Deep Optical Blood Analysis: COVID-19 Detection as a Case Study in Next Generation Blood Screening

Colin L. Cooke^a, Kanghyun Kim^b, Shiqi Xu^b, Amey Chaware^a, Xing Yao^b, Xi Yang^b, Jadee Neff^c, Patricia Pittman^c, Chad McCall^c, Carolyn Glass^c, Xiaoyin Sara Jiang^c, Roarke Horstmeyer^{a,b}

^a*Electrical and Computer Engineering Department, Duke University*

^b*Biomedical Engineering Department, Duke University*

^c*Pathology Department, Duke University Hospital*

Abstract

A wide variety of diseases are commonly diagnosed via the visual examination of cell morphology within a peripheral blood smear. For certain diseases, such as COVID-19, morphological impact across the multitude of blood cell types is still poorly understood. In this paper, we present a multiple instance learning-based approach to aggregate high-resolution morphological information across many blood cells and cell types to automatically diagnose disease at a per-patient level. We integrated image and diagnostic information from across 236 patients to demonstrate not only that there is a significant link between blood and a patient's COVID-19 infection status, but also that novel machine learning approaches offer a powerful and scalable means to analyze peripheral blood smears. Our results both backup and enhance hematological findings relating blood cell morphology to COVID-19, and offer a high diagnostic efficacy; with a 79% accuracy and a ROC-AUC of 0.90.

1. Introduction

Within hematology, the analysis of blood cell morphology plays a critical role in diagnosing and understanding various diseases.¹ A key tool for blood cell morphology assessment is the light microscope, which is often applied to examine peripheral blood smears (PBS).² In a typical procedure, a physician will visually examine white and red blood cells within a PBS on a glass slide at high microscope magnification (usually 100×). The nature of visual examination at high resolution limits the observable field-of-view (FOV) to contain just few white and red blood cells at a time, making analysis of multiple cells challenging and time consuming. Digital microscopes³ have emerged as an effective alternative to manual analysis. By automating the scanning process and presenting digitized images of PBSs to physicians on a computer, such digital microscopes are quickly becoming the predominate method of PBS analysis.

The digitization of PBS imagery has also led to new opportunities to apply advanced machine learning algorithms, such as deep learning methods, to examine blood data.^{4,5,6,7} However, thus far, most algorithmic methods have relied on preexisting and developed understandings of the morphological features of interest, to either facilitate the design of feature extraction techniques or to specify per-cell labeling criteria. This constraint limits the application of machine learning to automate the decisions that physicians are currently able to make, rather than providing fundamentally new capabilities or insights.

The limitations of current machine learning methods for blood cell morphology analysis were recently highlighted during the COVID-19 pandemic. There is a growing body of evidence that suggests COVID-19 and blood have a complex set of interactions that lead to significant morbidity and mortality,^{8,9} and there is a variety of clinically reported evidence that COVID-19 induces morphological changes to both white and red blood cells.^{10,11,12,13} However, our current limited understanding and agreement regarding such morphological impact has impeded development of effective blood-based diagnostic and prognostic screening tools.

In this work, we argue that a new way of analyzing blood, which we term *Deep Optical Blood Analysis* (DOBA), circumvents the need to pre-define features of interest or label individual cells within particular categories, and instead allows for an entirely data-driven analysis of blood using only patient-level information. DOBA uses deep learning to develop a mapping between images from a patient's PBS and their condition. In a typical digital PBS scan, images of hundreds of white and red blood cells are captured per-patient. It is therefore desirable to examine each image in detail, without requiring labels on the individual images. To accomplish this, we adopted a *Multiple Instance Learning* (MIL)¹⁴ technique to link a patient's COVID-19 diagnosis (obtained with a standard PCR-RT laboratory test) to their blood image data. Specifically, a recent approach¹⁵ paired an MIL attention mechanism with a convolutional neural network to simultaneously learn how to extract information from individual images and aggregate information across multiple images. We extended this work to form a novel hybrid MIL network based upon *model ensembling*,¹⁶ and applied our new algorithm to produce accurate final per-patient screening results directly from blood image data.

