

Towards Low-Cost and Efficient Malaria Detection

Waqas Sultani¹, Wajahat Nawaz^{1*}, Syed Javed^{1*}, Muhammad Sohail Danish^{1*}, Asma Saadia², Mohsen Ali¹
¹Intelligent Machines Lab, Information Technology University, Pakistan
²Central Park Medical College, Pakistan

Abstract

Malaria, a fatal but curable disease claims hundreds of thousands of lives every year. Early and correct diagnosis is vital to avoid health complexities, however, it depends upon the availability of costly microscopes and trained experts to analyze blood-smear slides. Deep learning-based methods have the potential to not only decrease the burden of experts but also improve diagnostic accuracy on low-cost microscopes. However, this is hampered by the absence of a reasonable size dataset. One of the most challenging aspects is the reluctance of the experts to annotate the dataset at low magnification on low-cost microscopes. We present a dataset to further the research on malaria microscopy over low-cost microscopes at low magnification. Our large-scale dataset consists of images of blood-smear slides from several malaria-infected patients, collected through microscopes at two different cost spectrums and multiple magnifications. Malarial cells are annotated for the localization and life-stage classification task on the images collected through the high-cost microscope at high magnification. We design a mechanism to transfer these annotations from the high-cost microscope at high magnification to the low-cost microscope, at multiple magnifications. Multiple object detectors and domain adaptation methods are presented as the baselines. Furthermore, a partially supervised domain adaptation method is introduced to adapt the object-detector to work on the images collected from the low-cost microscope. The dataset is available here: <http://im.itu.edu.pk/m5-malaria-dataset/>

1. Introduction

Every year, on average, about 226 million cases of malaria are reported in 87 countries and 425,600 of them result in fatalities¹ [30]. Around 67% of all malaria deaths, in 2019, were of children under the age of five [30]. Malaria, also named as a disease of the poor, mostly infects the population in resource deficient tropical and sub-tropical regions

* Authors contributed equally

¹average over the last five years

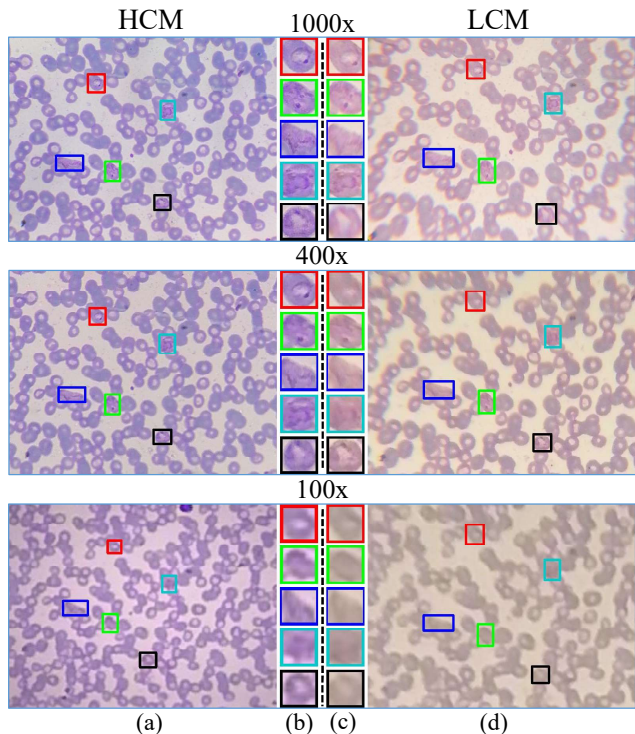


Figure 1. (a) shows images captured using high-cost microscope (HCM) at 1000x, 400x and 100x magnifications and (d) shows the same for low-cost microscope (LCM). (b) and (c) shows malarial cells from HCM and LCM respectively. Note that the images captured using HCM are clear and internal cells structures are more evident as compared to that of LCM. HCM cost 17 times more than LCM.

with a weak health care system. Although fatal, malaria is preventable and curable. Even in developed countries, the leading cause of death in malaria is considered to be a delay in diagnosis and treatment [2]. Early detection not only allows avoiding the medical complications but also prevents further spread [20].

Although the testing kits are becoming common, microscopic analysis of the stained blood slides is still considered a gold standard malaria diagnosis [6, 17, 27, 28]. However,

malaria microscopic analysis is a taxing process, requiring the availability of expensive microscopes and trained experts. All these factors limit the accessibility to the early and correct malaria diagnosis, specifically in remote and resource constraint areas. To tackle the subjectivity and shortage of doctors, several microscopic malarial image analysis methods have been proposed. This includes approaches using hand-crafted visual features [14, 23, 29, 39] and deep learning based approaches [17, 28, 33, 41, 42]. To achieve a good detection accuracy, these deep learning based algorithms require large datasets and their accuracy is dependent on the correctness of the annotation of these datasets. Although appealing, the previously proposed methods do not address the core limitations of malaria detection which is to detect malaria on low-cost microscopes (as malaria mostly occurs in resource-constraints areas) and on the low-magnifications lens (for efficient diagnostics). To achieve low-cost and efficient malaria detection using deep learning, it is vital to collect malaria data from the low-cost microscope and at lower magnifications.

