Title: iGWAS: image-based genome-wide association of self-supervised deep phenotyping of human medical images

Ziqian Xie[1,2]
Tao Zhang[1]
Sangbae Kim[1]
Jiaxiong Lu[1]
Wanheng Zhang[3]
Cheng-Hui Lin[4]
Man-Ru Wu[4]
Alexander Davis[4]
Roomasa Channa[5]
Luca Giancardo[2]
Han Chen[3,6]
Sui Wang[4]
Rui Chen[1,*]
Degui Zhi[2,*]

[1]Department of Molecular and Human Genetics, Baylor College of Medicine, Houston, Texas 77030, USA;
[2]School of Biomedical Informatics, University of Texas Health Science Center, Houston, Texas 77030, USA;
[3]School of Public Health, University of Texas Health Science Center, Houston, Texas 77030, USA;
[4]Department of Ophthalmology, Stanford University School of Medicine, Stanford, CA 94305, USA;
[5]Department of Ophthalmology and Visual Sciences, University of Wisconsin, Madison, Wisconsin 53726;
[6]Human Genetics Center, University of Texas Health Science Center, Houston, Texas 77030, USA.

[*]Corresponding authors: Rui Chen (ruichen@bcm.edu) and Degui Zhi (Degui.Zhi@uth.tmc.edu)

Keywords: Deep Learning, Artificial Intelligence, Retina image, Imaging genetics, GWAS

Abstract
Existing imaging genetics studies have been mostly limited in scope by using imaging-derived phenotypes defined by human experts. Here, leveraging new breakthroughs in self-supervised deep representation learning, we propose a new approach, image-based genome-wide association study (iGWAS), for identifying genetic factors associated with phenotypes discovered from medical images using contrastive learning. Using retinal fundus photos, our model extracts a 128-dimensional vector representing features of retinal vasculature that can be used as phenotypes. After training the model on 40,000 images from the EyePACS dataset, we generated phenotypes from 130,967 images of 65,629 British White participants in the UK Biobank. A GWAS of these vasculature phenotypes identified 34 associated loci with genome-wide significance. A functional follow up study verified the role of the *WNT7B* gene, a newly found locus, in retinal vessel development. Our results establish the feasibility of this new framework of genomic study based on self-supervised phenotyping of medical images.

Introduction

Although genome-wide association studies (GWAS) have successfully identified thousands of genetic associations, most existing GWAS are based on a set of predefined phenotypes. While these phenotypes encode valuable biomedical knowledge, they are also biased by current clinical practice and epidemiological studies. In addition, as the granularity of phenotype code is often limited, it is often not sufficient to capture the complexity of human physiology and pathology in their entirety. Therefore, deriving new phenotypes beyond expert curation would enable the discovery of new genetic associations.

Medical imaging is a rich resource for phenotype discovery. Through rapid technological advancements, modern medical imaging offers unprecedented details about a patient's physiological condition and can be a high-content phenotyping modality. Existing imaging GWASs have leveraged imaging-derived phenotypes (IDPs)[1–3]. These IDPs were typically designed by imaging experts and generated by special-purpose image processing pipelines. Recently, machine learning, especially supervised deep learning (DL), is used to automatically generate IDPs[4–6]. These methods were trained by learning from data labeled by experts and identified new loci in GWAS[2,7,8]. However, although supervised DL can vastly improve the efficiency of image labeling, it fails to provide phenotypes beyond those defined by experts. In addition, although these phenotypes are derived for medical practice, clinical decision processes, and natural-language-based reporting, they often do not comprehensively capture the imaging content. There are limitations to the amount of information a human eye can extract from images. Many meaningful imaging features, some of which might be used implicitly by physicians, may not be verbalized in medical reports. In addition, there may be physiologically informative features that are present in the image but are completely missed or ignored by readers. For example, Google's DL algorithm extracted novel features from retinal images, such as age, gender, and smoking status, that are not readily apparent to expert human graders[9]. Following studies identified features such as refractive error and anemia from retinal images[10,11]. These results suggest that additional information beyond human curation may be encoded within imaging data, and new methods are needed to extract such information.

Here, we have designed a new framework of genome-wide genotype-phenotype association study by performing unsupervised image-based genome-wide association studies (iGWAS). For phenotype discovery, instead of supervised learning that relies on labels from expert annotations, unsupervised deep learning is applied to an image to capture its intrinsic contents[12–15]. Endophenotypes generated by the deep learning model are then subjected to GWAS to identify associated genomic loci. We tested this new approach using human fundus images. In particular, we focus on the vasculature structure of the retina because retinal vasculature of healthy individuals remains relatively stable since development, making it a suitable phenotype for genetic association

studies. In addition, defects in the blood vessel in the retina lead to many human retinal diseases, such as diabetic retinopathy, and the feasibility of learning retinal vasculature representations as biomarkers has already been investigated[16,17]. We constructed a contrastive loss function over an Inception V3 architecture to learn a representation that captures the intrinsic vascular features of individuals. Using vessel segmentation masks generated from retinal fundus photos as inputs, we designed a phenotyper, a neural network model, that outputs 128 endophenotypes representing retinal vasculature. After training on 40,000 images from EyePACS, our model generated phenotypes from 130,967 images of 65,629 British White participants in the UK Biobank. A GWAS of these vasculature phenotypes identified 34 independent loci. A follow up functional assay verified the role of the *WNT7B* gene, a novel candidate locus, in retinal vascular development *in vivo*. This new iGWAS framework can be readily expanded to all types of medical images and enable unbiased discovery of genetic factors associated with human phenotypes.

