



# Deep-learning based, automated segmentation of macular edema in optical coherence tomography

CECILIA S. LEE,<sup>1</sup> ARIEL J. TYRING,<sup>1</sup> NICOLAAS P. DERUYTER,<sup>3</sup> YUE WU,<sup>1</sup> ARIEL ROKEM,<sup>4</sup> AND AARON Y. LEE<sup>1,2,4,\*</sup>

<sup>1</sup>Department of Ophthalmology, University of Washington, Seattle, Washington, USA

<sup>2</sup>Department of Ophthalmology, Puget Sound Veteran Affairs, Seattle Washington, USA

<sup>3</sup>University of Washington School of Medicine, Seattle, Washington, USA

<sup>4</sup>eScience Institute, University of Washington, Seattle, Washington, USA

\*leeay@uw.edu

**Abstract:** Evaluation of clinical images is essential for diagnosis in many specialties. Therefore the development of computer vision algorithms to help analyze biomedical images will be important. In ophthalmology, optical coherence tomography (OCT) is critical for managing retinal conditions. We developed a convolutional neural network (CNN) that detects intraretinal fluid (IRF) on OCT in a manner indistinguishable from clinicians. Using 1,289 OCT images, the CNN segmented images with a 0.911 cross-validated Dice coefficient, compared with segmentations by experts. Additionally, the agreement between experts and between experts and CNN were similar. Our results reveal that CNN can be trained to perform automated segmentations of clinically relevant image features.

© 2017 Optical Society of America

**OCIS codes:** (150.1135) Algorithms; (110.4500) Optical coherence tomography.

## References and links

1. J. B. Mitchell, "Machine learning methods in chemoinformatics," *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **4**(5), 468–481 (2014).
2. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature* **521**(7553), 436–444 (2015).
3. A. Kendall, V. Badrinarayanan, R. Cipolla, *Bayesian SegNet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding* [Internet]. arXiv [cs.CV]. 2015.
4. P. Korfiatis, T. L. Kline, and B. J. Erickson, "Automated segmentation of hyperintense regions in FLAIR MRI using deep learning," *Tomography* **2**(4), 334–340 (2016).
5. N. Hatipoglu and G. Bilgin, "Cell segmentation in histopathological images with deep learning algorithms by utilizing spatial relationships," *Med. Biol. Eng. Comput.* (2017).
6. M. D. Abrámov, Y. Lou, A. Erginay, W. Clarida, R. Amelon, J. C. Folk, and M. Niemeijer, "Improved automated detection of diabetic retinopathy on a publicly available data set through integration of deep learning," *Invest. Ophthalmol. Vis. Sci.* **57**(13), 5200–5206 (2016).
7. P. Burlina, K. D. Pacheco, N. Joshi, D. E. Freund, and N. M. Bressler, "Comparing humans and deep learning performance for grading AMD: A study in using universal deep features and transfer learning for automated AMD analysis," *Comput. Biol. Med.* **82**, 80–86 (2017).
8. R. Asaoka, H. Murata, A. Iwase, and M. Araie, "Detecting preperimetric glaucoma with standard automated perimetry using a deep learning classifier," *Ophthalmology* **123**(9), 1974–1980 (2016).
9. P. Prentašić, M. Heisler, Z. Mammo, S. Lee, A. Merkur, E. Navajas, M. F. Beg, M. Šarunic, and S. Loncaric, "Segmentation of the foveal microvasculature using deep learning networks," *J. Biomed. Opt.* **21**(7), 075008 (2016).
10. X. Gao, S. Lin, and T. Y. Wong, "Automatic feature learning to grade nuclear cataracts based on deep learning," *IEEE Trans. Biomed. Eng.* **62**(11), 2693–2701 (2015).
11. G. Coscas, J. Cunha-Vaz, and G. Soubrane, "Macular edema: definition and basic concepts," *Dev. Ophthalmol.* **47**, 1–9 (2010).
12. P. A. Keane and S. R. Sadda, "Predicting visual outcomes for macular disease using optical coherence tomography," *Saudi J. Ophthalmol.* **25**(2), 145–158 (2011).
13. A. Kumar, J. N. Sahni, A. N. Stangos, C. Campa, and S. P. Harding, "Effectiveness of ranibizumab for neovascular age-related macular degeneration using clinician-determined retreatment strategy," *Br. J. Ophthalmol.* **95**(4), 530–533 (2011).
14. D. J. Browning, A. R. Glassman, L. P. Aiello, R. W. Beck, D. M. Brown, D. S. Fong, N. M. Bressler, R. P. Danis, J. L. Kinyoun, Q. D. Nguyen, A. R. Bhavsar, J. Gottlieb, D. J. Pieramici, M. E. Rausser, R. S. Apte, J. I.