We chose COVID-19 as a case-study in developing DOBA due not only to the significance of the disease, but also the growing medical consensus regarding its connections to blood.¹⁷ Despite this consensus, there is no convergence on a particular expression of COVID-19 in blood, with responses ranging from Thrombocytopenia,¹⁸ to COVID-19 induced blood clots,¹⁹ to morphological abnormalities.^{13,12} Therefore, despite sufficient evidence to attempt to use blood cell morphology to detect COVID-19, there is no clear starting point in examining individual cells using standard supervised machine learning approaches (i.e. labelling each cell individually).

Our new method not only enables diagnosis of COVID-19 without requiring such labeling, but also sheds light on how this new disease affects blood, by automatically producing a statistical summary of which specific cells and cell types are more or less important for the COVID-19 diagnostic task. Further, by applying specific *perturbations* to our image datasets, we have also developed a procedure to highlight which spatial features of the acquired image data were more or less important to enable robust screening. Apart from enhancing our understanding

of the disease, these features also offer a window into algorithm operation to improve the explainability and reliability of our approach. We are hopeful that our new learning-based data aggregation strategy can serve as a starting point for future algorithmic strategies to elucidate the hematological impact of COVID-19 and other blood-related diseases.

2. Results and Discussion

2.1. Study Design

We investigated the diagnostic potential of PBS images for COVID-19 infection through a partnership with the Duke University Medical Center. Over a five month period (April 2020 - August 2020) we collected digital PBS image data from 236 patients, 53% of whom tested positive for COVID-19 by a separately administered PCR test. No other patient information was collected for this cohort. In addition, we collected PBS image data from 40 additional patients admitted to the medical intensive care unit who presented with acute respiratory illness, but were confirmed to COVID-19 negative using the same PCR testing method. We denote these two cohorts of patients as the *Standard* and *Challenge* groups throughout this work.

Figure 1: Workflow of automated COVID-19 infection analysis from peripheral blood smear (PBS) image data. **a)** Data collection procedure at Duke University Medical Center. Subjects randomly selected from patient population subset whom had both a COVID-19 RT-PCR test (ground truth labels) and images from a digital PBS scan (input data). All data was collected anonymously, retaining only the results of the RT-PCR test and digitized PBS images. **b)** Summary of pre-processing pipeline, where individual images of white blood cells are extracted from the full slide using a high-resolution oil-immersion microscope. **c)** Data analysis strategy. All cells from a patient are collectively analyzed by a deep neural network to produce both a COVID-19 diagnosis and per-cell importance score.

PBS image data was collected using a clinically approved digital slide scanner (*Cellvision DM9600*), which uses an oil immersion objective lens to capture multiple high resolution images per-patient centered upon stained (Wright-Giemsa) white blood cells (WBCs), with an average of 130 images captured per-patient. To preserve patient privacy, no additional data, such as demographic information, was collected.

2.2. Detecting COVID-19 from Peripheral Blood Smears

Figure 2: Overview of hybrid machine learning system. The SIL branch processes each image from a patient PBS scan individually. The outputs are then aggregated to produce a single prediction (using the median of the single image predictions). The MIL branch collectively analyzes all of a patient’s images simultaneously, producing one prediction per patient. This is accomplished by first extracting learned per-image features, and then feeding those features into an attention module. The attention module assigns weights (summing to one) to each learned feature. The weights are used to compute a weighted sum across image features, the result of which is passed into a classification module to produce the MIL patient prediction. These two strategies are combined through ensembling, where the outputs of each branch are averaged to produce the final outcome.

