MLPR

END-SEM

“Explainable Retinal Age Gap Prediction using RETFound and Vascular Biomarkers”

GROUP 39

Agamvir Singh, Hitesh Jindal, Subhi Sabharwal

PROBLEM STATEMENT

BACKGROUND

When doctors look into the human eye using a retinal camera, they are not just seeing the eye they are seeing a living map of the body's blood vessels.

The retina is the only place in the body where micro-blood vessels can be observed directly without surgery.
These tiny vessels supply oxygen and nutrients to retinal tissue and are connected to the same vascular system that supplies the brain and heart.

Because of this, changes in retinal blood vessels often reflect systemic aging and disease.

WHY EXAMINE RETINA???

Many diseases such as stroke, cardiovascular disease, and neurodegeneration originate from problems in the body's microvascular system.

However, studying blood vessels in major organs is difficult.

- Brain: requires expensive imaging such as MRI or CT scans to observe vascular structures.
- Heart: coronary vessels are usually examined using invasive procedures like angiography.

The retina is unique because it is the only place in the body where micro blood vessels can be directly observed non invasively using retinal imaging.

The retinal microvasculature shares strong anatomical and physiological similarities with the brain's vascular system, making the retina a natural window into systemic vascular health.

Changes in retinal vessels such as tortuosity, narrowing, and branching patterns often reflect aging and disease processes occurring throughout the body.

A PEEK INTO RETINAL FUNDUS

👀

Key Structures and Their Biological Role

Optic Disc:
The point where the optic nerve exits the eye and connects the retina to the brain.
It contains the entry and exit of the central retinal artery and vein.

Macula:
A specialized region responsible for central vision and detailed visual tasks like reading.

Fovea:
Located at the center of the macula.
Contains a very high concentration of photoreceptors, providing the sharpest vision.

Central Retinal Artery and Vein:
These vessels supply oxygen and nutrients to the retina and remove metabolic waste.

Retinal Vessels (Arteries and Venules):
Form a dense microvascular network across the retina.
Changes in their thickness, curvature, and branching patterns often reflect systemic vascular health.

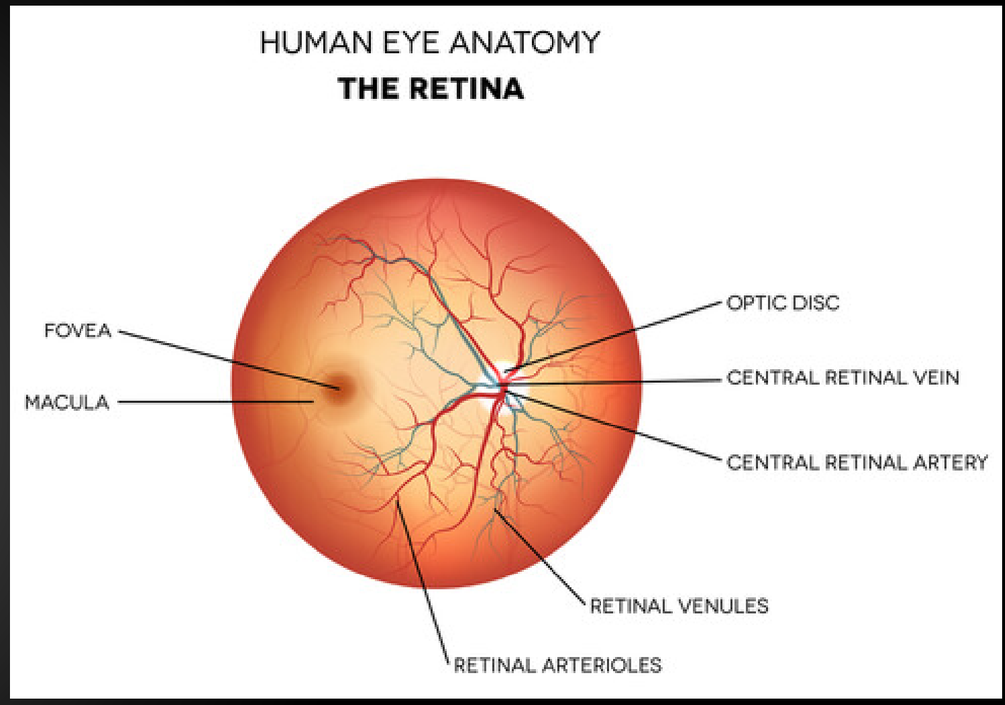

CORE PROBLEM

Although modern deep learning models can predict retinal age with good accuracy, there is a major limitation.

Current models operate as black box systems.

They take a retinal image as input and output a predicted age, but they do not explain:

Which retinal regions influenced the prediction
Which vascular structures contributed to the age estimate
Which biological features indicate accelerated aging

In medical applications, interpretability is extremely important because clinicians need to understand why the model made a prediction.

Without this explanation, it is difficult for doctors to trust the model or use it for clinical decision making.

Retinal Image
↓
Deep Learning Model
↓
Predicted Age

Cause?

Monitoring?

Condition?

Diagnosis??

APPLICATIONS+IMPACT

1. Clinically Trustworthy AI

Traditional models only output a predicted age.
Our model explains which retinal regions and vascular features influenced the prediction, making the system more trustworthy for clinicians.

2. Early Risk Detection for Systemic Diseases

Because retinal vessels reflect systemic microvascular health, interpretable predictions could help identify individuals at higher risk for:
stroke,cardiovascular disease,neurodegenerative disorders

while also explaining which vascular abnormalities are responsible.

3. Identification of Biological Markers of Aging

By linking predictions to measurable features such as:
vessel tortuosity
vessel thickness
branching patterns
the model can help researchers discover new vascular biomarkers of biological aging.

LITERATURE SURVEY

PAPER - 1

RetiphenoAge: Estimating Biological Aging from Retinal Images

RetiphenoAge is a deep learning framework developed to estimate biological aging using retinal fundus images. Instead of relying only on chronological age, the model predicts Phenotypic Age (PhenoAge), a clinically validated biomarker that reflects the physiological condition of an individual.

Phenotypic Age represents the biological state of the body and is considered a stronger indicator of mortality and disease risk than chronological age alone. It is computed using chronological age along with nine blood-based biomarkers that capture metabolic, inflammatory, and organ-function related health indicators.

Examples of these biomarkers include:
Albumin,Creatinine,Glucose,C-Reactive Protein (CRP),Alkaline Phosphatase,White Blood Cell Count

The core idea behind RetiphenoAge is that the retina acts as a window to systemic health, as retinal vasculature reflects changes in cardiovascular, metabolic, and neurological conditions. By learning patterns present in retinal images, the model attempts to infer the biological aging state of an individual non-invasively.

The predicted age is then compared to chronological age to estimate accelerated or decelerated biological aging, enabling identification of individuals at higher systemic health risk.

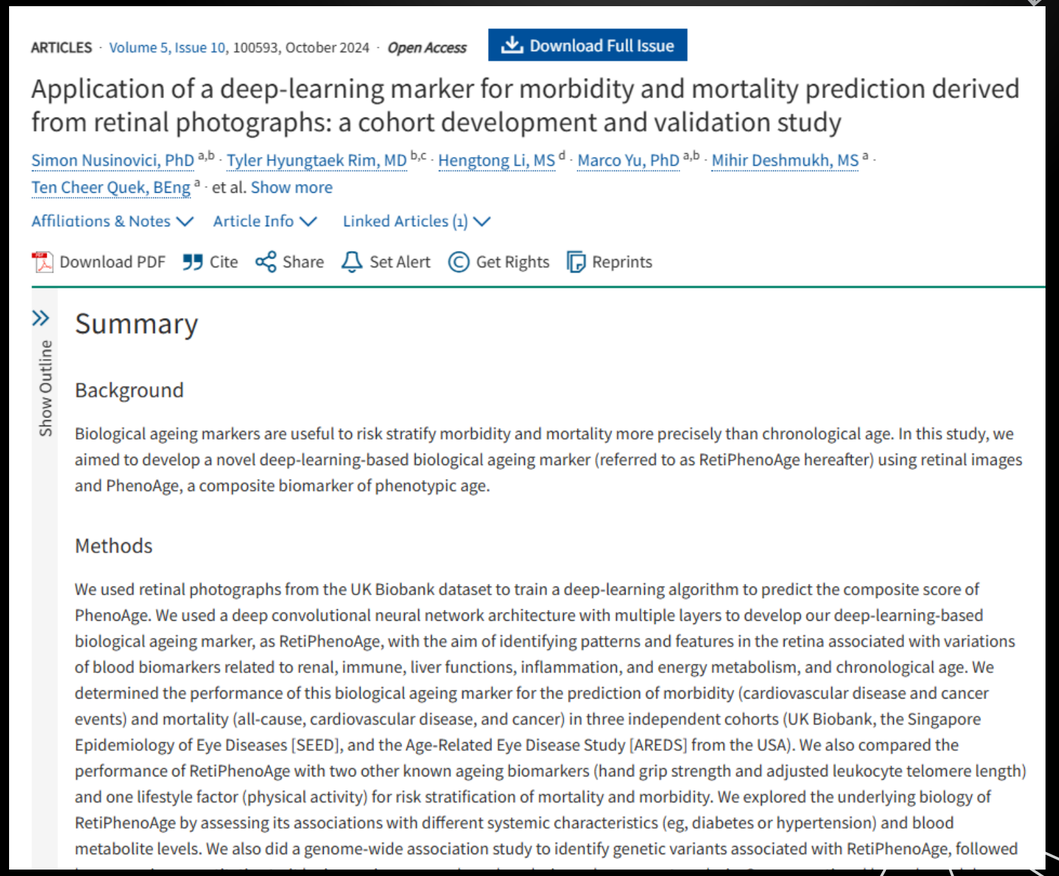

ARTICLES · Volume 5, Issue 10, 100593, October 2024 · *Open Access* Download Full Issue

Application of a deep-learning marker for morbidity and mortality prediction derived from retinal photographs: a cohort development and validation study

Simon Nusinovici, PhD [a,b] · Tyler Hyungtaek Rim, MD [b,c] · Hengtong Li, MS [d] · Marco Yu, PhD [a,b] · Mihir Deshmukh, MS [a] · Ten Cheer Quek, BEng [a] · et al. Show more

Affiliations & Notes · Article Info · Linked Articles (1)

Download PDF · Cite · Share · Set Alert · Get Rights · Reprints

Show Outline

Summary

Background

Biological ageing markers are useful to risk stratify morbidity and mortality more precisely than chronological age. In this study, we aimed to develop a novel deep-learning-based biological ageing marker (referred to as RetiPhenoAge hereafter) using retinal images and PhenoAge, a composite biomarker of phenotypic age.