RESULTS

Overall iGWAS framework

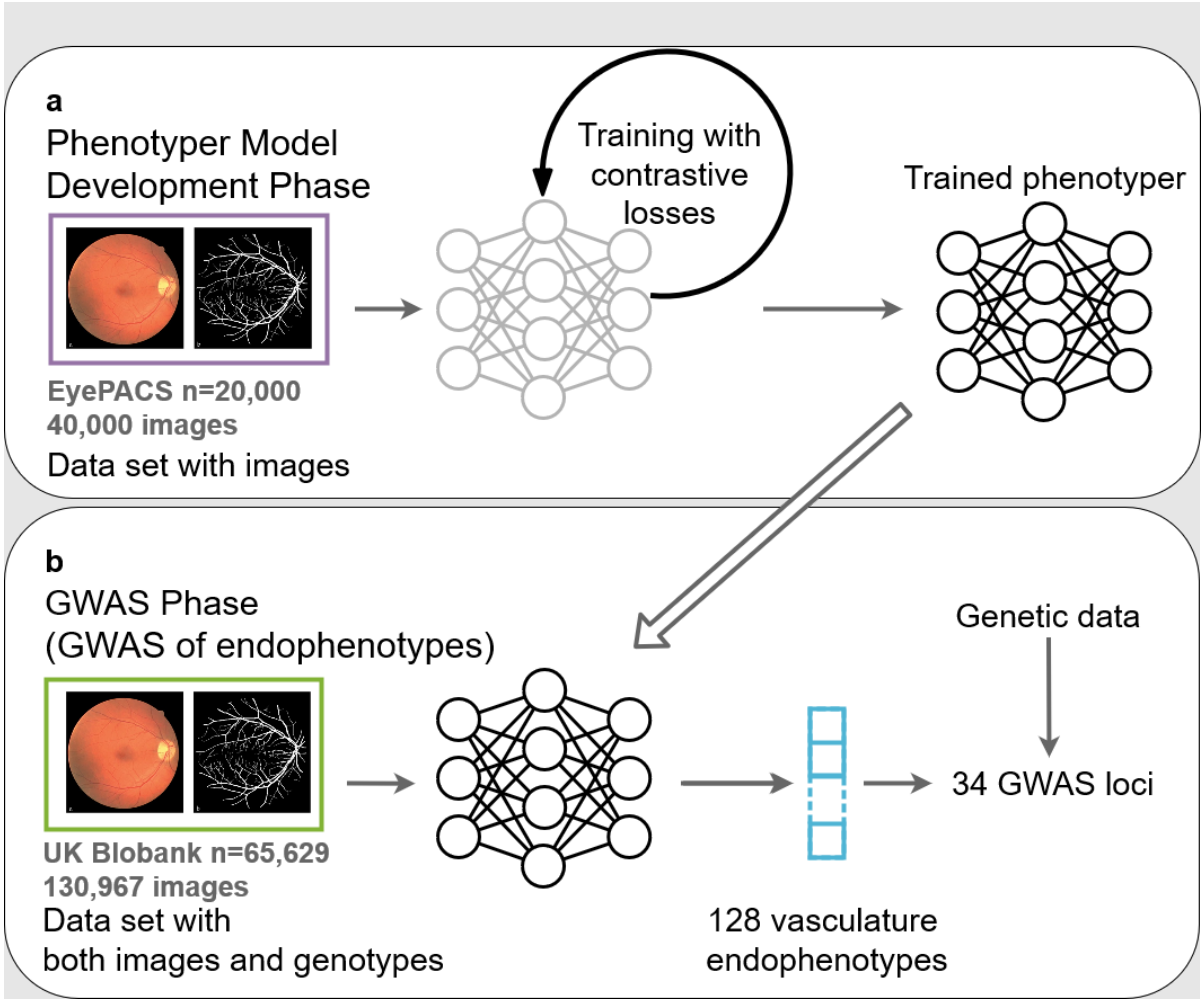

Figure 1: iGWAS of endophenotypes from vessel features of retinal fundus images. (**a**) Using segmented vasculature images derived from fundus images (segmentation network omitted for conciseness) in EyePACS, we developed a phenotyper neural network that optimizes contrastive losses; (**b**) Using the trained phenotyper, we generated 128 endophenotypes for each image in the UK Biobank vision cohort. GWAS of these endophenotypes identified 34 independent loci.

The core component of iGWAS is a phenotyping (encoder) neural network that generates endophenotypes, which are in turn associated with genotypes by GWAS (an example of iGWAS for retinal vasculature is shown in **Fig 1**). Distinct from traditional phenotypes labeled by experts or by AI trained via supervised learning, iGWAS's encoder network is trained by self-supervised learning to discover new phenotypes. We thus named it as **S**elf-**Su**pervised **P**henotyp**er** (SSuPer). Popular self-supervised learning losses, such as contrastive losses[12,14,18] and reconstruction losses[19], are used to extract coherent and biologically relevant features of individuals. The resulting "embedding vector," the output of the encoder, is treated as "endophenotypes" for downstream GWAS analysis.

The iGWAS approach is executed in two phases: the model development phase and the GWAS phase. In the model development phase, a "phenotype development set" is used to train the embedding network. The phenotype development set is a collection of images from individuals, whose genotype data are not needed. The result of the model development phase is a trained neural network model, SSuPer, that can transform an input image into a set of self-supervised image-derived phenotypes (SS-IDPs). In the GWAS phase, the trained SSuPer from the model development phase is used to generate SS-IDPs for images from the "GWAS set," a dataset containing both images and genotypes of a different cohort of individuals. The SS-IDPs are then tested for association with genome-wide markers.

Overall data analysis strategy for generating vasculature endophenotypes from fundus images

In this study, we designed and implemented the iGWAS approach to encode retinal vasculature from fundus images. For the phenotype development set, we used data from EyePACS, a large public collection of 88,702 fundus images (see Methods: dataset extraction). After quality control (see Methods: Image quality control), 40,000 top quality images were used (Supplementary Table 1). For the GWAS set, we used fundus images and genotype data of 65,629 British White UK Biobank participants. Although the demographics of the EyePACS and UK Biobank cohorts do not match exactly, we reasoned that some characteristics of their fundus images should be similar, so we expect the features learned from EyePACS can be generalized to UK Biobank.

The EyePACS fundus images were first fed into a vasculature segmentation network to generate vessel masks, which were then treated as the inputs to an

embedding network to generate the endophenotypes. We expected that the segmentation results to mainly contain vasculature information, so that the endophenotypes would mainly be related to the vasculature as well. However, we found some of the endophenotypes strongly correlated with the color and texture of the l fundus, possibly reflecting features from non-vascular factors, such as the retinal pigment epithelium layer. Therefore, to account for these confounding factors, the "retinal color," defined as the average intensities of the red, blue, and green channels of the central patch of the fundus image, was considered as a variable in subsequent analysis (see Methods). While the definition of retinal color may not fully account for change in illumination, texture of the retinal pigment epithelium, retinal lesions, and optic disk, it captures coarse-grained information of the retinal background. We conducted GWAS analyses for both the endophenotypes and the retina color, and significant endophenotype-associated GWAS loci were identified upon adjusting for the retinal color. To aid interpretation of the endophenotypes, we conducted univariate and correlation analyses among endophenotypes and between endophenotypes and relevant eye phenotypes. The overall pipeline is shown in **Supplementary Fig 1**.