Low-cost microscopes (LCM), although more than 70% cheaper than high-cost microscopes (HCM) suffer from a limited field of view (FOV) and less clear image due to the low-quality lens (see Figure 1). Due to these limitations, it is challenging and time-consuming for medical experts² to find the malaria cells at LCM and to confirm the malaria life-stage. On every microscope, the images could be captured at multiple magnifications such as 100x, 400x, and 1000x. Since 1000x makes the internal structure of cells very clear, it is a standard practice by doctors to analyze malarial slides at 1000x. FOV of 1000x is much smaller than that of 400x and 100x (see Figure 2 and Section 3). As compared to 1000x, at 100x or 400x a much larger portion of blood-smear slide is visible at one time and therefore the whole slide could be traversed in less time for automatic malaria detection. Thus a higher accuracy at lower-magnification will be beneficial for efficient analysis.

To cater the above-mentioned challenges with low-cost microscopes, we put forward a new large scale *multi-microscope multi-magnification malaria (M5)* dataset. Our dataset contains images of the *same* slide regions captured using HCM and LCM and at multiple magnifications (1000x, 400x, and 100x). As compared to the usual challenges faced in the data collection/annotations, computer-aided malaria data collection suffers additional challenges. First of all, the medical expert needs to manually locate all the malarial cells regions in the high-cost microscope at 1000x slides. Due to the mechanical parts and the properties of the lens, annotations collected in the high magnification are not easily transferable to low magnifications. A slight movement of the stage controls (due to manual manipulation) results in a large shift of the view, especially in high

²Verified after contacting multiple doctors in different major hospitals.

magnification. As we move from high magnification to low magnification, a portion of the slide in the center of the view (in high magnification) might not remain in the center of the view at the low magnification (see Figure 2 and Figure 3). Also, every time, we transverse the slide or change the magnifications, the slide needs refocusing. Capturing the same HCM slide regions in LCM is further challenging. As we move from the HCM to the LCM, FOV decreases resulting in loss of annotations, even if the center area is perfectly aligned (which itself is a challenging problem). Therefore, to recover malarial cells regions *manually* for collection of the dataset at multiple magnifications across multiple microscopes, we need to record and discover those locations manually at **millimeter** accuracy.

Our dataset can not only be used for the malarial cell localization and stage classification, but also for the across microscopes domain adaptation task. We provide baselines for both tasks on our dataset. Multiple object detectors [22, 35, 36, 40] are trained and their results are reported for HCM at multiple magnifications. As mentioned above, annotations for HCM are challenging to obtain but still comparatively easier than gathering annotations on LCM, therefore, we show adaptation results for detectors trained on HCM (source domain) to the LCM (target domain). Finally, we also present our partially supervised domain adaptation strategy and report state-of-the-art results. In summary, this paper has following contributions:

- 1) We collect the first large-scale multi-microscope multi-magnification malarial image dataset from the thin-blood smear slides. Dataset has been annotated by a medical expert with more than 20 years of experience, for both the malarial cell localization and life-stage classification tasks.
- 2) An annotation mechanism (consisting of multiple mobile cameras and microscopes) is designed that records the microscope's x-y transnational control knobs positions to speed up the data collection, handles issues of calibrations across magnifications and also across microscopes.
- 3) We compute the baseline results of several object detectors and domain adaptation methods on our dataset, creating a benchmark for the computer-aided malaria detection task.
- 4) Finally, utilizing the problem-specific constraints, a partially supervised domain adaptation mechanism is designed.

2. Related work

Recently, the vision community has shown a growing interest in collecting large-scale medical imaging datasets and bench-marking recent computer vision and deep learning models on them. To make chest X-rays analysis low-cost and scalable, Phillips et al., [32] presented a dataset contain-




Figure 2. Top row presents the relationship between different magnification lenses of microscope. The second shows that the patch can appear at different locations in the same 400x FOV due to manual movement of microscope’s stage.

ing pictures of chest X-rays captured by different mobile phone cameras. To address commonly occurring thoracic diseases, Wang et al., presented chest X-rays dataset [43] with image-level labels and provided weakly supervised baseline results. To facilitate Tuberculosis diagnosis, Liu et al., [24] collected Tuberculosis X-ray (TBX11K) dataset containing X-ray images along with corresponding bounding box annotations. They provided experimental results of different object detectors on their dataset. Recently, domain adaptation strategies have been applied to medical imaging [3, 21, 46].

Historically, taking a page from image analysis, several malaria detection techniques were introduced using hand-crafted visual features. Authors in [11, 12, 14, 34, 39] employed morphological operations to perform cell enhancement and segmentation followed by malarial cell classification using k-nearest neighbours, binary trees etc. Similarly, authors in [19, 23, 29] extracted different visual features such as SIFT [26], local binary patterns [4], and classified cells using support vector machines [10].

Recently, several deep learning-based approaches have been presented for automatic malaria detection. Hung et al., [17] proposed a two-stage malaria detection approach. In their first stage, they applied Faster-RCNN to discriminate healthy versus malarial cells. During the second stage, they employed AlexNet [18] for malarial stage classification. Vijayalakshmi et al., [42] proposed to extract CNN features from the VGG-16 network followed by SVM classification. Umer et al., [41] presented a stacked convolutional neural network for malaria detection in thin-film slides. Rajaraman et al., [33] and Mehanian et al., [28] analyzed different CNN architectures for malaria detection on thin blood smears and thick blood smears respectively. Both methods used traditional thresholding-based approaches for cell segmentation. Finally, some malaria datasets have been proposed recently [1, 5, 25, 38].

