Optimizing A Spacial Illumination Filter to Improve the Segmentation of Spirochaeta Bacteria

Mahish Kewalramani, Kaan Sahingur*
Department of Biomedical Engineering
Duke University
Durham, NC 27708
mahish.kewalramani@duke.edu, kaan.sahingur@duke.edu

Abstract

Segmentation of pathogens in blood is a vital task that is done in the biomedical community. It can provide quick insight into whether or not someone is infected by a virus or bacteria, which can help determine treatment. The present work explores the addition of a physical layer over a commonly used deep neural network for segmentation tasks, the U-Net, to build a model that can yield better results. The results show that the model with a physical layer does not do much better in segmenting the data. That being said, the model with the physical layer does not over fit to the train data as much, indicating promising results of including a physical layer in imaging blood samples.

1 Introduction

Throughout history, medical imaging has been a critical tool in both clinical and research settings. The proper collection and analysis of images has held paramount importance over the years, especially as technology has evolved. As imaging techniques and capabilities have drastically improved, the need to analyze increasing amounts of data has emerged. With the rise of machine learning, the ability to utilize models that help accurately assess these images can lead to incredible breakthroughs.

Bacteria detection in blood is one area that poses widespread potential research significance. The ability to identify invasive, harmful bacteria in the blood could promote new breakthroughs, both clinical and non-clinical. The Spirochaeta bacteria is one such bacteria that could be useful to identify. This type is characterised by its distinct rod shape and unique motility, seen in Figure 1 below [5]. In addition, Spirochaeta are quite invasive and responsible for several diseases, such as syphilis, Lyme disease, relapsing fever, and leptospirosis.

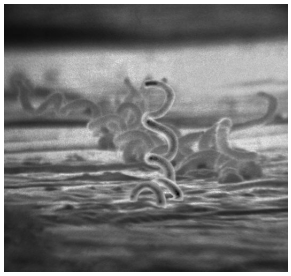

Figure 1: Spirochaeta bacteria

*Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies.

This project seeks to accurately classify Spirochaeta bacteria from segmented data. Data segmentation is an image processing task that is very common in medical research. It involves the process of taking data, dividing it up, and grouping based on selected parameters. The imaging of the bacteria was captured using dark-field microscopy. This microscopy technique produces a darker background and focused specimen than normal microscopy. With this data, it uses obstruction to get better imaging of the bacteria [3]. Unscattered light beams are excluded, resulting in a darker specimen field. When no sample is on the microscope stage, the entire field appears darker. When a sample is placed on the stage, the light at the apex of the cone strikes it. The rays scattered by the sample and captured in the objective lens make the image [3]. Figure 2 shows dark-field microscopy in action. This study aims to improve the imaging and classification of this bacteria by implementing a physical layer that optimizes how this light strikes the sample.

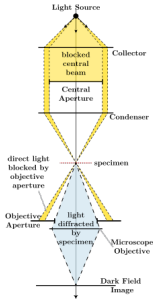

Figure 2: Dark-field microscopy

2 Related Works

In order to best optimize the deep learning of the segmentation data, a U-Net was utilized. A U-Net is a proven network architecture that is very commonly used when working with biomedical image segmentation [2]. It allows for high accuracy without requiring a large pool of data. This greatly benefited the data split of the segmentation data, with 292 train and 74 test images. The U-Net was first described in the paper *UNet: Convolutional Networks for Biomedical Image Segmentation*. The authors present a network and training strategy that strongly focuses on data augmentation to use annotated samples more efficiently [4]. It consists of both a contracting path that captures context and a symmetric expanding path that enables precise localization [4]. This allows for effective end-to-end training to occur with fewer images, and it inspires the "U"-net name. Figure 3 below shows the basic U-Net architecture diagram. The decision to use U-Net enabled for accurate segmentation to occur on the bacteria data.

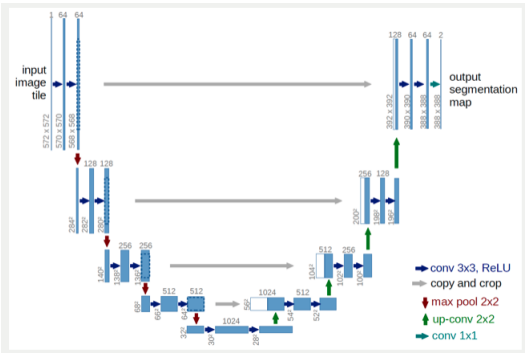

Figure 3: Basic U-Net architecture diagram

In addition, other papers have shown the usage of U-Net for similar biomedical image segmentation. A paper authored by J. B. Abraham in 2019 shows the use of U-Net for malaria parasite segmentation [1]. It looked at Plasmodium segmentation on thin blood smear images. The results show that U-Net can accurately perform this segmentation [1]. The success of parasite segmentation in blood provided confidence in using U-Net for bacteria segmentation in blood. Figure 4 shows examples of the test

data and train data. As you can see, the test is on the left and the train on the right. Blue is bacteria, green are erythrocytes, and red is the background.

