Using AI to Prevent Vision Loss for Millions Globally

Mr. Arun Shroff
Founder & CEO, Xtend.AI
Arun.shroff@gmail.com

A Global Healthcare Challenge:

As per the World Health Organization (WHO) the number of people with diabetes worldwide has almost quadrupled in the last several decades - from 108 million in 1980 to over 422 million people today. By 2040 that number is projected to grow to over 640 million. (1)

Diabetes leads to many complications including greater risk of heart disease, stroke and diabetic retinopathy (DR) which can lead to vision impairment and blindness. It is estimated that from a third to half of those with diabetes, or from 150 to 210 million people worldwide have some stage of DR. In addition, an estimated 10% or 40 million people with diabetes have DR that is considered vision-threatening (VTDR). This makes DR the leading cause of preventable blindness in working age adults worldwide.(2)

The risk of vision impairment and vision loss can be greatly reduced if DR is detected and treated in its early stages. Currently, detecting DR requires a trained eye doctor to examine color photos of the fundus region of the retina for tiny lesions and other vascular abnormalities caused by diabetes. This is a manual and time-consuming process, that requires expensive equipment and trained professionals. Unfortunately, many developing countries have an acute shortage of both specialists and resources to screen everyone at risk.

In India, for example, there are over 72 million people with diabetes and an estimated 25 million are afflicted with DR and 7 million with VTDR.(3) However, India a nation of 1.3 billion, only has 15,000 ophthalmologists - or a mere 9 specialists per million population. Similarly, Kenya, with a population of 48 million has less than 100 ophthalmologists, and Angola, with a population of 29 million has less than 20. In addition, to the shortage of trained professionals, many people afflicted with DR live in remote areas with little or no access to an eye care clinic or a screening center. (4)

As a result, diabetic retinopathy has become a global healthcare challenge. The good news is that AI offers a way to automate many aspects of DR detection and screening that can potentially help address this challenge.

How AI Can Help:

Recent advances in deep learning – a branch of AI that uses deep neural networks have been very successful in a large variety of tasks including voice and natural language processing, as well as image recognition and classification. A breakthrough in image classification was achieved in 2012, in the ImageNet Challenge which requires identification of 1000 different categories of objects in a dataset containing millions of images. Since then the best AI models have achieved an accuracy of over 96% in the image classification task – exceeding the best human accuracy levels.(5)

Many of these AI models have now been adapted successfully for use in a variety of medical image diagnosis tasks such as melanoma, breast, lung cancer detection and diabetic retinopathy.

In particular, a team at Google published results in 2016 of a study for detecting DR working with doctors in India and the US. The results show that their AI model's performance for DR detection and grading its severity was on-par with that of ophthalmologists. Their model had a combined accuracy score of 0.95, which was slightly better than the median of the 8 ophthalmologists consulted (measured at 0.91). (6)

Proposed Solution:

While an AI model can automate the screening process for DR, it is by no means intended to replace doctors. Trained professionals will always be essential for validating the AI model and for counseling, follow-up and treatment of patients.

The ideal system should combine AI, technology and human expertise in ways that can complement each other. In order to address the global DR challenge, the system design should satisfy the following objectives:

i) Provide the capability to be deployed in remote and underserved areas.
ii) Provide the capability to capture retinal images with a low-cost device.
iii) Use AI to automate routine screening of DR for majority of the cases.
iv) Allow remote diagnosis by trained professionals when required for cases that are difficult to diagnose with AI, allow experts to validate the AI diagnosis, and allow them to interact and counsel patients.

In 2018, we launched a project to use AI for early detection of DR in India to meet the above objectives. The project was proposed at the ITU/UN AI for Good Summit in Geneva in May 2018.(7) The proposal was aligned with the objective of the conference which was to use AI to accelerate the achievement of the UN's Sustainable Development Goals (SDGs), in the area of health. It was also presented at the ITU/WHO Focus Group on AI for Health (FGAI4H) at Columbia university in NYC, in Nov 2018, (8) and the proposal was accepted as one of the first eight use cases of AI for health by the ITU/WHO. (9)

Solution Overview:

An overview of the solution and its major components is as follows:

- Patient's retinal images are captured via fundus cameras at local screening centers or clinics and are uploaded via the web to a cloud-based server for further processing. A low-cost device that can be attached to a mobile phone for image capture of the retina is also being proposed for remote areas where no clinical facilities exist.

