

# **Uncertainty Quantification for Eosinophil Segmentation**

# Kevin Lin

School of Data Science University of Virginia Charlottesville, Virginia, USA pex7ps@Virginia.edu

## Sana Syed

Department of Pediatric Gastroenterology University of Virginia Charlottesville, Virginia, USA ss8xj@virginia.edu

# ABSTRACT

Eosinophilic Esophagitis (EoE) is an allergic condition increasing in prevalence. To diagnose EoE, pathologists must find 15 or more eosinophils within a single high-power field (400X magnification). Determining whether or not a patient has EoE can be an arduous process and any medical imaging approaches used to assist diagnosis must consider both efficiency and precision. We propose an improvement of Adorno et al's approach for quantifying eosinphils using deep image segmentation. Our new approach leverages Monte Carlo Dropout, a common approach in deep learning to reduce overfitting, to provide uncertainty quantification on current deep learning models. The uncertainty can be visualized in an output image to evaluate model performance, provide insight to how deep learning algorithms function, and assist pathologists in identifying eosinophils.

## **CCS CONCEPTS**

• Computing methodologies  $\rightarrow$  Supervised learning by classification; Neural networks; Cross-validation.

## **KEYWORDS**

Computer Vision, Bayesian Machine Learning, Domain Adaptation

## ACM Reference Format:

Kevin Lin, Donald Brown, Sana Syed, and Adam Greene. 2023. Uncertainty Quantification for Eosinophil Segmentation. In 2023 the 10th International Conference on Bioinformatics Research and Applications (ICBRA) (ICBRA 2023), September 22–24, 2023, Barcelona, Spain. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3632047.3632050

## **1 INTRODUCTION**

Eosinophilic esophagitis (EoE) is an inflammatory disease of the esophagus characterized by the prevalence of a type of white blood cell (eosinophil). Approximately 0.5-1.0 in 1,000 people have EoE



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

ICBRA 2023, September 22–24, 2023, Barcelona, Spain © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0815-2/23/09 https://doi.org/10.1145/3632047.3632050

## Donald Brown

School of Data Science University of Virginia Charlottesville, Virginia, USA deb@virginia.edu

## Adam Greene

Department of Pediatric Gastroenterology University of Virginia Charlottesville, Virginia, USA arg7ef@virginia.edu



Figure 1: Example Image Data from Gastroenterology Data Science Lab

and it can be seen in 2-7% of patients that undergo endoscopies [5]. Although the cause of EoE remains unclear, pathologists believe EoE to be triggered by a patient's diet. Furthermore, EoE is only increasing in prevalence [4] leading to an increased load on pathologists. Patients with EoE typically present with swallowing difficulties, food impaction, and chest pain [12]. To diagnose EoE, patients must undergo an endoscopy where eosinophils biopsy tissue samples are then evaluated for concentration of eosinophils. Pathologists diagnose the patient with EoE if at least one High-Power Field (HPF; 400× magnification adjustment) within a patient's tissue biopsy slide contains 15 or more eosinophils [6]. The dataset is obtained from the Gastroenterology Data Science Lab from UVA Hospital patient data. A sample image is given in Figure 1. Each image is 512x512x3 large and there are 514 images/masks in the dataset spanning 30 UVA Medical Center patients. We will keep the three channels [r,g,b] for the image but will import the masks as gravscale. All data is obtained from subjects under conditions of academic use only. No personal health information (PHI) is present in the data.