- Lim, and P. H. Miskala, Diabetic Retinopathy Clinical Research Network, "Relationship between optical coherence tomography-measured central retinal thickness and visual acuity in diabetic macular edema," *Ophthalmology* **114**(3), 525–536 (2007).
15. S. Srinivas, M. G. Nittala, A. Hariri, M. Pfau, J. Gasperini, M. Ip, and S. R. Sadda, "Quantification of intraretinal hard exudates in eyes with diabetic retinopathy by optical coherence tomography," *Retina* **1** (2017).
  16. I. Golbaz, C. Ahlers, G. Stock, C. Schütze, S. Schriefl, F. Schlanitz, C. Simader, C. Prünke, and U. M. Schmidt-Erfurth, "Quantification of the therapeutic response of intraretinal, subretinal, and subpigment epithelial compartments in exudative AMD during anti-VEGF therapy," *Invest. Ophthalmol. Vis. Sci.* **52**(3), 1599–1605 (2011).
  17. Y. Zheng, J. Sahni, C. Campa, A. N. Stangos, A. Raj, and S. P. Harding, "Computerized assessment of intraretinal and subretinal fluid regions in spectral-domain optical coherence tomography images of the retina," *Am. J. Ophthalmol.* **155**(2), 277–286 (2013).
  18. A. Lang, A. Carass, E. K. Swingle, O. Al-Louzi, P. Bhargava, S. Saidha, H. S. Ying, P. A. Calabresi, and J. L. Prince, "Automatic segmentation of microcystic macular edema in OCT," *Biomed. Opt. Express* **6**(1), 155–169 (2015).
  19. Z. Hu, G. G. Medioni, M. Hernandez, A. Hariri, X. Wu, and S. R. Sadda, "Segmentation of the geographic atrophy in spectral-domain optical coherence tomography and fundus autofluorescence images," *Invest. Ophthalmol. Vis. Sci.* **54**(13), 8375–8383 (2013).
  20. E. K. Swingle, A. Lang, A. Carass, O. Al-Louzi, S. Saidha, J. L. Prince, Segmentation of microcystic macular edema in Cirrus OCT scans with an exploratory longitudinal study, *Proc SPIE Int Soc Opt Eng.* 9417 (2015).
  21. S. J. Chiu, M. J. Allingham, P. S. Mettu, S. W. Cousins, J. A. Izatt, and S. Farsiu, "Kernel regression based segmentation of optical coherence tomography images with diabetic macular edema," *Biomed. Opt. Express* **6**(4), 1172–1194 (2015).
  22. L. Fang, D. Cuneffare, C. Wang, R. H. Guymner, S. Li, and S. Farsiu, "Automatic segmentation of nine retinal layer boundaries in OCT images of non-exudative AMD patients using deep learning and graph search," *Biomed. Opt. Express* **8**(5), 2732–2744 (2017).
  23. C. S. Lee, D. M. Baughman, and A. Y. Lee, "Deep learning is effective for classifying normal versus age-related macular degeneration optical coherence tomography images," *Ophthalmology Retina*. In press.
  24. J. Sahni, P. Stanga, D. Wong, and S. Harding, "Optical coherence tomography in photodynamic therapy for subfoveal choroidal neovascularisation secondary to age related macular degeneration: a cross sectional study," *Br. J. Ophthalmol.* **89**(3), 316–320 (2005).
  25. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (Springer, Cham; 2015). pp. 234–241.
  26. M. Esmaeili, A. M. Dehnavi, H. Rabbani, and F. Hajizadeh, "Three-dimensional segmentation of retinal cysts from spectral-domain optical coherence tomography images by the use of three-dimensional curvelet based K-SVD," *J. Med. Signals Sens.* **6**(3), 166–171 (2016).
  27. A. P. Zijdenbos, B. M. Dawant, R. A. Margolin, and A. C. Palmer, "Morphometric analysis of white matter lesions in MR images: method and validation," *IEEE Trans. Med. Imaging* **13**(4), 716–724 (1994).
  28. J. Wang, M. Zhang, A. D. Pechauer, L. Liu, T. S. Hwang, D. J. Wilson, D. Li, and Y. Jia, "Automated volumetric segmentation of retinal fluid on optical coherence tomography," *Biomed. Opt. Express* **7**(4), 1577–1589 (2016).