Design of vasculature-embedding network that captures coherent features of fundus images from the same person

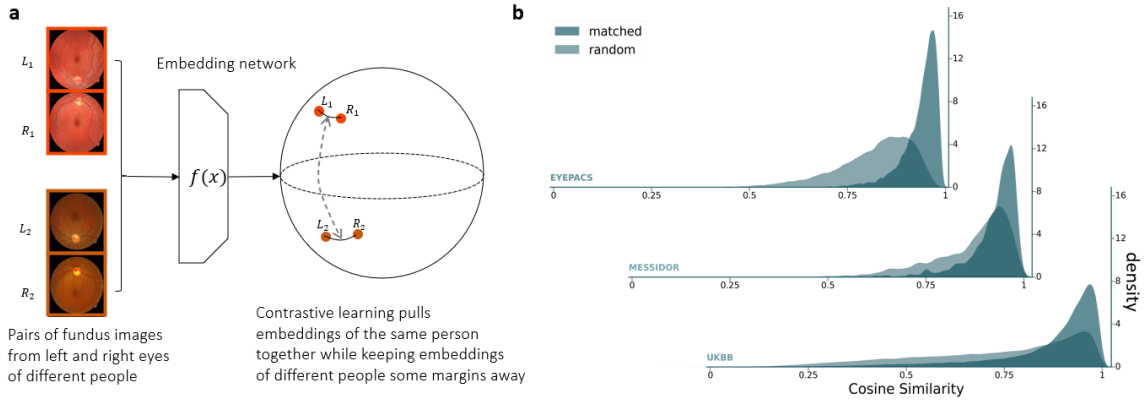

Figure 2: Contrastive loss for deriving phenotypes coherent across images from the same person. (a) Contrastive loss is designed to map images from the same person to be closer while keeping images from different persons apart. (**b**) The trained endophenotype vectors of the same persons reflect the design of contrastive learning in both the training set (EyePACS) and the test sets (Messidor and UKBB). The distributions of the matched pairs (images from the same person) and the random pairs are separated. The distributions were estimated using Scott's kernel with an additional multiplicative factor of 0.5 to smooth the curve.

The vessel segmentation network was designed as a preprocessing step to enrich vasculature-related features for the input to the main phenotyper network. Vessel segmentation of fundus images is a well-studied problem with mature methods[20–22]. Therefore, we choose a patch-based vessel segmentation network

with the classic U-net architecture[23] (See Methods: Fundus image segmentation). For 88,724 images from EyePACS, 54,992 passed our quality filter network (quality score > 0.5) (see Methods). 40,000 top quality images (quality score > 0.95) were selected as we reasoned that this balance point of sample size and sample quality is sufficient for training the main SSuPer network. The characteristics of the EyePACS dataset are shown in **Supplementary Table 1**.

To generate an embedding vector that represents the inherent biological features of an individual, we leverage a self-supervised metric learning approach that was described in ArcFace[24], a widely adopted algorithm that is used to extract features for developing human face recognition methods, with some technical modifications (see Methods: Embedding neural network). Inception v3, which has been demonstrated to be capable of capturing complex information within fundus images, was used as a backbone architecture for the metric learning[25]. The output of our embedding network was designed to be a 128-dimensional vector, based on previous works showing that 128-dimensional vectors are sufficient to represent complex datasets[26,27]. Our ArcFace loss function is a contrastive loss that first projects the embedding vector to the unit sphere and then optimizes the contrast between the embeddings from the eyes of the same person and the embeddings from different people by minimizing the angular distance between the embeddings of left and right retinas from the same individual while keeping the embeddings from different individuals at least some margins apart (**Fig 3**). We reasoned that if the trained model manages to capture real biologically relevant features, embeddings between an individual's left and right fundus images should be more similar than those from different individuals. Details of model design and training are described in the Methods: Embedding neural network.

To verify the performance of the SSuPer embedding network, we compared the matched pairs (left and right eye of the same person) and random pairs (**Fig 2b**). As expected, there is a clear separation in the distribution of cosine distance between matched pairs and random pairs. The Jensen-Shannon distances between the two distributions are 0.4194/0.1324/0.1073 and the Jaccard indices between the two distributions are 0.228/0.535/0.504 in the EyePACS, Messidor and UK Biobank datasets, respectively. Although less than that of EyePACS, the separation of matched and random pairs was clearly observed in UK Biobank (**Fig 2b**), indicating the segmentation and embedding models are transferable and indeed capture the intrinsic features of the fundus images. Therefore, we then directly applied the vasculature embedding network trained using EyePACS to the UK Biobank data without fine-tuning.

Descriptive analysis of vasculature endophenotypes in UK Biobank fundus photos

We conducted our analysis using of 65,629 British White participants from the UK Biobank who had available fundus images (see Methods: Dataset extraction).

For each participant, we chose the first image for each eye, resulting in 130,967 images. Basic demographic description of this dataset is shown in **Supplementary Table 2**. Retina colors were also extracted as phenotypes. The central patch of the image (the fovea region) was used because it has more pigment and of low vessel density, providing a cleaner estimate of the retinal color (see Methods: Color GWAS).

Univariate distributions of the 128 endophenotypes generated by our vasculature embedding network is shown in **Supplementary Fig 2**. Interestingly, we found that most endophenotypes have unimodal or bimodal distributions. Examining their pairwise correlations showed the 128 endophenotypes are grossly clustered into two groups (**Fig 3a**). The large cluster of about 100 endophenotypes is highly correlated (most with |r|>0.5) with red channel intensity (**Fig 3b**), and the small cluster of about 30 endophenotypes is correlated with blue channel intensity (**Fig 3b**). The endophenotypes showed a substantial correlation with fundus color although the vessel segmentation results did not explicitly contain fundus color information. One possible explanation for this observation is that the quality of vessel segmentation may be influenced by the overall retina color.

