

UCSF

UC San Francisco Previously Published Works

Title

Differentiation of Active Corneal Infections from Healed Scars Using Deep Learning

Permalink

<https://escholarship.org/uc/item/7dd4790t>

Journal

Ophthalmology, 129(2)

ISSN

0161-6420

Authors

Tiwari, Mo
Piech, Chris
Baitemirova, Medina
[et al.](#)

Publication Date

2022-02-01

DOI

10.1016/j.opthta.2021.07.033

Peer reviewed



Published in final edited form as:

Ophthalmology. 2022 February ; 129(2): 139–146. doi:10.1016/j.ophtha.2021.07.033.

Differentiation of Active Corneal Infections From Healed Scars Using Deep Learning

Mo Tiwari¹, Chris Piech, PhD¹, Medina Baitemirova², Namperumalsamy V. Prajna, DO, DNB³, Muthiah Srinivasan, DO, DNB³, Prajna Lalitha, DO, DNB³, Natacha Villegas, MD⁴, Niranjana Balachandar¹, Janice T. Chua⁵, Travis Redd, MD, MPH⁶, Thomas M. Lietman, MD⁷, Sebastian Thrun, PhD¹, Charles C. Lin, MD⁴

¹Department of Computer Science, Stanford University, Stanford, CA, USA

²Department of Biomedical Informatics, Stanford University, Stanford, CA, USA

³Aravind Eye Hospital, Madurai, India

⁴Byers Eye Institute, Stanford University, Stanford CA, USA

⁵University of California Irvine, School of Medicine, Irvine, CA, USA

⁶Department of Ophthalmology, Casey Eye Institute, Oregon Health and Science University, Portland, OR, USA

⁷Francis I. Proctor Foundation, University of California San Francisco, San Francisco, CA, USA

Abstract

Objective: To develop and evaluate an automated, portable algorithm to differentiate active corneal ulcers from healed scars using only external photographs.

Design: A convolutional neural network was trained and tested using photographs of corneal ulcers and scars.

Subjects: De-identified photographs of corneal ulcers were obtained from the Steroids for Corneal Ulcers Trial (SCUT), Mycotic Ulcer Treatment Trial (MUTT), and Byers Eye Institute at Stanford University.

Corresponding Author: Charles C. Lin, Clinical Associate Professor, Byers Eye Institute - Stanford University, 2452 Watson Court, Palo Alto, CA 94303, lincc@stanford.edu.

RESEARCH DESIGN: Mo Tiwari, Chris Piech, Medina Baitemirova, Niranjana Balachandar, Thomas M. Lietman, Sebastian Thrun, Charles C. Lin.

DATA ACQUISITION AND/OR RESEARCH EXECUTION: Mo Tiwari, Chris Piech, Medina Baitemirova, Namperumalsamy V. Prajna, Muthiah Srinivasan, Prajna Lalitha, Natacha Villegas, Travis Redd, Thomas M. Lietman, Sebastian Thrun, Charles C. Lin.

DATA ANALYSIS AND/OR INTERPRETATION: Mo Tiwari, Chris Piech, Medina Baitemirova, Niranjana Balachandar, Janice T. Chua, Thomas M. Lietman, Sebastian Thrun, Charles C. Lin.

MANUSCRIPT PREPARATION: Mo Tiwari, Chris Piech, Medina Baitemirova, Janice T. Chua, Travis Redd, Thomas M. Lietman, Sebastian Thrun, Charles C. Lin.

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 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.

Conflicts of Interest: None of the authors have a proprietary/financial interest to disclose.

Methods: Photographs of corneal ulcers (n=1313) and scars (n=1132) from the SCUT and MUTT trials were used to train a convolutional neural network (CNN). The CNN was tested on two different patient populations from eye clinics in India (n=200) and the Byers Eye Institute at Stanford University (n=101). Accuracy was evaluated against gold standard clinical classifications. Feature importances for the trained model were visualized using Gradient-weighted Class Activation Mapping (Grad-CAM).

Main Outcome Measure: Accuracy of the CNN was assessed via F_1 score. Area under the receiver operating characteristic curve (ROC) was used to measure the precision-recall trade-off.

