Artificial Intelligence in Ophthalmology

Reinventing the Eye Exam!

Pearse Keane
Moorfields Eye Hospital and
UCL Institute of Ophthalmology

Disclosure

Speaker Fees:

- *Topcon*
- *Heidelberg*
- *Haag-Streit*
- *Allergan*
- *Bayer*
- *Novartis*
- *Zeiss*

Consultancy:

- *DeepMind*
- *Optos*

Examining the Eye

Ophthalmoscopy

Ophthalmoscopy

Herman von Helmholtz

Ophthalmoscopy

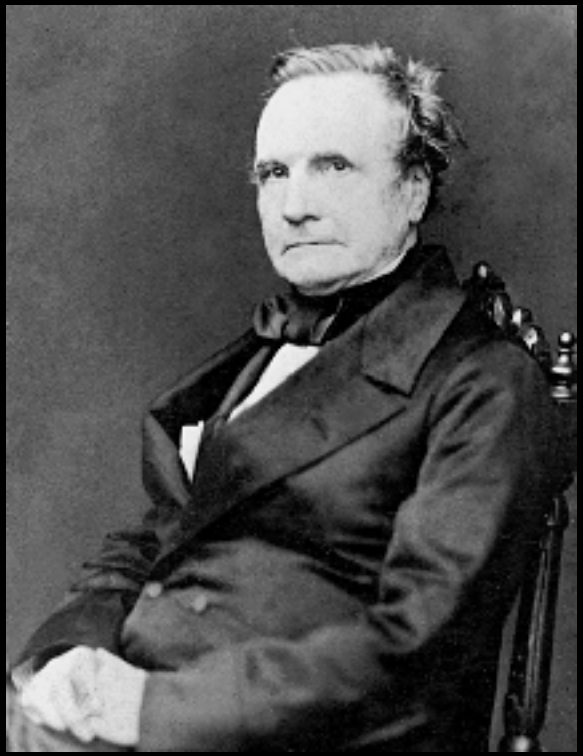

Charles Babbage

Ophthalmic Imaging

Ophthalmic Imaging

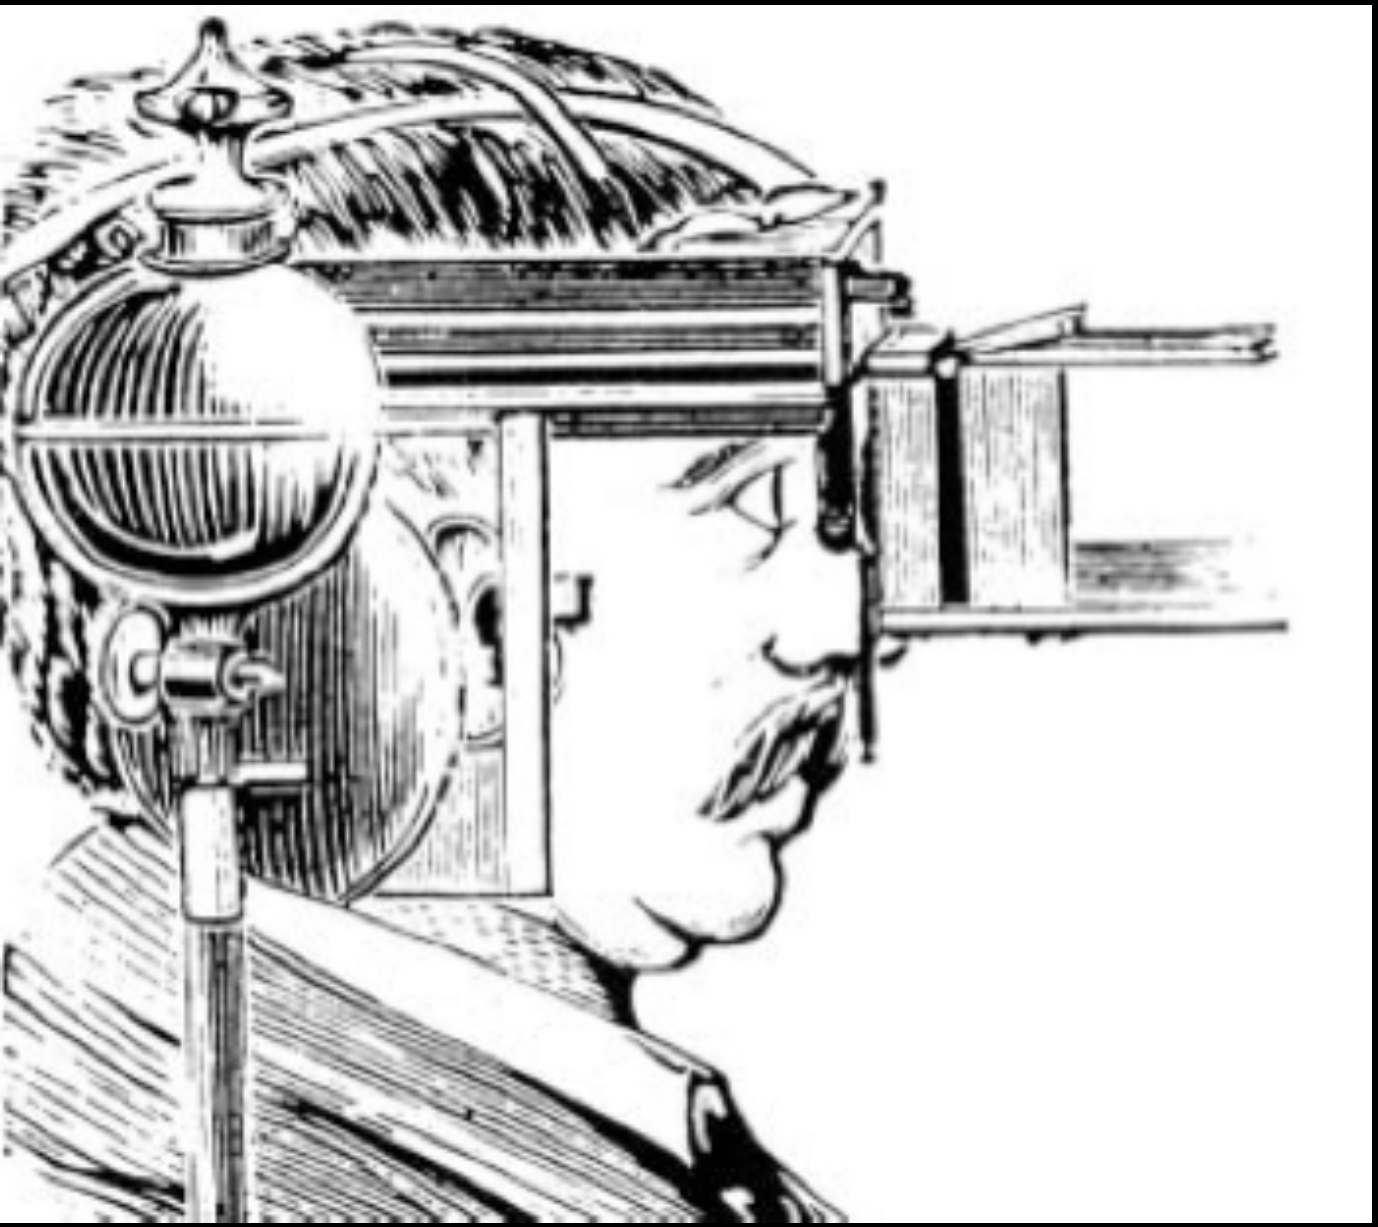

Jackman and Webster

Ophthalmic Imaging

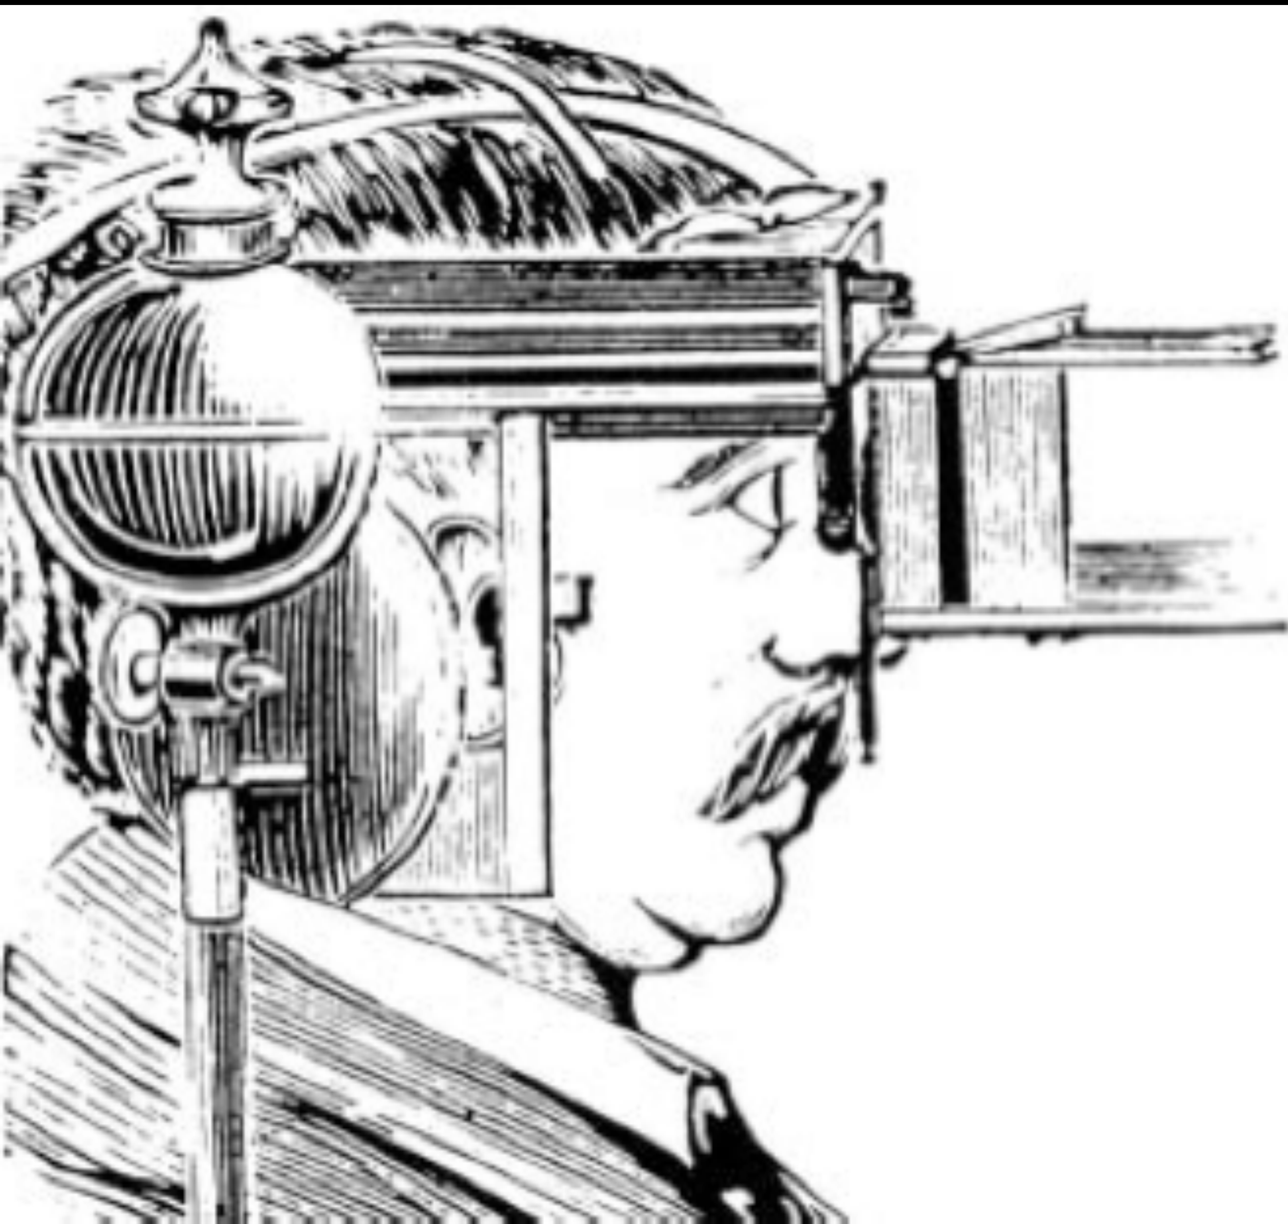

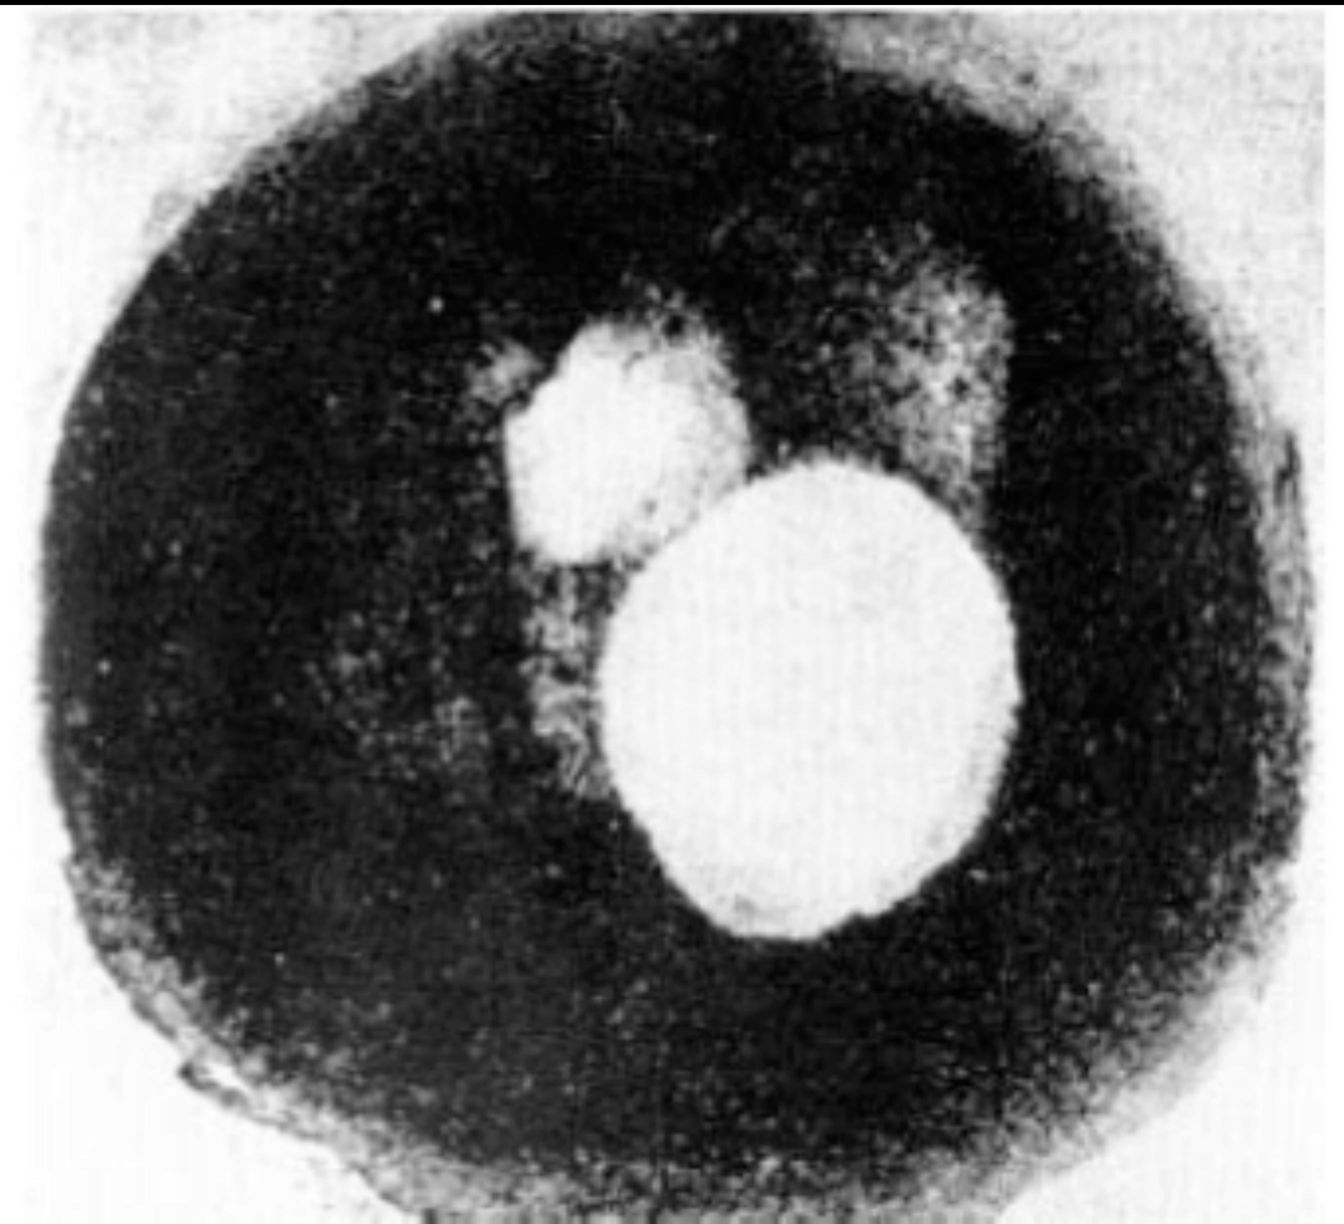

Jackman and Webster

Ophthalmic Photography

Ophthalmic Photography

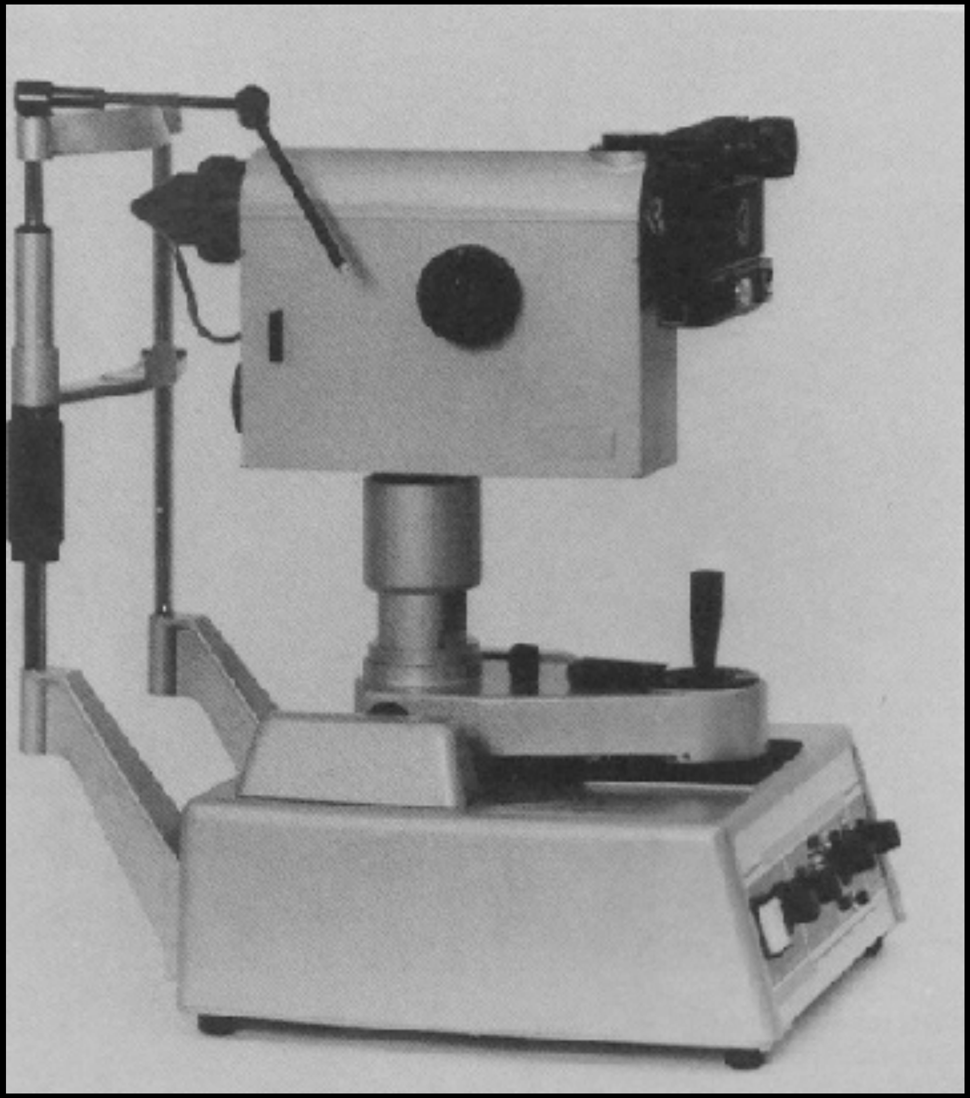

1960s

Ophthalmic Photography

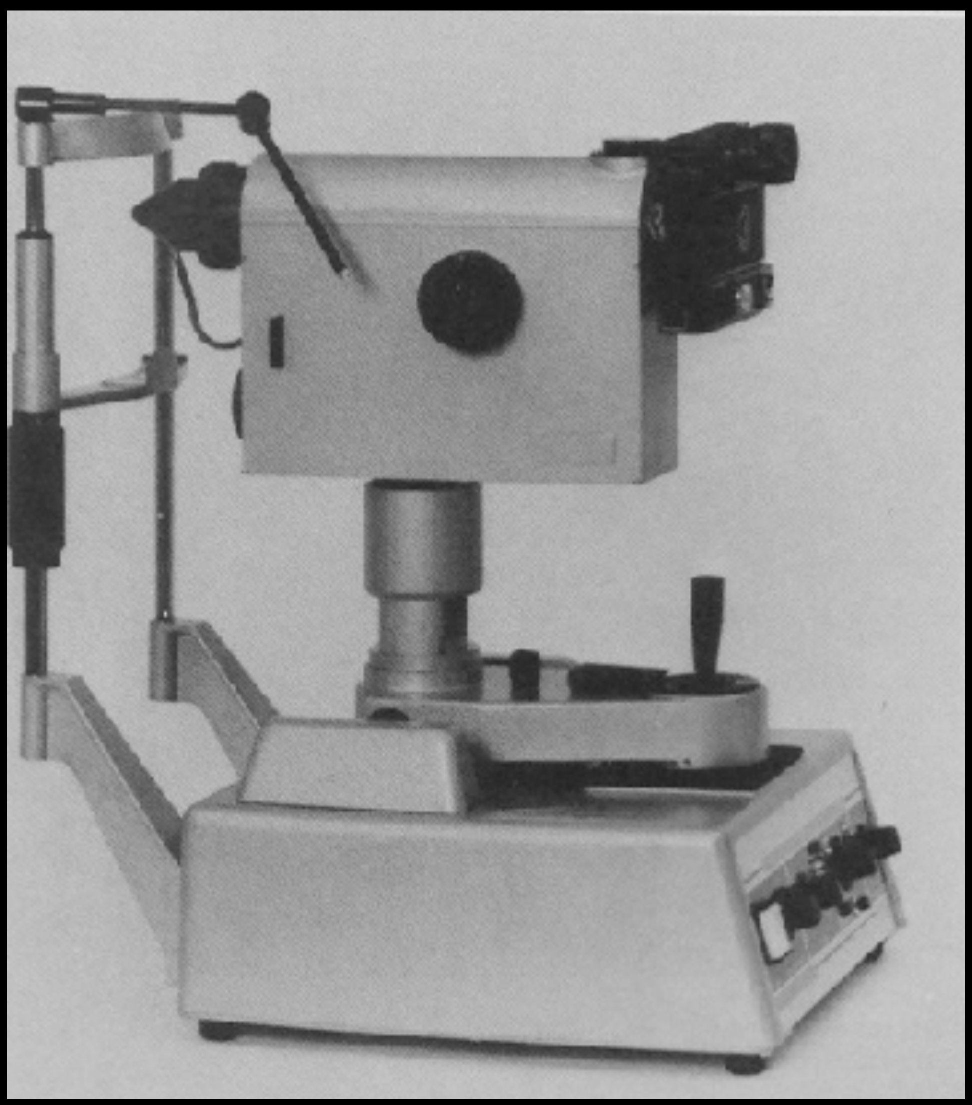

1960s

2000s

Ophthalmic Photography

Breakthrough

Breakthrough

Science 1991

Reports

Optical Coherence Tomography

David Huang, Eric A. Swanson, Charles P. Lin, Joel S. Schuman, William G. Stinson, Warren Chang, Michael R. Hee, Thomas Flotte, Kenton Gregory, Carmen A. Puliafito, James G. Fujimoto*

A technique called optical coherence tomography (OCT) has been developed for noninvasive cross-sectional imaging in biological systems. OCT uses low-coherence interferometry to produce a two-dimensional image of optical scattering from internal tissue microstructures in a way that is analogous to ultrasonic pulse-echo imaging. OCT has longitudinal and lateral spatial resolutions of a few micrometers and can detect reflected signals as small as $\sim 10^{-10}$ of the incident optical power. Tomographic imaging is demonstrated in vitro in the peripapillary area of the retina and in the coronary artery, two clinically relevant examples that are representative of transparent and turbid media, respectively.

Tomographic imaging techniques such as x-ray computed tomography (1), magnetic resonance imaging (2), and ultrasound imaging (3) have found widespread applications in medicine. Each of these techniques measures a different physical property and has a resolution and penetration range that prove advantageous for specific applications. In this report, we discuss OCT. With this technique it is possible to perform noninvasive cross-sectional imaging of internal structures in biological tissues by measuring their optical reflections.

Both low-coherence light and ultrashort laser pulses can be used to measure internal structure in biological systems. An optical signal that is transmitted through or reflected from a biological tissue will contain time-of-flight information, which in turn yields spatial information about tissue microstructure. Time-resolved transmission spectroscopy has been used to measure absorption and scattering properties in tissues and has been demonstrated as a noninvasive diagnostic measure of hemoglobin oxygenation in the brain (4). Optical ranging measurements of microstructure have been performed in the eye and the skin with femtosecond laser pulses (5). Time gating by means of coherent (6) as well as noncoherent (7) techniques has been used to preferentially detect directly transmitted light and obtain transmission images in turbid tissue. Low-coherence reflectometry has been used for ranging measurements in optical components (8), for surface contour mapping in integrated circuits (9), and for ranging measurements in the retina (10) and other eye structures (10, 11).

In contrast to time domain techniques, low-coherence reflectometry can be performed with continuous-wave light without the need for ultrashort pulse laser sources. Furthermore, recent technological advances in low-coherence reflectometry have allowed the construction of compact and modular systems that use diode light sources and fiber optics and have achieved micrometer spatial resolutions and high detection sensitivities (12).

We have extended the technique of low-coherence reflectometry to tomographic imaging in biological systems. In low-coherence reflectometry, the coherence property of light reflected from a sample provides information on the time-of-flight delay from the reflective boundaries and backscattering sites in the sample. The delay information is then used to determine the longitudinal location of the reflection sites. The OCT system performs multiple longitudinal scans at a series of lateral locations to provide a two-dimensional map of reflection sites in the sample. This mode of operation is analogous to ultrasonic pulse-echo imaging (ultrasound B-mode).

The optical sectioning capability of OCT is akin to that of confocal microscopic systems (13, 14). However, although the longitudinal resolution of confocal microscopy depends on the available numerical aperture (15), OCT's resolution is limited only by the coherence length of the light source. Thus, OCT can maintain high depth resolution even when the available aperture is small. This feature will be particularly useful for in vivo measurement of deep tissues, for example, in transpupillary imaging of the posterior eye and in endoscopic imaging.

The OCT scanner (Fig. 1) is an extension of previous low-coherence reflectometer systems (12). High-speed, continuous-motion longitudinal scanning is used to increase the data acquisition rate, and a transverse scanning mechanism makes possible two-dimensional imaging. The heart of the system is the fiber optic Michelson interferometer, which is illuminated by low-coherence light (830 nm wavelength) from a superluminescent diode (SLD). The tissue sample is placed in one interferometer arm, and sample reflections are combined with the reflection from the reference mirror. The amplitudes and delays of tissue reflections are measured by scanning the reference mirror

Fig. 1. Schematic of the OCT scanner. The SLD output is coupled into a single mode fiber and split at the 50/50 coupler into sample and reference arms. Reflections from

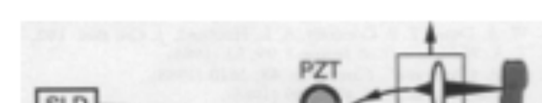

Optical Coherence Tomography

Optical Coherence Tomography

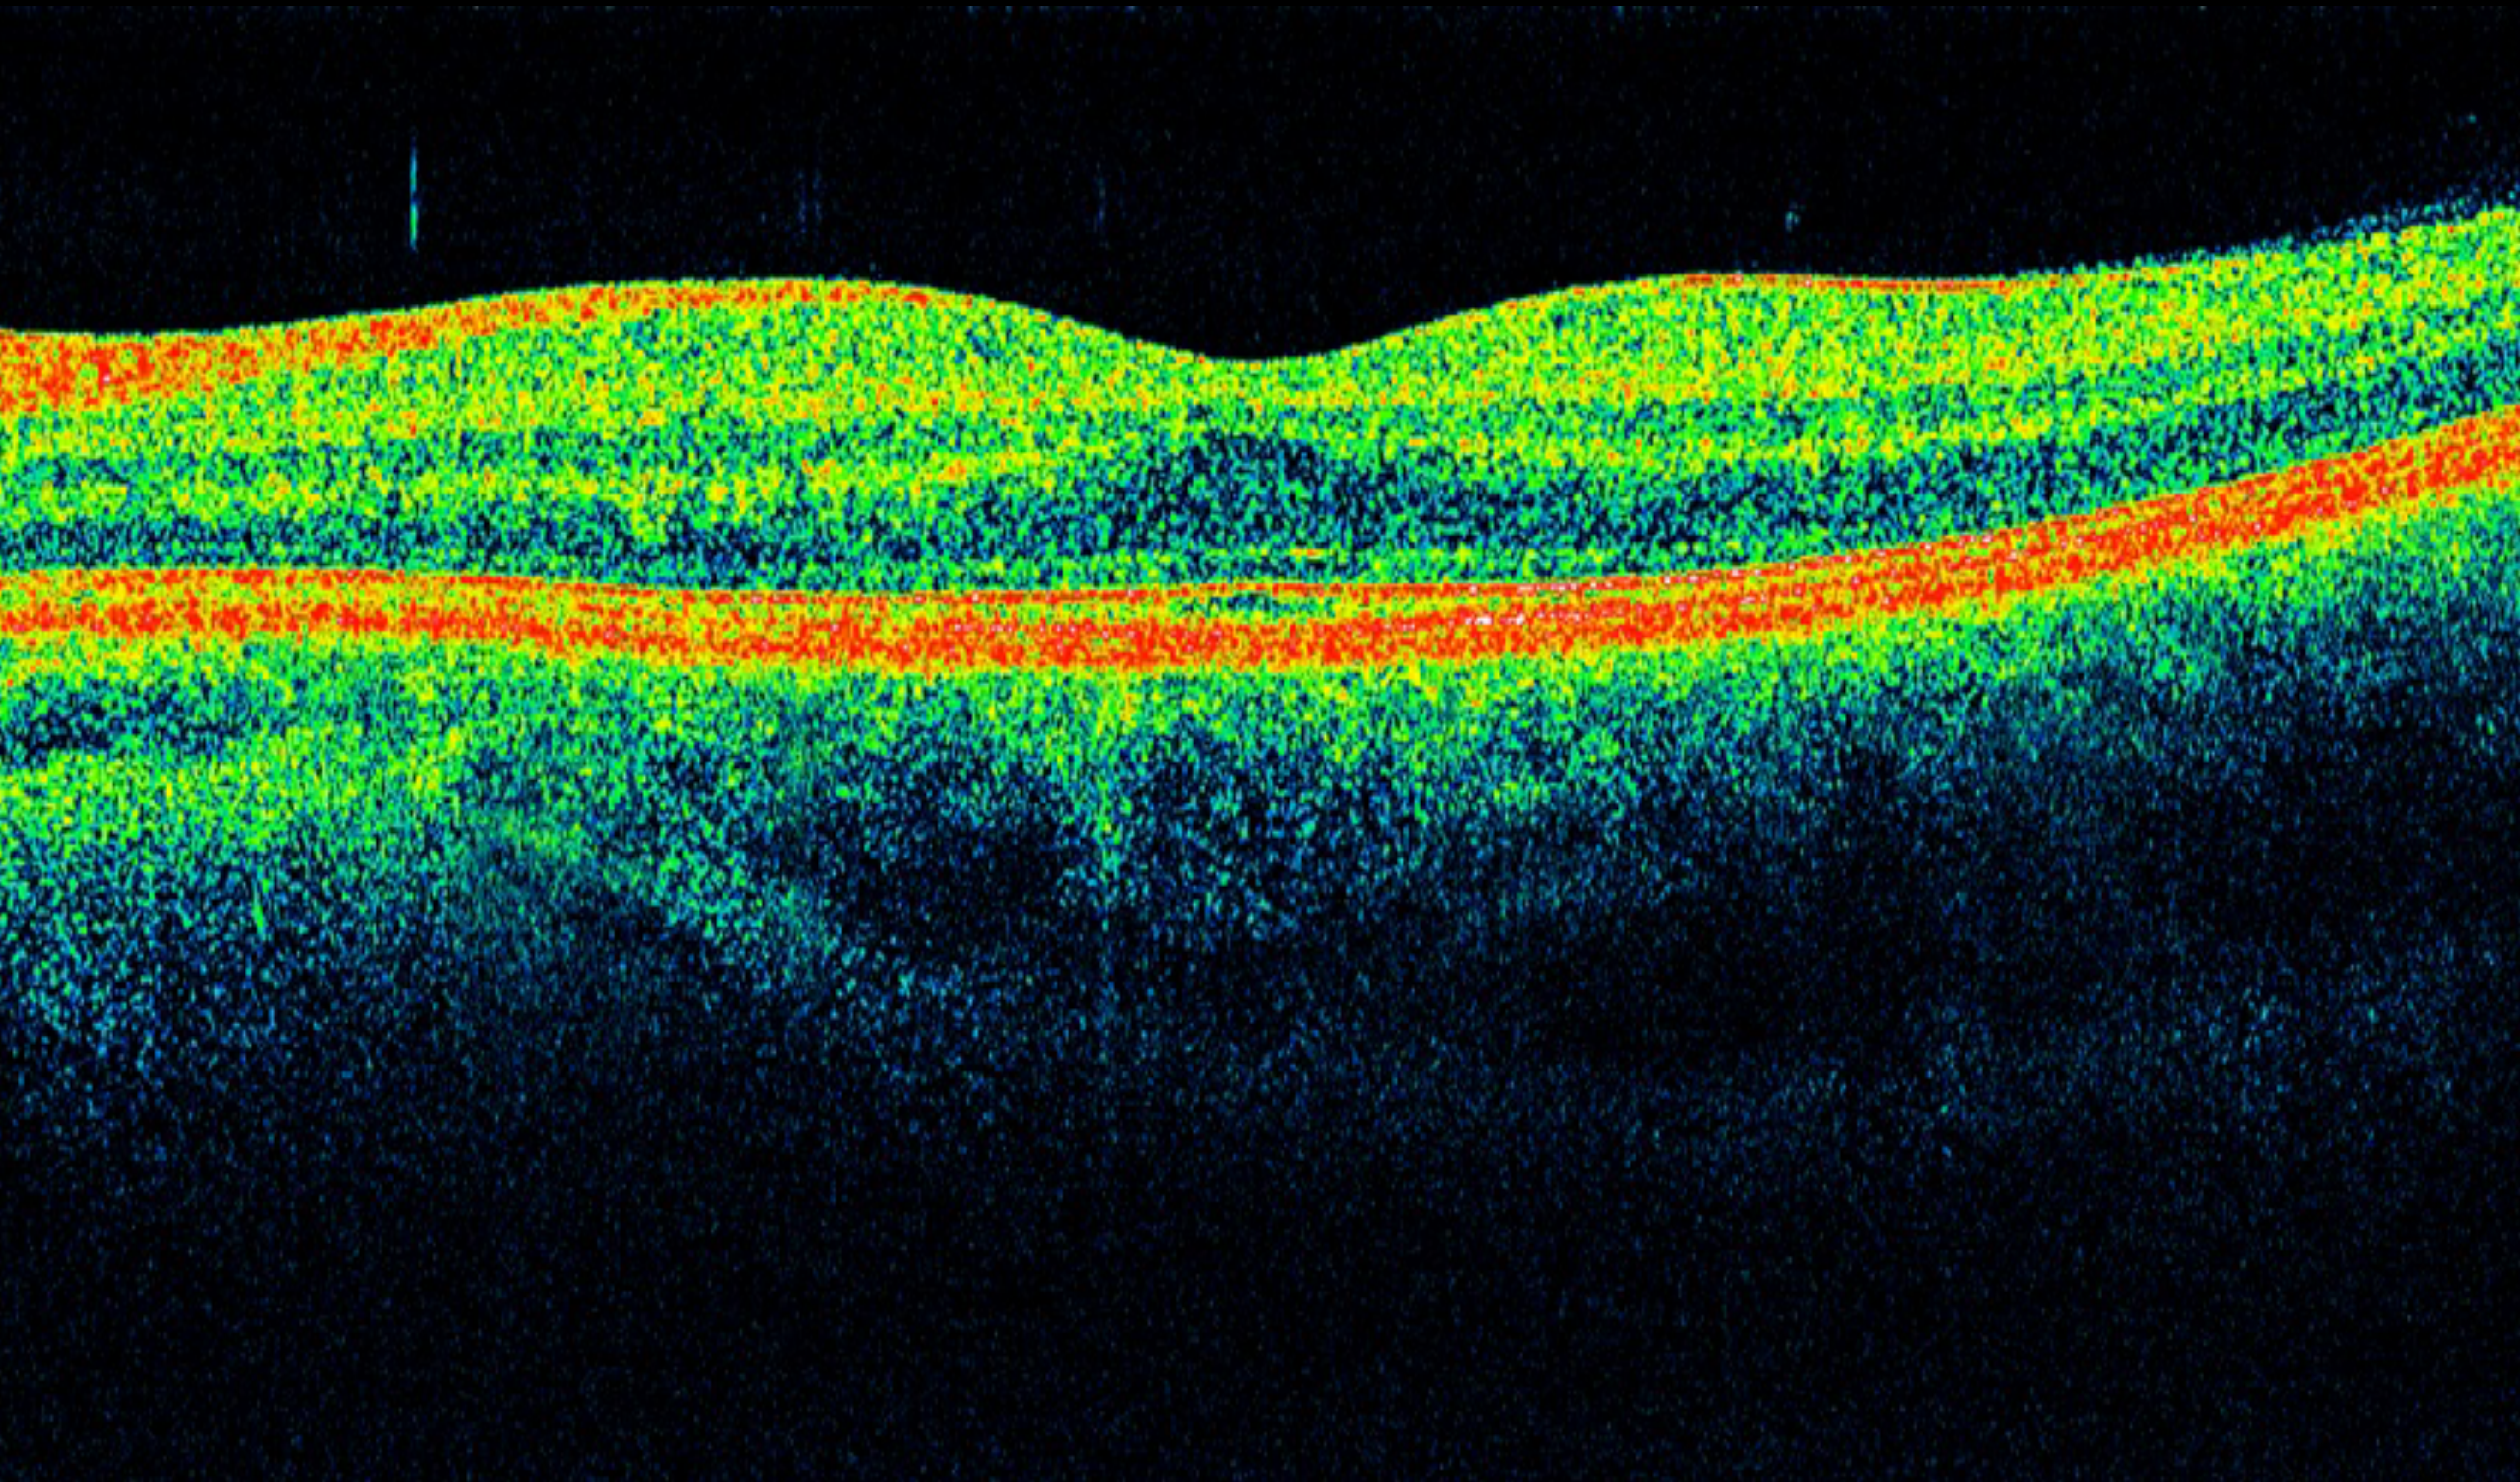

The Problem…


MOORFIELDS
EYE HOSPITAL
ROYAL LONDON OPHTHALMIC HOSPITAL

The Problem…

2017

The Problem...