- A cloud-based web application for patient registration and data entry, image capture and uploading, integration with the AI model, remote diagnosis by trained specialists, as well as patient reporting, messaging and notification.

- Automated DR Detection: The AI model runs on a remote server and automatically processes and classifies the image as gradable or not and if gradable whether it is referable DR or not along with the probability (or confidence level) of the classification.

- Remote Diagnosis: Eye-care professionals can login remotely to the web application to review and validate the AI diagnosis, add notes, provide referral to a specialist, follow-up and treatment options. The system design will also incorporate the ability to fine-tune the AI algorithm based on corrections of diagnosis errors by the specialists.

- Integrated administrative, reporting, and messaging for patient communication, system performance reports, and overall statistics.

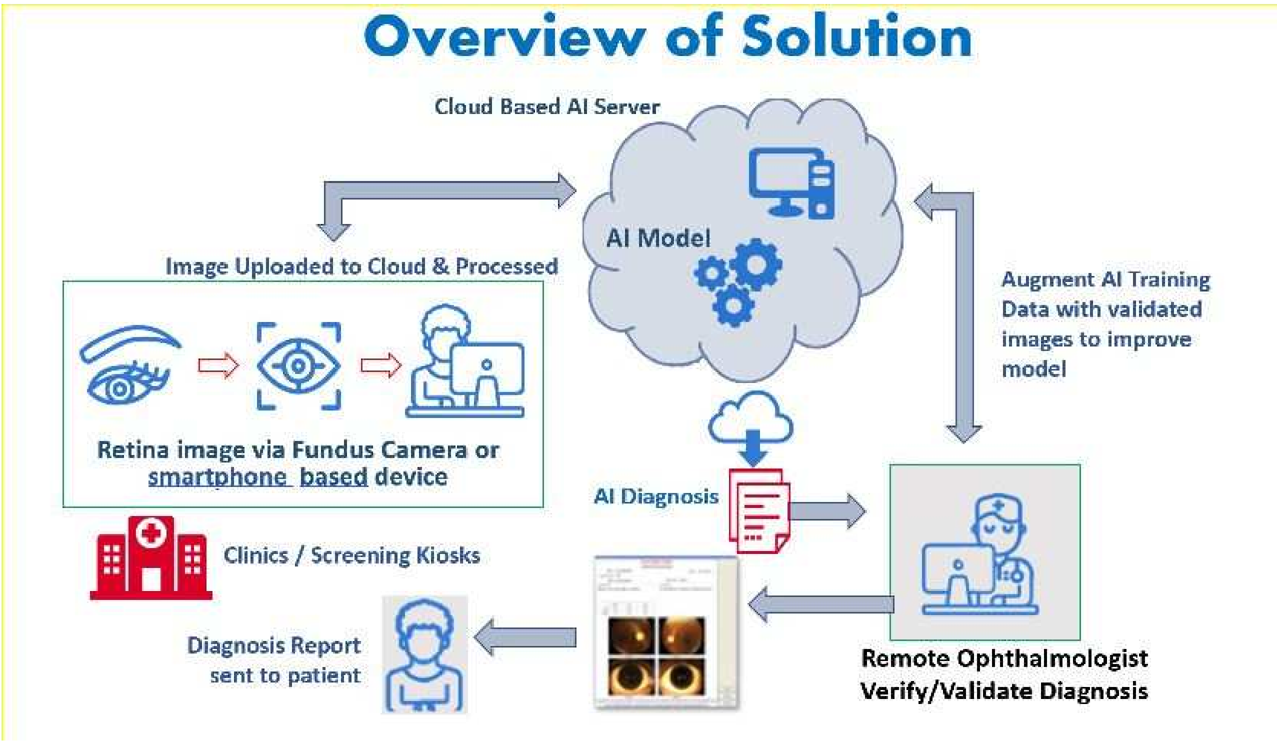

Implementation:

We partnered with a leading teleophthalmology company in India with over 275 screening centers in 22 states that screens over 25,000 patients a month for DR. It has a national network of ophthalmologists who perform the screening and diagnosis of DR remotely via a cloud-based application.

