



Published in final edited form as:

Ophthalmology. 2019 November ; 126(11): 1533–1540. doi:10.1016/j.ophtha.2019.06.005.

A deep learning approach for automated detection of geographic atrophy from color fundus photographs

Tiarnan D. Keenan, BM BCh, PhD^{1,*}, Shazia Dharssi^{1,2,*}, Yifan Peng, PhD^{2,*}, Qingyu Chen, PhD², Elvira Agrón, MA¹, Wai T. Wong, MD, PhD^{1,3}, Zhiyong Lu, PhD², Emily Y. Chew, MD¹

¹Division of Epidemiology and Clinical Applications, National Eye Institute (NEI), National Institutes of Health (NIH), Bethesda, MD, USA;

²National Center for Biotechnology Information (NCBI), National Library of Medicine (NLM), National Institutes of Health (NIH), Bethesda, MD, USA;

³Unit on Microglia, National Eye Institute (NEI), National Institutes of Health, Bethesda, MD, USA

Abstract

Purpose: To assess the utility of deep learning in the detection of geographic atrophy (GA) from color fundus photographs; secondary aim to explore potential utility in detecting central GA (CGA).

Design: A deep learning model was developed to detect the presence of GA in color fundus photographs, and two additional models to detect CGA in different scenarios.

Participants: 59,812 color fundus photographs from longitudinal follow up of 4,582 participants in the Age-Related Eye Disease Study (AREDS) dataset. Gold standard labels were from human expert reading center graders using a standardized protocol.

Methods: A deep learning model was trained to use color fundus photographs to predict GA presence from a population of eyes with no AMD to advanced AMD. A second model was trained to predict CGA presence from the same population. A third model was trained to predict CGA presence from the subset of eyes with GA. For training and testing, 5-fold cross-validation was employed. For comparison with human clinician performance, model performance was compared with that of 88 retinal specialists.

Corresponding Author(s): Emily Y. Chew, MD, National Eye Institute (NEI), National Institutes of Health (NIH), 9000 Rockville Pike, Bethesda, Maryland, United States, echew@nei.nih.gov, 301.496.6583. Zhiyong Lu, PhD, National Center for Biotechnology Information (NCBI), National Library of Medicine (NLM), National Institutes of Health (NIH), 9000 Rockville Pike, Bethesda, Maryland, United States, zhiyong.lu@nih.gov, 301.594.7089.

Conception and design: Keenan, Peng, Dharssi, Chew, Lu

Analysis and interpretation: Keenan, Peng, Dharssi, Chen, Agron, Wong, Chew, Lu

Data collection: Peng, Dharssi, Agrón, Chew, Lu

Drafting the work: Keenan, Peng, Dharssi, Chew, Lu

*These authors contributed equally to this work.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Conflict of Interest: no conflicting relationship exists for any author

Online Supplemental Materials: This article contains additional online-only material.

Main Outcome Measures: Area under the curve (AUC), accuracy, sensitivity, specificity, precision.

Results: The deep learning models (GA detection, CGA detection from all eyes, and centrality detection from GA eyes) had AUC of 0.933–0.976, 0.939–0.976, and 0.827–0.888, respectively. The GA detection model had accuracy, sensitivity, specificity, and precision of 0.965 (95% CI 0.959–0.971), 0.692 (0.560–0.825), 0.978 (0.970–0.985), and 0.584 (0.491–0.676), respectively, compared to 0.975 (0.971–0.980), 0.588 (0.468–0.707), 0.982 (0.978–0.985), and 0.368 (0.230–0.505) for the retinal specialists. The CGA detection model had equivalent values of 0.966 (0.957–0.975), 0.763 (0.641–0.885), 0.971 (0.960–0.982), and 0.394 (0.341–0.448), compared to 0.990 (0.987–0.993), 0.448 (0.255–0.641), 0.993 (0.989–0.996), and 0.296 (0.115–0.477). The centrality detection model had equivalent values of 0.762 (0.725–0.799), 0.782 (0.618–0.945), 0.729 (0.543–0.916), and 0.799 (0.710–0.888), compared to 0.735 (0.445–1), 0.878 (0.722–1), 0.703 (0.332–1), and 0.626 (0.273–0.979).

Conclusions: A deep learning model demonstrated high accuracy for the automated detection of GA. The AUC was non-inferior to that of human retinal specialists. Deep learning approaches may also be applied to the identification of CGA. The code and pretrained models are publicly available at <https://github.com/ncbi-nlp/DeepSeeNet>.

Précis

A deep learning model developed to automatically to detect geographic atrophy associated with age-related macular degeneration from color fundus images achieved high accuracy and was non inferior to determination by retinal specialists.

Taxonomy topics

deep learning; geographic atrophy; central geographic atrophy; Age-Related Eye Disease Study (AREDS); artificial intelligence (AI)

Introduction

Age-related macular degeneration (AMD) is the leading cause of irreversible vision loss in individuals older than 60 years of age in Western countries, and accounts for approximately 9% of blindness worldwide^{1–3}. Late AMD takes two forms, neovascular and atrophic, though these may coexist in the same eye. Geographic atrophy (GA) is the defining lesion of the atrophic form of late AMD. GA in AMD has been estimated to affect over 5 million people worldwide^{2,4}. Unlike for neovascular AMD, no drug therapies are available to prevent GA, slow its enlargement, or restore lost vision; this makes it an important research priority^{5,6}. Rapid and accurate identification of eyes with GA could lead to improved recruitment of eligible patients for future clinical trials and eventually to early identification of appropriate patients for proven treatments.