2017 **~1000/day!**

The Problem…

The Problem…

The Problem...

7000

The Problem...

7000

800

The Problem…

OT
Optometry Today

Blogs Jobs Advertise Join AOP MyAOP log in Subscribe | AOP website

CET | Science and vision | Industry | Professional support | In practice

OT ▸ Industry ▸ High Street ▸ OCT rollout in every Specsavers announced

OCT ROLLOUT IN EVERY SPECSAVERS ANNOUNCED

The multiple will ensure all 740 of its UK practices have an OCT device installed within the next two years

22 May 2017 by Emily McCormick

Category: Multiple, OCT

Specsavers has announced a multi-million pound plan to ensure that each of its 740 practices in the UK has an optical coherence tomography (OCT) device installed within the next two years.

The nationwide rollout will begin in June, the multiple said, confirming that 35 of its practices already have the machine in store.

May 2017

July 2016

July 2016

Moorfields Eye Hospital NHS
NHS Foundation Trust

DeepMind

DeepMind

King's Cross, London

DeepMind

King's Cross, London

>800 of the best AI researchers

DeepMind

King’s Cross, London

>800 of the best AI researchers

200+ peer-reviewed publications

Breakthrough

Breakthrough

ATARI® 2600™

Breakthrough

2015

Landmark Work

Landmark Work

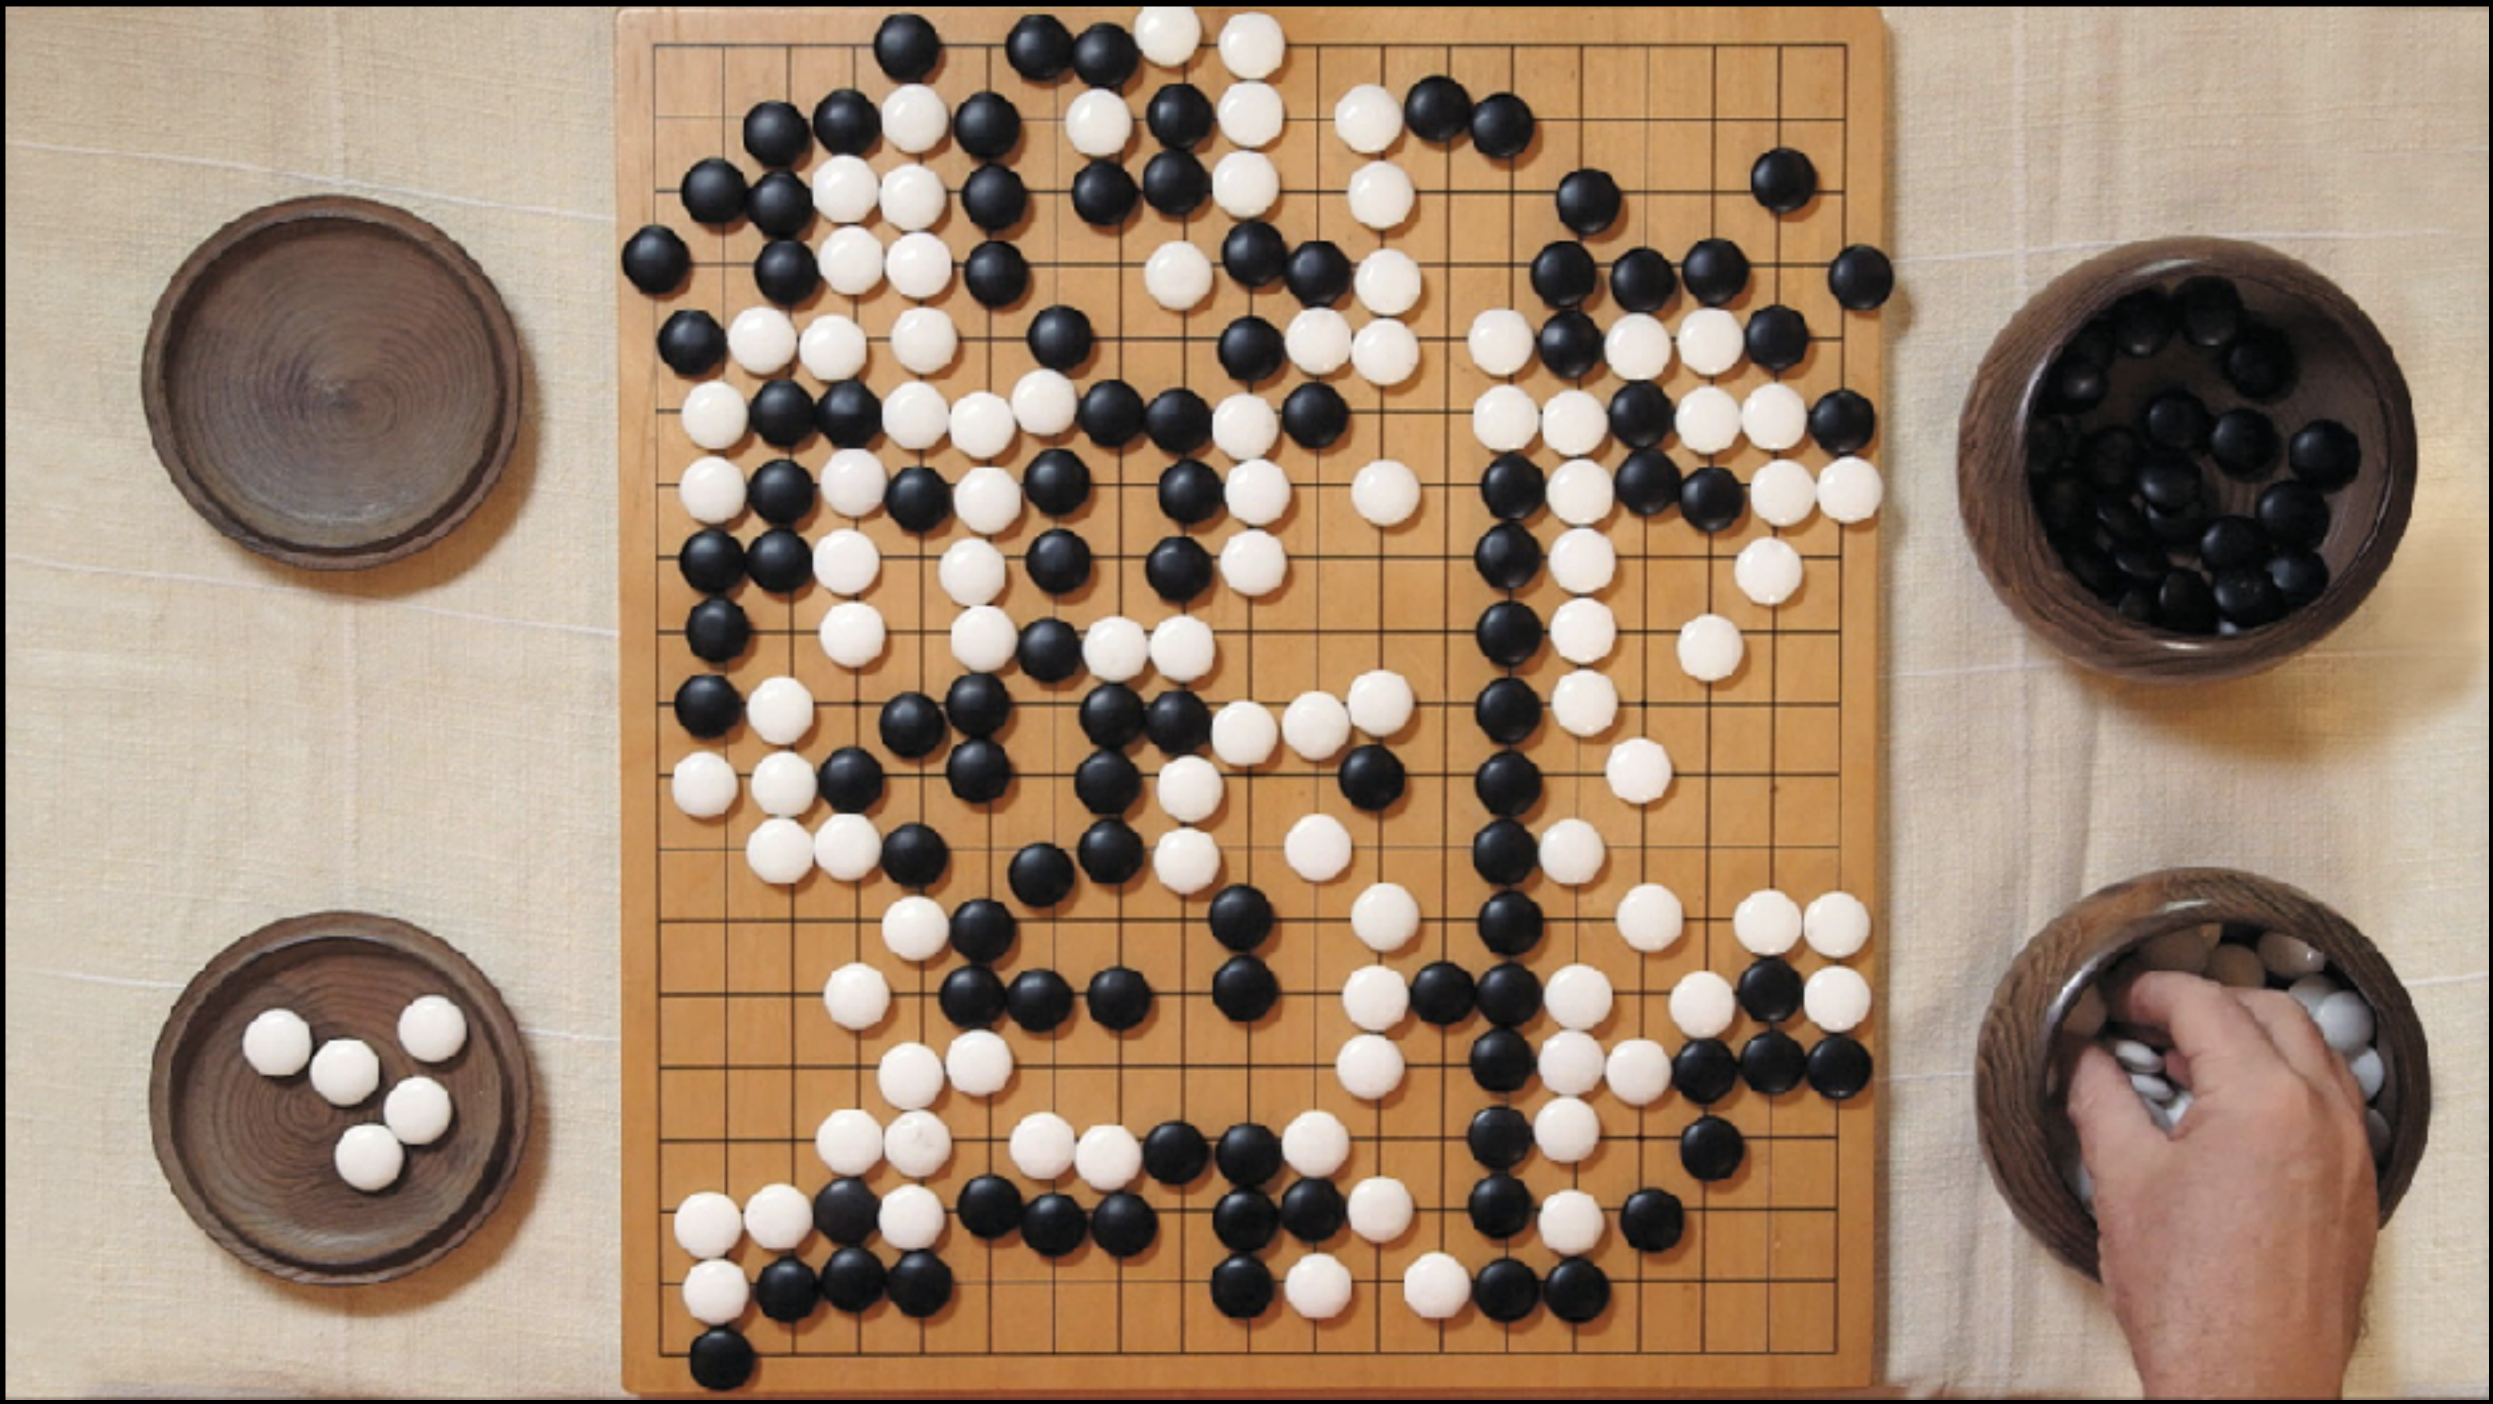

Landmark Work

Landmark Work

2016

March 2016

March 2016


KBA
Google DeepMind
Challenge Match
8 - 15 March 2016
AlphaGo
AlphaGo
Lee Sedol
Google DeepMind

March 2016

AlphaGo wins… 4-1 !!!

March 2016

December 2018

December 2018

December 2018

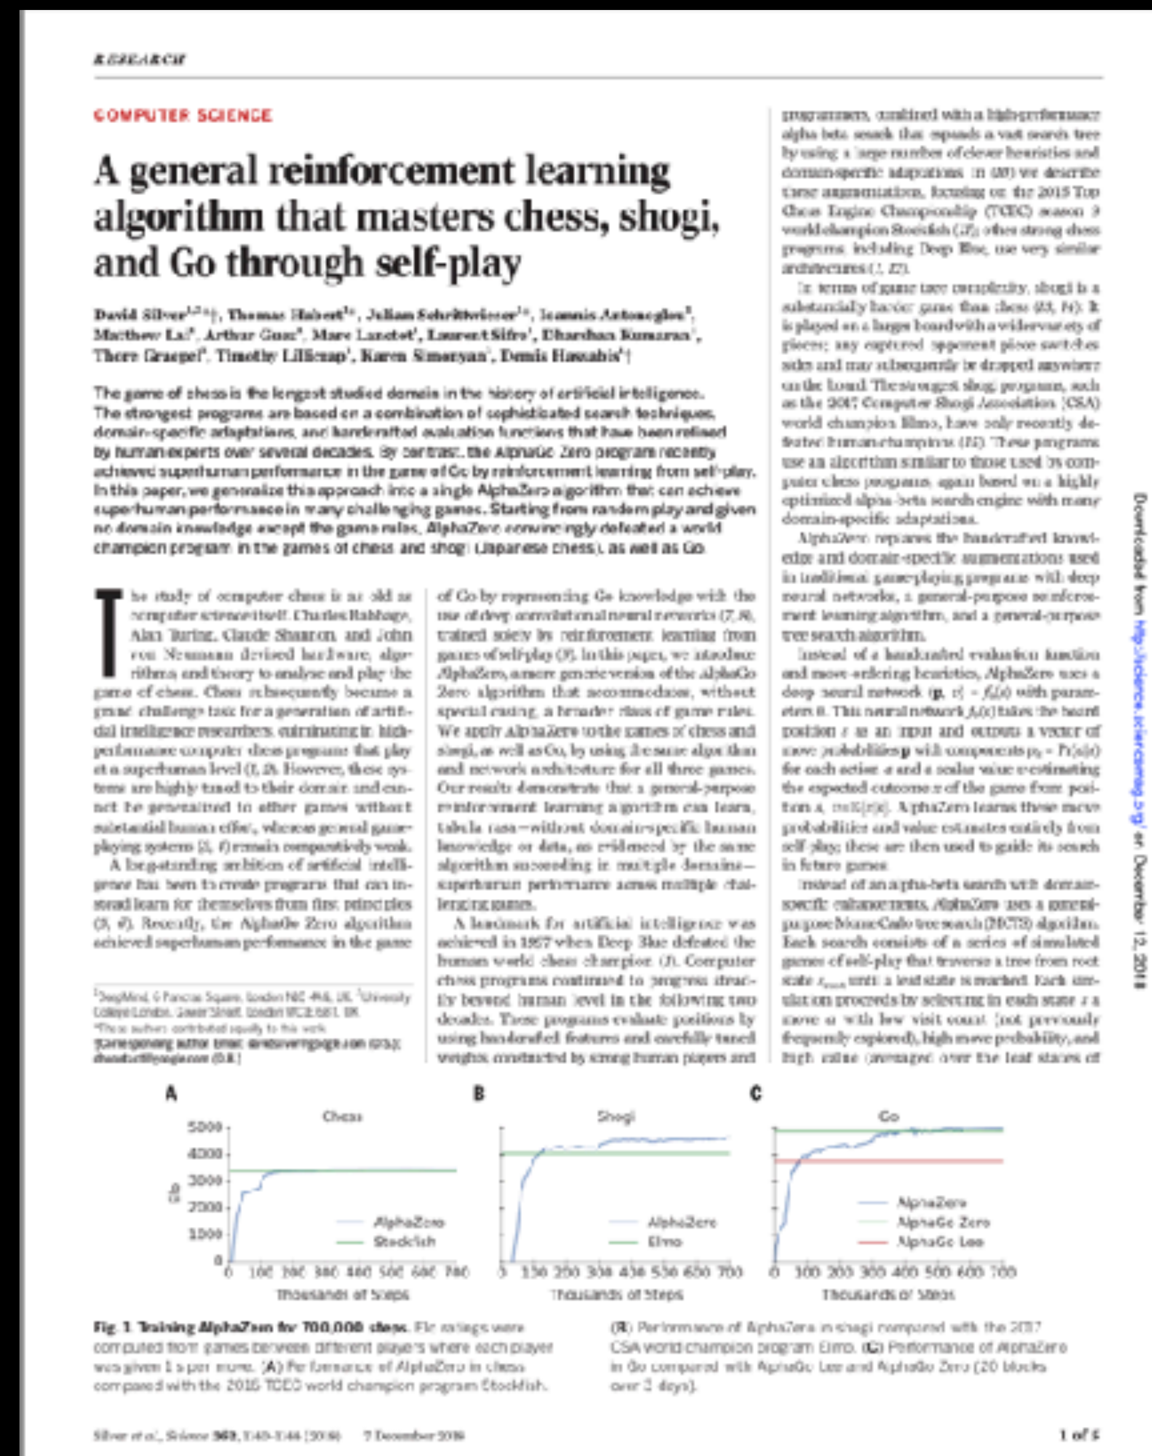
RESEARCH

COMPUTER SCIENCE

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play

David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, Demis Hassabis

The game of chess is the longest-studied domain in the history of artificial intelligence. The strongest programs are based on a combination of sophisticated search techniques, domain-specific adaptations, and handcrafted evaluation functions that have been refined by human experts over several decades. By contrast, the AlphaGo Zero program recently achieved superhuman performance in the game of Go by reinforcement learning from self-play. In this paper, we generalize this approach into a single AlphaZero algorithm that can achieve superhuman performance in many challenging games. Starting from random play and given no domain knowledge except the game rules, AlphaZero convincingly defeated a world champion program in the games of chess and shogi (Japanese chess), as well as Go.

Fig. 1. Training AlphaZero for 700,000 steps.

Silver et al., Science 362, 1140–1144 (2018) 7 December 2018

1 of 5

A New Artificial Intelligence “Spring”?

A New Artificial Intelligence “Spring”?

A New Artificial Intelligence "Spring"?

deep genomics

zebra™ MEDICAL.VISION

enlitic

Artificial Neural Networks

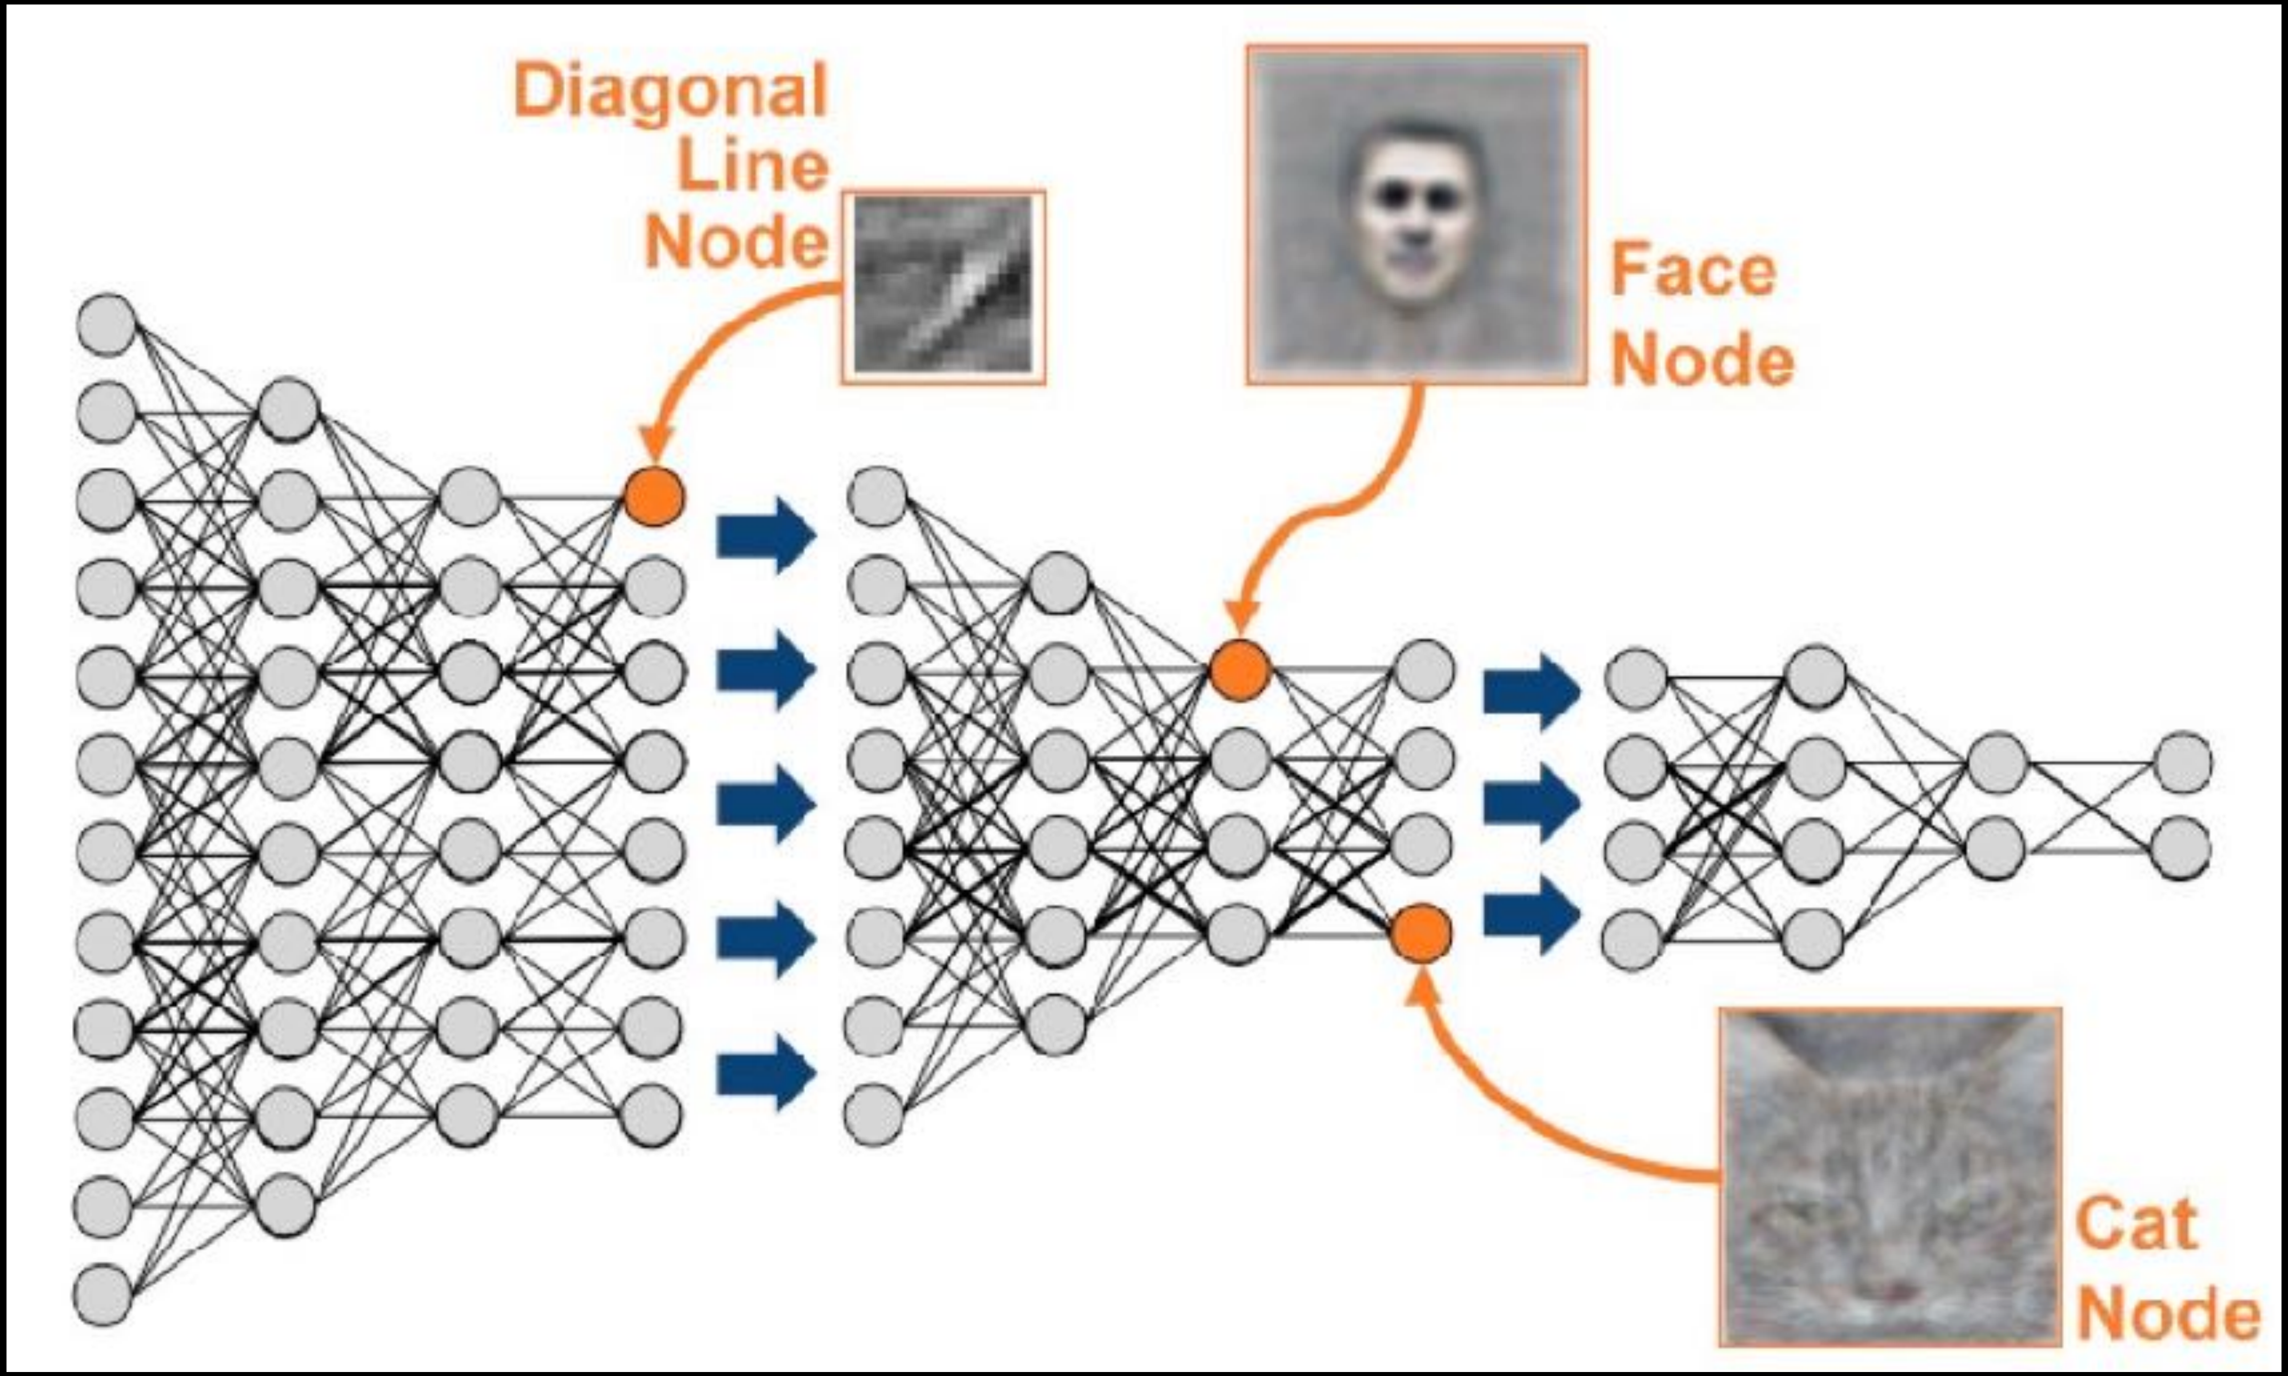

Learns from Experience

Learns from Experience

Not Pre-designed or Pre-Specified

“Deep Learning”

“Deep Learning”

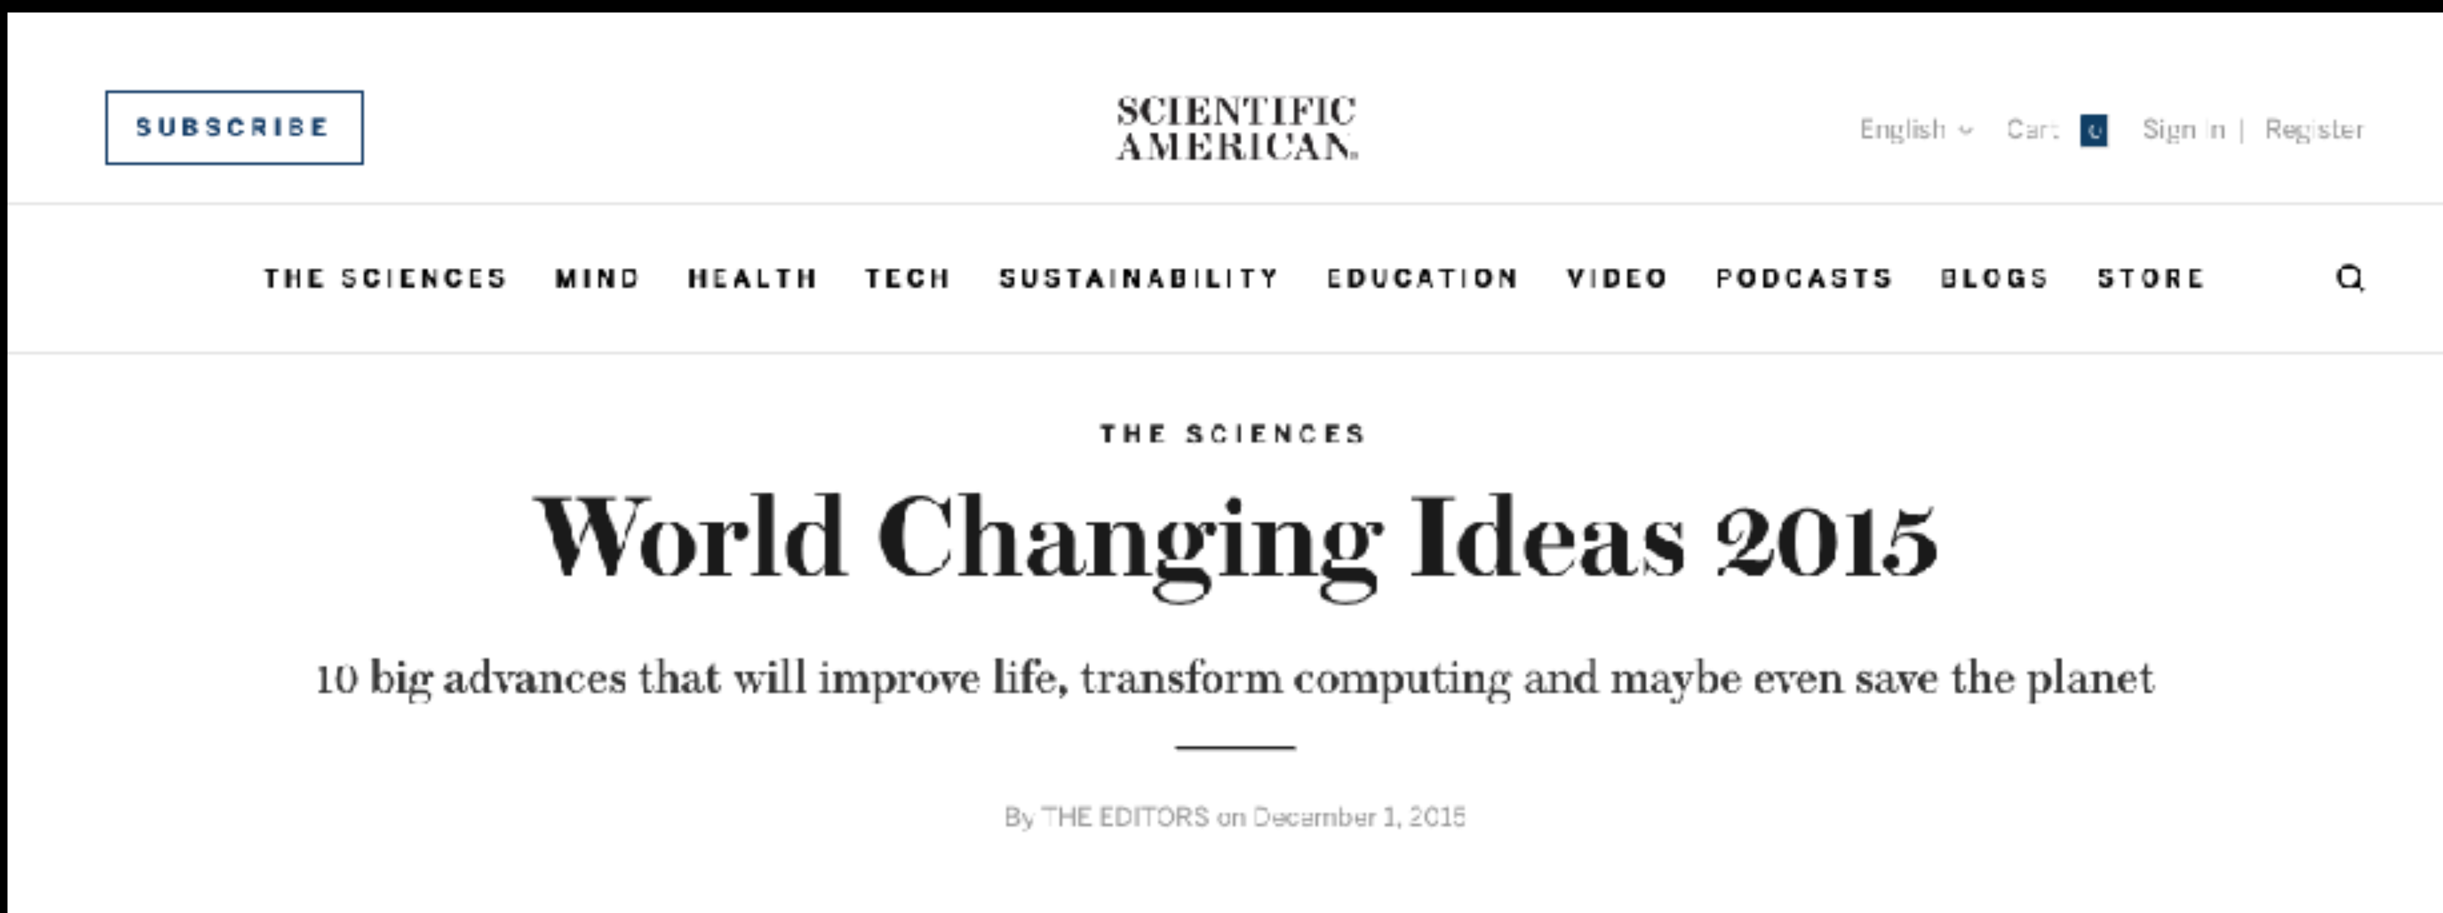
SUBSCRIBE

SCIENTIFIC AMERICAN

English Cart 0 Sign in | Register

THE SCIENCES MIND HEALTH TECH SUSTAINABILITY EDUCATION VIDEO PODCASTS BLOGS STORE

THE SCIENCES

World Changing Ideas 2015

10 big advances that will improve life, transform computing and maybe even save the planet

By THE EDITORS on December 1, 2015

Digging a Little Bit Deeper…

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

Digging a Little Bit Deeper…

Classical Statistics	Artificial Intelligence
Low dimensional data	High dimensional data (e.g., more than 1000 dimensions)
Lots of noise in the data	Noise is not sufficient to obscure the structure in the data if processed right
Not much structure in the data and what structure there is can be represented by a fairly simple model	A huge amount of structure in the data, but the structure is too complicated to be represented by a single model (e.g., the mapping of an OCT volume scan to a specific disease diagnosis)
Main problem is distinguishing true structure from noise	Main problem is figuring out how to represent the complicated structure in a way that allows it to be learned

Source: adapted from lecture by Professor Geoff Hinton, FRS, to the Royal Society, 2016
YouTube link: https://www.youtube.com/watch?v=VhmE_UXDOGs

High Dimensional Data

High Dimensional Data

High Dimensional Data

~65 million datapoints!!!

