



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Few-Shot Domain Adaptive Object Detection for Microscopic Images

Sumayya Inayat, Nimra Dilawar, Waqas Sultani, Mohsen Ali

Intelligent Machines Lab, Department of Artificial Intelligence,
Information Technology University, Pakistan
{sumayya.inayat, nimra.dilawar, waqas.sultani, mohsen.ali}@itu.edu.pk

Abstract. Currently, unsupervised domain adaptive strategies proposed to overcome domain shift, are handicapped by the requirement of large amount of target data. On the other hand medical imaging problems and datasets are often characterized not only by scarcity of labeled and unlabeled data but also class imbalance. Few-shot domain adaptive object detection (FSDAOD) addresses the challenge of adapting object detectors to target domains with limited labeled data. However, existing FSDAOD works struggle with randomly selected target domain images which might not represent the target distribution, resulting in overfitting and poor generalization. We propose a novel FSDAOD strategy for microscopic imaging to tackle high-class imbalance and localization errors due to foreground-background similarity. Our contributions include: a domain adaptive class balancing strategy for few shot scenario and label dependent cross domain feature alignment. Specifically, multi-layer instance-level inter and intra-domain feature alignment is performed by enhancing similarity between the instances of classes regardless of the domain and increasing dissimilarity between instances of different classes. In order to retain the features necessary for localizing and detecting minute texture variations in microscopic objects across the domain, the classification loss was applied at feature-map before the detection head. Extensive experimental results with competitive baselines indicate the effectiveness of our proposed approach, achieving state-of-the-art results on two public microscopic datasets, M5 [12] and Raabin-WBC [10]. Our method outperformed both datasets, increasing average mAP@50 by 8.3 points and 14.6 points, respectively. The project page is available here. ¹

Keywords: Few shot domain adaptive Object Detection · Second keyword · class-balancing-cut-paste.

1 Introduction

Deep-learning methods have shown substantial success for microscopic analysis [16,19,9,3,21,18,14], however, they require a lot of expert-annotated data, which is hard to obtain due to limited expert time and availability. This pose a significant limitation when required to train on one domain (source) and tested

¹ <https://im.itu.edu.pk/few-shot-DAODMI/>

on other (target), due to domain shift. Current unsupervised domain adaptation methods, have unrealistic assumption of access to large unlabeled target dataset. A more realistic scenario is having few labeled image samples from target domain, a challenge known as Few-Shot Domain Adaptive Object Detection (FSDAOD) [5,4,20,8].

The discrepancy between the source (abundant train data) and target (few-shot samples) domains arises due to differences in data acquisition protocols [15,20], such as microscopic quality, lighting conditions, microscopic-lens resolution, and camera-lens quality. This challenge is further compounded by extreme data imbalance in microscopic cells [6,12,13], visual similarity between the background and foreground, and intra-foreground visual similarity, leading to higher false-positive and false-negative rates. Recently, unsupervised and few-shot domain adaptive approaches have been proposed [4,5,20] to tackle this limitation. However, as our experimental results show (Table 1, 2, 3), these methods suffer from overfitting, leading to less generalization.

Since in few-shot domain adaptive object detection, only a few samples of any class are available from the target domain, a strategy for the feature alignment could be to enforce that, the representation of the same class across the domain is same. However, this results in over-adaptation to the smaller few-shot set. To address this challenge, we propose *intra and inter-Domain Feature Alignment* technique; I2DA, that addresses **(a)** the cross-domain shift between similar class cells by aligning the inter-domain feature level representations of cells coming from the same classes, and **(b)** Intra-Domain Feature Consistency at the cell level to learn distinguishable features for each class because the foreground cells in microscopic datasets possess high visual similarity with the background cells. This is especially challenging in the case of malarial-affected cells where the foreground classes are very similar to the background platelets, resulting in a higher rate of false positives (Supp. Figure 1). Secondly, we propose a *Domain-Generalized Class Balancing Cut-Paste strategy*; CBCP to tackle the extreme class imbalance when dealing with microscopic datasets. CBCP balances the overall count of the rare and abundant classes in the data. We propose generating samples of rare classes to match the abundant class. New samples are created by replicating existing ones with target domain-inspired visual augmentations to avoid over fitting. Simply stating, we crop the existing cells, apply visual augmentations, and paste them in carefully selected locations in the images where no other cells are present. The extensive experimentation on two public microscopic datasets M5-Malaria [15] and Raabin-WBC [11] demonstrate the effectiveness of our method by outperforming with an increase in average mAP@50 by 8.3 point and 14.6 points (respectively) as compared to other competitive baselines.

