Feature Space Translation Framework for Cross-Device Alignment and Derivation of Clinically Relevant Retinal Biomarkers from Foundation Model Embeddings

Sparsh Rastogi [† 1 2] **Akshat Bakshi** [1] **Jidugu Sriharisesh** [1] **Imran Razzak** [3] **Chinmaya Panigrahy** [1] **Jatin Bedi** [1] **Sahil Thakur** [4 5]

Abstract

Oculomics enables non-invasive assessment of systemic diseases through retinal imaging, offering a portable and cost-effective alternative to traditional diagnostic methods. However, its clinical deployment is hindered by device-induced variability, which limits generalization across heterogeneous imaging systems. Although retinal foundation models have recently achieved strong performance across diverse downstream tasks, their robustness across heterogeneous imaging devices remains limited, and their latent representations do not directly correspond to clinically meaningful retinal biomarkers. To address these challenges, we propose a unified feature-space translation framework based on tabular learning that learns feature-space alignments between heterogeneous imaging devices and retinal foundation-model representations, enabling cross-device retinal biomarker translation and clinically meaningful biomarker estimation from foundation-model embeddings in both same-device and cross-device settings. Experiments on the AI-READI dataset demonstrate accurate cross-device retinal biomarker alignment, recovery of clinically meaningful biomarkers from RETFound embeddings, and robust cross-device biomarker estimation, achieving up to 78.5%, 87.5%, and 76.0% R^2, respectively. These findings establish feature-space translation as an effective strategy for mitigating device-induced variability while improving the interpretability and clinical utility of retinal foundation models without requiring architectural modifications.

†Work conducted independently of Meesho. [1]Thapar Institute of Engineering and Technology, Patiala, India [2]Meesho, Bengaluru, India [3]Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), Abu Dhabi, UAE [4]Mediwhale, Seoul, South Korea [5]Singapore Eye Research Institute, Singapore. Correspondence to: Sparsh Rastogi <srastogi_be22@thapar.edu>, Sahil Thakur <sahil.thakur@mediwhale.edu>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

1. Introduction

Oculomics is a field that leverages retinal imaging to enable the non-invasive assessment of systemic disease burden by capturing microvascular changes through structural biomarkers (Irodi et al., 2026; Zhu et al., 2025; Sung et al., 2026). The retina is the only external organ in the human body where microvasculature can be directly visualized non-invasively (Inamullah et al., 2025; Honavar, 2022), providing a unique view of blood vessels that feed blood to the heart, kidneys, brain and other vital organs. Changes in retinal biomarkers such as vessel tortuosity, fractal dimension, vessel density, and the arteriovenous ratio (AVR) have demonstrated strong diagnostic relevance for chronic conditions (Estrada et al., 2015; Huang & Dashtbozorg, 2018; Srinidhi et al., 2019; Wang et al., 2026), including chronic kidney disease (Sasongko et al., 2012; Tan et al., 2024a), type 2 diabetes, cardiovascular disease (Wong et al., 2002; Cheung et al., 2007; Rastogi et al., 2025), and neurological disorders (Liew et al., 2021; Blose & Silverstein, 2026; Chan et al., 2026). These biomarkers provide a unique window into systemic physiology, as retinal vasculature shares anatomical and functional similarities with other microvascular beds, making it particularly suitable for early disease detection and risk stratification (Xu & Trucco, 2026; Rastogi et al., 2025; Majithia et al., 2023; Lim et al., 2020; Tan et al., 2024b; Gupta et al., 2020; Cheung et al., 2020).

Recent advances in deep learning, along with standardized extraction pipelines such as AutoMorph (Zhou et al., 2022), VAMPIRE, IVAN, SIVA (McGrory et al., 2018) have enabled large-scale, automated, and reproducible quantification of these retinal features from fundus images (Fraz et al., 2015; Perez-Rovira et al., 2011). In parallel, retinal foundation models (FMs) pretrained on large-scale datasets (Gharbi et al., 2026), such as RETFound (Zhou et al., 2023), RetFound Green (Engelmann & Bernabeu, 2025), VisionFM

