



Published in final edited form as:

JAMA Ophthalmol. 2020 October 01; 138(10): 1077–1078. doi:10.1001/jamaophthalmol.2020.3256.

Novel Low-Shot Deep Learning Approach for Retinal Image Classification with Few Examples

Matthew S. Hunt, BS^{1,2}, Yuka Kihara, MS¹, Aaron Y. Lee, MD MSCI¹

¹Department of Ophthalmology, School of Medicine, University of Washington, Seattle, WA

²Washington University, School of Medicine, St. Louis MO

Artificial intelligence (AI) stands to increasingly impact the practice of medicine as advances in deep learning for computer vision and natural language processing are translated to medical contexts. Applications of deep learning to retinal imaging have shown promise in recent years, beginning with a 2016 study by authors from Google demonstrating near-ophthalmologist accuracy in the classification of diabetic retinopathy from fundus photographs.¹ However, the challenges facing AI in retinal imaging parallel the challenges with applying AI to medicine in general. Many of these problems arise due to limitations in data. Large amounts of annotated data are traditionally required to “train” deep-learning models to perform well. Not only can high-volume data annotation be costly (the aforementioned diabetic retinopathy study required around 527 000 fundus-image evaluations by ophthalmologists), but sometimes sufficient data cannot be obtained, as may be the case for rare conditions. Additionally, an unbalanced distribution of characteristics, such as patient demographics, in the training data can result in poor performance in underrepresented populations. This new study in this issue of *JAMA Ophthalmology* by Burlina et al.² represents a key step toward combatting these challenges by applying cutting-edge, self-supervised, “low-shot” learning methods to improve performance in retina image classification under conditions of few training examples.

The authors studied the influence of the amount of training data on AI model performance in the classification of diabetic retinopathy from fundus photographs. They trained several different types of models on subsets of training data from a publicly available fundus image database. The images were categorized into two classes: either having referable diabetic retinopathy or not. Independent models of each type were trained using subsets of the database increasing in size from 10 images per class to 5120 images per class. All models were evaluated on the primary metrics of accuracy and area under the receiver operating characteristic curve (AUC) on the same test set taken from the same database. The models included a standard convolutional neural network (CNN) as a control (ResNet50), along with several different low-shot approaches (both standard and self-supervised) based on recent developments in AI research. As expected, all models tended to perform better with increased amounts of training data. At the maximum amount of training data (5120 images per class), the control CNN achieved an AUC of 0.8330, and the comparable self-

supervised, low-shot method achieved AUC of 0.8348. As the amount of training data decreased, the low-shot methods outperformed the standard control CNN, with the self-supervised approaches performing best.

Deep learning mitigates the need for human-designed features used in traditional image-processing techniques and, instead, allows the model to devise its own domain-specific features and correlate those features to classification based on exposure to training data. Because deep-learning models are highly flexible in terms of the features they can potentially learn, a high volume of training data is typically needed for a model to refine these features and make them relevant to the data the model will encounter at test time.

The low-shot methods presented in the article by Burlina et al² reduce the number of labeled training examples required to learn relevant features by restricting the model hypothesis (feature) space and leveraging internal associations within the data to enhance representations. The standard low-shot methods evaluated in this study combined “pretrained” CNNs with traditional classifiers (random forest, support vector machine, and K-nearest-neighbors), which help the model avoid overfitting at the fully connected layer. Second, the self-supervised approaches maximize the use of smaller amounts of labeled training data by learning about statistical regularities (eg. spatial relationships and orientation of images) that might be a cue to semantics of the underlying data. Self-supervised approaches use variations on the theme of using a portion of the data to predict something about another correlated portion of the data. The self-supervised approach used here (extension of Deep InfoMax)³ compares a model’s deepest representation of an entire transformed image with shallower representations of differently transformed portions of the same and different images. It requires no additional expert-provided labels, and in the process of analyzing these comparisons, it captures useful information about the data. The key contribution of the study by Burlina et al² was putting these low-shot methods to the test in the retinal-imaging domain, an important milestone in the translation of the latest deep-learning techniques into ophthalmology.