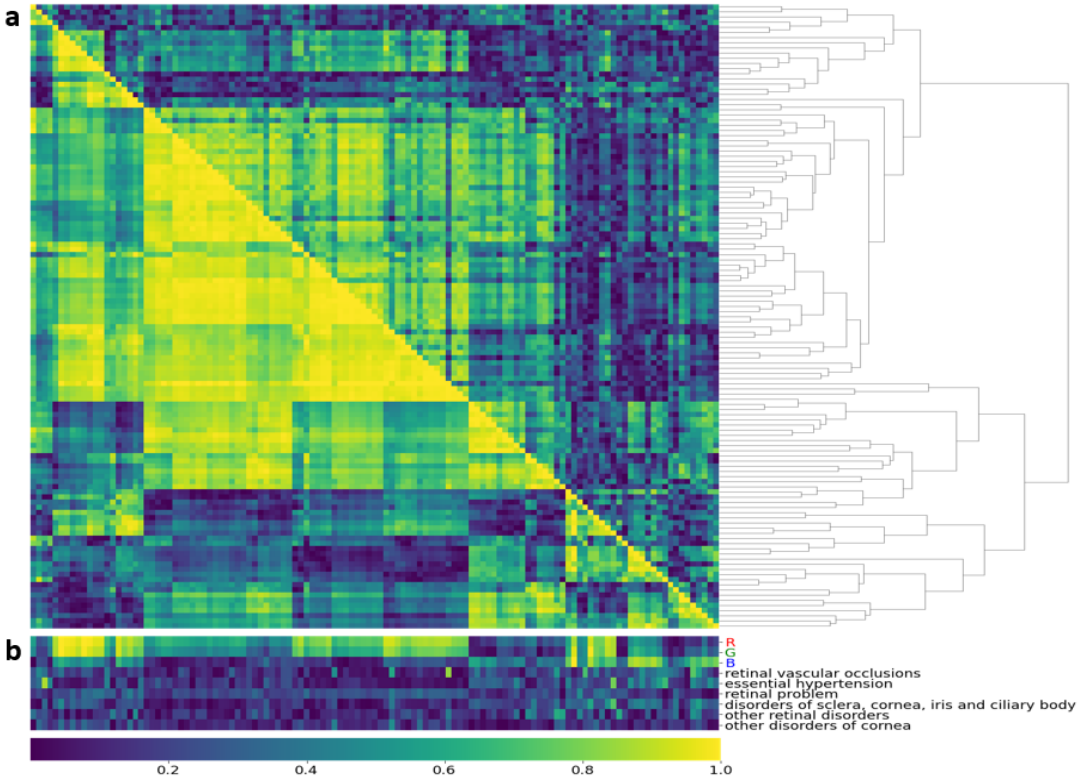

Figure 3. Interpretation of endophenotypes via correlations with known phenotypes. (**a**) absolute correlations among endophenotypes; cells in the upper triangle are phenotypic correlations and those in the lower triangles are genetic

correlations. (**b**) correlations between endophenotypes and retinal colors and 22 retina-related phenotypes. Phenotypes are sorted by the L2 norm of their correlation vector. See Supplementary Table 3 for detailed numeric values and the names of the endophenotype dimensions.

iGWAS: GWAS of vasculature endophenotypes

To identify genetic factors associated with endophenotypes, GWAS was performed for each of the 128 dimensions from all 130,967 images using linear mixed models as implemented by BOLT-LMM[28], adjusted by age, sex, and ancestral principal components (PCs) (see **Method: Vasculature endophenotype GWAS**). Analyses were conducted separately for the left and right retinal images. Their results were not meta-analyzed because the endophenotypes of the two eyes may be correlated due to training. Instead, we pooled the results from the two eyes, and only the more significant p-value between the two eyes was reported. Since the endophenotypes were derived without the direct use of any genetic information, we expected there to be minimal genomic inflation for the GWAS. Indeed, we observed that the genomic inflation factor was well-controlled ($\lambda_{GC} \approx 1$) (**Supplementary Fig 3**), though some endophenotypes had slightly higher (1.099) inflation factors, indicating potential polygenic genetic architecture.

As a result, we identified 4,011 association signals from 164 SNPs (**Supplementary Table 4**) showing genome-wide significance (p-value$<5\times10^{-8}$) (**Fig 4**). These SNPs were merged into 34 independent loci (**Table 1**) (see Methods: Vasculature endophenotype GWAS). Since the vessel endophenotypes correlated with retina color (**Fig 3b**), we also conducted a sensitivity analysis by adding the three color channel intensities as covariates in association analyses. After Bonferroni correction, 24 loci remained significantly associated with endophenotypes ($p<0.05/34$).

SNP	CHR	BP	REF/ALT	ALT_FREQ	PVAL	CA_PVAL	GENE
rs1129038	15	28,356,859	C/T	2.15E-01	1.00E-324	1.00E-324	*HERC2*
rs80308281	7	100,457,578	C/T	5.97E-03	3.50E-91	3.90E-66	*EPHB4*
rs12203592	6	396,321	T/C	2.05E-01	1.60E-33	3.00E-40	*IRF4*
rs188746966	7	99,618,147	A/G	3.76E-03	8.20E-46	1.20E-39	*ZSCAN21*
rs16891982	5	33,951,693	C/G	2.28E-02	6.20E-64	2.30E-36	*SLC45A2*
rs3791220	2	227,066	A/G	3.49E-01	3.50E-45	5.20E-36	*SH3YL1*
rs72928978	11	68,831,364	A/G	1.10E-01	7.80E-41	7.80E-26	*TPCN2*
rs941799	14	92,776,825	T/C	4.55E-01	2.10E-51	1.80E-21	*SLC24A1*
rs1126809	11	89,017,961	A/G	3.02E-01	2.30E-17	1.60E-16	*TYR*
rs10157145	1	205,261,963	C/T	4.97E-01	4.90E-28	7.90E-16	*TMCC2*
rs9561568	13	95,155,940	C/A	3.14E-01	2.30E-21	2.50E-15	*DCT*
rs12101582	15	29,034,720	C/G	4.60E-01	1.50E-17	2.70E-14	*PDCD6IPP2*
rs9906358	17	79,614,481	G/A	3.78E-01	5.40E-31	4.40E-13	*PDE6G*
rs11613220	12	20,663,103	G/T	2.46E-01	2.70E-14	5.30E-10	*PDE3A*
rs204781	8	82,588,492	C/T	2.19E-02	2.40E-13	8.00E-10	*IMPA1*
rs117514387	12	110,042,391	T/C	8.04E-02	2.60E-10	1.30E-09	*MVK*
rs1949482	11	16,214,262	A/G	4.67E-01	9.60E-12	1.10E-08	*SOX6*
rs9330813	22	46,364,161	A/G	3.14E-01	3.10E-18	1.40E-07	*WNT7B*
rs9376000	6	134,325,029	A/G	1.90E-01	6.50E-10	1.80E-06	*SLC2A12*
rs7045553	9	100,307,789	T/C	3.97E-01	5.00E-12	1.40E-04	*TMOD1*
rs208078	11	31,128,721	G/T	3.67E-01	3.90E-13	3.10E-04	*DCDC1*
rs76929107	3	195,300,740	T/C	1.31E-02	5.10E-09	9.70E-04	*APOD*
rs1794036	1	227,048,759	A/G	2.19E-01	2.30E-12	9.80E-04	*PSEN2*
rs11147339	13	29,168,148	C/A	1.65E-01	4.00E-23	1.20E-03	*FLT1*
rs10011821	4	1,777,512	A/G	1.66E-01	2.80E-09	2.80E-03	*FGFR3*
rs6142046	20	32,514,061	C/T	4.20E-01	2.20E-10	1.30E-02	*RALY*
rs12203518	6	41,510,291	C/G	2.78E-01	7.10E-15	2.50E-02	*FOXP4*
rs177441	3	4,537,135	T/C	1.21E-01	8.00E-14	8.40E-02	*ITPR1*
rs13209579	6	64,955,890	T/C	2.95E-01	1.20E-10	2.50E-01	*EYS*
rs10002734	4	47,366,395	T/G	4.48E-01	1.00E-13	3.00E-01	*GABRB1*
rs317523	11	110,891,223	C/T	3.80E-01	1.20E-22	3.50E-01	
rs6781242	3	41,300,601	G/A	4.73E-01	5.50E-13	5.80E-01	*CTNNB1*
rs2453581	17	19,438,560	T/C	3.89E-01	5.90E-13	7.00E-01	*SLC47A1*
rs12866475	13	73,103,649	C/T	2.67E-01	7.00E-13	7.80E-01	

Table 1: 34 loci significantly associated with vasculature endophenotypes. The PVAL column is the most significant p-value of the iGWAS among all endophenotypes. Loci are sorted by the p-values of the retinal-color-adjusted regression (CA_PVAL). 24 loci with significant CA_PVAL ($P<0.05/34$) are shown above the line. Candidate genes for each locus were annotated based on their distance from the leading SNPs and annotated function.