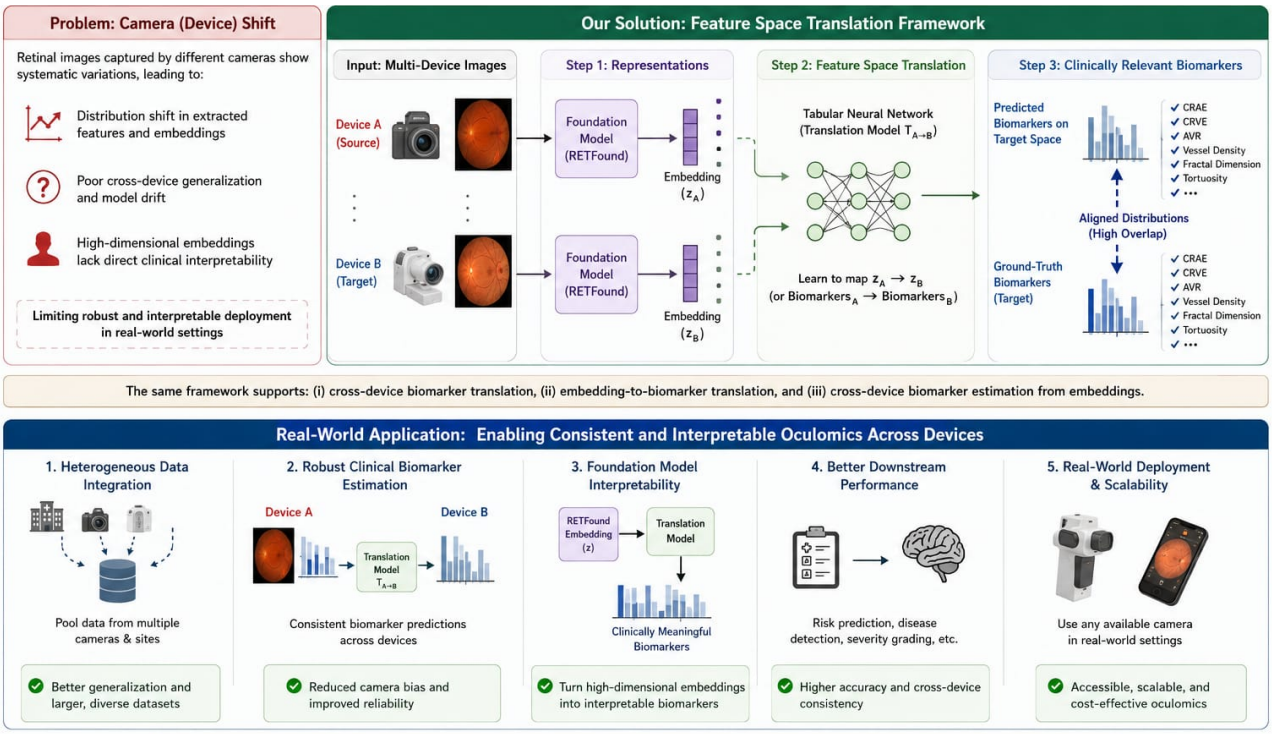

Figure 1. Overview of the proposed feature space translation framework. The training pipeline consists of three tasks: (i) cross-device biomarker translation, (ii) embedding-to-biomarker translation, and (iii) cross-device biomarker estimation from embeddings. The framework operates on retinal fundus images from multiple devices and is evaluated using both prediction accuracy and distributional alignment metrics.

(Qiu et al., 2023; 2024), EyeCLIP (Shi et al., 2025) have been proposed, leveraging extensive pretraining to shift the paradigm from handcrafted feature extraction to representation learning. These models learn rich, hierarchical representations that capture complex retinal patterns and achieve strong performance across diverse downstream tasks.

Despite this progress, the deployment of these approaches in real-world clinical settings remains limited due to two major constraints. First, out-of-distribution (OOD) shifts arising from camera and device heterogeneity introduce systematic variability in image characteristics and extracted features, thereby reducing robustness and leading to model drift across different devices, populations and datasets (Futoma et al., 2020; Cunnac et al., 2026). Although several studies have demonstrated robust generalization across population-level distribution shifts (Thakur et al., 2025; Tseng et al., 2023; Joo et al., 2023; Rim et al., 2021; Ting et al., 2017; Sabanayagam et al., 2020; Christopher et al., 2018), relatively limited work has focused on addressing distribution shifts introduced by differences in imaging devices and acquisition protocols (Rastogi & Thakur, 2025; Chen et al., 2024; Guo et al., 2024), which remain a major barrier to the development of generalizable and clinically deployable oculomic models (Cunnac et al., 2026; Kubin et al., 2024; Vilela et al., 2024; Abay et al., 2025; Ahn & Kim, 2024; Panwar et al., 2015). Second, despite their strong representational capacity, the high-dimensional latent embeddings produced by these models lack direct correspondence with clinically interpretable retinal biomarkers, limiting their transparency, clinical validation, and broader adoption in practice (An et al., 2025; Mensah et al., 2026).

To overcome these limitations, we propose a feature space translation framework that leverages tabular neural networks to map retinal biomarkers from the feature space of one imaging device to another, thereby mitigating device-induced variability and camera drift. Rather than introducing a new architectural design, our approach focuses on learning cross-device correspondences in feature space to improve robustness across heterogeneous imaging systems. Furthermore, we extend this framework to operate on the latent embeddings of retinal foundation models, such as RETFound (Zhou et al., 2023) to translate high-dimensional representations into clinically meaningful retinal biomarkers. In addition, the same formulation enables cross-device biomarker estimation directly from embeddings, allowing representations derived from one device to be calibrated and interpreted in the context of another. By decoupling representation learning from device-specific biases, our approach provides a simple yet effective way to improve the consistency, interpretability, and cross-device applicability

of oculomics pipelines. The key contributions of this work could be summarized as follows:

- We present a unified feature-space translation framework that aligns retinal biomarkers across imaging devices, translates latent representations from retinal foundation models into clinically meaningful biomarkers, and enables cross-device biomarker estimation directly from foundation-model embeddings.
- We benchmark a diverse set of state-of-the-art tabular neural networks alongside classical machine learning methods for feature-space translation, providing a comprehensive evaluation of their effectiveness for cross-device retinal biomarker alignment.
- Our framework is model-agnostic and can be seamlessly integrated into existing oculomics pipelines without architectural modifications, improving cross-device consistency and biomarker interpretability with minimal computational overhead.

2. Methodology

Problem Formulation. Let $\mathcal{D} = \{d_1, \ldots, d_K\}$ denote a set of imaging devices. For each device $d \in \mathcal{D}$, let $X^{(d)} \in \mathbb{R}^p$ represent the structured retinal biomarker space (e.g., AutoMorph features), and $E^{(d)} \in \mathbb{R}^q$ denote the corresponding latent representation obtained from a retinal foundation model such as RETFound. We aim to learn a set of parametric mappings that enable (i) cross-device biomarker translation (Eq. 1), (ii) embedding-to-biomarker translation within a device (Eq. 2), and (iii) cross-device biomarker estimation directly from embeddings (Eq. 3).

$$f_{d_i \to d_j} : \mathbb{R}^p \to \mathbb{R}^p, \quad X^{(d_j)} \approx f_{d_i \to d_j}(X^{(d_i)}) \tag{1}$$

$$g_d : \mathbb{R}^q \to \mathbb{R}^p, \quad X^{(d)} \approx g_d(E^{(d)}) \tag{2}$$

$$h_{d_i \to d_j} : \mathbb{R}^q \to \mathbb{R}^p, \quad X^{(d_j)} \approx h_{d_i \to d_j}(E^{(d_i)}) \tag{3}$$

where $f_{d_i \to d_j}$, g_d, and $h_{d_i \to d_j}$ correspond to tasks (i), (ii), and (iii), respectively. All mappings are learned using tabular predictive models by minimizing supervised reconstruction objectives over paired samples. For a given mapping $\phi \in \{f_{d_i \to d_j}, g_d, h_{d_i \to d_j}\}$, we define a general objective (Eq. 4) over input–target pairs $(u, v) \sim \mathcal{P}$, where $\ell(\cdot, \cdot)$ denotes a task-specific loss function.