Methods

We used retinal photographs from the UK Biobank dataset to train a deep-learning algorithm to predict the composite score of PhenoAge. We used a deep convolutional neural network architecture with multiple layers to develop our deep-learning-based biological ageing marker, as RetiPhenoAge, with the aim of identifying patterns and features in the retina associated with variations of blood biomarkers related to renal, immune, liver functions, inflammation, and energy metabolism, and chronological age. We determined the performance of this biological ageing marker for the prediction of morbidity (cardiovascular disease and cancer events) and mortality (all-cause, cardiovascular disease, and cancer) in three independent cohorts (UK Biobank, the Singapore Epidemiology of Eye Diseases [SEED], and the Age-Related Eye Disease Study [AREDS] from the USA). We also compared the performance of RetiPhenoAge with two other known ageing biomarkers (hand grip strength and adjusted leukocyte telomere length) and one lifestyle factor (physical activity) for risk stratification of mortality and morbidity. We explored the underlying biology of RetiPhenoAge by assessing its associations with different systemic characteristics (eg, diabetes or hypertension) and blood metabolite levels. We also did a genome-wide association study to identify genetic variants associated with RetiPhenoAge, followed

DATASET AND RETINAL IMAGE PROCESSING PIPELINE

UK Biobank retinal imaging dataset

The large scale of the dataset allows the model to learn diverse retinal patterns associated with systemic aging.

To ensure that the model generalizes well beyond the training data, the framework was also evaluated on external validation datasets, demonstrating consistent performance across different populations and imaging conditions.

MODEL PIPELINE

Image Acquisition
Retinal fundus images are captured using specialized ophthalmic imaging devices that photograph the interior surface of the eye.

Image Resizing
All images are resized to a uniform resolution so they can be efficiently processed by the neural network.

Pixel Normalization
Pixel intensity values are normalized to reduce variations in brightness and contrast across different images.

Color Normalization
Color variations caused by different imaging devices or illumination conditions are corrected.

Data Augmentation
Various transformations such as image rotation, flipping, and brightness adjustments are applied to increase dataset diversity and reduce model overfitting.

DEEP LEARNING ARCHITECTURE

The RetiPhenoAge framework uses a Convolutional Neural Network (CNN) to analyze retinal fundus images and learn patterns associated with biological aging. CNNs are widely used in medical imaging because they can automatically extract features from images.

As the image passes through the network:

- Early layers detect basic patterns such as edges, textures, and vessel boundaries.
- Intermediate layers identify retinal structures like vessel branching, tortuosity, and optic disc features.
- Deeper layers learn higher-level aging indicators such as microvascular abnormalities and structural retinal changes.

Using these features, the model predicts phenotypic age from retinal images and classifies individuals into different aging-risk groups.

Model Performance

Mean Absolute Error (MAE) ≈ 3.5–5 years in biological age prediction

However, a major limitation is that the model behaves as a black box.
It predicts biological age but does not explain which retinal regions or features drive the prediction.
This lack of interpretability limits clinical trust and motivates the need for explainable, attention-guided retinal aging models.

PAPER - 2

Retinal Age Gap: A Predictive Biomarker for Mortality Risk

This study explores the concept of retinal age gap, a biomarker derived from retinal fundus images that reflects the difference between an individual's predicted retinal age and their chronological age.

The central hypothesis of the study is that the retina captures microvascular and structural changes associated with systemic aging, making it possible to estimate a person's biological aging state from retinal images.

Retinal Age Gap Definition

Retinal Age Gap = Predicted Retinal Age - Chronological Age

Interpretation of the gap:

Positive Retinal Age Gap
Indicates accelerated biological aging and higher systemic health risk.

Negative Retinal Age Gap
Suggests healthier physiological aging compared to chronological age.

The study demonstrates that individuals with a larger positive retinal age gap show significantly higher mortality risk, suggesting that retinal imaging can act as a non invasive biomarker for predicting long-term health outcomes.

This approach highlights the potential of retinal fundus images as a scalable tool for population level health monitoring and early disease risk assessment.

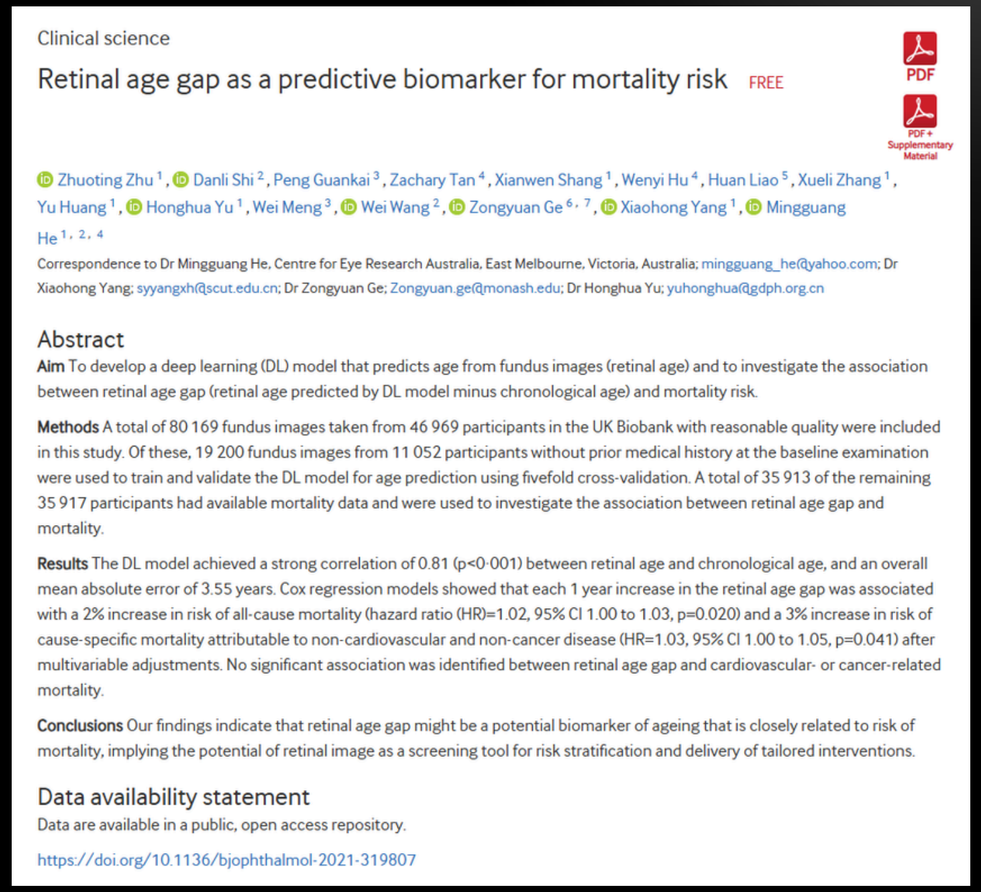

Clinical science

Retinal age gap as a predictive biomarker for mortality risk FREE

PDF

PDF + Supplementary Material

Zhuoting Zhu [1], Danli Shi [2], Peng Guankai [3], Zachary Tan [4], Xianwen Shang [1], Wenyi Hu [4], Huan Liao [5], Xueli Zhang [1], Yu Huang [1], Honghua Yu [1], Wei Meng [3], Wei Wang [2], Zongyuan Ge [6, 7], Xiaohong Yang [1], Mingguang He [1, 2, 4]

Correspondence to Dr Mingguang He, Centre for Eye Research Australia, East Melbourne, Victoria, Australia; mingguang_he@yahoo.com; Dr Xiaohong Yang; syyangxh@scut.edu.cn; Dr Zongyuan Ge; Zongyuan.ge@monash.edu; Dr Honghua Yu; yuhonghua@gdph.org.cn

Abstract

Aim To develop a deep learning (DL) model that predicts age from fundus images (retinal age) and to investigate the association between retinal age gap (retinal age predicted by DL model minus chronological age) and mortality risk.

Methods A total of 80 169 fundus images taken from 46 969 participants in the UK Biobank with reasonable quality were included in this study. Of these, 19 200 fundus images from 11 052 participants without prior medical history at the baseline examination were used to train and validate the DL model for age prediction using fivefold cross-validation. A total of 35 913 of the remaining 35 917 participants had available mortality data and were used to investigate the association between retinal age gap and mortality.

Results The DL model achieved a strong correlation of 0.81 ($p<0{\cdot}001$) between retinal age and chronological age, and an overall mean absolute error of 3.55 years. Cox regression models showed that each 1 year increase in the retinal age gap was associated with a 2% increase in risk of all-cause mortality (hazard ratio (HR)=1.02, 95% CI 1.00 to 1.03, p=0.020) and a 3% increase in risk of cause-specific mortality attributable to non-cardiovascular and non-cancer disease (HR=1.03, 95% CI 1.00 to 1.05, p=0.041) after multivariable adjustments. No significant association was identified between retinal age gap and cardiovascular- or cancer-related mortality.

Conclusions Our findings indicate that retinal age gap might be a potential biomarker of ageing that is closely related to risk of mortality, implying the potential of retinal image as a screening tool for risk stratification and delivery of tailored interventions.

Data availability statement

Data are available in a public, open access repository.

https://doi.org/10.1136/bjophthalmol-2021-319807

DATASET AND RETINAL IMAGE PROCESSING PIPELINE

UK Biobank retinal imaging dataset

The study utilizes retinal fundus images from the UK Biobank dataset, which contains a large population scale collection of retinal images along with extensive clinical and demographic information.

The dataset includes tens of thousands of participants, enabling the model to learn robust patterns linking retinal structure to biological aging and mortality risk.

To ensure the model generalizes well across populations, the trained model was also validated using external cohorts, confirming that retinal age predictions remain consistent across different imaging conditions and patient groups.

MODEL PIPELINE

Image Acquisition
Retinal fundus photographs are captured using specialized ophthalmic cameras that provide detailed visualization of retinal structures.

Image Resizing
All images are resized to a consistent resolution to allow efficient processing by the neural network.

Pixel Normalization
Pixel intensity values are normalized to reduce variations in brightness and contrast across images.

Color Normalization
Color variations introduced by different imaging devices or lighting conditions are corrected.

Data Augmentation
Various transformations such as rotations, flips, and brightness adjustments are applied to increase dataset diversity and reduce overfitting.

DEEP LEARNING ARCHITECTURE

The study uses a Convolutional Neural Network (CNN) to analyze retinal fundus images and estimate an individual's retinal age. CNNs are widely used in medical imaging because they automatically learn visual patterns from images.

As the image passes through the network:
Early layers: detect basic patterns such as edges and vessel boundaries.
Intermediate layers: capture retinal structures like vessel branching, tortuosity, and optic disc features.
Deeper layers: learn aging related biomarkers such as microvascular abnormalities and macular texture changes.

The model predicts retinal age, which is compared with chronological age to compute the Retinal Age Gap. Studies show that a larger positive age gap is associated with a higher risk of mortality, highlighting its potential as a non-invasive biomarker of biological aging.

Model Performance

Mean Absolute Error (MAE)
≈ 4.5 years in retinal age prediction

However, the model functions as a black box system. It predicts retinal age but does not explain which retinal regions or features influence the prediction, limiting clinical interpretability and motivating the need for explainable retinal aging models.

