

Adaptive Stain Normalization for Cross-Domain Medical Histology

Tianyue Xu^{1*}, Yanlin Wu^{1*}, Abhai K. Tripathi², Matthew M. Ippolito^{2,3†}, and Benjamin D. Haeffele^{4†}✉

¹ Dept. of Biomedical Engineering, Johns Hopkins University, Baltimore, MD, USA
{txu48,ywu289}@jh.edu

² Johns Hopkins Malaria Research Institute, W. Harry Feinstone Department of Molecular Microbiology and Immunology, Bloomberg School of Public Health, Baltimore, MD, USA atripat2@jhu.edu

³ Dept. of Medicine, Johns Hopkins University School of Medicine, Baltimore, MD, USA mippolito@jhu.edu

⁴ Dept. of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA haeffele@upenn.edu

Abstract. Deep learning advances have revolutionized automated digital pathology analysis. However, differences in staining protocols and imaging conditions can introduce significant color variability. In deep learning, such color inconsistency often reduces performance when deploying models on data acquired under different conditions from the training data, a challenge known as domain shift. Many existing methods attempt to address this problem via color normalization but suffer from several notable drawbacks such as introducing artifacts or requiring careful choice of a template image for stain mapping. To address these limitations, we propose a trainable color normalization model that can be integrated with any backbone network for downstream tasks such as object detection and classification. Based on the physics of the imaging process per the Beer-Lambert law, our model architecture is derived via algorithmic unrolling of a nonnegative matrix factorization (NMF) model to extract stain-invariant structural information from the original pathology images, which serves as input for further processing. Experimentally, we evaluate the method on publicly available pathology datasets and an internally curated collection of malaria blood smears for cross-domain object detection and classification, where our method outperforms many state-of-the-art stain normalization methods. Our code is available at <https://github.com/xutian Yue/BeerLaNet>.

Keywords: Domain Adaptation · Stain Normalization · Pathology

1 Introduction

Pathological examination plays a fundamental role in disease diagnosis, providing critical insights into tissue morphology, cellular abnormalities, and disease

* Equal contribution, co-first authors.

† Equal contribution, co-senior authors.

progression. However, manual assessment of digital pathology images is labor-intensive, time-consuming, and subject to inter- and intra-observer variability [20,19]. Deep learning-driven analysis offers a promising alternative but struggles with a major challenge for model generalizability: color inconsistency in digital pathology. Recorded images can have significant differences in stain appearance even with the same staining protocol, caused by (i) differences in chemical reactions and exposure times of dyes, (ii) variability in sample preparation, and (iii) variable imaging conditions across different scanning hardware [15]. While experienced pathologists can interpret slides despite stain variability, deep learning models are prone to performance degradation [3], leading to poor generalization across different staining protocols and acquisition settings. This challenge is closely related to domain adaptation, a key discipline in deep learning that aims to improve generalization across different data distributions. To address this challenge, a strategy called stain color normalization is widely used in domain adaptation methods specific to medical pathology.

Traditional stain color normalization approaches often attempt to match statistics of the color distribution of test images to those of images in the training domain [13]. Yet, the effectiveness of these methods strongly relies on selecting appropriate representative templates from the training domain to match, which can require prior knowledge of the domain. Moreover, these methods can be generic color transformations that do not account for the underlying physics of histological staining. Other methods leverage algorithms to decompose the color space of a reference image into different dye staining components [11,18]. While these methods account for the underlying physics of histological staining and image acquisition, they still rely on a predefined template which the method attempts to match and, in some cases, require additional prior knowledge, such as the absorption spectrum matrix of specific dyes or the number of distinct color components within the images [18]. More recently, deep learning-based approaches such as generative adversarial networks (GANs) [16] have been explored for stain style transfer. However, despite their success in aligning color distributions and overall appearance, it has been noted that they can often introduce synthetic artifacts or ‘hallucinate’ cellular structures, posing risks in medical diagnosis [12]. Other deep learning methods have also proposed to use attention mechanisms to identify relevant color transformations from a variety of standard normalized color spaces, but these methods are also largely generic and do not necessarily account for the relevant image formation process [10].