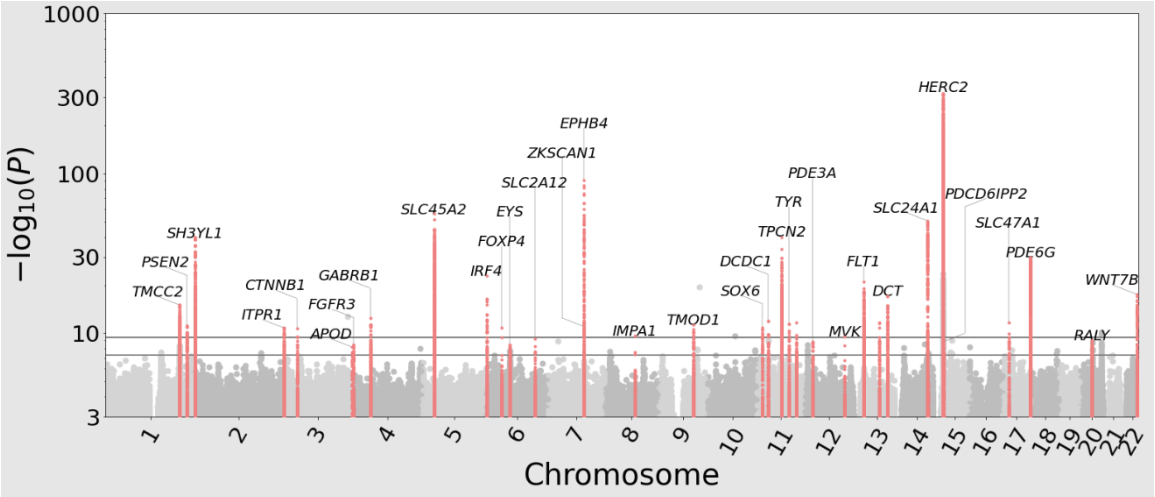

Figure 4: Manhattan plot of all 128 endophenotypes. The association P-values for the iGWAS without adjusting for color are shown. The two horizontal lines indicate significance levels set for individual GWAS (p=5×10^{-8}) and all phenotypes (p=5×10^{-8}/128). Among 34 significant loci, 24 remain significant after adjustment of retina color (P<0.05/34) and are shown in light red with gene name annotated.

Very recently, in a GWAS of retinal vessel tortuosity using fundus images, tortuosity was determined by a non-deep learning software. Interestingly, two of the 34 of loci identified in our study overlapped with the top 10 SNPs shown in the tortuosity study[29], at chr12:*PDE3A* and chr15:*HERC2/OCA2*. Furthermore, at least five candidate genes we identified, including *FLT1*, *EPHB4*, *WNT7B*, *SOX6*, and *APOD*, have been reported to have vessel-related functions. For example, highly significant SNPs (p-value = 4×10^{-23}) were found around *FLT1*, which is a negative regulator of *VEGF* and is known to play an important role in retinal vessel development[30]. Similarly, a strong association was found with erythropoietin-producing hepatocellular receptor B4 (*EPHB4*) (Supplementary Fig 4.1) , which is essential in vessel development[31]. Modulation of *Ephb4* activity in the mouse retina was found to alter retinal neovascularization[32,33]. Another identified locus was centered at *WNT7B* (Supplementary Fig 4.2), which is known to be important in blood brain barrier development although its role in retinal vessels has not been investigated[34,35]. Similarly, *SOX6* has been implicated in arteriole development in kidney[36] and *APOD* has been linked to chorioretinal blood vessel formation[37]. Strikingly, none of these loci have been associated with retinal vasculature phenotypes in previous GWAS studies[38,39] as documented in GWAS Catalog, further indicating that our framework is an effective, novel approach for identifying new genetic loci associated with phenotypes captured by imaging.

GWAS of retina color

Pigmentation of the human body, such as hair, skin, and iris, is strongly influenced by genetics. As the color of the human retina is influenced by factors such as the level of pigmentation of retinal pigment epithelium (RPE) and choroid blood vessels, we tested if genomic loci associated with retina color can be identified through genome association study of fundus images (**Method: Color GWAS**). While association of iris color has been conducted[40], no direct association studies of retinal color using fundus images have been conducted. In our study, significant genome-wide association ($p<1.67\times10^{-8}$, adjusting for the three color channels tested) was obtained for a total of 154 SNPs (Supplementary Table 5 and Supplementary Fig 5) from 35 independent loci (Supplementary Table 6). We found 13 out of the 35 loci overlapped with previously reported GWAS loci for "hair color"[41], "eye color"[42], and "skin pigmentation"[42] in the GWAS catalog (Supplementary Table 6, Supplementary Fig 7, see Supplementary Table 7 for details), supporting the validity our approach. Interestingly, many genes from unique loci identified in this study can be linked to pigmentation pathways (Supplementary Table 6). For example, mutations in *FGFR3* lead to familial acanthosis nigricans, which results in skin pigmentation abnormalities[43]. In addition to pigmentation, it is interesting to note that 7 of the 35 loci overlap with loci associated with macular thickness (p-value=2.2×10^{-16}, Supplementary Table 6), including *EPHB4*, *DCDC1*, *TPCN2*, *NCAM1*, *HERC2*, *PDE6G*, and *WNT7B*.

Genetic correlation analyses of endophenotypes

To further interpret these retinal vasculature endophenotypes, we correlated them with other traits that are related to retinal phenotypes. We conducted genetic correlation using summary statistics as they are easier to access and are suggested to be a good surrogate for phenotypic correlation[44,45]. We included traits that either have GWAS hits near the iGWAS loci (within 200kb) or are known to be related to retinal or corneal disorders, and whose genetic summary statistics for UK Biobank data are available (see **Methods: Genetic correlation**). Corneal phenotypes were included because they may affect eye refraction, which can have a detectable effect on the fundus images. Interestingly, as shown in **Fig 3b**, we found that some endophenotypes have high genetic correlation ($|r| > 0.5$) with diseases such as essential hypertension, retinal vascular occlusion, and disorders of the sclera, cornea, iris and ciliary body.