Results: The CNN correctly classified 115/123 active ulcers and 65/77 scars in corneal ulcer patients from India (F_1 score: 92.0% (95% CI: 88.2 – 95.8%), sensitivity: 93.5% (95% CI: 89.1 – 97.9%), specificity: 84.42% (95% CI: 79.42 – 89.42%), ROC (AUC=0.9731)). The CNN correctly classified 43/55 active ulcers and 42/46 scars in corneal ulcer patients from Northern California (F_1 score: 84.3% (95% CI: 77.2 – 91.4%), sensitivity: 78.2% (95% CI: 67.3 – 89.1%), specificity: 91.3% (95% CI: 85.8 – 96.8%), ROC (AUC=0.9474)). The CNN visualizations correlated with clinically relevant features such as corneal infiltrate, hypopyon, and conjunctival injection.

Conclusion: The CNN classified corneal ulcers and scars with high accuracy and generalizes to patient populations outside of its training data. The CNN focuses on clinically relevant features when it makes a diagnosis. The CNN demonstrates potential as an inexpensive diagnostic approach that may aid triage in communities with limited access to eye care.

Precis

A deep learning algorithm was developed to classify active cornea ulcers from healed scars. The algorithm performed with high accuracy and generalized well to different patient populations and imaging modalities.

Keywords

corneal ulcer; corneal scar; artificial intelligence; deep learning; infectious keratitis

Introduction

Despite advances in our understanding of how to treat corneal infections, corneal blindness remains the 4th most common cause of blindness in the world, accounting for over 5% of the total blind population.¹ Corneal opacities disproportionately affect developing countries, where they may constitute over 10% of preventable visual impairment cases. For example, in India, corneal ulcers affect around 2 million people every year.¹ In such developing countries, delayed or inappropriate treatment due to misdiagnosis is a primary driving factor. Appropriate triage of these patients through accurate diagnosis of a corneal ulcer is the first step towards addressing this problem, but this requires a trained eye care provider, which often is a limited resource in these areas. India, for example, has an estimated 25,000 ophthalmologists for its 1.3 billion residents.² This ratio of 1:52,000 is well below typical public health guidelines.³ In the absence of trained eye care providers, differentiating active corneal ulcers from healed scars may be a more challenging, less intuitive task.

A sophisticated form of pattern recognition – identifying key discriminating features of ulcers and scars – lies at the heart of proper diagnosis. For example, ophthalmologists look for key features such as corneal infiltrate, conjunctival injection, and hypopyon when making a diagnosis. Recent advancements in the fields of machine learning and computer vision have enabled a wide variety of new technologies, particularly in pattern recognition and image classification. In particular, deep learning - a specific subfield of machine learning that uses large neural networks - has recently been used in a variety of image classification tasks in medicine. Deep learning has been used to classify different forms of skin cancer,⁴ detect heart attacks,⁵ and diagnose neurological stroke,⁶ among other conditions. In ophthalmology, machine learning algorithms have been demonstrated to successfully detect age-related macular degeneration,⁷ classify diabetic retinopathy severity,⁸ and diagnose glaucoma.⁹ In this work, we train a convolutional neural network (CNN), a particular type of deep learning model, to distinguish corneal ulcers from scars.

Methods

External photographs were acquired from the Steroid for Corneal Ulcer Trial (SCUT) and the Mycotic Ulcer Treatment Trial (MUTT).^{10,11} Photographs from both studies were captured using a handheld Nikon D-series digital SLR camera with a 105-mm f/2.8D AF Micro Nikkor Autofocus Lens and a modified Nikon SB29s electronic flash or Nikon R1 Wireless Close-up Speedlight system. In the SCUT study, 500 patients with bacterial corneal ulcers had photographs taken at enrollment, 3-week, 3-month, and 12-months follow-up visits.¹⁰ In the MUTT study, 323 patients with fungal corneal ulcers had photographs taken at enrollment, 3-week, and 3-month follow-ups.¹¹ Both the SCUT and MUTT studies were conducted at the Aravind Eye Hospital in Tamil Nadu, India.

Inclusion criteria in both the SCUT and MUTT studies required an active, culture-proven bacterial or fungal corneal ulcer. As such, photographs taken at enrollment were labeled as active corneal ulcers. When clinical review indicated resolution of the ulcer, defined as complete epithelialization, images from the 3 month follow-up visits were labeled as healed corneal scars. The majority of ulcers healed well before 3 months and the mean elapsed time between clinical resolution to the 3 month follow-up visit was 76.9 days (standard deviation of 12.1 days). As some cases, especially fungal ulcers, can remain active despite epithelialization, if there was ambiguity regarding clinical resolution based on chart records, such as worsening with cessation of antimicrobials, these cases and photographs were excluded.