$$\mathcal{L}(\phi) = \mathbb{E}_{(u,v)\sim\mathcal{P}} \left[\ell\big(v, \phi(u)\big)\right] \tag{4}$$

$$\mathcal{L}_{\text{NLL}}(\phi) = -\mathbb{E}_{(u,v)\sim\mathcal{P}} \left[\log p_\phi\big(v \mid u\big)\right] \tag{5}$$

Under a probabilistic formulation, we model the conditional distribution $p_\phi(v \mid u)$ and minimize the negative log-likelihood (Eq. 5). Assuming a Gaussian likelihood with mean $\phi(u)$ and fixed variance, this reduces to a mean squared error objective. To account for distributional discrepancies induced by cross-device variability, we additionally consider a Wasserstein-distance-based objective (Eq. 6) that promotes alignment between the predicted and target distributions. The overall optimization problem is defined in Eq. 7, where the mappings are learned over paired samples. This formulation provides a unified framework for addressing device-induced variability while improving the interpretability of foundation model representations.

$$\mathcal{L}_{\text{W}}(\phi) = \mathbb{E}_{(u,v)\sim\mathcal{P}} \left[W\big(\delta_{\phi(u)}, \delta_v\big)\right] \tag{6}$$

$$\min_{\phi} \ \mathcal{L}(\phi), \quad \phi \in \{f_{d_i \to d_j}, g_d, h_{d_i \to d_j}\} \tag{7}$$

Data Curation. We use retinal fundus images from the AI-READI dataset (AI-READI Consortium, 2024b;a), comprising 1,067 participants imaged using multiple non-mydriatic fundus cameras, primarily the Topcon Triton and Topcon Maestro2. The paired multi-device acquisitions enable supervised learning of cross-device feature mappings while capturing realistic acquisition variability. For each image, we extract 12 clinically relevant retinal vascular biomarkers using AutoMorph, including measures of vessel caliber (e.g., CRAE, CRVE, and AVR), vascular complexity (e.g., fractal dimension and vessel density), and geometric properties such as vessel tortuosity. We additionally extract high-dimensional latent representations using the pretrained RETFound model, yielding complementary structured biomarker vectors and foundation-model embeddings for each sample.

Framework Overview. The overall pipeline is illustrated in Fig. 1. Our framework comprises three complementary feature-space translation tasks: (i) cross-device biomarker translation by learning mappings between AutoMorph-derived biomarkers extracted from paired Triton and Maestro2 images; (ii) embedding-to-biomarker translation by learning device-specific mappings from RETFound embeddings to clinically meaningful retinal biomarkers; and (iii) cross-device biomarker estimation by predicting target-device biomarkers directly from source-device RETFound embeddings. We investigate these tasks using eight tabular predictive models spanning both classical machine learning (Linear Regression, Ridge Regression, Gradient Boosting, and CatBoost) and modern tabular neural networks, including TabPFN (Hollmann et al., 2023), TabM (Gorishniy et al., 2025), RealMLP, and a lightweight MLP. Together, these models represent a broad spectrum of inductive biases, ranging from simple linear baselines to recent architectures that

Table 1. Cross-device biomarker translation performance from Maestro2 to Triton across eight tabular learning models. **Bold** indicates the best result, while underline denotes the second-best.

	TabPFN		Lightweight MLP		CatBoost		TabM		RealMLP		Ridge		Gradient Boosting		Linear Regression	
Biomarker	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE
AVR (Knudtson)	**59.70**	**0.05**	59.20	0.05	53.00	0.05	59.20	0.05	56.80	0.05	57.80	0.05	50.30	0.05	57.80	0.05
Artery Distance Tortuosity	**29.60**	**3.07**	25.60	3.16	26.60	3.13	27.80	3.11	24.60	3.18	25.60	3.16	22.10	3.23	25.30	3.16
Artery Fractal Dimension	**68.80**	**0.03**	66.90	0.03	63.30	0.03	68.30	0.03	68.20	0.03	67.50	0.03	64.50	0.03	67.50	0.03
Average Width	**67.40**	**3.58**	66.40	3.63	61.00	3.91	67.30	3.58	64.80	3.72	66.80	3.61	64.90	3.71	66.70	3.62
CRAE (Knudtson)	72.90	11.39	73.50	11.28	72.50	11.48	72.30	11.52	**73.70**	**11.22**	73.00	11.37	69.80	12.02	73.00	11.37
CRVE (Knudtson)	78.20	14.38	78.00	14.46	76.80	14.86	77.70	14.56	78.10	14.43	**78.50**	**14.30**	74.80	15.48	78.40	14.32
Distance Tortuosity	**42.90**	**0.01**	41.20	0.01	39.50	0.01	42.00	0.01	40.50	0.01	41.50	0.01	37.20	0.01	41.30	0.01
Fractal Dimension	**70.10**	**0.02**	69.40	0.02	67.50	0.02	69.90	0.02	69.20	0.02	69.70	0.02	66.10	0.02	69.60	0.02
Tortuosity Density	**51.20**	**0.02**	49.80	0.02	47.00	0.02	50.50	0.02	48.80	0.02	50.00	0.02	45.20	0.02	49.90	0.02
Vein Distance Tortuosity	**31.80**	**0.01**	30.20	0.02	28.10	0.02	31.00	0.01	29.50	0.02	30.50	0.01	26.00	0.02	30.30	0.02
Vein Fractal Dimension	**64.50**	**0.02**	63.20	0.02	60.10	0.02	64.00	0.02	62.80	0.02	63.60	0.02	59.00	0.02	63.50	0.02
Vessel Density	**58.90**	**0.03**	57.50	0.03	54.20	0.03	58.40	0.03	56.70	0.03	57.90	0.03	53.00	0.03	57.70	0.03

achieve state-of-the-art performance on tabular benchmarks such as TabArena (Erickson et al., 2025).