To overcome these limitations, we propose the Beer-Lambert Net (BeerLaNet). Our key contributions include: (1) **Adaptive Stain Disentanglement**: Unlike previous methods, which focus largely on hematoxylin and eosin-stained images, BeerLaNet extends to arbitrary staining protocols, learning stain-invariant representations of images without requiring any prior knowledge of the staining protocol. (2) **Trainable and Physics-Informed**: Built on non-negative matrix factorization (NMF) and algorithmic unrolling, BeerLaNet enables a data-driven, end-to-end stain decomposition process based on the imaging physics. (3) **Flexible Integration**: Designed as a plug-and-play module,

BeerLaNet can be combined with arbitrary backbone networks for downstream tasks, such as object detection or image classification.

2 Technical Background

Each stain can be considered a unique element that interacts with light in a specific way. In an RGB image, every stain has distinct absorption characteristics across the three spectral channels, influencing the observed color intensity. To mathematically describe how stains modify light intensity in histological images, we rely on the Beer-Lambert law [14], which models how incident light is attenuated by stains based on their concentration and absorption properties. Specifically, the absorption of stains by histological structures attenuates the incident light on stained biological samples in a specific color spectrum. Given an image $\bar{\mathbf{X}} = [\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_p] \in \mathbb{R}^{c \times p}$, where $\bar{\mathbf{x}}_i \in \mathbb{R}^c$ contains the c color intensities for the i^{th} pixel (c is typically 3 for RGB images) and p is the number of pixels in the image, then when light $\bar{\mathbf{x}}_0 \in \mathbb{R}^c$ illuminates the image, the color intensity at each pixel is given by the Beer-Lambert law;

$$\bar{\mathbf{X}} = (\bar{\mathbf{x}}_0 \mathbf{1}^\top) \odot e^{-\mathbf{S}\mathbf{D}^\top} \quad (1)$$

where $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_r] \in \mathbb{R}^{c \times r}$ is the color appearance matrix whose i^{th} column contains the color spectra of the i^{th} colored component (e.g., the color of a stain or a naturally colored material in the specimen, such as hemoglobin in red blood cells), r is the number of colored components, $\mathbf{D} = [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_r] \in \mathbb{R}^{p \times r}$ is the optical density matrix whose i^{th} column contains the optical density of the i^{th} colored component at each pixel, $\mathbf{1}$ denotes a vector of all ones, and \odot denotes the element-wise matrix product.

Matrix factorization for stain separation. Based on the Beer-Lambert law, prior approaches have proposed models for stain normalization based on matrix factorization models. In particular, taking an element-wise logarithm on both sides of (1) gives

$$\ln(\bar{\mathbf{X}}) = \ln(\bar{\mathbf{x}}_0) \mathbf{1}^\top - \mathbf{S}\mathbf{D}^\top \iff \ln(\bar{\mathbf{x}}_0) \mathbf{1}^\top - \ln(\bar{\mathbf{X}}) = \mathbf{S}\mathbf{D}^\top \quad (2)$$

so by estimating a low-rank approximation $\mathbf{S}\mathbf{D}^\top$ (recall the number of columns r in (\mathbf{S}, \mathbf{D}) corresponds to the number of stains/colored materials in the specimen) one could hope to have a reliable decomposition of the image into color and spatial optical density components. For example, the authors of [11] propose a model which estimates the image background \mathbf{x}_0 via a heuristic and the estimates \mathbf{S} and \mathbf{D} via principal component analysis (PCA). The estimates of \mathbf{S} and \mathbf{D} are then normalized via further heuristics before reforming a normalized version of the original image using the normalized $(\mathbf{x}_0, \mathbf{S}, \mathbf{D})$ parameters in (1).