The creation of an AI model requires a sufficiently large dataset of images labeled with the actual diagnosis classifications previously obtained from human experts. The labeled data is then used to train an AI model using supervised learning.

The model uses a convolutional neural network (CNN), and the training process fine tunes the weights of the network to minimize the error between the predicted classification and the ground truth. A part of the dataset, called the validation data, is not used for training, but used to validate the performance of the AI model on unseen data.

We started by first creating a dataset of approximately 90,000 images from our teleophthalmology partner, that had been graded by licensed ophthalmologists into one of the following categories: non-gradable, no retinopathy, mild NPDR, moderate NPDR, severe NPDR, and PDR, along with presence or absence of DME. The non-gradable category implied a low-quality image that could not be graded and assigned a diagnosis. All images were obtained with appropriate informed consent and anonymized prior to use in training.

For the purposes of the first phase of this project the scope was limited to determining if the DR was referable or not. Therefore, the images were regrouped as:

i) Non-gradable
ii) Non-referable DR – which included No retinopathy and mild NPDR and
iii) Referable DR – which included Moderate NPDR, Severe NPDR, PDR, and DME.

Approximately 80% of the dataset was used for training the model and 20% used for validation or testing of the model's performance.

The model was trained till it reached 92% accuracy on the validation dataset. The sensitivity of the model at the end of training was 92% and specificity 93%. Sensitivity refers to the proportion of positive (referable DR) cases that the model identifies correctly, and specificity refers to the proportion of negative (non-referable DR) cases that the model identifies correctly.(10)

The AI model is currently being further tested and validated with real data while it is undergoing field testing. The level of accuracy achieved by the model during testing is comparable to ophthalmologists and is considered acceptable for screening of DR. In addition to internal validation, the AI model has also been submitted to the ITU/WHO's Focus Group on AI for Health, which has the goal of benchmarking AI for health algorithms and provide a neutral, independent assessment of performance. (11)

Challenges Faced:

In designing the system, we faced several challenges:

i) Data Curation: The first challenge was to ensure that the data used for training the AI algorithm was clean. This is necessary to ensure that the trained model is accurate in making predictions on new images. Many large real-word datasets contain invalid data and labels. For example, in our case the data contained many anterior (outside) images of eyes. Images also varied in quality and included out-of-focus and low-quality images. We created a separate AI model to remove invalid and very low-quality images from the training dataset. This improved the overall accuracy, sensitivity and specificity of the AI model.

ii) Speed and Scalability: Design objectives for the AI system included fast prediction and scalability to hundreds of locations. By fine-tuning the model and parameters, we achieved a response time of less than a second on returning a prediction from the model. The system is hosted on Google Cloud Platform (GCP) infrastructure, to ensure scalability when deployed across hundreds of locations.

iii) Data Privacy, Security & Compliance: We designed the system to comply with Indian regulations, confidentiality and security by using informed consent, data privay anonymization, and encryption, as proposed by the Digital Information Security in Healthcare Act ("DISHA") and the Personal Data Protection Bill, 2018. India does not yet have formal regulations on using AI for health but requires all diagnostic reports to be reviewed and signed by a licensed doctor. Since the AI system will be used in assistive mode, all reports will be reviewed and validated by an ophthalmologist.

Results:

The AI solution has achieved clinically acceptable levels of accuracy in initial testing and field trials in India and is ready for deployment on a larger scale. While regulatory approval for fully automated screening is being sought it will be used in an assistive mode as follows:

i) Improve screening throughput: Currently 15% to 20% of images uploaded are non-gradable or invalid due to operator error. The AI system can instantly catch these errors and require the operator to capture a gradable image. This will reduce delays in diagnosis and improve overall throughput by 15% to 20%.

ii) Triage and prioritize screening: The AI system can identify higher risk cases and prioritize based on DR severity – with PDR, severe, and moderate cases receiving immediate attention and priority screening. This will improve overall level of care for those most at risk of vision loss.
iii) Quality Assurance: In cases where the AI and ophthalmologist's diagnosis differ, the case can be automatically assigned to a second ophthalmologist for further review. This will improve overall accuracy and decrease errors.

We also plan to conduct field trials of the model in countries in Africa, and the far east where there is an acute shortage of ophthalmologists. Future plans include obtaining FDA and EMA approvals for launch In USA and Europe.