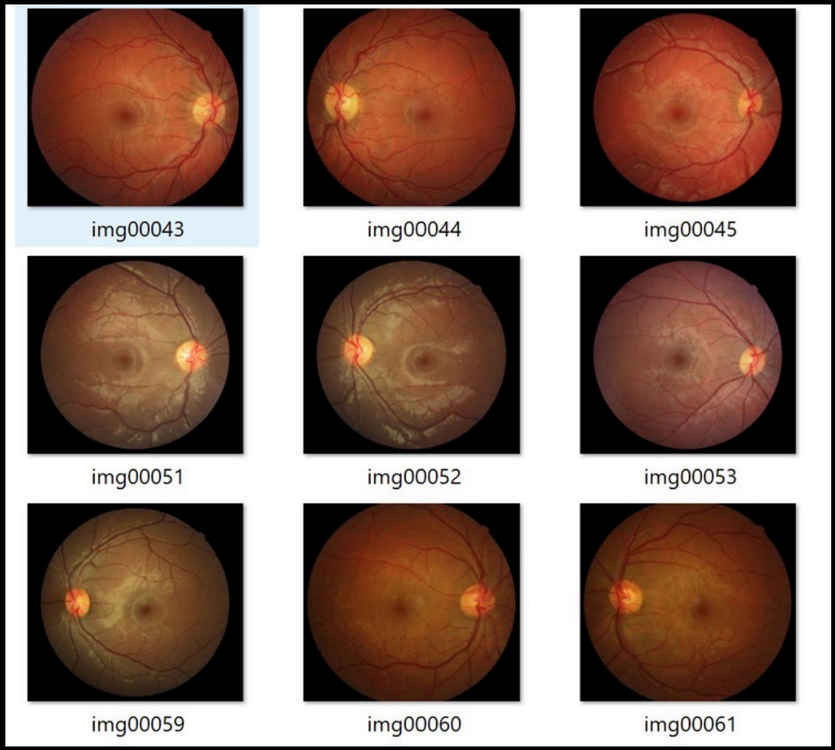

DATASETS

BRSET

The primary dataset used in this project is the BRSET dataset, a publicly available Brazilian retinal fundus image dataset collected for ophthalmological and clinical research purposes.

The dataset contains retinal fundus images along with clinical and demographic metadata obtained during medical examinations. Each image is associated with information such as patient age, gender, and medical annotations.

Dataset statistics:

Total patients: ~8,500

Total images: ~16,266

Each patient typically has two images (left and right eye)

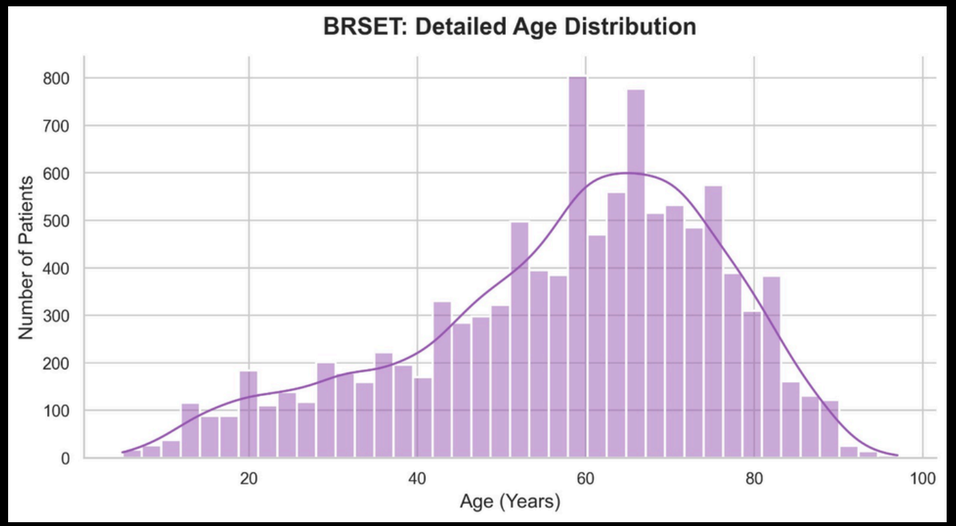

Observations:

→ Large and Diverse Retinal Fundus Dataset
→ Includes clinically relevant metadata
→ Variations in illumination and image quality
→ Suitable for retinal aging and biomarker analysis

DATASETS

ODIR-5K ~ Kaggle

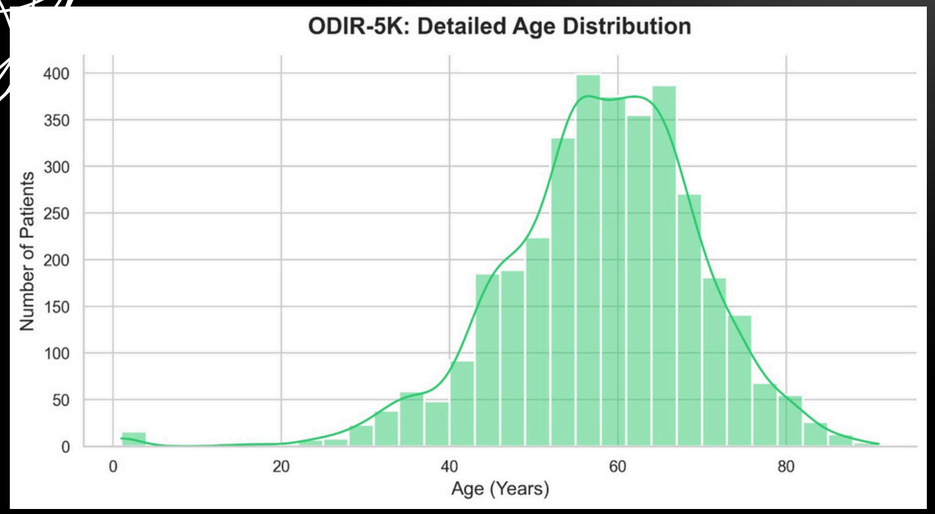

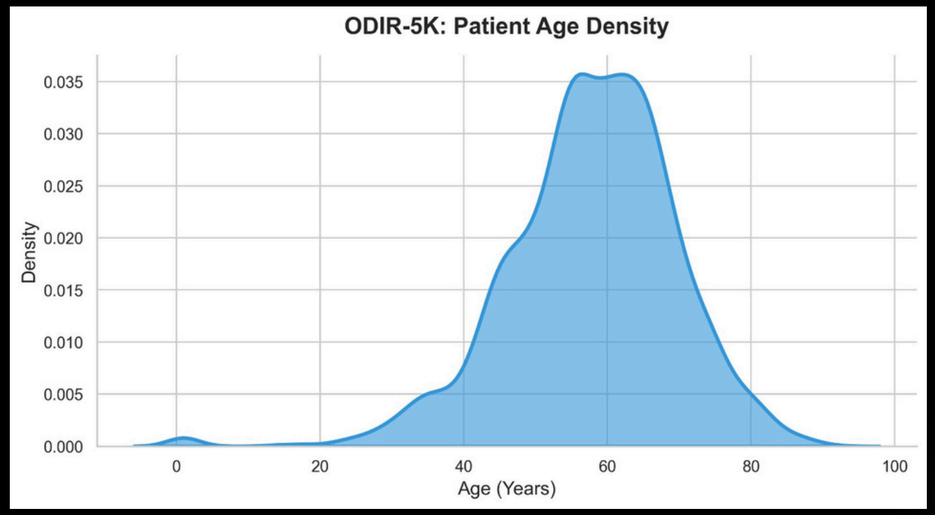

The second datasets we are using in this project is the ODIR-5K, a publicly available retinal fundus image dataset collected from clinical ophthalmology examinations.

The dataset contains retinal fundus images captured using standard fundus cameras and includes both left and right eye images of patients. Each image is associated with metadata such as patient ID, age, gender, and diagnostic labels.

Key dataset characteristics:

Approximately 3500 patients

Around 7000 retinal fundus images (left and right eyes)

The dataset is skewed toward middle aged and older individuals, which is expected because retinal diseases are more prevalent in older populations

Retinal images show variations in:

- illumination
- image quality
- retinal orientation
- disease conditions

DATASET PREPPROCESSING

BRSET Dataset Cleaning Pipeline

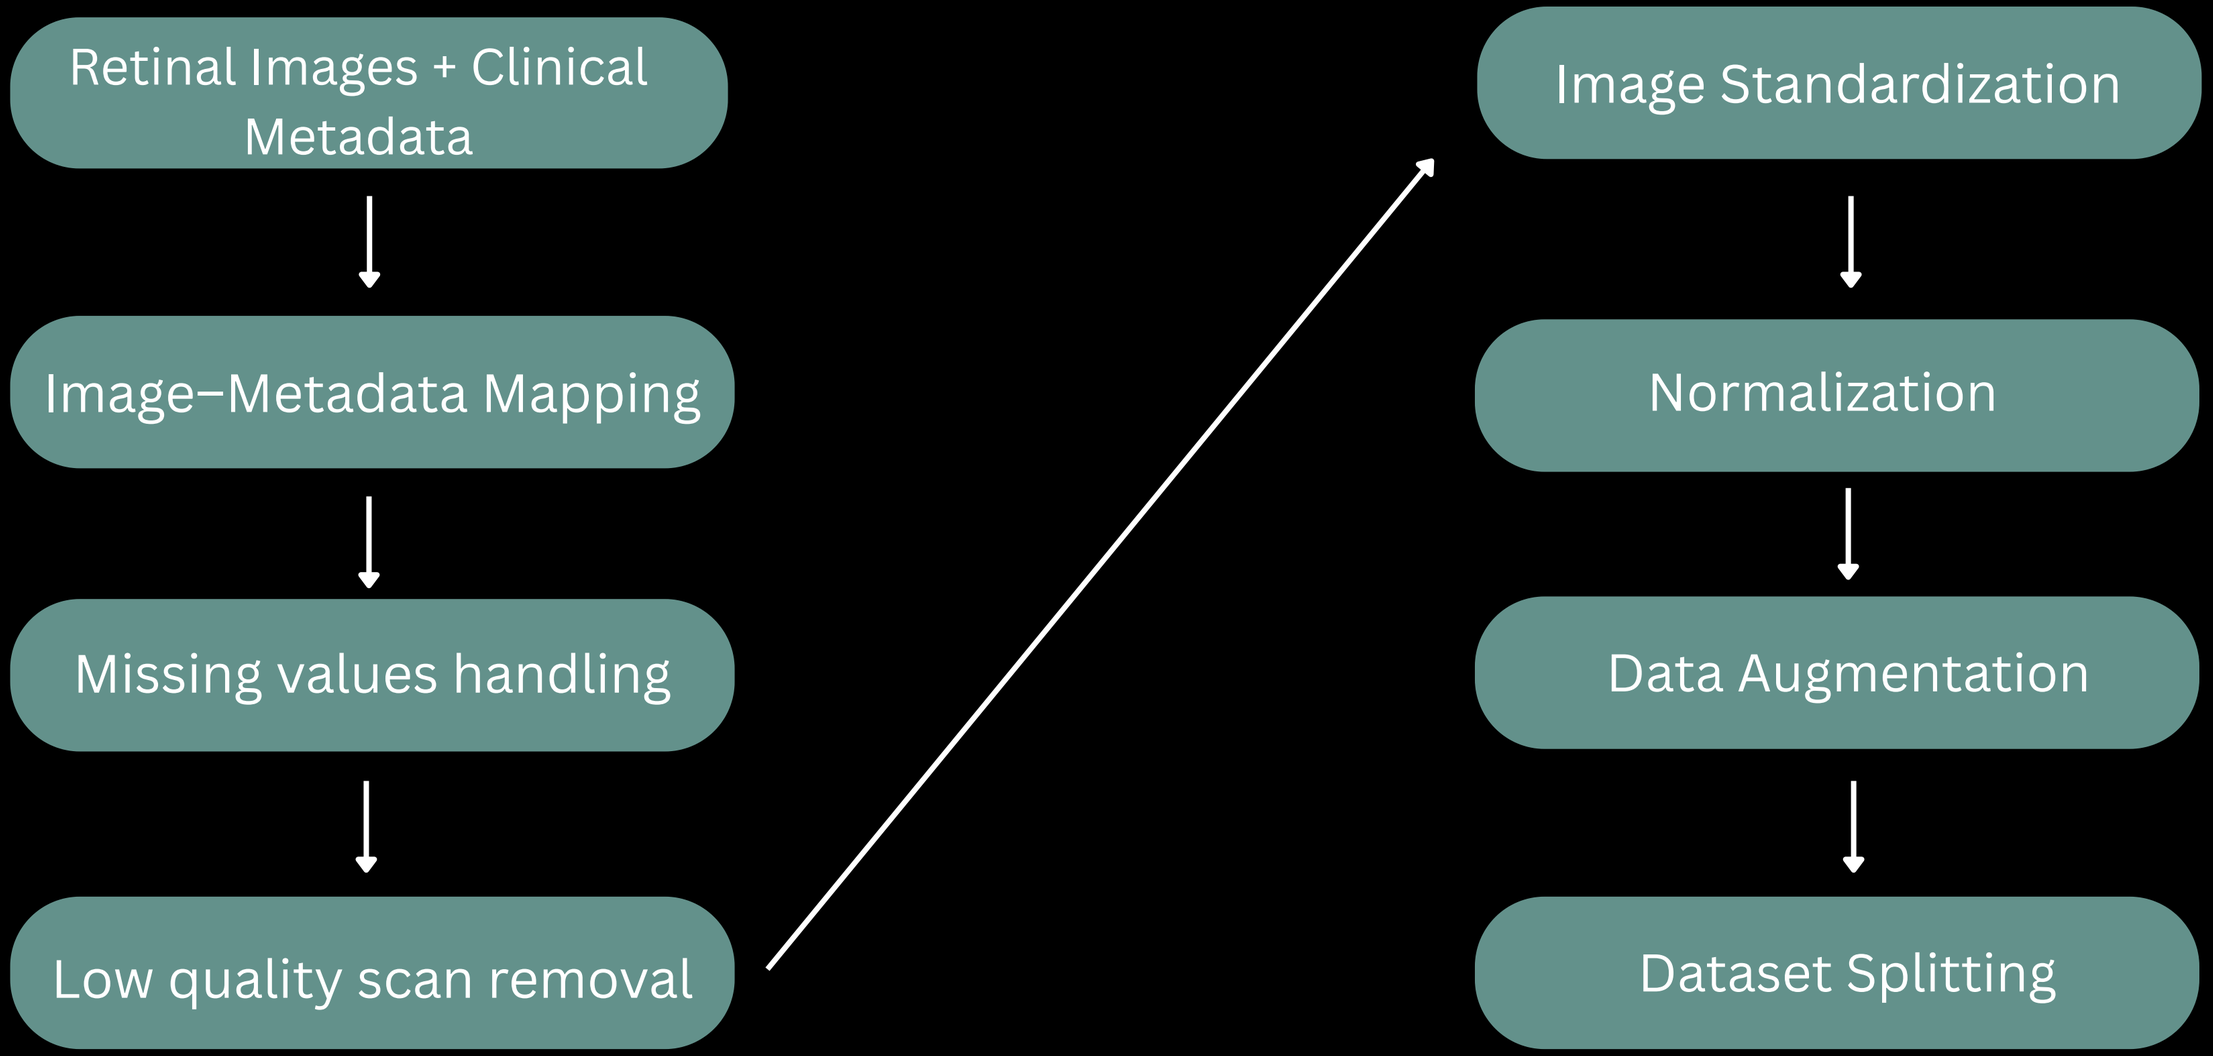

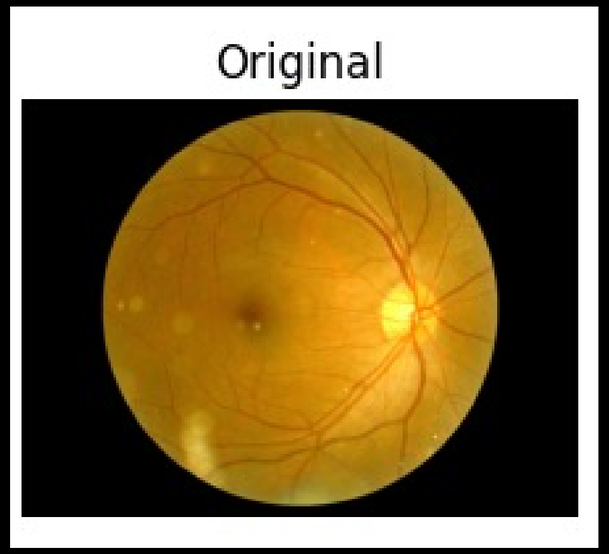

CLAHE

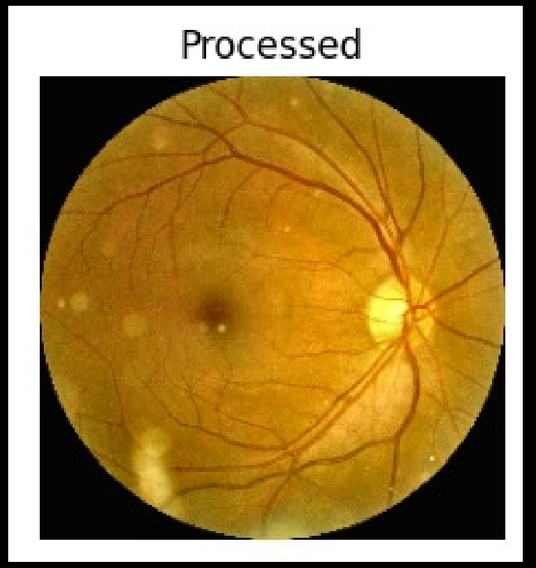

AUGMENTATION

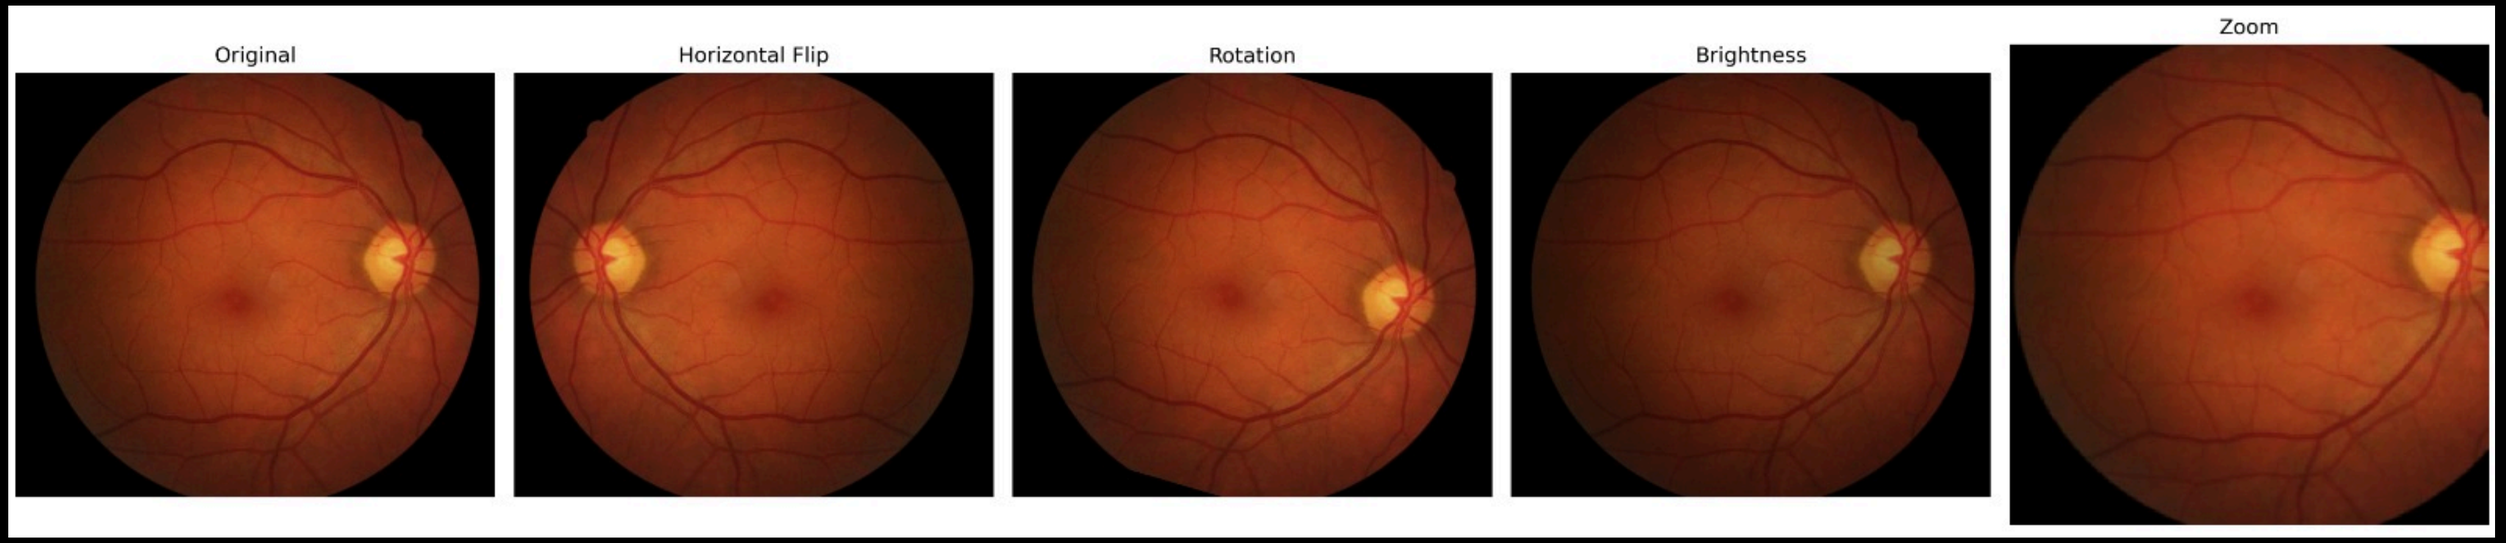

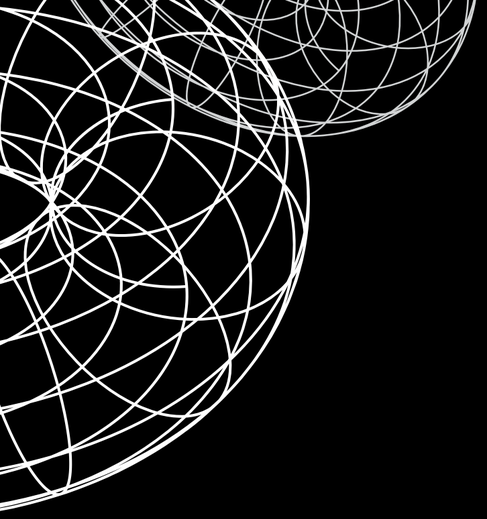

METHODOLOGY

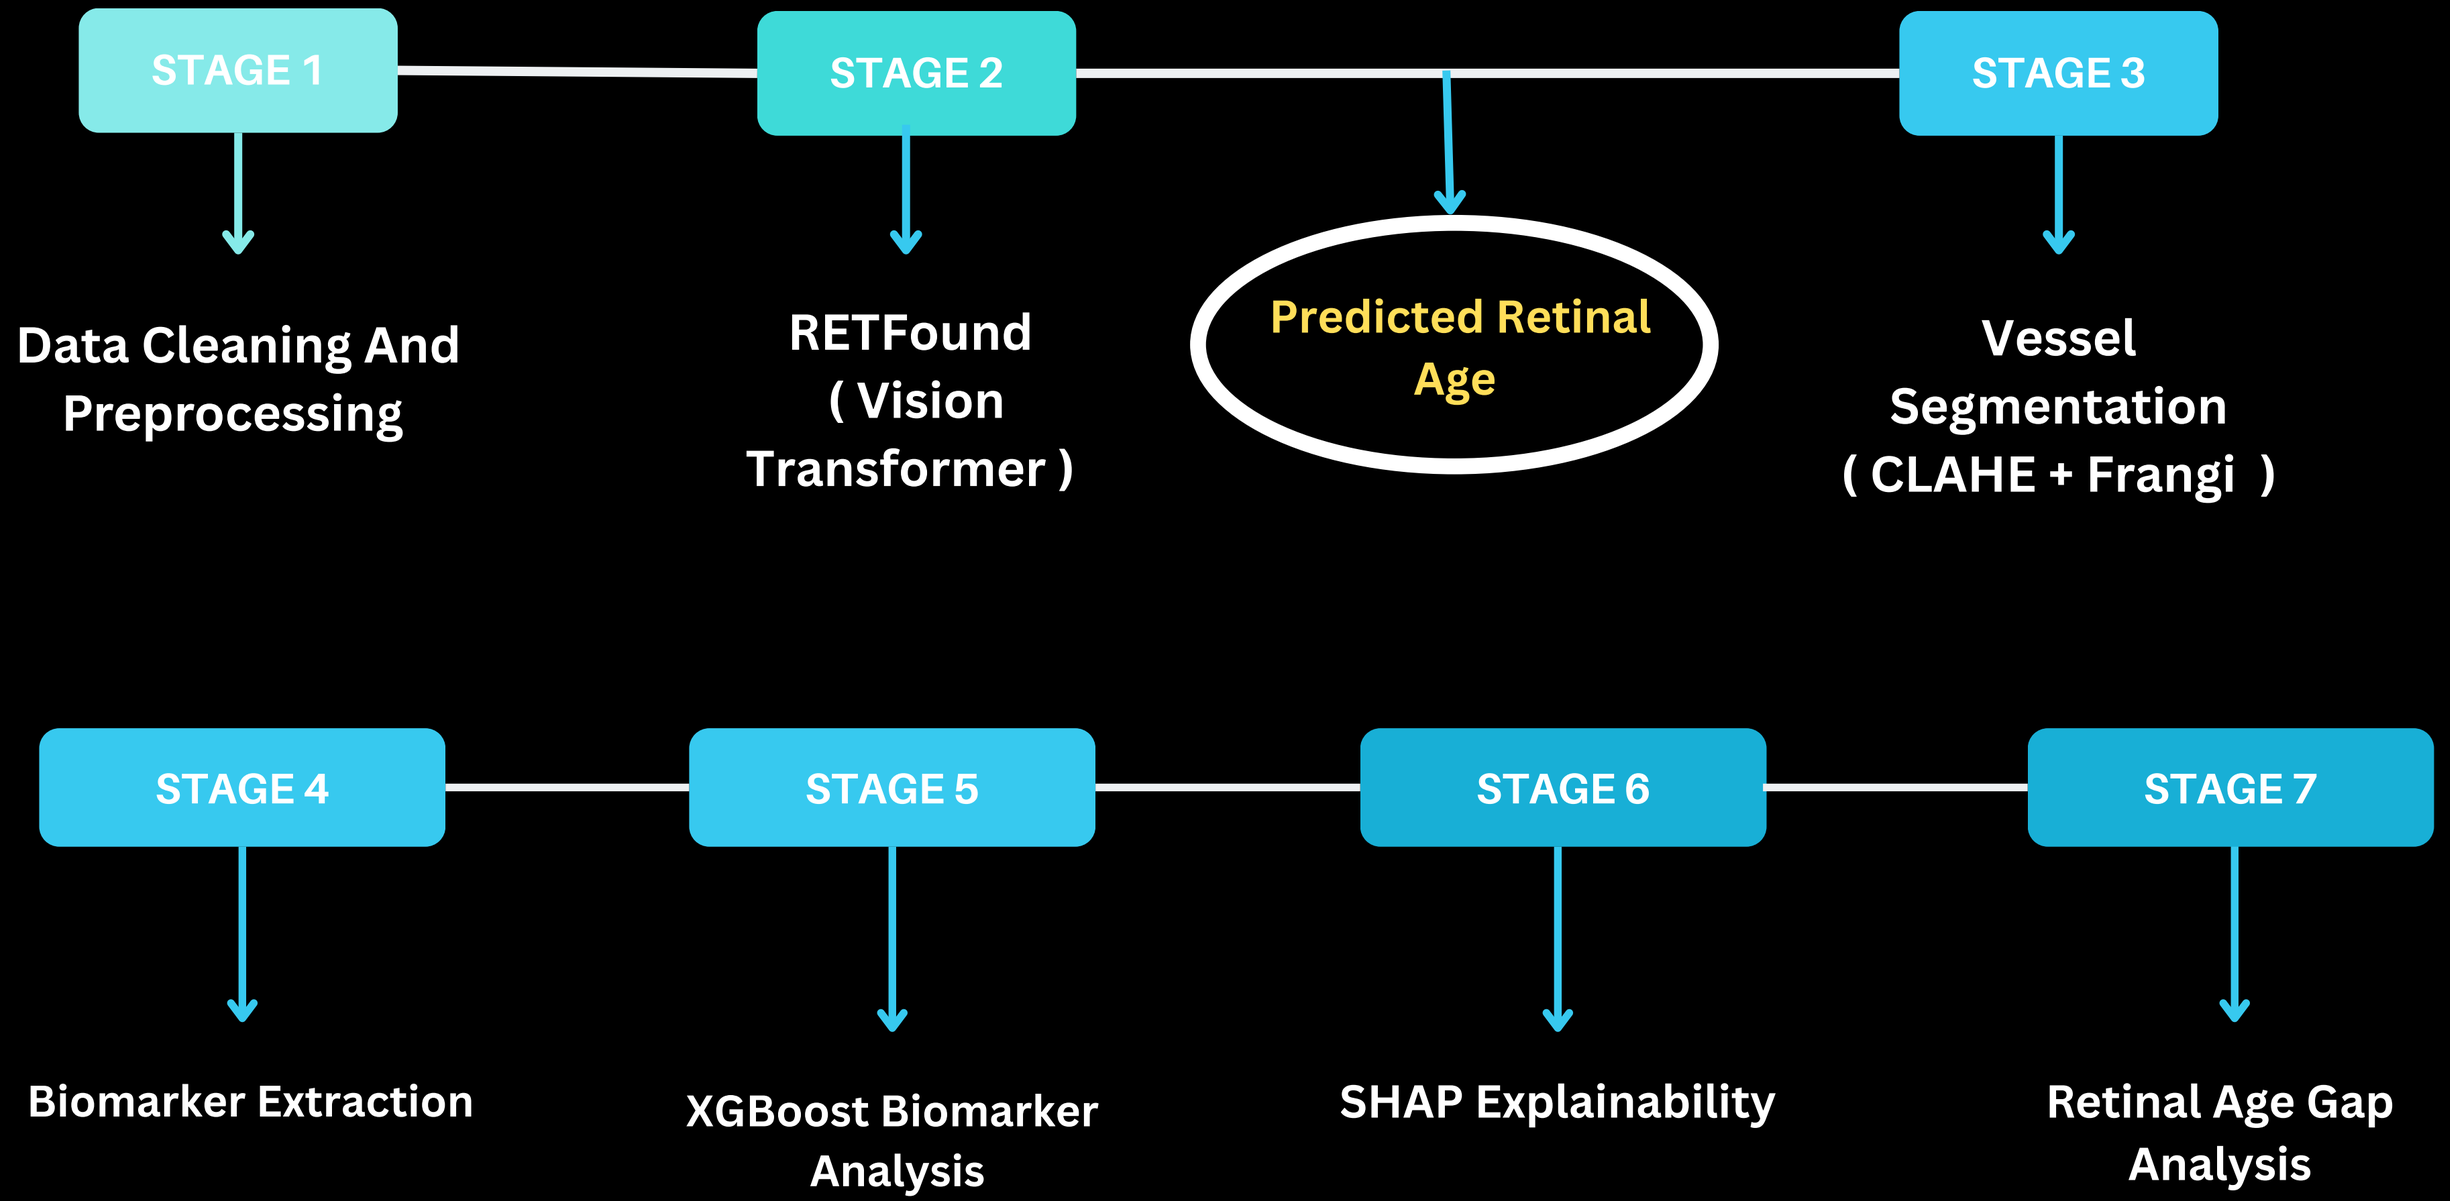
STAGE 1
STAGE 2
STAGE 3
Data Cleaning And Preprocessing
RETFound (Vision Transformer)
Predicted Retinal Age
Vessel Segmentation (CLAHE + Frangi)
STAGE 4
STAGE 5