While the above model has the benefit of being motivated by the image formation physics, the use of PCA for estimating \mathbf{S} and \mathbf{D} has clear deficiencies, as it requires the columns of the matrices to be orthogonal, which is clearly not




Fig. 1. Overview of our proposed BeerLaNet method.

realistic for the problem. As a result, the authors of [18] propose a sparse NMF model of the form:

$$\min_{\mathbf{x}_0, \mathbf{S}, \mathbf{D}} \frac{1}{2} \|\mathbf{x}_0 \mathbf{1}^\top - \mathbf{X} - \mathbf{S} \mathbf{D}^\top\|_F^2 + \lambda \sum_{i=1}^r \|\mathbf{d}_i\|_2 \quad \text{s.t.} \quad \|\mathbf{s}_i\|_2 = 1, \forall i, \quad \mathbf{S}, \mathbf{D} \geq \mathbf{0}. \quad (3)$$

The model in (3) is motivated by the fact that the color spectra and optical density matrices are known to be nonnegative and additionally adds sparse regularization on the columns of \mathbf{D} via the ℓ_1 norm based on the common assumption that the relevant histological features of interest are sparse in space. After estimating the parameters of the model in (3) the authors then produce a stain-normalized image by applying heuristics to normalize the columns of \mathbf{D} . They then use (1) with the normalized estimate of \mathbf{D} but discard the estimated \mathbf{S} matrix and instead use a fixed color spectra matrix which has been estimated from a template image whose color the method hopes to match. However, this presents challenges in practice as one must choose a suitable template image and corresponding color spectra matrix and the columns of the estimated \mathbf{S} and \mathbf{D} matrices could be permuted from the corresponding color spectra matrix, resulting in significant color distortions. Further, the method also requires one to carefully select both the sparse regularization strength, λ , and the rank of the matrix factorization r . In the following section we present our proposed BeerLaNet method and describe how it alleviates many of these problems.

3 Method

To derive our method, we begin with similar assumptions to the sparse NMF model described above, but we make several key modifications to mitigate the issues described above and to allow for our method to be incorporated in an end-to-end fashion with any desired deep network model for further downstream tasks. Specifically, we note that the NMF model in (3) is non-convex and requires one to solve a non-convex optimization problem due to the $\mathbf{S} \mathbf{D}^\top$ matrix product. Moreover, the number of columns in \mathbf{S} and \mathbf{D} needs to be specified *a priori*, which can be challenging in applications where the number of colored components

is unknown such as when the specimen itself contains colored components in addition to the applied stain (e.g., hemoglobin in red blood cells).

We address these issues by first leveraging literature on optimizing structured matrix factorization models [6,7] and modify (3) into the following formulation:

$$\min_{\mathbf{x}_0, \mathbf{S}, \mathbf{D}} \frac{1}{2} \|\mathbf{x}_0 \mathbf{1}^\top - \mathbf{X} - \mathbf{S} \mathbf{D}^\top\|_F^2 + \lambda \sum_{i=1}^r \|\mathbf{s}_i\|_2 (\gamma \|\mathbf{d}_i\|_1 + \|\mathbf{d}_i\|_2) \quad \text{s.t.} \quad \mathbf{S}, \mathbf{D} \geq \mathbf{0}. \quad (4)$$

Compared to the model in (3) the key difference is that we have incorporated an additional ℓ_2 regularization on the columns of \mathbf{S} and \mathbf{D} which is known to promote low-rank solutions due to connections with the variational form of the nuclear norm [6,7]. This allows us to initialize the number of colored components (r) larger than the expected number of components and adapt the rank of the solution to the data. We then derive our method from an algorithmic unrolling of (4) which can be integrated in an end-to-end manner with any deep network architecture for further downstream tasks as we describe below.

Unrolled Network Architecture. To derive our proposed method, we design an architecture via an algorithmic unrolling of (4) via alternating proximal gradient descent [4]. Specifically, we make the regularization parameters (γ, λ) and the initialization for \mathbf{S} (\mathbf{S}_{init}) learnable parameters (with \mathbf{x}_0 and \mathbf{D} initialized as all zeros), then our algorithm first updates \mathbf{x}_0 in closed form, followed by updates to \mathbf{D} and \mathbf{S} by taking a gradient descent step on the Frobenius norm term in (4), followed by solving the proximal operator of the regularization term and nonnegative constraints. This is then repeated for K unrolled layers, after which we reshape the recovered \mathbf{D} matrix into an r -channel image passed through a 1×1 convolution layer to map the image back to a 3-channel image to allow for passing the image to a backbone network (e.g., YOLO, ResNet) for downstream tasks, with the learnable parameters $(\mathbf{S}_{\text{init}}, \gamma, \lambda)$ being updated via backpropagation via supervision on the downstream task and with non-negative constraints. An overview of our approach is depicted in Fig. 1, and the full details of our method are given in Algorithm 1. Note that $(\cdot)_+$ denotes the element-wise ReLU operation and the proximal operator of the regularization/constraints on \mathbf{S} and \mathbf{D} can be shown to be computable by sequentially computing the proximal operator of the non-negativity constraints and ℓ_1 norm (for the update for \mathbf{D}) followed by the proximal operator of the ℓ_2 norm [6,7].