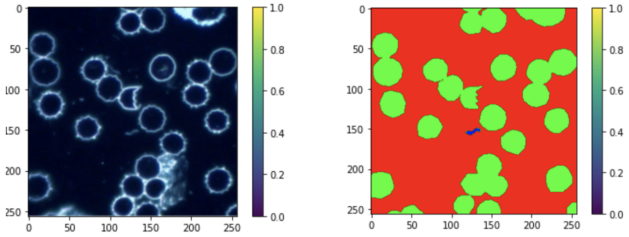

Figure 4: Test and train data examples

3 Methods

3.1 Dataset

The deep learning model utilized in this paper was developed, trained, and evaluated on was a data set of 366 dark-field microscopy images that contained Spirochaeta bacteria in blood from Hoschschule Heilbronn University. This data set had both images and manually annotated masks that could be used for segmentation purposes. There were three different classes that needed to be segmented in this model, the background, red blood cells, and Spirochaeta bacteria. In order to start working with this data set,the images were downloaded and imported it to Google Drive. After doing this the data was imported and pre-processed such that it could be trained. The prepossessing pipeline included scaling down the images to a 256x256 image using cv2.imread for both the images and labels. The labels were also one hot encoded such that the background was encoded by the color red, the red blood cells were encoded by the color green, and the bacteria were encoded by the color blue. Examples of the images and labels are shown in figure 4. An 80-20 Train-Test split was used for the data so that there was sufficient amount of images to train then model and see performance.

3.2 Neural Network

The goal for the model was to preform semantic segmentation on the images from the Spirochaeta bacteria dataset. This means that it has to go through the image and label each pixel based on what the model thinks is being represented. As discussed earlier, a U-Net architecture is a very popular neural network applied to these type of problems, specifically for biomedical data. For the encoder side of the U-Net - used to capture the context of the image - it started with two 2D convolution layers with 8 3x3 filters each followed by a batch normalization and 2D max pooling layer. The encoder repeated this framework for 16, 32, 64, and 128 filters. Next, there was the decoder part of the U-Net network, which determined where the information is using up-sampling layers. For the up-sampling layer, it utilized an up-sampling layer followed by a 2D convolution that used a 2x2 kernel. This up-sampling layer was followed by a skip connection between the the encoder and decoder sides, then two 2D convolution layers with 3x3 kernels each followed by a batch normalization layer. This up-sampling, skip connection, and 2 convolution layers were done with 64, 32, 16, and 8 filters. All the convolutions in the encoder and decoder parts of the U-Net used a relu activation and had a normal kernel initializer. Finally, the last layer in the convolution layer was a 2D convolution with 3 filters, had a 1x1 kernel, and had a softmax activation. This final layer gave 3 outputs corresponding to the three classes in the segmentation data to classify what each pixel corresponded to.

As a note, the U-Net model that was used as a basis did not have the batch normalization layers after each of the convolution layers. Without these batch normalization layers, the model initially predicted everything as background. This solution kept happening due to class imbalances in the data, as the occurrences of background were much higher than either the red blood cells or bacteria. With batch normalization, however, the model was able to perform the segmentation task much better.

3.3 Physical Layer

For the physical layer, a trainable mask that would optimize the spatial illumination of the dark-field microscope was made. Dark-field microscopes work through illuminating a field such that the sample gives signal and the background is dark based on refractive and reflective properties of the sample. In this way, even clear samples can be seen as long as they reflect and refract some light into the objective lens. The goal was to find the optimal illumination for the sample to get better classification of the images in the segmentation model. The mask was initialized as a 256x256 trainable variable that had a constant value of ones. This mask was then trained to in the network to get an optimal illumination.

4 Results

4.1 U-Net's capability to Segment Data

Two examples of the U-Net's ability to segment the data set without a physical layer is shown in the figures below:

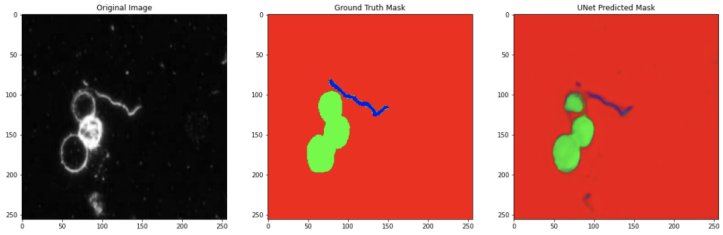

5a: Good Prediction With Our U-Net

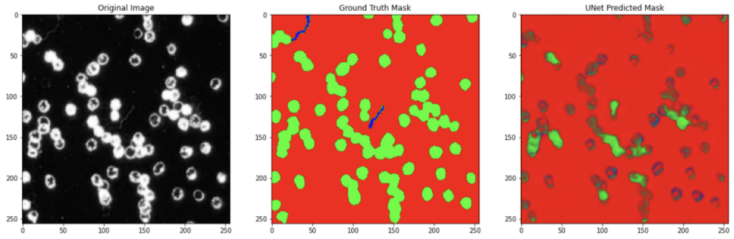

5b: Bad Prediction With Our U-Net

Figure 5: Segmentation Capabilities of Our U-Net

As can be seen from the two images above, our model is good at segmenting data that does not have two many different objects in it. This can be seen in figure 5a. When a lot of bacteria and red blood cells are in the frame of the image however, the predictive capabilities go way down as seen from figure 5b as the model clearly is not segmenting the data properly here.

4.2 Comparison of Our Two Models

At a first glance, it might seem like our model with a physical layer essentially gave the same results as our model without the physical layer in its ability to segment the data. Looking at both models losses after 25 epochs we saw the data shown in the table below:

Model	Train Loss	Validation Loss
No Physical Layer	0.1725	0.361
With Physical Layer	0.2967	0.351

Table 1: Comparison of the loss of both models

As can be seen the validation loss without the physical layer was 0.361 while the validation loss with the physical layer was 0.351. These losses are essentially the same and so does not differentiate the

physical layer model for the one without the physical layer. That being said, when looking at the training loss, you start to see that the model with the physical layer is indeed preforming better. The training loss with the physical layer was 0.2967 as opposed to 0.1725 without the physical layer. This means that there is less over fitting in the model without a physical layer. Plots of the losses for both models are shown in the figures below:

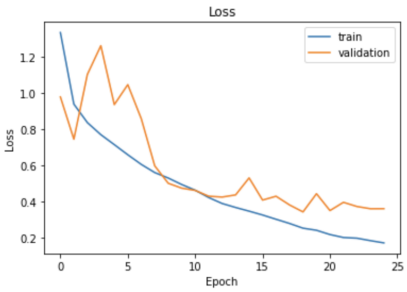

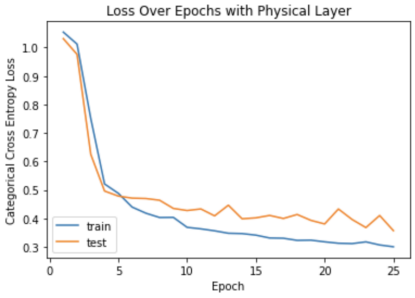

6a: Model Loss Without Physical Layer 6b: Model Loss With Physical Layer

Figure 6: Plots of Both Models Losses

5 Discussion

5.1 Comparison of Models

As mentioned in the results section, the training loss with the physical layer was 0.2967 and the training loss of the model without the physical layer was 0.1725. That being said, both models come to about the same loss for the validation data showing that there was less over fitting with our model with the physical layer. This is an interesting insight as this means that with images that utilize the illumination mask that was trained, predictions will be more generalized. This is important in the real world as a model that when taking images of blood there will be a lot of variability and so models which segment images from blood samples should try to avoid just trying to find patterns seen in the train data. Another thing seen between the two models is that the model with a physical layer

5.2 Physical Layer Insights

Some physical insights form our model can be seen in the figure below:

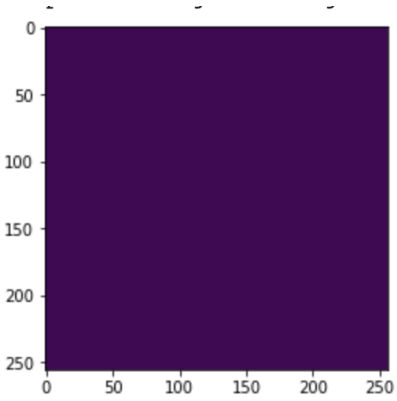

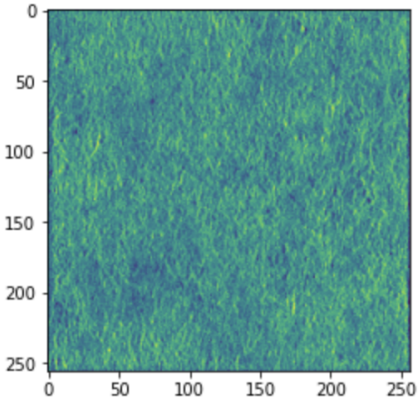

7a: Pre trained mask 7b: Mask After Training

Figure 7: Mask Before and After Training

As can be seen above, the mask before and after training is shown. The mask out uniformly initialized, and after training became what is shown on the right.