For both SCUT and MUTT studies, multiple photographs were taken at each study visit. Except for using different external flash devices, the SCUT and MUTT studies used the same cameras and aperture priority settings, centered and focused on the corneal apex. Due to differences in data management of the image collections, only 1 photograph from each study visit was available for SCUT patients whereas an average of 4.8 and 4.2 photographs for the enrollment and 3-month visits, respectively, were available for MUTT. To maximize the number of photographs that could be used, only exact duplicates and out of focus images were excluded.

Patients from the SCUT and MUTT studies were pooled into a single list, and randomly assigned to the train, validation, and test datasets according to a 60/20/20 split. The train, validation, and test datasets were stratified in this way (by patient) to ensure that no patient cross-contamination occurred between the test set and other sets. A total of 612, 200, and 200 patients were allocated to the training, validation, and test sets, respectively. For MUTT patients, multiple images from the same patient study visit could be used for the training set whereas for patients stratified to the test set, only 1 image each from their enrollment and 3 month visit could be included. This test set, called Test Set 1, contains patients and images from SCUT and MUTT that were not seen by the CNN during training.

A separate dataset of 101 slit lamp photographs was acquired from the Byers Eye Institute at Stanford University in Northern California (Canon 7D digital camera, Haag-Streit BX900) and comprised 55 active corneal ulcers and 46 corneal scars from 101 unique patients. These photographs were classified based on clinical interpretation by a cornea specialist (CCL). None of the images in this dataset, called Test Set 2, were used for training the CNN and were only used for testing.

All images were normalized to the training dataset and rescaled to 224 pixels by 224 pixels. The training images from the SCUT and MUTT studies were used to train a deep learning model to classify active corneal infections and corneal scars. We used the VGG-16 architecture for the CNN, a specific type of deep learning model.¹² The model architecture is depicted in Figure 1. The model was pretrained on the 2014 ImageNet Large Scale Visual Recognition Challenge, a database of approximately 1.28 million images of 1,000 object categories, which is currently the standard form of pretraining in PyTorch, a deep learning software package.^{13,14} The final 1,000-node layer typically used for ImageNet was replaced with a 2-node layer for our task.¹² The pretrained model was fine-tuned using the training dataset from the SCUT and MUTT studies. The model was trained to perform a binary classification task (active ulcer versus healed scar) using cross-entropy loss, optimized via Adam with an initial learning rate of 10^{-5} that decayed by a factor of 10 every 5 training epochs. Training was continued for 13 epochs, when the validation loss began to increase which indicated overfitting. When available, multiple photographs from each patient was used during fine-tuning.

After fine-tuning was completed, the performance of the model was measured on Test Set 1 and Test Set 2. A single photograph from each patient in each test set was used to measure final performance. Confidence intervals for final test performance were computed using the bootstrap method.

A precision-recall curve was generated to assess the false positive rate and false negative rate tradeoff as a function of model confidence. A receiver operating characteristic (ROC) curve to assess true positive rate and false positive rate tradeoff was also calculated. A calibration curve was created by binning the patient predictions into separate confidence bins of being an active ulcer, and then plotting the proportion of patients in each bin that were determined by the ophthalmologist to be active ulcers. The bins were selected to provide a sufficient level of granularity that may be relevant for clinical predictions, while still permitted at least 10 patients per bin.

As a metric of model performance, the F_1 score was chosen as it is a commonly used statistic in settings where test sets may be imbalanced. It is computed as the harmonic mean of the precision and recall. Recall is defined as the number of predicted and true positives divided by the number of true positives and is a measure of the model's ability to accurately identify active ulcers. Precision is the number of predicted and true positives divided by the number of false positives and is a measure of the model's ability to identify healed infections.

The deep learning model was also compared to a simple baseline model that classifies images as ulcers and scars based on the proportion of red pixels. The exact proportion of red pixels that served as the cutoff for classifying an image as an active ulcer is a hyperparameter that was tuned on the training dataset.

Gradient-weighted Class Activation Mapping (Grad-CAM) was used to visualize the model and determine salient regions of active ulcers. Grad-CAM uses the gradients of the target class (in this case, active ulcer) and creates a localization map that highlights regions of importance that were used to make the active ulcer versus healed scar prediction.¹⁵

A cornea specialist (CCL) manually annotated 40 random photographs with an active ulcer from Test Set 1 to highlight areas of clinical relevance. Pixels in each image were then manually marked as important according to the cornea specialist's annotations, as well as by thresholding the Grad-CAM heatmap. The similarity between Grad-CAM heatmaps and manually annotated photographs was calculated using the overlap coefficient, or Szymkiewicz–Simpson coefficient, defined as the number of pixels common to both importance maps divided by the number of pixels in the image with the smaller number of important pixels.¹⁶ This segmentation analysis was not used to augment our deep learning classification model.

