Identifying life-threatening uveal melanoma: A directed application of general-purpose AI

BY MAX JACKSON

Uveal melanoma (UM) is a rare but aggressive eye cancer, affecting approximately six people per million annually [1]. Uveal melanoma arises in three locations: the choroid, ciliary body, and iris. As the most common primary intraocular malignancy in adults, UM poses significant challenges, particularly due to its high metastatic potential, leading to liver metastasis in 50% of cases [2]. The prognosis following metastasis remains poor, with a median survival of less than one year, as such, early detection is critical [3]. Earlier detection of UM is key to better treatment and improved patient prognosis.

For choroidal and ciliary body melanoma, initial detection typically occurs at opticians. Fundus imaging is the primary modality used by opticians to assess eye health and this is where these tumours can be seen. Uveal melanoma are often symptomless until the later stages and remain undetected for long periods of time if the patient does not routinely attend check-ups. In contrast, iris melanomas are easily visible tumours and can easily be detected whilst looking in the mirror or someone else noticing the lesion. This project focused on the more insidious threat by choroidal and ciliary body melanoma, aiming to automate their detection using routine fundus photography.

Figure 1: (a) Ultra-wide fundus image of a choroidal melanoma shown by the bottom of the two black dots. (b) Ultra-wide fundus image of a choroidal freckle shown superotemporal by the darker lesion.

Diagnostic delay

There are several reasons for delays in tumour diagnosis. One of the primary factors, discussed in this article, is that UM shares many of the same clinical features as freckles [4]. For example, Figure 1 presents two fundus images from the Liverpool Ocular Oncology Centre (LOOC): one with choroidal melanoma and the other with a choroidal freckle (naevus). Could you tell the difference?

To the untrained eye, larger or darker lesions may look more sinister but to avoid ambiguity, all lesions are referred to ophthalmologists for diagnosis. Opticians may refer patients to local eye hospitals or directly to regional centres specialising in ocular oncology. Any suspicious lesions detected at the local hospitals are then referred to those regional centres for further investigation and management. In some cases, intraocular biopsies are required for genetic and pathological analysis to allow for a definitive diagnosis. Considerable time is lost along this referral pathway, despite prompt diagnosis and treatment being critical to optimise outcomes. What if there was a way to use opticians' fundus images to predict the likelihood that a lesion is a freckle or UM, allowing opticians to make immediate decisions to refer UM to a specialist centre and lower risk freckles to local eye hospitals? Emerging artificial intelligence (AI) models could be the answer.

Can AI help?

Integration of AI into healthcare has revolutionised various aspects, including early disease detection, treatment planning, and prognostication. However, in the realm of ocular oncology, the application of AI remains underexplored, particularly for rare cancers like UM. Despite this, promising studies have emerged, such as one by Xincheng Yao, et al. which used deep learning (DL) to differentiate between choroidal melanomas and naevi with a mean accuracy of 90% [5]. However, the same underlying problem that occurs with most rare cancers is the sample size. A total number of 798 images were used for training, test and validation of the DL model. Whilst this model shows great proof of concept, the relatively low number of images makes it hard to justify a real-world application of this tool. This is because, with such a low sample size, the variation in what the images look like is very low. A good variety is key to a more well-rounded, truer model as with more complicated or poor-quality cases, a low-variability model could struggle.

It is easy to see that a DL model, like the above, could be used in opticians to increase referral speed for high-risk patients whilst also reducing the number of unnecessary referrals to specialist centres. Clinical application of a model of this sort would require a vast number of images for the DL training and testing alongside validation to prove its capabilities. We propose a model of this calibre. Using the RETFound foundation model as a backbone, a DL model was fine-tuned for the differentiation between choroidal and ciliary body melanomas and freckles.

A new model

We fine-tuned the RETFound foundation model using 28,732 ultra-wide field fundus images sourced from the LOOC, encompassing 20,136 images of choroidal and ciliary body melanomas and 8596 of choroidal freckles. This model achieved an 82% accuracy in correct diagnosis, with a specificity of 88%, correctly identifying UM with 90% accuracy and freckles with 74% accuracy [6]. Whilst the overall accuracy of the model is modest, its ability and robustness to work with lower quality images is a significant strength.

Our model differs from previous models as it uses an already pre-trained foundation model for classification. RETFound has been previously trained on over 900,000 fundus images, giving it a good understanding of what a fundus image looks like. Having then been fine-tuned to 28,732 more images it should perform better on our dataset than a conventional convolutional neural network (CNN). Using the same images, a conventional CNN (ResNet50) was trained for the same classification task using the same images. We saw an overall accuracy of just 69% for this model suggesting that the foundation model resulted in a better classification. This is likely because of the model already understanding fundus imaging and being less distracted by artefacts such as or scarring in the image resulting in a more focused approach. More CCN's should be tested against RETFound for UM but these results are promising for use of a foundation model.

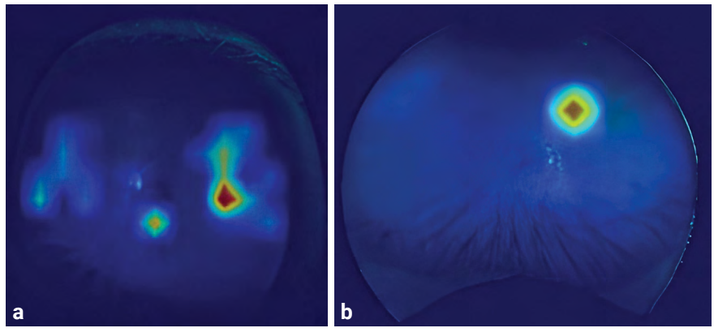

Figure 2: (a) Class activation map showing the models prediction of a tumour being present. (b) Class activation map showing the models prediction of a freckle being present.

What does the model see?

Class activation maps, in this case Grad-CAM [7], allows us to visualise the decision-making process of the model. In Figure 2, for instance, the model focuses on specific regions of the lesion that might be overlooked by the human eye. In Figure 2a, the model identifies a small area within the freckle, leading to an 81% certainty that it is benign. Conversely, Figure 2b shows an 82% certainty of UM, with the model also factoring in darker regions surrounding the fovea – an area that could hold key differentiating features. These maps are one of the many ways of regulating a real-world model with these being output at the time of diagnosis. A manual check by a clinician on these images would also help the model to make correct decisions.

Impact

A model that is robust and accurate having been trained on a vast number and variety of images could have a great impact in ocular oncology. A reduction in the number of unnecessary referrals to regional centres would result in less pressure on clinicians and earlier treatment for patients, improving prognosis. The results shown in this article are preliminary, proof of concept results, which demonstrates the value of foundation models. More work would need to be done to reduce the number of incorrect referrals and externally validate the model before implementation. On top of this, several steps towards safeguarding would be required to protect the patients that fell into the false negative categories. However, it demonstrates the capabilities of AI to aid in rare cancer diagnosis and treatment, even with the limited number of cases. A sample size this large is a first of its kind in UM and shows the great potential of foundation models. Applying these models to a relatively rare cancer can improve the classification of UM and naevi without the need for large datasets. However, when also applied to larger datasets, it still performs exceptionally well, exceeding the accuracy of narrow models such as ResNet. Foundation models are a step towards the application of AI, with better accuracy, to rare diseases.

References

1. Hope-Stone L, Brown SL, Heimann H, Damato B. Comparison between patient-reported outcomes after enucleation and proton beam radiotherapy for uveal melanomas: a 2-year cohort study. *Eye (Lond)* 2019;**33(9)**:1478–84.
2. Carvajal RD, Schwartz GK, Tezel T, et al. Metastatic disease from uveal melanoma: treatment options and future prospects. *Br J Ophthalmol* 2017;**101(1)**:38–44.
3. Nathan P, Hassel JC, Rutkowski P, et al. Overall Survival Benefit with Tebentafusp in Metastatic Uveal Melanoma. *N Engl J Med* 2021;**385(13)**:1196–206.
4. Bindewald-Wittich A, Holz FG, Ach T, et al. Fundus Autofluorescence Imaging in Patients with Choroidal Melanoma. *Cancers (Basel)* 2022;**14(7)**:1809.
5. Yao X, Dadzie A, Iddir S, et al. Color Fusion Effect on Deep Learning Classification of Uveal Melanoma. *Eye (Lond)* 2024;**38(14)**:2781–7.
6. Jackson M, Kalirai H, Hussain RN, et al. Differentiating Choroidal Melanomas and Nevi Using a Self-Supervised Deep Learning Model Applied to Clinical Fundoscopy Images. *Opthamol Sci* 2024;**5(2)**:100647.
7. Gildenblat J, various Contributors. PyTorch library for CAM methods. *Github* (2021). https://github.com/jacobgil/pytorch-grad-cam [Link last accessed February 2025]

TAKE HOME MESSAGES

- Foundation models show great potential to improve classification results especially on smaller datasets.
- Reducing the number of referrals to specialist eye cancer centres overall will result in better patient outcomes.
- Deep learning could improve the referral system by prioritising higher risk cases who need treatment urgently to prevent metastasis.
- Class activation maps can be a good way of trying to 'see' what the model sees and can give us a deeper understanding of what features it attempts to detect.

AUTHOR

Max Jackson,
Liverpool Ocular Oncology Research Group, Institute of Life Course and Medical Sciences (ILCaMS), University of Liverpool; Department of Eye and Vision Science, University of Liverpool, UK.

SECTION EDITORS

Nima Ghadiri,
Medical Ophthalmology Consultant and Honorary Senior Clinical Lecturer, Liverpool, UK.
nima.ghadiri@liverpoolft.nhs.uk

Arun James Thirunavukarasu,
Academic Foundation Doctor, Oxford University Hospitals NHS Foundation Trust; Clinical Research Fellow, Nuffield Department of Clinical Neurosciences & Big Data Institute, University of Oxford; Rising Leader Fellow, Aspen Institute, UK.
ajt205@cantab.ac.uk

Declaration of competing interests: None declared.

Acknowledgements: The author thanks North West Cancer Research for funding his PhD and this project.

Eye News | April/May 2025 | VOL 31 NO 6 | www.eyenews.uk.com