Check for updates

STUDY PROTOCOL

Automated analysis of retinal imaging using machine learning techniques for computer vision [version 1; referees: 2 approved]

Jeffrey De Fauw[1], Pearse Keane[1,2], Nenad Tomasev[1], Daniel Visentin[1], George van den Driessche[1], Mike Johnson[1], Cian O Hughes[1], Carlton Chu[1], Joseph Ledsam[1], Trevor Back[1], Tunde Peto[2], Geraint Rees[3], Hugh Montgomery[5], Rosalind Raine[4], Olaf Ronneberger[1], Julien Cornebise[1]

[1]Google DeepMind, London, EC4A 3TW, UK
[2]Moorfields Eye Hospital NHS Foundation Trust, London, EC1V 2PD, UK
[3]Alexandra House University College London, Bloomsbury Campus, London, WC1N 3AR, UK
[4]Department of Applied Heath Research, University College London, London, WC1E 7HB, UK
[5]Institute of Sport, Exercise and Health, London, W1T 7HA, UK

Abstract

There are almost two million people in the United Kingdom living with sight loss, including around 360,000 people who are registered as blind or partially sighted. Sight threatening diseases, such as diabetic retinopathy and age related macular degeneration have contributed to the 40% increase in outpatient attendances in the last decade but are amenable to early detection and monitoring. With early and appropriate intervention, blindness may be prevented in many cases.

Ophthalmic imaging provides a way to diagnose and objectively assess the progression of a number of pathologies including neovascular ("wet") age-related macular degeneration (wet AMD) and diabetic retinopathy. Two methods of imaging are commonly used: digital photographs of the fundus (the 'back' of the eye) and Optical Coherence Tomography (OCT, a modality that uses light waves in a similar way to how ultrasound uses sound waves). Changes in population demographics and expectations and the changing pattern of chronic diseases creates a rising demand for such imaging. Meanwhile, interrogation of such images is time consuming, costly, and prone to human error. The application of novel analysis methods may provide a solution to these challenges.

This research will focus on applying novel machine learning algorithms to automatic analysis of both digital fundus photographs and OCT in Moorfields Eye Hospital NHS Foundation Trust patients.

Through analysis of the images used in ophthalmology, along with relevant clinical and demographic information, Google DeepMind Health will investigate the feasibility of automated grading of digital fundus photographs and OCT and

Open Peer Review

Referee Status: ✓ ✓

	Invited Referees	
	1	**2**
version 1 published 05 Jul 2016	✓ report	✓ report

1 **Yit Yang**, Royal Wolverhampton NHS Trust UK

2 **Sandrine Zweifel**, University Hospital Zürich Switzerland

Discuss this article

Comments (0)

provide novel quantitative measures for specific disease features and for monitoring the therapeutic success.

Competing interests: Moorfields Eye Hospital NHS Foundation Trust administration time spent on this work will be paid to the trust. The Chief Investigator and some co-investigators are paid employees of Google DeepMind. Several co-investigators (PK, JL, GR, RR) are paid contractors for Google DeepMind.

Background

Age-related macular degeneration (AMD) is a degenerative retinal disease that can cause irreversible visual loss (Bressler, 2004). It is the leading cause of blindness in Europe and North America and accounts for over half of partially sighted or legally blind certifications in the UK (Bunce, 2010). Neovascular ("wet") AMD is an advanced form of macular degeneration that historically has accounted for the majority of vision loss related to AMD. It is characterised by abnormal blood vessel growth that can result in hemorrhage, fluid exudation and fibrosis, and thus to local macular damage and ultimately vision loss (Owen, 2012).

Diabetic retinopathy (DR) is the leading cause of blindness in working age populations in the developed world (Cheung, 2010). It is estimated that up to 50% of people with proliferative DR (characterised by neovascularisation) who do not receive timely treatment will become legally blind within 5 years (Shaw, 2010). Although up to 98% of severe visual loss due to DR can be prevented with early detection and treatment, once it has progressed vision loss is often permanent (Kollias, 2010). Indeed, 4,200 people in England every year are at risk of blindness caused by diabetic retinopathy and there are 1,280 new cases of blindness caused by diabetic retinopathy (Scanlon, 2008).