4 Experimental Performance

We evaluate our method on a variety of modalities for two diagnostic tasks: object detection and image classification. Here, we detail the modalities and their respective role in our study, and example images showing the variance in domain between training and testing data are displayed in Fig 2.

Malaria Parasite Detection and Classification. We used images of malaria blood smears for both detection and classification tasks. For detection, we used a public malaria parasite detection dataset by the authors of [5], which comprises

Algorithm 1 Beer-Lambert Net (BeerLaNet)

Input: Image $\mathbf{X} \in \mathbb{R}^{c \times p}$. # of unrolled iterations K . # of color components r .
Learnable Parameters: $\gamma \in \mathbb{R}$, $\lambda \in \mathbb{R}$, $\mathbf{S}_{\text{init}} \in \mathbb{R}^{c \times r}$.
Outputs: The optical density map matrix $\mathbf{D} \in \mathbb{R}^{p \times r}$.
Initialize: $\mathbf{S} = \mathbf{S}_{\text{init}}$, $\mathbf{D} = \mathbf{0}$, $\mathbf{x}_0 = \mathbf{0}$.
for $k = 1, \dots, K$ **do**
 $\mathbf{x}_0 = \frac{1}{p}(\mathbf{X} + \mathbf{S}\mathbf{D}^\top)\mathbf{1}$ % Closed form update for \mathbf{x}_0 .
 $\tau = \frac{1}{\|\mathbf{S}\|_F^2}$ % Compute step size for \mathbf{D} .
 $\mathbf{D} = \mathbf{D} - \tau[\mathbf{D}\mathbf{S}^\top\mathbf{S} + \mathbf{X}^\top\mathbf{S} - \mathbf{x}_0^\top\mathbf{S}]$ % Gradient step of $\|\cdot\|_F^2$ term w.r.t. \mathbf{D} .
 for $i = 1, \dots, r$ **do**
 $\mathbf{d}_i = (\mathbf{d}_i - \lambda\gamma\tau\|\mathbf{s}_i\|_2)_+$ % Proximal operator of ℓ_1 + non-negative constraints.

 $\mathbf{d}_i = \mathbf{d}_i \left(1 - \frac{\min\{\|\mathbf{d}_i\|_2, \lambda\gamma\tau\|\mathbf{s}_i\|_2\}}{\|\mathbf{d}_i\|_2}\right)$ % Proximal operator of ℓ_2 .
 end for
 $\tau = \frac{1}{\|\mathbf{D}\|_F^2}$ % Compute step size for \mathbf{S} .
 $\mathbf{S} = \mathbf{S} - \tau[\mathbf{S}\mathbf{D}^\top\mathbf{D} + \mathbf{X}\mathbf{D} - \mathbf{x}_0\mathbf{1}^\top\mathbf{D}]$ % Gradient step of $\|\cdot\|_F^2$ term w.r.t. \mathbf{S} .
 for $i = 0, 1, \dots, r$ **do**
 $\mathbf{s}_i = (\mathbf{s}_i)_+ \left(1 - \frac{\min\{\|(\mathbf{s}_i)_+\|_2, \lambda\gamma\tau(\gamma\|\mathbf{d}_i\|_1 + \|\mathbf{d}_i\|_2)\}}{\|(\mathbf{s}_i)_+\|_2}\right)$ % Proximal operator for \mathbf{S} .
 end for
end for

24,720 May Grunwald-Giemsa (MGG)-stained thin blood smear images with annotations across four categories: white blood cells, red blood cells, platelets, and parasites (including *Trypanosoma brucei*, and erythrocytes infected by *Plasmodium falciparum*, *P. ovale*, *P. vivax*, *P. malariase*, and *Babesia divergens*). We additionally collected and manually labeled a test dataset comprising 264 thin blood smear images captured from a Zeiss Axioscan microscope.