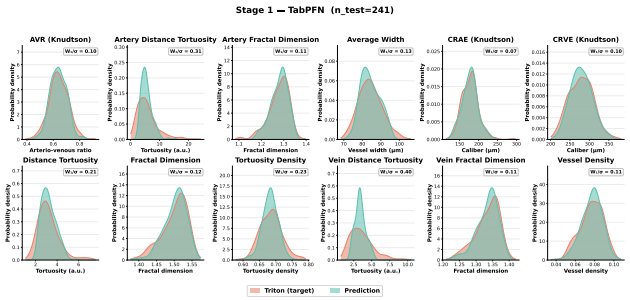

Figure 2. Distributional alignment for direct cross-device translation of AutoMorph-derived retinal biomarkers from Maestro2 to Triton.

Evaluation. We evaluate our framework across three complementary feature-space translation settings. First, we perform cross-device biomarker translation by learning mappings between AutoMorph-derived biomarkers extracted from paired Triton and Maestro2 images. Second, we perform embedding-to-biomarker translation by learning device-specific mappings from RETFound embeddings to the corresponding AutoMorph biomarkers, evaluating the extent to which retinal foundation-model representations encode clinically meaningful information. Third, we perform cross-device biomarker estimation by predicting target-device biomarkers directly from source-device RETFound embeddings, enabling device-invariant biomarker estimation without requiring target-device imaging. Models are trained on paired samples from participants imaged using multiple devices. The dataset is split at the patient level into training (70%), validation (15%), and test (15%) sets, ensuring subject-disjoint partitions. Models are selected using 5-fold cross-validation on the training set.Performance is evaluated using mean absolute error (MAE) and coefficient of determination (R^2), while cross-device distributional alignment is quantified using the Wasserstein distance, Kullback–Leibler (KL) divergence, and Jensen–Shannon (JS) divergence, providing a comprehensive assessment of both predictive accuracy and feature-space alignment.

3. Results and Discussion

Cross-device Biomarker Translation. Table 1 summarizes the results for cross-device biomarker translation from Maestro2 to Triton. Overall, the proposed framework achieves strong predictive performance, with morphological biomarkers such as CRAE, CRVE, fractal dimension, average vessel width, and vessel density attaining R^2 values of up to 78.5%, while also exhibiting close agreement between the predicted and target distributions with low Wasserstein distances (Fig. 2). In contrast, tortuosity-related biomarkers achieve lower R^2 values and show larger distributional discrepancies, suggesting greater sensitivity to device-specific acquisition variability. TabPFN consistently achieves the best or near-best performance across most biomarkers, while TabM and the lightweight MLP also perform competitively, highlighting the advantage of modern tabular learning methods over classical machine learning models for capturing complex cross-device relationships. Based on these observations, we restrict subsequent experiments to TabPFN, lightweight MLP, and CatBoost. Together, these results demonstrate that the proposed framework preserves both prediction accuracy and the underlying biomarker distributions, effectively mitigating device-induced variability.

Table 2. Embedding-to-biomarker translation on Maestro2, mapping RETFound embeddings to AutoMorph-derived retinal biomarkers. **Bold** indicates the best result, while underline denotes the second-best.

	TabPFN		Ridge		ElasticNet		PLS	
Biomarker	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE
AVR (Knudtson)	5.80	0.08	2.70	0.08	3.30	0.08	**8.30**	**0.08**
Artery Distance Tortuosity	-3.60	3.91	-34.20	4.45	0.00	3.84	**1.90**	**3.80**
Artery Fractal Dimension	74.90	0.02	75.10	0.02	**78.90**	**0.02**	75.40	0.02
Average Width	47.50	4.81	**57.60**	**4.33**	53.90	4.51	42.60	5.03
CRAE (Knudtson)	48.70	16.71	**53.70**	**15.88**	50.40	16.43	51.10	16.32
CRVE (Knudtson)	49.00	27.98	**54.00**	**26.56**	50.10	27.68	53.90	26.61
Distance Tortuosity	**23.30**	**0.82**	12.60	0.88	18.10	0.85	12.90	0.87
Fractal Dimension	85.70	0.02	**87.50**	**0.01**	83.50	0.02	83.00	0.02
Tortuosity Density	24.10	0.03	-3.50	0.03	14.90	0.03	**24.30**	**0.03**
Vein Distance Tortuosity	-4.60	1.44	-15.90	1.52	**3.90**	**1.38**	-3.10	1.43

Embedding-to-Biomarker Translation. Table 2 summarizes the results for embedding-to-biomarker translation on Maestro2. Overall, RETFound embeddings capture substantial clinically meaningful retinal information, with biomark-

Table 3. Cross-device biomarker estimation from embeddings (Maestro2 → Triton), where RETFound embeddings are first translated to biomarkers and then aligned across devices. **Bold** indicates the best result, while underline denotes the second-best.