To diagnose these conditions and monitor their progression (and response to treatment) the presence and precise location of the lesions must be determined. Two imaging modalities are commonly used for this purpose: digital photographs of the fundus (the 'back' of the eye) and Optical Coherence Tomography (OCT, a modality that uses light waves in a similar way to how ultrasound uses sound waves) (Huang, 1991).

Clinics such as these have generated very large datasets of both digital fundus and OCT images. They are also very busy; ophthalmology clinics often have long waits and average times to urgent first treatments can be greater than two weeks in the busiest clinics.

Machine learning algorithms make use of the rich, varied datasets to find high dimensional interactions between multiple data points (Murphy, 2012). Most machine learning methods can be thought of as a form of statistics. Some algorithms look for patterns in the data, identifying subgroups within the full sample (known as clustering). Others rely on specific clinical information against which the algorithm compares its predictions and adjusts accordingly (supervised learning).

There have been significant recent advances in the field of machine learning demonstrating algorithms able to learn how to accomplish tasks without instruction (Mnih, 2015; Silver, 2016). Recent healthcare applications of such algorithms have shed light on complex genetic interactions in autism (Uddin, 2014) and monitoring of physiological observations in intensive care (Clifton, 2013).

This study aims to combine traditional statistical methodology and machine learning algorithms to achieve automatic grading and quantitative analysis of both digital fundus photograph and OCT in Moorfields Eye Hospital NHS Foundation Trust patients (London, UK). Moorfields Eye Hospital is the leading provider of eye health services in the UK and a world-class centre of excellence for ophthalmic research and education. Should the research be successful, implementation of the outcomes would improve patient access to treatment and ease pressures on time and resources in ophthalmology clinics.

Aims and objectives

Primary objective

1.1
Exploratory study: investigate whether computer algorithms can detect and classify pathological features on eye imaging, including fundus digital photographs and OCT.

Secondary objectives

If the exploratory study is successful:

2.1
To provide novel image analysis algorithms to identify and quantify specific pathological features in eye imaging, using validated methods and expert clinical consensus.

2.2
To provide quantitative measurements disease progression, severity and to monitor the therapeutic success over time.

Study design

This is a retrospective, non-interventional exploratory study. Analyses performed in the study will be on fully anonymised (to achieve the primary objective) and pseudonymised (to achieve the secondary objective 2.2) retinal images (including fundus images and OCT). These images will contain no patient identifiable information.

Inclusion criteria

All patients attending Moorfields Eye Hospital NHS Foundation Trust sites between 01/01/2007 and 29/02/2016, and who had digital retinal imaging (including fundus digital photographs and OCT) as part of their routine clinical care, will be eligible for inclusion in this study.

Exclusion criteria

Hard copy examinations (i.e. physical photograph copies prior to digital storage) will be ineligible, and will be excluded by the nature of the original service evaluation requests from the Moorfields team. Data from patients who have previously manually requested that their data should not be shared, even for research purposes in anonymised form, and have informed Moorfields Eye Hospital of this, will be ineligible and removed by Moorfields Eye Hospital staff before research begins.

Sample size

Approximately 1 million examinations meet the above criteria.

Most recent machine learning algorithms benefit from large datasets on which to train (tens to hundreds of thousands of data instances (Silver, 2016)). Across all machine learning applications the predictive power (as percentage of data instances correctly classified) of the algorithm depends on the size and quality of the dataset.

The sample size is informed by the existing literature and by DeepMind's previous work in the field of machine learning (Mnih, 2015; Silver, 2016). Today's most powerful deep neural networks can have millions or billions of parameters, so large amounts of data are needed to automatically infer those parameters during learning. Most problems in the medical domain are highly complex as they arise as an interplay of many clinical, demographic, behavioural and environmental factors that are correlated in non-trivial ways. This is even more true for state-of-the art deep learning methodologies that are expected to give the best results (Szegedy, 2014).