## 1. Introduction

Over the past decade there has been increased interest in deep learning, a promising class of machine learning models that utilizes multiple neural network layers to rapidly extract data in a nonlinear fashion [1,2]. Trained with large data sets, deep learning has been used successfully for pattern recognition, signal processing and statistical analysis [2]. Additionally, deep convolutional neural networks (CNN) have facilitated breakthroughs in image processing and segmentation [1]. These results have important clinical implications, particularly with regards to interpretation of results from medical imaging procedures.

A major advance in computer vision was the development of SegNet, a CNN algorithm that learns segmentations from a training set of segmented images and applies pixel-wise segmentations to novel images [3]. Deep learning approaches based on SegNet have been used successfully for automated segmentation in medical imaging problems such as identifying abnormalities in the brain and segmentation of histopathological cells [4,5]. Although increasing, the application of deep learning has been limited in ophthalmology to detecting retinal abnormalities from fundus photographs [6,7], glaucoma from perimetry [8], segmentation of foveal microvasculature [9], and grading of cataracts [10].

Macular edema, characterized by loss of the blood retinal barrier in retinal microvasculature, leads to accumulation of intraretinal fluid (IRF) and decreased vision [11]. With the advent of optical coherence tomography (OCT), the measurement of macular edema inferred by central retinal thickness (CRT) has become one of the most important clinical endpoints in retinal diseases [12,13]. However, the extent and severity of IRF are not fully captured by CRT [14]. Thus, the segmentation of IRF on OCT images would enable more precise quantification of macular edema.

Nevertheless, manual segmentations, which require retinal expert identification, are extremely time consuming with variable interpretation, repeatability, and interobserver agreement [15,16]. Prior studies involving automated OCT segmentations were limited by additional manual input requirements and small validation cohorts [17,18]. In addition, semi-automated methods included artifact with segmentations [19], inaccurately detected microcysts [20], and could not distinguish cystic from non-cystic fluid accumulations [21]. CNN was recently used for automated segmentation of the anatomic retinal boundaries by Fang et al., however, it has not yet been applied for automated segmentation of IRF [22]. Additionally, we demonstrated that deep learning is effective in classifying OCT images of patients with age-related macular degeneration [23]. Using the same OCT electronic medical record database, we sought to apply the CNN for automated IRF segmentations and to validate our deep learning algorithm.

## 2. Materials and methods

This study was approved by the Institutional Review Board of the University of Washington (UW) and was in adherence with the tenets of the Declaration of Helsinki and the Health Insurance Portability and Accountability Act.

Macular OCT scans from the period 2006-2016 were extracted using an automated extraction tool from the Heidelberg Spectralis imaging database at the University of Washington Ophthalmology Department. All scans were obtained using a 61-line raster macula scan, and every image of each macular OCT was extracted. Clinical variables were extracted in an automated fashion from the Epic electronic medical records database in tandem with macular OCT volumes from patients that were identified as relevant to the study, because they had received therapy for diabetic macular edema, retinal vein occlusions, and/or macular degeneration. The central slice from these OCT volumes was randomly selected for manual segmentation.

A custom browser-based application was implemented (using HTML5) and a webpage was created to passively record the paths drawn by clinicians when segmenting the images. IRF was defined as “an intraretinal hyporeflective space surrounded by reflective septae” as used in other studies [17,24]. Each segmentation result was then reviewed by an independent retina-trained clinician and inaccurate segmentations were rejected with feedback and sent for resegmentation. These images were then divided: 70% were designated as a *training set* used for optimization of the CNN. The remaining 30% were designated as a *validation set*. These were used for cross-validation during the optimization procedure. An additional final *test set* of 30 images was randomly selected, ensuring that none of the patients from the training and validation sets was present in the test set. The images in the *test set* were segmented by four independent clinicians.