	TabPFN		CatBoost		Lightweight MLP		Ridge	
Biomarker	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE	R^2 (%)	RMSE
AVR (Knudtson)	4.20	0.09	1.70	0.09	**5.20**	**0.09**	4.00	0.09
Artery Distance Tortuosity	6.50	3.84	2.30	3.92	-0.60	3.98	**8.50**	**3.79**
Artery Fractal Dimension	57.30	0.03	57.00	0.03	56.70	0.03	**57.70**	**0.03**
Average Width	31.80	5.13	**33.80**	**5.05**	33.10	5.08	32.20	5.11
CRAE (Knudtson)	44.30	18.58	44.20	18.59	**46.10**	**18.26**	45.20	18.42
CRVE (Knudtson)	44.50	**27.07**	44.20	27.13	43.70	27.25	**54.00**	26.56
Distance Tortuosity	21.60	0.79	19.10	0.80	11.60	0.83	**19.60**	**0.80**
Fractal Dimension	71.50	0.02	70.30	0.02	**71.60**	0.02	70.70	**0.02**
Tortuosity Density	19.40	0.03	17.60	0.03	14.70	0.03	**19.50**	**0.03**
Vein Distance Tortuosity	-8.10	1.45	-11.00	1.47	-7.10	1.44	**-6.00**	**1.44**
Vein Fractal Dimension	71.90	0.02	69.90	0.02	**73.10**	**0.02**	71.90	0.02
Vessel Density	**76.00**	0.01	74.70	0.01	75.20	**0.01**	75.00	0.01

ers such as fractal dimension, artery fractal dimension, and vessel caliber (CRAE and CRVE) achieving R^2 values exceeding 75.00%. In contrast, tortuosity-related biomarkers are recovered less accurately, suggesting that these features are less explicitly encoded in the latent representation. This trend is consistent with the distributional analysis in Fig. 3, where biomarkers with higher R^2 values exhibit closer agreement between the predicted and target distributions and lower Wasserstein distances, while tortuosity measures show comparatively larger distributional discrepancies. Overall, these findings demonstrate that RETFound embeddings preserve clinically relevant vascular morphology, although the fidelity of the encoded information varies across biomarker types.

Cross-device Biomarker Estimation from Embeddings. Table 3 summarizes the results for compositional cross-device biomarker estimation from RETFound embeddings. As expected, performance decreases relative to the individual translation stages due to error propagation across the embedding-to-biomarker and cross-device translation mappings. Nevertheless, structural biomarkers, including fractal dimension, vessel density, and vessel caliber (CRAE and CRVE), retain strong predictive performance, indicating that clinically relevant vascular information is preserved throughout the compositional pipeline. This trend is further reflected in the distributional analysis (Fig. 4), where these biomarkers exhibit close agreement between the predicted and target distributions with comparatively low Wasserstein distances, whereas tortuosity-related biomarkers show larger distributional discrepancies. Together, these findings demonstrate that the proposed framework enables robust cross-device biomarker estimation directly from retinal foundation-model embeddings, supporting its potential for clinically interpretable and device-agnostic oculomics.

4. Conclusion

In this work, we introduced a unified feature-space translation framework for addressing device-induced variability in oculomics. The proposed framework unifies cross-device retinal biomarker translation with the translation of retinal foundation-model embeddings into clinically meaningful biomarkers, enabling robust biomarker estimation across both same-device and cross-device settings. Experiments on the AI-READI dataset demonstrate that the learned mappings preserve both predictive accuracy and biomarker distributions, establishing feature-space translation as an effective strategy for improving cross-device robustness while enhancing the clinical interpretability of retinal foundation models. Future work will extend the framework to bidirectional and multi-device translation, evaluate it on the expanded AI-READI release and additional retinal datasets, and improve robustness to lower-quality imaging to further support large-scale, device-agnostic clinical deployment.

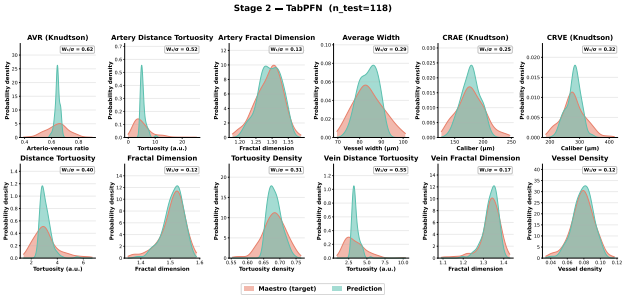

Figure 3. Distributional alignment between predicted biomarkers obtained from RETFound embeddings and ground-truth AutoMorph biomarkers. Strong overlap across key vascular features indicates that the embeddings encode clinically meaningful retinal information.

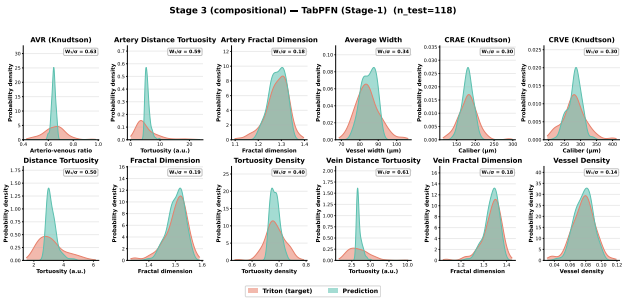

Figure 4. Distributional alignment for cross-device biomarker estimation from embeddings. RETFound embeddings are first translated to biomarkers and then aligned across devices. Predicted distributions remain well-aligned with target Triton biomarkers for key features.

References

Abay, S. G., Asmare, M. H., and Geurts, L. An affordable smartphone-based fundus imaging device for retinal examination. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, volume 00, pp. 1–6. Institute of Electrical and Electronics Engineers (IEEE), jul 2025. doi: 10.1109/embc58623.2025.11254224.

Ahn, S. J. and Kim, Y. H. Clinical applications and future directions of smartphone fundus imaging. *Diagnostics*, 14 (13):1395, jun 2024. doi: 10.3390/diagnostics14131395. URL https://www.mdpi.com/2075-4418/14/13/1395/pdf?version=1719732813.

AI-READI Consortium. Flagship dataset of type 2 diabetes from the ai-readi project (2.0.0) [data set], 2024a. URL https://doi.org/10.60775/fairhub.2.

AI-READI Consortium. Ai-readi: rethinking data collection, preparation and sharing for propelling ai-based discoveries in diabetes research and beyond. *Nature Metabolism*, 2024b. doi: 10.1038/s42255-024-01165-x. URL https://doi.org/10.1038/s42255-024-01165-x.