Data

For all patients meeting inclusion/exclusion criteria the following electronic health record data will be required to complete this project successfully:

(1) Digital Fundus Photographs

(2) Digital OCT images

In addition to image data, the anonymised dataset will contain additional information required to train an algorithm (objective 1.1):

- Demographic information shown to be associated with eye disease. This is because the retina differs by features such as age, as do the likelihoods, manifestations and progression of specifc disease states.

- Primary and secondary diagnostic labels describing what pathology is in the image (e.g., wet AMD, diabetic retinopathy) and the associated severity (e.g., grade of retinopathy/maculopathy in diabetic retinopathy).

- Treatment information describing aspects of management alongside pathology information and clinical data such as visual acuity.

- Model of imaging device.

A second dataset will be pseudonymised to allow further investigation of disease progression and treatment effects. This will include the above information, and add temporal information allowing pseudonymised data to be joined over time with knowledge of the elapsed time between each:

- Time elapsed between each data point and the first visit in weeks (Sunday to Sunday).

- The dates of each scan will be removed so it will be impossible to identify an individual date or year of patient's birth.

The anonymisation and pseudonymisation procedures adopted will remove any information not specified to further avoid transfer of patient identifiable information. All anonymisation and pseudonymisation will be formally checked by Moorfields Eye Hospital staff before transfer.

Algorithm development

In order to develop the algorithm, Google DeepMind will work with the eye images, split by the known diagnoses, using machine learning and Artificial Intelligence techniques including but not limited to: supervised and semi-supervised convolutional neural networks, recurrent neural networks, unsupervised clustering, reinforcement learning (Murphy, 2012).

For a selection of images additional manual labels will be produced by experts to allow investigation of the clinical benefit that can be achieved through the use of machine learning in eye imaging. Pathological and anatomical features will be annotated by trained graders at Moorfields Eye Hospital and overseen by a consultant ophthalmologist with over 10 years of experience.

Statistical analysis

Descriptive statistics will be used to describe specific pathological features extracted by the model on retinal images, both continuous and categorical (e.g. presence and size of an abnormality), and to compare them to an expert references.

In addition we will analyse the outcomes against recent Moorfields audit data on human performance to further understand the potential impacts of the model in clinical practice.

Data protection

Anonymisation and pseudonymisation

This study requires existing retrospective data only; no prospective data are needed nor will be collected from patients, hospitals or healthcare workers. No direct patient contact will occur and necessary data will be anonymised or pseudonymised from this source dataset.

Anonymisation and pseudonymisation of all image files and clinical information is performed and validated by Moorfields Eye Hospital staff at Moorfields Eye Hospital. No patient identifiable data will be transferred to Google DeepMind.

During validation of the anonymisation procedure it was noted that the current anonymisation tool (*TopCon IDConvert*) in use at Moorfields had the potential to leave information that may identify the patient in a small number of datasets. DeepMind collaborators worked with Moorfields Eye Hospital IT team to develop an anonymisation script that reliably deletes all of these information in a second step. This new script will be used in the project and will be run by Moorfields Eye Hospital NHS Foundation Trust staff at Moorfields Eye Hospital. Google DeepMind will not have access to patient identifiable information at any time.

Data storage

Google DeepMind Health has developed and established a state-of-the-art secure patient information handling service utilising Common Criteria EAL4 compliant firewalls and on-disk encryption (using Advanced Encryption Standard with a 256-bit key) of all research data, all housed within an ISO 27001 compliant data centre. After anonymisation data will be transferred to our London, UK data centre. This data handling facility conforms to NHS HSCIC Information Governance Statement of Compliance Toolkit (formally assessed at level 3).

Access will be granted by the custodian of data and no other members of the team. Only those working directly on the data in a research capacity will have access.

Data destruction
The data sharing agreement between Google DeepMind and Moorfields Eye Hospital NHS Foundation Trust lasts for 5 years. After this period the agreement will be reviewed should future work seek to build on this project. After the data sharing agreement expires all data used in the study will be destroyed. No modification will be made based on the data after destruction.