Example from Google Brain

Example from Google Brain

nature
biomedical engineering

ARTICLES

https://doi.org/10.1038/s41551-018-0195-0

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin[1,4], Avinash V. Varadarajan[1,4], Katy Blumer[1], Yun Liu[1], Michael V. McConnell[2,3], Greg S. Corrado[1], Lily Peng[1,4]* and Dale R. Webster[1,4]

Traditionally, medical discoveries are made by observing associations, making hypotheses from them and then designing and running experiments to test the hypotheses. However, with medical images, observing and quantifying associations can often be difficult because of the wide variety of features, patterns, colours, values and shapes that are present in real data. Here, we show that deep learning can extract new knowledge from retinal fundus images. Using deep-learning models trained on data from 284,335 patients and validated on two independent datasets of 12,026 and 999 patients, we predicted cardiovascular risk factors not previously thought to be present or quantifiable in retinal images, such as age (mean absolute error within 3.26 years), gender (area under the receiver operating characteristic curve (AUC) = 0.97), smoking status (AUC = 0.71), systolic blood pressure (mean absolute error within 11.23 mmHg) and major adverse cardiac events (AUC = 0.70). We also show that the trained deep-learning models used anatomical features, such as the optic disc or blood vessels, to generate each prediction.

February 2018

Example from Google Brain

nature biomedical engineering

ARTICLES

https://doi.org/10.1038/s41551-018-0195-0

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin[1,4], Avinash V. Varadarajan[1,4], Katy Blumer[1], Yun Liu[1], Michael V. McConnell[2,3], Greg S. Corrado[1], Lily Peng[1,4]* and Dale R. Webster[1,4]

Traditionally, medical discoveries are made by observing associations, making hypotheses from them and then designing and running experiments to test the hypotheses. However, with medical images, observing and quantifying associations can often be difficult because of the wide variety of features, patterns, colours, values and shapes that are present in real data. Here, we show that deep learning can extract new knowledge from retinal fundus images. Using deep-learning models trained on data from 284,335 patients and validated on two independent datasets of 12,026 and 999 patients, we predicted cardiovascular risk factors not previously thought to be present or quantifiable in retinal images, such as age (mean absolute error within 3.26 years), gender (area under the receiver operating characteristic curve (AUC) = 0.97), smoking status (AUC = 0.71), systolic blood pressure (mean absolute error within 11.23 mmHg) and major adverse cardiac events (AUC = 0.70). We also show that the trained deep-learning models used anatomical features, such as the optic disc or blood vessels, to generate each prediction.

February 2018

Example from Google Brain

nature biomedical engineering

ARTICLES

https://doi.org/10.1038/s41551-018-0195-0

Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning

Ryan Poplin[1,4], Avinash V. Varadarajan[1,4], Katy Blumer[1], Yun Liu[1], Michael V. McConnell[2,3], Greg S. Corrado[1], Lily Peng[1,4]* and Dale R. Webster[1,4]

Traditionally, medical discoveries are made by observing associations, making hypotheses from them and then designing and running experiments to test the hypotheses. However, with medical images, observing and quantifying associations can often be difficult because of the wide variety of features, patterns, colours, values and shapes that are present in real data. Here, we show that deep learning can extract new knowledge from retinal fundus images. Using deep-learning models trained on data from 284,335 patients and validated on two independent datasets of 12,026 and 999 patients, we predicted cardiovascular risk factors not previously thought to be present or quantifiable in retinal images, such as age (mean absolute error within 3.26 years), gender (area under the receiver operating characteristic curve (AUC) = 0.97), smoking status (AUC = 0.71), systolic blood pressure (mean absolute error within 11.23 mmHg) and major adverse cardiac events (AUC = 0.70). We also show that the trained deep-learning models used anatomical features, such as the optic disc or blood vessels, to generate each prediction.

February 2018

Clinical Applications?

Clinical Applications?

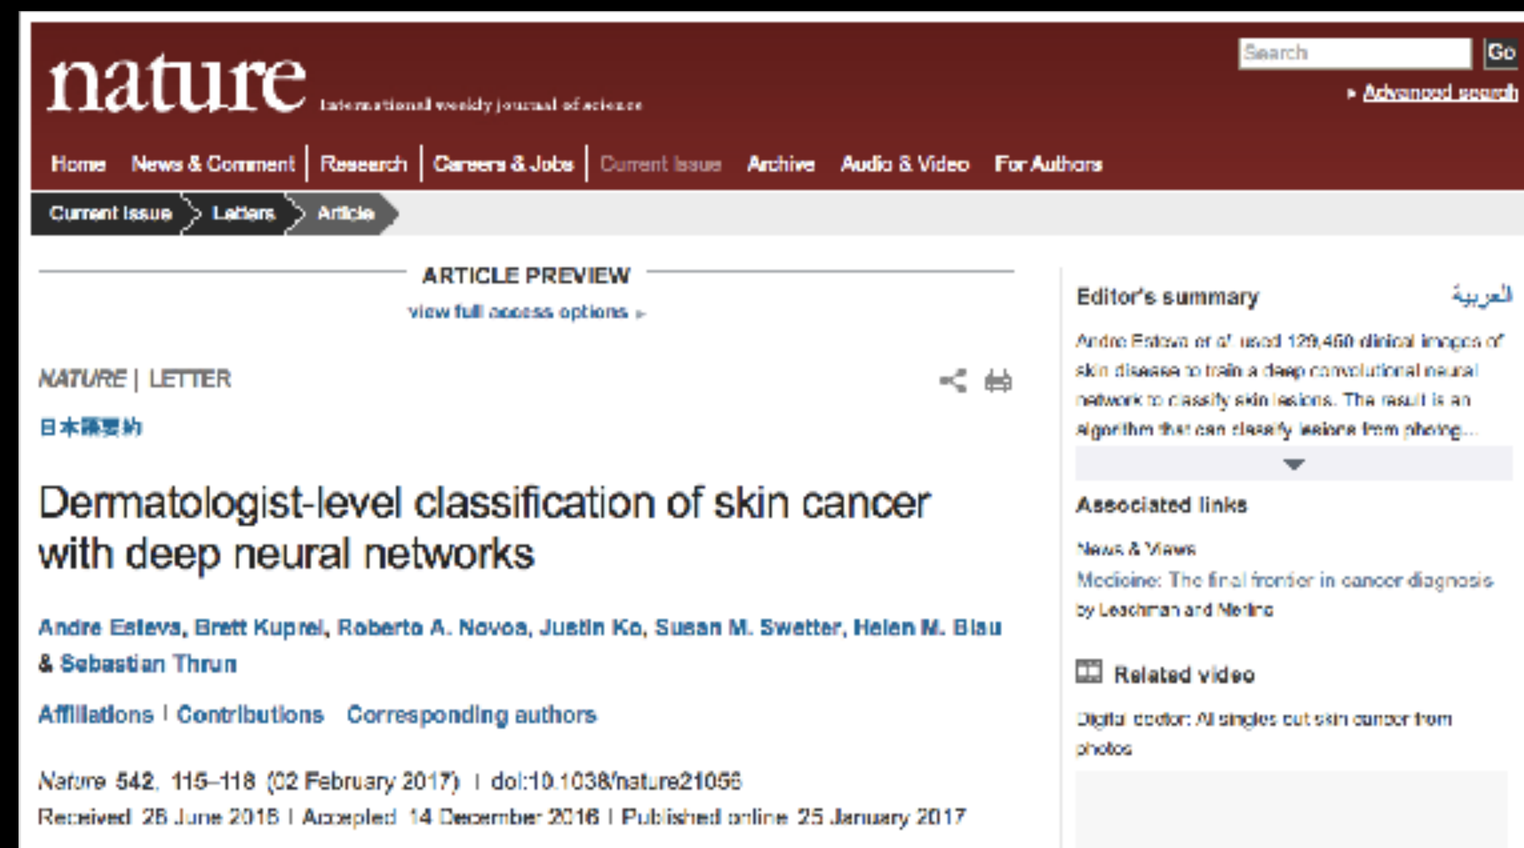

nature international weekly journal of science

Search Go

Advanced search

Home | News & Comment | Research | Careers & Jobs | Current Issue | Archive | Audio & Video | For Authors

Current issue > Letters > Article

ARTICLE PREVIEW

view full access options

NATURE | LETTER

日本語要約

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau & Sebastian Thrun

Affiliations | Contributions | Corresponding authors

Nature 542, 115–118 (02 February 2017) | doi:10.1038/nature21056

Received 28 June 2016 | Accepted 14 December 2016 | Published online 25 January 2017

Editor's summary

العربية

Andre Esteva et al. used 129,450 clinical images of skin disease to train a deep convolutional neural network to classify skin lesions. The result is an algorithm that can classify lesions from photog...

Associated links

News & Views

Medicine: The final frontier in cancer diagnosis by Leachman and Merlino

Related video

Digital doctor: AI singles out skin cancer from photos

January 2017

Clinical and Scientific Applications?

Cornell University Library

We gratefully acknowledge support from the Simons Foundation and member institutions

arXiv.org > cs > arXiv:1711.05225

Search or Article ID inside arXiv | All papers | Broaden your search using Semantic Scholar

(Help | Advanced search)

Computer Science > Computer Vision and Pattern Recognition

CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning

Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, Andrew Y. Ng

(Submitted on 14 Nov 2017)

We develop an algorithm that can detect pneumonia from chest X-rays at a level exceeding practicing radiologists. Our algorithm, CheXNet, is a 121-layer convolutional neural network trained on ChestX-ray14, currently the largest publicly available chest X-ray dataset, containing over 100,000 frontal-view X-ray images with 14 diseases. Four practicing academic radiologists annotate a test set, on which we compare the performance of CheXNet to that of radiologists. We find that CheXNet exceeds average radiologist performance on pneumonia detection on both sensitivity and specificity. We extend CheXNet to detect all 14 diseases in ChestX-ray14 and achieve state of the art results on all 14 diseases.

Subjects: Computer Vision and Pattern Recognition (cs.CV); Learning (cs.LG); Machine Learning (stat.ML)
Cite as: arXiv:1711.05225 [cs.CV]
(or arXiv:1711.05225v1 [cs.CV] for this version)

Submission history

From: Pranav Rajpurkar [view email]
[v1] Tue, 14 Nov 2017 17:58:50 GMT (16273kb,D)

Download:
- PDF
- Other formats

(license)

Current browse context:
cs.CV
< prev | next >
new | recent | 1711

Change to browse by:
cs
cs.LG
stat
stat.ML

References & Citations
- NASA ADS

Bookmark (what is this?)

November 2017

Clinical and Scientific Applications?

CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning

Pranav Rajpurkar *[1] Jeremy Irvin *[1] Kaylie Zhu [1] Brandon Yang [1] Hershel Mehta [1]
Tony Duan [1] Daisy Ding [1] Aarti Bagul [1] Curtis Langlotz [2] Katie Shpanskaya [2]
Matthew P. Lungren [2] Andrew Y. Ng [1]

Abstract

We develop an algorithm that can detect pneumonia from chest X-rays at a level exceeding practicing radiologists. Our algorithm, CheXNet, is a 121-layer convolutional neural network trained on ChestX-ray14, currently the largest publicly available chest X-ray dataset, containing over 100,000 frontal-view X-ray images with 14 diseases. Four practicing academic radiologists annotate a test set, on which we compare the performance of CheXNet to that of radiologists. We find that CheXNet exceeds average radiologist performance on pneumonia detection on both sensitivity and specificity. We extend CheXNet to detect all 14 diseases in ChestX-ray14 and achieve state of the art results on all 14 diseases.

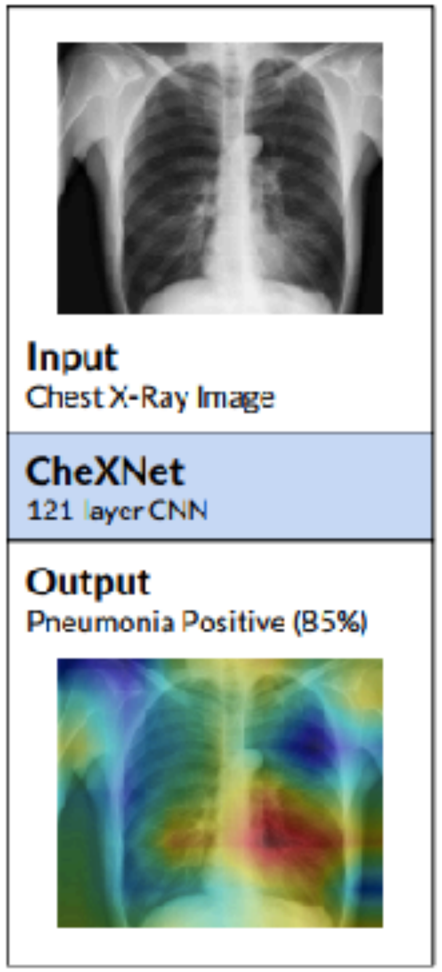

Figure 1. CheXNet is a 121-layer convolutional neural network that takes a chest X-ray image as input, and outputs

1. Introduction

More than 1 million adults are hospitalized with pneumonia and around 50,000 die from the disease every year in the US alone (CDC, 2017). Chest X-rays are currently the best available method for diagnosing pneumonia (WHO, 2001), playing a crucial role in clin-

Claim… and a Pushback!

Claim… and a Pushback!

"If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future."

Andrew Ng

Claim… and a Pushback!

"If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future."

Andrew Ng

"If a typical person can do a mental task with less than one second of thought, and we can gather an enormous amount of directly relevant data, we have a fighting chance — so long as the test data aren't too terribly different from the training data, and the domain doesn't change too much over time"

Gary Marcus

Applications in Ophthalmology?

Applications in Ophthalmology?

OCT

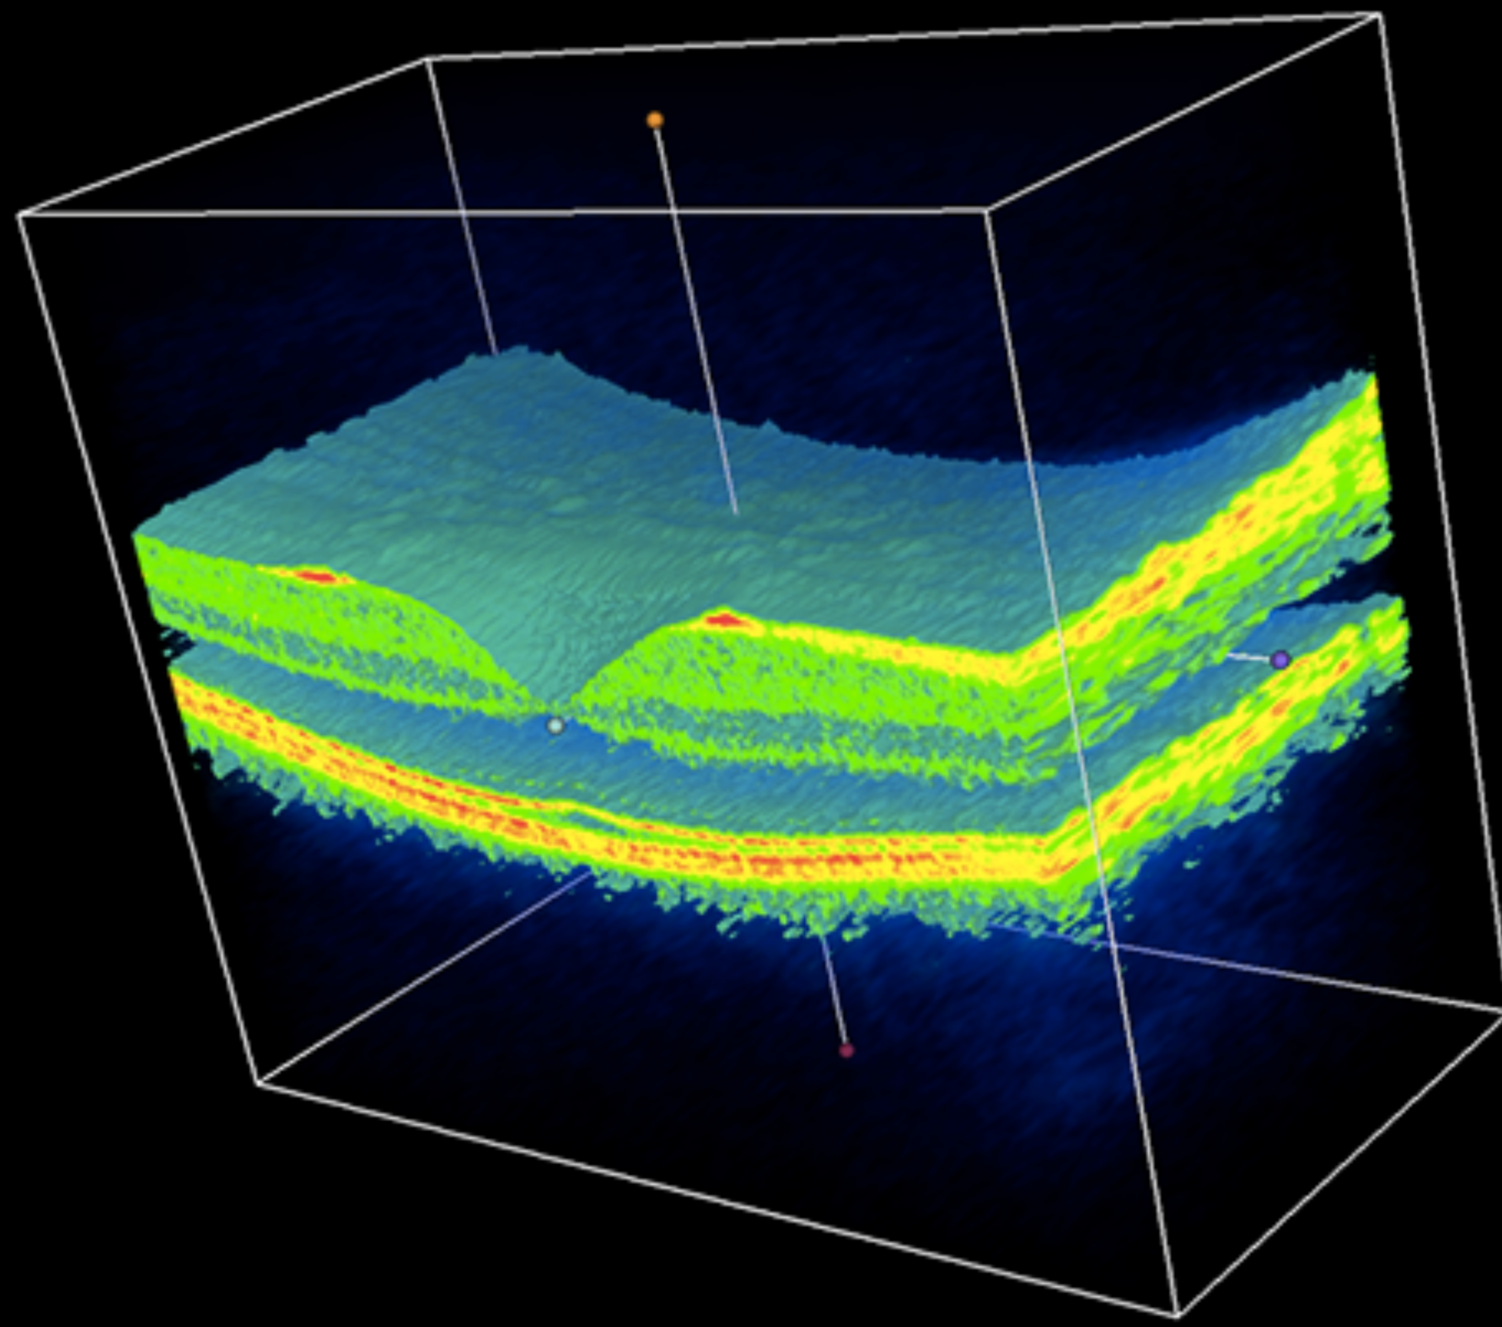

July 2015

July 2015

1. **Datasets Quantity and Quality**

July 2015

1. Datasets Quantity and Quality

2. Technical Feasibility

July 2015

1. Datasets Quantity and Quality
2. Technical Feasibility
3. Ethical / Governance Requirements

July 2015

1. Datasets Quantity and Quality
2. Technical Feasibility
3. Ethical / Governance Requirements

When all questions answered…

When all questions answered…

Research Collaboration Agreement

Website

Website

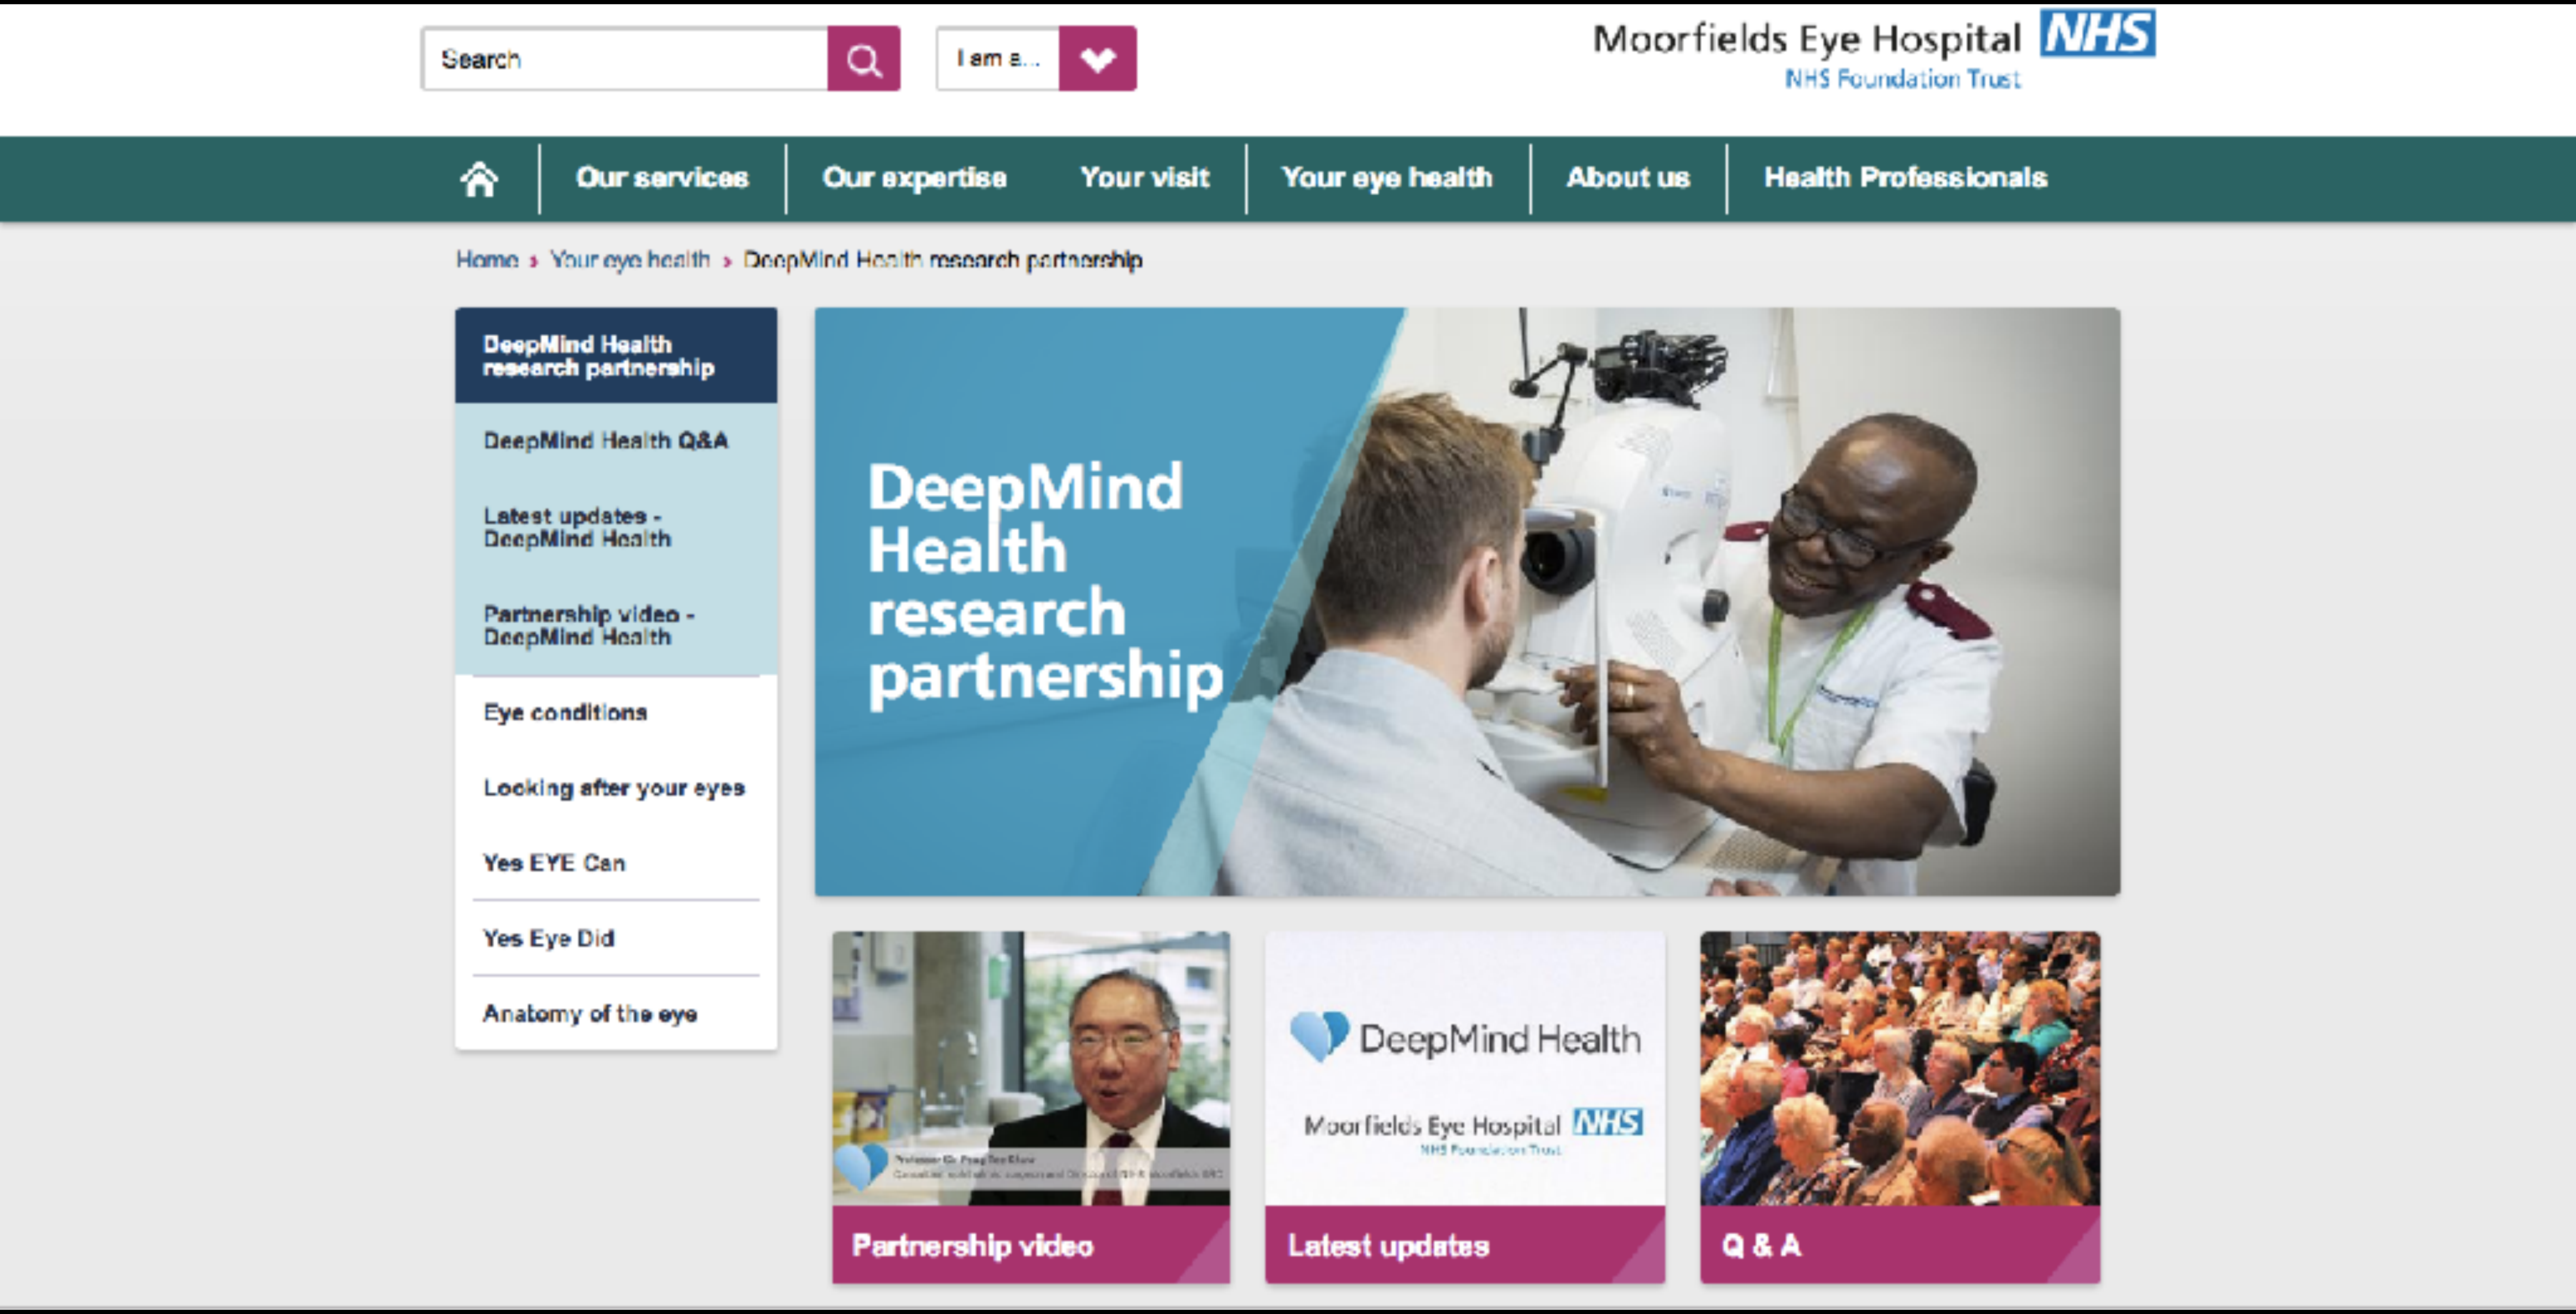
Search
I am a...
Moorfields Eye Hospital NHS
NHS Foundation Trust
Our services
Our expertise
Your visit
Your eye health
About us
Health Professionals
Home › Your eye health › DeepMind Health research partnership
DeepMind Health research partnership
DeepMind Health Q&A
Latest updates - DeepMind Health
Partnership video - DeepMind Health
Eye conditions
Looking after your eyes
Yes EYE Can
Yes Eye Did
Anatomy of the eye
DeepMind Health research partnership
DeepMind Health
Moorfields Eye Hospital NHS
Partnership video
Latest updates
Q & A

Website

Home > Your eye health > DeepMind Health research partnership > DeepMind Health Q&A

- DeepMind Health research partnership
- DeepMind Health Q&A
- Latest updates - DeepMind Health
- Partnership video - DeepMind Health
- Eye conditions
- Looking after your eyes
- Yes EYE Can
- Yes Eye Did
- Anatomy of the eye

DeepMind Health Q&A

- How did the partnership come about?
- What will the project involve?
- What is the project trying to achieve?
- How long will the project last?
- How much data has DeepMind been given access to?
- Do patients have to give their consent for their data to be used?
- What are the data protection measures in place for this project?
- Will any further patient information be shared between Moorfields and DeepMind in future?
- How can patients be sure that no identifiable data is being shared with DeepMind?
- What processes are in place to ensure the data transferred to DeepMind is only ever seen by the research team?
- What approvals has DeepMind been given for this research project?

Research Protocol

F1000Research

F1000Research 2016, null:null Last updated: 10 JUN 2016

CrossMark
click for updates

STUDY PROTOCOL

Automated analysis of retinal imaging using machine learning techniques for computer vision

Jeffrey De Fauw[1], Pearse Keane[1], 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

Patient-Centred Research

Patient-Centred Research

Patient Engagement Event, Sept 2016

Macular Society

Royal National Institute for the Blind

Fight For Sight UK

August 2018

August 2018

nature medicine

ARTICLES

https://doi.org/10.1038/s41591-018-0107-6

Clinically applicable deep learning for diagnosis and referral in retinal disease

Jeffrey De Fauw[1], Joseph R. Ledsam[1], Bernardino Romera-Paredes[1], Stanislav Nikolov[1], Nenad Tomasev[1], Sam Blackwell[1], Harry Askham[1], Xavier Glorot[1], Brendan O'Donoghue[1], Daniel Visentin[1], George van den Driessche[1], Balaji Lakshminarayanan[1], Clemens Meyer[1], Faith Mackinder[1], Simon Bouton[1], Kareem Ayoub[1], Reena Chopra[2], Dominic King[1], Alan Karthikesalingam[1], Cían O. Hughes[1,3], Rosalind Raine[3], Julian Hughes[2], Dawn A. Sim[2], Catherine Egan[2], Adnan Tufail[2], Hugh Montgomery[3], Demis Hassabis[1], Geraint Rees[3], Trevor Back[1], Peng T. Khaw[2], Mustafa Suleyman[1], Julien Cornebise[1,3,4], Pearse A. Keane[2,4*] and Olaf Ronneberger[1,4*]

The volume and complexity of diagnostic imaging is increasing at a pace faster than the availability of human expertise to interpret it. Artificial intelligence has shown great promise in classifying two-dimensional photographs of some common diseases and typically relies on databases of millions of annotated images. Until now, the challenge of reaching the performance of expert clinicians in a real-world clinical pathway with three-dimensional diagnostic scans has remained unsolved. Here, we apply a novel deep learning architecture to a clinically heterogeneous set of three-dimensional optical coherence tomography scans from patients referred to a major eye hospital. We demonstrate performance in making a referral recommendation that reaches or exceeds that of experts on a range of sight-threatening retinal diseases after training on only 14,884 scans. Moreover, we demonstrate that the tissue segmentations produced by our architecture act as a device-independent representation; referral accuracy is maintained when using tissue segmentations from a different type of device. Our work removes previous barriers to wider clinical use without prohibitive training data requirements across multiple pathologies in a real-world setting.