An, S., Teo, K., McConnell, M. V., Marshall, J., Galloway, C., and Squirrell, D. Ai explainability in oculomics: How it works, its role in establishing trust, and what still needs to be addressed. *Progress in Retinal and Eye Research*, 106:101352, mar 2025. doi: 10.1016/j.preteyeres.2025.101352. URL https://doi.org/10.1016/j.preteyeres.2025.101352.

Blose, B. A. and Silverstein, S. M. Retinal biomarkers in schizophrenia spectrum disorders: evidence and implications for the neurodevelopmental and neurodegenerative models. *Frontiers in Medicine*, 12, January 2026. ISSN 2296-858X. doi: 10.3389/fmed.2025.1697871. URL http://dx.doi.org/10.3389/fmed.2025.1697871.

Chan, J. W., Wang, Z., Xu, E., Abboud, I., Alhasany, A., Xu, S., Wu, X., Astraea, N., Jiang, F., and Hu, Z. J. Inner retinal layer thickness reflects plasma biomarkers in preclinical alzheimer's disease. *Frontiers in Medicine*, 13, May 2026. ISSN 2296-858X. doi: 10.3389/fmed.2026.1800287. URL http://dx.doi.org/10.3389/fmed.2026.1800287.

Chen, T., Bai, Y., Mao, H., Liu, S., Xu, K., Xiong, Z., Ma, S., Yang, F., and Zhao, Y. Cross-modality transfer learning with knowledge infusion for diabetic retinopathy grading. *Frontiers in Medicine*, 11:1400137, may 2024. doi: 10.3389/fmed.2024.1400137. URL https://www.frontiersin.org/articles/10.3389/fmed.2024.1400137/pdf?isPublishedV2=False.

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., Farouque, O., Li, L. J., Tan, G., Ting, D. S. W., Hsu, W., Lee, M. L., and Wong, T. Y. 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, oct 2020. doi: 10.1038/s41551-020-00626-4.

Cheung, N., Bluemke, D. A., et al. Retinal arteriolar narrowing and left ventricular remodeling. In *Journal of the American College of Cardiology*, pp. 48–55, 2007.

Christopher, M., Belghith, A., Bowd, C., Proudfoot, J. A., Goldbaum, M. H., Weinreb, R. N., Girkin, C. A., Liebmann, J. M., and Zangwill, L. M. Performance of deep learning architectures and transfer learning for detecting glaucomatous optic neuropathy in fundus photographs. *Scientific Reports*, 8(1):16685, nov 2018. doi: 10.1038/s41598-018-35044-9. URL https://www.nature.com/articles/s41598-018-35044-9.pdf.

Cunnac, P., Boussat, B., Hogg, S., Geiser, M., MacGillivray, T., Trucco, E., and Chiquet, C. Effect of using different fundus cameras and image resolutions on automatic measurements of retinal vascular parameters. *Scientific Reports*, 16(1):14366, mar 2026. doi: 10.1038/s41598-026-44196-y. URL https://www.nature.com/articles/s41598-026-44196-y_reference.pdf.

Engelmann, J. and Bernabeu, M. O. Training a high-performance retinal foundation model with half-the-data and 400 times less compute. *Nature Communications*, 16(1):6862, jul 2025. doi: 10.1038/s41467-025-62123-z. URL https://www.nature.com/articles/s41467-025-62123-z.pdf.

Erickson, N., Purucker, L., Tschalzev, A., Holzmüller, D., Desai, P. M., Salinas, D., and Hutter, F. Tabarena: A living benchmark for machine learning on tabular data. *arXiv preprint arXiv:2506.16791*, 2025.

Estrada, R., Allingham, M. J., et al. Retinal artery-vein classification via topology estimation. In *IEEE Transactions on Medical Imaging*, pp. 2518–2534, 2015.

Fraz, M. M., Welikala, R. A., et al. Quartz: Quantitative analysis of retinal vessel morphology. In *Expert Systems with Applications*, pp. 7221–7234, 2015.

Futoma, J., Simons, M., et al. The myth of generalisability in clinical machine learning. In *Lancet Digital Health*, pp. e489–e492, 2020.

Gharbi, K., Wijngaarden, P. v., and Hadoux, X. Foundation models for ophthalmic imaging. *Survey of Ophthalmology*, 71(4):1129–1147, jan 2026. doi: 10.1016/j.survophthal.2026.01.004.

Gorishniy, Y., Kotelnikov, A., and Babenko, A. Tabm: Advancing tabular deep learning with parameter-efficient ensembling. In *International Conference on Learning Representations (ICLR)*, 2025.

Guo, Y., Yan, H., Guo, C., and Wang, J. Comparative analysis of two retinal fractures with ultrabroad-angle fundus photography systems. *Biomedical Reports*, 20(2):34, jan 2024. doi: 10.3892/br.2024.1722. URL http://www.spandidos-publications.com/10.3892/br.2024.1722/download.

Gupta, P., Man, R. E. K., Fenwick, E. K., Aravindhan, A., Gan, A. T., Thakur, S., Soh, B. L. P., Wood, J. M., Black, A. A., Chan, A., Ng, D., Hean, T. K., Goh, E., Mary, C. F.-F., Loo, J., Forde, C. G., Sabanayagam, C., Cheng, C.-Y., Wong, T. Y., and Lamoureux, E. L. Rationale and methodology of the population health and eye disease profile in elderly singaporeans study [pioneer]. *Aging and Disease*, 11(6):1444–1458, dec 2020. doi: 10.14336/ad.2020.0206. URL http://www.aginganddisease.org/EN/article/downloadArticleFile.do?attachType=PDF&id=147947.

Hollmann, N., Müller, S., Eggensperger, K., and Hutter, F. Tabpfn: A transformer that solves small tabular classification problems in a second. In *International Conference on Learning Representations (ICLR)*, 2023.

Honavar, S. G. Oculomics – the eyes talk a great deal. *Indian Journal of Ophthalmology*, 70(3):713–713, February 2022. ISSN 1998-3689. doi: 10.4103/ijo.ijo_474_22. URL http://dx.doi.org/10.4103/ijo.IJO_474_22.