Data destruction will involve the deletion of the encryption/decryption keys for all project volumes, and 8-way random data write to all physical disks within the Google DeepMind Health data infrastructure. A certificate of destruction will be provided to the Trust.

The algorithms developed during the study will not be destroyed. Google DeepMind Health knows of no way to recreate the patient images transferred from the algorithms developed. No patient identifiable data will be included in the algorithms.

Ethical considerations
Ethical approvals
The research on anonymised data received formal approval from the Moorfields Eye Hospital research and development office, responsible for approving studies working with anonymised data, on 16 Oct 2015 (reference 15/050).

The research on the pseudonymised dataset received formal Research Ethics Committee approval on 03 Jun 2016 (reference 16/EE/0253). Health Research Authority approval for this is in progress and no research on pseudonymised data will begin until this is approved.

Consent
No patient will be approached directly and this work will include no direct patient contact. Only anonymised or pseudonymised retrospective data collected as part of routine clinical care are included. No patient identifiable data will be collected as part of this study. In such cases the ICO code of practice states that explicit consent is not generally required (ICO, 2012).

Adverse events
The project is non-interventional and does not involve any direct patient contact. We do not anticipate any change to patient management. For pseudonymised data should expert assessment during model validation highlight a clinical error this will be raised to the appropriate clinical team. The primary point of call for adverse events should they arise will be the Trust information governance lead, and the responsible clinical team at Moorfields Eye Hospital will be notified.

Monitoring
The study will be monitored both internally and externally. Internally Google DeepMind managers (TB, JL) will oversee and monitor progress on a day-to-day basis, ensuring the protocol is adhered to and no compliance issues arise.

Clinical and methodological experts (PK, GR, RR) are working with Google DeepMind to further oversee the ethical, clinical and methodological considerations of the project and will advise on at least a weekly basis to ensure no deviation from the described protocol.

Externally the information governance team at the Moorfields Eye Hospital will be consulted before commencing data collection, and weekly thereafter to ensure no deviation from the described protocol.

Access
Google DeepMind has access to the required data to support the research aims of this study. To ensure compliance with the common law principle of data confidentiality, Google DeepMind will only receive anonymised or pseudonymised data from the Trust. DeepMind works with Moorfields Eye Hospital to ensure accuracy and clarity in the data to allow useful and consistent interpretation at all times.

Dissemination
The results will be disseminated through normal academic channels, initially focusing on conference proceedings and the indexed peer reviewed literature relevant to the fields of machine learning, artificial intelligence, ophthalmology and clinical research. Patient and Public Involvement representatives will be involved at each stage of the research.

Conclusion
We propose an exploratory study and initial testing of machine learning algorithms to analyse and quantify pathological and anatomical features in eye images. The results will be compared to expert annotations.

Should development be successful we plan to submit additional applications for permission to refine the model if required and ultimately investigate the performance in a clinical implementation setting.

Author contributions

References

Bressler NM: **Age-related macular degeneration is the leading cause of blindness.** *JAMA*. 2004; **291**(15): 1900–1.
PubMed Abstract

Bunce C, Xing W, Wormald R: **Causes of blind and partial sight certifications in England and Wales: April 2007-March 2008.** *Eye (Lond)*. 2010; **24**(11): 1692–9.
PubMed Abstract | **Publisher Full Text**

Cheung N, Mitchell P, Wong TY: **Diabetic retinopathy.** *Lancet*. 2010; **376**(9735): 124–136.
PubMed Abstract | **Publisher Full Text**

Clifton L, Clifton DA, Pimentel MA, *et al.*: **Gaussian processes for personalized e-health monitoring with wearable sensors.** *IEEE Trans Biomed Eng*. 2013; **60**(1): 193–197.
PubMed Abstract | **Publisher Full Text**