This study was approved by the Institutional Review Board at Stanford University and adhered to the tenets of the Declaration of Helsinki.

Results

On Test Set 1, comprised of held out images from the SCUT and MUTT studies, the model accurately classified 115 out of 123 active corneal ulcers and 65 out of 77 corneal scars, with an F_1 score of 92.0% (95% CI: 88.2 – 95.8%). The sensitivity, or the ability of our model to correctly identify active corneal ulcers, was 93.5% (95% CI: 89.1 – 97.9%). The specificity, or the ability of our model to identify true negatives, was 84.42% (95% CI: 79.42 – 89.42%). Figure 2a shows the receiver operating characteristic (ROC) curve, which plots the true positive rate (TPR) against the false positive rate (FPR) for the trained model (AUC = 0.9731). Figure 2b shows the precision-recall (PR) curve, which indicates the tradeoff between the false positive rate (FPR) and false negative rates (FNR) as a function of decision threshold (AUC = 0.9811). More formally, the PR curve plots $TPR/(TPR + FPR)$ against $TPR/(TPR+FNR)$. Figure 2c shows the calibration curve for Test Set 1.

On Test Set 2, comprised of images from the Stanford Byers Eye Institute, the model correctly classified 43 out of 55 active corneal ulcers and 42 out of 46 corneal scars in

patients, with an F_1 score of 84.3% (95% CI: 77.2 – 91.4%), sensitivity of 78.2% (95% CI: 67.3 – 89.1%), and specificity of 91.3% (95% CI: 85.8 – 96.8%). Figures 2d, 2e, and 2f show the ROC (AUC = 0.9474), PR (AUC = 0.9524), and calibration curves for Test Set 2, respectively.

The baseline classifier, which used the proportion of red pixels (defined as having a hue value between 0 and 10) in each achieved a lower accuracy of 76.0% (95% CI: 74.9 – 77.1%) on the training set.

Figure 3 demonstrates a sample of heatmaps created by Grad-CAM. The heatmaps identify areas of the images that are relevant to the trained model's classification decisions. These areas align with clinical features associated with an active corneal ulcer, including conjunctival injection and corneal infiltrate (Fig. 4).

Among the 40 randomly selected photographs with an active ulcer manually annotated by a cornea specialist to indicate areas of clinical relevance, some annotations overlapped significantly with areas highlighted by heatmaps (Fig. 4), while in others, there was little or no overlap (Fig. 5). The average overlap, or Szymkiewicz–Simpson coefficient was 0.48. In cases of significant overlap (Szymkiewicz–Simpson coefficient >0.7), the model accurately classified 10 out of 12 ulcers (accuracy 83%) whereas in cases of little overlap (Szymkiewicz–Simpson coefficient <0.3), the model accurately classified only 9 out of 16 ulcers (accuracy 56%). As such, greater overlap is associated with better algorithm performance.

Discussion

In this work, we developed a deep learning model that can accurately differentiate between active corneal infections and healed scars. In particular, the model achieved a high level of accuracy in classifying corneal ulcers and scars when trained on photographs from the SCUT and MUTT studies. These results compare favorably with the performance of human experts in completing this task. A recent study from the University of Michigan evaluating cornea specialists' accuracy in diagnosing corneal ulcers from images captured from portable cameras reported a sensitivity of 82 – 94%.¹⁷ Results from our model are similar, with sensitivities of 93.5 \pm 4.4% and 78.2 \pm 8.0% in Test Sets 1 and 2, respectively. These are promising results when taken in light of the proposed accuracy threshold of 80% sensitivity for telemedicine screening for conditions such as diabetic retinopathy.¹⁸

Crucial to the assessment of any machine learning model is its ability to “generalize”, or perform well on data which was not used in the training dataset. Our model not only generalized to patients in Test Set 1 that it had never seen before, but also to a different patient population from the United States in Test Set 2. Notably, the Northern California population had significantly different patient demographics and microbiological epidemiology than the training data from the Indian clinics. The SCUT and MUTT studies enrolled patients predominantly from the Tamil Nadu state in South India, where a leading cause of corneal ulcers is agricultural trauma. In this tropical environment, fungal ulcers

are more common than bacterial.¹⁹ In contrast, in the coastal temperate climate of Northern California, the majority of corneal ulcers are bacterial in origin and are most often caused by contact lens wear.²⁰ The generalizability of our model indicates that the algorithm identifies salient features of corneal ulcers and scars that are present despite epidemiological differences.