After collecting high-resolution blood image data, our subsequent goal was to test the accuracy of a novel MIL algorithm to predict patient infection status. Unlike standard supervised machine learning problems that aim to establish a mapping between a single input image and a known output, this problem presents a somewhat unique challenge of accurately mapping a variable number of blood cell images to a single prediction of disease state. In a series of initial tests, we found that while predicting infection from a single image had poor performance (ROC-AUC of ~ 0.7), processing and averaging the predictions from multiple images of unique cells per patient dramatically improved algorithm accuracy. Based upon this key insight, we hypothesized that poor performance at the single image level stemmed from the fact that not every image has the requisite indicators to detect, or rule out, a COVID-19 infection, and that it would be beneficial to jointly optimize a cross-cell predictor aggregation strategy to increase diagnostic accuracy. Accordingly, our final machine learning system used a hybrid of two different multiple instance learning (MIL) methods to map a patient blood imagery to a single diagnostic score (Figure 2). Each system branch offered unique theoretic benefits and in tandem formed an effective classifier that holistically examines patient PBS data.

Our final system allowed us to both effectively predict COVID-19 infection status, and identify which cells were most relevant to disease state prediction. We used a *Receiver Operator*

Characteristic (ROC) curve to quantify the trade-off between diagnostic sensitivity and specificity of our new network, and obtained an area under the curve (ROC-AUC) of 0.90 and 0.89 on the standard and challenge groups respectively, shown in Figure 3. Additionally, we evaluated the accuracy of our predictions by assigning a pre-defined threshold (0.5 confidence score) to the outputs and calculating the percentage of subjects classified correctly. The accuracy of our network was 79% for subjects within the standard group, and 82% within the challenge group. To evaluate the *challenge* group ROC-AUC, we randomly selected an equal number of COVID-19 positive patients from the *standard* group for comparison. We found that performance was roughly equivalent across both cohorts of data, suggesting that the detection mechanism determined by our automated system is not influenced by hidden variables that may change as a function of scan time or patient cohort, but instead is correlated with the underlying disease.

Figure 3: Performance of COVID-19 diagnosis from blood cell morphology analysis as measured by the receiver operator characteristic (ROC). Classification accuracy was 79% and 82% for the standard and challenge cohorts respectively. a) Results reported are the average across the entire dataset, k-fold cross validation was used to maintain independence between the training and test sets. b) COVID-19 positive patients were randomly selected from the *standard* cohort to counterbalance the COVID-19 negative patients from the challenge cohort (ROC-AUC calculation requires both positive and negative examples).

2.3. Cell Importance

In the process of predicting a patient's infection status, our system's MIL model uses an attention mechanism to generate a per-image importance score (trained jointly with the neural network). During a forward model pass, the importance score is used to create a weighted sum of the feature vectors from each image, which is then processed by a classifier module to generate a single diagnostic score per patient. Relative image importance scores may thus be directly examined during patient infection status inference, by translating each into a percentile

score, where a 100th percentile cell image is the most important for the system to reach its COVID-19 diagnosis. With such an approach, we can jointly categorize cell images into their respective cell types (e.g., monocyte, basophil, etc.), and then compute statistical distributions of importance scores within and across cell types (Figure 4a).

From this analysis, we can draw several initial conclusions about the mechanisms used by our machine learning algorithm for COVID-19 detection. First, it is clear that neutrophils had a consistently higher importance value than other cell types. Second, cells classified as monocytes, platelets, and smudged cells (cells likely destroyed during slide preparation), had the lowest average scores and thus had less diagnostic use. Finally, the remaining cell types are moderately important, and likely contribute to an accurate classification, but less so than neutrophils (see example images in Figure 4b).

While these measures of importance alone are not enough to identify the exact mechanisms that our algorithm is using to perform diagnosis, they act as a rough guide to inform us which cells are more or less diagnostically relevant. Due to the inherently complex and non-linear nature of deep neural networks, it is difficult to identify precisely how classification decisions are made. However, our findings - that aspects of neutrophil morphology are important to identify a COVID-19 infection - are well supported by existing literature. Broad findings have recently connected COVID-19 to neutrophil-based abnormalities such as increased amounts of activated neutrophils in the bloodstream^{20,21} and elevated levels of neutrophil extracellular traps,²² among others.²³

Figure 4: **a)** Box plots of cell importance by detected white blood cell type. A higher score indicates higher importance. Neutrophils are the most consistently important cell type, where monocytes, smudged cells, and platelets are less important. The remaining cell types are moderately important and similar to each other. **b)** Examples of white blood cell images from three randomly sampled patients, split by importance level. Highly important cells are the cells of the patient with the highest score, medium importance cells were drawn from cells with scores closest to the 50th percentile of scores, and low importance cells were the four cells from each patient with the lowest scores.