#### 1.1 Related Work

The frequency of use of ResNet and UNet models to detect diseases has been increasing in recent years; classification and segmentation of medical images using these models has shown efficient results [8][13]. Namely, Adorno, et al.[1] demonstrated impressive performance in segmentation of Eosinophils to assist diagnosis of Eosinophilic Esophagitis. However, applications of these models in deep neural networks have caused concern among scholars about their inability to measure uncertainty [3]. There are two common approaches to quantifying uncertainty in machine learning: aleatoric and epistemic. Aleatoric uncertainty refers to the variability in outcomes with random effects, and epistemic uncertainty refers to a lack of information in applying models [9]. For example, while one could include more weights in the neural network to address epistemic uncertainty, aleatoric uncertainty cannot be addressed, only identified. Here we attempted to quantify aleatoric uncertainty using MCD UNet, which has proven effective in other applications without changing or optimizing architectures [2]. Motivation for this work stems from the drive in medical image analysis to increase efficiency and precision in results. Increasing the efficiency of a model and thereby reducing the time needed to train disease detection models can reduce the load on medical providers and lead to better overall patient outcomes while increasing the precision in the results can reduce the probability of error. In the case where a model is attempting to detect a disease and potentially affect a patient's treatment, there is little room for error.

#### 2 METHODOLOGY

## 2.1 Mathematical Linkage Between Problem and Method

We will treat the data as discrete since the image is made up of a three dimensional tuple [r, g, b]. Since we have a very large dataset, we will use variational approximation, specifically minimizing KL divergence, in order to approximate the distribution of the affected cells. KL divergence is given by

$$\begin{aligned} & \mathbf{KL}(q(\mathbf{Z}||p(\mathbf{Z}|D))) = \\ & -(E_q(\log p(D,\mathbf{Z})) - E_q(\log q(\mathbf{Z}))) + \log p(D) \end{aligned}$$

Since maximizing the evidence lower bound (ELBO) is largely impractical, we will be using a Monte-Carlo approximation. We will use a UNet convolutional neural network architecture and use Monte Carlo (MC) Dropout as a Bayesian approximation to identify which image segments correspond to different cells [7]. In Bayesian neural networks, each weight is represented by a probability distribution which we will assume is Gaussian instead of just a number. The learning aspect corresponds to Bayesian inference which we will use MC Sampling. We will measure uncertainty for every pixel using cross-entropy over two classes of "background" (C = 0) and "foreground"(C = 1):

$$U = -(p_{C=0} \cdot \ln(p_{C=0}) + p_{C=1} \cdot \ln(p_{C=1}))$$

We will also explore if we can use aleatoric uncertainty to measure performance. This is given by

$$E_{p(z|D)}H[p(y|z,x)] = -\int p(z|D)\left(\sum_{y\in Y} p(y|z,x)\log p(y|z,x)dw\right)$$

#### 2.2 Bayesian Methods

Work from Gal and Ghahramani in 2016 [7], suggest that a Monte Carlo Dropout UNet is equivalent to the deep Gaussian process used Lin et al.



**Figure 2: UNet Architecture** 

in Bayesian Neural Networks. Essentially, we can minimize the KL divergence using approximation through Monte Carlo integration to get an unbiased estimate. Minimizing the KL divergence between the approximate posterior q(w) and the posterior of the full deep Gaussian Process p(w|X, Y) is given by the objective function:

$$-\int q(w)\log p(Y|X,w)dw + KL(q(w)||p(w))$$

The first and second term can be represented by a sum and approximated by Monte Carlo integration. For the Monte Carlo Dropout UNet, we apply dropout before every weight where dropout is defined as switching off neurons at each training step.

#### 3 RESULTS

The UNet is convolutional network architecture for fast and precise segmentation of images [11]. It has been shown to outperform what was previously considered the best method (a sliding-window convolutional network) on the ISBI challenge for segmentation of neuronal structures in electron microscopic stacks. For the UNet model, we used 23 convolutional layers with batch normalization. In both the encoder and decoder steps, we used a ReLU activation function, and for the final layer we used a sigmoid activation function. Loss was computed using binary cross entropy and ADAM was used as the optimizer. To prevent overfitting, we implemented early stopping and data augmentation. The four models we used are MCD UNet, UNet, DenseNet, and ResNet50. All training was done on 4 NVIDIA A100 GPUs with 300GB of RAM in TensorFlow/Keras 2.7. Each model run was ran for 100 epochs with a learning rate of 0.001.