Since the original description of GA by Gass⁷, clinical definitions have varied between research groups⁸. In the Age-Related Eye Disease Study (AREDS), it was defined as a sharply demarcated, usually circular zone of partial or complete depigmentation of the retinal pigment epithelium (RPE), typically with exposure of underlying large choroidal

blood vessels, at least as large as grading circle I-1 (1/8 disc diameter in diameter)⁹. Sensitivity of the retina to light stimuli is markedly decreased (i.e., dense scotomata) in areas affected by GA. The natural history of the disease involves progressive enlargement of GA lesions over time, with visual acuity decreasing markedly as the central macula becomes involved⁵.

The identification of GA by ophthalmologists conducting dilated fundus examinations is sometimes challenging. This may be particularly true for cases with early GA, comprising smaller lesions with less extensive RPE atrophy. In addition, the increasing prevalence of GA (through aging populations in many countries) will translate to greater demand for retinal services. As such, deep learning approaches involving retinal images, obtained perhaps using telemedicine-based devices, might support GA detection and diagnosis. However, these approaches would require the establishment of evidence-based and ‘explainable’ systems that have undergone extensive validation and demonstrated performance metrics that are at least non-inferior to those of clinical ophthalmologists in routine practice.

Recent studies have demonstrated the utility of deep learning in the field of ophthalmology^{10–18}. Deep learning is a branch of machine learning that allows computers to learn by example. As applied to image analysis, deep learning enables computers to perform classification directly (unsupervised) from images, rather than through the recognition of features prespecified by human experts. This is achieved by training algorithmic models on images with accompanying labels (e.g. color fundus photographs categorized manually for the presence or absence of GA), such that these models can then be used to classify new images with similar labels. The models are neural networks that are constructed of an input layer (which receives, for example, the color fundus photograph), followed by multiple layers of nonlinear transformations, to produce an output (e.g. GA present or absent). One recent deep learning model, DeepSeeNet¹⁷, performed patient-based AMD severity classification with a level of accuracy higher than a group of human retinal specialists. Indeed, it also did this in an ‘explainable’ way, by simulating the human grading process, which may improve levels of acceptability to ophthalmologists and patients. However, to the best of our knowledge, few studies have focused specifically on GA^{19–22}, and the majority of these have concentrated on image segmentation tasks rather than automated detection of GA.

The primary aim of this study was to assess the utility of deep learning for the detection of GA from color fundus images. We also conducted experiments to assess the potential utility of deep learning for the detection of central GA (CGA) from color fundus images.

Methods

Dataset

The dataset used for this study (training and testing) was from the AREDS. The AREDS was a multi-center, prospective cohort study of the clinical course of AMD (and age-related cataract), as well as a phase III randomized clinical trial of nutritional supplementation for treatment of AMD and cataract²³. Its primary outcome was the development of advanced

AMD, defined as CGA or neovascular AMD. The study design has been described previously²³. In short, 4,757 participants aged 55 to 80 years were recruited between 1992 and 1998 at 11 retinal specialty clinics in the USA. Based on color fundus photographs, best-corrected visual acuity, and ophthalmologic evaluations, participants were enrolled into one of several AMD categories. Institutional review board approval was obtained at each clinical site and written informed consent for the research was obtained from all study participants. The research was conducted under the Declaration of Helsinki.

As described previously, at baseline and annual study visits, comprehensive eye examinations were performed by certified study personnel using standardized protocols. At the baseline, 2-year, and annual study visits thereafter, stereoscopic color fundus photographs were taken of both eyes (field 2, i.e., 30° imaging field centered at the fovea). Because of the inherent redundancy in a pair of stereoscopic photographs, for each eye, only one of the pair of photographs was used in the current study. In general, the left image of the pair was used, unless it was missing from the database, in which case the right image was used instead (~ 0.5%). As a result, a total of 59,812 color fundus images from 4,582 participants were extracted from the AREDS dataset (Table 1).

The ground truth labels used for both training and testing were the grades previously assigned to each color fundus photograph by expert human graders at the University of Wisconsin Fundus Photograph Reading Center. (These grading experts did not overlap at all with the 88 retinal specialists described elsewhere). The Reading Center workflow is described in detail in AREDS Report number 6⁹. In brief, a senior grader (grader 1) performed initial grading of the photograph for AMD severity using a 4-step scale (where step 3 includes non-central GA and step 4 includes CGA), and a junior grader (grader 2) performed detailed grading for multiple AMD-specific features (including CGA and non-central GA). In the case of discrepancy between the two graders, a senior investigator at the Reading Center would adjudicate the final grade. All photographs were graded independently, that is, graders were masked to the photographs and grades from previous visits. In addition, a rigorous process of grading quality control was performed at the Reading Center, including the assessment for the inter-grader and intra-grader agreement overall and according to specific AMD features²⁴. Inter-grader agreement was substantial for CGA (kappa 0.73). Analyses for potential “temporal drift” were conducted by having all graders re-grade in a masked fashion the same group of images annually for the duration of the study. The presence of GA was defined as a sharply demarcated, usually circular, zone of partial or complete depigmentation of the RPE, typically with exposure of underlying large choroidal blood vessels, at least as large as grading circle I-1 (1/8 disc diameter in diameter)⁹. If present, GA was categorized as central or non-central according to involvement of the ‘center point’ (Figure 1). In this study, images where GA was present alongside signs of neovascular AMD (i.e., image graded positive for both GA and neovascular AMD) were excluded from the dataset.