For classification, we assembled a diverse dataset from four microscopy platforms, encompassing both manual and high-throughput imaging systems with varying numerical apertures (Hamamatsu NanoZoomer, Zeiss Axioscan, Olympus CX43, Morphle Hemolens). The source materials included Giemsa-stained thin blood smears obtained from multiple origins: a clinical trial, laboratory-prepared slides with spiked parasites, and de-identified slides from a malariology course collection. Classification labels were independently assigned by two authors (AKT, MMI), with any disagreements resolved through joint review by both. The dataset consists of single-cell cropped images infected by *P. falciparum*, with a training set with 2,486 samples, and two test sets with 343 and 261 samples, each from a distinct domain (imaging platform). The classification task was to label detected parasites by lifestage: early ring, middle ring, late ring, trophozoite, schizont, and gametocyte, as shown in Fig. 2(a).

Breast Cancer Classification. The Camelyon17-WILDS dataset [2] is a large-scale histopathology dataset for evaluating domain adaptation in medical image analysis. It consists of 96×96 image patches extracted from lymph node whole slide images, with tumor presence labeled in the central 32×32 region. The

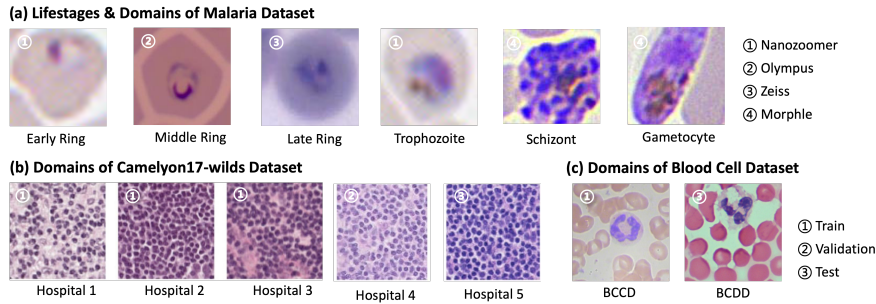


Fig. 2. Example images from our tested datasets.

dataset spans five hospitals, introducing domain shifts in staining and imaging conditions. It contains tumor/non-tumor classifications, with 302,436 training, 34,904 validation, and 85,054 test samples. Here we train on the 3 training datasets and perform inference on the testing and validation datasets.

Whole Blood Cell Detection. We used two additional public datasets for detecting various blood cell types from whole blood smears: the Blood Cell Count and Detection (BCCD) dataset [17] containing 366 blood smear images for training, and the Blood Cell Detection Dataset (BCDD) [1] with 100 blood smear images for testing.

Implementation Details. For BeerLaNet, we set the number of colored component r to 8 and internal iterations to 10. For detection tasks, we implemented our method with a YOLOv8 backbone [9]. The training parameter batch size was set to 8, and the number of epochs was set to 50. The optimizer was automatically chosen by YOLO with an initial learning rate of 0.01. For classification, we trained the model using Adam optimization with a batch size of 128 and a ResNet-18 backbone [8]. The *P. falciparum* malaria parasite dataset was trained for 20 epochs and Camelyon-17 WILDS for 50 epochs, both with a learning rate of $1e-4$. For the malaria classification dataset, we also applied a two-step denoising pipeline due to compression artifacts from the imaging platform: a median filter with kernel size of 11, followed by a Gaussian blur with kernel size of 21 and $\sigma = 1$. All experimental results are reported as averages across 3 random seeds.

Comparison Methods & Evaluation Metrics. We compared our proposed method with 3 classical histology normalization methods—Reinhard [13], Macenko [11], and Vanhadane [18]—and two deep learning-based methods—StainGAN [16] and LStainNorm [10]. For those methods requiring a template, we randomly select one image from the training dataset as the template for normalization. The inference results are evaluated by mAP_{50} and mAP_{50-95} on the detection task. For classification we report classification accuracy (Acc), and for the malaria parasite classification task we also report a relaxed accuracy (RAcc) which considers a classification successful if the manual and predicted labels are

Table 1. Comparison of Stain Normalization Techniques. The best and the second-best results are **boldfaced** or starred (*), respectively. (C17 denotes Camelyon17-WILDS)