STAGE 6
STAGE 7
Biomarker Extraction
XGBoost Biomarker Analysis
SHAP Explainability
Retinal Age Gap Analysis

RETFOUND VISION TRANSFORMER ARCHITECTURE

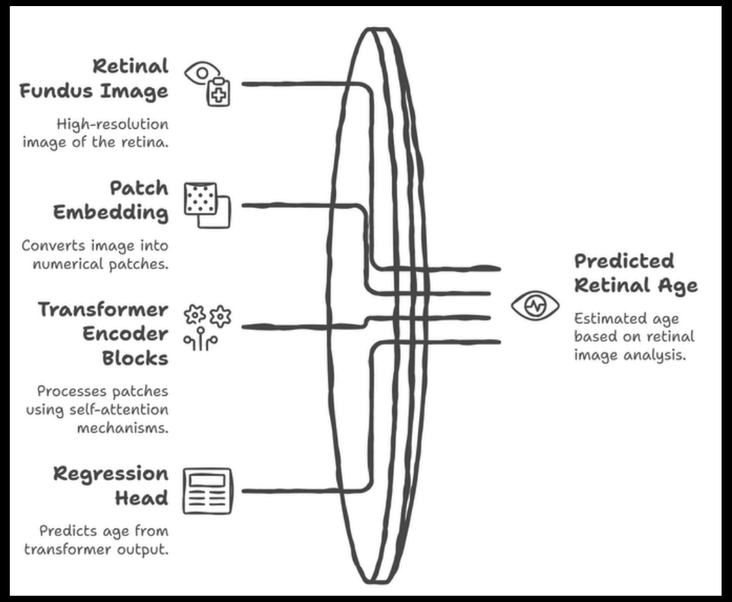

Model Architecture

- Vision Transformer (ViT)-based retinal foundation model
- Pretrained RETFound weights
- Input size: 224 × 224 retinal fundus images
- Fine-tuned for retinal age prediction

Training Objective

- Predict retinal biological age
- Learn global retinal vascular representations

Output

Predicted Retinal Age

Transformer attention captures complex retinal aging patterns.