As compared to the above-mentioned approaches, our main focus in this paper is to collect a multi-microscope multi-magnification malaria dataset and evaluate different CNN object detectors and domain adaptation methods.

3. Preliminaries

Malaria: Due to low cost, high sensitivity, and specificity, an optical microscope is a gold standard for diagnosing malaria by examining thin blood smears [31]. This is in contrast to Polymerase Chain Reaction (PCR) [14] and Rapid Diagnostic Test (RDT) [13], which are either expensive or have low accuracy. As compared to other malaria types, malarial parasites named *Plasmodium vivax* (*P. vivax*) are more common and fatal, therefore we collected data from *P.vivax* infected patients. *P.vivax* has four stages (classes) of development in the human body. These stages are ring, trophozoite, schizont, and gametocyte. Classification of a cell infected by malaria into these four classes is crucial to identify the seriousness of the disease and for the prescription of the best medication.

Relationship between different microscopes: The quality and price of microscopes are governed by their field of view, precise movement of the mechanical stage, and most importantly the quality of the magnification lenses (e.g., plan achromatic versus non-plan achromatic). The price range for good quality microscopes is around several thousand dollars. In our setup, the price of the high-cost microscope (Olympus CX23) is around 3,000 USD and the price of the low-cost microscope (XSZ-107BN) is around 160 USD. The high-cost microscopes, although preferred by hematologists and lab technicians, are too expensive to be readily available at all locations. At the same time, resource-constrained areas have very few highly trained experts available. These two constraints result in the situation where automatic methods for the microscopic analysis of the slides, especially for the low-cost microscopes are highly needed.

Relationship between different microscopic magnifications: The magnification of a microscope is measured by eyepiece magnification \times objective lens magnification. Most of optical microscopes are equipped with three magnification scales: 100x, 400x and 1000x which makes cell/tissues look 100 times, 400 times, and 1000 times bigger respectively. Note that 1000x is also called immersion oil magnification since oil must be placed on blood smear to see clear cell structure at 1000x. In Figure 2 (top row), we depict the relationship between different magnifications. A FOV on 100x covers around 20 FOVs of 400x and almost 180 FOVs of 1000x. For clarity purposes, we manually moved the patch of interest to the center location. In reality, the *same* patch appear may appear in different locations at different magnifications. The patch correspondence between magnifications is different for different microscopes and is discovered during the microscope calibration step.

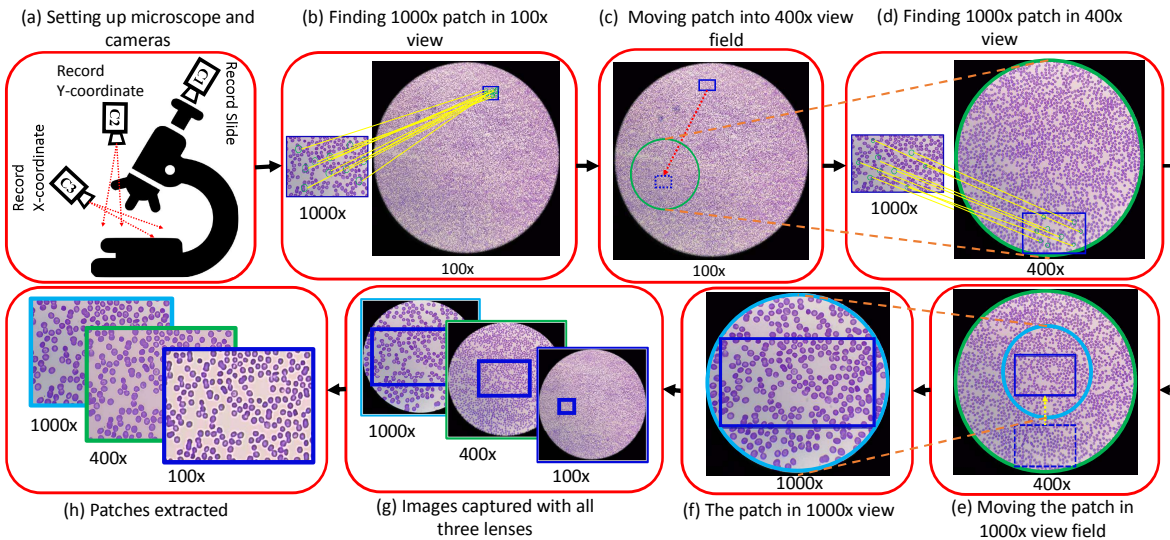


Figure 3. Steps to capture malarial region images at multiple magnification on HCM. (a) Initial image captured by hematologist at 1000x along microscope’s stage x-y locations pictures, (b) captured image is detected at 100x view, (c) 100x view image is tracked to a new location to be viewed at 400x, (d) The original 1000x patch is found in the 400x view followed by (e) tracking to new location to be found at 1000x, (f) finally the original 1000x patch is found at 1000x. (g) and (h) shows the final extracted patches (See supplementary material).

As shown in the second row of Figure 2, due to the manual setting (in millimeter) of the microscope’s stage, on repetition, the same patch can appear at different locations.

4. Dataset Collection Process

In the process of understanding the problem and data collection, we contacted multiple health professionals and multiple hospitals. Collecting the slides of malarial infected patients and annotating them, was a challenge. During this process, it was discovered that, due to the low quality of visual information from the LCM, experts were reluctant to use LCM for diagnosis or annotation. Therefore we decided to only annotate data using HCM at 1000x. To tackle the above-mentioned challenges, we propose to capture images of malarial regions and annotations on HCM at 1000x and then transfer these annotations to multiple magnifications in and across microscopes.