In addition, the model generalizes well to a different imaging modality, suggesting greater potential for a portable telemedicine implementation. The model was trained on images captured with a handheld camera yet performed well on photographs captured with a slit lamp mounted camera. A robust evaluation of the performance of a model beyond simply overall accuracy is important as the ramifications of misdiagnosis are significant. For example, misclassifying an active infection as a scar (false negative) can lead to corneal blindness and classifying a scar as an active infection (false positive) can lead to unnecessary evaluations and treatment. The PR, ROC, and calibration curves in Figure 2 provide additional information to evaluate the model's performance. The high AUCs of the PR and ROC curves suggests that the FPR and FNR are relatively low at reasonable decision thresholds. In practice, the decision threshold can be varied to balance the disparate costs associated with false negatives and false positives.

Furthermore, Figure 2 also contains calibration plots for the model's predictions on Test Sets 1 and 2. The Expected Calibration Error (ECE), a standard metric of calibration error, for both Test Sets 1 and 2 is approximately 7 percentage points, which is promising though may demonstrate mild overfitting in data-scarce regions. The 5 bins in the calibration plots were chosen to provide sufficient granularity in predicted confidence level to clinicians while ensuring more than 10 datapoints per bin.

The more sophisticated CNN outperformed the simple baseline classifier based solely on the number of red pixels in an image. This suggests that the CNN identifies features that are not accessible to this baseline classifier, such as the arrangement of redness (e.g. around the limbus). Furthermore, the Grad-CAM heatmaps (Figures 4 and 5) demonstrate that the model focuses on features beyond ciliary flush and conjunctival injection. Indeed, the heatmaps demonstrate the model's emphasis on clinically relevant features of an active infection, such as the presence of a hypopyon or purulent corneal infiltrate. Additionally, the generation of Grad-CAM visualizations does not require retraining or otherwise modifying the CNN, in contrast with some other CNN visualization techniques.¹⁵ Such heatmaps may be presented to a medical provider for real-time human oversight of an automated diagnosis. Notably, the salient regions of an image in a classification prediction may not always be readily interpretable or even visible to the naked human eye. For example, features such as the amount of blood in sparsely or densely clustered blood vessels in an eye are harder to visually quantify. For these reasons, the regions highlighted by the model should be further reviewed by medical professionals. Discrepancies amongst model predictions and ophthalmologist predictions may allow for novel insights into clinical features of corneal ulcers. Furthermore, the human supervision provided by the ophthalmologist could be used to improve the model towards more fully-automated AI diagnosis.

The trained model is a step towards a portable system that can differentiate active corneal ulcers from healed infections, which may be helpful in areas with limited access to eye care. This model only requires low-resolution (224 pixels by 224 pixels) photographs, which is well within the capabilities of most smartphones. The incorporation of this model into a smartphone-enabled application could provide a mobile diagnostic aid to assist triage efforts by a healthcare provider. While much work lies ahead, including training providers to use the developed technology, there is potential to make a significant impact in developing countries that are disproportionately impacted by corneal blindness. In India, for example, there is only 1 ophthalmologist for every 52,000 residents but almost 1 in every 3 people has a smartphone.^{2,21} This model may help the triage process for primary care providers, who generally have limited experience with eye care. While patient history can provide important clues to providers regarding diagnosis, there are many conditions that can mitigate symptoms of an active corneal ulcer such as a neurotrophic cornea, treatment with topical steroids that mask inflammation, and an infection in its early stages. Patient reported symptoms also vary considerably. For example, a scar can lead to debilitating photophobia for 1 patient and be relatively asymptomatic for another. Visualizing corneal pathology is crucial to accurate diagnosis and our algorithm aims to support the diagnostic process through analysis of photographs. If a primary care provider flags a potential infection with this model, the patient could then be referred to an eye specialist. As telemedicine becomes more prevalent, especially in the era of COVID-19, mobile diagnostic tools will likely become even more important.

Furthermore, the trained model could potentially be used to help patients determine treatment efficacy. For example, patients could take serial photos of their corneal ulcer to assess disease resolution. An ulcer that is not responding to treatment may be flagged as still active and prompt additional consultation.

There are a few limitations to this study. First, there is a slight drop in generalization performance when testing the model on Test Set 2, a patient population geographically and epidemiologically different from the training set. In practice, this could be remediated by continually retraining the model on new patients from each location in which the model is to be deployed. Second, all images in this study were obtained using high resolution digital cameras. In these photographs, ambient lighting was relatively uniform and out-of-focus images were excluded. Our model may not be as accurate with poorer quality photographs where camera conditions such as lighting and focus are inconsistent. In practice, this may be ameliorated by publishing standardization guidelines for taking photographs to be classified by our model. Finally, our model was trained only on bacterial and fungal eye infections, and may perform worse when tested on parasitic or viral infections. Future work could include additional data from such cases to improve our model.