August 2018

nature medicine
SEPTEMBER 2018 VOL 24 NO 9
AI accelerates diagnosis
NAD+ biosynthesis and high-risk hospitalizations
Targeted microbiome therapy for thrombosis

August 2018

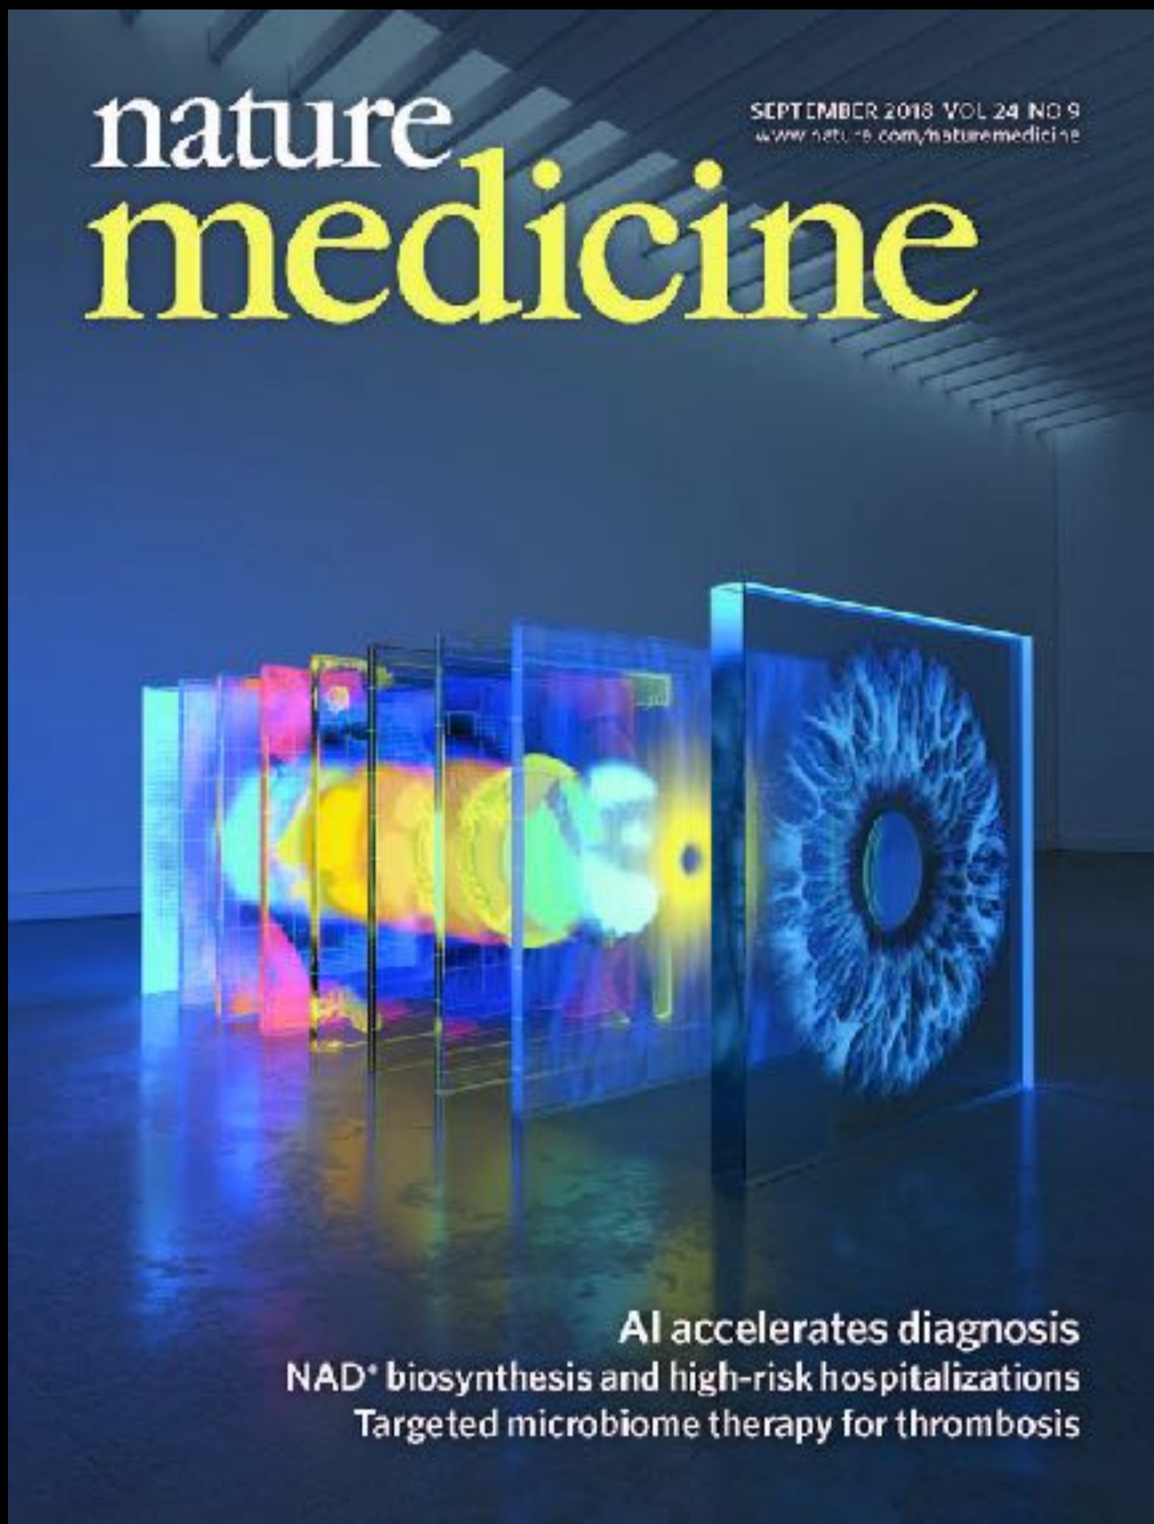


Evening Standard

DIER: WE'RE READY TO SILENCE OUR CRITICS

SPURS MIDFIELDER SAYS TEAM WILL SHINE DESPITE NO SIGNINGS » SPORT PAGE 52

FLEX APPEAL — GET STRONG WITH YOGA

COMPUTER DIAGNOSIS COULD SAVE SIGHT OF MILLIONS

» PIONEERING TECH TRIALLED AT MOORFIELDS RIVALS WORLD-LEADING EXPERTS FOR SPOTTING EYE DISEASE

KATIE PIPER REVEALED AS FIRST STRICTLY ST

» INSIDE TODAY

Boris should focus on Sadiq, say PM's allies

Pray for Aretha: singer is 'gravely ill with cancer'

No new cash for rough sleepers, minister admits

High Level Support!

High Level Support!

Novel Framework

Novel Framework

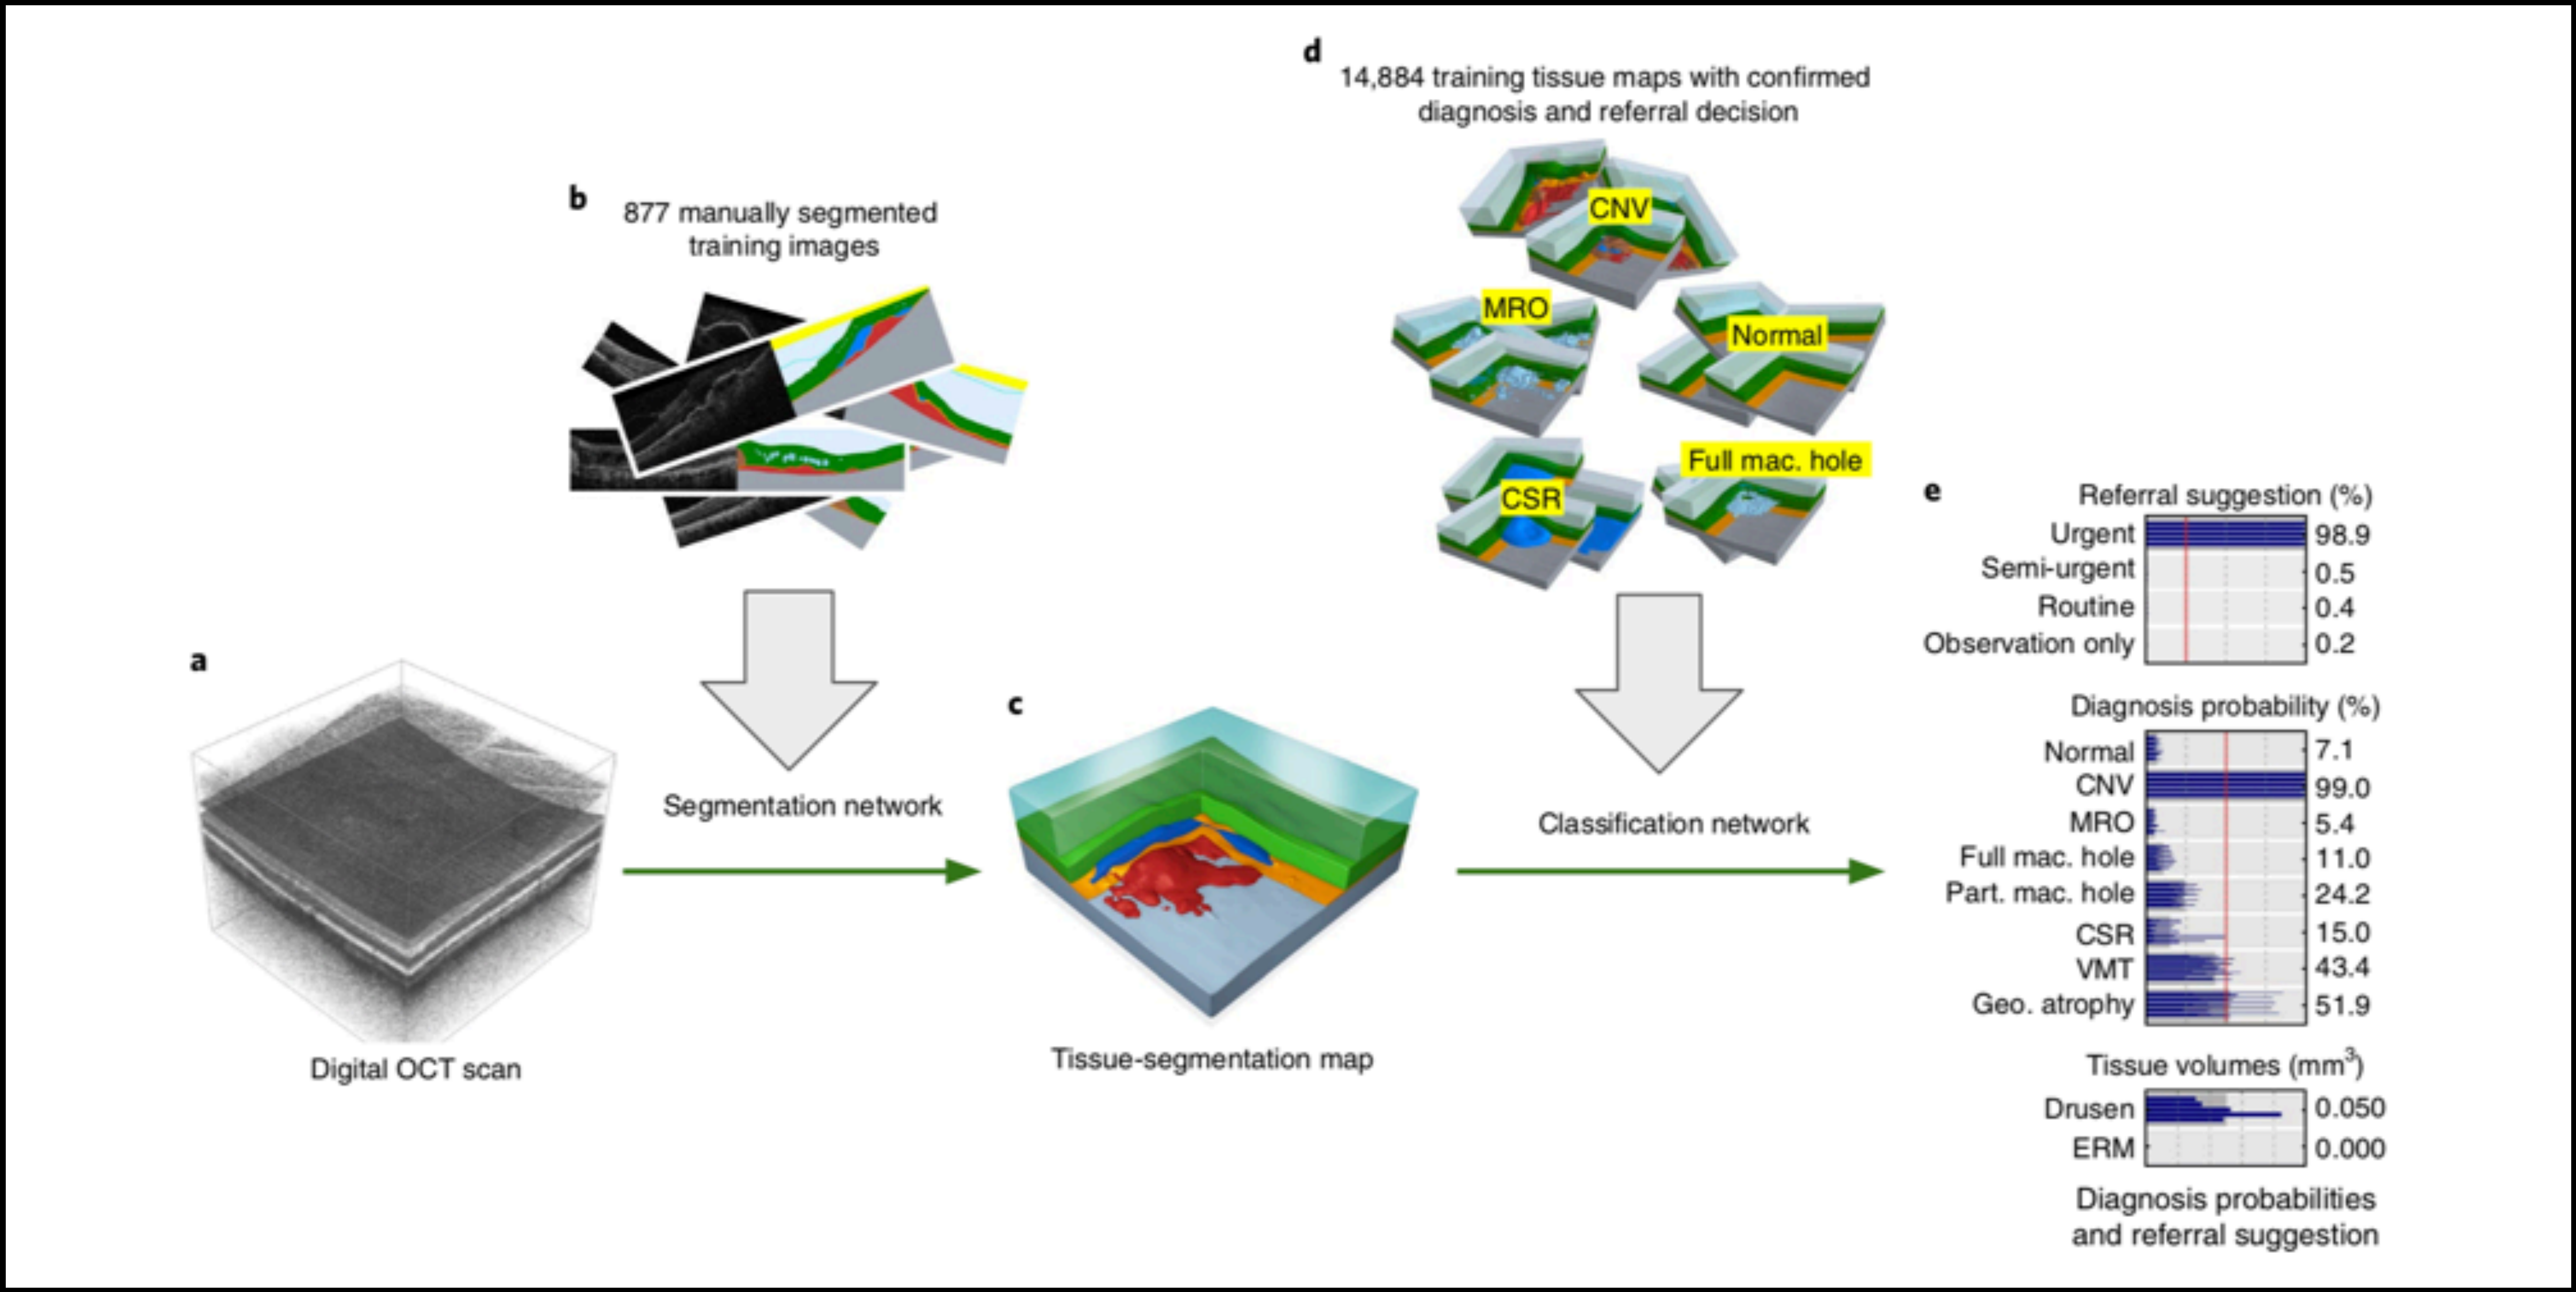

Novel Framework

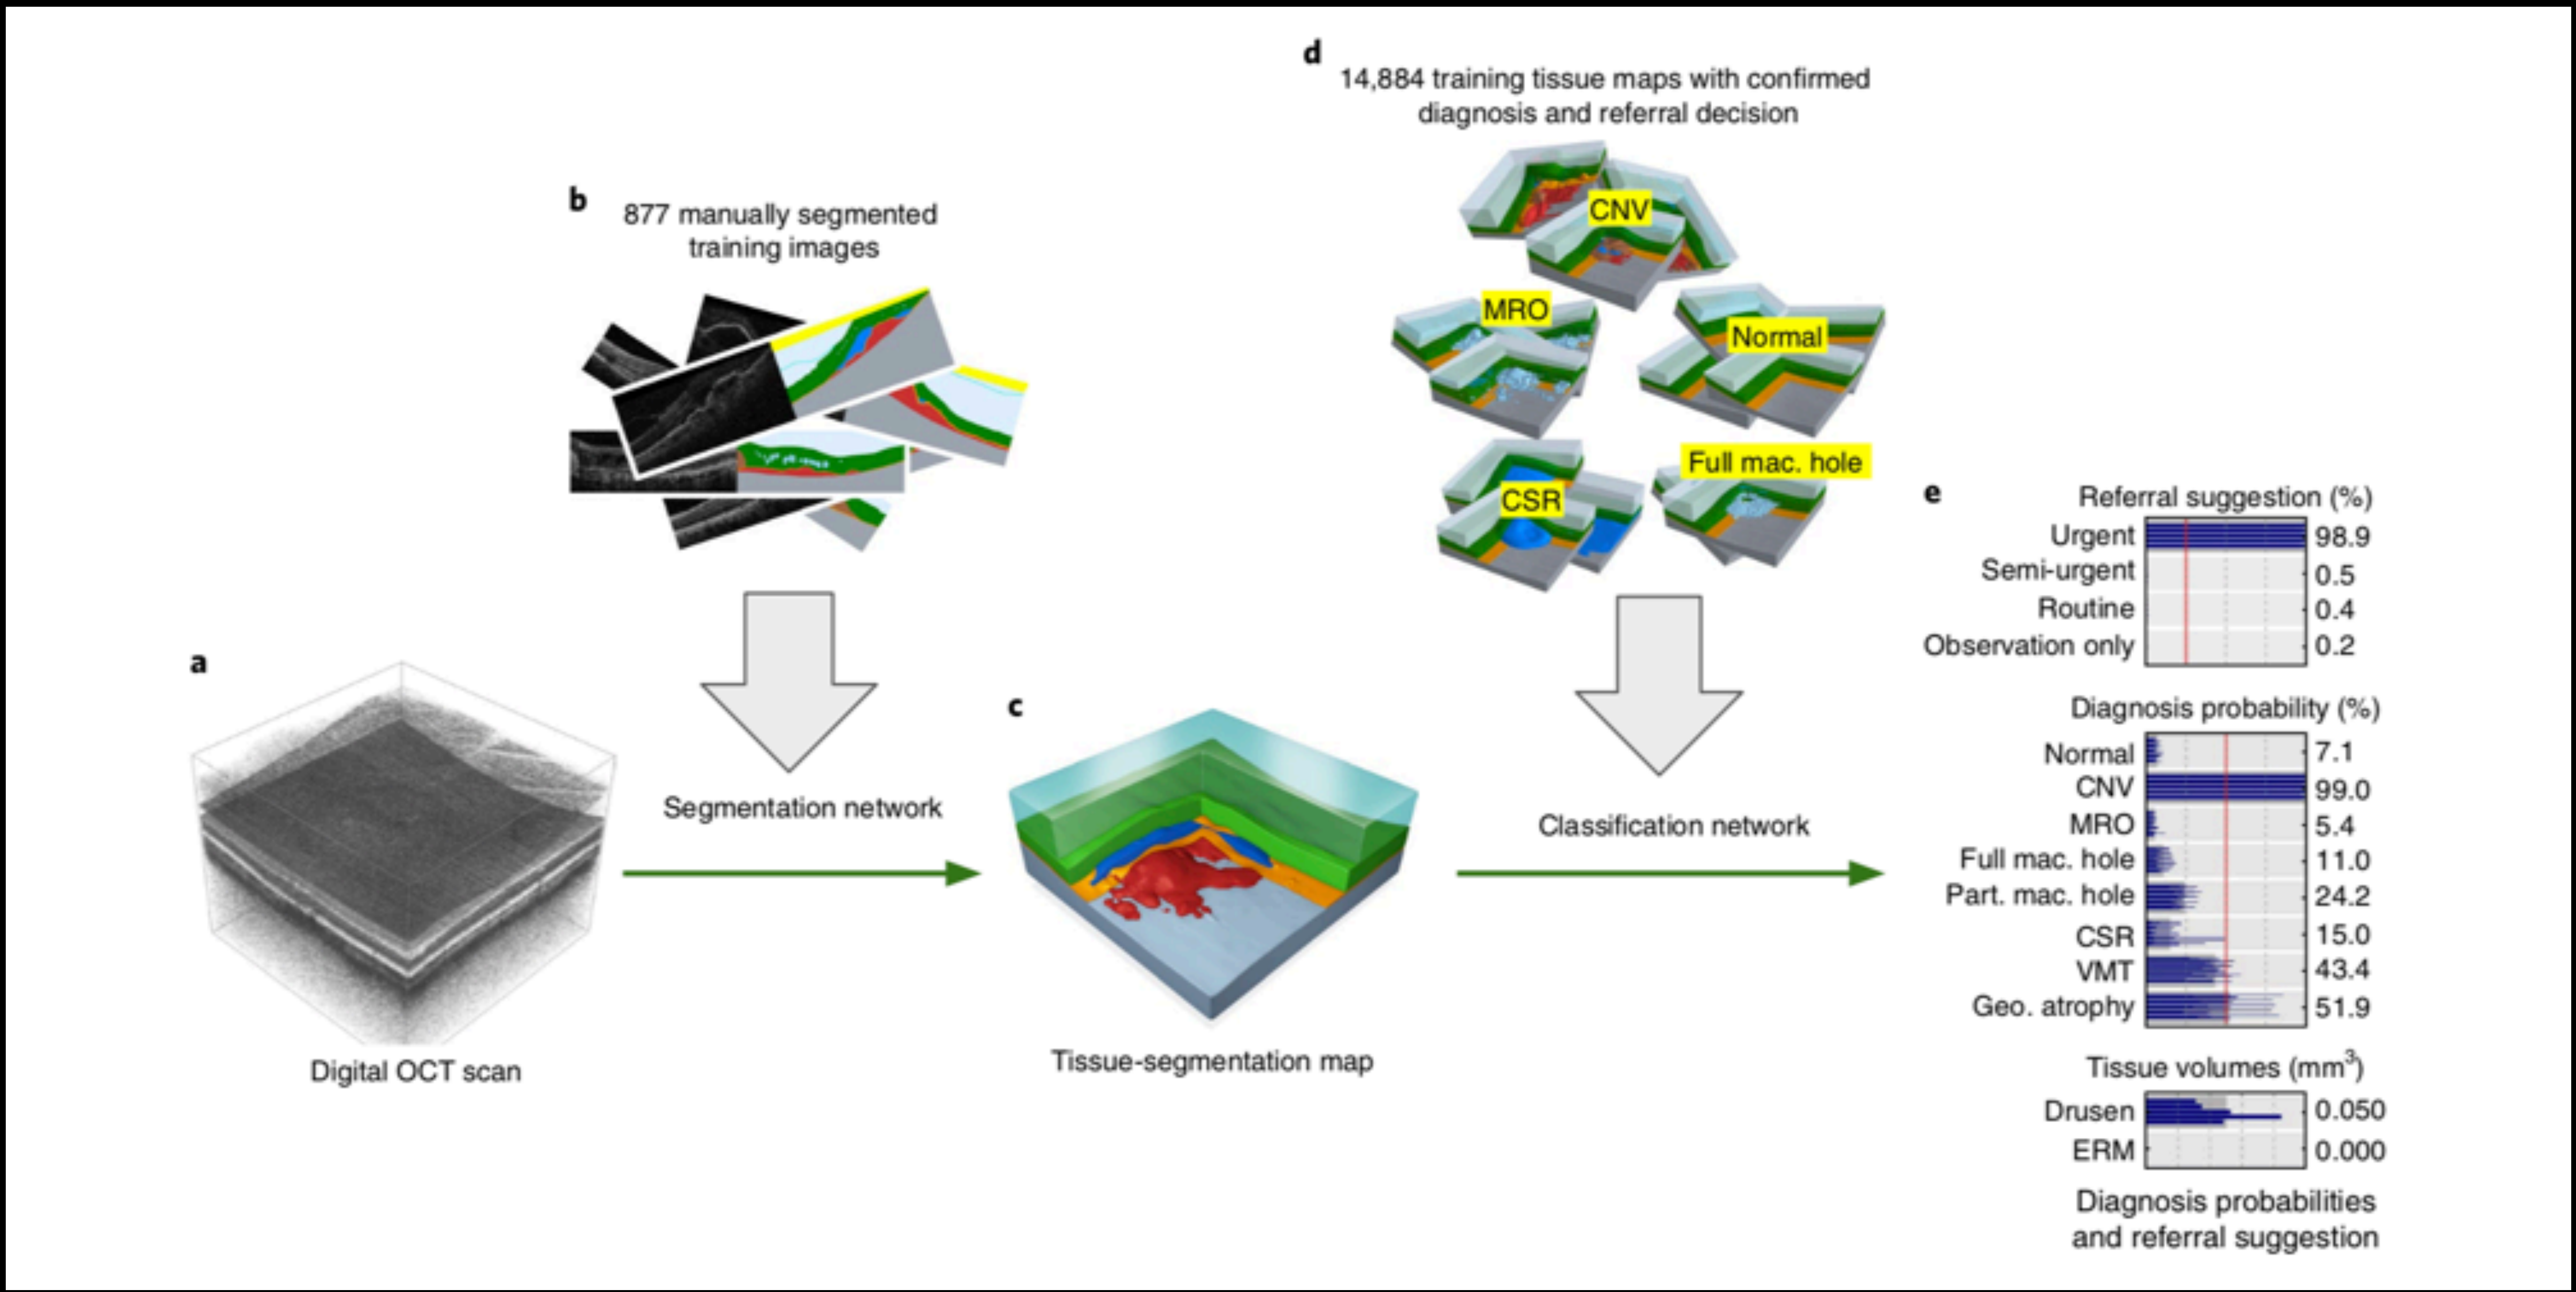

3D Model

Segmentation Outputs

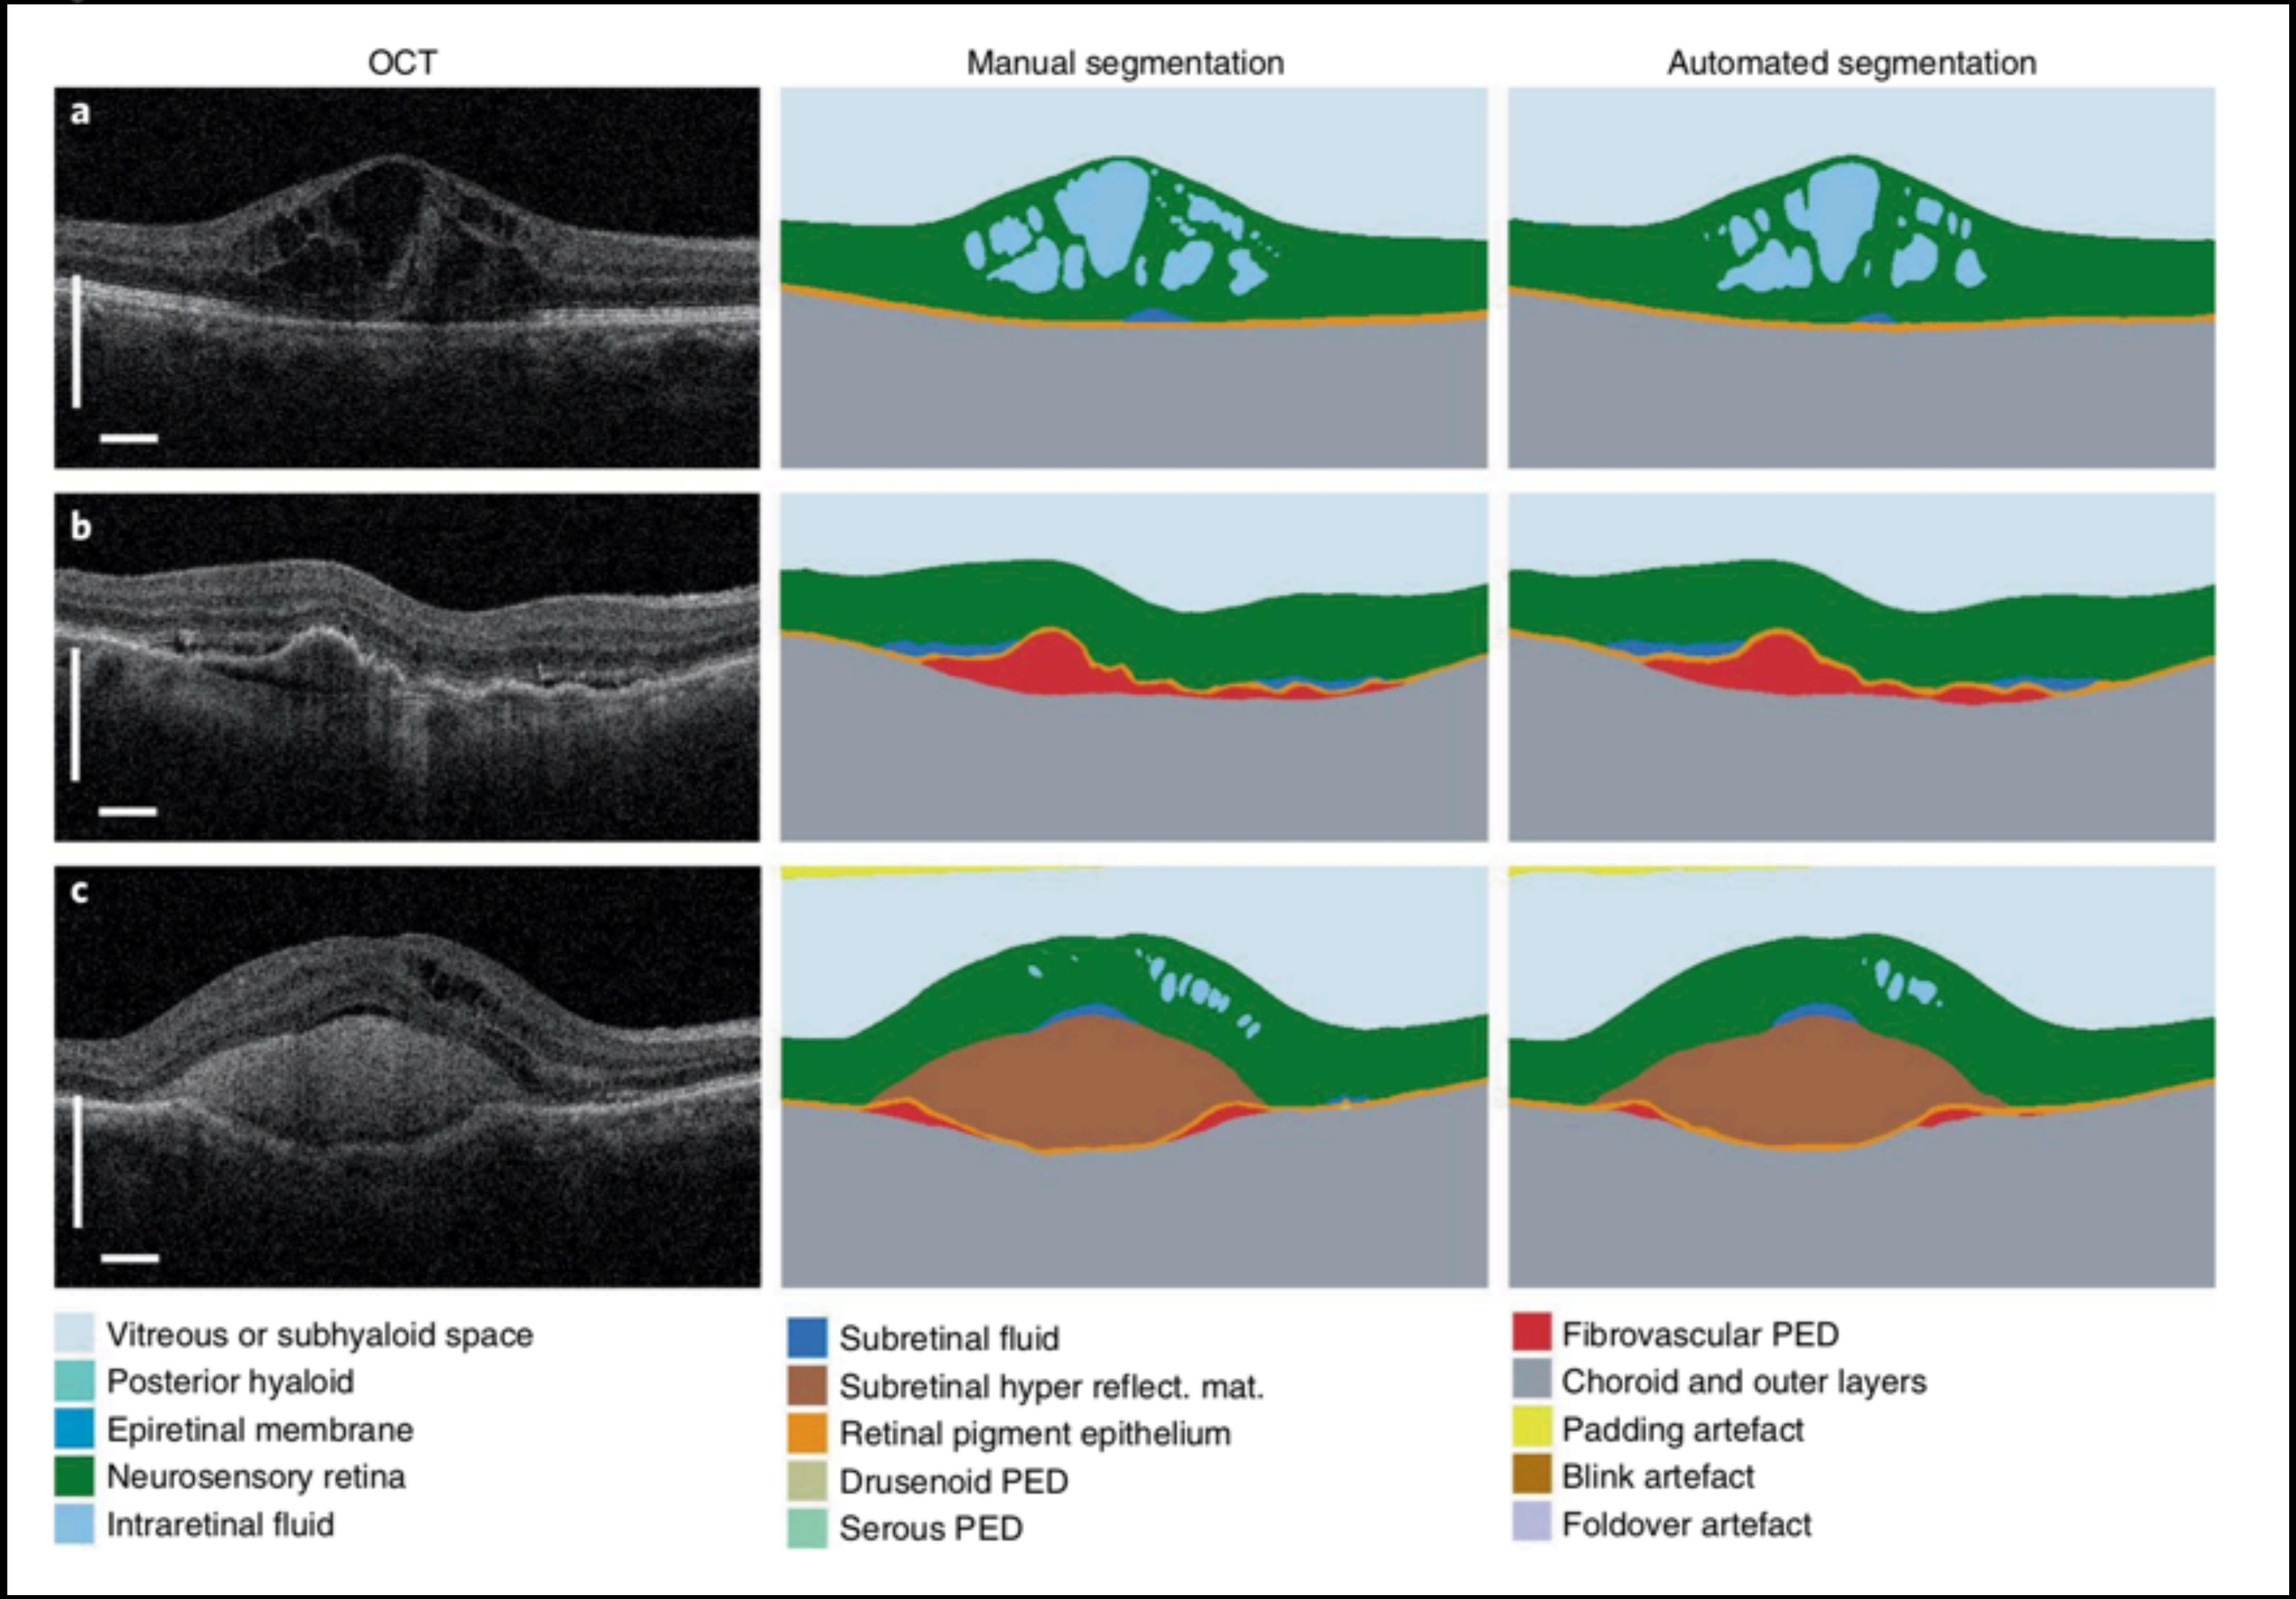

Referral Categories

Referral Categories

Referral Category	Definition
Urgent	All causes of choroidal neovascularization, including age related macular degeneration, high myopia, central serous retinopathy, inherited retinal dystrophies (e.g., angioid streaks), posterior uveitis (e.g., multiple choroiditis), and post traumatic choroidal rupture.
Semi-urgent	Referable edema classed as semi-urgent included diabetic maculopathy, retinal vein occlusion, postoperative (Irvine-Gass syndrome), uveitis, Coat's disease, radiation and miscellaneous other cases.
Routine	All other non-urgent cases with a large variety, from uncomplicated central serous retinopathy to more rare conditions such as Macular Telangiectasia (MacTel) type 2.
Observation only	The absence of pathology classes described above.