Dataset collection consists of multiple steps and multiple people. Collecting the blood-smear slides from the malarial-infected patients is itself a challenging process. First of all, images of malarial regions are captured by the expert hematologist, by traversing the slide on HCM at 1000x, along with images of the exact microscope’s x-y transnational control knobs coordinates. In addition, hematologist provides helps localize (dropping a pin on) the malaria-infected cells and identify their life-stage labels. The operator then uses this information to draw the bounding boxes. After that, we locate these malarial regions at multiple magnifications on both HCM and LCM. These localizations are manually verified. The complete dataset is

finalized after intense efforts of several months. All personal information of patients has been removed to make the dataset publicly available. Table 1 demonstrates the comparison of our dataset with existing malaria datasets. None of the existing datasets capture malaria from different microscopes or at multiple magnifications.

4.1. Initial Malarial Image Collection

To annotate and collect images, we seek help from a hematologist having more than 20 years of experience. An annotation rig is designed, consisting of three mobile phone cameras and microscopes. Expert traverses the slides to identify both the images with the malarial cells and localize and classify these cells. The mobile camera is mounted on a microscope eyepiece using a cell phone stand (C1 in Fig. 3 (a)), to capture the image. Since we are aiming at a low-cost solution, we use a nominal price (<200 USD) mobile phone camera (Honor 9x Lite) to capture images without any post-processing. To transfer these annotations from one magnification to other and across the microscope, two cameras are set up (C2 & C3 in Fig. 3 (a)) around the microscope to record the microscope’s stage x-y locations. We need two cameras to avoid occlusion from microscopic nose-piece objective lenses. This allows the more efficient annotation and helps save experts time. Note that each image is captured and each malarial cell is localized by the hematologist.

4.2. Image Capturing on HCM

Our next task is to collect malarial images on different microscopic magnification scales on HCM. As shown in

Figure 3 (b), after extracting the maximum size rectangular patch that covers most of the cells at 1000x, we search that patch on 100x FOV. 100x FOV covers almost 180 FOV of 1000x, as discussed in the Section 3 under *Relationship between different microscopic magnifications*. We do not search the patch on 400x directly from 1000x as the small change in the stage location (due to the manual movement) drastically changes the view and makes it almost impossible to find the exact malarial patch. To search the patch in 100x FOV, we employ keypoint detection (Speeded-Up Robust Features (SURF) [7]) followed by RANSAC [15].

After cropping the patch in the 100x, we want to find the patch in the 400x. However, the center visible in 100x magnification does not map to the center of the view when magnification is increased while keeping the x-y location fixed (Figure 3 (c)). We perform cross magnification calibration to identify where the region from the lower magnification is shifted in higher magnification. A graphical interface is designed to guide the operator to traverse the slide manually such that the detected patch could move to the location that can be seen after switching to a 400x magnification (look at the green circle in Figure 3 (c)). Note that this step is accomplished by simultaneous manual movement of the slide and real-time tracking of the patch. Global image alignment is performed to track the patch across the frames as the operator transverse the slide. Once the microscope magnifications is changed to 400x, 1000x patch is searched again in the new FOV (Figure 3(d)). We need this search since the patch may exist anywhere within 400x FOV due to manual movement of the microscope’s stage by the operator.

Above steps are repeated as we move from 400x to 1000x magnification³ (Figure 3 (e), and (f)). Figure 3 (g) and (h) shows the final patch extracted at 100x, 400x and 1000x. Note that we need to apply immersion oil at 1000x and rinse it off at 100x and 400x from the slide for every single image. The procedure is repeated for all malarial regions.

4.3. Image Capturing on LCM

After capturing images at multiple magnifications on HCM, our next task is to capture the same malarial regions on LCM. Since we already have recorded stage x-y coordinates of malarial regions of the slide on HCM, we manually set the LCM stage to those coordinates. This step requires across microscope *stage calibrations* as stage coordinate for both microscopes for the same slide-location are shifted.

The overall procedure for finding the malarial patches on LCM is similar to that of HCM except for a few changes which we discuss below (see supplementary material). At step (b), (d) and (f) shown in Figure 3, instead of patch from 1000x magnification, 100x, 400x, and 1000x magnification

³We search 1000x again to have a good quality picture of 1000x which is sometime difficult in front of doctor due to time constraints.




Figure 4. Transferring annotations across magnifications and across HCM and LCM. The green boxes shows the malarial cells that are not in the LCM FOV (as LCM has smaller FOV than HCM). The black box’s annotations were missed during automatic transformations and were included during manual verification.

patches from HCM are searched at 100x, 400x, and 1000x FOV of LCM. The rest of the steps (c), (e) remain the same.

4.4. Malarial Cell Annotations on HCM and LCM

For any given malarial region annotated by an expert hematologist with bounding box and class labels at 1000x on HCM, these annotations are automatically transferred to multiple magnifications across the microscope. To make this transfer efficient, we first compute homography between the images using SURF feature point matching along with RANSAC. Figure 4 summarizes the annotation transfer process. Specifically, on HCM, we transfer the bounding box annotations from 1000x to 400x and from 400x to 100x. Annotations are transferred from 100x image of HCM to 100x image of LCM. Finally, on LCM, the annotations are transferred from 100x to 400x and from 400x to 1000x. Note that due to the difference in resolutions, image quality, and FOVs, we often get poor key points matching or insufficient features points to be matched to compute homography. Therefore, after automatic annotation transfer, we manually verified each image to correct shifted bounding boxes and put missing bounding boxes.