In this work, we developed a deep learning model to classify corneal ulcers and scars that performs well on out-of-sample data. Our model offers a promising, inexpensive diagnostic approach with potential to aid triage in communities with limited access to eye care. In addition, the model may also be used as an adjunctive tool to help monitor clinical progression and prompt re-evaluation in refractory cases.

Financial Support:

This work was supported by core grant P30-026877 from the National Eye Institute and Research to Prevent Blindness, NIH K12EY027720 (Redd), NIH U10EY015114 (SCUT study, Lietman), and NIH U10EY018573 (MUTT study, Lietman). Medina Baitemirova is supported by a graduate fellowship award from Knight-Hennessy Scholars at Stanford University. This research was funded in part by JPMorgan Chase & Co. Any views or opinions expressed herein are solely those of the authors listed, and may differ from the views and opinions expressed by JPMorgan Chase & Co. or its affiliates. This material is not a product of the Research Department of J.P. Morgan Securities LLC. This material should not be construed as an individual recommendation for any particular client and is not intended as a recommendation of particular securities, financial instruments or strategies for a particular client. This material does not constitute a solicitation or offer in any jurisdiction.

Abbreviations and Acronyms:

SCUT	Steroids for Corneal Ulcers Trial
MUTT	Mycotic Ulcer Treatment Trial
CNN	convolutional neural network
Grad-CAM	Gradient-weighted Class Activation Mapping
ROC	operating characteristic
AUC	area under the curve
PR	precision-recall
TPR	true positive rate
FPR	false positive rate
FNR	false negative rate

References

1. Austin A, Lietman T, Rose-Nussbaumer J. Update on the Management of Infectious Keratitis. *Ophthalmology*. 2017;124(11):1678–1689. [PubMed: 28942073]
2. India - number of ophthalmologists 2005–2020. <https://www.statista.com/statistics/953172/india-number-of-ophthalmologists/>. Accessed December 22, 2020.
3. Access to eye care, uptake of services are issues in India. <https://www.healio.com/news/ophthalmology/20120225/access-to-eye-care-uptake-of-services-are-issues-in-india>. Accessed December 22, 2020.
4. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks [published correction appears in *Nature*. 2017 Jun 28;546(7660):686]. *Nature*. 2017;542(7639):115–118. [PubMed: 28117445]
5. Strodthoff N, Strodthoff C. Detecting and interpreting myocardial infarction using fully convolutional neural networks. *Physiol Meas*. 2019;40(1):015001. [PubMed: 30523982]
6. Zech J, Pain M, Titano J, et al. Natural Language-based Machine Learning Models for the Annotation of Clinical Radiology Reports. *Radiology*. 2018;287(2):570–580. [PubMed: 29381109]
7. 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]
8. 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]

9. Li Z, He Y, Keel S, Meng W, Chang RT, He M. Efficacy of a Deep Learning System for Detecting Glaucomatous Optic Neuropathy Based on Color Fundus Photographs. *Ophthalmology*. 2018;125(8):1199–1206. [PubMed: 29506863]
10. Srinivasan M, Mascarenhas J, Rajaraman R, et al. Corticosteroids for bacterial keratitis: the Steroids for Corneal Ulcers Trial (SCUT). *Arch Ophthalmol*. 2012;130(2):143–150. [PubMed: 21987582]
11. Prajna NV, Krishnan T, Mascarenhas J, et al. The mycotic ulcer treatment trial: a randomized trial comparing natamycin vs voriconazole. *JAMA Ophthalmol*. 2013;131(4):422–429. [PubMed: 23710492]
12. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. <http://arxiv.org/abs/1409.1556>. Accessed December 20, 2020.
13. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. IEEE; 2009.
14. Paszke A, Gross S, Massa F, et al. PyTorch: An imperative style, high-performance deep learning library. <http://arxiv.org/abs/1912.01703>. Accessed December 22, 2020.
15. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. 2017 IEEE International Conference on Computer Vision (ICCV). 2017:618–626.
16. Vijaymeena MK, Kavitha K. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal (MLAIJ)*. 2016;3(1):19–28.
17. Woodward MA, Musch DC, Hood CT, et al. Tele-ophthalmic Approach for Detection of Corneal Diseases: Accuracy and Reliability. *Cornea*. 2017;36(10):1159–1165. [PubMed: 28820791]
18. British Diabetic Association. Retinal Photography Screening for Diabetic Eye Disease: A British Diabetic Association Report. London: British Diabetic Association; 1997.
19. Lin CC, Lalitha P, Srinivasan M, et al. Seasonal trends of microbial keratitis in South India. *Cornea*. 2012;31(10):1123–1127. [PubMed: 22868629]
20. Jeng BH, Gritz DC, Kumar AB, et al. Epidemiology of ulcerative keratitis in Northern California. *Arch Ophthalmol*. 2010;128(8):1022–1028. [PubMed: 20697003]
21. Vaibhav A Smartphone penetration in India 2014–2022. <https://www.statista.com/statistics/257048/smartphone-user-penetration-in-india/> Accessed December 22, 2020.