Models

The deep learning models for predicting GA and CGA were built using the methods described in DeepSeeNet¹⁷. Specifically, three separate deep learning models were built: (i)

one ('GA model') for identifying GA from a population of eyes with no AMD to advanced AMD, irrespective of central involvement, (ii) one ('CGA model') for identifying CGA from a population of eyes with no AMD to advanced AMD, and (iii) one ('centrality detector model') for identifying CGA from a population of eyes with GA (i.e., a model that assesses for central involvement, when GA presence is already known). All three models were binary classification models. As shown in Table 1, there were 2,585 positive instances for the former (GA model) and 1,455 for the latter two (CGA model and centrality detector model). The convolutional neural network architecture used for all models was Inception-v3²⁵, which is a state-of-the-art convolutional neural network for image classification. It contains 317 layers, comprising a total of over 21 million weights (learnable parameters) that are subject to training.

The three models were pre-trained using ImageNet, an image database of over 14 million natural images with corresponding labels, using methods described previously¹⁷. (This very large dataset is often used in deep learning to pre-train models. In a process known as transfer learning, pre-training on ImageNet is used to initialize the layers/weights, leading to the recognition of primitive features (e.g., edge detection), prior to subsequent training on the dataset of interest). During subsequent training, for all three models, the input to the model was a color fundus photograph cropped to a square and scaled to 512 pixels. Cropping was done by decreasing the longer dimension of the image to match the shorter dimension (keeping the center point unchanged). We also applied a Gaussian filter to normalize the color balance^{26,27}. We trained the deep learning models using two commonly used libraries: Keras²⁸ and TensorFlow²⁹. During the training process, we updated the model parameters using the Adam optimizer (learning rate of 0.0001) for every minibatch of 32 images. Model convergence was measured when the loss on the development set started to increase. The training was stopped after 5 epochs (passes of the entire training set) once the accuracy values no longer increased or started to decrease. All experiments were conducted on a server with 32 Intel Xeon CPUs, using a NVIDIA GeForce GTX 1080 Ti 11Gb GPU for training and testing, with 512Gb available in RAM memory.

Evaluation

For training and testing, we used the standard machine learning approach of five-fold cross-validation. This consists of subdividing the total data into five equally-sized folds. We used three folds for training, one fold for development (to optimize the hyperparameters and measure the convergence), and the remaining one fold for testing. This process was repeated five times, reserving a different fold each time as the testing set. These settings avoid overfitting the test set, as well as having the test set large enough to be representative. The cross-validation splitting (into the five folds) was performed randomly, at the participant level (i.e., without stratification by disease status). All images from a single participant were present in only one fold. In addition, image augmentation procedures were used, as follows, in order to increase the dataset artificially (i.e., using additional synthetically modified data): (i) rotation, (ii) width shift up to 0.1, (iii) height shift up to 0.1, (iv) horizontal flip, and (v) vertical flip.

Each model was evaluated against the gold standard Reading Center grades. We obtained the following metrics: overall accuracy, sensitivity, specificity, and precision (with 95% confidence intervals). Outcomes were also evaluated using receiver operating characteristic (ROC) curves and their corresponding area under the curve (AUC).

The performance of the models was compared with the performance of 88 human retinal specialists (in terms of their pooled gradings). The retinal specialists had graded all color fundus photographs (as part of a qualification survey used to determine AMD severity) at AREDS baseline; the grading was non-overlapping (i.e., each image was graded by one retinal specialist only) and was independent of the Reading Center grading, which was used as the ground truth.

Results

The ROC curves for the deep learning models are shown in Figure 2. For each model, five ROC curves are displayed on the same graph, with each curve representing one of the folds tested (in the five-fold cross-validation procedure). For comparison, the performance of the human retinal specialists (in terms of their pooled gradings) on the same fold is shown as a single point, such that one point accompanies each of the five curves. For the GA model, the five AUC values were 0.933, 0.952, 0.962, 0.964, and 0.976. For the CGA model, the five AUC values were 0.939, 0.941, 0.955, 0.968, and 0.976. For the centrality detector model, the five AUC values were 0.827, 0.833, 0.864, 0.884, and 0.888. Based on the ROC analysis, the performance levels of the GA model were slightly superior to those of the retinal specialists; those of the CGA model were slightly superior or similar, and those of the centrality detector model were generally inferior.

The accuracy, sensitivity, specificity, and precision of the three models are shown in Table 2. The GA model demonstrated high accuracy (0.965, 95% CI 0.959–0.971), attaining sensitivity of 0.692 (0.560–0.825) and specificity of 0.978 (0.970–0.985). The sensitivity was substantially higher and the specificity slightly lower than those of the human retinal specialists. The CGA model also demonstrated high accuracy (0.966, 0.957–0.975), with sensitivity of 0.763 (0.641–0.885) and specificity of 0.971 (0.960–0.982). Again, the sensitivity was substantially higher and the specificity slightly lower than those of the retinal specialists. The centrality detector model had accuracy of 0.762 (0.725–0.799), sensitivity of 0.782 (0.618–0.945), and specificity of 0.729 (0.543–0.916). The accuracy and specificity were slightly higher than those of the retinal specialists, while the sensitivity was substantially lower. For all three models, the precision was higher than that of the retinal specialists.