Below, you can see how the image looks when using the mask that did not undergo training versus the how the image looks when using a mask that does undergo training. Along with that, for reference the label of the image is also there:

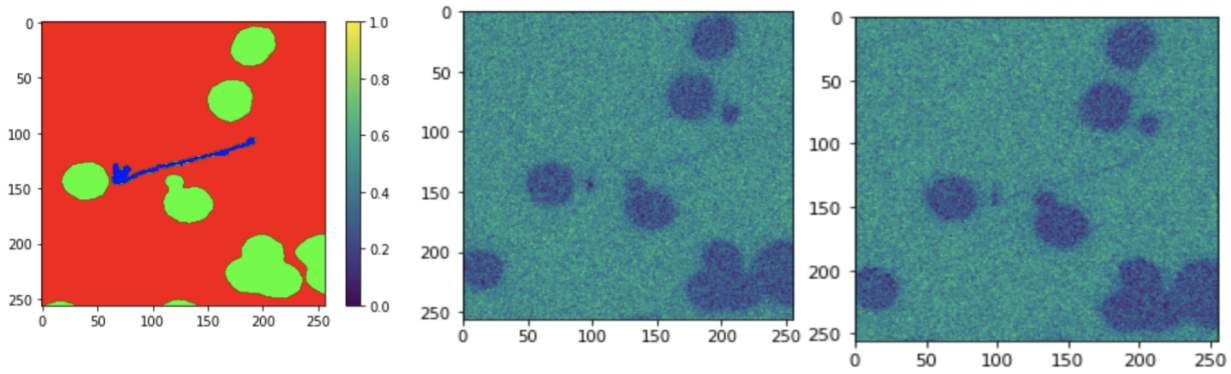

Figure 9: Label(left) image with non trained mask (center) and image with trained mask (right)

As can be seen from the image above, one can more clearly see the bacteria with the mask that is trained than with the mask that is not trained. This shows that our mask was indeed being optimized for an overall better segmentation. By more clearly labeling the bacteria our model will be able to better segment the data in a correct manner.

5.3 Future Work

For future work, one thing that we can do to improve segmentation of our neural network is to input loss weights. As stated earlier there is a much higher representation of the background than the blood cells, and there is even loss bacteria present in the images. This made it hard for the model to segment the data properly. A way to combat that after looking online is to implement loss weights based on the effective number of samples of each class from the labels [6]. This can be the first step towards improving our model. For the physical layer, we can attempt to optimize a phase layer for our image. Dark image microscopy heavily relies on the phase of the incoming light to get proper images. By implementing a phase layer we further improve our model in optimizing the physical parameters of how the image should be taken.

6 Conclusion

The ability to recognize pathogens in blood, which in this case is Spirochaeta bacteria, is an important tool to use to diagnose patients that come in with things like illnesses. In this project, we imposed a physical layer that was supposed to give an optimal illumination pattern for dark field microscope imaging to to improve a CNN image segmentation model. Our results indicated that while the performance of our model did not really improve with the physical layer, the model with the physical layer was over-fitting less to the data which is an important aspect to consider. We believe further work needs to be done in optimizing other parts of the physical space to get even better segmentation.

Acknowledgments

Thank you to Dr. Roarke Horstmeyer, Colin Cooke, and the Duke University Biomedical Engineering Department for all of their time, guidance, and encouragement.

References

[1] Abraham, J. B. (2019). Malaria parasite segmentation using U-Net: Comparative study of loss functions. Communications in Science and Technology, 4(2), 57-62. https://doi.org/10.21924/cst.4.2.2019.128.

[2] Cai, L., Gao, J., Zhao, D. (2020). A review of the application of deep learning in medical image classification and segmentation. Annals of translational medicine, 8(11), 713. https://doi.org/10.21037/atm.2020.02.44.

[3] Fakhrullin, R., Nigamatzyanova, L., Fakhrullina, G. (2021). Dark-field/hyperspectral microscopy for detecting nanoscale particles in environmental nanotoxicology research, Science of The Total Environment, 772, https://doi.org/10.1016/j.scitotenv.2021.145478.

[4] Ronneberger, O., Fischer, P., Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention. 2015. https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/.

[5] Wirth, R., Ugele, M., Wanner, G. (2016). Motility and Ultrastructure of Spirochaeta thermophila. Frontiers in microbiology, 7, 1609. https://doi.org/10.3389/fmicb.2016.01609.

[6] Y. Cui, M. Jia, T. Lin, Y. Song and S. Belongie, "Class-Balanced Loss Based on Effective Number of Samples," 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 9260-9269, doi: 10.1109/CVPR.2019.00949.