The underlying architecture of the UNet model is shown in Figure 2. At first, the encoder is used to obtain and normalize the transformation of the input volume, using a Leaky ReLU activation function at each layer. At the bottleneck of this architecture, the volume will be in the size of  $2 \times 2 \times 2$  which represents the reduction of dimensionality prior to using a sigmoid activation function for segmentation. The decoder then up-samples this transformed  $2 \times 2 \times 2$  volume to reconstruct the image with this segmentation.

Our main evaluation of the algorithm is through a combination of the Binary Cross Entropy Loss and Sørensen-Dice Loss given by the Dice Coefficient. For each segmentation method, we will tune the hyperparameters to minimize our loss function.

#### 3.1 Sorensen-Dice Coefficient

The dice coefficient is the industry and academic standard for evaluating medical image segmentation and classification results. Evaluation is given through a scale of 0 to 1. Given two images *X* and *Y*, Uncertainty Quantification for Eosinophil Segmentation

Test Set Dice Coefficient Statistics						
Model	Size	Median	Min	Max		
U-Net	4.9M	0.628	0.594	0.685		
U-Net	8.6M	0.660	0.632	0.698		
U-Net	10.9M	0.665	0.600	0.701		
Res. U-Net	2.7M	0.632	0.588	0.697		
Res. U-Net	4.7M	0.656	0.609	0.696		
Res. U-Net	7.4M	0.557	0.541	0.645		
R2U-Net	3.4M	0.634	0.606	0.686		
R2U-Net	6.0M	0.614	0.530	0.647		
R2U-Net	9.4M	0.631	0.572	0.666		
Attn. U-Net	3.1M	0.517	0.439	0.627		
Attn. U-Net	4.5M	0.529	0.465	0.586		

Figure 3: Adorno et.al. Dice Values [1]

a zero dice coefficient indicates that there is no similarity between two images while a one dice coefficient indicates that the images are exactly the same down to the pixel. The equation for the metric is given by

$$DCS = \frac{2|X \cap Y|}{|X| + |Y|}$$

From the literature review, most state-of-the-art dice scores on medical images range from 0.3 to 0.7. Much like a regression metric, a dice score of 0 or 1 are usually causes for concern since no two images from real-world data are truly the same. Although accuracy is often used as an evaluation metric, it is important to note that the dice score may be low even if the training and validation accuracy are high on a given dataset.

Once again, our dataset is 514 images from the GI Data Science Lab and sized 512x512x3. To improve upon Adorno, et. al. [1], we augment the data through flipping and rotation for each image. This means that for each input image, we have created a flipped and rotated image as well, effectively tripling our input dataset. Data augmentation here makes the model generalize better due to the larger amount of training data. For reference, the results from Adorno, et. al. [1] are shown in Figure 3. The size field in Figure 3 refers to the total number of parameters of each model. For comparison, our UNet results are shown in Table 1.

Model	Size	Median	Min	Max
MCD UNet	494K	0.591	0.470	0.650
UNet	494K	0.598	0.442	0.657
DenseNet	7.2M	0.592	0.369	0.645
ResNet	24M	0.612	0.505	0.651

Table 1: Test Results Dice Score

Most significantly, we can see that our approach has only  $\approx$  494,000 parameters while the smallest model in Adorno et al. had 3.1 million parameters and the largest had 10.9 million parameters. Thus, our work is at least one order of magnitude less in size than Adorno et al. which means our model is less complex. Given that the same UNet approach, we can assume, holding all other factors such as GPU availability constant, our approach runs faster

ICBRA 2023, September 22-24, 2023, Barcelona, Spain



Figure 4: Output comparing the true, predicted, and uncertainty values for an image patch

and more efficiently. Despite the significant difference is model size, our performance is similar to Adorno et. al with dice scores around 0.6, matching their Residual UNet and R2UNet results while outperforming their Attention UNet results.

Possibly the strongest result of this approach is the visualization possible due to the quantification of our model's uncertainty shown in Figure 4.