5. Evaluation Methods

We provide an initial benchmark by evaluating two types of methods on our dataset. To compute the accuracy of malarial cells detection as well as malarial-cell stage clas-

Malaria Dataset	Across Micros.	Multi Magn.	Malarial cells	BBX	No. of Images
BBBC041 [1]	No	No	2452	Yes	1364
Malaria655 [38]	No	No	557	No	655
MPIDB [25]	No	No	840	No	229
IML [5]	No	No	529	Yes	345
Our dataset	Yes	Yes	20,331*	Yes	7542*

Table 1. Comparison of malaria datasets. Medical expert traversed multiple slides identifying 1257 malarial regions using HCM at 1000x. Images corresponding to these locations were collected at all the magnifications in both HCM and LCM microscopes resulting in the dataset of size 7542 ($= 2 \times 3 \times 1257$) and total number of malarial cells 3×3624 (HCM) + 3×3153 (LCM) = 20331.

sification, we apply recent object detectors. Our experimental results (Table 2) show that detector trained on HCM performs poorly on LCM even on the same magnification. Therefore, we evaluate a few recent domain adaptation approaches to improve detection accuracy across the microscopes. Furthermore, we propose to use problem-specific constraints to improve cross microscope detection accuracy. Our key findings are that these problem-specific constraints can be plugged-in into any domain adaptation approach and can help improve cross-microscopes accuracy.

5.1. Object Detectors

To generate benchmark results on our dataset, we perform experiments using recent popular object detectors including Faster R-CNN [36], YOLO-v3 [35], RetinaNet [22], and FCOS [40]. YOLO [35], is a single-stage detector that divides the image into grids to achieve a fast inference rate. RetinaNet [22] employs novel focal loss to cater to the class imbalance. Instead of relying on a set of pre-defined anchor boxes to compute proposals, FCOS [40] uses the center point of objects to define whether a location is positive and regresses the four distances from the center to the object boundary. YOLO, RetinaNet, and FCOS are single-stage whereas Faster R-CNN is a two-stage detector.

5.2. Domain adaptation Methods

Our benchmark object detectors report reasonable results over the HCM (Table 2), however, their accuracy drops when tested on LCM. Due to several factors including the lens differences, images taken from the LCM exhibit considerable domain shift from the ones taken from HCM. In Figure 1, we present some of the sample images representing distribution shift. This shift is exhibited not only in terms of change in color distribution but also that LCM images appeared to be more blurred and less clear, with the fine-grain texture information reduced when compared to the HCM images. We explore different domain adaptation strategies to overcome domain shift.

Since Faster-RCNN has the highest mAP in comparison

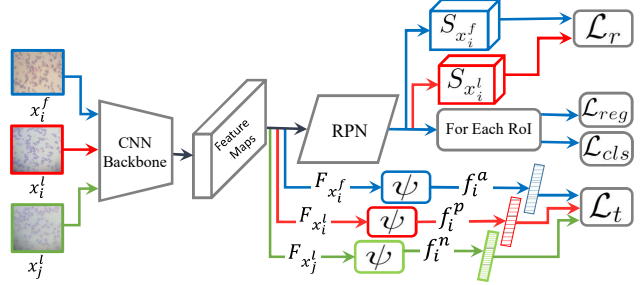


Figure 5. Partially supervised domain adaptation method. For details, see Section 5.2

to other detectors (Table 2), we evaluate the object detector domain adaptation strategies [9, 37, 44] that are based on Faster RCNN on our dataset. Xu et al., [44] propose Graph-induced prototype alignment for the domain adaptation and report results on cityscapes to foggy-cityscapes datasets, a case similar to ours. Chen et al., [9] controls the image and instance-level shifts to avoid domain discrepancy at an optimal level whereas [37] suggests strongly aligning the local features and weakly aligning the global features to retain the unique features of each domain that may help in better inferences. However, all these methods are unable to tackle the domain shift in our case, hence we propose a domain adaptation method specific to our dataset.

Partially Supervised Domain adaptation Approach: Let $D_h = \{(x_i^h, y_i^h)\}_{i=1}^{N_h}$ be the set of images obtained from the HCM along with their annotations, and $D_l = \{(x_i^l)\}_{i=1}^{N_l}$ be the images obtained from LCM. We assume that both D_h (source domain) and D_l (target domain) have same magnification and $N_h = N_l$. Note that we do not have the annotations for the D_l , however, we do have the correspondence across the datasets. That is $x_i^h \in D_h$ and $x_i^l \in D_l$ are obtained by centering microscopes to the same region of the slide. However, due to the decrease in FOV in the LCM, regions captured in both microscopes might not be exactly the same. In fact, x_i^l might have slightly less number of cells than the x_i^h . We use this correspondence information to come up with the *partially-supervised domain adaptation* strategy.