Task	Detection (YOLOv8)					Classification (ResNet18)				
Dataset	Malaria		Whole Blood Cells			Malaria	C17 _{test}	C17 _{val}		
Metrics (%)	mAP ₅₀	mAP ₅₀₋₉₅	mAP ₅₀	mAP ₅₀₋₉₅	APU	Acc	RAcc	Acc	Acc	APU
Baseline	91.03	52.00	65.20	36.70	18.10	21.32	45.59	85.21	83.38	31.70
Reinhard	90.17	51.47	79.53	44.03	11.17	29.34	62.30	94.55	91.36*	18.36
Macenko	72.03	39.27	65.43	35.27	29.27	29.59	73.54	95.92	85.77	16.27
Vahadane	92.10*	55.87*	57.57	31.30	20.76	38.81*	81.04*	95.85*	86.43	9.31*
StainGAN	81.67	37.70	89.60	53.97	12.02	31.46	70.94	94.28	90.77	15.12
LStainNorm	91.80	53.67	85.83	50.00	5.25*	21.17	61.82	93.23	92.56	22.72
BeerLaNet	95.07	57.10	86.80*	51.33*	2.00	48.66	90.33	91.36	90.09	1.86

within one lifestage of each other for the early/middle/late ring stages (e.g., if the manual label is ‘early ring’, then predictions of ‘early ring’ or ‘middle ring’ are considered correct) to reflect the fact that the parasite growth is a continuous process and the discretization into 3 stages is somewhat arbitrary. In addition, to compute an average performance across all tasks and metrics, we also introduce the average percent underperformance (APU) which computes the percent difference between the current method and the best-performing method for a particular task+metric (percent underperformance) and then averages this across all task+metric combinations for either detection or classification. For example, for each entry in Table 1, we compute the percent difference with the maximum value in the column and then average the percent differences across the 4 entries in the row for detection and classification.

Results. Our experiment results are presented in Table 1. Here we note that BeerLaNet achieves the best performance for both the malaria parasite detection and classification tasks and the second highest performance on the whole blood cell detection task, while also having excellent consistency across all tasks and datasets. For example, on the Camelyon-17 WILDS classification task, the Macenko method achieves the best performance on one dataset (C17_{test}) but drops 10 percentage points on the other Camelyon-17 WILDS dataset (C17_{val}). Likewise, for virtually all comparison methods, a large drop in performance can be observed for one of the tasks/datasets. In contrast, even when BeerLaNet is not the top method for a particular dataset, it still performs very competitively, which is reflected in the APU metric, where BeerLaNet significantly outperforms the comparison methods and quantifies the consistently strong performance of BeerLaNet. We conjecture this discrepancy in performance across tasks/datasets observed in comparison methods can potentially be attributed to greater domain variations in the malaria dataset compared to Camelyon17-WILDS and whole blood cell datasets (see Fig. 2), datasets with relatively minor color shifts compared to those observed in the malaria dataset. Methods with a generic design that do not incorporate stain-specific characteristics may perform well on small domain shift data (Camelyon17-WILDS, whole blood cells) but fail to generalize to the large color shifts observed in the malaria dataset. This highlights the flex-

ible nature of BeerLaNet to perform robustly in a variety of conditions without prior knowledge of imaging or staining protocols.

5 Discussion and Conclusion

Here we have presented BeerLaNet, an adaptive stain normalization method that enhances cross-domain generalization in digital pathology through physics-based principles and algorithmic unrolling. Our experimental results demonstrate that BeerLaNet consistently outperforms traditional normalization techniques across multiple tasks and datasets. The key advantages of BeerLaNet include end-to-end trainability, template-free operation, physics-informed decomposition, and flexibility with various staining protocols. Potential future work could focus on exploring additional tasks such as segmentation and other histopathological domains.