Discussion

The deep learning model showed relatively robust performance for the detection of GA from a population of eyes with a wide spectrum of disease, from no AMD to advanced AMD. The ROC analysis demonstrated that its performance was non-inferior to that of human retinal specialists. The results of this study highlight the potential utility of deep learning models in identifying GA, based on simple color fundus photographs without additional imaging

modalities or other information. In addition, potential utility was demonstrated for deep learning-based detection of CGA in two different scenarios, i.e., from a population of eyes with a wide spectrum of disease and from a population of eyes already known to have GA. Each of these scenarios might be relevant in different clinical or research settings.

To probe further the performance of the GA model that identifies GA presence/absence, particularly in light of lower sensitivity than specificity, we performed an error analysis of false negative cases according to GA area. The hypothesis was that the model performance may have been superior for large GA lesions and inferior for small GA lesions, such that missed cases with small GA lesions might help explain suboptimal model sensitivity. As part of the original Reading Center grading, GA area (within the AREDS grid) was quantified into one of seven categories using the AREDS grading circles⁹. We analyzed the proportions of false negatives (GA cases missed by the model) according to the seven area categories. Table 3 demonstrates that questionable and small GA lesions had a much higher false negative rate (62.9% for questionable and 78.9% for area < circle I-2), whereas large GA lesions had a much lower false negative rate (19.8% for area \geq 2 DA), with evidence of a dose-response effect. Hence, the model's performance had (perhaps understandably) lower sensitivity for small and questionable GA lesions. One contributing factor to this is likely to be the relatively low number of instances during training, since the smaller GA categories accounted for low proportions of the training set.

To probe further the performance of the CGA model that identifies CGA presence/absence, again in light of lower sensitivity than specificity, we performed an error analysis of false negative cases according to the two categories shown in Figure 1 (i.e., 'center point with definite GA' versus 'center point with questionable GA but central subfield with definite GA'). As before, we analyzed the proportions of false negatives (CGA cases missed by the model) according to the two categories (Table 4). As expected, the latter category had a higher false negative rate (31.6%) than the former category (21.5%). Hence, the model's performance had (perhaps understandably) lower sensitivity for cases with questionable involvement of the center point.

In future studies, possible approaches to improve the false negative rate in images with small GA lesions might include the use of additional datasets, image augmentation, and generation of synthetic images. In particular, further training on the AREDS2 dataset is likely to improve this aspect of performance, since the dataset includes 1616 eyes with GA (of which 1118 had GA area <0.75 disc areas at first appearance)⁵. In the near future, deep learning might also be able to generate synthetic retinal images with GA (where small lesions could be specified), for improved model training³⁰.

In addition, for the false negative cases, we analyzed the fundus photographs and their accompanying saliency maps (explained in the Supplementary information) for a random sample of 20 images. We graded the photographs qualitatively using pre-specified criteria: (i) image quality, (ii) GA size, (iii) degree of depigmentation in GA lesion(s), and (iv) other factors. Image quality was extremely poor in 5/20 and moderately poor in 10/20. GA size was very small in 4/20 and extremely large (occupying almost all of the photograph) in 2/20. Degree of depigmentation was very mild in 3/20. The most common other factor (in 11/20)

was GA lesion features suggestive of prior neovascular AMD (even though the Reading Center grading for these images was negative for neovascular AMD); indeed, examination of previous Reading Center gradings for the same eyes confirmed that, in 10 of these eyes, neovascular AMD had been graded positive at one or more previous study visits. Hence, one or more potential reasons for a false negative grading (very poor image quality, very small GA size, very mild depigmentation, or features of previous neovascular AMD) was found in 19/20 cases. The remaining case had a combination of moderately poor image quality and moderately small GA lesion size. In addition, in 5/20 cases, the saliency maps demonstrated increased signal concentrated in the GA lesion(s), despite the negative prediction.

We consider that these models, and future models arising from them, may be useful in several different scenarios, including assistance in (i) clinical care, (ii) recruitment for clinical trials, (iii) population-based screening for clinical care, and (iv) population-based screening for epidemiological research. Notably, GA has been estimated to affect over 5 million people worldwide^{2,4}, but it is likely that many of these individuals remain undiagnosed, especially in countries with lower levels of access to retinal specialists. As regards clinical care, the ongoing incorporation of deep learning into radiology and pathology is through augmentation (rather than replacement) of human diagnosis. In ophthalmic care, it might be helpful for an optometrist, general ophthalmologist, or even retinal specialist to have software like this available alongside traditional clinical tools. The software would assist human diagnosis and decision-making, rather than replacing it. For example, it could be used (with settings selected for high sensitivity) by an optometrist, alongside normal examination procedures. Deep learning-assisted GA detection could lead to referral for consultation with a retinal specialist. This approach might also lend itself well to telemedicine approaches, particularly in remote areas. Ultimately, an integrated suite of tools like this might be helpful for augmented human evaluation of a wide range of ophthalmic pathologies, such as AMD, diabetic retinopathy, and glaucoma. However, it would be very important to perform prospective studies to compare the performance of clinicians with and without assistance from deep learning systems.