This study has a number of strengths. The authors properly implemented a recently developed self-supervised learning algorithm that uses general knowledge from intranet images, and they demonstrated that this knowledge can be usefully transferred from natural images to retinal images. Several low-shot classification methods were evaluated against a control CNN, and full experimental results were presented in detail along with 95% CIs. Another major strength of this study is its method, which included using a large, open-source database, partitioning images at the patient level, and balancing classes across all experiments. This method strengthens the finding that the self-supervised approaches exceeded the performance of the industry standard in the low-data regime. These results lay the foundation for the use of low-shot learning in retinal imaging, but several important questions remain.

One limitation of the study stems from its limited range of training data sizes. In the 2016 diabetic retinopathy study, the low-training-data threshold for optimal model performance was around 60 000 images.¹ At best performance, the AUCs in that study were around 0.99. Given that this study used, at most, 10 240 training examples, it is expected that the AUCs

would be worse than the 2016 study. However, testing lower amounts of training data limited the findings in 2 main ways. Although the superiority of low-shot methods in the very-low-data regime shows promise for certain applications, it is not clear whether these results are acceptable for clinical use in domains where sufficient data are available without further studies. Second, the restricted representational capacity of some low-shot methods may limit their performance compared with traditional methods in the untested high-data regime. The self-supervised, low-shot methods may be most useful in addressing the issue of training data set imbalances. Future work could evaluate their performance as the training data set becomes more imbalanced with respect to disease category or patient characteristics.

Low-shot learning is particularly beneficial to the success of medical AI when studying uncommon diseases or requiring high-quality, clinician-derived labels, as opposed to training with functional targets.^{4–6} Burlina et al² have accomplished a critical step toward realizing the benefits of low-shot learning by demonstrating the advantages of self-supervised methods in low-shot retinal imaging. Future research could quantify where these approaches can be most clinically relevant and evaluate whether they can address the important issues associated with data set bias.

Financial Support:

NIH/NEI K23EY029246, CDA from Research to Prevent Blindness (RPB). The sponsors/funding organizations had no role in the design or conduct of this research.

NIH/NEI K23EY029246, CDA from Research to Prevent Blindness (RPB). The sponsors/funding organizations had no role in the design or conduct of this research.

Dr. A. Lee reports support from the US Food and Drug Administration, grants from Santen, Carl Zeiss Meditec, and Novartis, personal fees from Genentech, Topcon, and Verana Health, outside of the submitted work; This article does not reflect the opinions of the Food and Drug Administration.

Dr. A. Lee reports support from the US Food and Drug Administration, grants from Santen, Carl Zeiss Meditec, and Novartis, personal fees from Genentech, Topcon, and Verana Health, outside of the submitted work; This article does not reflect the opinions of the Food and Drug Administration.

References

1. Gulshan V, Peng L, Coram M, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA*. 2016;316(22):2402–2410. [PubMed: 27898976]
2. Burlina P, Paul W, Mathew P, Joshi N, Pacheco K, Bressler N. Low-Shot Deep Learning of Diabetic Retinopathy with Potential Applications to Address Artificial Intelligence Bias in Retinal Diagnostics and Rare Ophthalmic Diseases. *JAMA Ophthalmol*. In Press.
3. Bachman P, Hjelm RD, Buchwalter W. Learning Representations by Maximizing Mutual Information Across Views. In: Wallach H, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox E, Garnett R, eds. *Adv Neural Inf Process Syst* 32. Curran Associates, Inc.; 2019:15535–15545.
4. Kihara Y, Heeren TFC, Lee CS, et al. Estimating Retinal Sensitivity Using Optical Coherence Tomography With Deep-Learning Algorithms in Macular Telangiectasia Type 2. *JAMA Netw Open*. 2019;2(2):e188029. [PubMed: 30735236]
5. Wen JC, Lee CS, Keane PA, et al. Forecasting future Humphrey Visual Fields using deep learning. *PLoS One*. 2019;14(4):e0214875. [PubMed: 30951547]
6. Lee CS, Tyring AJ, Wu Y, et al. Generating retinal flow maps from structural optical coherence tomography with artificial intelligence. *Sci Rep*. 2019;9(1):5694. [PubMed: 30952891]