In this example image, white indicates high amounts of uncertainty while black indicates low amounts of uncertainty. We can see that the pixels are nearly white around the borders of the eosinophils essentially "highlighting" them in the output image. The high uncertainty on boundary pixels around each eosinophil indicate that until the model sees more of the eosinophil, the model hesitates to classify the pixels as an eosinophil. We can see this verified that once the model passes the boundary pixels of each eosinophil, the uncertainty drops significantly and the interior of the eosinophil is nearly black. This dramatic shift in values from the outside of the eosinophil which is black, to the border of the eosinophil which is white, to the interior of the eosinophil which is black again creates these "rings" of white in the resulting image. In areas where multiple eosinophils are clustered, the model seems to struggle to differentiate between the eosinophils leading to a "cloud" of high uncertainty around the area. However, we can see that this output provides valuable information to pathologists as the highlighting the general area of interest and reducing the visual load compared to the original input on the far left.

#### **4** ANALYSIS AND INTERPRETATION

For analysis, we plotted the Dice scores of all models in Figure 5. The boxplots overlapped, which indicates the results from each neural net are not statistically significantly better than others. This is important because we can obtain uncertainty visualizations from the MCD UNet, as shown in Figure 4, whereas UNet and ResNet50 do not allow for this. Monte Carlo Dropout allows us to visualize the uncertainty through the dropout layers since it randomly turns off neurons during training, which adds this stochastic element.

As the novel piece in this work, we compared the model's uncertainty first through an overview of the raw data and through a boxplot for visualization. We can see that the uncertainty of almost all model is comparatively similar with only the DenseNet having an outlier uncertainty value of 0.05. Considering the significant difference in the order of magnitude between the sizes of the models, the relative similarity in the uncertainty values demonstrate that at least at a sufficient size, model uncertainty stays consistent regardless of the size of the model chosen.



Figure 5: Boxplots of Model Dice Scores

Model	Size	Median	Min	Max
MCD UNet	494K	0.007	0.004	0.016
UNet	494K	0.007	0.005	0.013
DenseNet	7.2M	0.009	0.006	0.05
ResNet	24M	0.008	0.005	0.01







## **5 CONCLUSION AND FUTURE DIRECTIONS**

As stewards of patient data, researchers in medical image analysis must ensure all models perform efficiently and accurately. This work addresses both concerns demonstrating MCD UNet's comparable performance at an order of magnitude smaller than current models and introducing model uncertainty as an evaluation metric. All models in this work had comparable values in uncertainty indicating that at least after a model reaches a certain size, the aleatoric uncertainty will stay constant. We then illustrated the model's uncertainty through a visualization which highlighted the eosinophils in the resulting image. Scaling this approach with uncertainty to full size biopsy images can help pathologists quickly identify eosinophils while also reducing the mental load. Compared to a screen full of cells and color, the black and white "rings" circling the eosinophils can at least narrow down eosinophil locations while having the potential to count all eosinophils and output a mask showing their exact locations. One of the most difficult parts of this work was working with limited patient data which is a common challenge in the field of medical image analysis. A potential improvement to this work would be to incorporate few-shot learning. While our dataset exists in a high-dimensional space, we only have a limited number of samples available to us due to the cumbersome nature of acquiring annotated histology images. Few-shot learning has significantly improved classification accuracy in medical imaging datasets [10] and likely would produce competitive results.

## ACKNOWLEDGMENTS

The authors would like to thank the medical researchers of the UVA Gastroenterology (GI) Data Science Laboratory for obtaining the EoE dataset through biopsies of UVA Medical Center patients. Additionally, the staff of the GI Data Science Laboratory tirelessly provided direct medical feedback on our results. Research reported in this publication was supported by National Institutes of Health (NIH) through the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK) under award numbers K23DK117061-01A1 (Syed) and R01DK132369 (Syed).