Functional validation of candidate gene involved in retinal vessel development

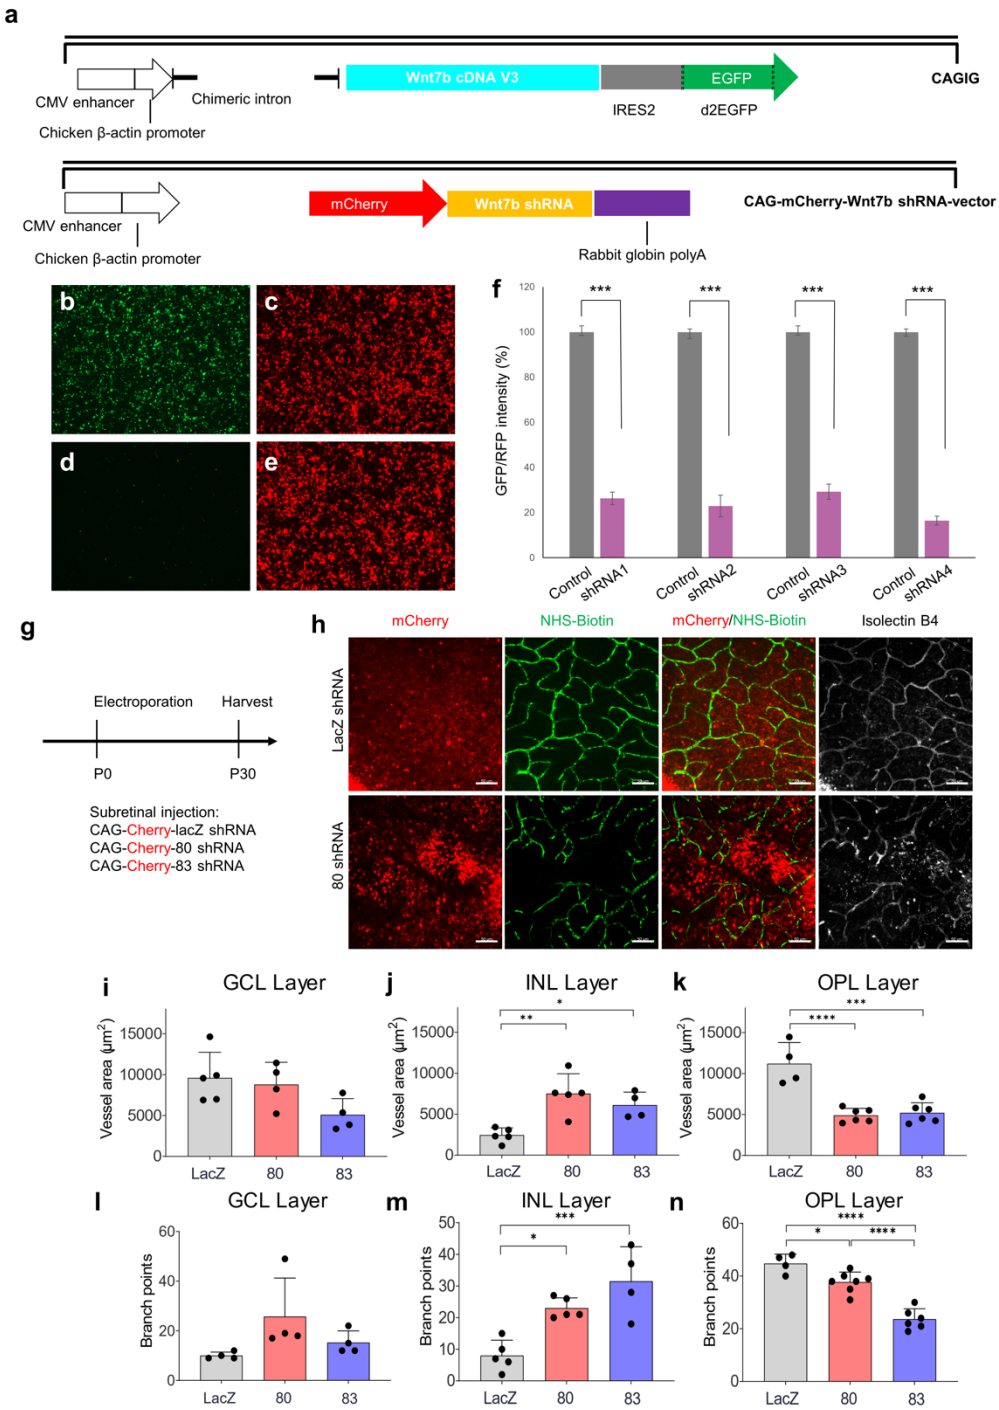

Figure 5: Depletion of *Wnt7b* in developing retina induced defects in

vascular development.

(a) Schematic constructs of *Wnt7b* cDNA- and shRNA-expressing vectors. **(b-e)** Figures demonstrating HEK293T cells under GFP and mCherry channels after 24-hours' transfections. **b:** *Wnt7b* cDNA only. **c:** *Wnt7b* shRNA80 only. **d:** GFP channel of *Wnt7b* cDNA + shRNA80. **e:** mCherry channel of *Wnt7b* cDNA + shRNA80. **(f)** Knockdown efficiencies of each shRNA after 24 hours transfection was calculated. Gray column: The ratio of cDNA only GFP fluorescent intensity over shRNA only mCherry fluorescent intensity. Magenta column: the ratio of GFP over the mCherry fluorescent intensity after a combo of cDNA and shRNA transfections. Statistical differences were determined by Two-tailed T-test. *** represents a significant difference P-value <0.001. **(g).** The experimental design. **(h).** Knock down of *Wnt7b* (shRNA 80) led to abnormal vessels in the retina in the outer plexiform layer (OPL). NHS-Biotin: tracer for vessels. Isolectin B4: vascular cell marker. **(i-n).** Quantification of vessels in superficial, intermediate and deep vascular plexuses in the retina. Vessel area and branch points were quantified. Data are mean ± SEM (n = 3 in control group, six lobes of retina were included; n = 2 in shRNA 80 group, four lobes of retina were included; n = 2 in shRNA 83 group, four lobes of retina were included). Statistical differences were determined by one-way ANOVA with a post Tukey's multiple comparisons test, comparing the shRNA 80 and 83 groups with control group. *p < 0.05, **p < 0.01, *p < 0.001 and ****p < 0.0001. Scale bar: 50 μm.