Since retinal OCTs images have variable widths, a vertical section of 432 pixels tall and 32 pixels wide was chosen *a priori* with the plan to slide the window across the images and collect the predicted probabilities from the CNN. The outputted probabilities for each pixel were then averaged to create a probability distribution map. Probability distribution maps were visualized to ascertain the certainty of the model in labeling pixels as intraretinal fluid. For binary segmentation, if the averaged probability was greater than 0.5, then the pixel was labeled as having intraretinal fluid.

In the training set, data augmentation was performed by varying the position of the window on the OCT image. Care was taken in the validation set to ensure that no overlap was present in any of the validation images. In addition, the training and validation sets were balanced such that 50% of the images contained at least one pixel marked as intraretinal fluid while the other half were from images without any intraretinal fluid. Prior to the input of the model, mean value of overall intensity of the training set was subtracted and then divided by the standard deviation to normalize the input into the model.

A smoothed Dice coefficient was used as the loss function and as the assessment of the final held-out test set with the smoothness parameter set to 1.0.

$$Dice = \frac{2 \cdot TruePositives + smooth}{2 \cdot TruePositives + FalsePositives + FalseNegatives + smooth}$$

The CNN used was a modified version of the U-net autoencoder architecture [25], with a total of 18 convolutional layers (Fig. 1). The final activation function was a sigmoid function that generates a prediction of a binary segmentation map.

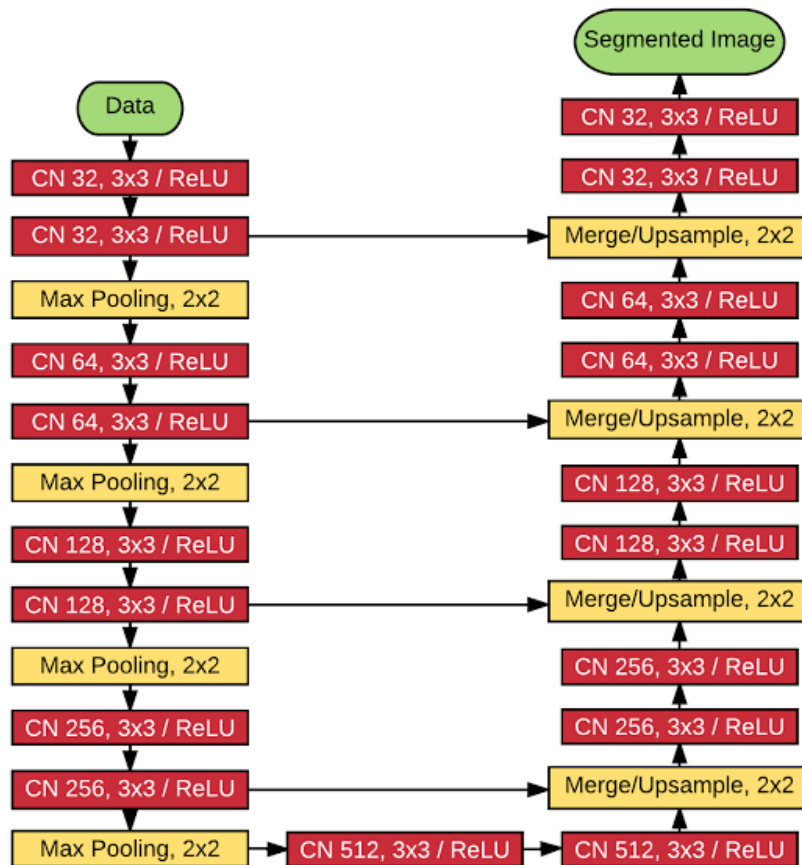


Fig. 1. Schematic of deep learning module using a total of 18 convolutional layers. CN, convolutional neural network layer. ReLU, rectified linear unit. The first number after CN refers to the number of filters and the second set of number refers to the filter sizes. The Max Pooling and Upsampling windows are also shown. Stride was set to 1 on all layers, and merge was performed using concatenation.

