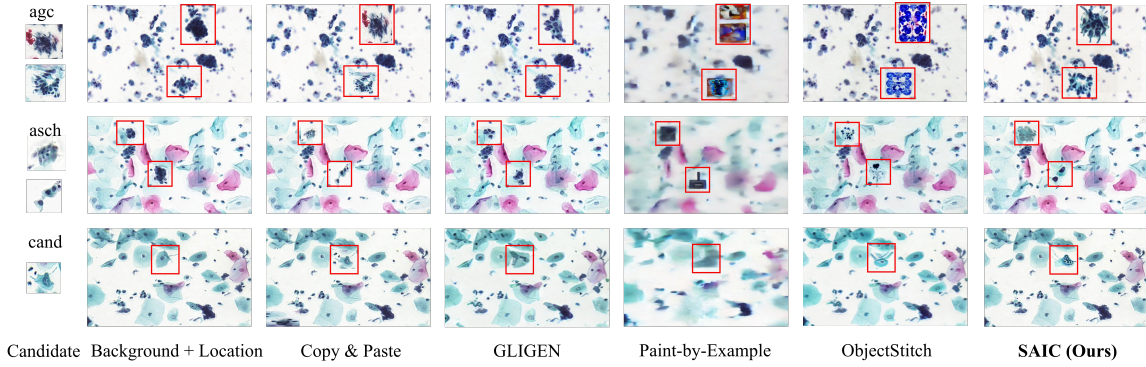


Figure 8: Qualitative comparisons across augmentation methods on agc, asch, and cand.

## Appendix D. More Qualitative Comparisons

Figure 8 shows more examples of qualitative comparisons. For candidate cells with various attributes (category, type, and area) and different background images, our SAIC consistently synthesizes images with high fidelity and rich informational density, outperforming other augmentation methods.

## Appendix E. Ablation Study

Figure 9 shows the qualitative results of the ablation study. The first and last columns show the real background image (Real) and the synthesis results of the full framework (Full), respectively. The intermediate columns illustrate the effects of omitting core strategies: (1) None: No strategies applied, resulting in random synthesis with disrupted cell distribution patterns; (2) w/ Stage 1: Attribute-based Selection strategy applied alone, yielding decent synthesis quality; (3) w/ Stage 2: Style-aligned Composition strategy applied alone, resulting in a lack of harmony between the foreground and background.

Note that in our experiments, through LVLm-based filtration, 64% of the synthetic images produced by our SAIC are derived from background-style alignment, while 36% are derived from self-style alignment.

## Appendix F. Cross Domain Application

Figure 10 briefly shows the application of our SAIC in the domain of circulating tumor cells. Although circulating tumor cells are less complex than the diverse cervical cancer cells, considering that our SAIC requires no fine-tuning, it still demonstrates strong cross-domain applicability and its potential to positively impact a variety of downstream applications.



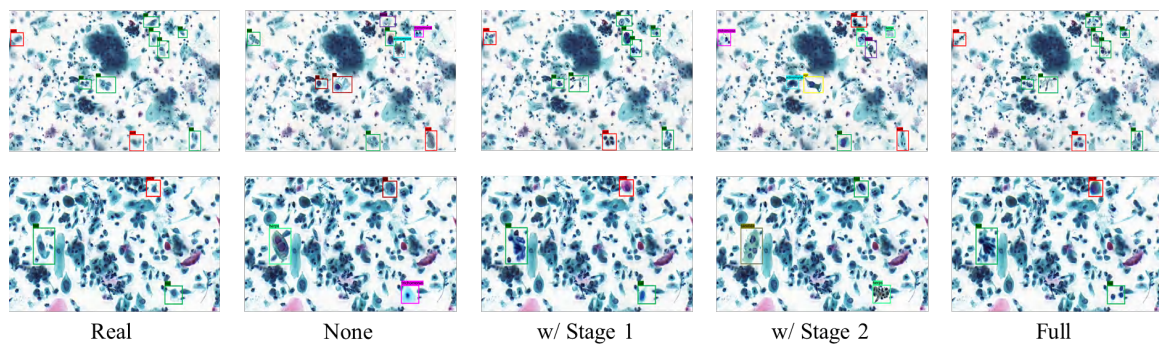


Figure 9: Qualitative results of the ablation study.




Figure 10: Demonstration of cross domain application.