As to assistance in recruitment for clinical trials, the number of clinical trials of drugs to slow GA enlargement is likely to increase over the next decade. Models like these would enable rapid screening of large numbers of individuals, to identify potential qualifying participants. If treatments that successfully slow GA enlargement are identified and approved, the identification of patients with GA who may stand to benefit from treatment may be aided by these models. Finally, as regards assistance in population-based screening for epidemiological research, accurate data on GA prevalence and characteristics are important for understanding AMD pathophysiology and planning service provision. However, few modern population-based studies have these data readily available. If models like this were used for rapid population-based screening, large quantities of epidemiological data could be produced. For example, the UK Biobank study contains color fundus photographs on 68,544 adult participants (with detailed accompanying data on medical history, lifestyle, physical measures, and blood samples). Applying this model to the dataset would generate a large volume of data on GA prevalence, with powerful links to the medical characteristics of the affected participants. Similar approaches might be possible in the future, e.g. using the Million Veterans Program and others.

There are several limitations to this study. First, the proportions of positive test cases (i.e., images with GA and CGA) were relatively low. This limitation may be reflected in the higher rates of specificity than sensitivity for the first two models. Second, model training and testing were carried out (necessarily) in the context of a high prevalence of AMD, which is a controlled clinical trial rather than a real-world scenario. The ideal setting for evaluating model performance in real-world scenarios (i.e., within a realistic background of retinal diseases and findings) would be a population-based dataset, in which the presence of GA and of other retinal findings had been rigorously graded by human reading center graders. However, the AREDS participants did encounter ocular and systemic disease typical for an elderly population and were not excluded from follow-up for these reasons. Therefore, many AREDS images contain signs of diabetic retinopathy and other pathology (though this was towards the milder end of the disease spectrum, particularly in earlier study years). Hence, model training and testing were performed in the presence of co-existing retinal pathology, perhaps not unlike that expected in population-based studies. Despite this, the GA detection model's specificity was very high, suggesting that the algorithm did not confuse non-AMD lesions with GA. As regards sensitivity, of the sample of 20 false negative cases, we observed no cases where coexisting retinal pathology led to GA being missed.

However, we recommend that the present use case for these models (as currently configured) be limited to research settings, such as in highlighting cases with the potential for enrollment into clinical trials. Before clinical deployment, as with any novel diagnostic aid or test, external validation is required, and should be conducted in the intended use setting(s). The performance of the models for populations like the AREDS, with a very high prevalence of AMD and a lower (but non-zero) prevalence of other retinal diseases and lesions, is expected to be similar to that reported here. We anticipate that the performance may be similar in adult population-based studies, since these also have a relatively low prevalence of non-AMD retinal conditions, though performance might depend on the racial characteristics of the study population. For these reasons, first, we have made freely available the models and code, so that other research groups may perform testing in their own diverse datasets. Second, we are planning external validation studies using well-curated population-based studies from different continents. Following demonstrations of robust external validity and generalizability, these models (or newer models based on them) may be ready for deployment to assist healthcare practitioners in specific clinical settings.

Third, this study is based on one imaging modality only. We used color fundus photographs for GA detection because this imaging modality is the historical standard; color fundus photographs were used in, for example, the AREDS and AREDS2 to define and classify AMD, i.e., to derive classification and prognostic systems including the AREDS 9-step and simplified severity scales^{31,32}. As such, color fundus photographs are the most highly validated tools for AMD classification and the prediction of disease progression³³. They have the additional advantages of (i) being the closest correlates to biomicroscopy, (ii) consisting of simple datasets (unlike complex 3-dimensional datasets from optical coherence tomography (OCT)), and (iii) (unlike fundus autofluorescence) being widely available and permitting assessment of GA central involvement (without interference from macular pigment in the foveal region).

While OCT may be used increasingly to define GA in the future³⁴, few large standardized datasets currently exist that have followed eyes sequentially for long periods using OCT with detailed accompanying reading center grades and clinical information. Deep learning approaches to OCT datasets hold promise³⁵ but sophisticated software is required to process complex 3-dimensional datasets, and automated segmentation of OCT data can be error-prone in disease. Ultimately, we anticipate that deep learning approaches to AMD may be developed using multiple imaging modalities, first using each independently, then combined in a multimodal fashion. In the meantime, the development and validation of deep learning models applied to individual imaging modalities is an important first step.

In conclusion, deep learning demonstrated relatively robust performance for the automated detection of GA. The performance was non-inferior to that of human retinal specialists. In addition, potential utility was observed in the automated identification of CGA. The results of this study may serve as a benchmark for deep learning-based identification of GA from color fundus photographs. We are making the code and pretrained models publicly available at <https://github.com/ncbi-nlp/DeepSeeNet>. In this way, the work may help stimulate the future development of automated retinal image analysis tools, not only in the detection of GA, but also in the analysis of GA features including central involvement, multifocality, and area quantification.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Financial Support:

Supported by the intramural program funds and contracts from the National Center for Biotechnology Information/National Library of Medicine/National Institutes of Health, the National Eye Institute/National Institutes of Health, Department of Health and Human Services, Bethesda Maryland (contract HHS-N-260-2005-00007-C; ADB contract NO1-EY-5-0007). Funds were generously contributed to these contracts by the following NIH institutes: Office of Dietary Supplements, National Center for Complementary and Alternative Medicine; National Institute on Aging; National Heart, Lung, and Blood Institute, and National Institute of Neurological Disorders and Stroke. The sponsor and funding organization participated in the design and conduct of the study; data collection, management, analysis and interpretation; and the preparation, review and approval of the manuscript.