Figure 1.
Architecture diagram for convolutional neural network VGG-16.

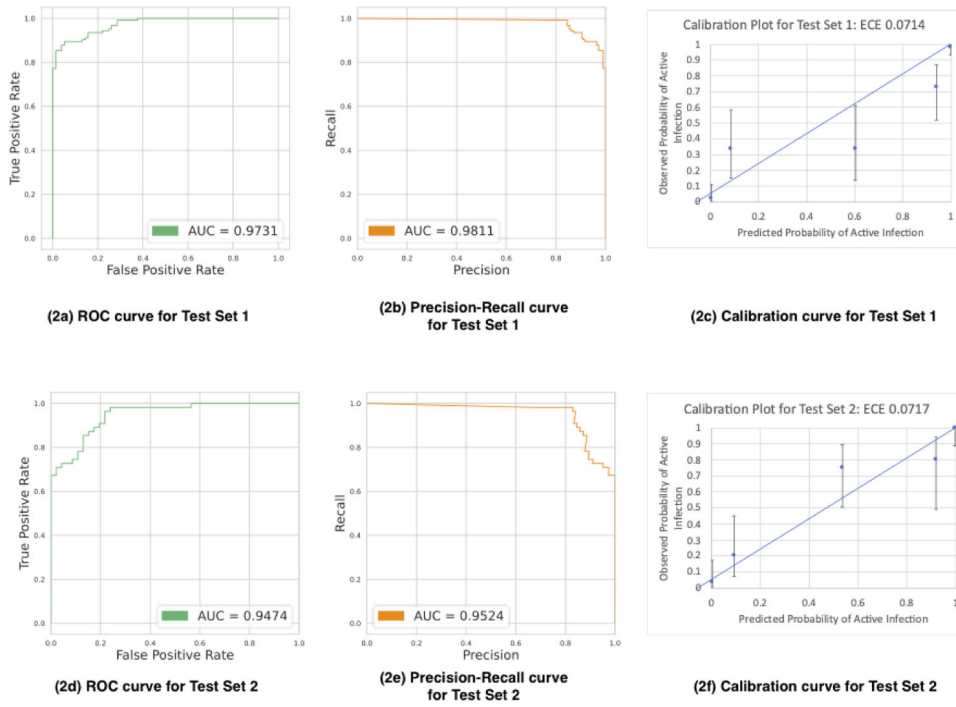


Figure 2. Test Set 1 Receiver operating characteristic curve (2a), Precision-Recall curve (2b), and calibration curve (2c). Test Set 2 Receiver operating characteristic curve (2d), Precision-Recall curve (2e), and calibration (2f) curve.