To validate the candidate genes identified from our vasculature GWAS, we sought to test if *Wnt7b* functions in retinal vessel development since it was not considered related to vessel development in the retina. Four shRNAs (shRNA80-83) targeting *Wnt7b* were designed and tested in cell culture for their knockdown efficiency. Greater than 70% knockdown efficiency was observed for all shRNAs with shRNA83 having the highest knockdown efficiency (Fig 5a-f). To test the function of *Wnt7b* in the retina, shRNA80 and shRNA83 were electroporated into mouse retinas *in vivo* at postnatal day 0 (P0) as described[46] (Fig 5g). Thirty days after electroporation, the retinas were harvested following injection of NHS-biotin vessel tracer[47]. When *Wnt7b* was knocked down in the retina, the total vessel area was significantly increased in the intermediate vascular plexus and reduced in the deep vascular plexus (Fig 5h-k). Similarly, the vessel branch-points were significantly increased in the intermediate and superficial plexuses and decreased in the deep vascular plexus when *Wnt7b* was depleted from the developing retina (Fig 5l-n). These results demonstrate that *Wnt7b* plays an important role in normal retinal vascular development.

Discussion

Our work is the first proof-of-concept of a self-supervised learning-based phenotype discovery method for imaging GWAS. With no expert supervision, our method was able to extract endophenotypes and identify genes relevant to retinal vessel development, including a new locus that was validated experimentally.

While there have been previous imaging GWAS studies on DL-based phenotyping, they either used expert-defined phenotypes[7] or clustering of dense representational vectors into subtypes. We directly use the dense vectors, which contain more information than the subtype cluster labels, as phenotypes. There are a handful of studies that use final or intermediate layers of the neural network as phenotypes, but these networks were trained in a supervised fashion using external labels (e.g., age[48] or eye diseases[49]).

Our iGWAS framework is flexible and can be adapted and extended in various ways. To study retinal vasculature, our current approach includes a segmentation step that generates vessel masks, and the embeddings are subsequently derived from these predicted vessel masks. To capture additional information in fundus images, such as the morphology of the optic disc, hemorrhages, exudates, or the pigmentation level, alternative preprocessing/segmentation steps may be applied, or this process can be completely skipped. Also, while the pair of eyes of an individual are natural "biological replicates" for our ArcFace-like approach, our approach may be extended to images without replicates, via current approaches for contrastive learning[14]. Furthermore, to inject labels to make more specific phenotypes, one can use a hybrid approach that minimizes both supervised and self-supervised losses.

Admittedly, while we prioritize the proof-of-concept, there is room for further methodological improvements. For example, it is not completely optimized to use the 128-dimensional vector as phenotypes. Additionally, the vasculature phenotypes are also correlated with retina color and with each other. Retina color may influence the performance of the automatic processing pipeline. For example, two recent GWAS[8,50] on AI-based optic nerve head morphology from fundus images and the GWAS on retinal vessel tortuosity[29] also identified the *HERC2*/*OCA2* locus. In addition, other factors such as retinal background texture were not considered. More sophisticated representation learning with disentanglement may be used to control for these correlations[51,52]. Also, our retina color GWAS uses RGB color, which is susceptible to change in illumination. Defining retina color in other color spaces may further improve the detection power.

The benefit of self-supervised-learning-derived phenotypes is that no external training labels are required. This frees up the burden of complicated and expensive labeling and makes our approach applicable to any large collection of

images. As we leverage big data sets to improve our understanding of diseases, self-supervised methods are needed to efficiently extract meaningful information from medical images. We predict that iGWAS as a general phenotype discovery approach will be a fruitful research avenue.

Methods

Data set extraction

The DRIMDB dataset[53], was downloaded on 2018/11/26 from https://www.researchgate.net/publication/282641760_DRIMDB_Diabetic_Retinopathy_Images_Database_Database_for_Quality_Testing_of_Retinal_Images. We used it as part of the training set to train the quality control network as it contains images with quality labels. It contains 69 bad quality fundus images and 125 good quality fundus images.

Multiple datasets (accessed in 2018/11/02), including ARIA[54] (sample number: 143), CHASEDB1[55] (sample number: 28), DRIVE[55,56] (sample number: 40), HRF[57] (sample number: 45), IOSTAR[57,58] (sample number: 30) and STARE[57–60] (sample number: 40), containing fundus images and their corresponding vessel segmentation masks, were used to train the vessel segmentation network.

The Messidor dataset[61] (accessed in 2018/11/02) contains 1200 eye fundus color images that were acquired using a color video 3CCD camera mounted on a Topcon TRC NW6 non-mydriatic retinograph with a 45-degree field of view. The dataset was used to validate the performance of our embedding method.

The EyePACS dataset (accessed in 2018/11/02) was downloaded from Kaggle.com. It contains fundus images from both healthy subjects and subjects with different grades of diabetic retinopathy. 35,126 Kaggle training set images and 53,576 Kaggle test set images were combined. The demographic characteristics including age, sex, and ethnicity of individual images were undisclosed.

The UK Biobank data was accessed via approved project 24247. We conducted our analysis on over 65,629 British White (self-reported white British (field: 21000) and genetically identified as Caucasian (field: 22006)) participants from the UK Biobank who had fundus images available (field: 21015 and 21016). For each participant, we chose the first image for each eye, resulting in 130,967 images. Genetic data as genotyped by Applied Biosystems UK BiLEVE Axiom Array (field: 22438) and imputed (field: 22828)[62] were downloaded. The fundus images in the UK Biobank data were taken using the TOPCON 3D OCT 1000 Mk2 alongside with the optical coherence tomography (OCT) imaging data. The

data were collected in two phases: the initial assessment visit (2006-2010) at which participants were recruited and consent given and the first repeat assessment visit (2012-13). The size of each fundus image is 1536x2048 pixels.

Image quality control

We trained a neural network to automatically assess the quality of the fundus images. Since the DRIMDB does not contain enough labelled images, we manually labeled 1,000 fundus images of good and bad quality from the EyePACS dataset and combined them with the DRIMDB dataset as the training set. An Inception v3 network[25] pretrained on ImageNet was downloaded and fine-tuned to classify qualities of different samples with early stopping. The quality assessment network outputs a score between 0 (bad) and 1 (good) to indicate the quality of the image, and it was trained using cross entropy loss. An image was defined as good quality if the output quality score of the network from that image was greater than 0.5.

The performance of the quality assessment network was validated on a subset of UK Biobank fundus images taken from white British subjects with diabetes mellitus (n=7,683). A previously validated procedure was used to determine DM status based on self-reported DM diagnosis, use of DM medications and presence of DM complications[63]. We also used HbA1c $\geq$ 6.5% as a criterion for identifying DM. Two ophthalmologists were asked to grade the image for the stage of diabetic retinopathy and determine if an image is of bad quality. A fundus image in this subset was classified as bad quality if both graders agreed that the quality of the image is poor. Comparing with this ground truth, the quality assessment network reached an AUC ROC of 92.14%. At 0.5 threshold, the positive predictive value was 0.9832, the negative predictive value was 0.4916, the sensitivity was 0.7155, and the specificity was 0.9574.