Training was performed on a single server with dual Intel Xeon 3.40 GHz processor, 256 GB of RAM, SSD-based RAID 6, and 8xNVIDIA P100. Batch size was set to 1,192 and parallelized across the 8 GPUs such that for each GPU, the batch was 149. All weights were initialized using a random uniform distribution. Adam optimizer for stochastic gradient descent was used set with an initial learning rate of  $1e-7$ . All training occurred using Keras (<https://keras.io>) and Tensorflow (<https://www.tensorflow.org/>) using NVIDIA Tesla P100 Graphics Processing Units (GPU) with NVIDIA cuda (v8.0) and cu-dnn (v5.5.1) libraries (<http://www.nvidia.com>). Statistics were performed using R (<http://www.r-project.org>). The model, along with the trained final weights and the code to run the model, have been published as an open source project: (<https://github.com/uw-biomedical-ml/irf-segmenter>).

In assessing the final held-out test set, the CNN performed inference using a sliding window generating an array of probabilities for each pixel. Then this output was compared using the above Dice coefficient to baseline segmentations performed by clinicians.

In addition, the CNN was compared against four clinicians. For this comparison, an individual clinician was chosen as the ground truth, and the Dice coefficient of the other three clinicians as well as the CNN were computed from this groundtruth. This was repeated for all four clinicians so that Dice coefficients with each clinician as baseline were computed. The coupled Dice coefficients showed no statistical difference between the clinicians and the CNN.

### 3. Results

Manual segmentations of 1,289 OCT macular images were used for training and cross validation. A total of 1,919,680 sections of 934 manually segmented OCT images were designated as a training set: these were presented to the CNN, and the segmentation errors from these images were used to adjust the weights in different layers of the networks. Another 1,088 sections of 355 images were used for cross-validation: they were presented to the CNN during training and segmentation errors were recorded, but were not used to adjust the weights in the network layers

The model was trained with 200,000 iterations and periodically assessed against the validation set (cross-validation). Training time was 7.25 days using parallelized training across 8 x NVIDIA P100 GPUs. Inference time on the test set was 13.2 seconds per OCT B scan image with a sliding window using a single NVIDIA P100 GPU.

Final trained parameters of the first convolutional layer are shown in Fig. 2. The learning curve showing the Dice coefficients of the training iterations and the validation set are shown in Fig. 3. The model achieved a maximal cross-validation Dice coefficient of 0.911.




Fig. 2. Visualization of trained weights of the first layer of convolutional filters.




Fig. 3. Learning curve of the neural network training and validation.

The weights that produced the highest cross-validated coefficient were then used for comparison against manual segmentations in a held out test set. Example deep learning segmentations of intraretinal fluid compared with manual segmentations from the held out test set are shown in Fig. 4.




Fig. 4. Example segmentations of intraretinal fluid (IRF) from the held out test set by deep learning. A, an example optical coherence tomography (OCT) image with intraretinal fluid. B, manual segmentation of IRF C, automated segmentation by deep learning. D, an example optical coherence tomography (OCT) image with intraretinal fluid. E, manual segmentation of IRF. F, automated segmentation by deep learning. G, an example OCT image with IRF and pigment epithelial detachment (PED). H, manual segmentation of IRF avoiding PED. I, deep learning correctly segments IRF cysts but not PED. J, an example optical coherence tomography (OCT) image with IRF and shadowing under vasculature. K, manual segmentation of IRF. L, deep learning correctly segments IRF cysts but avoids the shadowing under the retinal vasculature.

To assess reliability of the deep learning algorithm we compared the Dice coefficients for deep learning automated segmentations against manual segmentations by four independent clinicians. The pairwise comparison of Dice coefficients are shown in Table 1.

**Table 1. Mean dice coefficients with standard deviations.**

	Clinician 2	Clinician 3	Clinician 4	Deep Learning
Clinician 1	0.744 (0.185)	0.771 (0.155)	0.809 (0.164)	0.754 (0.135)
Clinician 2		0.710 (0.190)	0.747 (0.189)	0.703 (0.182)
Clinician 3			0.722 (0.209)	0.725 (0.117)
Clinician 4				0.730 (0.176)

The clinician in the first column is the baseline from whom Dice coefficients are computed for each row.

The mean Dice coefficient for human interrater reliability and deep learning were 0.750 and 0.729, respectively, reaching a generally accepted value of excellent agreement (Fig. 5) [26]. No statistically significant difference was found between clinicians and deep learning ( $p = 0.247$ ).