2.4. Perturbation Studies

To understand the spatial factors influencing our ability to detect COVID-19 from PBS images, we conducted a set of *Perturbation Studies*, where we manipulated aspects of the digital PBS image data in a controlled manner during neural network training and inference. In the interest of computational efficiency, these studies were only performed on a representative subset (three of the six folds used for k-fold cross-validation) of our data, and only for the *single-image* branch of our system. As noted above, PBS image data for each patient consists of 130 (on average) cropped images centered on a WBC, typically containing RBCs around the periphery. Accordingly, we varied three unique aspects of these image datasets to help elucidate important factors for accurate COVID-19 diagnosis: the number of unique images per-patient, the amount of occlusion within the central image area that typically contains a white blood cell, and the amount of occlusion of the image periphery that contains red blood cells (see Figure 5). Occlusions were applied by zeroing pixel values in the same manner to all images within each patient's PBS image dataset.

While we can determine how the number of images per-patient influences performance simply by manipulating this value during inference, the latter two perturbations influence how the neural network processes images. Therefore, to effectively understand their impact, we re-

trained networks from scratch using occluded images to create a unique network per occlusion experiment. To jointly evaluate how the *quantity* of images per-patient influenced our ability to screen for COVID-19, we also varied the number of images available for COVID-19 diagnosis inference by randomly selecting up to N images from each patient. This full process was repeated three unique experiments with unmodified, center-occluded, and outer-occluded image data to produce the results summarized in Figure 6.

Figure 5: Visual image perturbations applied during training. Center masking occludes the white blood cell (which is always centered within the image). Outer masking occludes the red blood cells, which are in the area surrounding the white blood cell (towards the outside of the image).

Across all configurations, we observe that a larger number of images per-patient (i.e., morphological data about a larger number of cells) leads to a higher quality screening result, with diminishing returns. This trend supports the notion that morphological indicators for COVID-19 infection are spread across many images (i.e., multiple blood cells). Somewhat surprisingly, with relatively few (~ 16) images, the *original* configuration (no occlusions) reaches close to maximum performance, suggesting that while not every image has the requisite indicators to detect COVID-19, they are moderately prevalent within our dataset.

Examining the effect of the occlusions on the accuracy of our COVID-19 diagnosis predictions, we observe a negative relationship between occlusion size and prediction performance (Figure 6). Contrary to expectations, the system continues to perform fairly well (ROC-AUCs of ~ 0.8) under significant occlusion, if many images are used to make a prediction. These results point to several insights of interest. First, note that red blood cells are completely occluded for nearly all images in the extreme version of outer masking, and white blood cells are completely occluded for nearly all images in the extreme version of center masking. Accordingly, it appears possible to at least weakly predict COVID-19 infection from either information about the white or red blood cells alone. Second, small "glimpses" of information that our model may see across hundreds of images per patient enables fairly accurate diagnosis. When only a few occluded images are available with data on a small number of cells, performance suffers greatly.

Finally, while predictions *can* be made with only the red or white blood cell data, the best performance is consistently achieved when information about both cell types is jointly available.

Figure 6: Results of perturbation experiments. Average ROC-AUC of screening predictions compared to number of images (randomly selected across five trials) used to make the prediction.

These findings are reinforced by related preliminary evidence from the clinical domain. The impact of COVID-19 infection on white blood cell morphology has been observed in a number of studies.^{24, 12, 13, 25} Our finding that red blood cells can morphologically change in response to COVID-19 infection is consistent with recent studies suggesting that red blood cell distribution width (RDW) is a significant predictor of illness,¹¹ and that digital holographic videos of red blood cells can be used to assist with prediction of COVID-19 infection.²⁶ In conclusion, our new approach can jointly provide diagnostic predictions and lead to novel insights into how disease processes impact blood cell morphology. By aggregating trends across many hundreds of cells in a holistic manner, our DOBA pipeline offers a promising new direction for large-scale analysis of hematological and potentially alternative cytopathological image data in future tasks.