First, using the CycleGAN [45] $G_{h \rightarrow l}$, we translate the images $x_i^h \in D_h$ to D_l . Thus creating an intermediate fake-LCM domain $D_f = \{(x_i^f = G_{h \rightarrow l}(x_i^h), y_i^f = y_i^h)\}_{i=1}^{N_h}$. Different strategies are used to perform domain alignment at different stages of Faster-RCNN. Since we have access to the ground-truth for the intermediate fake-LCM domain D_f , we can use the standard loss functions (\mathcal{L}_{reg} and \mathcal{L}_{cls}) to bridge the gap between the D_l and D_h . We introduce ranking loss and triplet loss to align the feature extraction and the region proposal network (RPN) across domains.

Let F_x be the feature volume being input to the RPN and $S_x \in \mathbb{R}^{(H \times W \times K)}$ be the objectness score output by

Training Magnification	FCOS			RetinaNet			YOLO			Faster R-CNN		
	Test Magnification											
	1000x	400x	100x	1000x	400x	100x	1000x	400x	100x	1000x	400x	100x
1000x	36.8	13.5	0.0	43.1	29.7	0.0	62.8	36.7	0.0	66.8	31.3	0.0
400x	31.4	29.1	1.9	32.9	34.0	1.8	55.2	56.6	4.5	56.9	61.1	1.4
100x	9.4	14.8	8.9	10.2	15.4	16.3	10.5	3.9	20.1	25.4	31.9	31.5

Table 2. Experimental results of models trained and tested on different magnifications on HCM. The evaluation metric is mAP.

the RPN for K anchors when we input image x . Since, $x_i^f = G_{h \rightarrow l}(x_i^h)$ and x_i^l are corresponding images across the domain, RPN should have a similar response for both. In order to minimize the difference between objectness score of the proposed region from the RPN for fake-LCM and its corresponding LCM image, we use the following ranking loss:

$$\mathcal{L}_r(x_i^f, x_i^l) = \max(0, \text{avg}(S_{x_i^f}) - \text{avg}(S_{x_i^l}) - \beta), \quad (1)$$

where β is the margin. Similarly based on the correspondence across domains, we introduce the *triplet loss* to maximize the similarity between x_i^f and x_i^l . Let x_i^f , x_i^l and x_j^l ($i \neq j$) be anchor, positive and negative samples for the triplet loss. Let $f_i^a = \psi(F_{x_i^f})$, $f_i^p = \psi(F_{x_i^l})$ and $f_i^n = \psi(F_{x_j^l})$, where ψ is global average pooling function, the triplet loss is defined as:

$$\mathcal{L}_t = \max(0, \|f_i^p - f_i^a\|_2^2 - \|f_i^a - f_i^n\|_2^2 + \alpha), \quad (2)$$

where α is the margin. The source domain model is adapted using the standard Faster-RCNN loss-function for the $x_i^f \in D_f$ (images from the intermediate domain). For the images from the D_l , the \mathcal{L}_t and \mathcal{L}_r help in domain alignment.

6. Experimental Results

6.1. Implementation details

All the object detectors are trained using *mmdetection* library [8]. ResNet-50 [16] (pre-trained on Imagenet) with feature pyramid network is used as backbone except for the YOLO [35] which uses the Darknet53 [35]. We use stochastic gradient descent with a base learning rate of 0.01 and momentum of 0.9 for all detectors except for the FCOS [40] where we set the base learning rate to 0.001. All models are trained for 30 epochs except 100 epochs for the YOLO [35] and the learning rate is divided by 10 after the eighth and twelfth epoch. Faster-RCNN based methods domain adaptation [9, 37, 44] are evaluated using the author’s provided codes. β and α are chosen 0 and 1, respectively.

Dataset splits: Our train set contains 66.5% of the total images and also around 66% of each malaria class; our test set contains 30% of the total images and from 27% to 30% of each malaria class and finally, our validation set contains 3.5% of the total images and from 3.5% to 5.5% of each

malarial class. We have same division across all the magnifications (i.e., 100x, 400x and 1000x) and across the microscope (HCM and LCM), with corresponding images being in the same split across magnifications and microscopes.

6.2. Experimental results of Object Detectors

Table 2 shows results of above mentioned detectors when trained separately for all three magnification levels on HCM. Faster R-CNN exhibits overall better mAP (mean average precision). In almost all the experiments, the best results are obtained when both training and testing are done on the same magnification level. For all detectors, the detection accuracy is quite low at 100x even when the model is trained on 100x. We believe that this is because even on HCM, the internal structure of cells (which is important for malarial cell detection and stage classification) is not clearly visible at 100x (Figure 1). Due to low accuracy at 100x, we decided to focus only on 1000x and 400x for later experiments. Where models trained on 400x magnification give reasonable results on the 1000x (even comparable to models trained on 1000x), vice versa is not true. This could be attributed to the fact that supervised deep learning models are infamous for learning shortcuts and over-fitting the domain. Whereas noise and less fine-grained texture information in 400x images allow the model to be more robust. Since the FOV at 400x covers around 20 FOVs of 1000x, we can scan the blood smear slide and perform malarial cell detection about 20 times faster at 400x as compared to scanning at 1000x.

6.3. Experimental results of Domain adaptations

The experimental results of adapting from one magnification in HCM (source) to a similar magnification in the

Methods	Source = HCM, Target = LCM	
	1000x → 1000x	400x → 400x
Xu et al. [44]	15.5	21.6
Saito et al. [37]	24.8	21.4
Chen et al. [9]	17.6	21.5
Source only	17.1	26.7
Fine Tuning on fake-LQM	33.3	31.8
Ranking loss	35.7	32.4
Triplet loss	37.2	32.2
Ranking+Triplet loss	37.5	33.8

Table 3. Results (mAP %) of domain adaptation methods.