2 Methodology

Preliminaries: Let $\{(x_i^s, \mathbf{y}_i^s)\}_{i=1}^{N_s} \in \mathcal{D}_s$ be dataset of images from source domain and $\{(x_i^t, \mathbf{y}_i^t)\}_{i=1}^{N_t} \in \mathcal{D}_t$ be few-shot dataset from the target domain. Where, $(x_i, \mathbf{y}_i)$ is image and ground-truth pair for object detection task. Note that,




Fig. 1: Class Balancing Cut Paste strategy (CBCP): We first compute the meta-data and increment-stats from the abundant source dataset and few target images and then increment cells to the images with less pre-existing cells.

$|\mathcal{D}_s| \gg |\mathcal{D}_t|$, and label space in both the domains is same. The N be the total number of classes, and both datasets suffer from class imbalance. Objective is to train the object detector $\mathcal{F}$, using the source dataset $\mathcal{D}_s$ and overcome domain shift challenges by using few shot target dataset $\mathcal{D}_t$.

2.1 Domain Generalized Class Balancing Cut-Paste

To address the limitations of existing image resampling techniques [6] we propose to increase the count of rare class instances to match the maximum instances class count. Instead of matching the number of instances per class exactly, our goal is to achieve a more realistic and balanced distribution. In the text following we will use the terminology of class-size to indicate number of instances per class **Class Balancing Cut Paste Strategy (CBCP)**: Given, $\mathcal{D}_s$ and $\mathcal{D}_t$, we construct a new dataset $\mathcal{D}_{aug}$ from their combination, such that $|\mathcal{D}_{aug}| = |\mathcal{D}_s|$ but has balanced class-sizes (Fig. 1). From $\mathcal{D}_s$ we collect following data statistics: (1) $\mathcal{O}_c$: number of instance per class c , (2) $\mathcal{I}_f = \{(x_i^s, \mathbf{y}_i^s) \text{ s.t. } |\mathbf{y}_i^s| < r \ \& \ (x_i^s, \mathbf{y}_i^s) \in \mathcal{D}_s\}$ (3) $\mathcal{I}_g = \{(x_i^s, \mathbf{y}_i^s) \text{ s.t. } |\mathbf{y}_i^s| \geq r \ \& \ (x_i^s, \mathbf{y}_i^s) \in \mathcal{D}_s\}$ (4) $\mathcal{P}_s^c$ images with c^{th} class present, $\mathcal{P}_s^c \subset \mathcal{D}_s$. We need to determine how many times each instance has to be replicated, $\mathcal{S}^c = (\max(\{\mathcal{O}_c\}_{c=1}^N) - \mathcal{O}_c) / \mathcal{O}_c$, to balance the distribution (see supplementary Algorithm 1). However, these replicated instances are only added in the images $x_i \in \mathcal{I}_f$ thus trying to keep distribution of instances in each image about same. Specifically, for each instance in $\mathcal{P}_s^c$, we create its multiple copies as per $\mathcal{S}^c$, by applying random visual augmentations. Thus creating the list P_l of candidates to be used for balancing. Next, for each $x_i \in \mathcal{I}_f$ using $\mathbf{y}_i^s$ we identify regions where no object is present. Objects from P_l are randomly selected and




Fig. 2: Proposed approach: We first build our class-wise balanced dataset through our proposed Class Balancing Cut Paste (CBCP) augmentation strategy (Fig 1), then train the model with our proposed *inter-intra-domain alignment* (I2DA); inter-domain instance feature-level alignment and intra-domain instance feature-level consistency. We extract multi-layer neck features and upsample them to a common size, followed by the extraction of pooled object-level features, which are then passed to the similarity-dissimilarity and classification module.

placed at non-overlapping locations, taking in consideration that distribution of classes is balanced.

To achieve a domain generalized we extract a random object from a random image $x_j \in \mathcal{D}_t$ and paste it in $x_z \in \mathcal{D}_{aug}$ and update the corresponding annotations. Note that for random visual augmentation, we choose random color intensity variation and random Gaussian blurring (see supplementary). We only increment cells to $\mathcal{I}_f$ to ensure more realistic real world simulation. The overall flow of the process is shown in Fig 1.

2.2 Inter-domain Alignment and Intra-domain Class Consistency

Generally feature level adversarial loss is used to align the features across domain. However, such strategies gets overwhelmed by large background information and fail to overcome domain shift at object level. Applying object level contrastive loss [2] for domain adaptive object detection, results in high false positive when there is high visual similarity between foreground and background. This is especially true for malarial peripheral blood smear slides.