3. Methods

3.1. Data Collection

We collected digital anonymized PBS images from patients at the Duke Medical Center (IRB Protocol 00105472). We preserved patient anonymity by only collecting PBS image data and COVID-19 infection status within the *standard* group of tested patients. The patients from the *challenge* group were selected by collecting PBS data from patients admitted to the medical intensive care unit with acute respiratory illness (from pneumonia or other acute respiratory failure) who tested negative for COVID-19. While used for cohort formation, this diagnostic information was not present during analysis. The SARS-CoV-2 infection test was performed with a Nasopharyngeal Swab based PCR test. All PBS were imaged with a *CellaVision DM9600* optical slide scanning system. The system captured high-resolution image segments (estimated $0.44\mu m$ optical resolution) over a 360×360 px image area, of which the inner 240×240 pixels

were used in our analysis. The images were typically centered on WBCs and almost always contained red blood cells. An average of 130 images were captured from each PBS.

The *standard* cohort included 236 patients, 125 of whom tested positive for COVID-19. Using k-fold cross-validation ($k = 6$) we repeatedly split patients into training and test sets (equal proportion COVID-19 positive patients within each set). There was no cross-over of patient data from set-to-set. Unless otherwise indicated, all performance metrics are reported as the average test-set performance across all six folds, where multiple independent models were trained from scratch exclusively for individual folds. This strategy enabled us to test our system on all available data while isolating the test data during the training process.

3.2. Machine Learning System

Our machine learning system is a novel hybrid of two complementary multiple instance learning (MIL) approaches (shown in Figure 2). In a first branch, we used DenseNet-121,²⁷ a convolutional neural network (CNN), to process each image from a patient PBS image set independently. A single diagnostic label per patient was distributed across all of their individual PBS images. The CNN output a per-image confidence score, which were then combined across all images per patient by measuring the proportion of images classified as COVID-19 positive (0.5 threshold). This strategy may initially seem counter-intuitive, as not every image can be used to predict COVID-19 infection. If we consider these mislabelled images as a source of label noise, then we can reconcile the effective performance of our technique despite this phenomenon with the fact that deep learning systems often succeed despite large amounts of label noise.²⁸ By training on individual images, the number of samples within our training dataset increases by several orders of magnitude to minimize overfitting and generalization issues, but at the expense of being unable learn a method to integrate information *across* images.

To address this latter issue, we implemented a second MIL branch that adopted an attention based version of MIL, first shown by Ilse et al.¹⁵ In this strategy, a CNN extracts a feature vector from each patient images and then applies an attention mechanism, implemented using a multilayer perceptron (MLP), to combine these features together. The combined feature vector is then propagated through a final MLP to produce a final patient-level classification. The entire pipeline is differentiable, so we trained this using only patient-wide labels. We used the ResNet50²⁹ CNN architecture to generate feature vectors and MLPs with one hidden layer to both generate attention scores and perform final classifications. While training with patient-level data reduces the number of unique training examples (one per-patient rather than per-image), it allows our model to learn relationships between per-patient cell images.

As a single data point contains N images instead of one within our system's second MIL branch, a forward pass of our model becomes N times computationally larger. To accommodate this overhead, we adopted a simple modification: instead of inputting the entire set of images from each patient, we input a randomly selected subset of images. Based upon our image count findings (see Figure 6), we chose the size of this random subset to be 16. In addition to reducing memory requirements, we believe this sampling strategy acted as a regularization

method to prevent the model from relying on relatively few images per patient for diagnosis formation.

Model ensembling¹⁶ was employed both within and between MIL branches. The output of each branch is the average of three independently trained models, and the output of our total system is the average of both branches. Ensembling substantially reduced average output error and could be additionally scaled up in the future for additional performance gains.