References:

1. Congdon N, O'Colmain B, Klaver CC, et al. Causes and prevalence of visual impairment among adults in the United States. *Arch Ophthalmol*. 2004;122(4):477–485. [PubMed: 15078664]
2. Wong WL, Su X, Li X, et al. Global prevalence of age-related macular degeneration and disease burden projection for 2020 and 2040: a systematic review and meta-analysis. *Lancet Glob Health*. 2014;2(2):e106–116. [PubMed: 25104651]
3. Quartilho A, Simkiss P, Zekite A, Xing W, Wormald R, Bunce C. Leading causes of certifiable visual loss in England and Wales during the year ending 31 March 2013. *Eye (Lond)*. 2016;30(4):602–607. [PubMed: 26821759]
4. Rudnicka AR, Jarrar Z, Wormald R, Cook DG, Fletcher A, Owen CG. Age and gender variations in age-related macular degeneration prevalence in populations of European ancestry: a meta-analysis. *Ophthalmology*. 2012;119(3):571–580. [PubMed: 22176800]
5. Keenan TD, Agron E, Domalpally A, et al. Progression of Geographic Atrophy in Age-related Macular Degeneration: AREDS2 Report Number 16. *Ophthalmology*. 2018;125(12):1913–1928. [PubMed: 30060980]

6. Rosenfeld PJ. Preventing the Growth of Geographic Atrophy: An Important Therapeutic Target in Age-Related Macular Degeneration. *Ophthalmology*. 2018;125(6):794–795. [PubMed: 29784088]
7. Gass JD. Drusen and disciform macular detachment and degeneration. *Arch Ophthalmol*. 1973;90(3):206–217. [PubMed: 4738143]
8. Schmitz-Valckenberg S, Sadda S, Staurengi G, et al. GEOGRAPHIC ATROPHY: Semantic Considerations and Literature Review. *Retina*. 2016;36(12):2250–2264. [PubMed: 27552292]
9. Age-Related Eye Disease Study Research Group. A randomized, placebo-controlled, clinical trial of high-dose supplementation with vitamins C and E, beta carotene, and zinc for age-related macular degeneration and vision loss: AREDS report no. 8. *Arch Ophthalmol*. 2001;119(10):1417–1436. [PubMed: 11594942]
10. Burlina PM, Joshi N, Pekala M, Pacheco KD, Freund DE, Bressler NM. Automated Grading of Age-Related Macular Degeneration From Color Fundus Images Using Deep Convolutional Neural Networks. *JAMA Ophthalmol*. 2017;135(11):1170–1176. [PubMed: 28973096]
11. Lam C, Yu C, Huang L, Rubin D. Retinal Lesion Detection With Deep Learning Using Image Patches. *Invest Ophthalmol Vis Sci*. 2018;59(1):590–596. [PubMed: 29372258]
12. Gargeya R, Leng T. Automated Identification of Diabetic Retinopathy Using Deep Learning. *Ophthalmology*. 2017;124(7):962–969. [PubMed: 28359545]
13. 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]
14. Wong TY, Bressler NM. Artificial Intelligence With Deep Learning Technology Looks Into Diabetic Retinopathy Screening. *JAMA*. 2016;316(22):2366–2367. [PubMed: 27898977]
15. Venhuizen FG, van Ginneken B, Liefers B, et al. Deep learning approach for the detection and quantification of intraretinal cystoid fluid in multivendor optical coherence tomography. *Biomed Opt Express*. 2018;9(4):1545–1569. [PubMed: 29675301]
16. Asaoka R, Murata H, Iwase A, Araie M. Detecting Preperimetric Glaucoma with Standard Automated Perimetry Using a Deep Learning Classifier. *Ophthalmology*. 2016;123(9):1974–1980. [PubMed: 27395766]
17. Peng Y, Dharssi S, Chen Q, et al. DeepSeeNet: A Deep Learning Model for Automated Classification of Patient-based Age-related Macular Degeneration Severity from Color Fundus Photographs. *Ophthalmology*. 2019;126(4):565–575. [PubMed: 30471319]
18. Chen Q, Peng Y, Keenan TD, et al. A multi-task deep learning framework for the classification of Age-related Macular Degeneration. Paper presented at: AMIA 2019 Informatics Summit 2019.
19. Hu Z, Medioni GG, Hernandez M, Hariri A, Wu X, Sadda SR. Segmentation of the geographic atrophy in spectral-domain optical coherence tomography and fundus autofluorescence images. *Invest Ophthalmol Vis Sci*. 2013;54(13):8375–8383. [PubMed: 24265015]
20. Feeny AK, Tadarati M, Freund DE, Bressler NM, Burlina P. Automated segmentation of geographic atrophy of the retinal epithelium via random forests in AREDS color fundus images. *Comput Biol Med*. 2015;65:124–136. [PubMed: 26318113]
21. Ramsey DJ, Sunness JS, Malviya P, Applegate C, Hager GD, Handa JT. Automated image alignment and segmentation to follow progression of geographic atrophy in age-related macular degeneration. *Retina*. 2014;34(7):1296–1307. [PubMed: 24398699]
22. Treder M, Lauermaun JL, Eter N. Deep learning-based detection and classification of geographic atrophy using a deep convolutional neural network classifier. *Graefes Arch Clin Exp Ophthalmol*. 2018;256(11):2053–2060. [PubMed: 30091055]
23. Age-Related Eye Disease Study Research Group. The Age-Related Eye Disease Study (AREDS): design implications. AREDS report no. 1. *Control Clin Trials*. 1999;20(6):573–600. [PubMed: 10588299]
24. Age-Related Eye Disease Study Research Group. The Age-Related Eye Disease Study system for classifying age-related macular degeneration from stereoscopic color fundus photographs: the Age-Related Eye Disease Study Report Number 6. *Am J Ophthalmol*. 2001;132(5):668–681. [PubMed: 11704028]

25. Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z. Rethinking the Inception Architecture for Computer Vision. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2015.
26. Grassmann F, Mengelkamp J, Brandl C, et al. A Deep Learning Algorithm for Prediction of Age-Related Eye Disease Study Severity Scale for Age-Related Macular Degeneration from Color Fundus Photography. *Ophthalmology*. 2018;125(9):1410–1420. [PubMed: 29653860]
27. Bergholm F Edge focusing. *IEEE Trans Pattern Anal Mach Intell*. 1987;9(6):726–741. [PubMed: 21869435]
28. Chollet F, others. Keras 2015; <https://github.com/keras-team/keras>.
29. Abadi M, Agarwal A, Barham P, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. arXiv. 2016.
30. Burlina PM, Joshi N, Pacheco KD, Liu TYA, Bressler NM. Assessment of Deep Generative Models for High-Resolution Synthetic Retinal Image Generation of Age-Related Macular Degeneration. *JAMA Ophthalmol*. 2019.
31. Ferris FL, Davis MD, Clemons TE, et al. A simplified severity scale for age-related macular degeneration: AREDS Report No. 18. *Arch Ophthalmol*. 2005;123(11):1570–1574. [PubMed: 16286620]
32. Davis MD, Gangnon RE, Lee LY, et al. The Age-Related Eye Disease Study severity scale for age-related macular degeneration: AREDS Report No. 17. *Arch Ophthalmol*. 2005;123(11):1484–1498. [PubMed: 16286610]
33. Burlina PM, Joshi N, Pacheco KD, Freund DE, Kong J, Bressler NM. Use of Deep Learning for Detailed Severity Characterization and Estimation of 5-Year Risk Among Patients With Age-Related Macular Degeneration. *JAMA Ophthalmol*. 2018;136(12):1359–1366. [PubMed: 30242349]
34. Sadda SR, Guymer R, Holz FG, et al. Consensus Definition for Atrophy Associated with Age-Related Macular Degeneration on OCT: Classification of Atrophy Report 3. *Ophthalmology*. 2018;125(4):537–548. [PubMed: 29103793]
35. De Fauw J, Ledsam JR, Romera-Paredes B, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat Med*. 2018;24(9):1342–1350. [PubMed: 30104768]

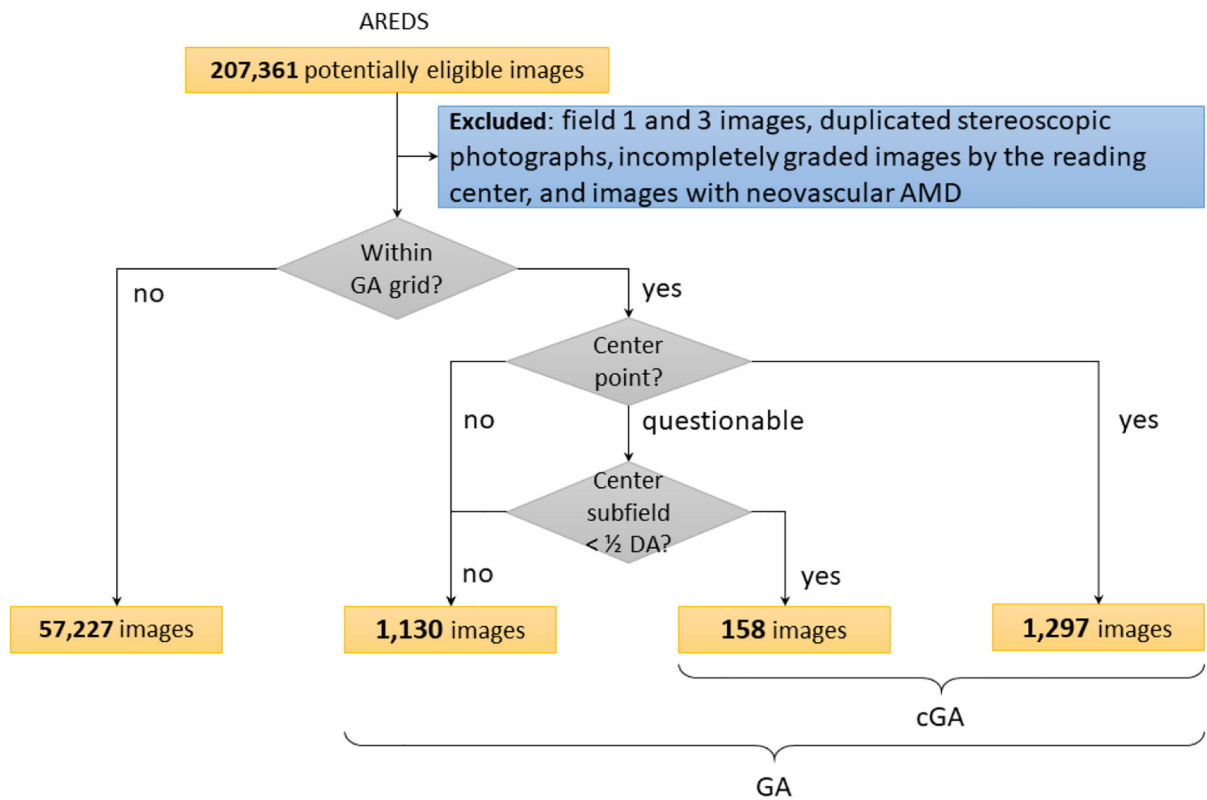


Figure 1. Distribution of color fundus photographs from the Age-Related Eye Disease Study categorized according to the presence and absence of geographic atrophy (GA) and central geographic atrophy (CGA).