To enforce feature alignment across domain, we propose maximizing the dissimilarity between instances of different classes while minimizing the similarity between instances of the same class. Specifically to achieve we formulate a novel solution, (see Fig 2), aiming to learn more robust instance features. The similarity and dissimilarity must be performed within the domain as well as across

Table 1: Results in mAP@50(%) on Malaria[15] test set.

Malaria-HCM-1000x → Malaria-LCM-1000x															
Method	mAP@50(%)			Gametocyte			Schizont			Trophozoit			Ring		
Source	19.9			3.9			0.5			55.9			19.3		
Oracle	43.7			33.3			4.3			81.6			55.7		
Shots	2	3	5	2	3	5	2	3	5	2	3	5	2	3	5
FsDet[17]	11.2	11.5	12.9	11.8	11.2	12.3	0.0	0.0	0.0	27.7	28.3	30.7	5.4	6.5	8.6
VFA [7]	8.9	6.5	8.1	9.1	3.2	24.5	0.1	3.3	0.2	17.8	19.4	6.5	12.8	0.2	1.3
FDP[1]	14.6	14.7	20.5	16.4	8.1	33.7	1.20	10.1	1.3	27.4	24.4	33.3	13.5	16.4	14.0
AsyFOD[4]	26.0	29.1	33.5	14.9	23.3	36.8	1.20	7.0	2.80	<u>59.4</u>	60.7	<u>64.7</u>	<u>28.7</u>	31.8	<u>30.9</u>
AcroFOD[5]	<u>32.9</u>	<u>42.5</u>	<u>39.1</u>	<u>27.6</u>	<u>50.9</u>	<u>62.9</u>	17.6	22.1	<u>5.40</u>	58.7	<u>62.7</u>	61.3	27.8	34.4	27.0
Ours	44.7	45.9	48.9	71.4	66.0	68.2	<u>11.4</u>	<u>18.1</u>	30.4	66.9	66.7	65.6	29.3	<u>32.6</u>	31.5

domain to learn the diversity of intra-class features and capture inter-class variations.

We extract multi-layer instance-level features from the neck of the detector $\mathcal{F}$. The reason behind choosing multi-layer neck features is to take the most representative small, medium, and large-size object-level features. We up-sample these features to $(S \times S)$ and extract all the cell features corresponding to its ground truth, followed by average pooling.

Let, n_c be the total instances of objects belonging to the class c in the batch. Let $\mathbf{v}_{ik}$ and $\mathbf{v}_{il}$ denote the feature vector of object k and l of c^{th} class. We compute pairwise cosine similarity, denote by $\text{sim}(\mathbf{v}_{ik}, \mathbf{v}_{il})$, between features of all the objects of same class. We compute the similarity loss for each level of features. We sum up all the losses to compute the overall similarity loss L_{sim} .

$$L_{sim} = \sum_c \frac{1}{\binom{n_c}{2}} \sum_{k=1}^{n_c-1} \sum_{l=k+1}^{n_c} \text{sim}(\mathbf{v}_{ck}, \mathbf{v}_{cl}). \quad (1)$$

To improve the classification between classes during the domain shift, we maximize dissimilarity loss. Instead of increasing dissimilarity between instances, we push for increasing dissimilarity between the mean representation of each class. Let $\bar{\mathbf{v}}_{c_i}$ be the mean of the n_{c_i} instances of c_i^{th} class.

We compute pair-wise similarity between mean-feature representations of classes. Specifically, similarity d is computed using cosine similarity function between pairs of mean-feature representations of two classes. If $s_d < m$, set $d = 0$. Finally, the total dissimilarity loss L_{dis} can be obtained by summing up the resulting d values for all pairwise combinations of mean feature vectors:

$$L_{dis} = \sum_{c_i=1}^{C-1} \sum_{c_j=c_i+1}^C \max \left(0, \left(\frac{\bar{\mathbf{v}}_{c_i} \cdot \bar{\mathbf{v}}_{c_j}}{\|\bar{\mathbf{v}}_{c_i}\| \|\bar{\mathbf{v}}_{c_j}\|} \right) - m \right). \quad (2)$$

We boost the learning with a classifier to learn more robust features for each class. We compute the C class-wise classification losses, let l_c^k denote the class

Table 2: Results in mAP@50(%) on Raabin-WBC [11] test set.

Raabin-WBC-HCM → Raabin-WBC-LCM															
Method	mAP50(%)			Large Lymph			Neutrophil			Small Lymph			Monocyte		
Source	27.2			25.1			59.6			22.9			1.0		
Oracle	75.0			90.9			98.1			83.2			27.7		
Shots	2	3	5	2	3	5	2	3	5	2	3	5	2	3	5
FsDet[17]	26.5	28.5	30.1	41.5	24.3	38.1	17.5	31.8	44.1	23.7	34.4	29.6	<u>23.1</u>	<u>23.7</u>	8.7
VFA[7]	30.3	33.2	45.2	44.6	28.6	59.8	25.7	59.8	66.0	<u>48.7</u>	27.7	<u>42.6</u>	2.40	16.7	12.5
FDP[1]	35.9	32.4	44.3	28.2	36.2	59.1	60.7	39.8	66.8	46.2	34.0	35.2	8.60	19.9	<u>15.9</u>
AsyFOD[4]	35.6	38.1	26.3	37.2	42.2	39.3	58.3	48.7	43.1	46.5	51.1	22.4	0.3	10.4	0.3
AcroFOD [5]	<u>44.9</u>	<u>47.2</u>	<u>61.2</u>	<u>50.5</u>	<u>64.1</u>	82.1	88.1	89.1	95.9	37.6	30.7	59.6	3.5	5.1	7.3
Ours	64.2	62.6	70.8	74.1	76.0	<u>75.2</u>	<u>87.2</u>	<u>86.0</u>	<u>94.3</u>	54.7	<u>46.8</u>	42.5	40.6	41.6	71.3

loss of instance k in c^{th} class, Let n_c be the number of instances in class c , then the instance level classification loss L_{cls} is given by:

$$L_{cls} = \sum_{c=1}^C \left(\frac{1}{n_c} \sum_{k=1}^{n_c} l_c^k \right) \quad (3)$$

Finally, we compute the mean similarity (Equation 1), dissimilarity (Equation 2), and classification (Equation 3) losses for the three levels, followed by multiplication with weights λ_1 , λ_2 and λ_3 with the similarity, dissimilarity, and class mean losses respectively. We add up the losses as our final $I2DA$ loss, L_{I2DA} is represented as:

$$L_{I2DA} = \lambda_1 L_{sim} + \lambda_2 L_{dis} + \lambda_3 L_{cls} \quad (4)$$

3 Experiments and Results

Datasets: M5[15] is a large-scale malarial domain adaptive cell detection dataset captured from two different microscopes, one high cost, and one low cost, and the corresponding images captured from three different resolution levels. We utilized their standard train val test splits for training whereas for few-shot settings, we randomly sampled a set of 8 images as per [4,5] while also selected images as per our strategy (described in section below) for 2-shot, 3-shot, and 5-shot. We consider the shots as the number of images per a specific category.

Raabin-WBC (R-WBC) is a white blood cell dataset comprising images captured by a mobile phone camera. 17819 images belong to the high-cost microscope and 3167 images belong to the low-cost microscope. The authors did not provide any standard train, val, or test splits for the detection task, hence we first extracted center cropped images and as per ‘Label2’, for the four following classes, Large Lymph, Neutrophil, Small Lymph, and Monocyte. We then made

Table 3: mAP@50(%) on [15] & [11] test sets on 8 random few-target images.

Data	Malaria					Raabin-WBC				
	mAP@50	Gamet.	Schizo.	Troph.	Ring	mAP@50	L-Lymp.	Neutro.	S-Lymp.	Mono.
AsyFOD	30.2	23.8	1.3	<u>61.8</u>	33.9	33.7	28.4	48.6	56.2	1.4
AcroFOD	<u>33.1</u>	<u>46.8</u>	3.8	56.9	24.9	<u>48.9</u>	<u>69.1</u>	90.7	27.7	<u>7.9</u>
Ours	40.3	62.3	<u>2.0</u>	64.4	<u>32.6</u>	55.7	71.3	<u>80.6</u>	<u>49.6</u>	23.9

Table 4: Ablation of each component of our method on a 5-shot set

CBCP	L_{sim}	L_{dis}	L_{cls}	M5	R-WBC
✓	✗	✗	✗	42.1	<u>69.2</u>
✗	✓	✗	✗	40.3	63.8
✗	✗	✓	✗	41.9	65.2
✗	✗	✗	✓	41.9	59.3
✗	✓	✓	✗	42.4	68.5
✗	✓	✗	✓	44.3	67.6
✗	✗	✓	✓	43.9	65.4
✗	✓	✓	✓	<u>45.4</u>	66.5
✓	✓	✓	✓	48.9	70.8

random splits of the train (2052 images), val (150 images), and test (450 images) for both the microscope data and chose the few-shot samples similar to M5.