Huang D, Swanson EA, Lin CP, *et al.*: **Optical coherence tomography.** *Science*. 1991; **254**(5035): 1178–81.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Information Commissioner's Office: **Anonymisation: managing data protection risk code of practice.** 2012; Accessed 2016.
Reference Source

Kollias AN, Ulbig MW: **Diabetic retinopathy: Early diagnosis and effective treatment.** *Dtsch Arztebl Int*. 2010; **107**(5): 75–84.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Mnih V, Hinton G: **Learning to Detect Roads in High-Resolution Aerial Images.** European Conference on Computer Vision. 2010; **6316**: 210–223.
Publisher Full Text

Murphy KP: **Machine Learning: A Probabilistic Perspective.** *Adaptive Computation and Machine Learning Series*. Cambridge, Mass.: MIT Press. 2012.
Reference Source

Owen CG, Jarrar Z, Wormald R, *et al.*: **The estimated prevalence and incidence of late stage age related macular degeneration in the UK.** *Br J Ophthalmol*. 2012; **96**(5): 752–56.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Scanlon PH: **The English national screening programme for sight-threatening diabetic retinopathy.** *J Med Screen*. 2008; **15**(1): 1–4.
PubMed Abstract

Shaw JE, Sicree RA, Zimmet PZ: **Global estimates of the prevalence of diabetes for 2010 and 2030.** *Diabetes Res Clin Pract*. 2010; **87**(1): 4–14.
PubMed Abstract | **Publisher Full Text**

Silver D, Huang A, Maddison CJ, *et al.*: **Mastering the game of Go with deep neural networks and tree search.** *Nature*. 2016; **529**(7587): 484–489.
PubMed Abstract | **Publisher Full Text**

Szegedy C, Liu W, Jia Y, *et al.*: **Going deeper with convolutions.** arXiv:1409.4842v1. 2014.
Reference Source

Uddin M, Tammimies K, Pellecchia G, *et al.*: **Brain-expressed exons under purifying selection are enriched for** *de novo* **mutations in autism spectrum disorder.** *Nature Genetics*. 2014; **46**(7): 742–747.
PubMed Abstract | **Publisher Full Text**

Open Peer Review

Current Referee Status: ✔ ✔

Version 1

Referee Report 10 October 2016

Sandrine Zweifel

Department of Ophthalmology, University Hospital Zürich, Zürich, Switzerland

Thank you very much for the opportunity to review the manuscript by De Fauw and co-workers. The topic is of great interest. Moorefields recently announced their research partnership with DeepMind Health. In this manuscript the project plan is detailed. One million examinations (digital fundus photographs and OCT scans) which met the inclusion criteria will be analysed using machine learning and Artificial Intelligence techniques.

The authors did not specify why they did not include angiography as an additional imaging modality which is usually used at baseline or at follow-up exams in patients with AMD and diabetic retinopathy. Please add this information.

Since data security is an important issue to be discussed especially when evaluating such large data sets the authors need to provide more information regarding patient consent. Does everyone who is examined at Moorefields Eye Hospital give a general consent for evaluating their data? Patients are not specifically informed about the project with DeepMind Health, aren't they? So data of patients are only excluded if patients previously (independent of this project) requested that there data should not be shared.

Although there might be some risk regarding "data security" it outweighs the potential of earlier detection and treatment of million of patients.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Competing Interests: No competing interests were disclosed.

Referee Report 19 July 2016

Yit Yang

Royal Wolverhampton NHS Trust, Wolverhampton, UK

There is a good introduction of the two main conditions, AMD and DR and their burden on public health.

Very little information on machine learning to convince the reader that the fundamental processes and pieces of the jigsaw are already available to apply to AMD and DR or in simpler healthcare scenarios.

How does MEH plan to locate the ID of those who have given instructions for their images NOT to be used? Is there a database for this specific parameter?

Good attention to data security and data protection.

Not very much information on specific data outcomes expectations. Eg how detailed will be grading be of a wet AMD lesion.

Overall, a novel concept and worth exploring as it will be able to replace human workforce if successful.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Competing Interests: No competing interests were disclosed.