Figure 6. Each row represents different malarial regions of LCM 1000x. \square , \square , \square represent gametocyte, ring, and trophozoite respectively. (a) ground truths, (b) results of source only model, (c) results for HCM trained model, fine-tuned on fake LCM with ranking loss. (d) shows results when model trained on HCM, fine-tuned on fake LCM with triplet loss, and (e) shows results with both Ranking & Triplet loss.

LCM (target) are detailed in Table 3. The method proposed in [37], outperforms the other two off-the-shelf algorithms. One reason could be that it has explicit local alignment and malaria cell localization is less dependent upon the context. However, results are still lacking.

The experimental results with *ranking loss* and *triplet loss* show a significant improvement as compared to other methods. Triplet loss works better on 1000x as compared to ranking loss and ranking works better at 400x as compared to triplet loss. Both losses have a complementary effect as their combination improves the overall results. For comparison, we present results of FasterRCNN trained on HCM (source only) and also model fine-tuned on the fake-LCM (since we have ground truth for that). Note that we do not have ground truth for the training part of LCM. We evaluated our method for the across magnifications adaptation. For For FasterRCNN, HCM (1000x) to LCM (400x) with our approach mAP improves from 1.4 to 13.2. Adaptation was not successful for HCM (400x) to LCM (100x).

6.4. Qualitative Analysis

Some of qualitative examples when model is trained on HCM (1000x) and tested on LCM (1000x) with and without domain adaptation are shown in Figure 6. The first row shows that source only model in (b) successfully detects trophozoite but is unable to detect the ring. Equipped with triplet loss, our model is able to predict the missing ring successfully. The second row shows that although ranking and triplet loss give a false positive at different locations, their combination results in false-positive removal, on the

other hand, the source-only model does not provide any detection. Finally, in the third row, we showed a failure case. As compared to the source model, our model is able to detect several rings, however, it misses the ring at the bottom. Also, neutrophil has been miss-classified into trophozoite due to its appearance similarities. We consider these results quite encouraging since adapting from the image in the clear domain to the ones with less fine-grain information is a challenging task.

7. Conclusion

We have attempted making malaria detection low-cost and efficient. To this end, we have collected a large-scale multi-magnification, multi-microscope malaria dataset. Furthermore, we benchmark different popular object detectors and domain adaptation methods on our dataset. We believe that this dataset will pave the way for future research in developing algorithms that can work with low-cost microscopes. The malaria-cycle stage contains inherent class imbalance because most blood samples are usually gathered at the middle of the malaria life cycle when patients show symptoms. Our future work will be to focus on designing detectors that handle class imbalance.

Acknowledgement: We thank Qazi Ammar Arshad and Muhammad Affan for their help and contribution in the project. The project is partially supported by an unrestricted gift award from Facebook, USA.

References

- [1] Bbbc041, malaria dataset. <https://bbbc.broadinstitute.org/BBBC041>. 3, 6
- [2] Centers for disease control and prevention. https://www.cdc.gov/malaria/diagnosis_treatment/diagnosis.html. 1
- [3] Part of the endoscopic vision challenge. <https://www.synapse.org/#!Synapse:syn22083820/wiki/603200>. 3
- [4] Timo Ahonen, Abdenour Hadid, and Matti Pietikäinen. Face recognition with local binary patterns. In Tomás Pajdla and Jiří Matas, editors, *Computer Vision - ECCV 2004*, 2004. 3
- [5] Qazi Ammar Arshad, Mohsen Ali, Saeed-ul Hassan, Chen Chen, Ayisha Imran, Ghulam Rasul, and Waqas Sultani. A dataset and benchmark for malaria life-cycle classification in thin blood smear images. *arXiv preprint arXiv:2102.08708*, 2021. 3, 6
- [6] Sania Ashraf, Angie Kao, Cecilia Hugo, Eva M Christophel, Bayo Fatunmbi, Jennifer Luchavez, Ken Lilley, and David Bell. Developing standards for malaria microscopy: external competency assessment for malaria microscopists in the asia-pacific. *Malaria journal*, 11(1):352, 2012. 1
- [7] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer Vision and Image Understanding*, pages 346–359, 2008. 5
- [8] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 7
- [9] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 6, 7
- [10] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20:273–297, 1995. 3
- [11] Ishan R Dave. Image analysis for malaria parasite detection from microscopic images of thick blood smear. In *2017 International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET)*, pages 1303–1307. IEEE, 2017. 3
- [12] Cecilia Di Ruberto, Andrew Dempster, Shahid Khan, and Bill Jarra. Analysis of infected blood cell images using morphological operators. *Image and vision computing*, 20(2):133–146, 2002. 3
- [13] Gloria Díaz, Fabio A González, and Eduardo Romero. A semi-automatic method for quantification and classification of erythrocytes infected with malaria parasites in microscopic images. *Journal of Biomedical Informatics*, 42(2):296–307, 2009. 3
- [14] Tehreem Fatima and Muhammad Shahid Farid. Automatic detection of plasmodium parasites from microscopic blood images. *Journal of Parasitic Diseases*, 44(1):69–78, 2020. 2, 3
- [15] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24:381–395, 1981. 5
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 7
- [17] Jane Hung and Anne Carpenter. Applying faster r-cnn for object detection on malaria images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 1, 2, 3
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012. 3
- [19] Saravana Kumar Kumarasamy, SH Ong, and Kevin SW Tan. Robust contour reconstruction of red blood cells and parasites in the automated identification of the stages of malarial infection. *Machine Vision and Applications*, 22(3):461–469, 2011. 3
- [20] Jordi Landier, Daniel M Parker, Aung Myint Thu, Verena I Carrara, Khin Maung Lwin, Craig A Bonnington, Sasithon Pukrittayakamee, Gilles Delmas, and François H Nosten. The role of early detection and treatment in malaria elimination. *Malaria journal*, 15(1):1–8, 2016. 1
- [21] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex C Kot. Domain generalization for medical imaging classification with linear-dependency regularization. *arXiv preprint arXiv:2009.12829*, 2020. 3
- [22] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar. Focal loss for dense object detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE Computer Society, 2017. 2, 6
- [23] Nina Linder, Riku Turkki, Margarita Walliander, Andreas Mårtensson, Vinod Diwan, Esa Rahtu, Matti Pietikäinen, Mikael Lundin, and Johan Lundin. A malaria diagnostic tool based on computer vision screening and visualization of plasmodium falciparum candidate areas in digitized blood smears. *PLoS One*, 9(8), 2014. 2, 3
- [24] Yun Liu, Yu-Huan Wu, Yunfeng Ban, Huifang Wang, and Ming-Ming Cheng. Rethinking computer-aided tuberculosis diagnosis. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2643–2652, 2020. 3
- [25] Andrea Loddo, Cecilia Di Ruberto, Michel Kocher, and Guy Prod'Hom. Mp-idb: The malaria parasite image database for image processing and analysis. In *Sipaim-Miccai Biomedical Workshop*, pages 57–65. Springer, 2018. 3, 6
- [26] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, pages 91–110, 2004. 3
- [27] Blaine Mathison and Bobbi Pritt. Update on malaria diagnostics and test utilization. *Journal of Clinical Microbiology*, 55, 04 2017. 1