3.3. Evaluation of Cell Importance

The multi-image MIL branch assigned an importance score to each image (centered on a unique WBC) via its attention mechanism. While used collectively when predicting a patient's infection status, scores are functionally computed independently. To directly output per-image importance scores for additional analysis, we created a new sub-model, consisting of a subset of the trained multi-image branch (the feature extraction module and a portion of the attention module). All images across all patients were input into this sub-model and assigned importance scores before translating the distribution of values into percentiles ranging from *0th* (lowest) to *100th* (highest). To assess image importance as a function of cell type, we used a standard cell-type classifier to sort all images into one of nine standard categories: platelets, eosinophils, neutrophils, immature granulocytes, lymphocytes, monocytes, basophils, erythroblasts, and smudged cells (see additional details in Supplementary Information).

3.4. Perturbation Experiments

Spatial perturbations were introduced by modifying all images before being input into the neural network model. It was important to ensure perturbations were in place for the entirety of the training process (i.e., not just applied to test data, but also included during network training). Accordingly, all perturbation results are from independently trained models. We masked out a fixed number of pixels either from the center of the image within a given diameter, or the surrounding area outside of this central circle. For center masking, we used circles centered within the image with diameters of 60px, 120px, 180px, and 240px, for the minor, medium, major, and extreme configurations respectively. For the outer masking experiment, we followed the opposite approach, masking out *all but* the center of the image within the fixed diameters of 240px, 180px, 120px, and 60px for the minor, medium, major, and extreme configurations respectively.

Across all perturbation studies, we examined how the number of images being used per-patient influenced performance by reducing the number of images via random sub-selection. For example, if we wanted to test our system performance for a quantity of 16 images, we would randomly select 16 images from each patient to be included within the analysis, discarding all others. To ensure statistical significance with the random image sub-selection process, we repeated all analyses five times per model. Reported results include the average and standard deviation across all five trials.

4. Data Sharing

Due to the conditions and agreements under which this data was collected, no data will be made publicly available.

References

- ¹ Jones KW. Evaluation of cell morphology and introduction to platelet and white blood cell morphology. *Clinical hematology and fundamentals of hemostasis*. 2009:93–116.
- ² Gulati G, Song J, Florea AD, Gong J. Purpose and criteria for blood smear scan, blood smear examination, and blood smear review. *Annals of laboratory medicine*. 2013;33(1):1.
- ³ Ceelie H, Dinkelaar R, van Gelder W. Examination of peripheral blood films using automated microscopy; evaluation of Diffmaster Octavia and Cellavision DM96. *Journal of clinical pathology*. 2007;60(1):72–79.
- ⁴ Radakovich N, Nagy M, Nazha A. Artificial Intelligence in Hematology: Current Challenges and Opportunities. *networks*. 2020;2:6.
- ⁵ Shafique S, Tehsin S. Acute lymphoblastic leukemia detection and classification of its subtypes using pretrained deep convolutional neural networks. *Technology in cancer research & treatment*. 2018;17:1533033818802789.
- ⁶ Kimura K, Tabe Y, Ai T, Takehara I, Fukuda H, Takahashi H, et al. A novel automated image analysis system using deep convolutional neural networks can assist to differentiate MDS and AA. *Scientific reports*. 2019;9(1):1–9.
- ⁷ Xu M, Papageorgiou DP, Abidi SZ, Dao M, Zhao H, Karniadakis GE. A deep convolutional neural network for classification of red blood cells in sickle cell anemia. *PLoS computational biology*. 2017;13(10):e1005746.
- ⁸ Izcovich A, Ragusa MA, Tortosa F, Lavena Marzio MA, Agnoletti C, Bengolea A, et al. Prognostic factors for severity and mortality in patients infected with COVID-19: A systematic review. *PloS one*. 2020;15(11):e0241955.
- ⁹ Debuc B, Smadja DM. Is COVID-19 a new hematologic disease? *Stem cell reviews and reports*. 2021;17(1):4–8.
- ¹⁰ Berzuini A, Bianco C, Migliorini AC, Maggioni M, Valenti L, Prati D. Red blood cell morphology in patients with COVID-19-related anaemia. *Blood Transfusion*. 2021;19(1):34.
- ¹¹ Lippi G, Henry BM, Sanchis-Gomar F. Red blood cell distribution is a significant predictor of severe illness in coronavirus disease 2019. *Acta Haematologica*. 2021;144(4):4–8.
- ¹² Singh A, Sood N, Narang V, Goyal A. Morphology of COVID-19-affected cells in peripheral blood film. *BMJ Case Reports CP*. 2020;13(5):e236117.
- ¹³ Nazarullah A, Liang C, Villarreal A, Higgins RA, Mais DD. Peripheral blood examination findings in SARS-CoV-2 infection. *American journal of clinical pathology*. 2020;154(3):319–329.