Prototype OCT Viewer

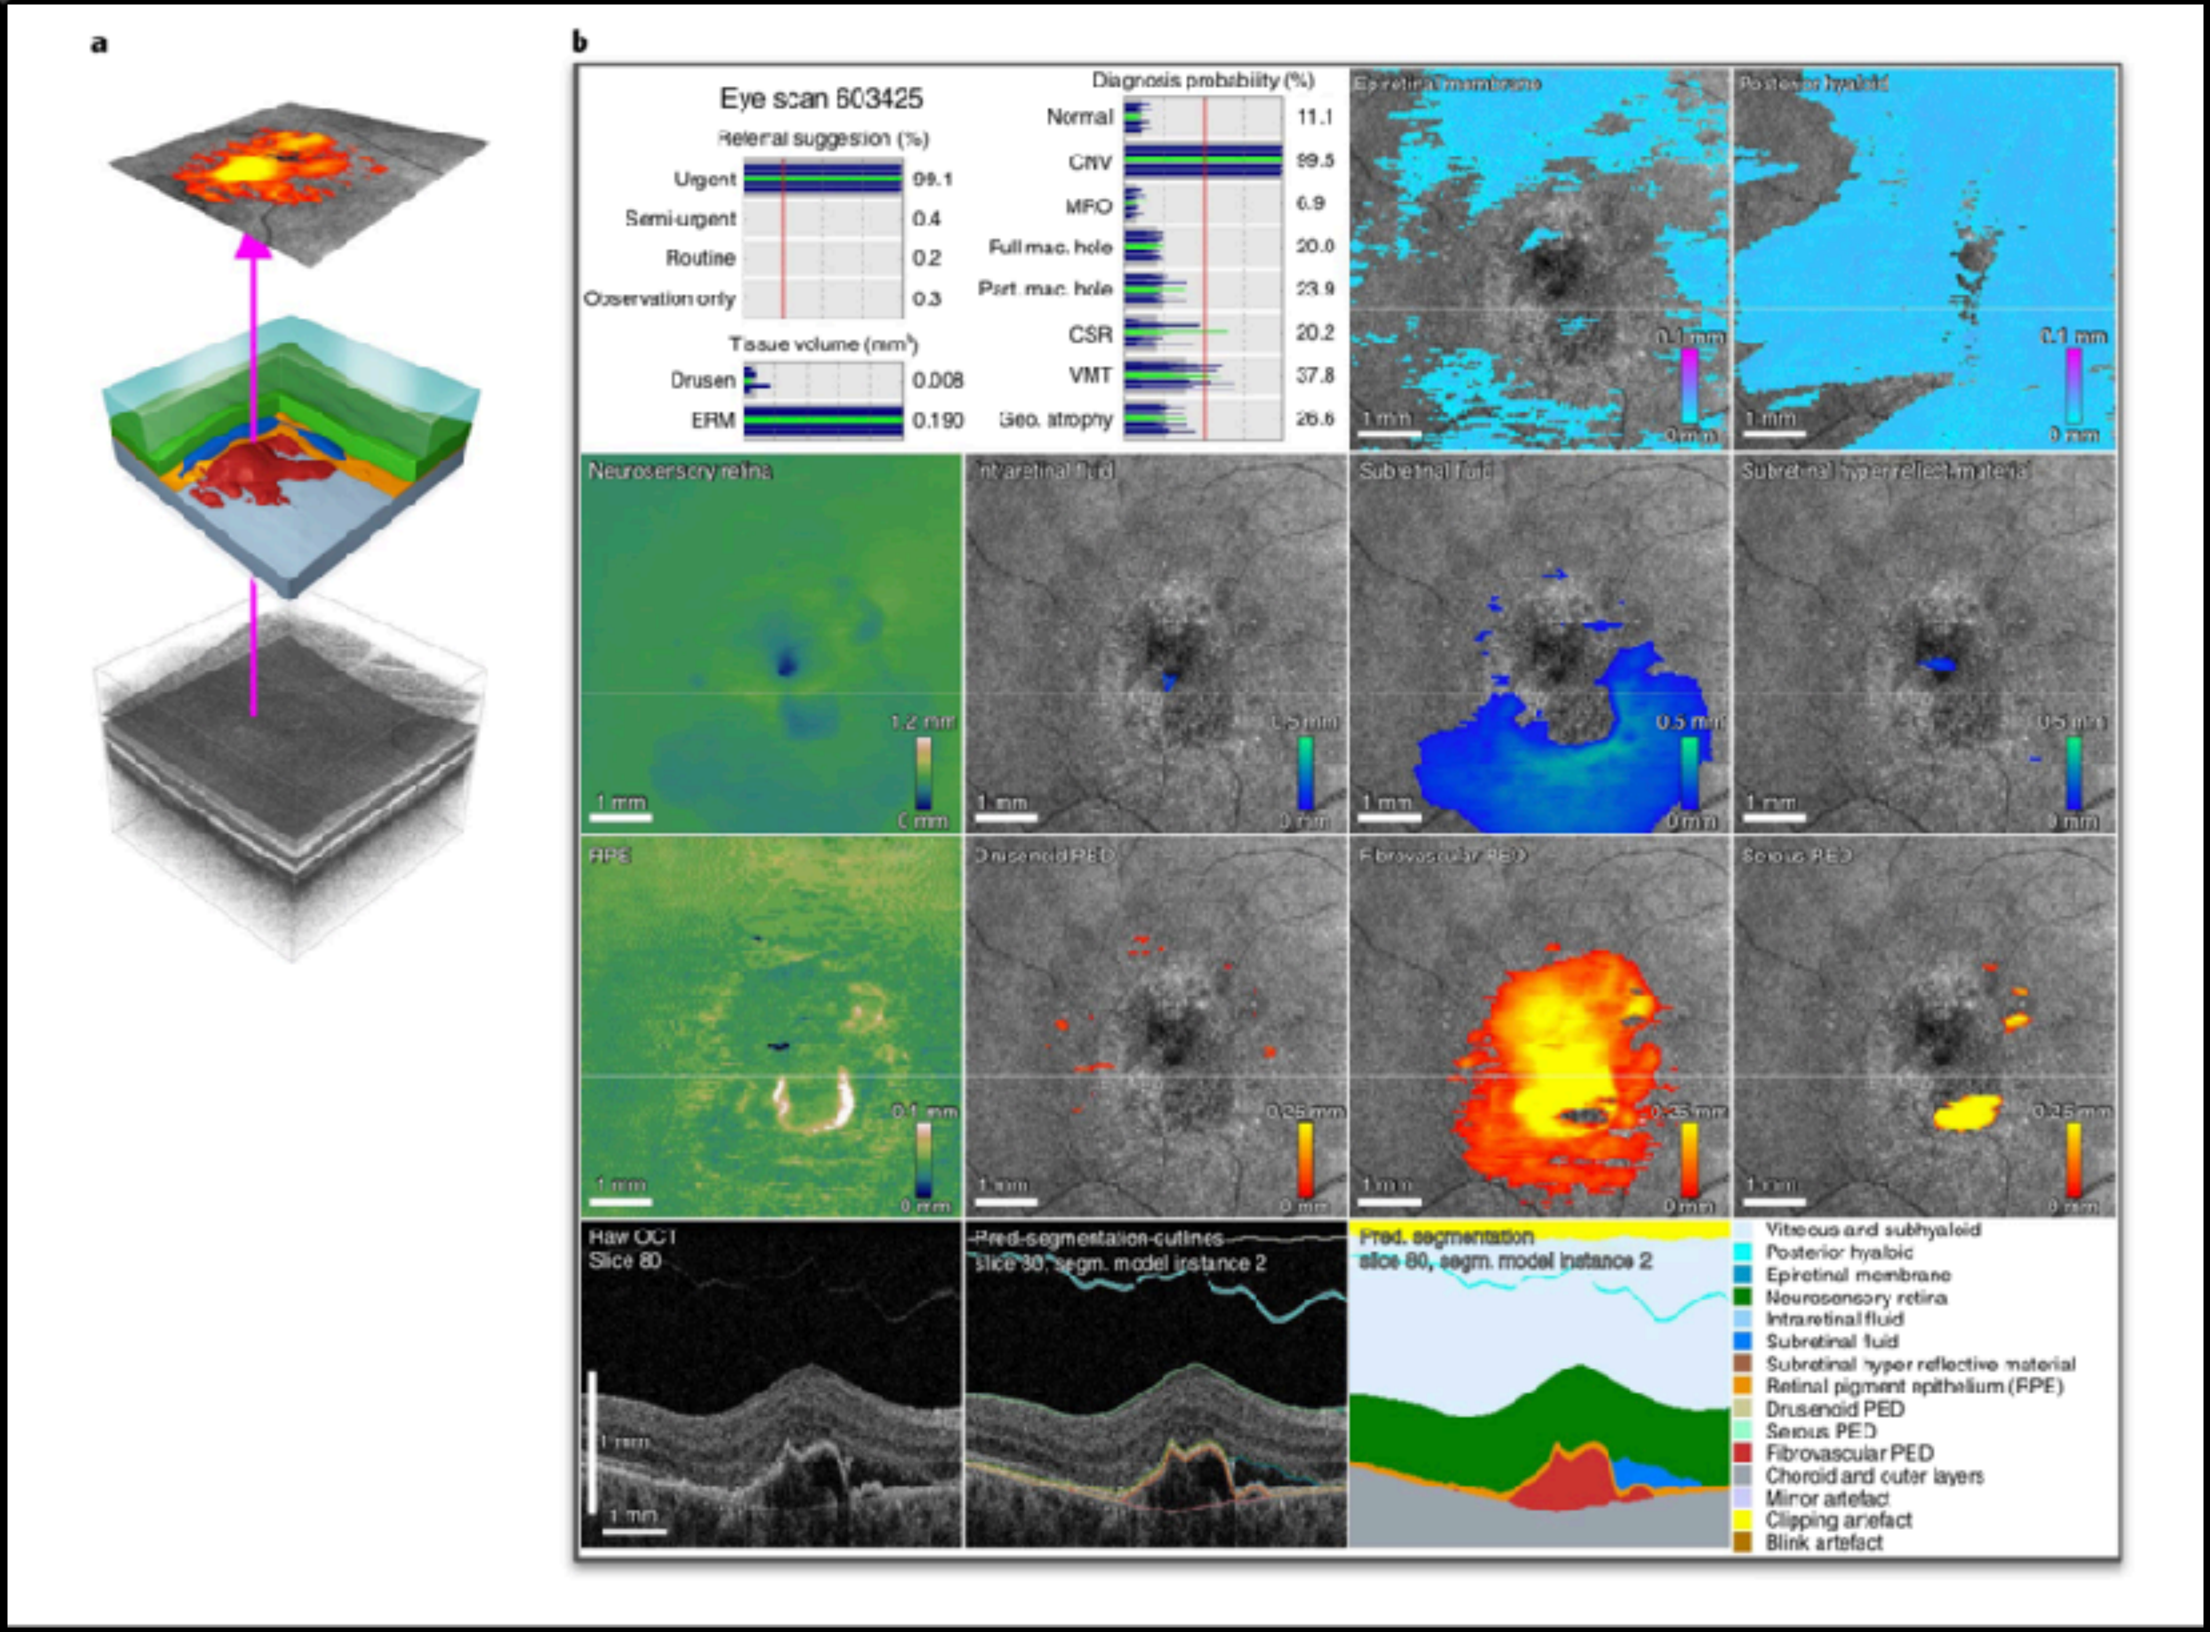

Examples - CMO

Examples - CMO

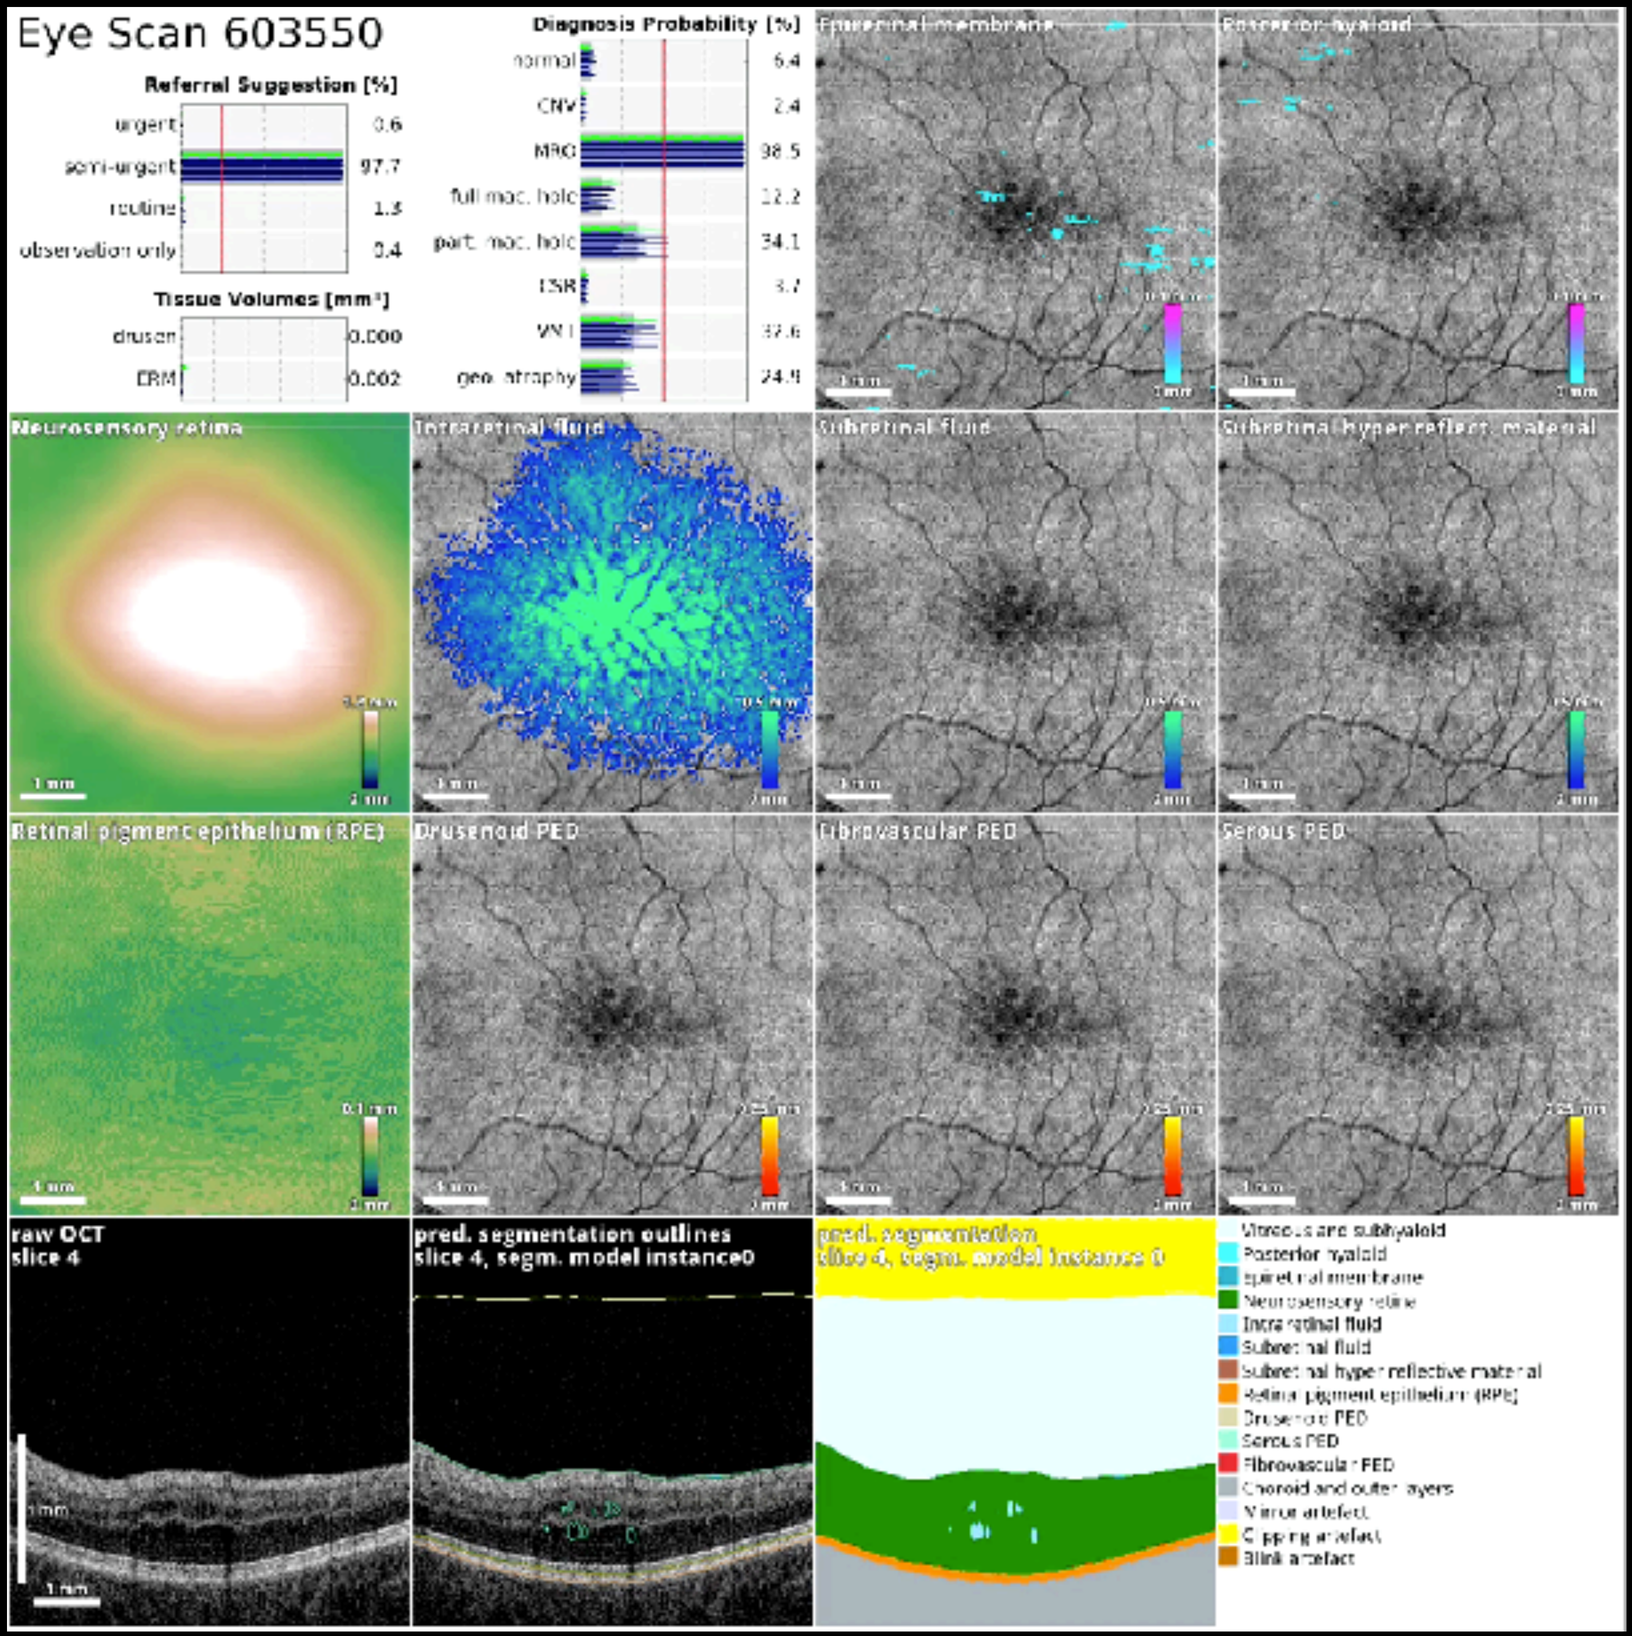

Examples - CSCR/CNV

Examples - CSCR/CNV

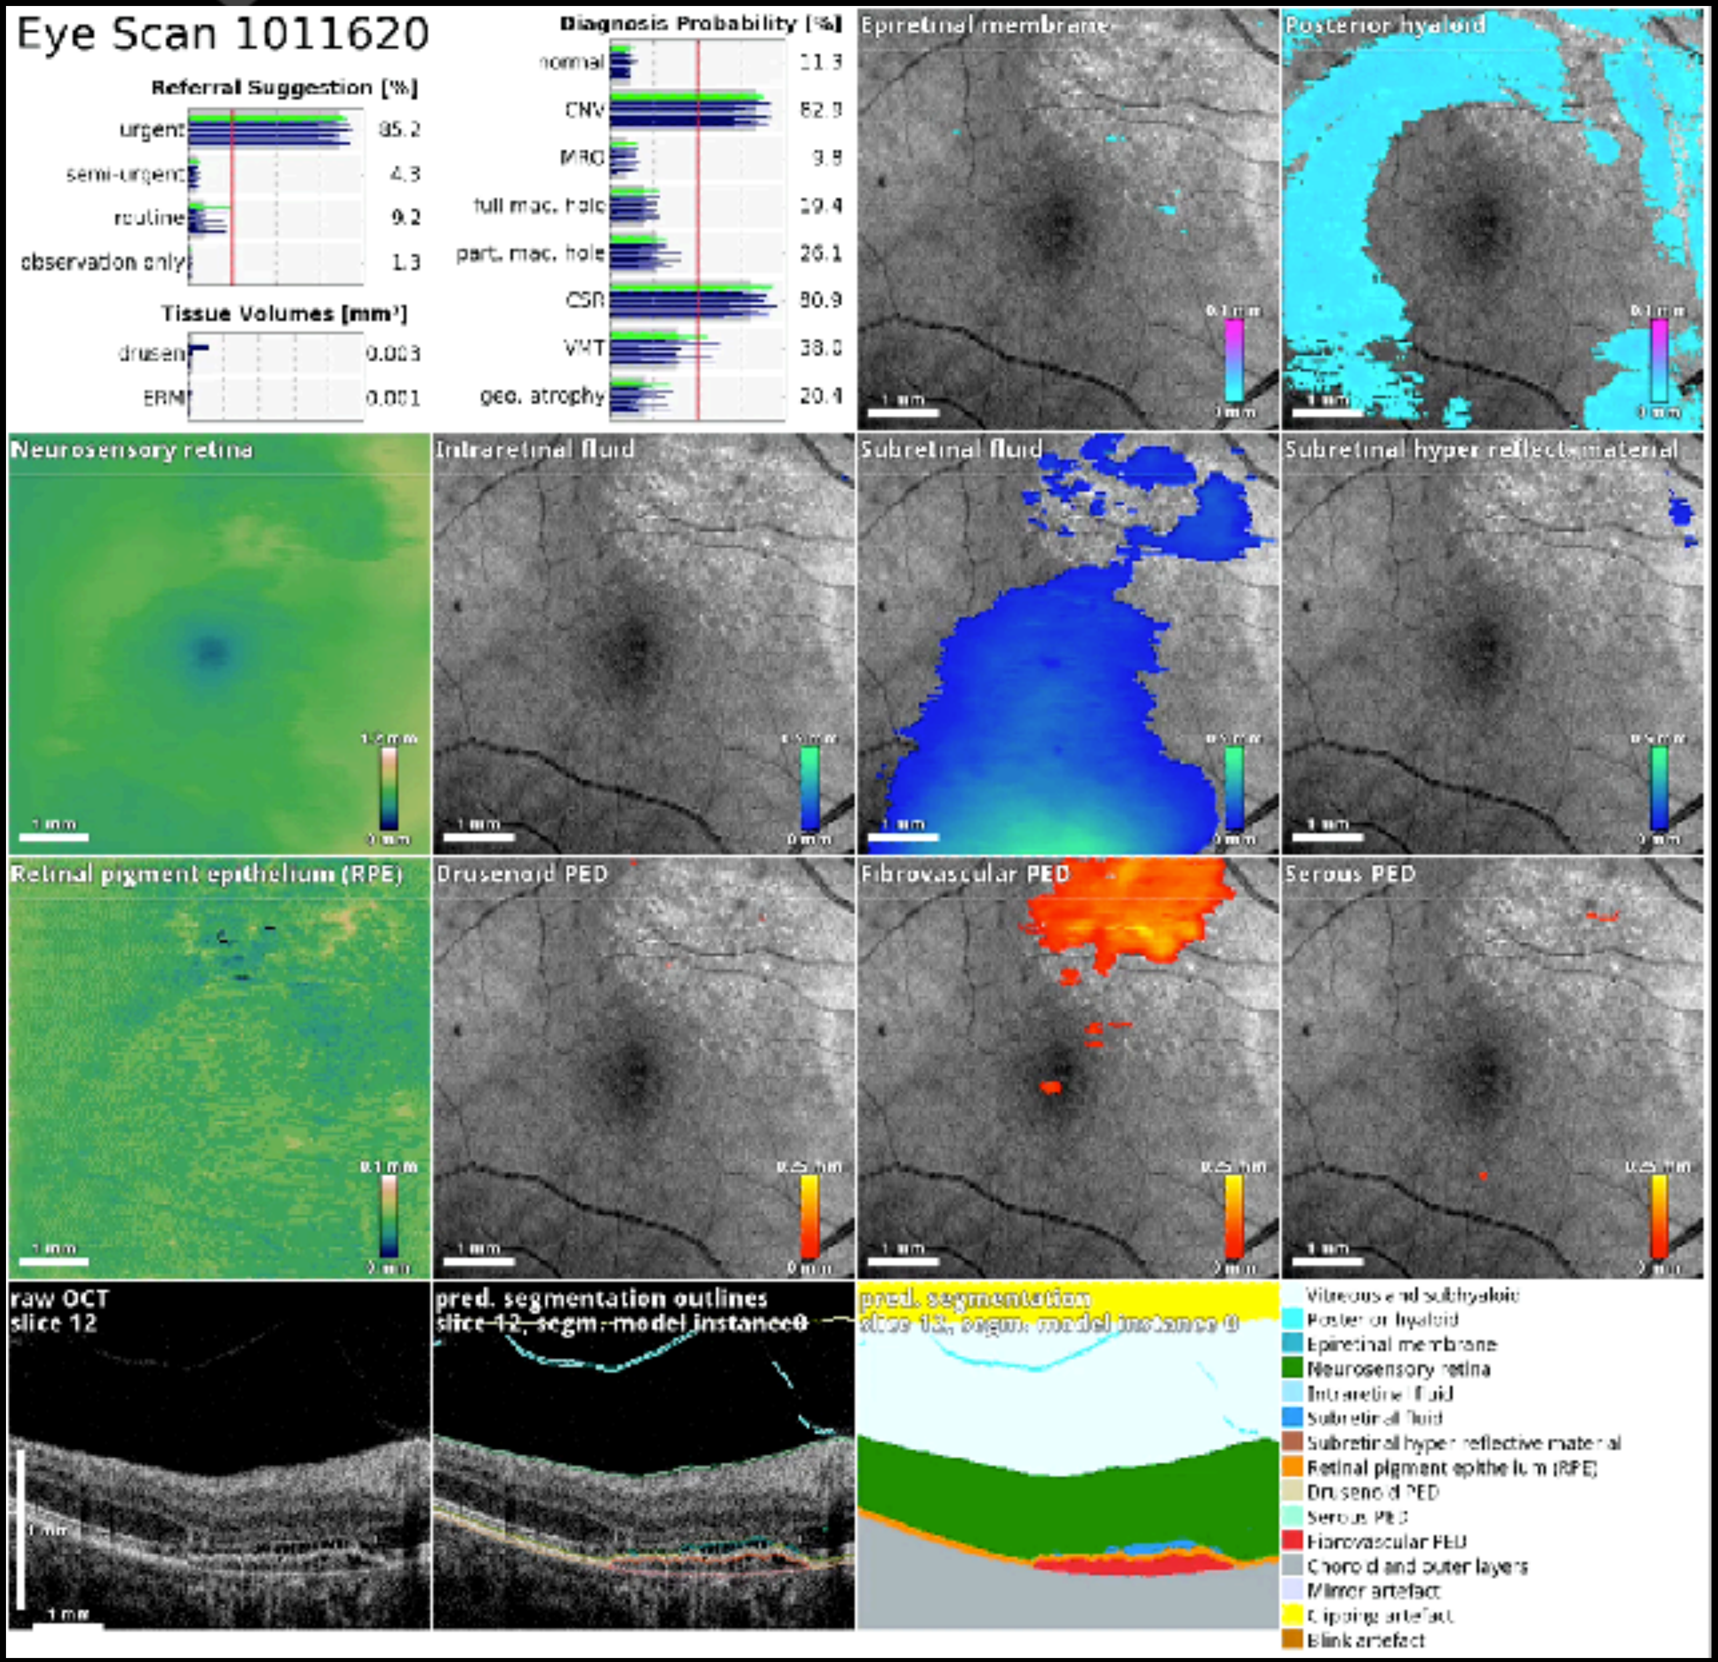

Examples - CSCR/CNV

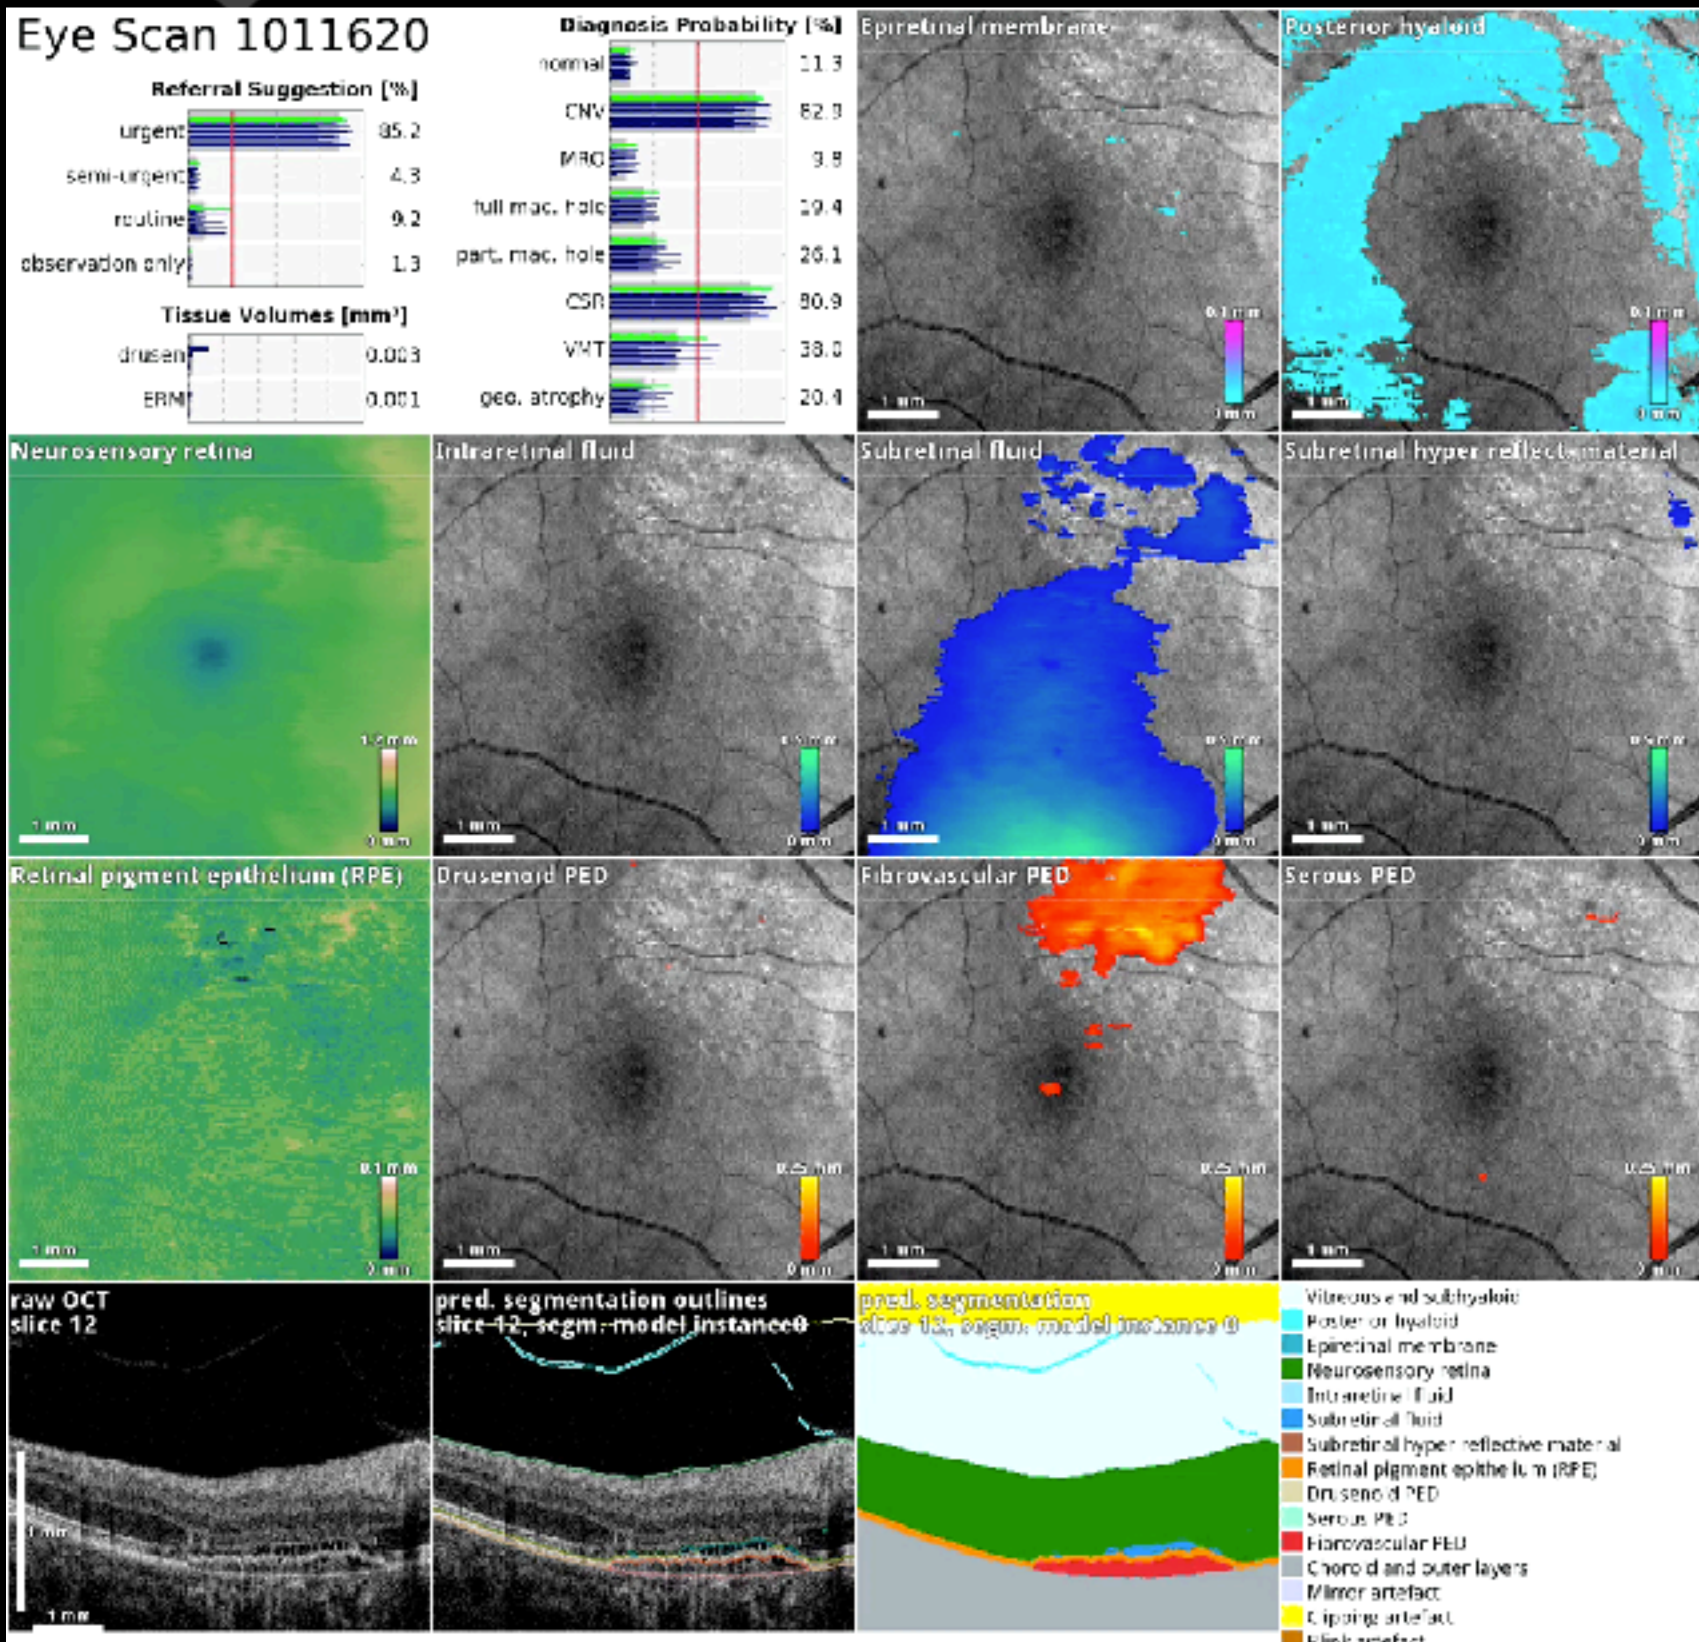

Multi-class classification

Examples - GA

Examples - GA

Ambiguous
Cases

Examples - GA