Huang, F. and Dashtbozorg, B. Artery/vein classification in fundus images. In *Machine Vision and Applications*, pp. 23–34, 2018.

Inamullah, I., Razzak, I., and Jameel, S. The eye as a window to systemic health: A survey of retinal imaging from classical techniques to oculomics. *The Journal of Precision Medicine: Health and Disease*, pp. 100023, 2025.

Irodi, A., Wang, X., Wang, Y. X., Tham, Y. C., Lee, C. S., Cheung, C. Y., Zhu, Z., Grzybowski, A., and Wong, T. Y. Evolution of oculomics: From naked eye observations to artificial intelligence over 100 years. *Ophthalmology Science*, 6(4):101079, jan 2026. doi: 10.1016/j.xops.2026.101079. URL https://www.ophthalmologyscience.org/article/S2666-9145(26)00017-5/pdf.

Joo, Y. S., Rim, T. H., Koh, H. B., Yi, J., Kim, H., Lee, G., Kim, Y. A., Kang, S.-W., Kim, S. S., and Park, J. T. Noninvasive chronic kidney disease risk stratification tool derived from retina-based deep learning and clinical factors. *npj Digital Medicine*, 6(1):114, jun 2023. doi: 10.1038/s41746-023-00860-5. URL https://www.nature.com/articles/s41746-023-00860-5.pdf.

Kubin, A.-M., Huhtinen, P., Ohtonen, P., Keskitalo, A., Wirkkala, J., and Hautala, N. Comparison of 21 artificial intelligence algorithms in automated diabetic retinopathy screening using handheld fundus camera. *Annals of Medicine*, 56(1):2352018, may 2024. doi: 10.1080/07853890.2024.2352018. URL https://www.tandfonline.com/doi/pdf/10.1080/07853890.2024.2352018?needAccess=true.

Liew, G., Gopinath, B., et al. Retinal vasculature fractal and stroke mortality. In *Stroke*, pp. 1276–1282, 2021.

Lim, Z. W., Chee, M.-L., Da Soh, Z., Cheung, N., Dai, W., Sahil, T., Tao, Y., Majithia, S., Sabanayagam, C., Chen, C. L.-H., Wong, T. Y., Cheng, C.-Y., and Tham, Y.-C. Association between visual impairment and decline in cognitive function in a multiethnic asian population. *JAMA Network Open*, 3(4):e203560, apr 2020. doi: 10.1001/jamanetworkopen.2020.3560. URL https://jamanetwork.com/journals/jamanetworkopen/articlepdf/2764909/lim_2020_oi_200166.pdf.

Majithia, S., Chong, C. C. Y., Chee, M. L., Yu, M., Da Soh, Z., Thakur, S., Lavanya, R., Rim, T. H., Nusinovici, S., Koh, V., Sabanayagam, C., Cheng, C.-Y., and Tham, Y.-C. Associations between chronic kidney disease and thinning of neuroretinal layers in multiethnic asian and white populations. *Ophthalmology Science*, 4(1):100353, jun 2023. doi: 10.1016/j.xops.2023.100353. URL http://www.ophthalmologyscience.org/article/S2666914523000854/pdf.

McGrory, S., Taylor, A. M., Pellegrini, E., Ballerini, L., Kirin, M., Doubal, F. N., Wardlaw, J. M., Doney, A. S. F., Dhillon, B., Starr, J. M., Trucco, E., Deary, I. J., and MacGillivray, T. J. Towards standardization of quantitative retinal vascular parameters: Comparison of siva and vampire measurements in the lothian birth cohort 1936. *Translational Vision Science & Technology*, 7(2):12–12, mar 2018. doi: 10.1167/tvst.7.2.12. URL https://tvst.arvojournals.org/arvo/content_public/journal/tvst/936792/i2164-2591-7-2-12.pdf.

Mensah, S. O., Roa, C., Djoumessi, K., and Berens, P. Towards interpretable foundation models for retinal fundus images. *arXiv preprint arXiv:2603.18846*, 2026.

Panwar, N., Huang, P., Lee, J., Keane, P. A., Chuan, T. S., Richhariya, A., Teoh, S., Lim, T. H., and Agrawal, R. Fundus photography in the 21st century—a review of recent technological advances and their implications for worldwide healthcare. *Telemedicine and e-Health*, 22(3):198–208, aug 2015. doi: 10.1089/tmj.2015.0068. URL https://www.ncbi.nlm.nih.gov/pmc/articles/4790203.

Perez-Rovira, A., MacGillivray, T., et al. Vampire: Vessel assessment platform for retina. In *IEEE EMBS Conference*, 2011.

Qiu, J., Wu, J., Wei, H., Shi, P., Zhang, M., Sun, Y., Li, L., Liu, H., Liu, H., Hou, S., et al. Visionfm: a multi-modal multi-task vision foundation model for generalist ophthalmic artificial intelligence. *arXiv preprint arXiv:2310.04992*, 2023.

Qiu, J., Wu, J., Wei, H., Shi, P., Zhang, M., Sun, Y., Li, L., Liu, H., Liu, H., Hou, S., et al. Development and validation of a multimodal multitask vision foundation model for generalist ophthalmic artificial intelligence. *NEJM AI*, 1(12):AIoa2300221, 2024.

Rastogi, S. and Thakur, S. Enabling cross-device retinal biomarker alignment for generalizable oculomics using supervised space translation. *KDD UMC 2025*, 2025.

Rastogi, S., Modi, K., Thakur, S., and Arora, V. *Early CHD Detection from Retinal Fundus Scans Using a Spatial Context-Aware Hierarchical Attention Framework*, pp. 147–157. Springer Nature Switzerland, November 2025. ISBN 9783032103512. doi: 10.1007/978-3-032-10351-2_15. URL http://dx.doi.org/10.1007/978-3-032-10351-2_15.