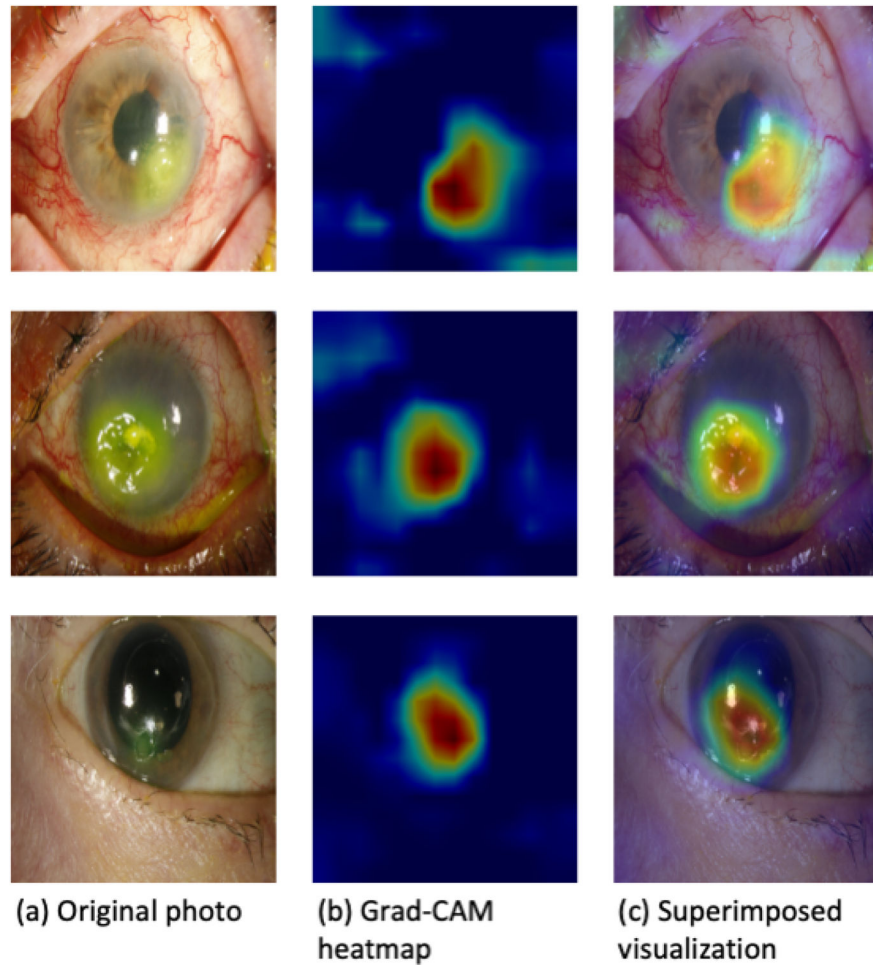


Figure 3.
Examples of Grad-CAM visualizations.

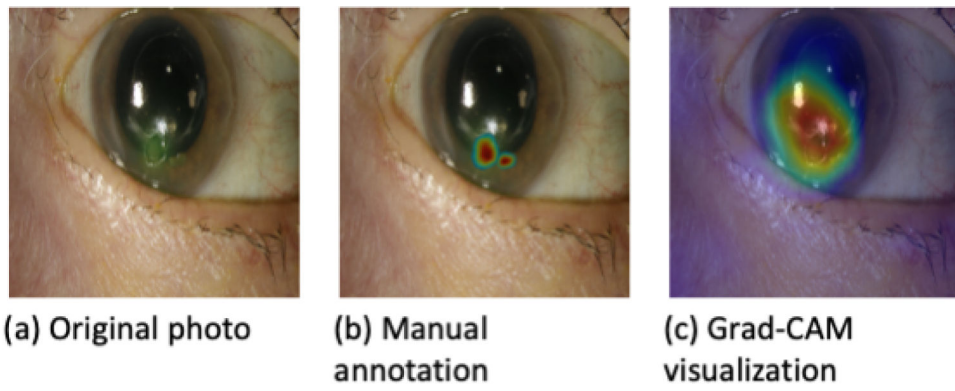


Figure 4. Comparison of a manually annotated photograph of a corneal ulcer and its corresponding Grad-CAM heatmap with good agreement.

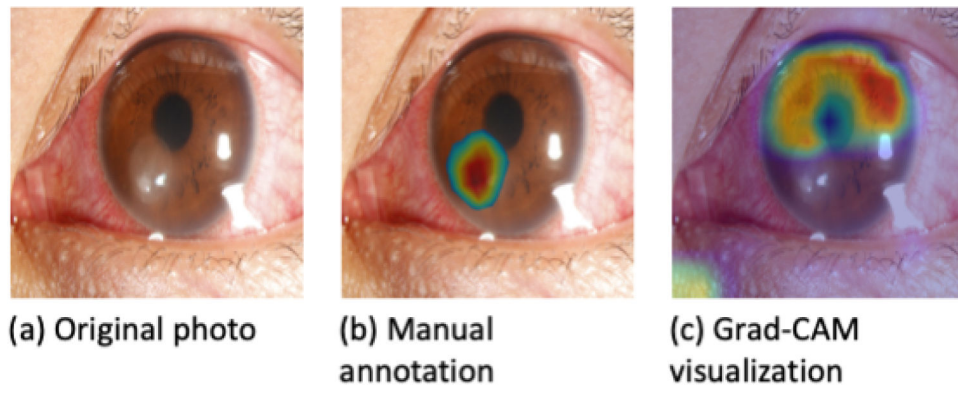


Figure 5. Comparison of a manually annotated photograph of a corneal ulcer and its corresponding Grad-CAM heatmap with poor overlap.