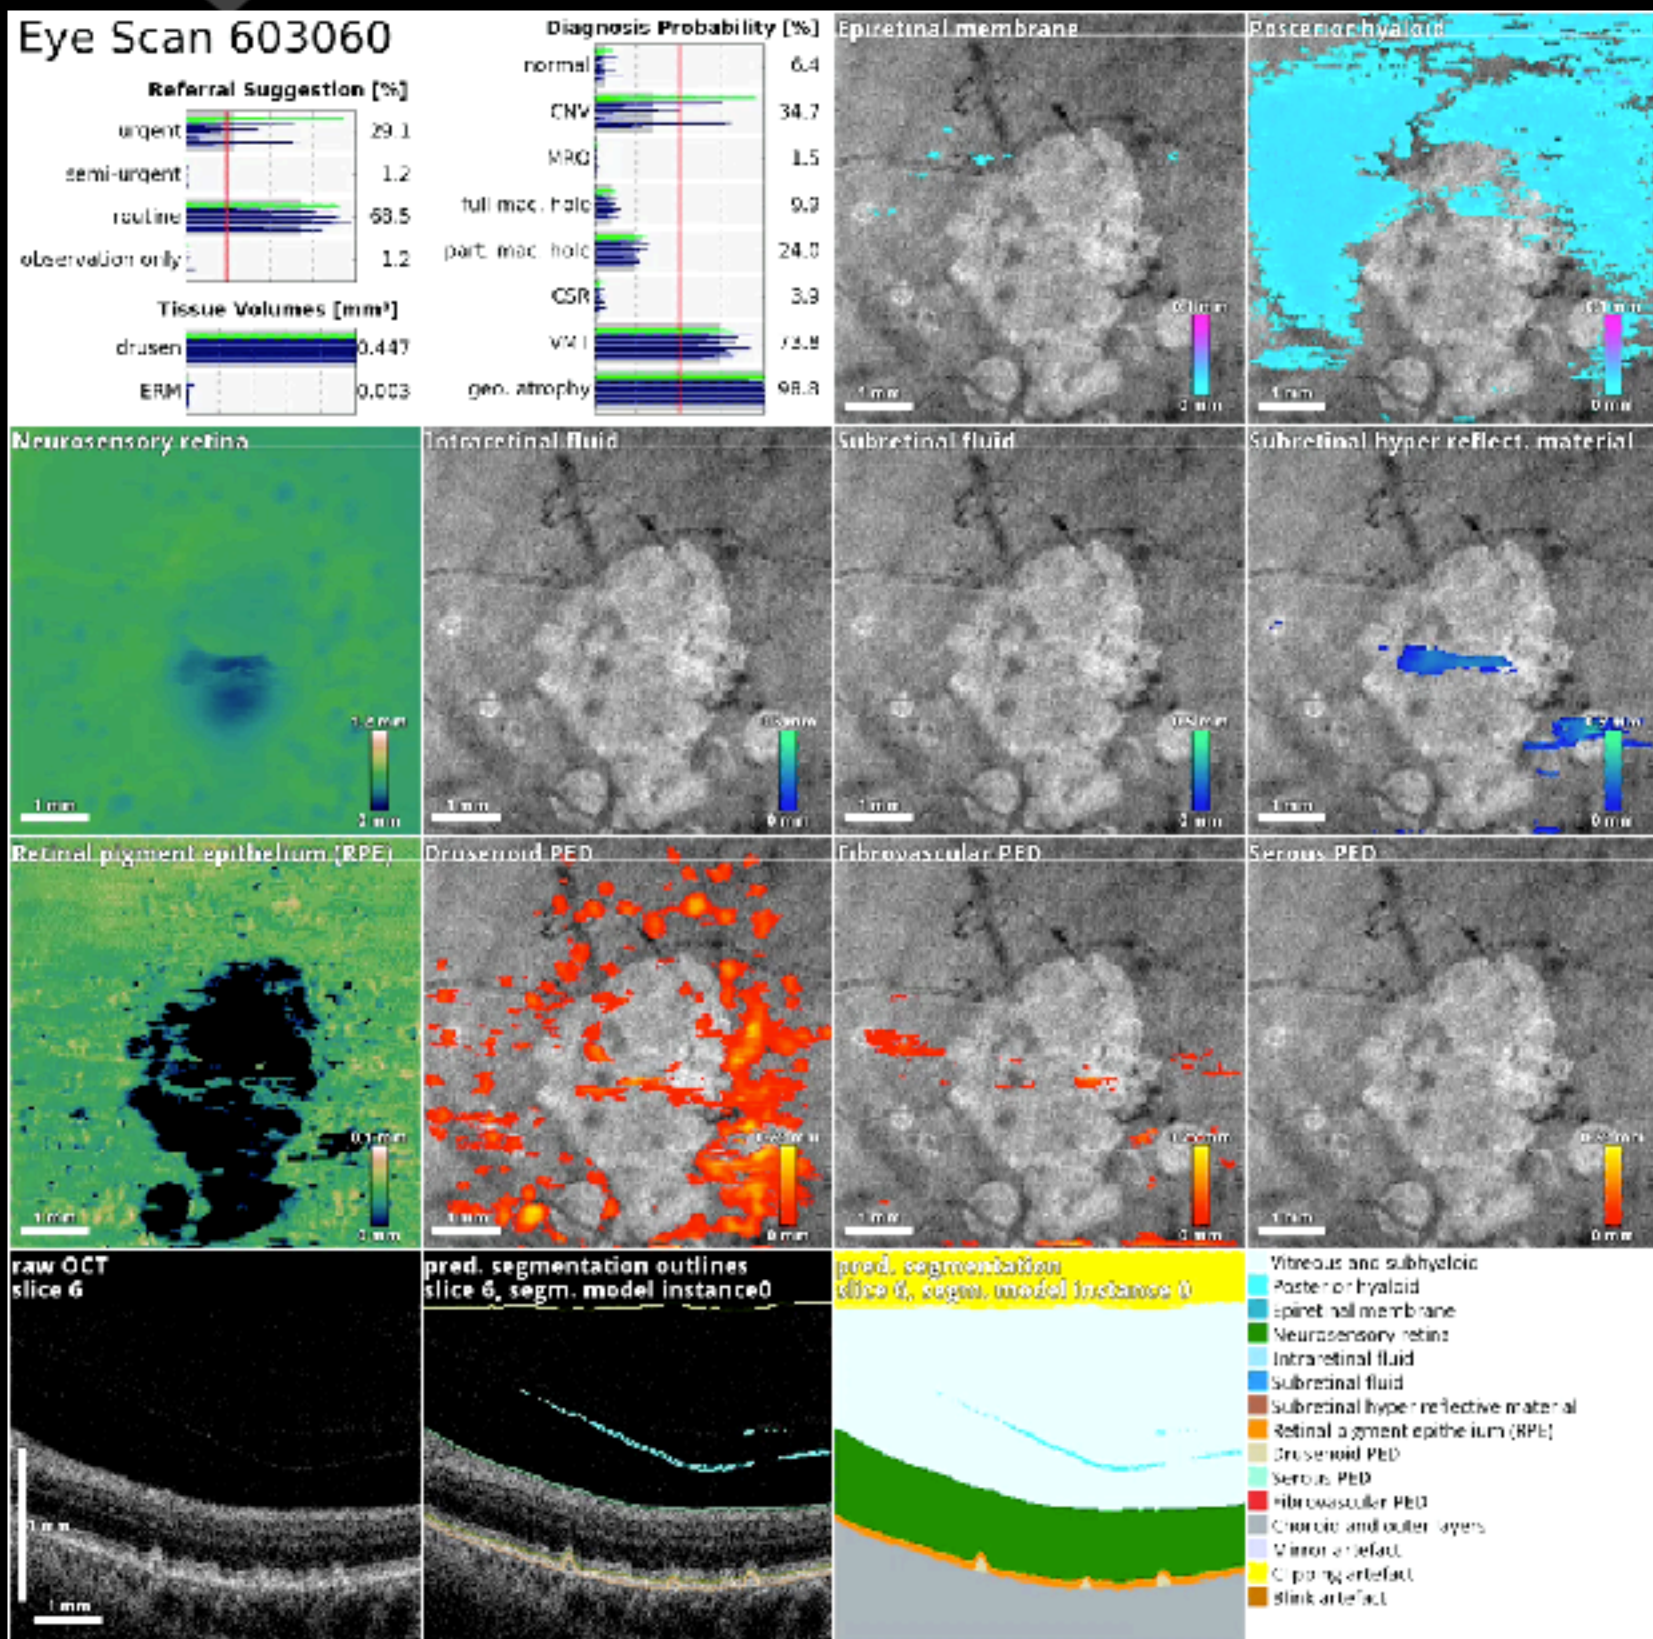

Ambiguous Cases

Examples - GA

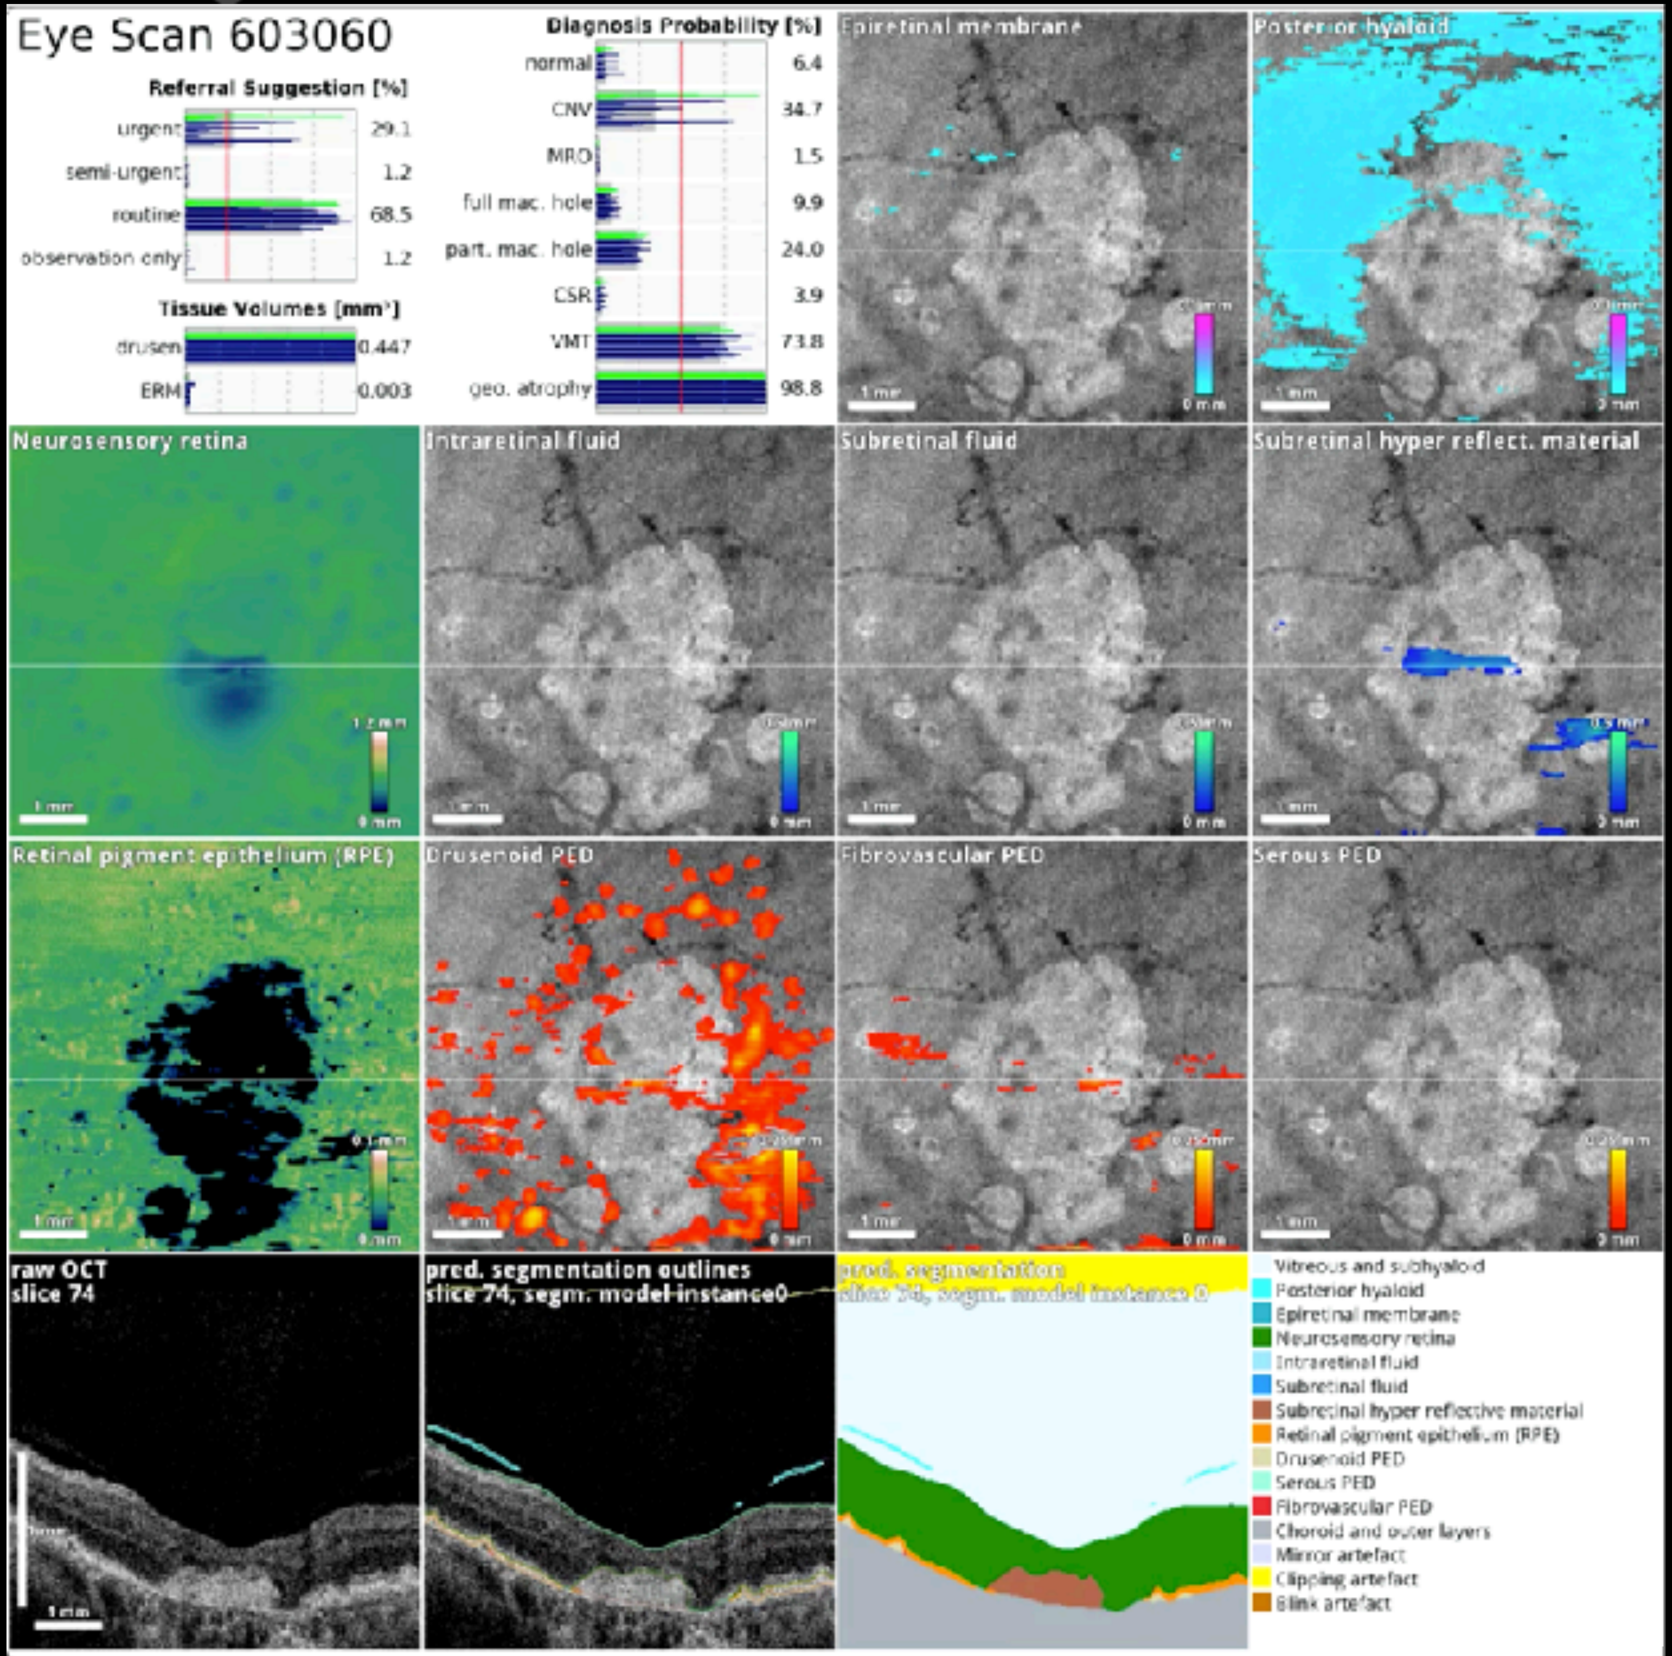

Ambiguous Cases

Examples - GA

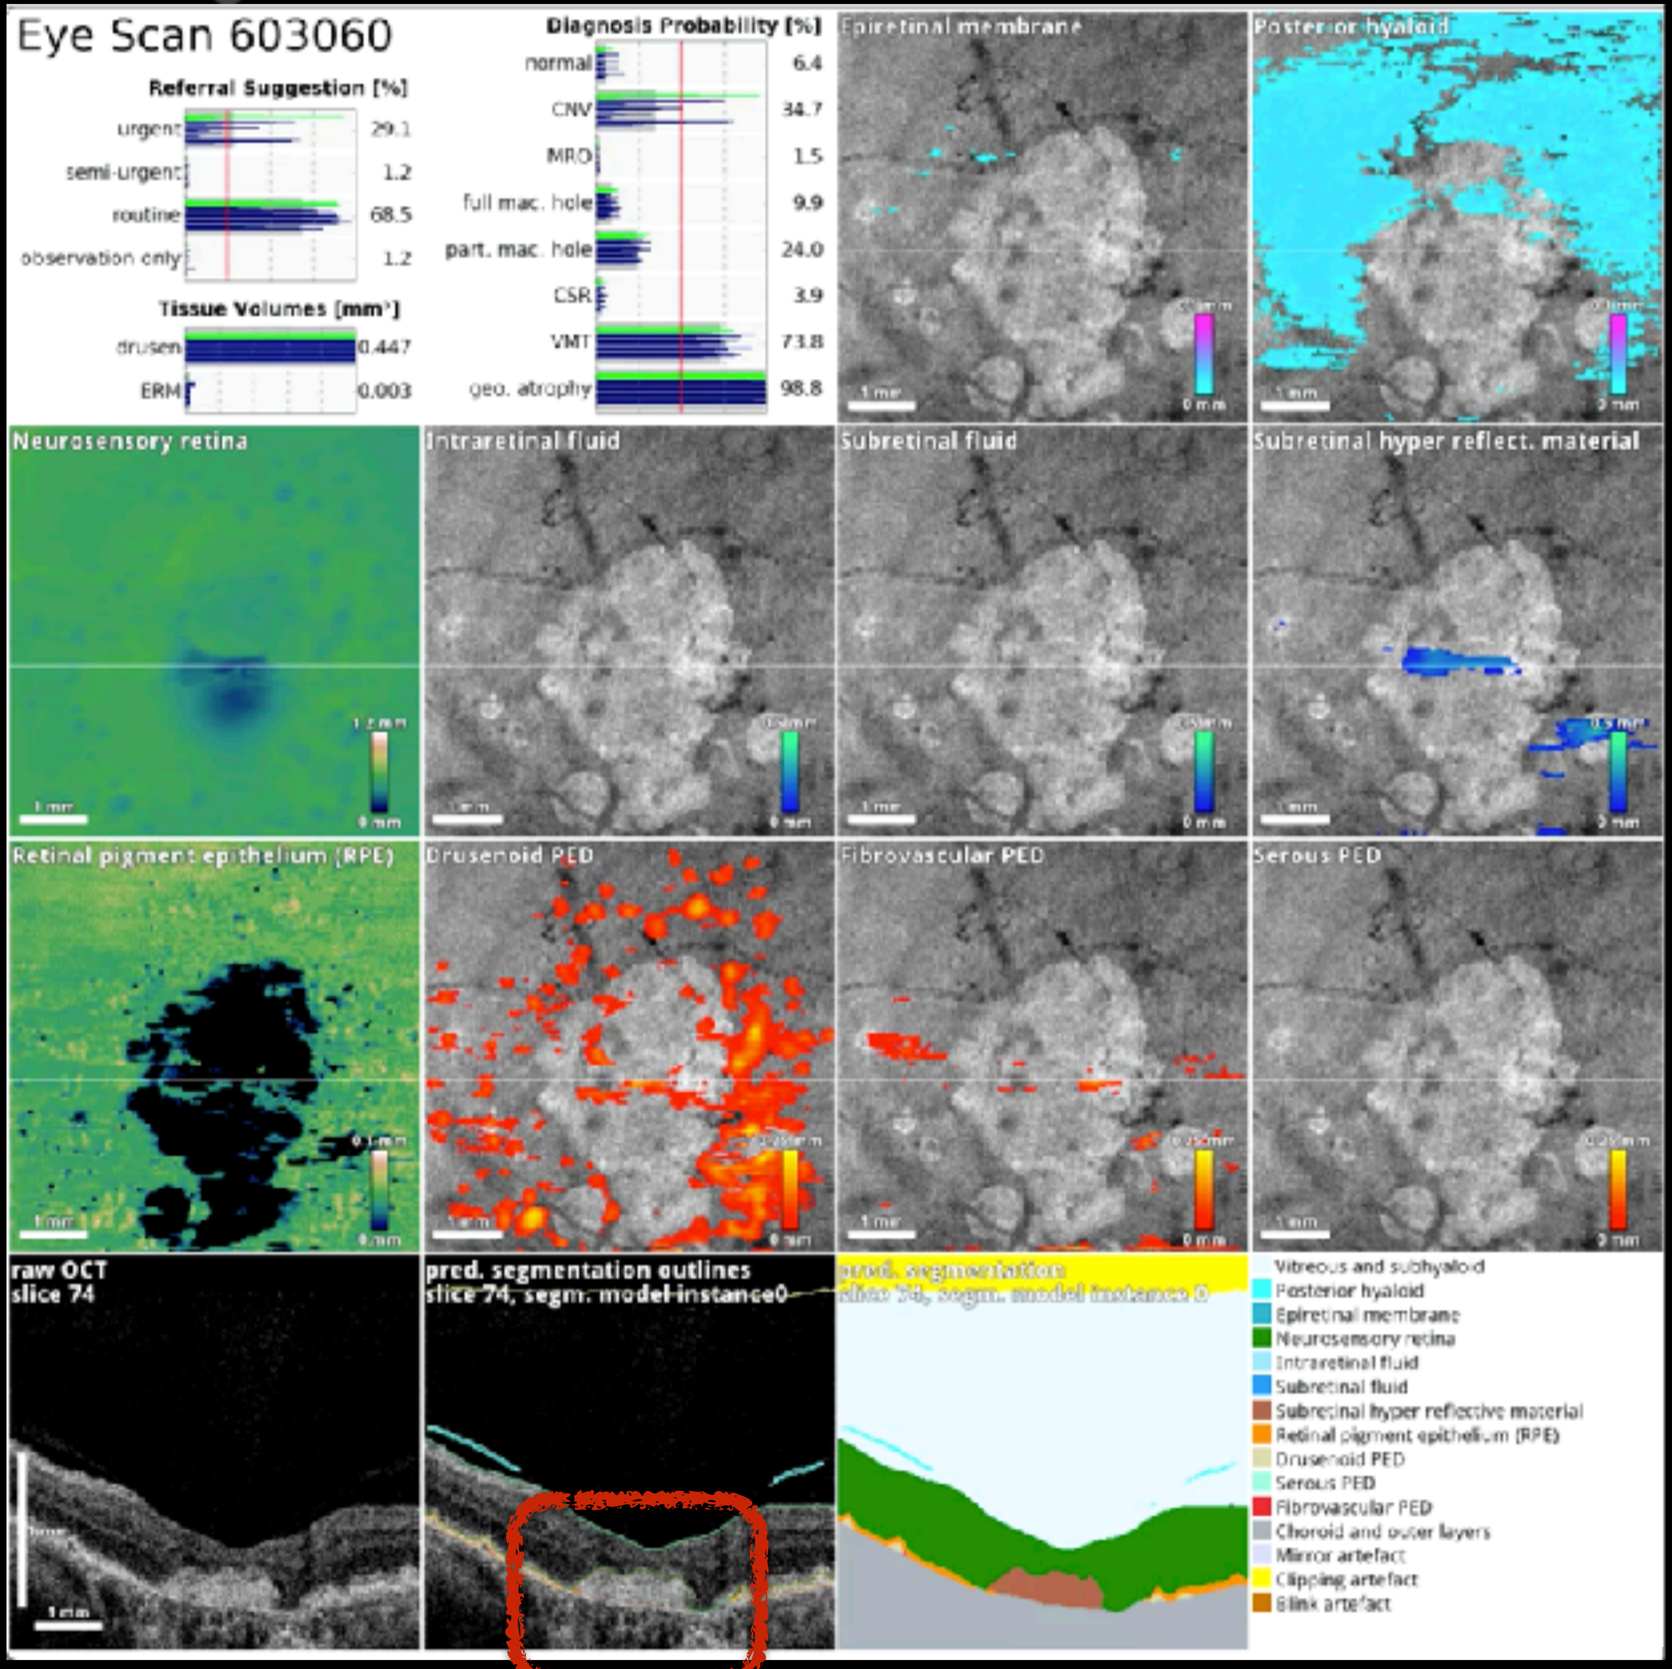

Ambiguous Cases

Examples - GA

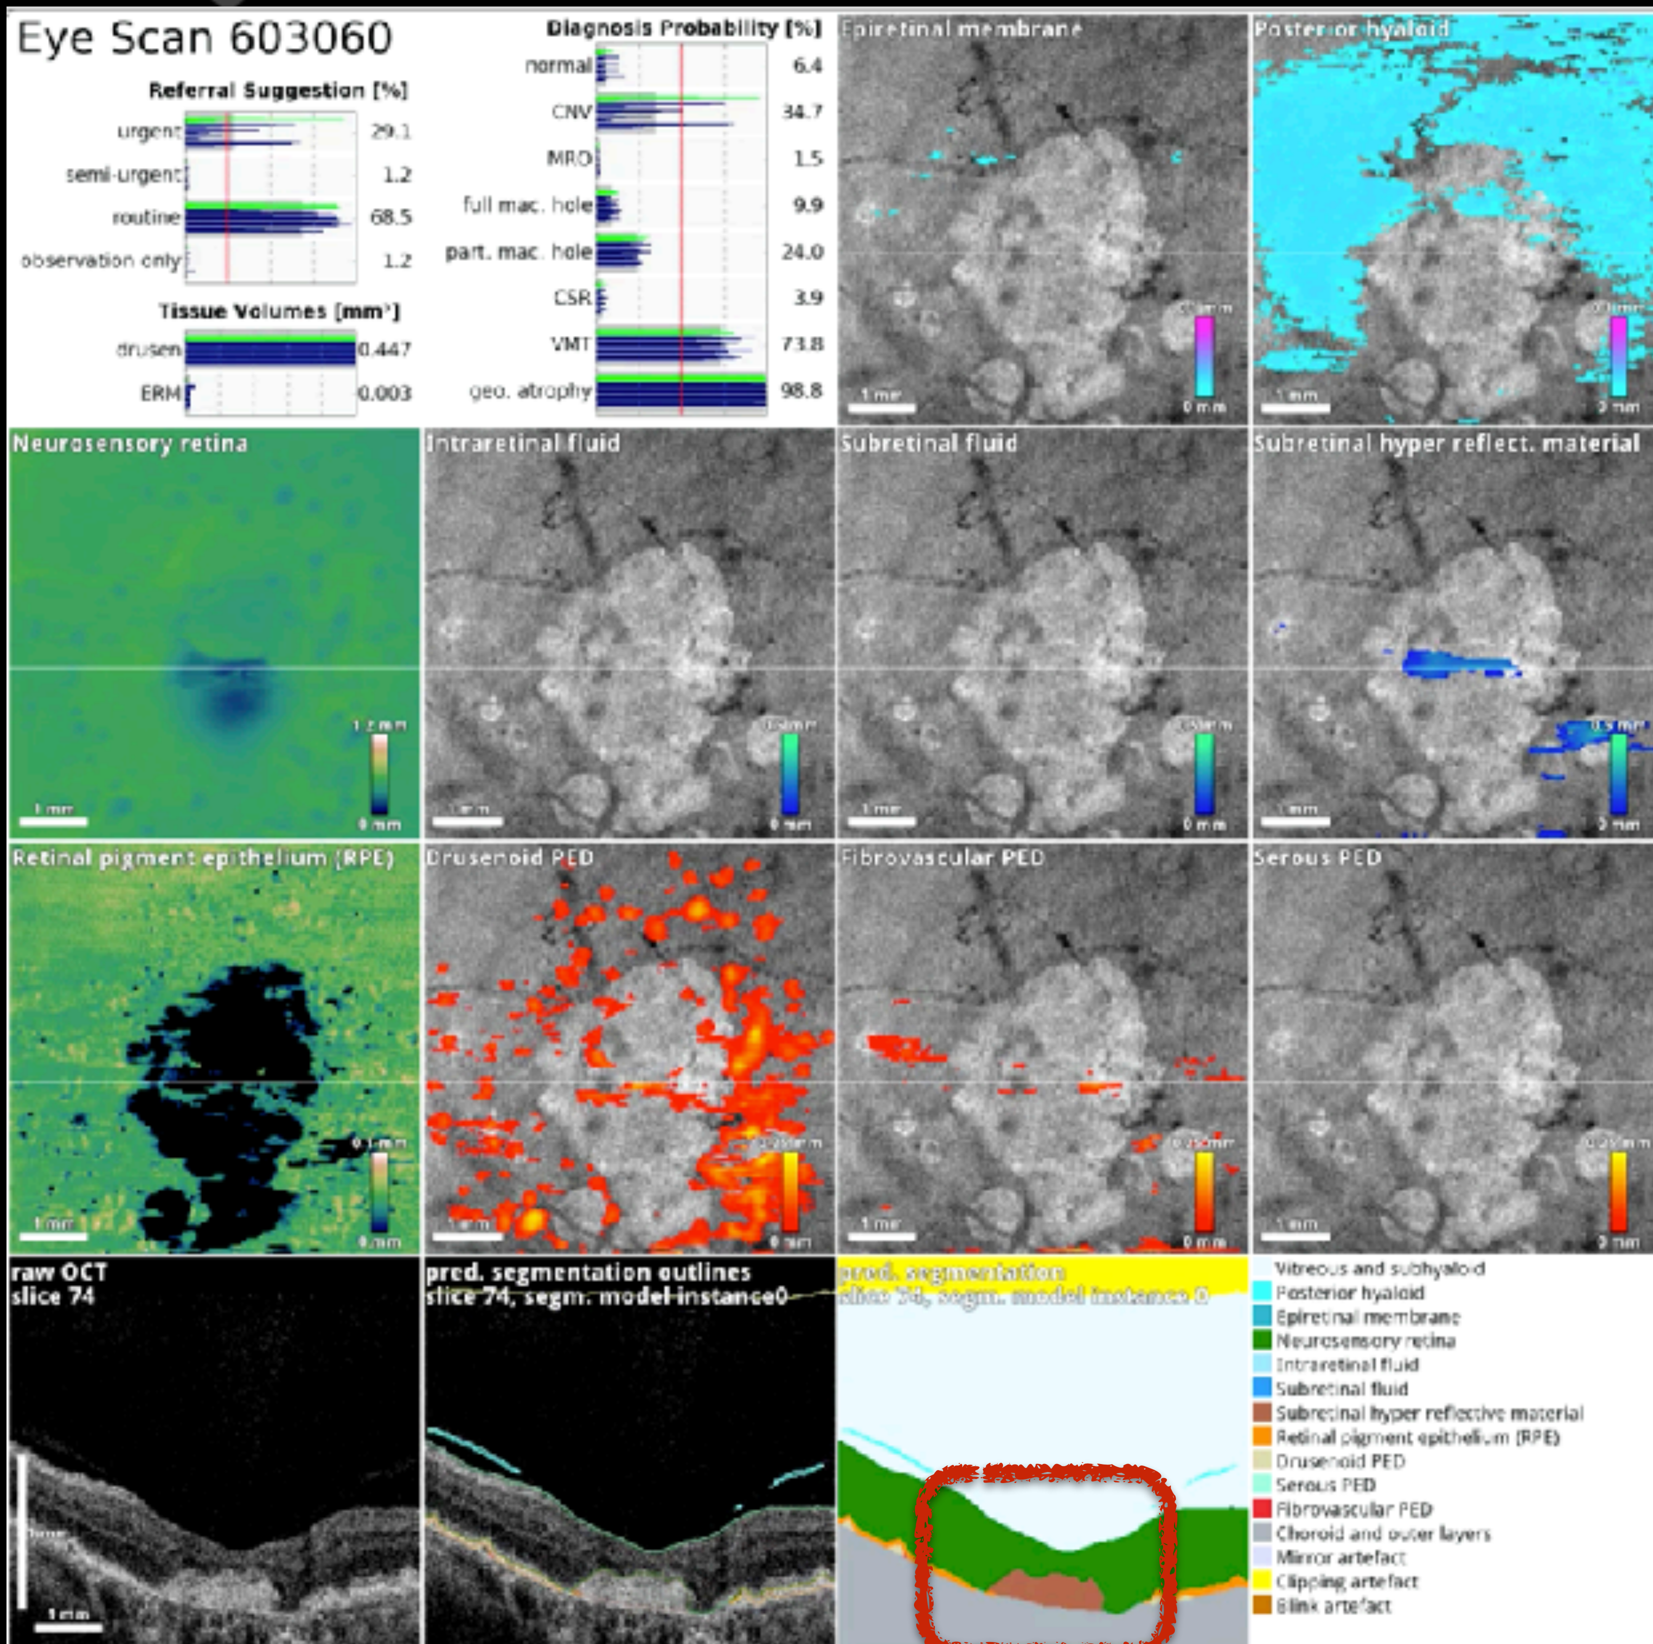

Ambiguous Cases

Evaluation of Performance

Evaluation of Performance

1000 “New” Patients Moorfields

Evaluation of Performance