VESSEL SEGMENTATION

- CLAHE enhances vessel visibility
- Frangi filtering extracts vascular structures
- Skeletonization preserves vessel topology

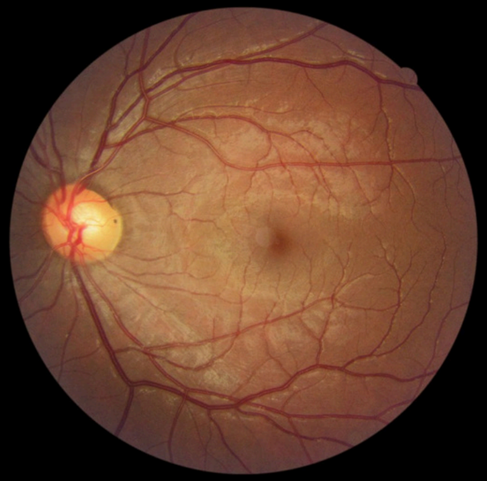

ORIGINAL IMAGE

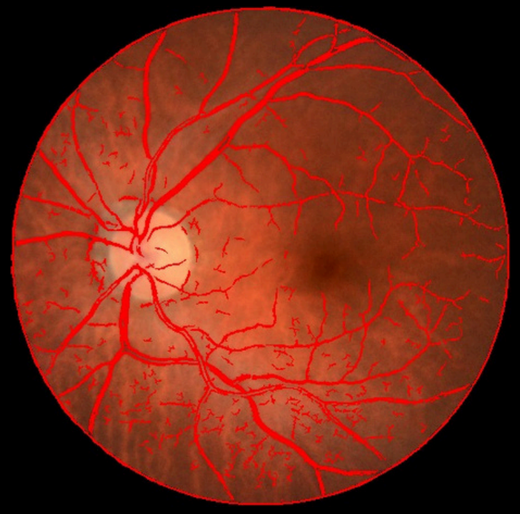

VESSEL SEGMENTATION
(CLAHE ENHANCEMENT
+
FRANGI VESSEL FILTERING)

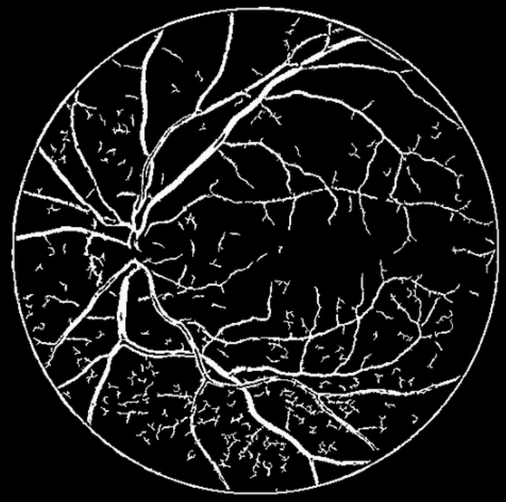

BINARY VESSEL MASK

VASCULAR BIOMARKER EXTRACTION

	image	vessel_density	vessel_length	branching_points	fractal_dimension	tortuosity	avg_thickness
2	1005_left_overlay.png	0.094844818	5791	2464	1.651847984	15.88573868	4.293386289
3	1005_left_vessel.png	0.080551147	11326	2178	1.538446665	22.96980362	1.864382836
4	1005_right_overlay.png	0.229026794	10236	4024	1.724639164	21.75903197	5.8653771
5	1005_right_vessel.png	0.086368561	12070	2519	1.551092753	24.44612428	1.875807788
6	1008_left_overlay.png	0.038921356	2486	811	1.483578394	6.282647419	4.104183427
7	1008_left_vessel.png	0.091217041	12317	2419	1.56053935	25.05806254	1.94138183
8	1008_right_overlay.png	0.016307831	1046	265	1.320477755	3.491927922	4.086998088
9	1008_right_vessel.png	0.090499878	11956	2436	1.552990297	24.31880278	1.984275677
10	1015_left_overlay.png	0.027088165	2478	931	1.484318621	7.499181184	2.865617433
11	1015_left_vessel.png	0.125087738	15648	4407	1.668882817	33.81243732	2.095539366
12	1015_right_overlay.png	0.050411224	3685	1317	1.681324771	15.05235608	3.586160109
13	1015_right_vessel.png	0.126663208	15986	4411	1.676333329	34.44610705	2.077067434
14	1018_left_overlay.png	0.006778717	330	117	1.36571723	5.240736739	5.384848485
15	1018_left_vessel.png	0.11240387	15792	4424	1.663608835	34.72279063	1.865881459
16	1018_right_overlay.png	0.015422821	961	336	1.560940634	6.962123196	4.207075963
17	1018_right_vessel.png	0.111995697	15496	4084	1.660262695	34.0550867	1.894617966
18	101_left_overlay.png	0.209701538	11241	3608	1.801882132	24.56352473	4.89031225
19	101_left_vessel.png	0.130268097	15573	4667	1.666769072	33.61233972	2.192833751
20	101_right_overlay.png	0.176151276	11122	3813	1.782930177	25.32801885	4.151861176

Segmented Vessel Mask
↓
Skeletonization & Morphological Analysis
↓
Quantitative Vascular Biomarkers

Extracted features:

Vessel Density : Vessel Coverage
Vessel Length : Total Vessel Span
Branching Points : Vascular Complexity
Fractal Dimension : Structural Irregularity
Tortuosity : Vessel Curvature
Avg Thickness : Vessel Caliber

SHAP BASED EXPLAINABILITY

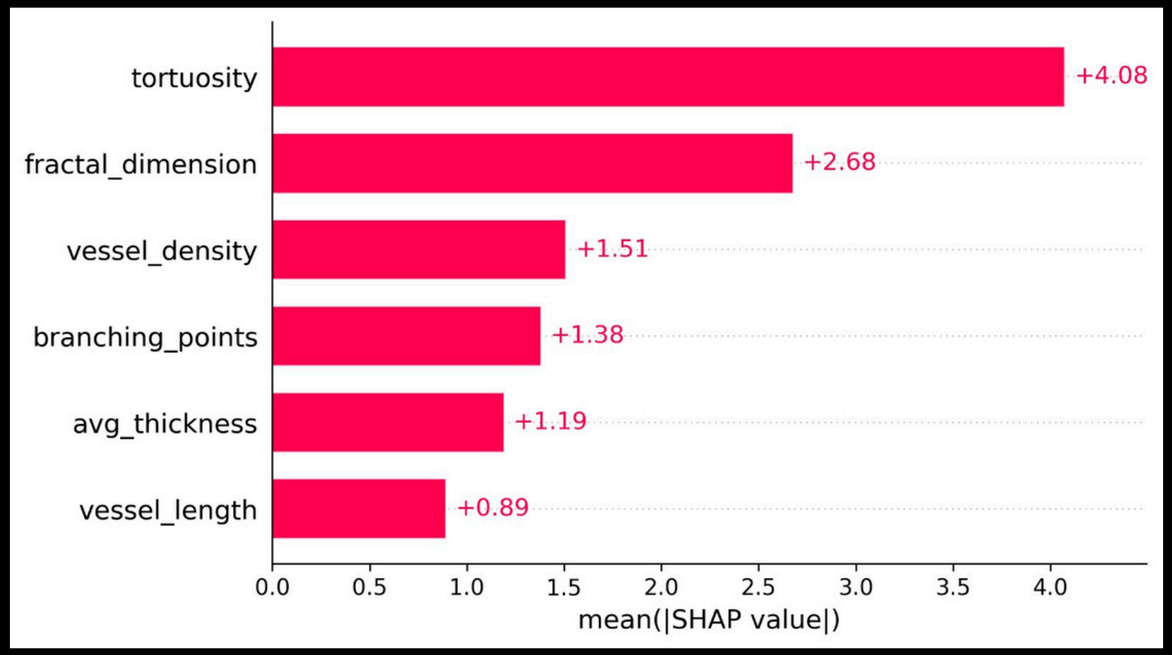

Feature Contribution Analysis

- Explains model decisions
- Quantifies feature importance
- Improves clinical interpretability

- Higher tortuosity increased retinal age prediction
- Lower vessel density contributed to vascular aging indication
- SHAP quantifies biomarker contributions to model predictions

SHAP transforms the model from a black-box predictor into an interpretable clinical tool.

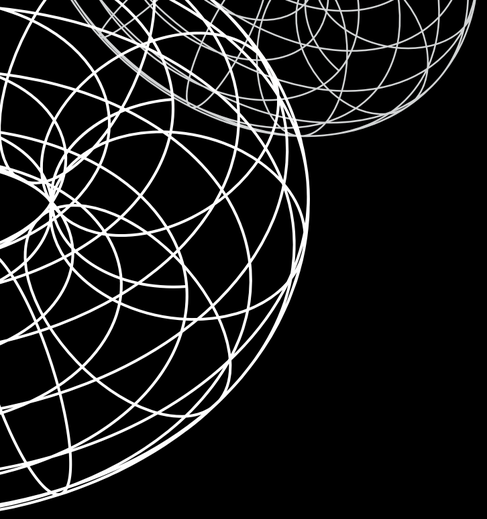

MODELS TRAINED

ATTEMPT 1- EFFICIENT NET

Key Features

- CNN-based deep learning architecture
- Efficient compound scaling of depth, width, and resolution
- Captures local retinal textures and vascular patterns
- Computationally efficient and lightweight

Why EfficientNet?

- Strong image feature extraction capability
- Faster training with fewer parameters
- Effective for retinal image analysis tasks

Limitations of EfficientNet

- CNNs primarily focus on local spatial features
- Limited ability to capture global retinal relationships
- Harder to interpret diffuse retinal aging patterns