Acknowledgments. This work was supported by NIH NIAID grant R21AI16936, Bloomberg Philanthropies through the Johns Hopkins Malaria Research Institute, and the Johns Hopkins Institute for Data-Intensive Engineering and Science.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdüssamet: Blood cell detection dataset. <https://www.kaggle.com/datasets/draaslan/blood-cell-detection-dataset> (2020), kaggle
2. Bandi, P., Geessink, O., Manson, Q., Van Dijk, M., Balkenhol, M., Hermesen, M., Bejnordi, B.E., Lee, B., Paeng, K., Zhong, A., et al.: From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging* (2018)
3. Ciompi, F., Geessink, O., Bejnordi, B.E., De Souza, G.S., Baidoshvili, A., Litjens, G., Van Ginneken, B., Nagtegaal, I., Van Der Laak, J.: The importance of stain normalization in colorectal tissue classification with convolutional networks. In: 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017). pp. 160–163. IEEE (2017)
4. Combettes, P.L., Pesquet, J.C.: Proximal splitting methods in signal processing. *Fixed-point algorithms for inverse problems in science and engineering* pp. 185–212 (2011)
5. Guemas, E., Routier, B., Ghelfenstein-Ferreira, T., Cordier, C., Hartuis, S., Marion, B., Bertout, S., Varlet-Marie, E., Costa, D., Pasquier, G.: Automatic patient-level recognition of four plasmodium species on thin blood smear by a real-time detection transformer (rt-detr) object detection algorithm: a proof-of-concept and evaluation. *Microbiology Spectrum* **12**(2), e01440–23 (2024)
6. Haeffele, B., Young, E., Vidal, R.: Structured low-rank matrix factorization: Optimality, algorithm, and applications to image processing. In: *International conference on machine learning*. pp. 2007–2015. PMLR (2014)

7. Haeffele, B.D., Vidal, R.: Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *IEEE transactions on pattern analysis and machine intelligence* **42**(6), 1468–1482 (2019)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
9. Jocher, G., Qiu, J., Chaurasia, A.: Ultralytics YOLO (Jan 2023), <https://github.com/ultralytics/ultralytics>
10. Ke, J., Zhou, Y., Shen, Y., Guo, Y., Liu, N., Han, X., Shen, D.: Learnable color space conversion and fusion for stain normalization in pathology images. *Medical Image Analysis* **101**, 103424 (2025)
11. Macenko, M., Niethammer, M., Marron, J.S., Borland, D., Woosley, J.T., Guan, X., Schmitt, C., Thomas, N.E.: A method for normalizing histology slides for quantitative analysis. In: *2009 IEEE international symposium on biomedical imaging: from nano to macro*. pp. 1107–1110. IEEE (2009)
12. Rahman, A., Rahman, M.S., Mahdy, M.R.C.: 3c-gan: class-consistent cyclegan for malaria domain adaptation model. *Biomedical Physics & Engineering Express* **7**(5), 055002 (2021)
13. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Computer graphics and applications* **21**(5), 34–41 (2001)
14. Ruifrok, A.C., Johnston, D.A., et al.: Quantification of histochemical staining by color deconvolution. *Analytical and quantitative cytology and histology* **23**(4), 291–299 (2001)
15. Salvi, M., Michielli, N., Molinari, F.: Stain color adaptive normalization (scan) algorithm: Separation and standardization of histological stains in digital pathology. *Computer methods and programs in biomedicine* **193**, 105506 (2020)
16. Shaban, M.T., Baur, C., Navab, N., Albarqouni, S.: Staingan: Stain style transfer for digital histological images. In: *2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019)*. pp. 953–956. IEEE (2019)
17. Shenggan Cheng, N.C.: Bccd (blood cell count and detection) dataset (Feb 2019), https://github.com/Shenggan/BCCD_Dataset
18. Vahadane, A., Peng, T., Sethi, A., Albarqouni, S., Wang, L., Baust, M., Steiger, K., Schlitter, A.M., Esposito, I., Navab, N.: Structure-preserving color normalization and sparse stain separation for histological images. *IEEE transactions on medical imaging* **35**(8), 1962–1971 (2016)
19. Veloso, S.G., Lima, M.F., Salles, P.G., Berenstein, C.K., Scaloni, J.D., Bambirra, E.A.: Interobserver agreement of gleason score and modified gleason score in needle biopsy and in surgical specimen of prostate cancer. *International braz j urol* **33**, 639–651 (2007)
20. Yeo, T., Ong, S., Sinniah, R., et al.: Autofocusing for tissue microscopy. *Image and vision computing* **11**(10), 629–639 (1993)