- ¹⁴ Maron O, Lozano-Pérez T. A framework for multiple-instance learning. *Advances in neural information processing systems*. 1998:570–576.
- ¹⁵ Ilse M, Tomczak J, Welling M. Attention-based deep multiple instance learning. In: *International conference on machine learning*. PMLR; 2018. p. 2127–2136.
- ¹⁶ Dietterich TG. Ensemble methods in machine learning. In: *International workshop on multiple classifier systems*. Springer; 2000. p. 1–15.
- ¹⁷ Rahman A, Niloofa R, Jayarajah U, De Mel S, Abeysuriya V, Seneviratne SL. Hematological Abnormalities in COVID-19: A Narrative Review. *The American Journal of Tropical Medicine and Hygiene*. 2021;104(4):1188.
- ¹⁸ Lippi G, Plebani M, Henry BM. Thrombocytopenia is associated with severe coronavirus disease 2019 (COVID-19) infections: a meta-analysis. *Clinica chimica acta*. 2020;506:145–148.
- ¹⁹ Biswas S, Thakur V, Kaur P, Khan A, Kulshrestha S, Kumar P. Blood clots in COVID-19 patients: Simplifying the curious mystery. *Medical Hypotheses*. 2021;146:110371.
- ²⁰ Parackova Z, Zentsova I, Bloomfield M, Vrabcova P, Smetanova J, Klocperk A, et al. Disharmonic inflammatory signatures in COVID-19: augmented neutrophils' but impaired monocytes' and dendritic cells' responsiveness. *Cells*. 2020;9(10):2206.
- ²¹ Wang J, Li Q, Yin Y, Zhang Y, Cao Y, Lin X, et al. Excessive neutrophils and neutrophil extracellular traps in COVID-19. *Frontiers in immunology*. 2020;11:2063.
- ²² Middleton EA, He XY, Denorme F, Campbell RA, Ng D, Salvatore SP, et al. Neutrophil extracellular traps contribute to immunothrombosis in COVID-19 acute respiratory distress syndrome. *Blood, The Journal of the American Society of Hematology*. 2020;136(10):1169–1179.
- ²³ Cavalcante-Silva LHA, Carvalho DCM, de Almeida Lima É, Galvão JG, da Silva JSdF, de Sales-Neto JM, et al. Neutrophils and COVID-19: The road so far. *International immunopharmacology*. 2021;90:107233.
- ²⁴ Merad M, Martin JC. Pathological inflammation in patients with COVID-19: a key role for monocytes and macrophages. *Nature Reviews Immunology*. 2020:1–8.
- ²⁵ Chong VC, Lim KGE, Fan BE, Chan SS, Ong KH, Kuperan P. Reactive lymphocytes in patients with Covid-19. *British Journal of Haematology*. 2020;189(5):844–844.
- ²⁶ O'Connor T, Shen JB, Liang BT, Javidi B. Digital holographic deep learning of red blood cells for field-portable, rapid COVID-19 screening. *Optics Letters*. 2021;46(10):2344–2347.

- ²⁷ Huang G, Liu Z, Van Der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2017. p. 4700–4708.
- ²⁸ Rolnick D, Veit A, Belongie S, Shavit N. Deep learning is robust to massive label noise. arXiv preprint arXiv:1705.10694. 2017.
- ²⁹ 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; 2016. p. 770–778.