Implementation Details: Our proposed approach is object detector agnostic, however, for a fair comparison with [4,5], we have used [10] as a base model. For our experiments (conducted on GTX1080 GPU), batch-size was set to 4. We develop customized batches for each epoch such that each batch of the extracted features contains $n \geq 1$ object from the few-shot target set. For each batch, we select 1 image from the few-shot set, 2 images from the real source, and 1 image from the augmented source. However, for training with larger batch sizes, we suggest allocating 2% of the batch size to the real target, 68% to the real source, and 30% to the augmented source dataset. The λ_1 λ_2 and λ_3 values are set to 0.005, 0.005, and 0.001 respectively. We chose low weights to scale the complete I2DA loss with the YOLOv5 detection loss.

3.1 Results

We perform two sets of experiments, one with 8 random images as per [4,5] that may have any number of images per class or even miss a rare class. The other set of experiments is performed on our variation of k-shot settings. We define our shots as k images per class. Malaria results of the baselines [17,7,1,4,5] and our work on 2-shot, 3-shot, and 5-shot images are shown in Table 1. We evaluated our models based on mAP@50 because all the given baselines [17,7,1,4,5] yielded results in mAP@50. But for [4,5] we additionally evaluated the models on mAP@50:95 and average precision and average recall. Note that we performed 100 epochs of adaptation and reported results over the test dataset for the model obtained after 100 epochs. The same experimental settings are used for all the




Fig. 3: Qualitative results of LCM Malarial defected regions after adaptation.

comparative methods. As shown in Table 1, our work outperforms the existing competitive baseline by a good margin in each shot setting. Table 2 shows the results of the Raabin-WBC test set and proves that our method works well for large-size objects as well. Table 3 shows our results obtained on 8 random target domain images and as visible our method outperforms in these settings as well. Please refer to the supplementary for detailed results and comparisons. Fig. 3 shows the qualitative results of AcroFOD [5] and our methods. AcroFOD has given some wrong predictions and was comparatively less confident in the correct predictions. In contrast, ours is more confident in the correct predictions and the false-positive rate is comparatively less.

3.2 Ablation

Table 4 shows the ablation studies of each of the I2DA loss components and the CBCP module on 5-shot sets of M5 and Raabin-WBC, respectively. The results validate the effectiveness of each component of the I2DA loss and the CBCP module. I2DA and CBCP complement each other hence obtaining the best results when applied together. We also performed ablation of our CBCP augmentation with AcroFOD alignment algorithm. AcroFOD alignment results w/o augmentation (M5:37.8 | RWBC:58.7) are much less than our alignment-only I2DA results (M5:44.2 | RWBC:66.5). AcroFOD alignment prioritizes target-similar examples, but we believe that small few-shot sets can't fully represent the whole target population. Our approach extracts moderate knowledge from the target set, ensuring generalizability to larger test sets while optimizing performance across the majority of few-shot scenarios.

4 Conclusion

We have put forward a novel solution to tackle **FSDAOD** in few shot settings in microscopic imaging. The intra-class feature space variation is minimized and

inter-class variation is maximized irrespective of domains to boost the performance with a specialized feature-level instance classifier. To handle the extreme class imbalance in microscopic datasets, especially in domain adaptive few-shot settings, we devise a novel strategy to balance the skewed data distribution with our cut-paste augmentation strategy. Extensive experimentation validate the effectiveness of our method as compared to the existing competitive baselines. Our method outperform the competitive baselines on average 8.3 points on M5 dataset and 14.7 points on Raabin-WBC demonstrating its capability to handle variable cell sizes.

Acknowledgments. This research was partially funded by the Google Research award.

Disclosure of Interests. The authors have no competing interests.