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., Kim, H. C., Kim, S. S., Park, S., Cheng, C.-Y., and Wong, T. Y. Deep-learning-based cardiovascular risk stratification using coronary artery calcium scores predicted from retinal photographs. *The Lancet Digital Health*, 3(5): e306–e316, may 2021. doi: 10.1016/s2589-7500(21)00043-1. URL http://www.thelancet.com/article/S2589750021000431/pdf.

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., and Wong, T. Y. A deep learning algorithm to detect chronic kidney disease from retinal photographs in community-based populations. *The Lancet Digital Health*, 2(6): e295–e302, may 2020. doi: 10.1016/s2589-7500(20)30063-7. URL http://www.thelancet.com/article/S2589750020300637/pdf.

Sasongko, M. B., Wong, T. Y., et al. Retinal arteriolar tortuosity and kidney dysfunction. In *American Journal of Ophthalmology*, pp. 176–183, 2012.

Shi, D., Zhang, W., Yang, J., Huang, S., Chen, X., Xu, P., Jin, K., Lin, S., Wei, J., Yusufu, M., Liu, S., Zhang, Q., Ge, Z., Xu, X., and He, M. A multimodal visual–language foundation model for computational ophthalmology. *npj Digital Medicine*, 8(1):381, jun 2025. doi: 10.1038/s41746-025-01772-2. URL https://www.nature.com/articles/s41746-025-01772-2.pdf.

Srinidhi, C. L., P., A., and Rajan, J. Automated retinal artery/vein separation. In *IEEE Transactions on Image Processing*, pp. 2705–2718, 2019.

Sung, J., Midorikawa-Inomata, A., Fujio, K., and Inomata, T. Expanding the horizon of oculomics: A narrative review on the emerging roles of the anterior segment and mobile health. *Survey of Ophthalmology*, 71(4):1074–1089, 2026. ISSN 0039-6257. doi: 10.1016/j.survophthal.2026.02.003. URL http://dx.doi.org/10.1016/j.survophthal.2026.02.003.

Tan, Y., Ma, Y., et al. Deep learning for detection of chronic kidney disease from retinal images. In *European Journal of Ophthalmology*, pp. 502–509, 2024a.

Tan, Y. Y., Kang, H. G., Lee, C. J., Kim, S. S., Park, S., Thakur, S., Da Soh, Z., Cho, Y., Peng, Q., Lee, K., Tham, Y.-C., Rim, T. H., and Cheng, C.-y. Prognostic potentials of ai in ophthalmology: systemic disease forecasting via retinal imaging. *Eye and Vision*, 11(1):17, may 2024b. doi: 10.1186/s40662-024-00384-3. URL https://eandv.biomedcentral.com/counter/pdf/10.1186/s40662-024-00384-3.

Thakur, S., Rastogi, S., Modi, K., Rukmini, A., Kim, S., Seo, W., Park, J., Nam, D., Cho, J., Seo, J., et al. Retinal vascular and renal biomarkers predict systemic cardiovascular and kidney disease in the brazilian multilabel ophthalmological dataset (brset). *Circulation*, 152(Suppl_3): A4369645–A4369645, 2025.

Ting, D. S. W., Cheung, C. Y.-L., Lim, G., Tan, G. S. W., Quang, N. D., Gan, A., Hamzah, H., Garcia-Franco, R., San Yeo, I. Y., Lee, S. Y., Wong, E. Y. M., Sabanayagam, C., Baskaran, M., Ibrahim, F., Tan, N. C., Finkelstein, E. A., Lamoureux, E. L., Wong, I. Y., Bressler, N. M., Sivaprasad, S., Varma, R., Jonas, J. B., He, M. G., Cheng, C.-Y., Cheung, G. C. M., Aung, T., Hsu, W., Lee, M. L., and Wong, T. Y. Development and validation of a deep learning system for diabetic retinopathy and related eye diseases using retinal images from

multiethnic populations with diabetes. *JAMA: The Journal of the American Medical Association*, 318(22): 2211–2223, dec 2017. doi: 10.1001/jama.2017.18152. URL https://jamanetwork.com/journals/jama/articlepdf/2665775/jama_ting_2017_oi_170140.pdf.

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., Kim, H., Lee, G., Yu, M., Tham, Y.-C., Bakhai, A., Leeson, P., Lip, G. Y., Wong, T. Y., and Cheng, C.-Y. Validation of a deep-learning-based retinal biomarker (reti-cvd) in the prediction of cardiovascular disease: data from uk biobank. *BMC Medicine*, 21(1): 28, jan 2023. doi: 10.1186/s12916-022-02684-8. URL https://bmcmedicine.biomedcentral.com/counter/pdf/10.1186/s12916-022-02684-8.

Vilela, M. A. P., Arrigo, A., Parodi, M. B., and da Silva Mengue, C. Smartphone eye examination: Artificial intelligence and telemedicine. *Telemedicine and e-Health*, 30(2):341–353, feb 2024. doi: 10.1089/tmj.2023.0041.

Wang, Y., Li, Y., Zhang, J., Niu, B., Jiang, Y., and Xiao, Y. Case report: Optic nerve infiltration associated with sequential crvo and crao in isolated cns relapse of aml. *Frontiers in Medicine*, 13, 2026. ISSN 2296-858X. doi: 10.3389/fmed.2026.1853833. URL http://dx.doi.org/10.3389/fmed.2026.1853833.

Wong, T. Y., Klein, R., et al. Retinal arteriolar narrowing and risk of coronary heart disease. In *JAMA*, 2002.

Xu, Y. and Trucco, E. The eye as a digital biomarker: transforming systemic disease assessment, 2026.

Zhou, Y., Wagner, S. K., et al. Automorph: Automated retinal vascular morphology quantification. In *Translational Vision Science and Technology*, 2022.

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., et al. A foundation model for generalizable disease detection from retinal images. *Nature*, 622(7981):156–163, 2023.

Zhu, Z., Wang, Y., et al. Oculomics: Current concepts and evidence. In *Progress in Retinal and Eye Research*, 2025.