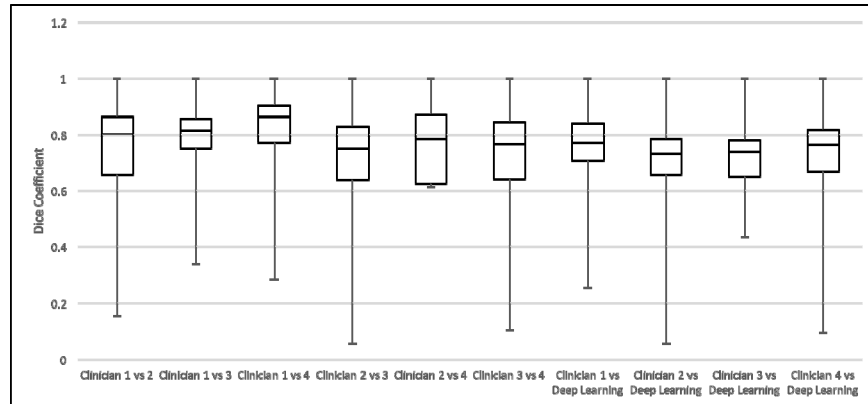


Fig. 5. Pairwise comparison of Dice coefficients from the held out test set. The mean Dice coefficient for human interrater reliability and deep learning were 0.750 and 0.729, respectively.

#### 4. Discussion

The concept of automated segmentation using deep learning has become increasingly popular in medical imaging; however, comparatively little has been applied in ophthalmology and the validation of proposed models were limited to small sample size. Zheng et al. evaluated semi-automated segmentation of intra- and subretinal fluid in OCT images of 37 patients [17]. Initial automated segmentation was performed followed by coarse and fine segmentation of fluid regions. Then an expert manually selected each region of potential IRF and SRF. This semi-automated segmentation had a Dice coefficient of 0.721 to 0.785 compared to manual segmentation, which is similar to our study. The advantage of our algorithm includes being a fully automated system with no human input. Esmaeili et al. used recursive Gaussian filter and 3D fast curvelet transform to determine curvelet coefficients, which were denoised with a 3D sparse dictionary. The average Dice coefficients for the segmentation of intraretinal cysts in whole 3D volume and within central 3 mm diameter were 0.65 and 0.77, respectively [27]. However, the validation of their study included only four 3D-OCT images, a much smaller sample than our study. Chiu et al. reported a dice coefficient of 0.53 using an automated detection of IRF in OCT images of 6 patients with diabetic macular edema. Both Dice coefficient and intermanual segmenter correlations (0.58) were much lower than our study [21]. Wang et al. recently applied fuzzy level set for automated segmentation of IRF with results similar to human graders [28]. However, their study consisted of only 10 patients and was limited by segmentation artifacts requiring integration with OCT angiography for verification. A benefit of our method is that it does not require OCT angiography, which is not readily available in most practices.

Our study shows that CNN can be trained to identify clinical features that are diagnostic, similar to how clinicians interpret an OCT image. The deep learning algorithm successfully identified areas of IRF by generating pixelwise probabilities which were then averaged to create a segmentation map. Our model was trained with more than 200,000 iterations without overfitting. Additional advantages of this study include the use of a publicly available model and training strategy for precise localization and segmentation of images. Furthermore, the deep learning model is fast and requires no additional human input, such as with semi-automated segmentation programs. To our knowledge, the use of a validated CNN for accurate segmentation of IRF has never been shown.



Our training data sets originated from images from a single academic center, thus our results may not be generalizable. However, our training data set was randomly selected from all consecutive retinal OCT images obtained at our institution during a 10-year period, therefore a wide variety of pathologies likely have been included, increasing its overall applicability in other data sets. In summary, deep learning shows promising results in accurate segmentation and automated quantification of intraretinal fluid volume on OCT imaging and our open-source algorithm may be useful for future clinical and research applications

**Funding**

National Eye Institute (NEI) (K23EY02492); Latham Vision Science Innovation Grant; Research to Prevent Blindness; The Gordon & Betty Moore Foundation; Alfred P. Sloan Foundation.

**Disclosures**

The authors declare that there are no conflicts of interest related to this article.

The contents do not represent the views of the U.S. Department of Veterans Affairs or the United States Government.

**Acknowledgments**

This work was supported by Research to Prevent Blindness. The hardware used in the completion of this project was donated by NVIDIA.