#### REFERENCES

- [1] William Adorno, Alexis Catalano, Lubaina Ehsan, Hans Vitzhum von Eckstaedt, Barrett Barnes, Emily McGowan, Sana Syed, and Donald E. Brown. 2021. Advancing eosinophilic esophagitis diagnosis and phenotype assessment with deep learning computer vision. BIOIMAGING 2021 - 8th International Conference on Bioimaging; Part of the 14th International Joint Conference on Bioimedical Engineering Systems and Technologies, BIOSTEC 2021 (2021), 44–55. https://doi.org/10.5220/0010241900440055
- [2] Mehmet Yigit Avci, Ziyu Li, Qiuyun Fan, Susie Huang, Berkin Bilgic, and Qiyuan Tian. 2021. Quantifying the uncertainty of neural networks using Monte Carlo dropout for deep learning based quantitative MRI. (2021). http://arxiv.org/abs/ 2112.01587
- [3] Edmon Begoli, Tanmoy Bhattacharya, and Dimitri Kusnezov. 2019. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence* 1, 1 (2019), 20–23. https://doi.org/10.1038/s42256-018-0004-1
- [4] Stuart Carr, Edmond S. Chan, and Wade Watson. 2018. Eosinophilic esophagitis. Allergy, Asthma and Clinical Immunology 14, s2 (2018), 1–11. https://doi.org/10. 1186/s13223-018-0287-0
- [5] Evan S. Dellon. 2014. Epidemiology of eosinophilic esophagitis. Gastroenterology Clinics of North America 43, 2 (2014), 201–218. https://doi.org/10.1016/j.gtc.2014. 02.002
- [6] Glenn T. Furuta, Chris A. Liacouras, Margaret H. Collins, Sandeep K. Gupta, Chris Justinich, Phil E. Putnam, Peter Bonis, Eric Hassall, Alex Straumann, and Marc E. Rothenberg. 2007. Eosinophilic Esophagitis in Children and Adults: A Systematic Review and Consensus Recommendations for Diagnosis and Treatment. Sponsored by the American Gastroenterological Association (AGA) Institute and North American Society of Pediatric Gastroenterol. *Gastroenterology* 133, 4 (2007), 1342–1363. https://doi.org/10.1053/j.gastro.2007.08.017
- [7] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. 33rd International Conference on Machine Learning, ICML 2016 3 (2016), 1651–1660.
- [8] Evan Hann, Ricardo A. Gonzales, Iulia A. Popescu, Qiang Zhang, Vanessa M. Ferreira, and Stefan K. Piechnik. 2021. Ensemble of Deep Convolutional Neural Networks with Monte Carlo Dropout Sampling for Automated Image Segmentation Quality Control and Robust Deep Learning Using Small Datasets. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 12722 LNCS, July (2021), 280–293. https://doi.org/10.1007/978-3-030-80432-9{\_}22
- [9] Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. Vol. 110. Springer US. 457–506 pages. https://doi.org/10.1007/s10994-021-05946-3
- [10] Mijung Kim, Jasper Zuallaert, and Wesley De Neve. 2017. Few-shot learning using a small-sized dataset of high-resolution FUNDUS images for glaucoma diagnosis. MMHealth 2017 - Proceedings of the 2nd International Workshop on

Uncertainty Quantification for Eosinophil Segmentation

Multimedia for Personal Health and Health Care, co-located with MM 2017 (2017), 89–92. https://doi.org/10.1145/3132635.3132650
[11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolu-

- [11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. International Conference on Medical image computing and computer-assisted intervention (5 2015), 234–241. http://arxiv.org/abs/1505.04597
- [12] Thomas M. Runge, Swathi Eluri, Cary C. Cotton, Caitlin M. Burk, John T. Woosley, Nicholas J. Shaheen, and Evan S. Dellon. 2017. Causes and Outcomes of

Esophageal Perforation in Eosinophilic Esophagitis. *Journal of Clinical Gastroenterology* 51, 9 (2017), 805–813. https://doi.org/10.1097/MCG.000000000000718 [13] Devvi Sarwinda, Radifa Hilya Paradisa, Alhadi Bustamam, and Pinkie Anggia.

[13] Devvi Sarwinda, Radifa Hilya Paradisa, Alhadi Bustamam, and Pinkie Anggia. 2021. Deep Learning in Image Classification using Residual Network (ResNet) Variants for Detection of Colorectal Cancer. *Procedia Computer Science* 179, 2019 (2021), 423–431. https://doi.org/10.1016/j.procs.2021.01.025