Conclusion:

An integrated system using AI can be deployed at scale and be effective for early detection and screening of DR, a major cause of preventable blindness worldwide AI based systems for detection of DR offer the following potential benefits:

- Bridge the shortage of healthcare professionals and provide access to screening where none exists.
- Increase overall efficiency and scalability of current screening methods.
- Provide earlier detection of DR thereby preventing vision loss for millions.
- Decrease overall health-care costs via earlier interventions when it is easier and less expensive to treat these diseases.

AI based systems would be very useful in countries such as India and developing nations in Africa and elsewhere which lack professionals and infrastructure to screen everyone at risk of vision loss.

With over 420 million people afflicted with diabetes worldwide and 148 million with DR, and the numbers increasing each year, AI powered systems will be critical to address the global healthcare challenge of DR and prevent vision loss for millions globally.

References:

(1) World Health Organization, WHO Global Report on Diabetes.
https://www.who.int/diabetes/global-report/en/

(2) Yau JW, Rogers SL, Kawasaki R, Lamoureux EL, Kowalski JW, Bek T, et al. Global prevalence and major risk factors of diabetic retinopathy. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3322721/

(3) International Diabetes Federation
https://idf.org/our-network/regions-members/south-east-asia/members/94-india.html

(4) International Council of Ophthalmology. Number of Ophthalmologists in Practice and Training Worldwide.
http://www.icoph.org/ophthalmologists-worldwide.html

(5) IMAGENET. Large Scale Visual Recognition Challenge 2017 (ILSVRC2017).
http://image-net.org/challenges/LSVRC/2012/index

(6) Varun Gulshan, PhD; Lily Peng, MD, PhD; Marc Coram, PhD; et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs.
https://jamanetwork.com/journals/jama/fullarticle/2588763

(7) Presentation at AI for Good Summit, AI for Primary Care and Service Delivery, May 2018, Geneva,
https://www.itu.int/en/ITU-T/AI/2018/Pages/programme.aspx

(8) Presentation at ITU/WHO Focus Group on AI for Health, Columbia University, NY Nov 2018,
https://www.itu.int/en/ITU-T/Workshops-and-Seminars/20181114/Pages/programme.aspx

(9) ITU News, Artificial Intelligence for Health: ITU and WHO accept 8 new use cases,
https://news.itu.int/artificial-intelligence-health-new-use-cases/

(10) Wikipedia, Sensitivity and specificity, https://en.wikipedia.org/wiki/Sensitivity_and_specificity

(11) ITU Website, Focus Group on "Artificial Intelligence for Health,
https://www.itu.int/en/ITU-T/focusgroups/ai4h/Pages/default.aspx

About the Author:

Arun Shroff is a serial entrepreneur who has founded and grown multiple technology companies in USA and India. He s founder/CEO of Xtend.AI, a startup using AI to solve global health challenges in health and other domains. He is also a co-founder/CTO at Medindia.net, one of India's top health websites, empowering consumers, patients and doctors with information and tools to better manage health. Arun serves as a topic lead at the ITU/WHO Focus Group on AI for Health - a UN / WHO initiative for benchmarking AI and creating a global community on AI for health. Arun is also a Director of Technology & Innovation at STAR Associates, USA providing strategic and advisory services to early stage technology companies in AI, Robotics, Blockchain, cybersecurity, health and telemedicine.

Arun has been an invited speaker and expert on AI and technology at conferences globally including at the AI for Good Summit in Geneva, at the ITU/WHO Focus Group in New York, at Wilton Park, UK, at Apollo Hospitals, Chennai, and IIT Madras. Arun also volunteers with many non-profits including MOHAN USA, where is a founding trustee and an advisor at MOHAN Foundation - both promoting ethical organ donation and IIMPACT, an education NGO for girl children in India's poorest villages. Arun has as a B.Tech from IIT Madras, an MBA from IIM, Ahmedabad, and M.S. in Computer Science from Penn State University, and certifications in Machine Learning & AI/Deep Learning from Stanford & Coursera.