- [28] Courosh Mehanian, Mayoore Jaiswal, Charles Delahunt, Clay Thompson, Matt Horning, Liming Hu, Travis Ostbye, Shawn McGuire, Martha Mehanian, Cary Champlin, et al. Computer-automated malaria diagnosis and quantitation using convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 116–125, 2017. 1, 2, 3
- [29] Angel Molina, Santiago Alf3rez, Laura Bold3, Andrea Acevedo, Jos3 Rodellar, and Anna Merino. Sequential classification system for recognition of malaria infection using peripheral blood cell images. *Journal of Clinical Pathology*, 2020. 2, 3
- [30] World Health Organization. *World malaria report 2020: 20 years of global progress and challenges*. World Health Organization, 2020. 1
- [31] World Health Organization et al. World malaria report 2019. 2019. 3
- [32] Nick A. Phillips, Pranav Rajpurkar, Mark Sabini, Rayan Krishnan, Sharon Zhou, Anuj Pareek, Nguyet Minh Phu, Chris Wang, Mudit Jain, Nguyen Duong Du, Steven QH Truong, Andrew Y. Ng, and Matthew P. Lungren. Chex-photo: 10,000+ photos and transformations of chest x-rays for benchmarking deep learning robustness. In *Proceedings of the Machine Learning for Health NeurIPS Workshop*, volume 136. PMLR, 2020. 2
- [33] Sivaramakrishnan Rajaraman, Sameer K Antani, Mahdiah Poostchi, Kamolrat Silamut, Md A Hossain, Richard J Maude, Stefan Jaeger, and George R Thoma. Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection in thin blood smear images. *PeerJ*, 6:e4568, 2018. 2, 3
- [34] KNR Moh-ana Rao, AG Dempster, Bill Jarra, and Shahid Khan. Automatic scanning of malaria infected blood slide images using mathematical morphology. 2002. 3
- [35] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. 2018. 2, 6, 7
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, 2015. 2, 6
- [37] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019. 6, 7, 8
- [38] F Boray Tek, Andrew G Dempster, and Izzet Kale. Images of thin blood smears with bounding boxes around malaria parasites (malaria-655). 3, 6
- [39] F Boray Tek, Andrew G Dempster, and Izzet Kale. Malaria parasite detection in peripheral blood images. *BMVA*, 2006. 2, 3
- [40] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *Proc. Int. Conf. Computer Vision (ICCV)*, 2019. 2, 6, 7
- [41] Muhammad Umer, Saima Sadiq, Muhammad Ahmad, Saleem Ullah, Gyu Sang Choi, and Arif Mehmood. A novel stacked cnn for malarial parasite detection in thin blood smear images. *IEEE Access*, pages 93782–93792, 2020. 2, 3
- [42] A Vijayalakshmi et al. Deep learning approach to detect malaria from microscopic images. *Multimedia Tools and Applications*, pages 1–21, 2019. 2, 3
- [43] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *CVPR*, 2017. 3
- [44] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12355–12364, 2020. 6, 7
- [45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017. 6
- [46] Aneeq Zia, Kiran Bhattacharyya, Xi Liu, Ziheng Wang, Satoshi Kondo, Emanuele Colleoni, Beatrice van Amsterdam, Razeen Hussain, Raabid Hussain, Lena Maier-Hein, Danail Stoyanov, Stefanie Speidel, and Anthony Jacr. Surgical visual domain adaptation: Results from the miccai 2020 surgvisdom challenge, 2021. 3