References

1. Fine-grained prototypes distillation for few-shot object detection
2. Cao, S., Joshi, D., Gui, L.Y., Wang, Y.X.: Contrastive mean teacher for domain adaptive object detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23839–23848 (2023)
3. Fujita, S., Han, X.H.: Cell detection and segmentation in microscopy images with improved mask r-cnn. In: Proceedings of the Asian conference on computer vision (2020)
4. Gao, Y., Lin, K.Y., Yan, J., Wang, Y., Zheng, W.S.: AsyfoD: An asymmetric adaptation paradigm for few-shot domain adaptive object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3261–3271 (2023)
5. Gao, Y., Yang, L., Huang, Y., Xie, S., Li, S., Zheng, W.S.: AcrofoD: An adaptive method for cross-domain few-shot object detection. In: European Conference on Computer Vision. pp. 673–690. Springer (2022)
6. Gosain, A., Sardana, S.: Handling class imbalance problem using oversampling techniques: A review. In: 2017 international conference on advances in computing, communications and informatics (ICACCI). pp. 79–85. IEEE (2017)
7. Han, J., Ren, Y., Ding, J., Yan, K., Xia, G.S.: Few-shot object detection via variational feature aggregation. arXiv preprint arXiv:2301.13411 (2023)
8. Han, L., Zhai, J., Yu, Z., Zheng, B.: See you somewhere in the ocean: few-shot domain adaptive underwater object detection. *Frontiers in Marine Science* **10**, 1151112 (2023)
9. Huang, J., Shen, Y., Shen, D., Ke, J.: Ca 2.5-net nuclei segmentation framework with a microscopy cell benchmark collection. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII 24. pp. 445–454. Springer (2021)
10. Jocher, G., Stoken, A., Borovec, J., NanoCode012, ChristopherSTAN, Changyu, L., Laughing, tkianai, Hogan, A., lorenzomammama, yxNONG, AlexWang1900, Diaconu, L., Marc, wanghaoyang0106, ml5ah, Doug, Ingham, F., Frederik, Guilhen, Hatovix, Poznanski, J., Fang, J., , L.Y., changyu98, Wang, M., Gupta, N., Akhtar, O., PetrDvoracek, Rai, P.: ultralytics/yolov5: v3.1 - Bug Fixes and Performance

- Improvements (Oct 2020). <https://doi.org/10.5281/zenodo.4154370>, <https://doi.org/10.5281/zenodo.4154370>
11. Kouzehkhanan, Z.M., Saghari, S., Tavakoli, S., Rostami, P., Abaszadeh, M., Mirzadeh, F., Satlsar, E.S., Gheidishahran, M., Gorgi, F., Mohammadi, S., et al.: A large dataset of white blood cells containing cell locations and types, along with segmented nuclei and cytoplasm. *Scientific reports* **12**(1), 1123 (2022)
 12. 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)
 13. Liu, M., Li, X., Gao, X., Chen, J., Shen, L., Wu, H.: Sample hardness based gradient loss for long-tailed cervical cell detection. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 109–119. Springer (2022)
 14. Sudre, C.H., Van Wijnen, K., Dubost, F., Adams, H., Atkinson, D., Barkhof, F., Birhanu, M.A., Bron, E.E., Camarasa, R., Chaturvedi, N., et al.: Where is valdo? vascular lesions detection and segmentation challenge at miccai 2021. *Medical Image Analysis* **91**, 103029 (2024)
 15. Sultani, W., Nawaz, W., Javed, S., Danish, M.S., Saadia, A., Ali, M.: Towards low-cost and efficient malaria detection. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 20655–20664. IEEE (2022)
 16. Thomas, R.M., John, J.: A review on cell detection and segmentation in microscopic images. In: *2017 International Conference on Circuit, Power and Computing Technologies (ICCPCT)*. pp. 1–5. IEEE (2017)
 17. Wang, X., Huang, T.E., Darrell, T., Gonzalez, J.E., Yu, F.: Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957* (2020)
 18. Wang, X., Zhang, J., Yang, S., Xiang, J., Luo, F., Wang, M., Zhang, J., Yang, W., Huang, J., Han, X.: A generalizable and robust deep learning algorithm for mitosis detection in multicenter breast histopathological images. *Medical Image Analysis* **84**, 102703 (2023)
 19. Yang, S., Fang, B., Tang, W., Wu, X., Qian, J., Yang, W.: Faster r-cnn based microscopic cell detection. In: *2017 international conference on security, pattern analysis, and cybernetics (SPAC)*. pp. 345–350. IEEE (2017)
 20. Zhang, J., Chao, H., Dhurandhar, A., Chen, P.Y., Tajer, A., Xu, Y., Yan, P.: Spectral adversarial mixup for few-shot unsupervised domain adaptation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 728–738. Springer (2023)
 21. Zhao, Z., Pang, F., Liu, Z., Ye, C.: Positive-unlabeled learning for cell detection in histopathology images with incomplete annotations. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII* 24. pp. 509–518. Springer (2021)