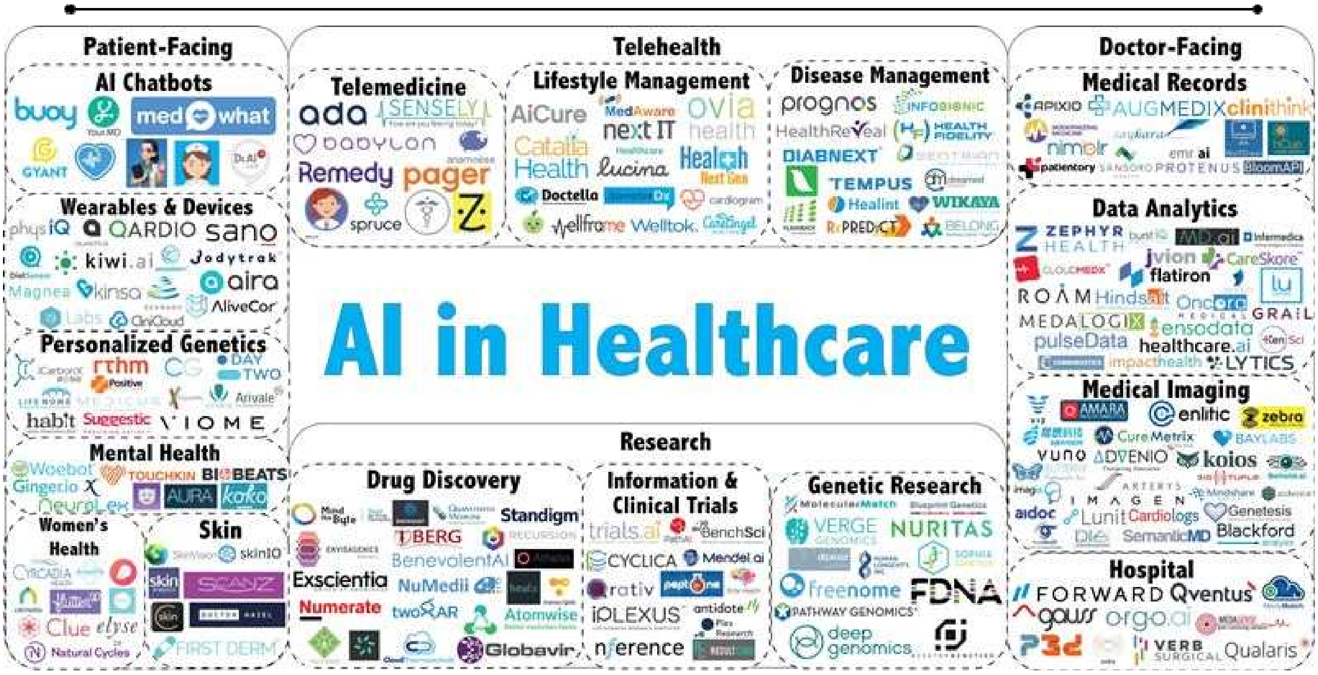

Source & Courtesy: AI in Healthcare: Industry Landscape
https://techburst.io/ai-in-healthcare-industry-landscape-c433829b320c

Artificial intelligence in healthcare: past, present and future: Artificial intelligence (AI) aims to mimic human cognitive functions. It is bringing a paradigm shift to healthcare, powered by increasing availability of healthcare data and rapid progress of analytics techniques. We survey the current status of AI applications in healthcare and discuss its future. AI can be applied to various types of healthcare data (structured and unstructured). Popular AI techniques include machine learning methods for structured data, such as the classical support vector machine and neural network, and the modern deep learning, as well as natural language processing for unstructured data. Major disease areas that use AI tools include cancer, neurology and cardiology. We then review in more details the AI applications in stroke, in the three major areas of early detection and diagnosis, treatment, as well as outcome prediction and prognosis evaluation. We conclude with discussion about pioneer AI systems, such as IBM Watson, and hurdles for real-life deployment of AI. http://bit.ly/2kYdvoj

Artificial Intelligence in Healthcare: the Ultimate Guide: Artificial intelligence (AI) plays a role in many industries, from banking and cybersecurity to product design and healthcare. The benefits of leveraging technology in healthcare have the power to impact both your facility and patients. Some implementations include diagnostic capabilities and predicting disease, customized treatment plans, enhanced electronic health records and more. This guide has everything you need to know about artificial intelligence in healthcare. http://bit.ly/2mulRoj