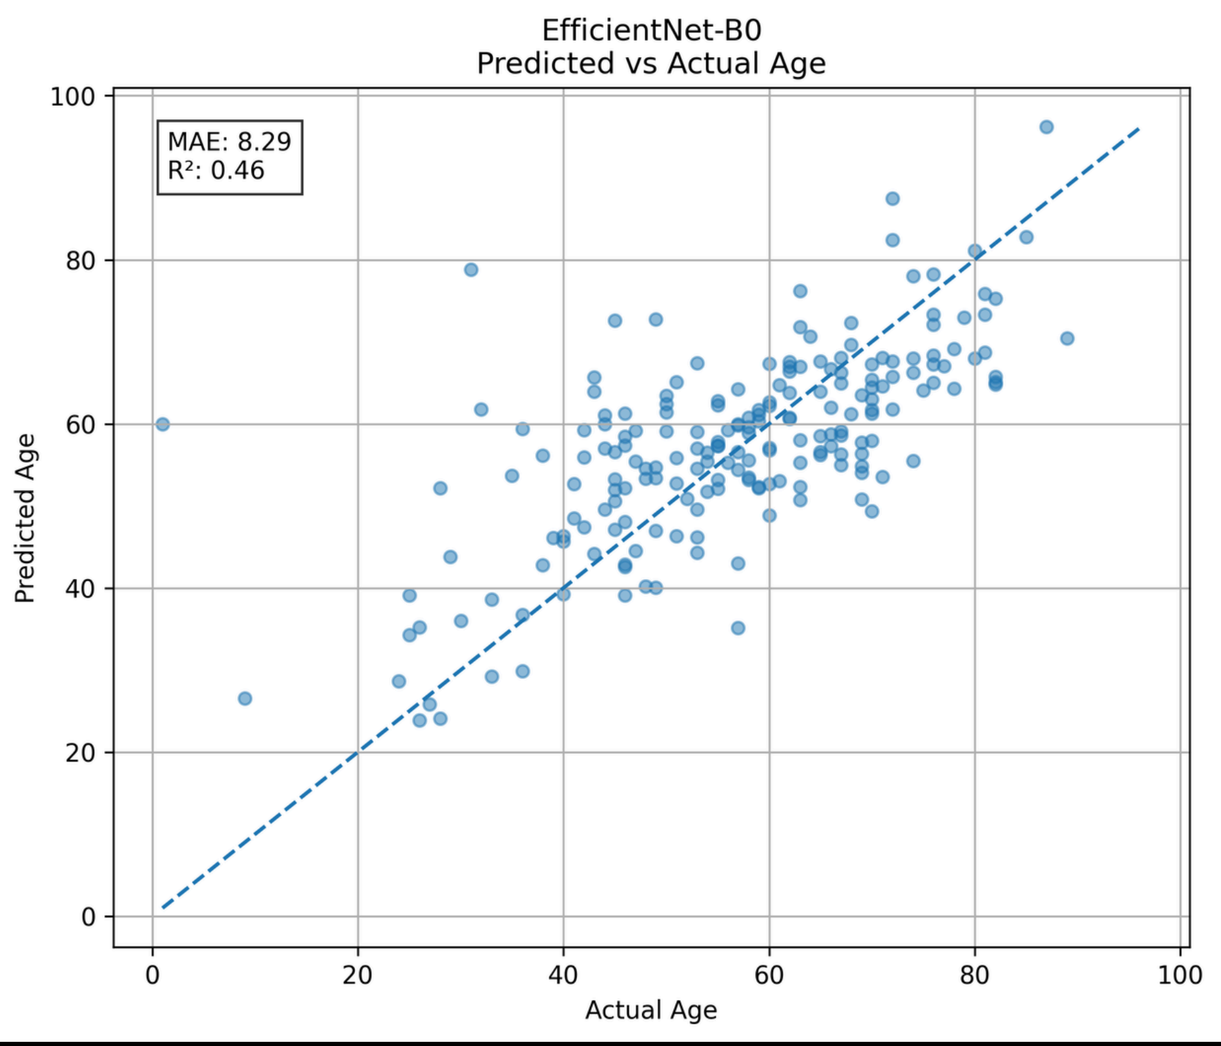

EfficientNet learns hierarchical retinal features associated with biological aging patterns.

ATTEMPT 1(B)-
B3 EFFICIENT NET

Why EfficientNet-B3?

- Deeper and wider architecture
- Better feature representation
- Improved retinal texture learning

Improvement Over B0

- Better MAE
- Improved retinal feature extraction
- Improved retinal feature representation

Limitation

- Still limited in modeling global retinal dependencies

B0 vs B3 mini comparison

Model	MAE
B0	8.29
B3	7.73

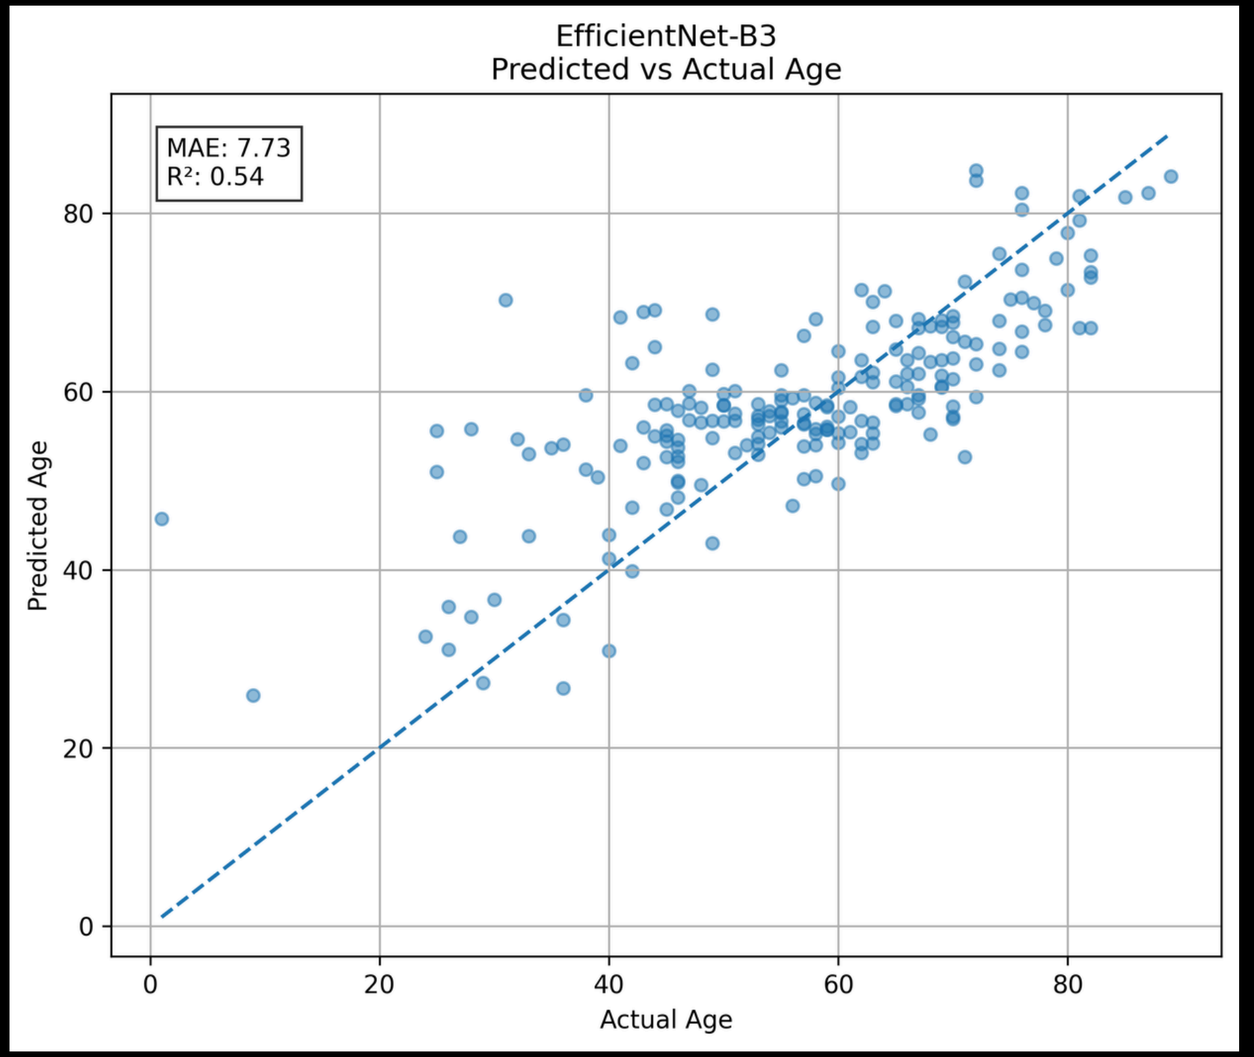

ATTEMPT 2-
RETFound Linear Probe

Linear Probe Strategy

- Pretrained RETFound encoder kept frozen
- Only final linear regression layer trained
- Preserves pretrained retinal representations
- Reduces overfitting on smaller medical datasets
- Computationally efficient adaptation strategy

Why Linear Probe?

- Faster training compared to full fine-tuning
- Leverages pretrained retinal feature knowledge
- Improves stability on limited retinal datasets

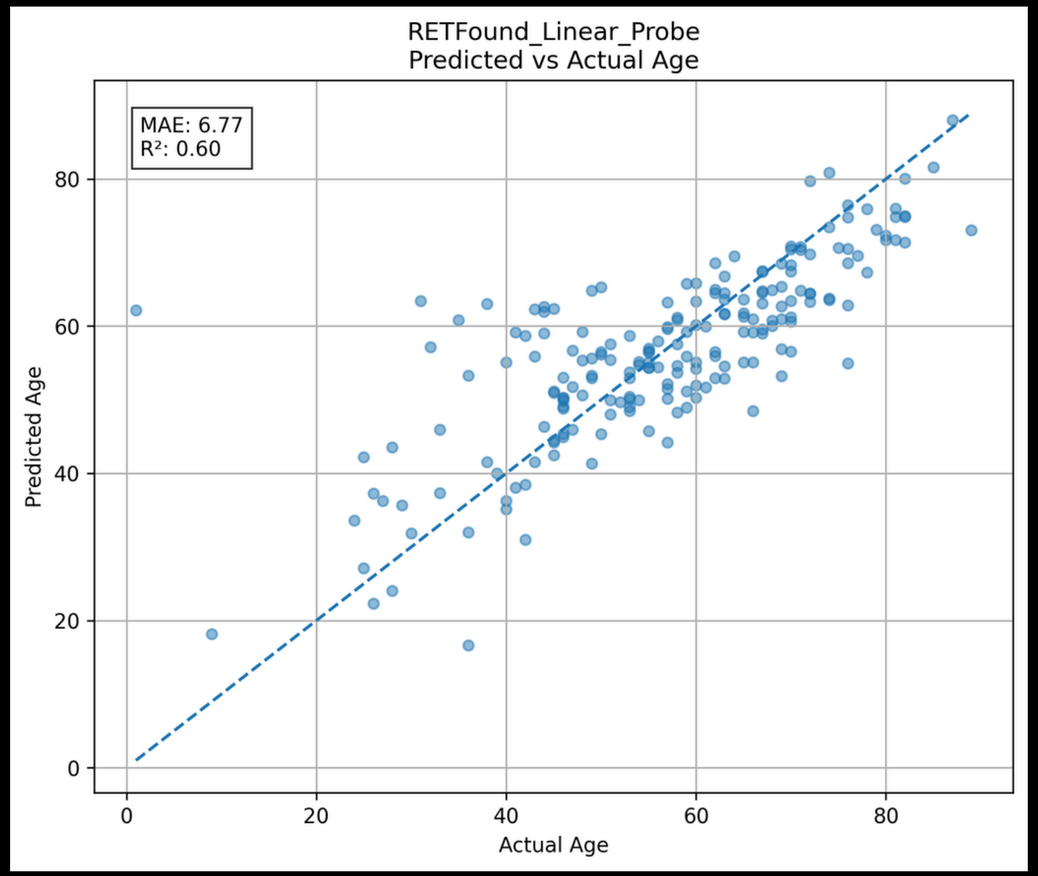

Linear probing enables efficient adaptation of pretrained RETFound representations for retinal age prediction tasks.

ATTEMPT 2(B)-
RETFOUND

Key Features

- Vision Transformer (ViT)-based architecture
- Pretrained retinal foundation model
- Learns global retinal dependencies using self-attention
- Strong contextual representation learning

Why RetFound?