1000 “New” Patients Moorfields

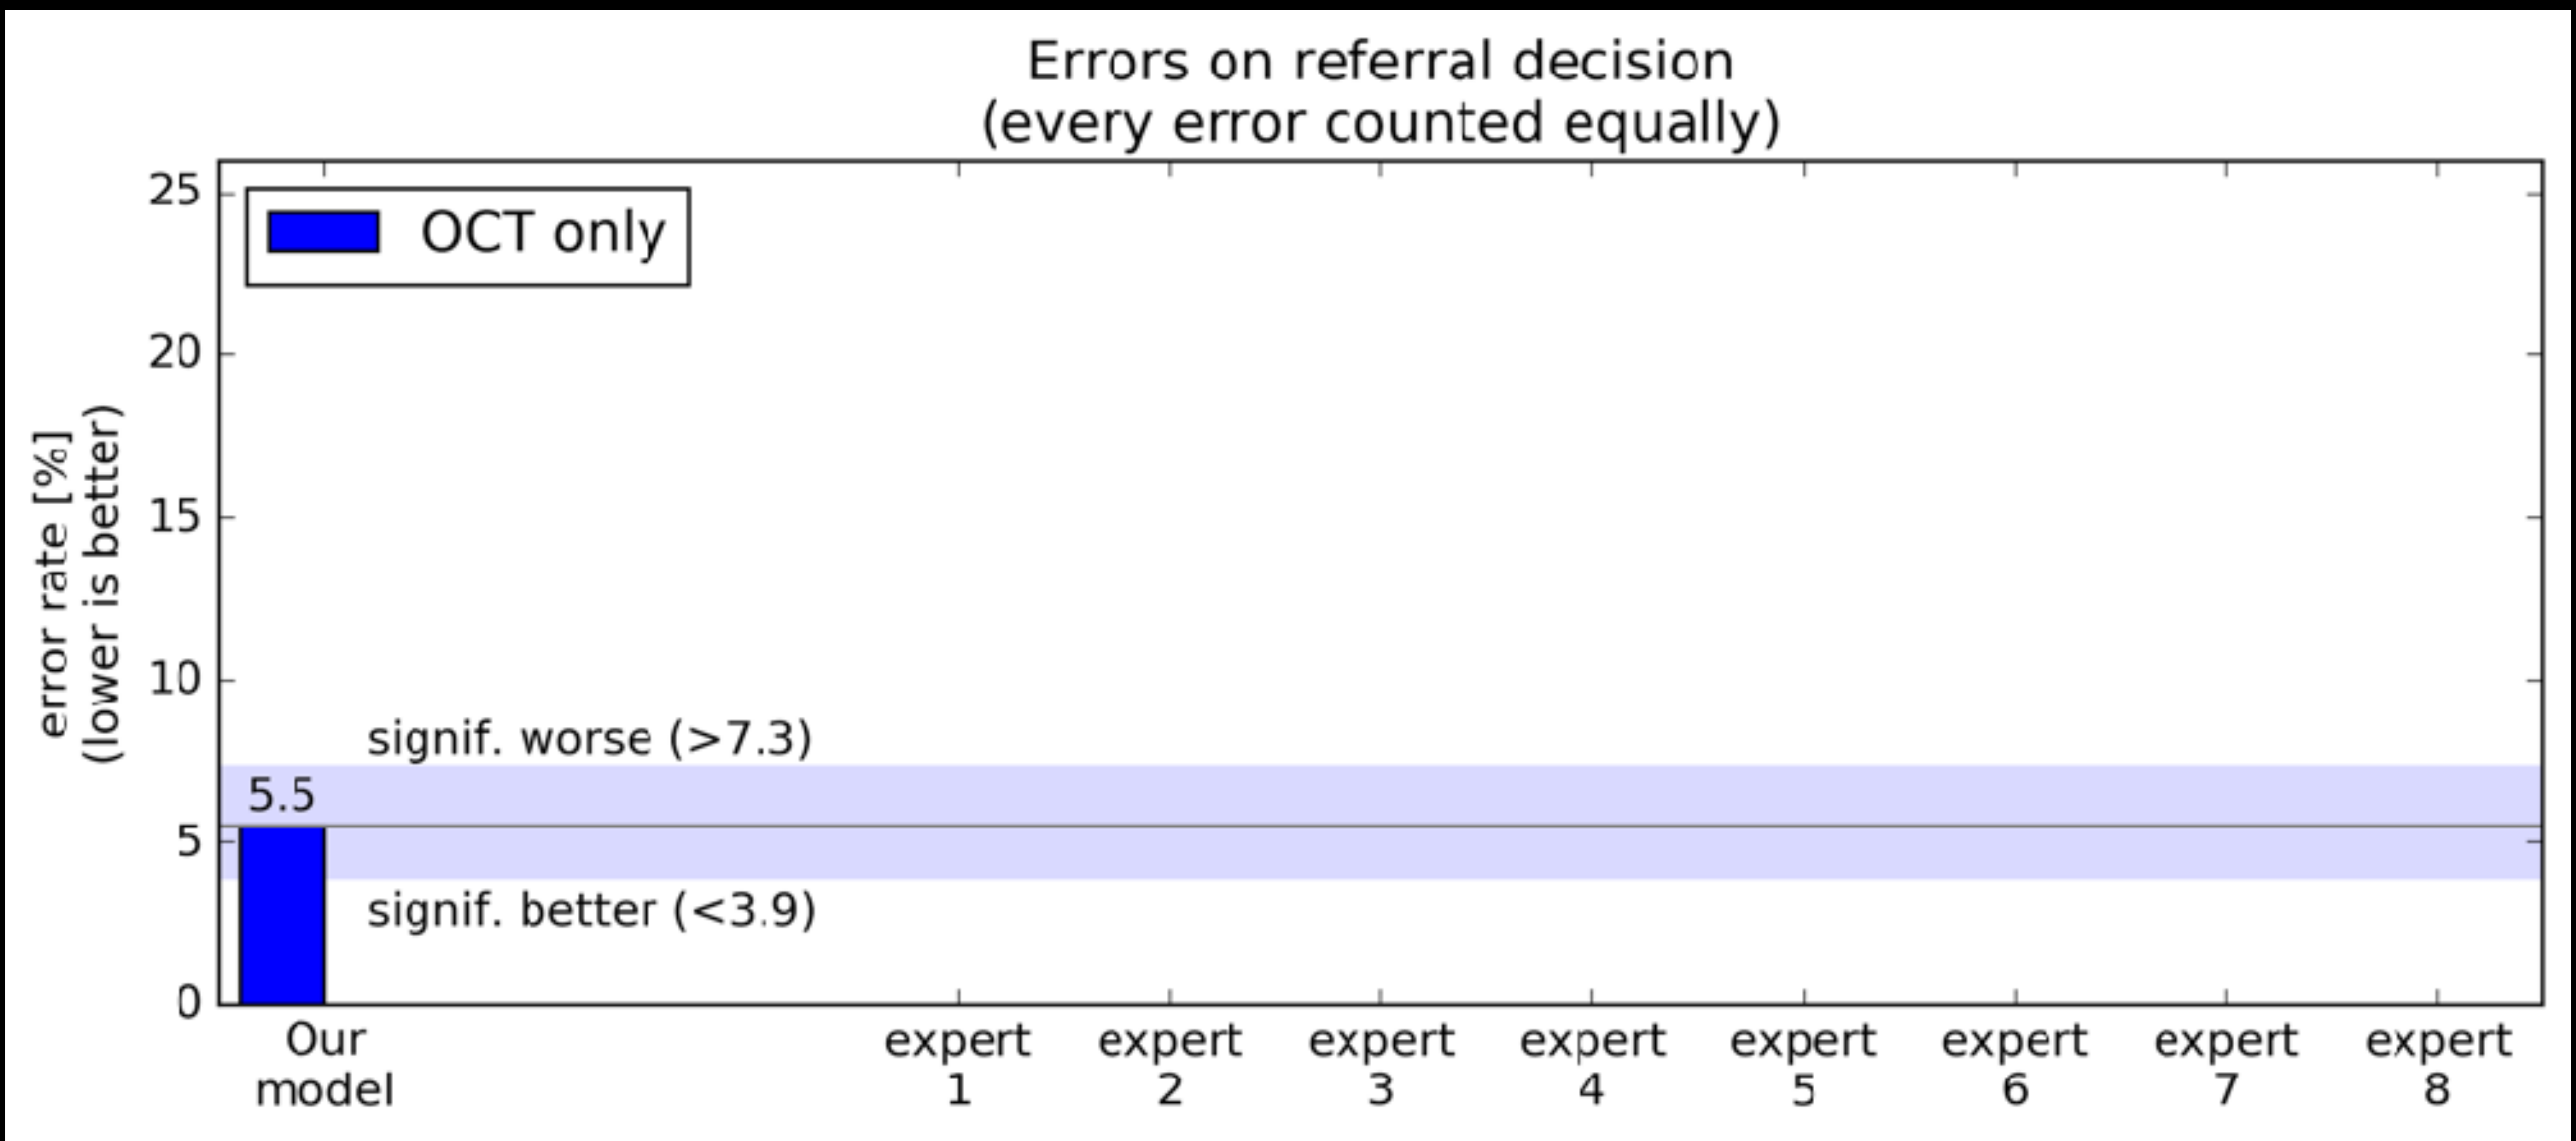

Evaluation of Performance

1000 “New” Patients Moorfields

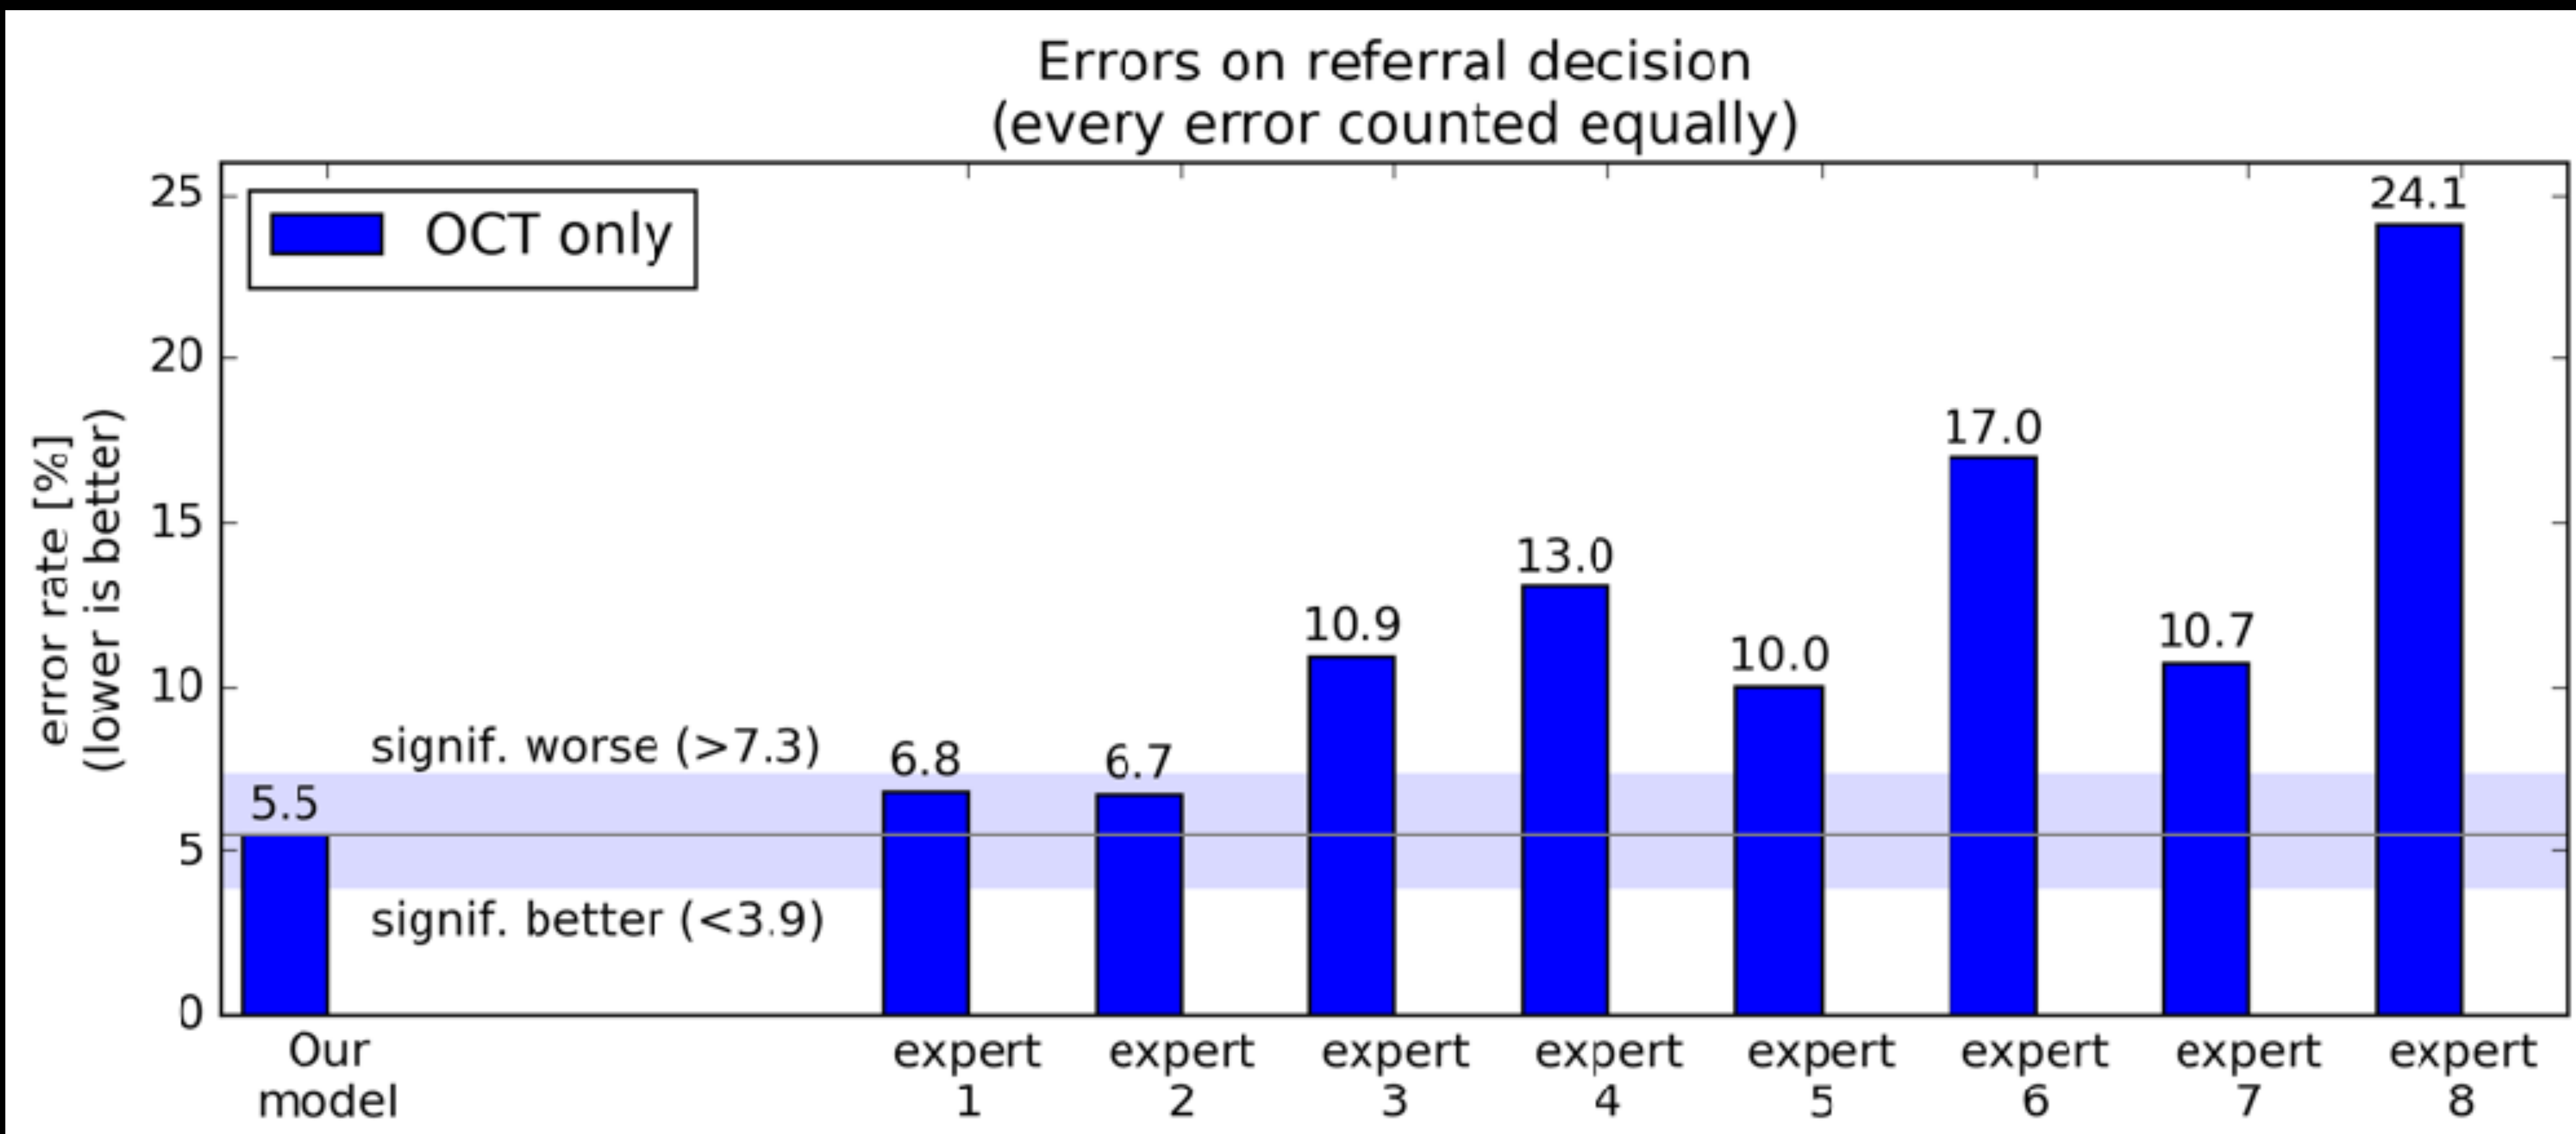

Evaluation of Performance

1000 “New” Patients Moorfields

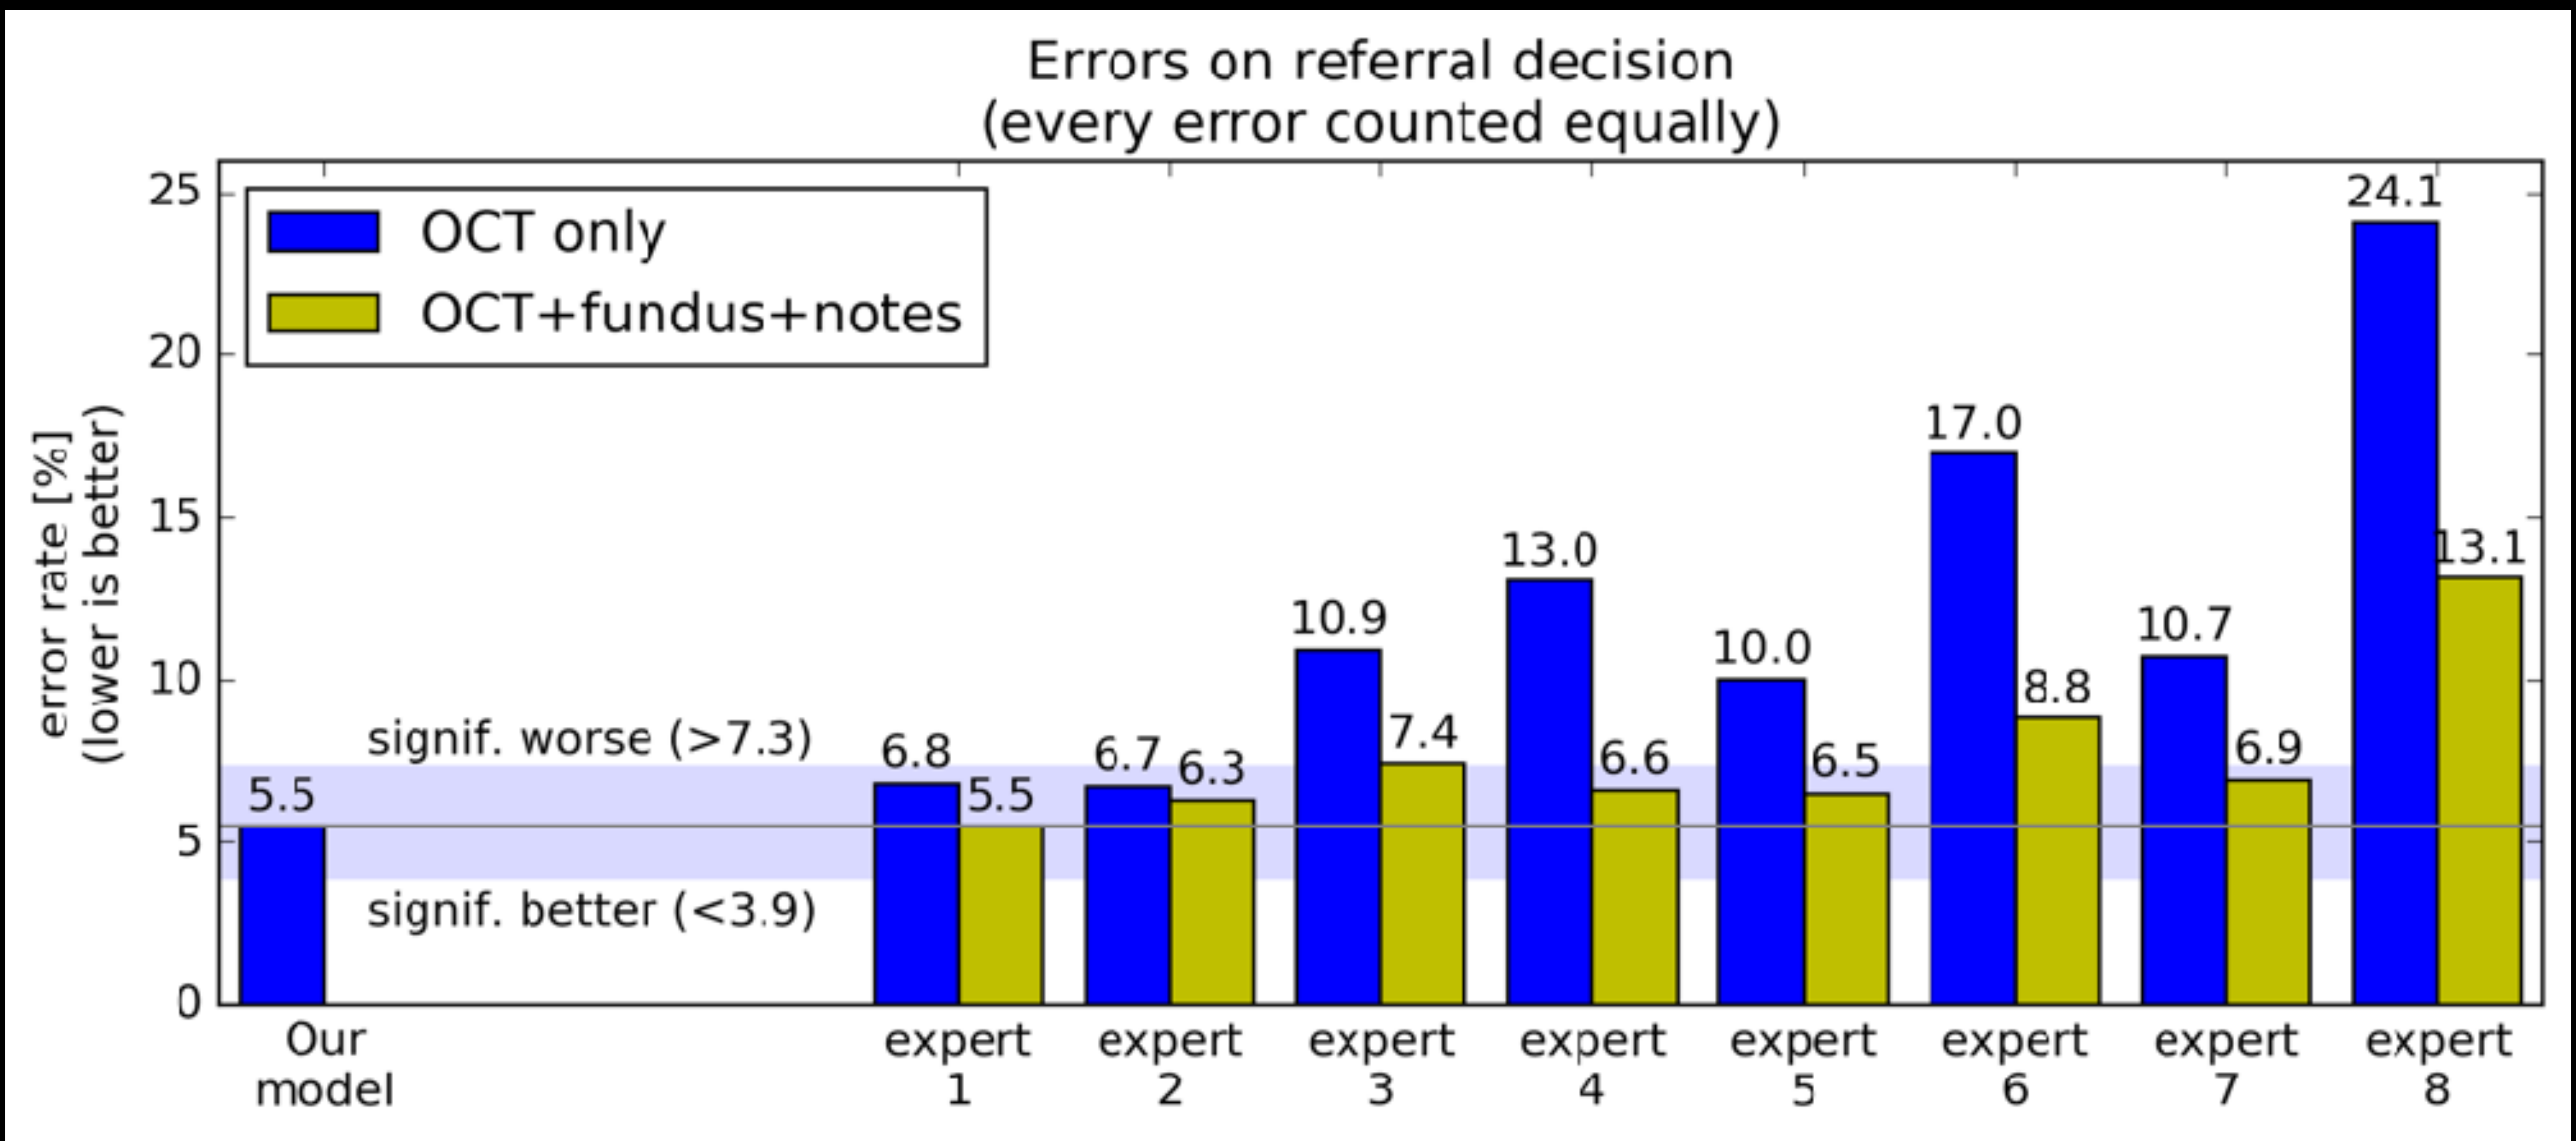

What did it get wrong?

What did it get wrong?

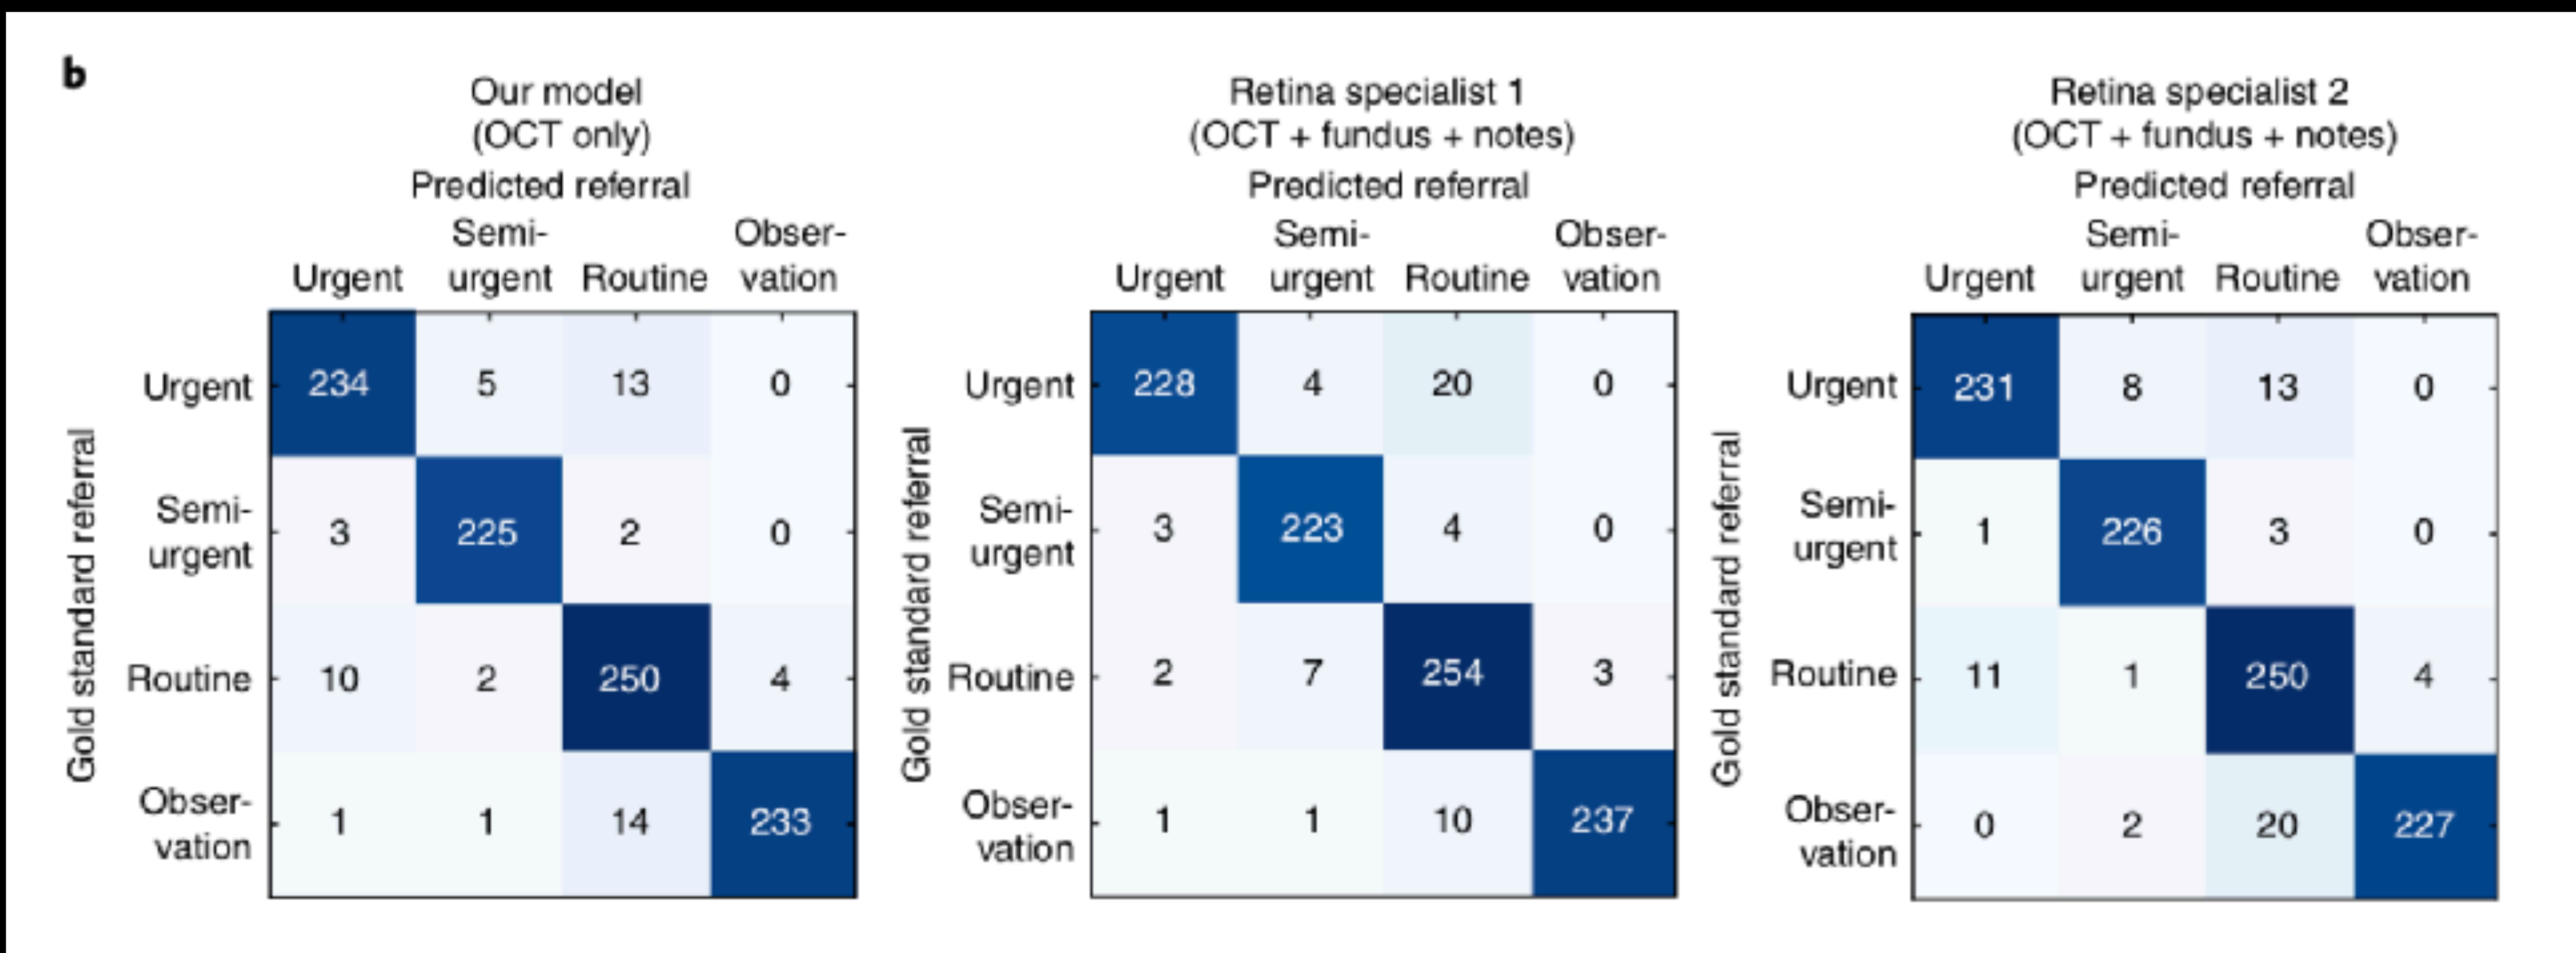

What did it get wrong?