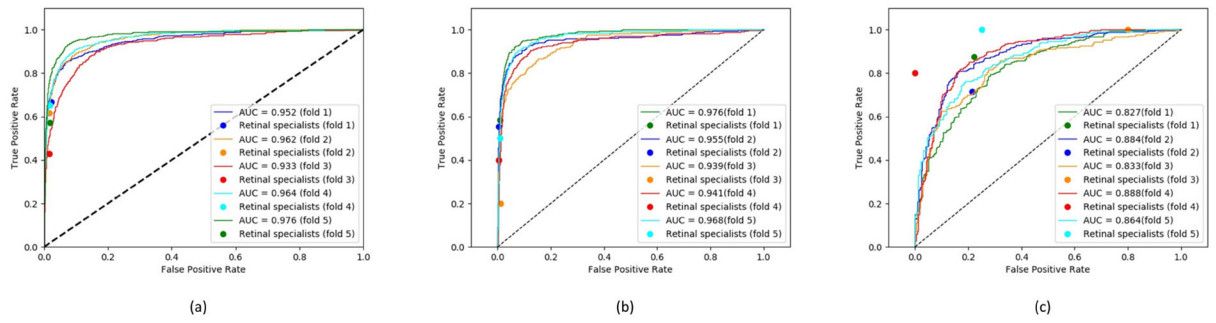


Figure 2. Receiver operating characteristic curves for (a) the geographic atrophy model, (b) the central geographic atrophy model, and (c) the centrality detector model.

Characteristics of images and participants in each partition according to the presence of geographic atrophy and central geographic atrophy.

Table 1.

	Total	Partition				
		1	2	3	4	5
Images	59,812	11,933	11,814	11,882	12,056	12,127
Participants	4,582	924	907	915	919	917
Age: Mean, years (SD)	69.4 (5.1)	69.5 (5.1)	69.4 (5.2)	69.4 (5.2)	69.6 (5.0)	69.2 (5.0)
Gender (% male)	44.3%	44.9%	44.4%	45.2%	44.7%	42.2%
Geographic atrophy	4.3% (n=2,585)	4.9% (n=588)	4.8% (n=571)	3.5% (n=420)	4.9% (n=586)	3.5% (n=420)
Central geographic atrophy	2.4% (n=1,455)	2.9% (n=344)	2.9% (n=347)	1.8% (n=216)	2.7% (n=325)	1.8% (n=223)

Table 2.

Performance of the three deep learning models: (i) geographic atrophy model ('GA model'), (ii) central geographic atrophy model ('cGA' model), and (iii) centrality detector model.

	GA detection		cGA detection from all eyes		cGA detection from eyes with GA	
	GA model (95% CI)	Retinal specialists (95% CI)	cGA model (95% CI)	Retinal specialists (95% CI)	Centrality detector model (95% CI)	Retinal specialists (95% CI)
Accuracy	0.965 (0.959, 0.971)	0.975 (0.971, 0.980)	0.966 (0.957, 0.975)	0.990 (0.987, 0.993)	0.762 (0.725, 0.799)	0.735 (0.445, 1.000)
Kappa	0.611 (0.533, 0.689)	0.438 (0.304, 0.571)	0.501 (0.451, 0.552)	0.344 (0.156, 0.532)	0.513 (0.442, 0.584)	0.533 (0.176, 0.891)
Sensitivity	0.692 (0.560, 0.825)	0.588 (0.468, 0.707)	0.763 (0.641, 0.885)	0.448 (0.255, 0.641)	0.782 (0.618, 0.945)	0.878 (0.722, 1.000)
Specificity	0.978 (0.970, 0.985)	0.982 (0.978, 0.985)	0.971 (0.960, 0.982)	0.993 (0.989, 0.996)	0.729 (0.543, 0.916)	0.703 (0.332, 1.000)
Precision	0.584 (0.491, 0.676)	0.368 (0.230, 0.505)	0.394 (0.341, 0.448)	0.296 (0.115, 0.477)	0.799 (0.710, 0.888)	0.626 (0.273, 0.979)

Table 3.

False negatives of geographic atrophy detection categorized by the geographic atrophy area within grid.

Area category	Whole test set	False negatives	
Questionable	275	173	62.9%
Definite, area < circle I-2	38	30	78.9%
Area I-2 but < O-2	125	76	60.8%
Area O-2 but < ½ DA	192	90	46.9%
Area ½ DA but < 1 DA	297	79	26.6%
Area 1 DA but < 2 DA	403	95	23.6%
Area 2 DA	1,255	248	19.8%

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4.

False negatives of central geographic atrophy detection categorized by involvement of the center point and center subfield.

Area category	Whole test set	False negatives	
Definite, center point < I-2	1,297	279	21.5%
Questionable center point, but center subfield area < ½ disc area	158	50	31.6%

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript