International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	AI-DRIVEN OCULOMICS: RETINAL IMAGING BIOMARKERS FOR SYSTEMIC DISEASE - FROM CARDIOVASCULAR RISK TO NEURODEGENERATION AND FOUNDATION MODELS
DOI	https://doi.org/10.31435/ijitss.3(51).2026.5896
RECEIVED	30 May 2026
ACCEPTED	04 September 2026
PUBLISHED	16 September 2026
LICENSE	The article is licensed under a **Creative Commons Attribution 4.0 International License.**

AI-DRIVEN OCULOMICS: RETINAL IMAGING BIOMARKERS FOR SYSTEMIC DISEASE - FROM CARDIOVASCULAR RISK TO NEURODEGENERATION AND FOUNDATION MODELS

Adrianna Buż *(Corresponding Author, Email: adriannabuz@gmail.com)*
John Paul II Municipal Hospital in Rzeszów, Rycerska 4, 35-241 Rzeszów, Poland
ORCID ID: 0009-0001-8202-5628

Nicol Baran
Medical Center in Łańcut, Ignacego Paderewskiego 5, 37-100 Łańcut, Poland
ORCID ID: 0009-0003-3607-6744

Marta Bajkowska-Piterak
Medical Center in Łańcut, Ignacego Paderewskiego 5, 37-100 Łańcut, Poland
ORCID ID: 0009-0000-6689-0788

Jerzy Buszko
Medical Care Center in Jarosław, 3-go Maja 70, 37-500 Jarosław, Poland
ORCID ID: 0009-0008-2708-1222

Damian Danilczuk
John Paul II Municipal Hospital in Rzeszów, Rycerska 4, 35-241 Rzeszów, Poland
ORCID ID: 0009-0005-5755-1115

Wiktor Adamiec
John Paul II Municipal Hospital in Rzeszów, Rycerska 4, 35-241 Rzeszów, Poland
ORCID ID: 0009-0002-3322-5479

Monika Szlachta-Gubernat
Hospital of the Ministry of Interior and Administration in Rzeszów, Krakowska 16, 35-111 Rzeszów, Poland
ORCID ID: 0009-0005-6994-9761

Przemysław Piterak
Hospital of the Ministry of Interior and Administration in Rzeszów, Krakowska 16, 35-111 Rzeszów, Poland
ORCID ID: 0009-0002-3553-7353

Michał Dyś
Health Care Complex in Dębica, Krakowska 91, 39-200 Dębica, Poland
ORCID ID: 0009-0000-2644-7626

ABSTRACT

Introduction and aim: The retina offers a unique, non-invasive window into systemic vascular and neurological health. Deep-learning algorithms applied to retinal photographs and optical coherence tomography (OCT) can predict cardiovascular events, detect chronic kidney disease, estimate biological age, and identify neurodegenerative diseases - a field termed oculomics. This review synthesizes current evidence on AI-driven oculomics across cardiovascular, cardio-renal-metabolic, neurodegenerative, and aging domains.

Materials and methods: A narrative review was conducted in PubMed, Scopus, and Google Scholar for studies published 2018-2026. Original research, validation studies, and major reviews on AI-based prediction of systemic disease from retinal imaging were included.

Results: Deep-learning models trained on fundus photographs predict major adverse cardiovascular events and ten-year atherosclerotic cardiovascular disease risk with areas under the receiver operating characteristic curve (AUROCs) of 0.70-0.89. Chronic kidney disease detection achieves AUROCs of 0.85-0.94. Alzheimer's disease classifiers reach AUROCs of 0.93 in development cohorts and 0.73-0.91 in external validation, while Parkinson's disease detection achieves AUROCs of 0.77-0.92. Retinal age-gap independently predicts mortality. Self-supervised foundation models (RETFound, VisionFM, EyeFound, RETFound-Green, MIRAGE, V-JEPA) show state-of-the-art performance, with volumetric OCT (V-JEPA) reaching AUROCs of 0.94.

Conclusions: AI-driven oculomics is a rapidly maturing field with strong translational potential. Despite encouraging accuracies, prospective validation, algorithmic bias, regulatory approval, and clinical integration remain key challenges. Future work should prioritize multi-ethnic cohort studies, cost-effectiveness analyses, and regulatory pathways.

KEYWORDS

Oculomics, Retinal Biomarkers, Artificial Intelligence, Deep Learning, Cardiovascular Disease, Foundation Models

CITATION

Buż, A., Baran, N., Bajkowska-Piterak, M., Buszko, J., Danilczuk, D., Adamiec, W., Szlachta-Gubernat, M., Piterak, P., & Dyś, M. (2026). AI-driven oculomics: Retinal imaging biomarkers for systemic disease - from cardiovascular risk to neurodegeneration and foundation models. *International Journal of Innovative Technologies in Social Science*, 3(51). https://doi.org/10.31435/ijitss.3(51).2026.5896

Introduction

The human retina has long been recognized as one of the few sites in the body where the microvasculature can be directly observed and photographed without invasive procedures. Ophthalmologists have, for over a century, used this anatomical privilege to diagnose not only ocular pathologies but also systemic conditions - hypertensive retinopathy and diabetic retinopathy being perhaps the most familiar examples. What has changed in recent years, however, is the scale and sophistication with which retinal images are being analyzed. The convergence of large biobank datasets, advances in convolutional neural networks, and the increasing availability of high-resolution fundus cameras and optical coherence tomography (OCT) devices has given rise to a discipline now commonly referred to as "oculomics" (Wagner et al., 2020).

Wagner et al. (2020) formalized the concept of oculomics in a review, defining it as the use of ophthalmic imaging to extract biomarkers for systemic disease. The underlying rationale is both anatomical and developmental: the retina shares embryological origins with the central nervous system, and its microvasculature mirrors the hemodynamic and metabolic state of organs such as the heart, kidneys, and brain (Wagner et al., 2020; Zhu et al., 2025). What makes the approach particularly attractive from a public health standpoint is that retinal photography is rapid, inexpensive, and widely available - even in primary care settings and, increasingly, via smartphone-based devices.

The landmark study by Poplin et al. (2018) at Google provided the first large-scale demonstration that deep-learning models could predict cardiovascular risk factors - including age, sex, smoking status, blood pressure, and major adverse cardiac events - directly from retinal fundus photographs, achieving performance levels that surprised even the investigators. That work opened a floodgate of research. Over the subsequent years, investigators have developed retinal-based algorithms for coronary artery calcium scoring (Rim et al., 2021), ten-year atherosclerotic cardiovascular disease (ASCVD) risk prediction (Vaghefi et al., 2024), chronic

kidney disease (CKD) detection (Sabanayagam et al., 2020; Zhang et al., 2021), biological age estimation (Nusinovici et al., 2022; Zhu et al., 2023), and the identification of neurodegenerative diseases including Alzheimer's disease (AD) (Cheung et al., 2022) and Parkinson's disease (PD) (Richardson et al., 2024; Tran et al., 2024).

In parallel, the field of retinal artificial intelligence (AI) has undergone a methodological transformation. Early studies relied on supervised deep-learning models trained from scratch on relatively modest, single-center datasets. More recently, self-supervised foundation models - trained on millions of unlabeled retinal images - have emerged as a powerful new paradigm. RETFound, published in Nature by Zhou et al. (2023), demonstrated that a single pre-trained model could be adapted to a wide range of ocular and systemic disease prediction tasks, outperforming task-specific baselines with minimal labeled data. Subsequent models, including VisionFM (Qiu et al., 2024), EyeFound (Shi et al., 2024), RETFound-Green (Engelmann & Bernabeu, 2025), MIRAGE (Morano et al., 2025), and V-JEPA (Judkiewicz et al., 2026), have further advanced the state of the art, particularly for OCT-based applications.

Despite these encouraging developments, the translation of oculomics research into routine clinical practice remains in its early stages. Concerns about algorithmic bias, the lack of prospective validation in diverse populations, uncertain cost-effectiveness, and the absence of clear regulatory pathways continue to temper enthusiasm (Zhu et al., 2025). This review aims to provide a comprehensive yet accessible synthesis of the current evidence, organized around the major clinical domains where AI-driven retinal biomarkers have shown the greatest promise.

Aim of the study

The primary aim of this narrative review is to synthesize and critically appraise the available evidence on AI-driven oculomics published between 2018 and 2026. Specifically, this work seeks to: (1) summarize the performance and clinical relevance of deep-learning models for predicting cardiovascular disease, cardio-renal-metabolic conditions, and neurodegenerative disorders from retinal images; (2) evaluate the emerging role of retinal age-gap estimation as a surrogate marker for biological aging; (3) assess the impact of self-supervised foundation models on the field; and (4) discuss the translational barriers - including validation, bias, regulation, and cost - that must be addressed before oculomics can be integrated into population-level screening programs.

Methodology

A narrative literature review was carried out using PubMed, Scopus, and Google Scholar. The search covered the period from January 2018 through April 2026 and employed the following search terms, used individually and in combination: "oculomics," "retinal biomarker," "deep learning retina," "fundus imaging cardiovascular," "retinal age," "retinal OCT systemic disease," "foundation model retina," "RETFound," "OCTA systemic," "Alzheimer retina AI," and "chronic kidney disease retinal." The reference lists of identified systematic and narrative reviews were additionally hand-searched to identify further relevant publications.

Studies were included if they described AI-based (predominantly deep-learning) models applied to retinal imaging modalities - color fundus photography (CFP), OCT, or OCT angiography (OCTA) - for the detection, prediction, or risk stratification of systemic (non-ocular) diseases. Both original research articles reporting model development or validation and major review articles providing conceptual frameworks were considered. Conference abstracts, preprints that had not subsequently appeared in peer-reviewed journals, and studies focused exclusively on ocular disease without a systemic endpoint were excluded. No formal quality scoring was applied, consistent with the narrative nature of this review; however, particular attention was given to sample size, the availability of external validation cohorts, and the diversity of study populations in assessing the strength of each body of evidence.

Results

Cardiovascular disease prediction from retinal images

The prediction of cardiovascular disease (CVD) from retinal photographs represents the most extensively studied application of oculomics. The foundational work in this area was the study by Poplin et al. (2018), who trained convolutional neural networks on 284,335 patients from the UK Biobank and EyePACS datasets. Their models inferred age with a mean absolute error (MAE) of 3.26 years, predicted sex with an area under the receiver operating characteristic curve (AUROC) of 0.97, and - most strikingly - predicted the occurrence of major adverse cardiovascular events (MACE) with an AUROC of 0.70. While an AUROC of 0.70 may appear modest in isolation, it was achieved from a single retinal photograph without any clinical history, which underscored the richness of cardiovascular information encoded in the retinal vasculature.

Subsequent studies have built on these findings with increasingly refined models and larger validation cohorts. Rim et al. (2021) demonstrated that deep-learning algorithms could predict coronary artery calcium (CAC) scores - an established marker of subclinical atherosclerosis - from retinal photographs at a population scale. This was a notable advance because CAC scoring traditionally requires computed tomography, which involves ionizing radiation and is not available in most primary care settings. The prospect of obtaining a proxy CAC assessment from a simple fundus photograph is therefore highly appealing from a screening perspective.

Tseng et al. (2023) externally validated a deep-learning retinal biomarker designated Reti-CVD in the UK Biobank cohort and showed that it added predictive value beyond the QRISK3 score conventionally used in UK clinical practice. Similarly, Diaz-Pinto et al. (2022) developed a multi-modal model combining UK Biobank fundus images with minimal demographic data, achieving an AUC of 0.80 for myocardial infarction (MI) prediction. Vaghefi et al. (2024) reported a model for ten-year ASCVD risk prediction that achieved an AUROC of 0.89 with sensitivity and specificity values of 84% and 90%, respectively, using both UK Biobank and EyePACS 10K datasets.

A somewhat different approach was taken by Cheung et al. (2021), who developed a fully automated deep-learning system for measuring retinal vessel caliber - a traditional epidemiological risk marker - and confirmed that it could reproduce the established associations between vessel caliber and cardiovascular risk. Chang et al. (2020) provided early evidence that a deep-learning-derived fundoscopic atherosclerosis score stratified cardiovascular mortality independently of conventional risk factors.

Taken together, these studies indicate that retinal AI models can capture cardiovascular risk information that is complementary to, and in some cases additive to, traditional risk assessment tools. However, it bears noting that the majority of this evidence has been generated retrospectively from biobank data - predominantly the UK Biobank - and prospective, interventional trials demonstrating that retinal-based screening alters patient outcomes remain absent as of 2026.

Cardio-renal-metabolic disease and chronic kidney disease

The interconnection between the retinal microvasculature and renal function has a sound physiological basis: both the retina and the kidney possess end-organ microvasculatures that are exposed to similar hemodynamic and metabolic insults in conditions such as hypertension and diabetes (Farrah et al., 2020). This shared vulnerability has motivated a growing body of work on retinal AI for chronic kidney disease (CKD) and broader cardio-renal-metabolic (CRM) risk prediction.

Zhang et al. (2021) used a dataset of 115,344 fundus images from 57,672 patients to develop deep-learning models for the simultaneous detection and incidence prediction of CKD and type 2 diabetes (T2D), reporting AUROCs in the range of 0.85 to 0.93. Sabanayagam et al. (2020) independently developed a deep-learning algorithm for CKD detection from retinal photographs in community-based populations, with development in the Singapore Epidemiology of Eye Diseases (SEED) cohort and external validation in the SP2 and Beijing Eye Study cohorts. This multi-cohort validation added credibility to the generalizability of the approach.

More recently, Joo et al. (2023) introduced Reti-CKD, a retinal imaging-derived score that was externally validated and shown to predict incident kidney failure independently of the estimated glomerular filtration rate (eGFR). Wu et al. (2025) went a step further with their KIDS system, which not only screens for CKD but also attempts to predict pathological classification non-invasively using retinal images supplemented with clinical data. Bhak et al. (2025) explored a multimodal strategy combining retinal images with urine dipstick data; in their multicenter analysis, the multimodal eGFR-MMDL model achieved high diagnostic

performance, although in external validation a simpler urine-dipstick-only model (eGFR-UDLR) performed comparably, indicating that the added value of retinal imaging may depend on context. A summary of the principal AI-driven oculomics studies for cardiovascular and renal disease prediction is presented in Table 1.

In the broader metabolic domain, Zhang et al. (2020) demonstrated the joint prediction of hypertension, hyperglycemia, and dyslipidemia from fundus photographs, while Yip et al. (2020) provided an important methodological contribution by cataloguing the technical and imaging factors that influence deep-learning performance in diabetic retinal AI applications - findings that are directly relevant to all CRM oculomics pipelines.

Collectively, these results suggest that retinal imaging holds considerable promise as an adjunctive screening modality for renal and metabolic disease, particularly in settings where traditional laboratory-based screening (e.g., serum creatinine, urine albumin) is not routinely available. The multimodal approaches combining retinal data with simple clinical measurements appear especially compelling, though independent replication across geographically and ethnically diverse cohorts remains a priority.

Table 1. Selected AI-driven oculomics studies for cardiovascular and renal disease prediction.

Study (Year)	Cohort	Sample size	Approach / outcome	Performance
Poplin (2018)	UK Biobank, EyePACS	284,335 patients	CNN; cardiovascular risk factors and MACE	MACE AUROC 0.70; age MAE 3.26 y
Chang (2020)	Korean retrospective cohort	~32,000	DL fundoscopic atherosclerosis score; CV mortality	Independent stratifier of CV mortality
Cheung (2021)	Multicohort (Asian + European)	~71,000 images	Automated DL retinal-vessel calibre measurement	Reproduces classical caliber-CVD associations
Rim (2021)	Multicohort (Korean, Singaporean, UK)	~216,000	DL prediction of coronary artery calcium score	AUROC ~0.83 for elevated CAC
Diaz-Pinto (2022)	UK Biobank	~5,663	Multimodal CNN (fundus + minimal demographics)	MI prediction AUC 0.80
Tseng (2023)	UK Biobank	~71,000	Reti-CVD score added to QRISK3	Improved CVD risk stratification beyond QRISK3
Vaghefi (2024)	UK Biobank, EyePACS 10K	~110,000	DL 10-year ASCVD risk	AUROC 0.89; sens 84%, spec 90%
Sabanayagam (2020)	SEED, SP2, Beijing Eye Study	~58,000 images	DL CKD detection from fundus photographs	AUROC 0.85-0.93 across cohorts
Zhang (2021)	Chinese multicentre cohort	115,344 images / 57,672 pts	DL detection and incidence prediction of CKD and T2D	AUROC 0.85-0.93
Joo (2023)	Korean cohort + external validation	~120,000	Reti-CKD score for incident kidney failure	Predicts ESKD independently of eGFR
Wu (2025)	Multicentre Chinese cohorts	>100,000 images	KIDS - CKD screening + pathological subtype identification	Multitask DL combining retina + clinical data
Bhak (2025)	Korean multicentre	65,082 (dev) + 58,284 (ext)	eGFR-MMDL - retina + urine dipstick multimodal DL	High internal AUROC; ext. validation comparable to UDLR alone

Abbreviations: ASCVD - atherosclerotic cardiovascular disease; AUROC - area under the receiver operating characteristic curve; CAC - coronary artery calcium; CKD - chronic kidney disease; CNN - convolutional neural network; CVD - cardiovascular disease; DL - deep learning; ESKD - end-stage kidney disease; MACE - major adverse cardiovascular events; MAE - mean absolute error; MI - myocardial infarction; SEED - Singapore Epidemiology of Eye Diseases.

Neurodegenerative disease detection

The retina and the brain share a common embryological origin in the neural ectoderm, and the retinal nerve fiber layer (RNFL) is, in essence, an unmyelinated extension of the central nervous system. This anatomical continuity has long prompted speculation that retinal changes might mirror neurodegenerative processes occurring in the brain, and recent AI-driven studies have provided increasingly strong evidence in support of this hypothesis.

For Alzheimer's disease (AD), the most influential study to date is the multinational, retrospective case-control analysis by Cheung et al. (2022), in which a deep-learning model trained on bilateral fundus images achieved an AUROC of 0.93 in the development cohort, with AUROCs ranging from 0.73 to 0.91 across five external testing datasets, including three with amyloid-PET confirmation. The authors used unsupervised domain adaptation with batch normalization to address data discrepancies across centers, which substantially improved generalizability. Wagner et al. (2022) described the AlzEye dataset - a longitudinal record-level linkage of ophthalmic imaging and hospital admissions for 353,157 patients in London - which has since become a foundational resource for retinal-AI dementia research.

Several additional studies have explored retinal markers for various stages of cognitive decline. Tian et al. (2021) demonstrated that quantified retinal vasculature features alone could stratify AD risk using modular machine learning. Hao et al. (2024) achieved particularly strong results with Eye-AD, a graph-representation deep-learning model applied to 5,751 OCTA images from 1,671 participants, reporting AUROCs of 0.94 (internal validation) and 0.90 (external validation) for early-onset AD detection, and 0.86/0.80 for mild cognitive impairment (MCI). Cheung et al. (2021) provided an authoritative review of OCT and OCTA findings in AD - including RNFL and ganglion cell-inner plexiform layer (GC-IPL) thinning, and parafoveal vessel density loss - that anchor most subsequent AI models.

For Parkinson's disease (PD), the evidence is somewhat more limited but nonetheless encouraging. Richardson et al. (2024) developed a multimodal model combining ultra-wide-field (UWF) imaging and OCTA that achieved AUROCs of 0.86 to 0.92 for PD detection on an independent test set. Tran et al. (2024) used UK Biobank fundus images to predict both prevalent and incident PD, reporting an AUROC of 0.77 and 68% accuracy, and notably included explainability heatmaps to shed light on the retinal regions most informative for PD classification. Selected studies of retinal AI for neurodegenerative disease detection are summarized in Table 2.

Beyond AD and PD, Sim et al. (2025) introduced a deep-learning retinal aging biomarker as a predictor of cognitive decline and incident dementia in a Singapore memory-clinic cohort with Caucasian replication, showing that a retinal aging biomarker can capture neurocognitive risk across ethnic groups.

While these findings are undeniably exciting, several caveats deserve emphasis. Sample sizes in many AD and PD studies remain small by the standards of cardiovascular oculomics, and the reliance on clinical rather than neuropathological gold standards introduces diagnostic uncertainty. The substantial overlap between cerebrovascular and neurodegenerative pathology further complicates interpretation, as shared risk factors may drive the observed retinal-brain associations. Nonetheless, the accessibility and low cost of retinal imaging make it a compelling candidate for early, population-level neurodegenerative screening - a prospect that the next generation of prospective studies is expected to test.

Table 2. Selected studies of retinal artificial intelligence for neurodegenerative disease detection.

Study (Year)	Cohort / dataset	Approach	Disease & key findings	Performance
Tian (2021)	Quantitative retinal vasculature features	Modular ML pipeline	Alzheimer's disease classification	Accuracy 82.4%
Cheung (2022)	Multicentre Asian + European case-control	DL fundus + unsupervised domain adaptation	Alzheimer's disease (incl. amyloid-PET cohorts)	AUROC 0.93 dev / 0.73-0.91 ext.
Wagner (2022)	AlzEye, London (record-linkage); n = 353,157	Longitudinal ophthalmic + hospital data resource	AD/dementia - foundational dataset	- (resource paper)
Hao (2024)	OCTA cohort	Eye-AD graph-representation DL	Early-onset AD and MCI detection	EOAD 0.94/0.90; MCI 0.86/0.80

Study (Year)	Cohort / dataset	Approach	Disease & key findings	Performance
Sim (2025)	Singapore + Caucasian replication	Retinal aging biomarker	Cognitive decline and incident dementia	Cross-ethnic neurocognitive risk stratification
Richardson (2024)	Multimodal UWF + OCTA cohort	Multimodal CNN	Parkinson's disease detection	AUROC 0.86-0.92 (test set)
Tran (2024)	UK Biobank	DL fundus + explainability heatmaps	Prevalent and incident PD	AUROC 0.77; accuracy 68%

Abbreviations: AD - Alzheimer's disease; AUROC - area under the receiver operating characteristic curve; DL - deep learning; EOAD - early-onset Alzheimer's disease; MCI - mild cognitive impairment; ML - machine learning; OCTA - optical coherence tomography angiography; PD - Parkinson's disease; UWF - ultra-wide-field.

Retinal age estimation and biological aging

The concept of biological age - as distinct from chronological age - has gained traction across multiple domains of medicine, and the retina has emerged as a promising organ from which to derive it. The premise is straightforward: a deep-learning model is trained to predict chronological age from retinal photographs of healthy individuals, and the difference between the predicted and actual age (the "retinal age gap") is then used as a proxy for accelerated or decelerated biological aging.

The seminal study in this area was published by Zhu et al. (2023), who trained a deep-learning model on 19,200 fundus images from healthy UK Biobank participants and found that each additional year of retinal age gap was associated with a 2% increase in the risk of all-cause mortality. Nusinovici et al. (2022) introduced RetiAGE, showing that participants in the highest quartile of the retinal age-gap distribution had 67% higher all-cause mortality, 142% higher CVD mortality, and 60% higher cancer mortality compared with those in the lowest quartile - associations that remained significant after adjustment for chronological age.

The same research group subsequently developed RetiPhenoAge (Nusinovici et al., 2024), a second-generation model trained on PhenoAge ground truth rather than chronological age, which outperformed RetiAGE in mortality stratification. Yu et al. (2025) extended this work across populations, employing self-supervised longitudinal pre-training and label-distribution learning across the UK Biobank and three Chinese cohorts (n = 34,433), achieving an MAE of 2.79 years and demonstrating that the retinal age gap predicted both all-cause mortality and a range of age-related diseases. The principal retinal-age and biological aging models discussed above are summarized in Table 3.

These retinal age models are conceptually appealing because they offer a single, easily obtained summary metric of systemic aging. However, it is important to recognize that several of these models originate from overlapping research groups and datasets, and truly independent replication - particularly in low- and middle-income countries - is still limited. Whether interventions that reduce the retinal age gap (e.g., blood pressure control, smoking cessation) translate into improved clinical outcomes has not yet been established.

Table 3. Retinal-age and biological aging models predicting morbidity and mortality.

Model (Year)	Authors	Cohort / sample	Key outcome	Headline result
Retinal age gap (2023)	Zhu et al.	UK Biobank - 19,200 healthy participants	All-cause mortality vs retinal-chronological age gap	+2% mortality / year
RetiAGE (2022)	Nusinovici et al.	UK Biobank + SEED	All-cause, CVD and cancer mortality (highest vs lowest quartile)	+67% / +142% / +60%
RetiPhenoAge (2024)	Nusinovici et al.	Cohort development + external validation	Morbidity and mortality prediction (PhenoAge ground truth)	Outperforms RetiAGE
Cross-population aging (2025)	Yu et al.	UK Biobank + 3 Chinese cohorts (n = 34,433)	Self-supervised longitudinal pre-training; age-related disease risk	MAE 2.79 y; predicts mortality and age-related disease across populations

Abbreviations: CVD - cardiovascular disease; MAE - mean absolute error; PhenoAge - composite phenotypic age based on clinical biomarkers; SEED - Singapore Epidemiology of Eye Diseases.

Foundation models for retinal imaging

The advent of self-supervised foundation models has arguably been the single most transformative methodological development in retinal AI since the initial applications of deep learning to fundus photography. Unlike conventional supervised models, which require large volumes of labeled training data, foundation models are pre-trained on vast quantities of unlabeled images using self-supervised objectives and can then be fine-tuned for specific downstream tasks with comparatively little labeled data.

RETFound, developed by Zhou et al. (2023), was the first dedicated retinal foundation model. It was trained on 1.6 million unlabeled color fundus photographs (CFPs) and OCT images using a masked autoencoder framework. Upon fine-tuning, RETFound outperformed task-specific baselines for a range of ocular disease detection tasks and, crucially, for three-year incident systemic disease prediction - including heart failure, myocardial infarction, and Parkinson's disease. This result demonstrated that a single, generalizable representation learned from unlabeled retinal data could capture both ocular and systemic pathophysiology.

Several competing and complementary models have since appeared. VisionFM (Qiu et al., 2024) was trained on 3.4 million images spanning eight imaging modalities and outperformed RETFound on several systemic disease prediction tasks. EyeFound (Shi et al., 2024), pre-trained on 2.78 million images across eleven modalities, offered greater flexibility in modality input. RETFound-Green (Engelmann & Bernabeu, 2025) achieved comparable performance to the original RETFound while requiring only half the training data and 400 times less compute - a significant advance for accessibility and sustainability. A side-by-side comparison of the principal retinal foundation models is provided in Table 4.

A head-to-head comparison by Hou et al. (2025) between RETFound and DINOv2 - a general-purpose vision foundation model developed by Meta AI that was not specifically trained on medical images - produced a nuanced result: DINOv2-Large achieved an AUROC of 0.892 versus 0.846 for RETFound in multiclass eye disease classification from CFPs. Notably, however, RETFound retained superior performance over DINOv2 for systemic disease prediction tasks (heart failure AUROC 0.796, myocardial infarction 0.732, ischemic stroke 0.754). These findings suggest that large-scale natural-image foundation models can be highly competitive for ocular disease classification when appropriately fine-tuned, while domain-specific pre-training may remain valuable for systemic - but perhaps not ocular - applications.

Two very recent models have pushed the boundaries further. MIRAGE (Morano et al., 2025) is a multimodal foundation model designed for OCT image analysis that incorporates scanning laser ophthalmoscopy (SLO) data alongside OCT and outperforms both DINOv2 and RETFound on segmentation and classification benchmarks. V-JEPA (Judkiewicz et al., 2026) represents perhaps the most conceptually innovative approach to date: rather than treating each OCT B-scan as an independent two-dimensional image, it applies a video-based joint-embedding predictive architecture to the full three-dimensional OCT volume, achieving average AUROCs of 0.94 compared with 0.90 for prior models. This volumetric strategy captures spatial relationships across slices that are inevitably lost when individual B-scans are analyzed in isolation.

The rapid proliferation of retinal foundation models raises important practical questions. Which model should clinicians and researchers adopt? The honest answer, as of early 2026, is that no single model has established clear superiority across all tasks, modalities, and populations. The comparative results between RETFound, DINOv2, MIRAGE, and V-JEPA are still evolving, and conclusions about universal best performance should be presented with appropriate hedging. Nonetheless, the overall trajectory is unmistakable: foundation models are replacing task-specific architectures as the default approach in retinal AI, and this shift is likely to accelerate further as training datasets and computational resources continue to grow.

Table 4. Comparison of self-supervised foundation models for retinal imaging.

Model (Year)	Authors	Pre-training data	Modalities	Distinctive feature
RETFound (2023)	Zhou et al.	1.6 M unlabeled images	CFP + OCT	First retinal foundation model; masked-autoencoder; generalizable to ocular and systemic tasks
VisionFM (2024)	Qiu et al.	3.4 M images	8 ophthalmic modalities	Multi-task generalist for ophthalmic AI; outperforms RETFound on several systemic tasks
EyeFound (2024)	Shi et al.	2.78 M images	11 modalities	Greatest modality breadth; flexible input handling
RETFound-Green (2025)	Engelmann & Bernabeu	≈ half of RETFound's data	CFP + OCT	Comparable performance with ~400× less compute; accessibility / sustainability
DINOv2 vs RETFound (2025)	Hou et al.	Natural images (DINOv2)	CFP	Head-to-head: DINOv2-Large 0.892 vs RETFound 0.846 for eye disease classification; RETFound superior for systemic disease prediction
MIRAGE (2025)	Morano et al.	Multimodal OCT + SLO corpus	OCT + SLO	OCT-focused multimodal benchmark; outperforms DINOv2 and RETFound on segmentation/classification
V-JEPA (2026)	Judkiewicz et al.	OCT volumes	3D OCT	Video-JEPA on volumetric OCT (rather than single B-scans); avg AUROC 0.94 vs 0.90

Abbreviations: AUROC - area under the receiver operating characteristic curve; CFP - color fundus photograph; JEPA - joint-embedding predictive architecture; OCT - optical coherence tomography; SLO - scanning laser ophthalmoscopy.

OCT and OCTA biomarkers for systemic disease

While much of the oculomics literature has focused on color fundus photography, optical coherence tomography (OCT) and OCT angiography (OCTA) offer distinct and complementary advantages. OCT provides high-resolution, cross-sectional images of retinal layers - enabling quantification of the retinal nerve fiber layer (RNFL), ganglion cell-inner plexiform layer (GC-IPL), and choroidal thickness - while OCTA generates depth-resolved maps of the retinal microvasculature without the need for intravenous dye injection.

Farrah et al. (2020) published a comprehensive review of OCT and OCTA biomarkers across the cardio-renal axis, highlighting the shared vulnerability of the retinal and renal microvasculatures to hypertensive and diabetic injury. Kellner et al. (2024) extended this work with a narrative review specifically focused on OCTA-derived microvascular biomarkers - including vessel density, perfusion density, and the foveal avascular zone (FAZ) - as indicators of hypertension, atherosclerosis, and heart failure.

In the neurodegenerative domain, OCT and OCTA findings have been particularly informative. Cheung et al. (2021) reviewed the literature on retinal imaging in Alzheimer's disease, documenting characteristic patterns of RNFL and GC-IPL thinning and parafoveal vessel density loss that distinguish AD patients from healthy controls. These structural and vascular changes appear to precede clinical cognitive decline in some studies, raising the possibility of early, presymptomatic detection. Cheung et al. (2022) used a deep-learning system to extract retinal vessel caliber from fundus photographs and demonstrated that these CFP-derived measurements were associated with the risk of cognitive decline and incident dementia in a prospective cohort.

The automated extraction of OCTA biomarkers has been facilitated by computational pipelines such as that described by Lu et al. (2024), which provides standardized metrics for FAZ area, vessel dispersion, and other quantitative features used in downstream oculomics analyses. Ran et al. (2021) provided a comprehensive review of how deep learning has been applied to OCT for glaucoma assessment, including approaches for automated retinal layer segmentation that have been adopted in many systemic disease studies.

A notable trend in recent years has been the integration of OCT and OCTA data into multimodal models. The Eye-AD system by Hao et al. (2024), for instance, used OCTA images processed through a graph-representation learning framework to detect early AD with impressive accuracy. V-JEPA (Judkiewicz et al., 2026) analyzed full three-dimensional OCT volumes rather than individual slices, capturing inter-slice spatial relationships that single-slice analysis inevitably misses. These multimodal and volumetric approaches are likely to become increasingly important as the field moves beyond two-dimensional fundus photography toward richer, three-dimensional representations of retinal structure and function.

Discussion

The evidence reviewed here leaves little doubt that AI-driven oculomics has progressed, in barely half a decade, from a provocative proof of concept to a field with genuine translational potential. Deep-learning models can extract medically meaningful information from retinal images - information that, in many cases, was invisible to unaided human interpretation. Yet the gap between research performance and clinical deployment remains substantial, and several critical issues warrant discussion.

First, the question of external validity looms large. A striking proportion of the high-performing models reviewed in this article were developed or validated - or both - using the UK Biobank. The UK Biobank is, without question, an extraordinary resource: its combination of standardized retinal imaging, extensive phenotyping, and long-term follow-up is unmatched. However, its participants are predominantly of European ancestry, disproportionately middle-aged, and healthier on average than the general population. Models that perform well in this cohort may not generalize to populations with different ethnic compositions, disease burdens, or imaging equipment. The efforts by Sabanayagam et al. (2020), Cheung et al. (2022), and Yu et al. (2025) to validate across Asian and multi-ethnic cohorts are encouraging, but much more work of this kind is needed, particularly in sub-Saharan Africa, South Asia, and Latin America.

Second, the issue of algorithmic bias deserves sustained attention. Retinal images encode information about age, sex, ethnicity, and smoking status (Poplin et al., 2018), and any model trained on demographically imbalanced data risks perpetuating or amplifying existing health disparities. The oculomics community has yet to adopt standardized fairness metrics or reporting frameworks analogous to those emerging in other areas of clinical AI (Zhu et al., 2025).

Third, the regulatory landscape for oculomics remains underdeveloped. As of 2026, no regulatory agency has approved an AI system specifically for the prediction of systemic (non-ocular) disease from retinal images. The regulatory frameworks that exist - such as the FDA's De Novo pathway for autonomous diabetic retinopathy screening devices - were designed for ocular applications and may require adaptation for systemic disease prediction, where the clinical decision pathway is less clearly defined (Zhu et al., 2025).

Fourth, cost-effectiveness has been insufficiently studied. While the low marginal cost of analyzing a retinal photograph with an AI algorithm is frequently cited as an advantage, the total cost of a screening program encompasses equipment, infrastructure, training, quality assurance, and downstream diagnostic workup. A handful of economic analyses have been conducted in the context of diabetic retinopathy screening (Zhu et al., 2025), but comparable analyses for cardiovascular, renal, or neurodegenerative oculomics remain scarce.

Fifth, the relationship between foundation models and clinical practice is still taking shape. The rapid succession of RETFound, VisionFM, EyeFound, RETFound-Green, DINOv2 fine-tuning, MIRAGE, and V-JEPA has demonstrated impressive technical progress, but it has also created a fragmented landscape in which no single model has been adopted as a community standard. From a clinical perspective, what matters is not which model achieves the highest AUROC on a particular benchmark but whether the model has been prospectively validated, integrated into a clinical workflow, and shown to improve patient outcomes. None of these criteria has been met for any foundation model as of mid-2026.

Finally, it is worth reflecting on the broader clinical context. Oculomics does not exist in isolation; it must compete for clinical attention with established screening modalities - blood tests, imaging studies, clinical risk scores - that already occupy the clinical workflow. The strongest case for oculomics arises in scenarios where conventional screening is unavailable, impractical, or insufficiently predictive. Population-level cardiovascular screening in low-resource settings, opportunistic renal screening during routine eye examinations, and early detection of neurodegenerative disease before symptoms are clinically apparent are all scenarios where retinal AI could plausibly fill unmet needs. Proving this case, however, will require the large-scale, pragmatic clinical trials that the field has so far lacked.

Limitations of this review

Several limitations of the present review should be acknowledged. First, this is a narrative - rather than systematic - synthesis; no formal protocol was registered, no risk-of-bias assessment was applied, and no quantitative meta-analysis was performed. Selection of studies, although guided by predefined search terms and inclusion criteria, inevitably reflects a degree of subjective judgment on the part of the authors. Second, the literature search was restricted to three databases (PubMed, Scopus, and Google Scholar) and to publications in English; relevant work indexed elsewhere or published in other languages may therefore have been missed. Third, the field of AI-driven oculomics is evolving rapidly, and several preprints, conference proceedings, and very recent publications were not included; the conclusions of this review reflect the state of the literature as of April 2026 and are likely to require updating as new validation studies, foundation models, and regulatory decisions emerge. Fourth, performance metrics reported across primary studies (most often AUROC) are highly dependent on the underlying cohort, imaging modality, and reference standard, which limits the comparability of headline figures cited throughout this review. Finally, by design this review focuses on AI-driven approaches and does not exhaustively cover the broader ecosystem of conventional retinal imaging biomarkers, ophthalmic genetics, or molecular oculomics - domains that remain complementary and clinically important.

Conclusions

AI-driven oculomics has established itself as a vibrant and rapidly evolving field at the intersection of ophthalmology, internal medicine, neurology, and computer science. The evidence accumulated between 2018 and 2026 demonstrates that deep-learning models applied to retinal fundus photographs and OCT images can predict cardiovascular events, detect chronic kidney disease and metabolic syndrome, identify early neurodegenerative changes, and estimate biological age - often with diagnostic performance that rivals or complements traditional clinical tools.

The emergence of self-supervised foundation models has further accelerated progress by enabling high-performance prediction with minimal labeled data and across multiple imaging modalities. Volumetric approaches such as V-JEPA, which analyze full three-dimensional OCT data, represent a particularly promising direction.

Nonetheless, the field faces substantial translational hurdles. Prospective clinical validation in diverse, real-world populations is urgently needed. Algorithmic fairness must be systematically evaluated and reported. Regulatory frameworks must evolve to accommodate the novel category of systemic disease prediction from ophthalmic imaging. Cost-effectiveness analyses must be conducted to justify the integration of oculomics into existing screening infrastructures.

If these challenges are met, oculomics has the potential to transform how systemic diseases are detected and managed - leveraging a simple, non-invasive retinal photograph as a gateway to personalized, data-driven healthcare.

Author Contributions: Conceptualization: Adrianna Buż. Methodology: Adrianna Buż, Nicol Baran, Marta Bajkowska-Piterak, Jerzy Buszko, Damian Danilczuk, Wiktor Adamiec, Monika Szlachta-Gubernat, Przemysław Piterak, Michał Dyś. Formal analysis: Adrianna Buż, Nicol Baran, Marta Bajkowska-Piterak, Jerzy Buszko, Damian Danilczuk, Wiktor Adamiec, Monika Szlachta-Gubernat, Przemysław Piterak, Michał Dyś. Resources: Adrianna Buż, Nicol Baran, Marta Bajkowska-Piterak, Jerzy Buszko, Damian Danilczuk, Wiktor Adamiec, Monika Szlachta-Gubernat, Przemysław Piterak, Michał Dyś. Writing - original draft preparation: Adrianna Buż, Nicol Baran, Marta Bajkowska-Piterak. Writing - review and editing: Nicol Baran, Marta Bajkowska-Piterak. Visualization: Adrianna Buż. Supervision: Nicol Baran, Marta Bajkowska-Piterak. Project administration: Adrianna Buż. All authors have read and agreed to the published version of the manuscript.

Funding Statement: This research received no external funding.

Data Availability Statement: No new data were generated in the course of this review. All cited data are available from the original publications.

Conflict of Interest Statement: The authors declare no conflicts of interest.

Declaration of the Use of Generative AI and AI-Assisted Technologies in the Writing Process: During the preparation of this work, the authors used ChatGPT (OpenAI) exclusively for language editing and grammatical correction of the manuscript, and Claude 4.7 Opus (Anthropic) to format and generate the tables and figures based on the provided data. Tools were not involved in data collection, analysis, interpretation of results, or the formulation of conclusions. All substantive content was generated, reviewed, and approved by the human authors, who accept full responsibility for the integrity of the publication.

REFERENCES

1. Wagner, S. K., Fu, D. J., Faes, L., Liu, X., Huemer, J., Khalid, H., Ferraz, D., Korot, E., Kelly, C., Balaskas, K., Denniston, A. K., & Keane, P. A. (2020). Insights into systemic disease through retinal imaging-based oculomics. *Translational Vision Science & Technology, 9*(2), 6. https://doi.org/10.1167/tvst.9.2.6
2. Zhu, Z., Wang, Y., Qi, Z., Hu, W., Zhang, X., Wagner, S. K., Wang, Y., Ran, A. R., Ong, J., Waisberg, E., Masalkhi, M., Suh, A., Tham, Y. C., Cheung, C. Y., Yang, X., Yu, H., Ge, Z., Wang, W., Sheng, B., . . . Wong, T. Y. (2025). Oculomics: Current concepts and evidence. *Progress in Retinal and Eye Research, 106*, 101350. https://doi.org/10.1016/j.preteyeres.2025.101350
3. Zhu, Z., Hu, W., Chen, R., Xiong, R., Wang, W., Shang, X., Chen, Y., Kiburg, K., Shi, D., He, S., Huang, Y., Zhang, X., Tang, S., Zeng, J., Yu, H., Yang, X., & He, M. (2022). Retinal age gap as a predictive biomarker of stroke risk. *BMC Medicine, 20*(1), 466. https://doi.org/10.1186/s12916-022-02620-w
4. Ghenciu, L. A., Dima, M., Stoicescu, E. R., Iacob, R., Boru, C., & Haţegan, O. A. (2024). Retinal imaging-based oculomics: Artificial intelligence as a tool in the diagnosis of cardiovascular and metabolic diseases. *Biomedicines, 12*(9), 2150. https://doi.org/10.3390/biomedicines12092150
5. Poplin, R., Varadarajan, A. V., Blumer, K., Liu, Y., McConnell, M. V., Corrado, G. S., Peng, L., & Webster, D. R. (2018). Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering, 2*(3), 158–164. https://doi.org/10.1038/s41551-018-0195-0
6. Rim, T. H., Lee, C. J., Tham, Y. C., Cheung, N., Yu, M., Lee, G., Kim, Y., Ting, D. S. W., Chong, C. C. Y., Choi, Y. S., Yoo, T. K., Ryu, I. H., Baik, S. J., Kim, Y. A., Kim, S. K., Lee, S. H., Lee, B. K., Kang, S. M., Wong, E. Y. M., . . . Wong, T. Y. (2021). Deep-learning-based cardiovascular risk stratification using coronary artery calcium scores predicted from retinal photographs: A multicohort study. *The Lancet Digital Health, 3*(5), e306-e316. https://doi.org/10.1016/S2589-7500(21)00043-1
7. Tseng, R. M. W. W., Rim, T. H., Shantsila, E., Yi, J. K., Park, S., Kim, S. S., Lee, C. J., Thakur, S., Nusinovici, S., Peng, Q., Li, H., Xue, C. C., Zhang, L., Cheung, C. Y., Lip, G. Y. H., Wong, T. Y., & Cheng, C. Y. (2023). Validation of a deep-learning-based retinal biomarker (Reti-CVD) in the prediction of cardiovascular disease: Data from UK Biobank. *BMC Medicine, 21*, 28. https://doi.org/10.1186/s12916-023-02732-x
8. Diaz-Pinto, A., Ravikumar, N., Attar, R., Suinesiaputra, A., Zhao, Y., Levelt, E., Dall'Armellina, E., Lorenzi, M., Chen, Q., Keenan, T. D. L., Agrón, E., Chew, E. Y., Lu, Z., Gale, C. P., Gale, R. P., Plein, S., & Frangi, A. F. (2022). Predicting myocardial infarction through retinal scans and minimal personal information. *Nature Machine Intelligence, 4*(1), 55–61. https://doi.org/10.1038/s42256-021-00427-7
9. Vaghefi, E., Squirrell, D., Yang, S., An, S., Xie, L., Durbin, M. K., Hou, H., Marshall, J., Shreibati, J., McConnell, M. V., & Budoff, M. (2024). Development and validation of a deep-learning model to predict 10-year atherosclerotic cardiovascular disease risk from retinal images using the UK Biobank and EyePACS 10K datasets. *American Journal of Preventive Cardiology, 17*, 100633. https://doi.org/10.1016/j.ajpc.2023.100633
10. Chang, J., Ko, A., Park, S. M., Choi, S., Kim, K., Kim, S. M., Yun, J. M., Kang, U., Shin, I. H., Shin, J. Y., Ko, T., Lee, J., Oh, B. L., & Park, K. H. (2020). Association of cardiovascular mortality and deep learning-funduscopic atherosclerosis score derived from retinal fundus images. *American Journal of Ophthalmology, 217*, 121–130. https://doi.org/10.1016/j.ajo.2020.03.027
11. Cheung, C. Y., Xu, D., Cheng, C. Y., Sabanayagam, C., Tham, Y. C., Yu, M., Rim, T. H., Chai, C. Y., Gopinath, B., Mitchell, P., Poulton, R., Moffitt, T. E., Caspi, A., Yam, J. C., Tham, C. C., Jonas, J. B., Wang, Y. X., Song, S. J., Burrell, L. M., . . . Wong, T. Y. (2021). A deep-learning system for the assessment of cardiovascular disease risk via the measurement of retinal-vessel calibre. *Nature Biomedical Engineering, 5*(6), 498–508. https://doi.org/10.1038/s41551-020-00626-4
12. Zhang, K., Liu, X., Xu, J., Yuan, J., Cai, W., Chen, T., Wang, K., Gao, Y., Nie, S., Xu, X., Qin, X., Su, Y., Xu, W., Olvera, A., Xue, K., Li, Z., Zhang, M., Zeng, X., Zhang, C. L., . . . Wang, G. (2021). Deep-learning models for the detection and incidence prediction of chronic kidney disease and type 2 diabetes from retinal fundus images. *Nature Biomedical Engineering, 5*(6), 533–545. https://doi.org/10.1038/s41551-021-00745-6
13. Sabanayagam, C., Xu, D., Ting, D. S. W., Nusinovici, S., Banu, R., Hamzah, H., Lim, C., Tham, Y. C., Cheung, C. Y., Tai, E. S., Wang, Y. X., Jonas, J. B., Cheng, C. Y., Lee, M. L., Hsu, W., & Wong, T. Y. (2020). A deep learning algorithm to detect chronic kidney disease from retinal photographs in community-based populations. *The Lancet Digital Health, 2*(6), e295-e302. https://doi.org/10.1016/S2589-7500(20)30063-7
14. Joo, Y. S., Rim, T. H., Koh, H. B., Yi, J., Kim, H., Lee, G., Kim, Y. A., Kang, S. W., Kim, S. S., & Park, J. T. (2023). Non-invasive chronic kidney disease risk stratification tool derived from retina-based deep learning and clinical factors. *npj Digital Medicine, 6*, 114. https://doi.org/10.1038/s41746-023-00860-5
15. Wu, Q., Li, J., Zhao, L., Liu, D., Wen, J., Wang, Y., Wang, Y., Huang, N., Jiang, L., Liu, Q., Lin, H., Wan, P., Yang, S., Chen, W., Ye, H., Hassan, M. H. R., Nur, A. H., Dai, Z., Guo, J., . . . Lin, H. (2025). A noninvasive model for chronic kidney disease screening and common pathological type identification from retinal images. *Nature Communications, 16*, 62273. https://doi.org/10.1038/s41467-025-62273-0

16. Bhak, Y., Lee, Y. H., Kim, J., Lee, K., Lee, D., Jang, E. C., Jang, E., Lee, C. S., Kang, E. S., Park, S., Han, H. W., & Nam, S. M. (2025). Diagnosis of chronic kidney disease using retinal imaging and urine dipstick data: Multimodal deep learning approach. *JMIR Medical Informatics, 13*, e55825. https://doi.org/10.2196/55825
17. Yip, M. Y. T., Lim, G., Lim, Z. W., Nguyen, Q. D., Chong, C. C. Y., Yu, M., Bellemo, V., Xie, Y., Lee, X. Q., Hamzah, H., Ho, J., Tan, T. E., Sabanayagam, C., Grzybowski, A., Tan, G. S. W., Hsu, W., Lee, M. L., Wong, T. Y., & Ting, D. S. W. (2020). Technical and imaging factors influencing performance of deep learning systems for diabetic retinopathy. *npj Digital Medicine, 3*, 40. https://doi.org/10.1038/s41746-020-0247-1
18. Zhang, L., Yuan, M., An, Z., Zhao, X., Wu, H., Li, H., Wang, Y., Sun, B., Li, H., Ding, S., Zeng, X., Chao, L., Li, P., & Wu, W. (2020). Prediction of hypertension, hyperglycemia and dyslipidemia from retinal fundus photographs via deep learning: A cross-sectional study of chronic diseases in central China. *PLOS ONE, 15*(5), e0233166. https://doi.org/10.1371/journal.pone.0233166
19. Cheung, C. Y., Ran, A. R., Wang, S., Chan, V. T. T., Sham, K., Hilal, S., Venketasubramanian, N., Cheng, C. Y., Sabanayagam, C., Tham, Y. C., Schmetterer, L., McKay, G. J., Williams, M. A., Wong, A., Au, L. W. C., Lu, Z., Yam, J. C., Tham, C. C., Chan, J. C., . . . Wong, T. Y. (2022). A deep-learning model for detection of Alzheimer's disease based on retinal photographs: A retrospective, multicentre case-control study. *The Lancet Digital Health, 4*(11), e806-e815. https://doi.org/10.1016/S2589-7500(22)00169-8
20. Wagner, S. K., Hughes, F., Cortina-Borja, M., Pontikos, N., Struyven, R., Liu, X., Montgomery, H., Alexander, D. C., Topol, E., Petersen, S. E., Balaskas, K., Hindley, J., Petzold, A., Rahi, J. S., Denniston, A. K., & Keane, P. A. (2022). AlzEye: Longitudinal record-level linkage of ophthalmic imaging and hospital admissions of 353,157 patients in London, UK. *BMJ Open, 12*(3), e058552. https://doi.org/10.1136/bmjopen-2021-058552
21. Tian, J., Smith, G., Guo, H., Liu, B., Pan, Z., Wang, Z., Xiong, S., & Fang, R. (2021). Modular machine learning for Alzheimer's disease classification from retinal vasculature. *Scientific Reports, 11*, 238. https://doi.org/10.1038/s41598-020-80312-2
22. Hao, J., Kwapong, W. R., Shen, T., Fu, H., Xu, Y., Lu, Q., Liu, S., Zhang, J., Liu, Y., Zhao, Y., Zheng, Y., Frangi, A. F., Zhang, S., Qi, H., & Zhao, Y. (2024). Early detection of dementia through retinal imaging and trustworthy AI (Eye-AD). *npj Digital Medicine, 7*, 294. https://doi.org/10.1038/s41746-024-01304-4
23. Tham, Y. C., Anees, A., Zhang, L., Goh, J. H. L., Rim, T. H., Nusinovici, S., Hamzah, H., Chee, M. L., Tjio, G., Li, S., Xu, X., Goh, R., Lamoureux, E., Sabanayagam, C., Cheung, C. Y., Tang, F., Wang, Y. X., Nangia, V., Jonas, J. B., . . . Liu, Y. (2021). Referral for disease-related visual impairment using retinal photograph-based deep learning: A proof-of-concept, model development study. *The Lancet Digital Health, 3*(1), e29-e40. https://doi.org/10.1016/S2589-7500(20)30226-0
24. Richardson, A., Kundu, A., Henao, R., Lee, T., Scott, B. L., Grewal, D. S., & Fekrat, S. (2024). Multimodal retinal imaging classification for Parkinson's disease using a convolutional neural network. *Translational Vision Science & Technology, 13*(8), 23. https://doi.org/10.1167/tvst.13.8.23
25. Tran, C., Shen, K., Liu, K., Ashok, A., Ramirez-Zamora, A., Chen, J., Li, Y., & Fang, R. (2024). Deep learning predicts prevalent and incident Parkinson's disease from UK Biobank fundus imaging. *Scientific Reports, 14*, 3637. https://doi.org/10.1038/s41598-024-54251-1
26. Sim, M. A., Tham, Y. C., Nusinovici, S., Quek, T. C., Yu, M., Xue, C. C., Chee, M. L., Peng, Q. S., Tan, E. S. J., Chan, S. P., Cai, Y., Chong, E. J. Y., Tan, B. Y., Venketasubramanian, N., Hilal, S., Lai, M. K. P., Choi, H., Richards, A. M., Cheng, C. Y., & Chen, C. L. H. (2025). A deep-learning retinal aging biomarker for cognitive decline and incident dementia. *Alzheimer's & Dementia, 21*(3), e14601. https://doi.org/10.1002/alz.14601
27. Zhu, Z., Shi, D., Guankai, P., Tan, Z., Shang, X., Hu, W., Liao, H., Zhang, X., Huang, Y., Yu, H., Meng, W., Wang, W., Ge, Z., Yang, X., & He, M. (2023). Retinal age gap as a predictive biomarker for mortality risk. *British Journal of Ophthalmology, 107*(4), 547–554. https://doi.org/10.1136/bjophthalmol-2021-319807
28. Nusinovici, S., Rim, T. H., Yu, M., Lee, G., Tham, Y. C., Cheung, N., Chong, C. C. Y., Soh, Z. D., Thakur, S., Lee, C. J., Sabanayagam, C., Lee, B. K., Park, S., Kim, S. S., Kim, H. C., Wong, T. Y., & Cheng, C. Y. (2022). Retinal photograph-based deep learning predicts biological age, and stratifies morbidity and mortality risk (RetiAGE). *Age and Ageing, 51*(4), afac065. https://doi.org/10.1093/ageing/afac065
29. Nusinovici, S., Rim, T. H., Li, H., Yu, M., Deshmukh, M., Quek, T. C., Lee, G., Chong, C. C. Y., Peng, Q., Xue, C. C., Zhu, Z., Chew, E. Y., Sabanayagam, C., Wong, T. Y., Tham, Y. C., & Cheng, C. Y. (2024). Application of a deep-learning marker for morbidity and mortality prediction derived from retinal photographs (RetiPhenoAge): A cohort development and validation study. *The Lancet Healthy Longevity, 5*(8), e514-e523. https://doi.org/10.1016/S2666-7568(24)00110-1
30. Yu, Z., Chen, R., Gui, P., Wang, W., Razzak, I., Alinejad-Rokny, H., Zeng, X., Shang, X., Zhang, L., Yang, X., Yu, H., Huang, W., Lu, H., van Wijngaarden, P., He, M., Zhu, Z., & Ge, Z. (2025). A cross-population study of retinal aging biomarkers with longitudinal pre-training and label distribution learning. *npj Digital Medicine, 8*, 331. https://doi.org/10.1038/s41746-025-00331-6
31. Zhou, Y., Chia, M. A., Wagner, S. K., Ayhan, M. S., Williamson, D. J., Struyven, R. R., Liu, T., Xu, M., Lozano, M. G., Woodward-Court, P., Kihara, Y., UK Biobank Eye & Vision Consortium, Altmann, A., Lee, A. Y., Topol, E. J., Denniston, A. K., Alexander, D. C., & Keane, P. A. (2023). A foundation model for generalizable disease detection from retinal images (RETFound). *Nature, 622*(7981), 156–163. https://doi.org/10.1038/s41586-023-06555-x

32. Qiu, J., Wu, J., Wei, H., Shi, P., Zhang, M., Sun, Y., Li, L., Liu, H., Liu, H., Hou, S., Zhao, Y., Shi, X., Xian, J., Qu, X., Zhu, S., Pan, L., Chen, X., Zhang, X., Jiang, S., . . . Yuan, W. (2024). Development and validation of a multimodal multitask vision foundation model for generalist ophthalmic artificial intelligence (VisionFM). *NEJM AI, 1*(12), AIoa2300221. https://doi.org/10.1056/AIoa2300221
33. Shi, D., Zhang, W., Chen, X., Liu, Y., Yang, J., Huang, S., Tham, Y. C., Zheng, Y., & He, M. (2024). EyeFound: A multimodal generalist foundation model for ophthalmic imaging. arXiv. https://doi.org/10.48550/arXiv.2405.11338
34. Engelmann, J., & Bernabeu, M. O. (2025). Training a high-performance retinal foundation model with half the data and 400 times less compute (RETFound-Green). *Nature Communications, 16*, 62123. https://doi.org/10.1038/s41467-025-62123-z
35. Hou, Q., Zhou, Y., Goh, J. H. L., Zou, K., Yew, S. M. E., Srinivasan, S., Wang, M., Lo, T. W. S., Lei, X., Wagner, S. K., Chia, M. A., Yang, G. D., Jiang, H., Ran, A. R., Santos, R., Somfai, G. M., Zhou, J. H., Chen, H., Chen, Q., . . . Tham, Y. C. (2025). Can a natural image-based foundation model outperform a retina-specific model in detecting ocular and systemic diseases? Head-to-head DINOv2 vs RETFound. *Ophthalmology Science, 5*(6), 100764. https://doi.org/10.1016/j.xops.2025.100764
36. Morano, J., Fazekas, B., Sükei, E., Fecso, R., Emre, T., Gumpinger, M., Faustmann, G., Oghbaie, M., Schmidt-Erfurth, U., & Bogunović, H. (2025). MIRAGE: Multimodal foundation model and benchmark for comprehensive retinal OCT image analysis. *npj Digital Medicine, 8*, 1852. https://doi.org/10.1038/s41746-025-01852-3
37. Judkiewicz, R., Berkowitz, E., Meisel, M., Michaeli, T., & Behar, J. A. (2026). Shifting the retinal foundation models paradigm from slices to volumes for OCT (V-JEPA). *npj Digital Medicine, 9*, 2496. https://doi.org/10.1038/s41746-026-02496-7
38. Farrah, T. E., Dhillon, B., Keane, P. A., Webb, D. J., & Dhaun, N. (2020). The eye, the kidney, and cardiovascular disease: Old concepts, better tools, and new horizons. *Kidney International, 98*(2), 323–342. https://doi.org/10.1016/j.kint.2020.01.039
39. Kellner, R. L., Harris, A., Ciulla, L., Guidoboni, G., Verticchio Vercellin, A., Oddone, F., Carnevale, C., Zaid, M., Antman, G., Kuvin, J. T., & Siesky, B. (2024). The eye as the window to the heart: Optical coherence tomography angiography biomarkers as indicators of cardiovascular disease. *Journal of Clinical Medicine, 13*(3), 829. https://doi.org/10.3390/jcm13030829
40. Ran, A. R., Tham, C. C., Chan, P. P., Cheng, C. Y., Tham, Y. C., Rim, T. H., & Cheung, C. Y. (2021). Deep learning in glaucoma with optical coherence tomography: A review. *Eye, 35*(1), 188–201. https://doi.org/10.1038/s41433-020-01191-5
41. Cheung, C. Y., Wong, W. L. E., Hilal, S., Kan, C. N., Gyanwali, B., Tham, Y. C., Schmetterer, L., Xu, D., Lee, M. L., Hsu, W., Venketasubramanian, N., Tan, B. Y., Wong, T. Y., & Chen, C. P. L. H. (2022). Deep-learning retinal vessel calibre measurements and risk of cognitive decline and dementia. *Brain Communications, 4*(4), fcac212. https://doi.org/10.1093/braincomms/fcac212
42. Lu, B., Li, Y., Xie, L., Chiu, K., Hao, X., Xu, J., Luo, J., & Sham, P. C. (2024). Computational retinal microvascular biomarkers from an OCTA image in clinical investigation. *Biomedicines, 12*(4), 868. https://doi.org/10.3390/biomedicines12040868
43. Cheung, C. Y., Mok, V., Foster, P. J., Trucco, E., Chen, C., & Wong, T. Y. (2021). Retinal imaging in Alzheimer's disease. *Journal of Neurology, Neurosurgery & Psychiatry, 92*(9), 983–994. https://doi.org/10.1136/jnnp-2020-325347