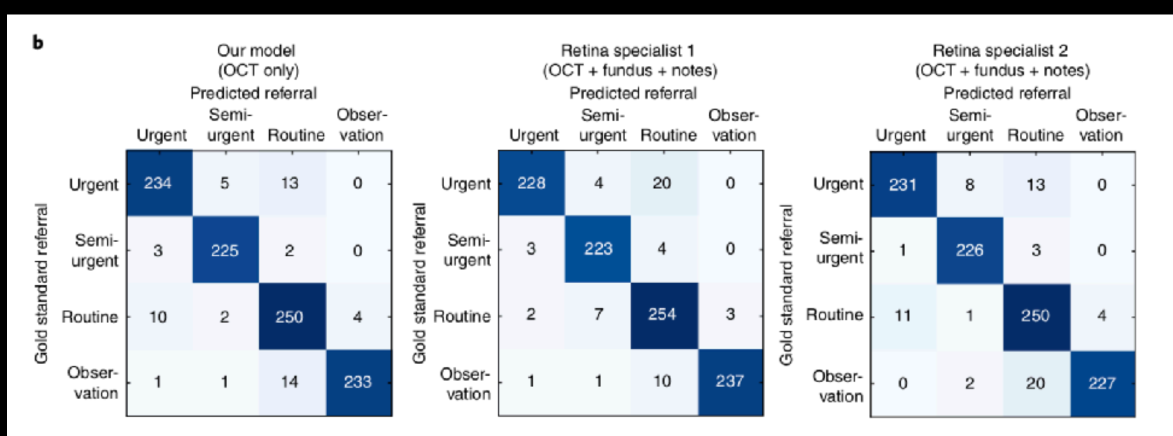

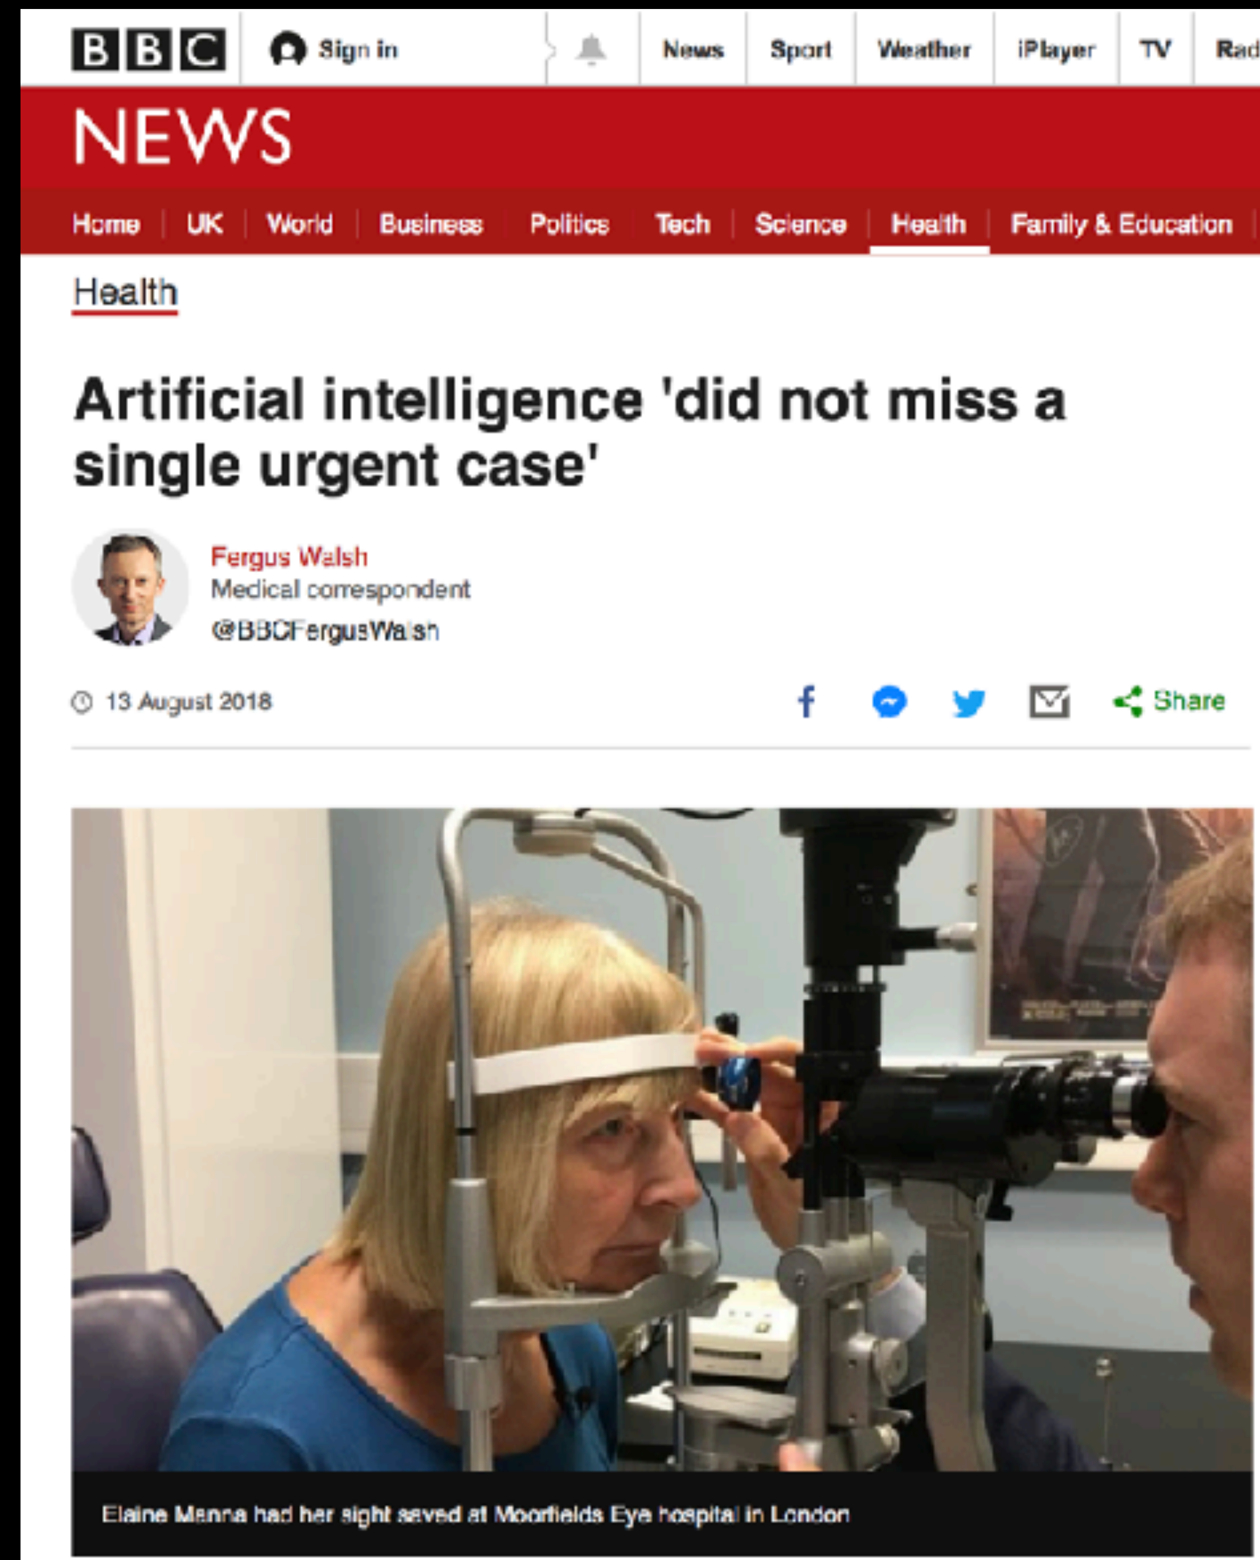
BBC
Sign in
News Sport Weather iPlayer TV Radio
NEWS
Home UK World Business Politics Tech Science Health Family & Education
Health

Artificial intelligence 'did not miss a single urgent case'

Fergus Walsh
Medical correspondent
@BBCFergusWalsh

13 August 2018

Share

Elaine Manna had her sight saved at Moorfields Eye hospital in London

Other OCT Systems?

Other OCT Systems?

Other OCT Systems?

Other OCT Systems?

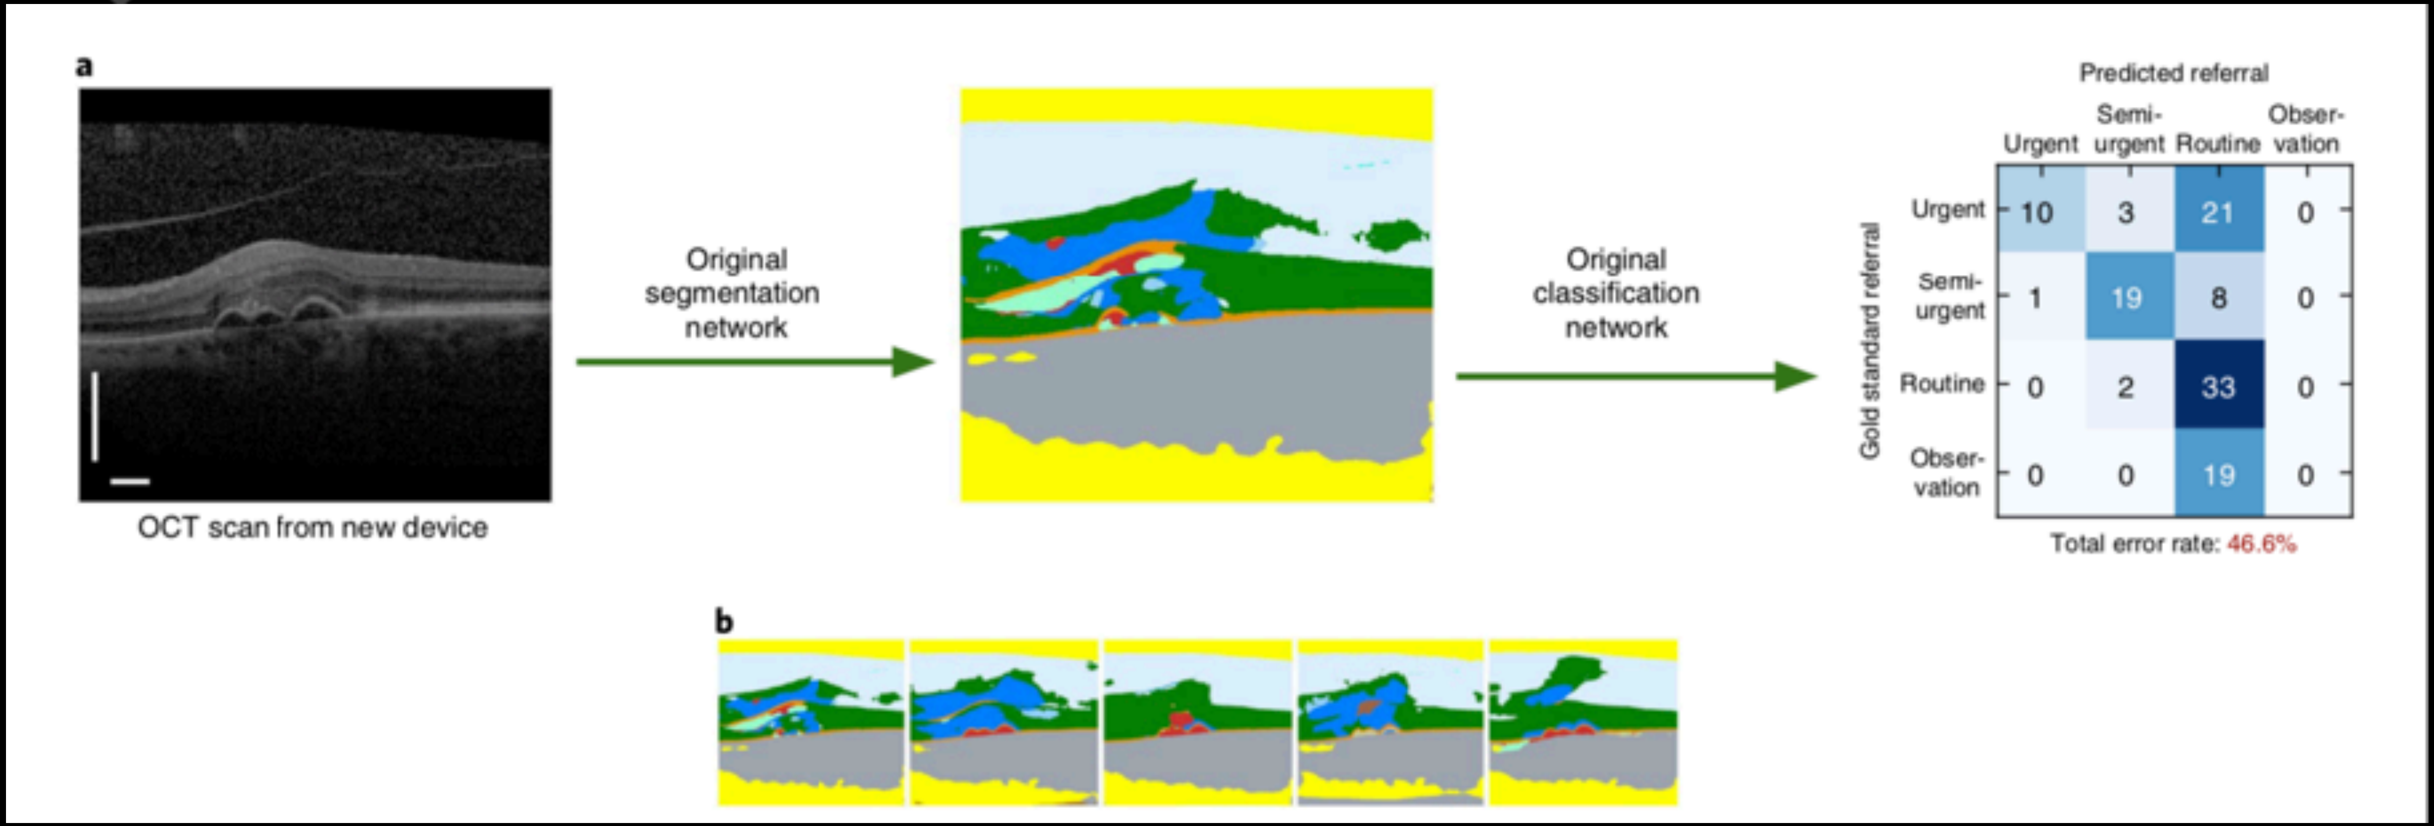

Other OCT Systems?

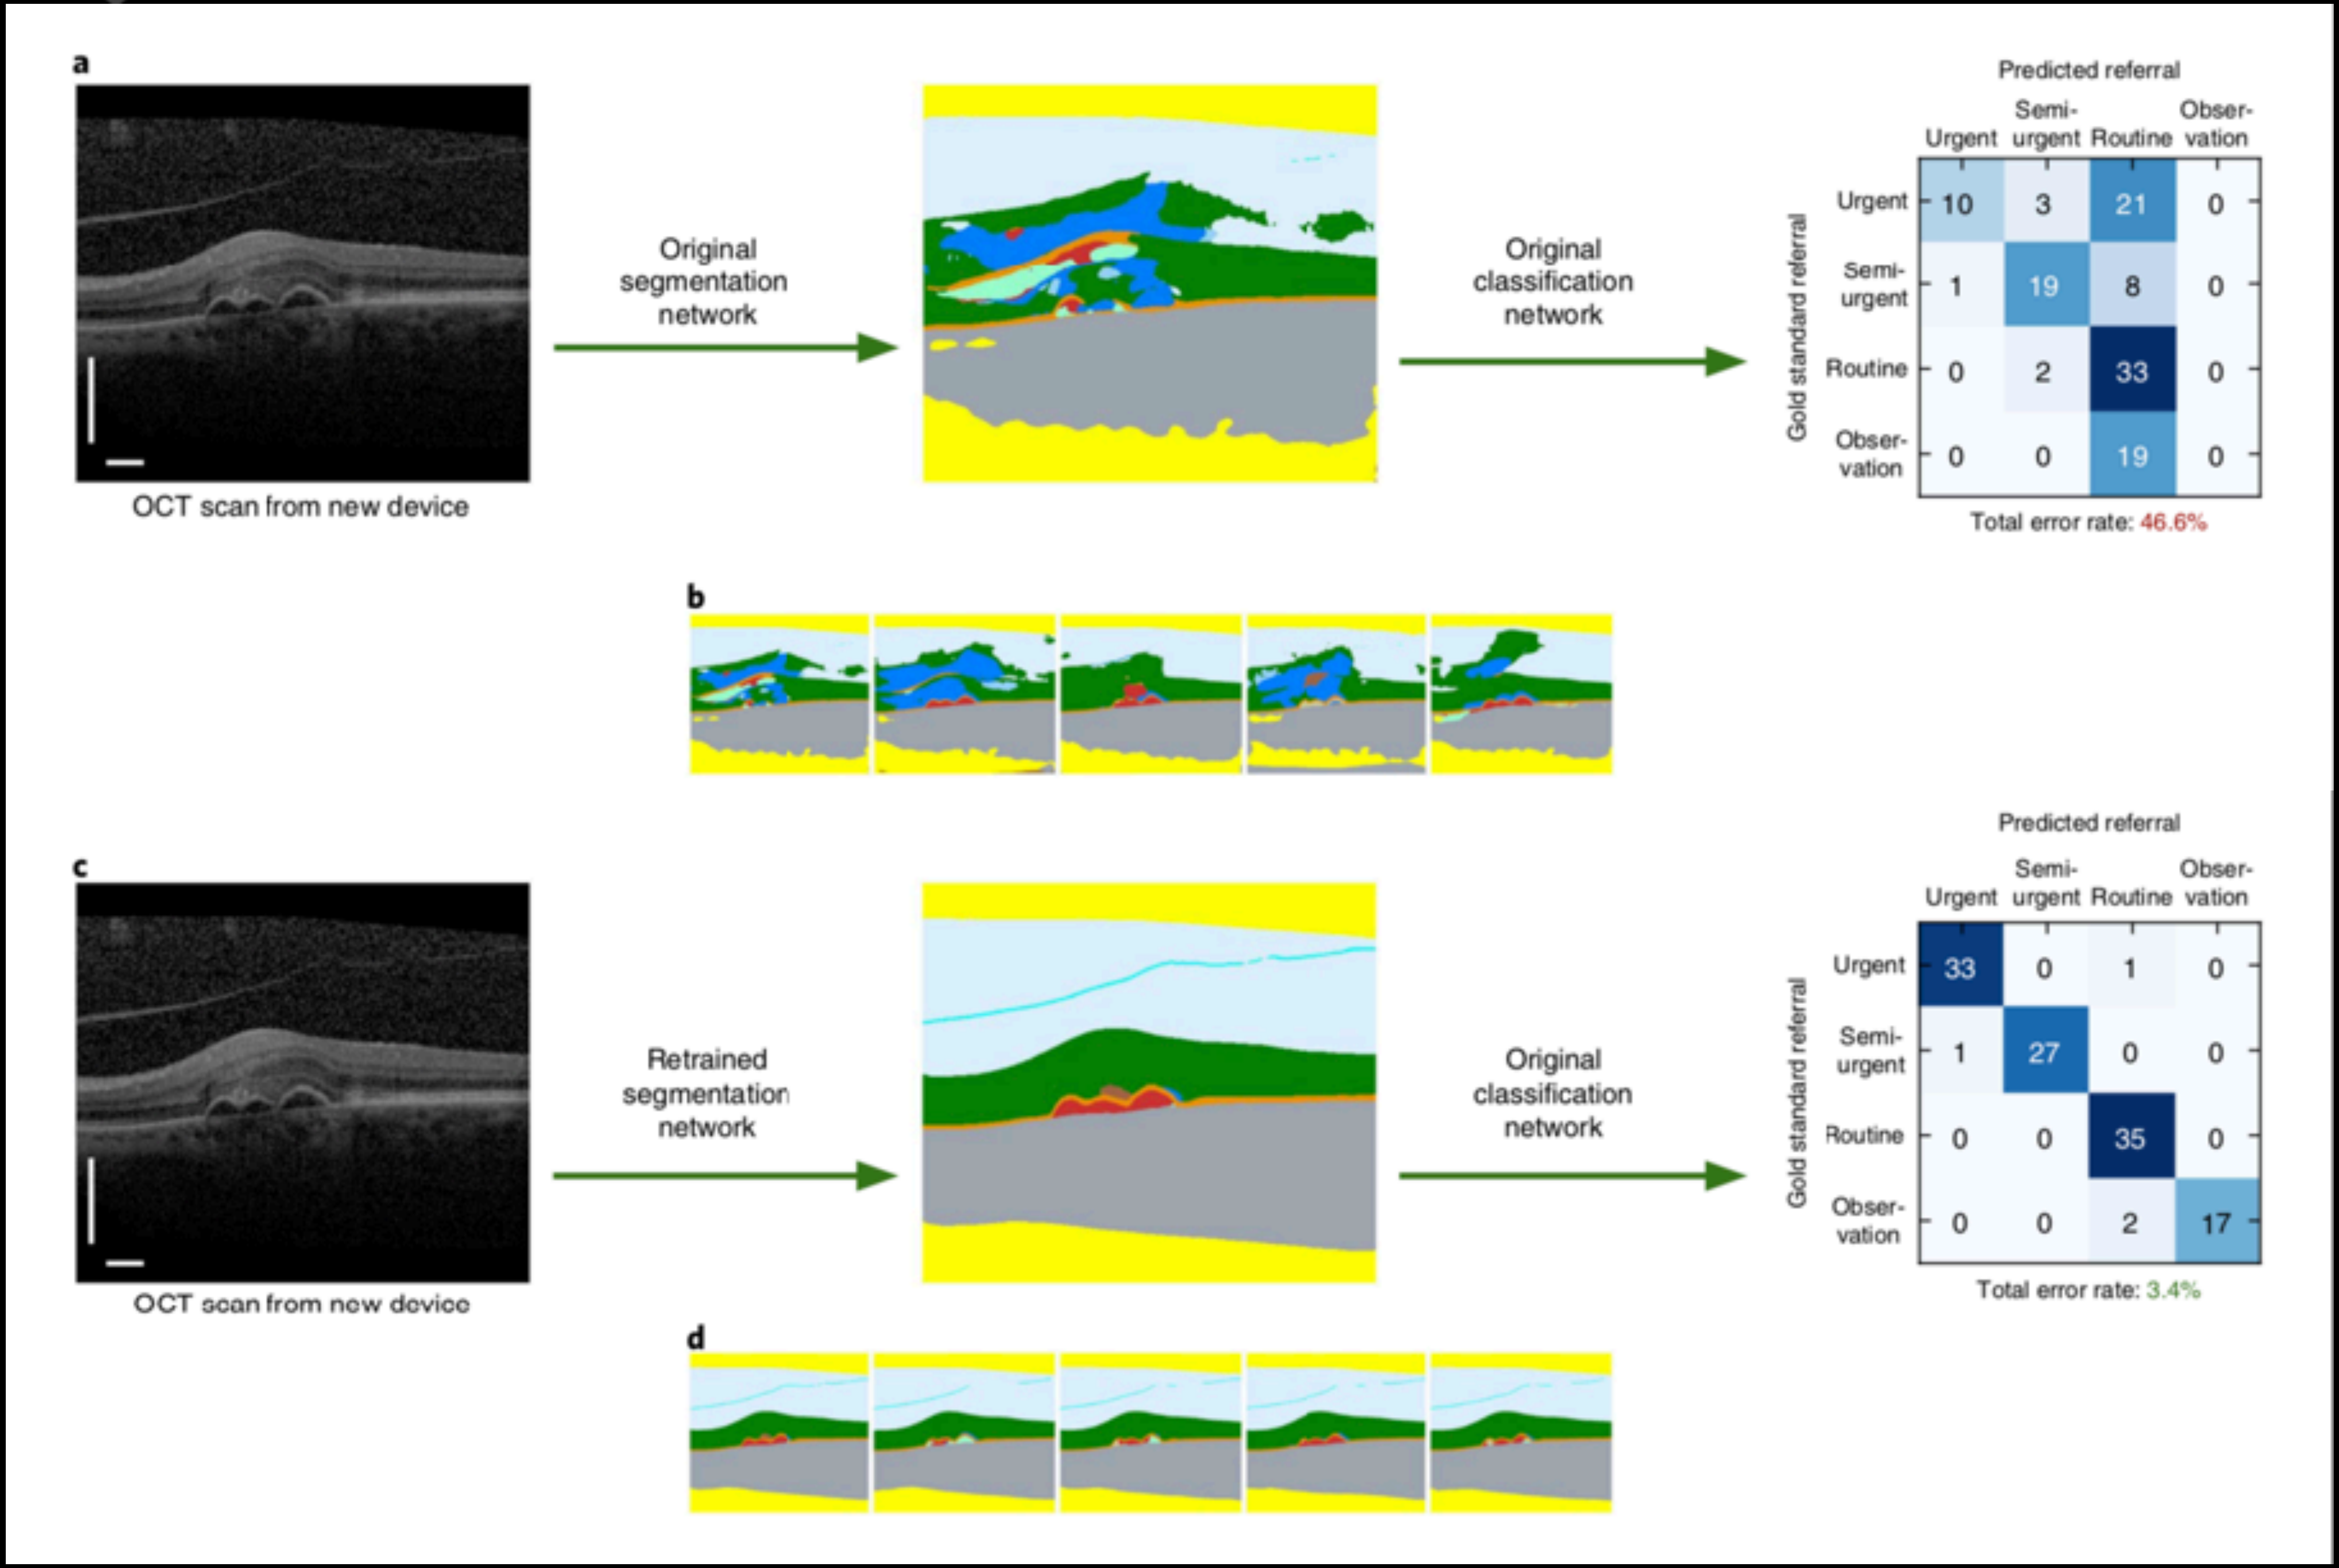

Medical Education?

Medical Education?

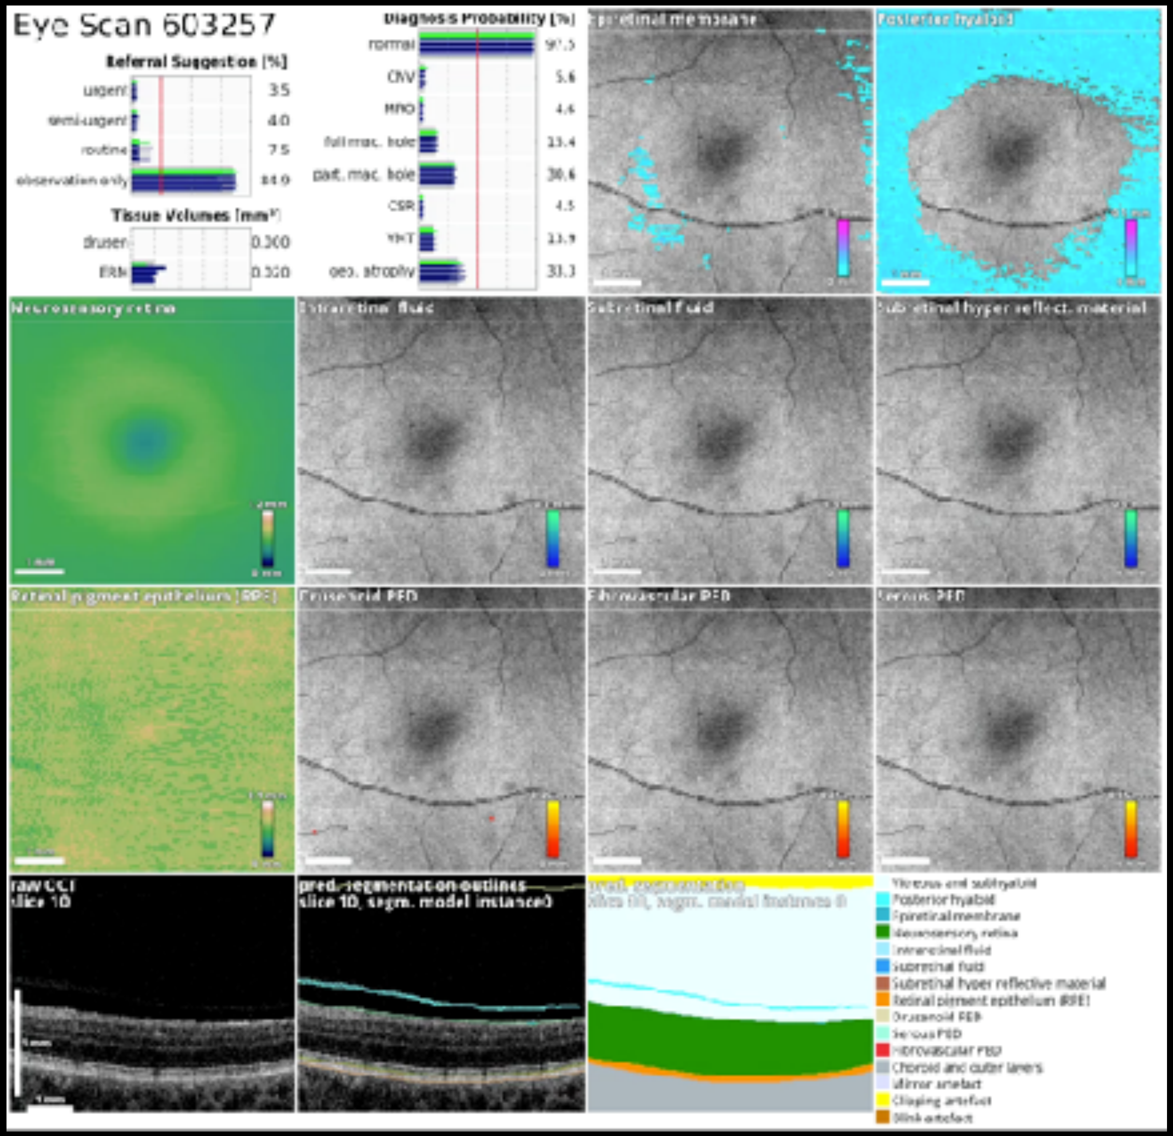

Medical Education?

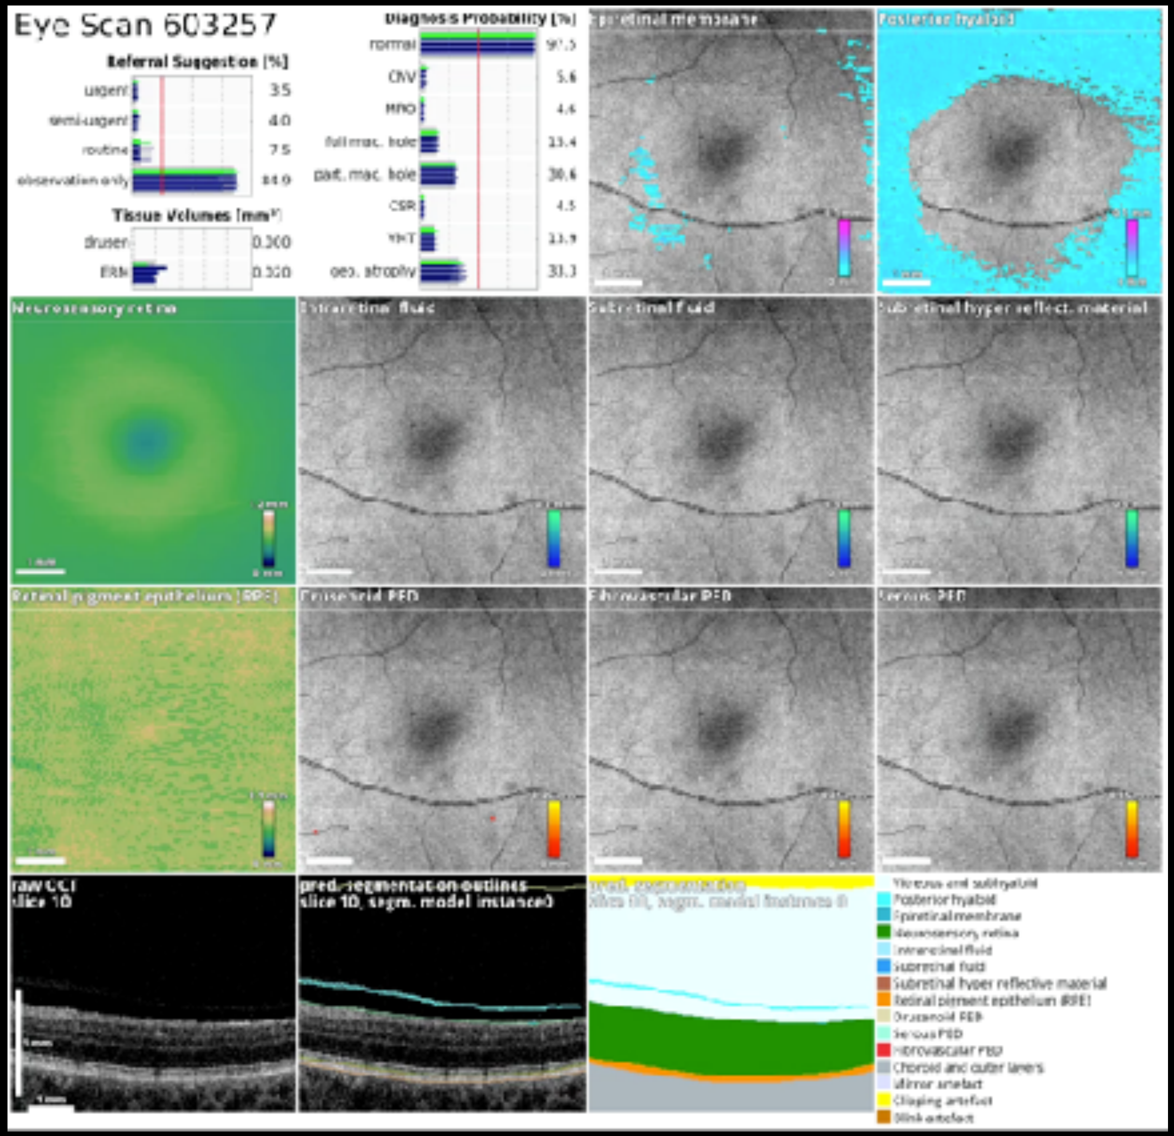

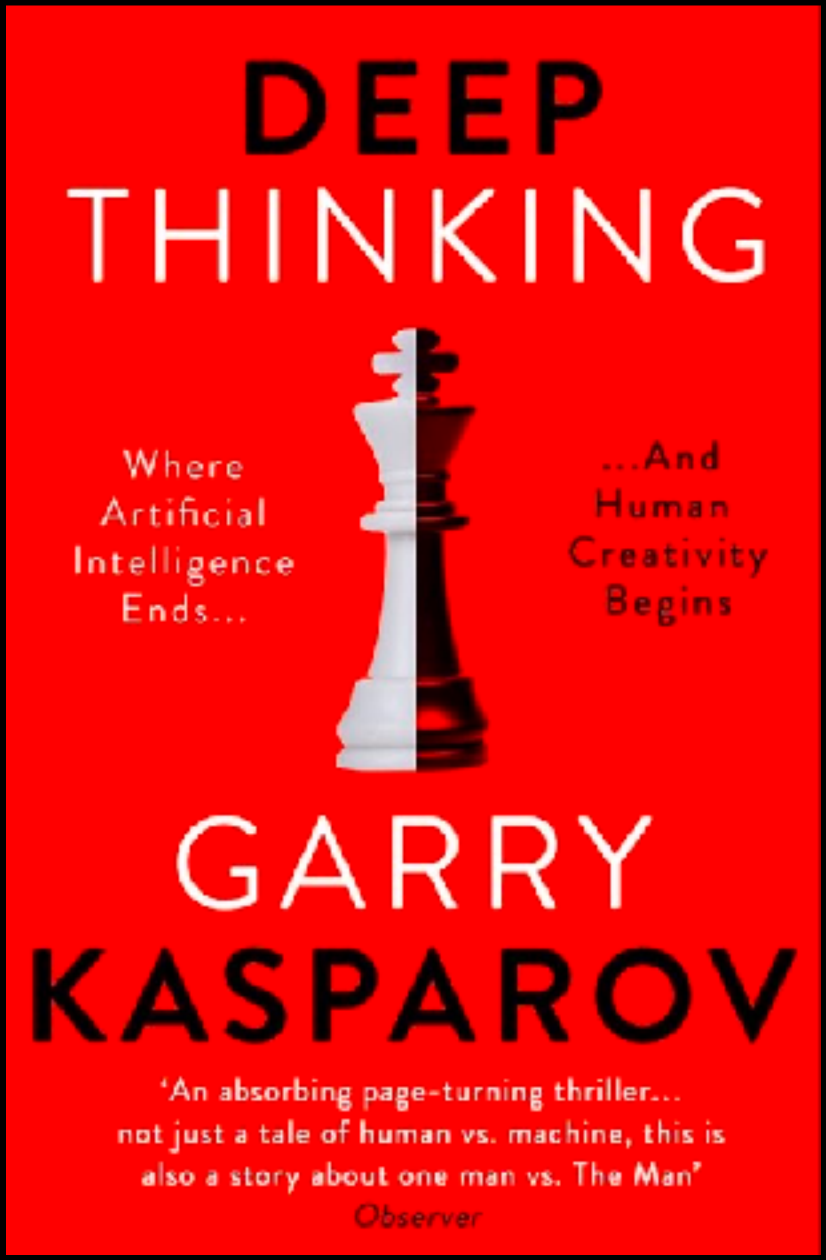

Next Steps?

Next Steps?

1. Prospective clinical trials

Next Steps?

1. Prospective clinical trials

2. AI-assisted Science

Next Steps?

1. Prospective clinical trials

2. AI-assisted Science

3. Reinventing the Eye Exam!

Next Steps?

1. Prospective clinical trials

2. AI-assisted Science

3. Reinventing the Eye Exam!

pearse.keane1@nhs.net