- Captures long range retinal relationships
- Better understanding of complex vascular structures
- Enhanced contextual understanding of retinal vascular structures

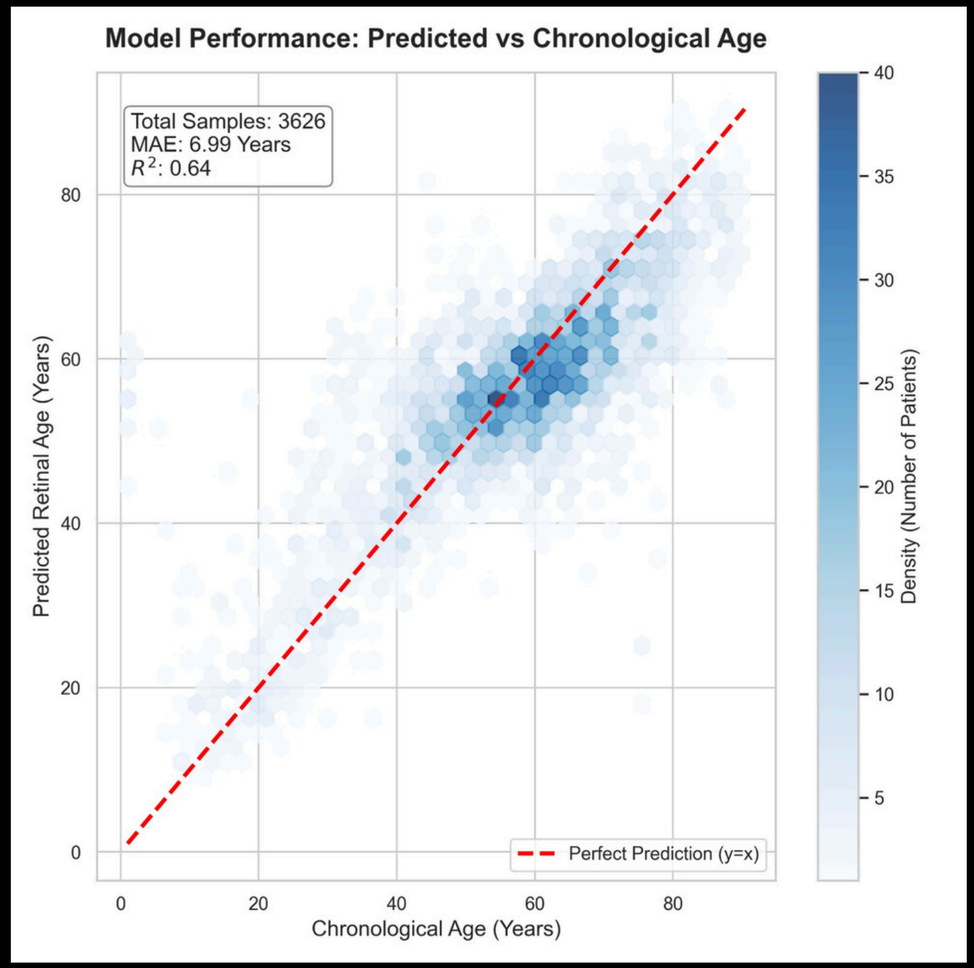

RETFound fine-tuning improved contextual retinal representation learning for vascular aging analysis.

MODEL COMPARISON

Model	Architecture	MAE	Key Strength
EfficientNet-B0	CNN	8.29	Strong local retinal feature learning
EfficientNet-B3	CNN	7.73	Improved retinal feature representation
RETFound Linear Probe	Vision Transformer (ViT)	6.77	Best numerical performance with reduced overfitting
RETFound Fine-Tuned	Vision Transformer (ViT)	6.99	Better contextual retinal representation learning

RETFound-based architectures demonstrated superior retinal representation learning compared to CNN-based approaches. Linear probing achieved the lowest MAE, while full fine-tuning enabled richer contextual retinal understanding for interpretability analysis.”

END-TO-END RETINAL AGING ANALYSIS

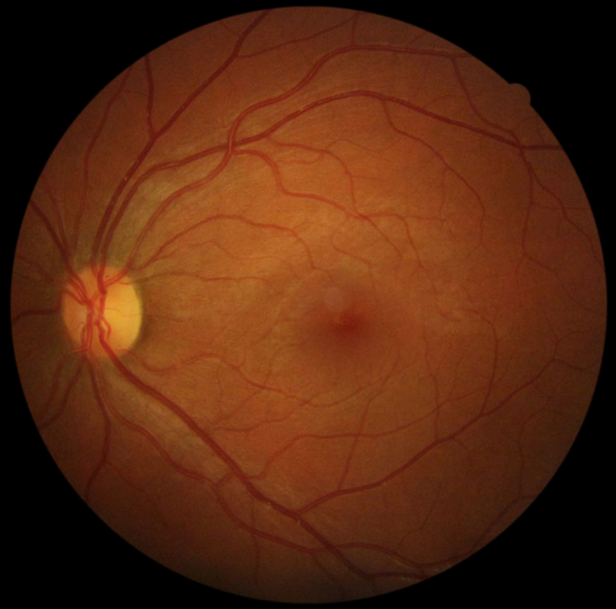

ORIGINAL IMAGE

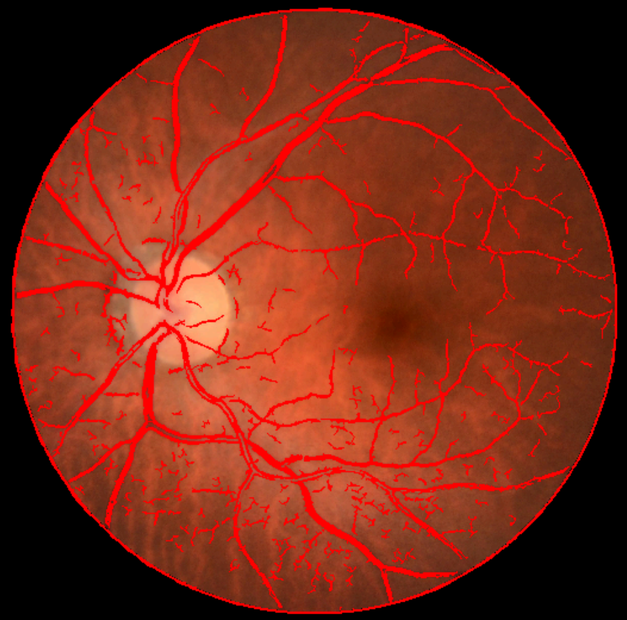

VESSEL SEGMENTATION
(CLAHE ENHANCEMENT
+
FRANGI VESSEL FILTERING)

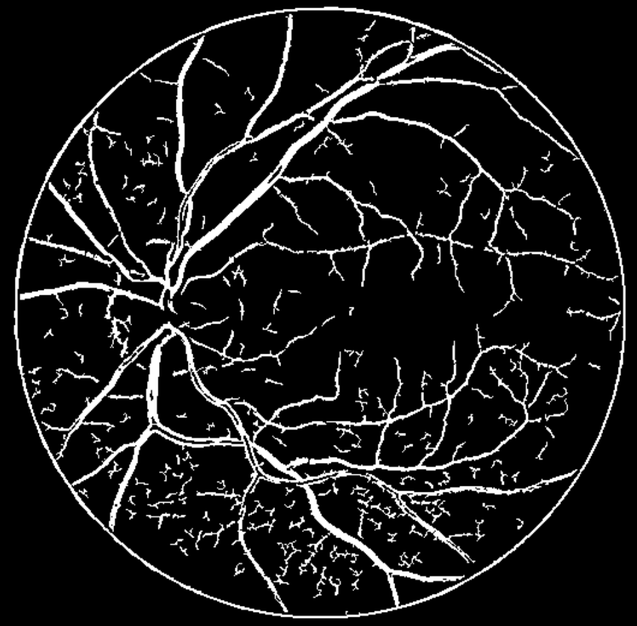

BINARY VESSEL MASK

✅ Actual Age: 51.0
✅ Predicted Age: 57.4
✅ Retinal Age Gap: 6.4

📊 BIOMARKER ANALYSIS

	Biomarker	Value	Status	SHAP Contribution	Effect
0	vessel_density	0.0798	Low	2.71	Increased Age
1	vessel_length	12838.0000	Normal	-1.43	Decreased Age
2	branching_points	3077.0000	Normal	-1.58	Decreased Age
3	fractal_dimension	1.6027	Normal	-2.23	Decreased Age
4	tortuosity	26.2535	Normal	-2.74	Decreased Age
5	avg_thickness	1.6293	Low	5.56	Increased Age

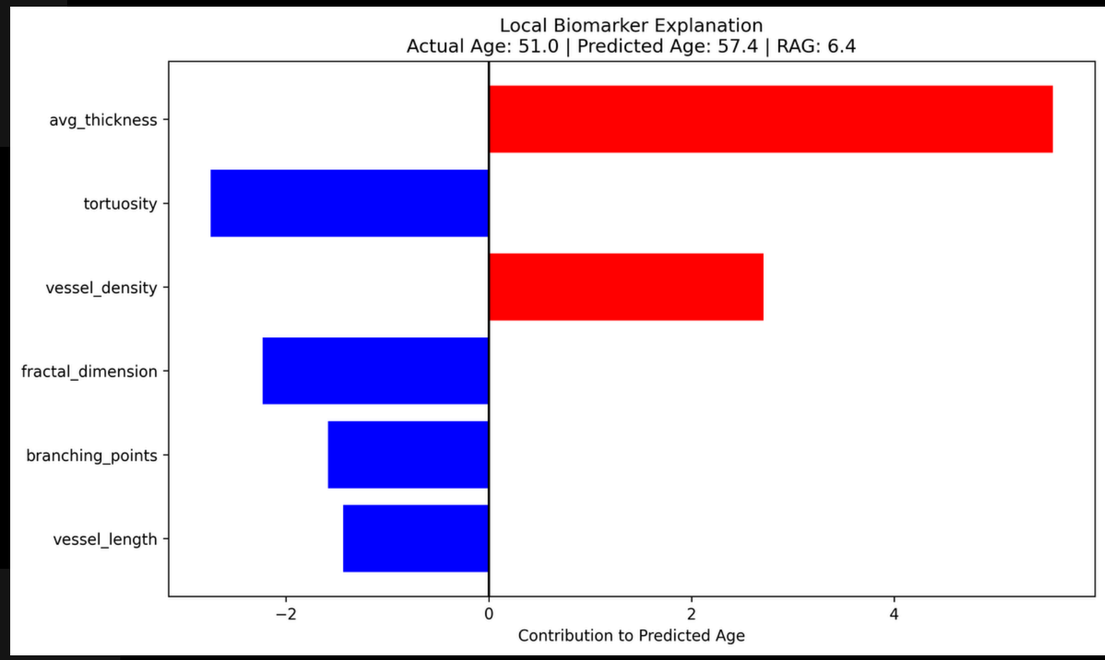

🩺 MEDICAL INTERPRETATION

Retinal age appears significantly elevated relative to chronological age.
vessel_density (Low) contributed to increased retinal age prediction.
avg_thickness (Low) contributed to increased retinal age prediction.

Limitations of our Model:

- Biomarker extraction depends on segmentation quality
- Simplified handcrafted vascular feature estimation
- Variability in retinal image quality and illumination conditions
- Vessel segmentation noise affecting biomarker extraction accuracy
- Model generalization to real-world clinical settings requires further validation

THANK YOU