Fundus image segmentation

The graph cut algorithm[64] from OpenCV[65] was implemented to separate the foreground and background, and the foreground image was cropped and then resized to 224×224. The segmentation network is a specialized U-net, ESPNet[66]. ESPNet is a lightweight memory friendly architecture that allows a whole batch of images to be processed at the same time. We modified the ESPNet for the vessel segmentation task. While a typical segmentation output is a binary image mask of the original image size, we chose to use an un-binarized grey-scale image output in order to retain more details of the small vessels. The quality of the segmentation was confirmed by visual inspection. The segmentation model was trained on a combined dataset (n=306) of ARIA[54], CHASEDB1[55], DRIVE[55,56], HRF[57], IOSTAR[57,58] and STARE[57–60] for 500 epochs using the Adam optimizer[67], with a learning rate of 1×10^{-4} using the cross-entropy loss. Then, the model with the best test Dice score[68] was selected. On the held-out test set, the Dice score of the trained model reached 0.77.

Embedding neural network

The retinal vessel segmentation images were fed to a network that also uses the Inception v3[25] backbone to produce a 128-dimensional embedding vector. The final fully connected layer of the Inception v3 network was replaced to produce a 128-dimensional vector. We adopted an approach similar to ArcFace[24]: Each subject is assigned a template embedding and the network is trained to minimize the angle between embeddings of different photos of a subject and his/her template while maintaining a margin between embeddings of a specific photo and templates of different subjects. Specifically, our loss function is:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\log\frac{e^{s\cos(\theta_i+m)}}{e^{s\cos(\theta_i+m)} + \sum_{j=1,j\neq i}^{N} e^{s\cos\theta_j}}$$

where N is the number of samples, j is the angle between the output of the network and the template of the jth sample, m is the margin, and s is the inverse temperature scaling factor. In our study, m is set to be 30 and s is set to be 0.5, which are the best performing hyperparameters on multiple face recognition datasets.

The embedding network was trained using 40,000 images from the EyePACS database (https://www.eyepacs.com/). The quality control network was used to score each image, and the top-ranked 40,000 images were taken. The right eye images were flipped for preprocessing, and random augmentations were applied to add robustness. The training-testing split was 80/20. We also trained the embedding network with an additional task of classifying the grade of diabetic retinopathy. The weight ratio of these two tasks was 10 to 1. The network was trained using Adam optimizer[67] with a learning rate of 1×10^{-4} for 500 epochs, and the model with the lowest test set loss was selected. Source code is available at https://github.com/ZhiGroup/iGWAS.

Vasculature endophenotype GWAS

The genome-wide scans for UK Biobank were conducted over 658,720 SNPs that were directly genotyped by arrays (field: 22438). To control for confounding factors due to ethnicity, we only included individuals of British white ethnicity (self-reported white British (field: 21000) and genetically identified as Caucasian (field: 22006)). The sample size was 65,629. We used all 130,967 images from this cohort without applying image quality control because fewer loci were identified otherwise (data not shown). The GWAS was performed with BOLT-LMM (Version 2.3.4)[28] on all 128 dimensions of the embedding vector using the linear mixed model association method (BOLT_LMM_INF) with age, sex, and the first 10 ancestral principal components as covariates. Each variant had 256 p-values, 128 for the left and 128 for the right fundus images. A variant was

selected if the minimum of the left 128 p-values and the minimum of the right 128 p-values both passed a threshold of 5×10^{-8}. The association signals from selected variants from both eyes and across endophenotypes that were merged into single independent loci if they are in linkage disequilibrium (r^2>0.2) or within 200kb from each other. For each loci group, the variant with the most significant p-value was deemed the lead SNP, and the imputed genotypes of SNPs (field: 22828) at the hit region, 200kb upstream and 200kb downstream of the lead SNP, were extracted for dense mapping. To generate consistent association results without rerunning the imputed SNPs across the entire genome, the imputed SNPs at each hit region were concatenated with the non-imputed SNPs on other chromosomes then fed into the BOLT-LMM. LocusZoom (legacy version, accessed between 2018 and 2021)[69] was used to generate localized annotations of the significant loci groups.

Extracting retina background color and color GWAS

The traits of retina color were created as follows. The size of each UK Biobank fundus image is 1536x2048 pixels. Right eye fundus images were first flipped before cropping. The center patch of size 400x400 pixels around the fovea region, [600:1000, 800:1200], was cropped, and the average intensities of each of 3 channels (red, green and blue) in this patch were taken as the quantitative traits. Since the fundus images of UK Biobank are mostly aligned as they are taken with unified protocol, the patches at the same location were comparable. The 3 color variables were adjusted for in color-adjusted iGWAS. In addition, the GWAS analyses were done on the same cohorts and using the same pipeline as in vasculature endophenotype GWAS.

Genetic correlation

The genetic correlations were calculated using LDSC software (v1.0.1)[70]. Besides computing genetic correlation within the endophenotypes, we additionally selected several traits (Supplementary Table 7) to probe the endophenotypes. We counted the number of iGWAS loci that overlapped with traits from the GWAS catalog (Supplementary Fig 6). The selection criteria were: (1) The previous GWAS hits of the trait fall within any iGWAS loci more than twice or the traits are related to retinal or corneal disorders, and (2) The summary statistics of the trait are available from either https://alkesgroup.broadinstitute.org/UKBB/UKBB_409K/ or http://geneatlas.roslin.ed.ac.uk. To our knowledge, these are the only publicly available summary statistics computed by linear mixed models.

Querying GWAS catalog

For each independent locus, the range from the first to the last significant SNP was first transformed using LiftOver, then a range query was performed with the range plus 200kb flanking regions on the GWAS Catalog database to identify previous associations.

Functional validation of candidate gene involved in retina vessel development

Electroporation was performed as previously described[71]. To label retinal vessels, mice were deeply anesthetized by a sustained flow of isoflurane (3% isoflurane at 2 L/minute mixed with pure oxygen). Sulfo-NHS-LC-biotin molecules (~226 Da, Thermo Fisher, 21335) were injected (0.5mg/g body weight) into the left ventricle of mouse hearts with a 31-gauge insulin syringe. The heartbeat was continuous for 10 minutes to ensure good circulation of the tracer. Eyes were enucleated and fixed in fresh 4% PFA for 30 minutes on ice. Retinas were dissected out in PBS, fixed overnight in 4% PFA, and washed three times in PBS the following day. Retinas were incubated with blocking buffer (containing 3% Triton X-100, 0.5% Tween 20, 1% BSA, and 10% donkey serum in PBS) at 4°C overnight. For immunohistological staining, retinas were incubated with isolectin B4 conjugated to Alexa Fluor® 647 (Molecular Probes™; I32450; 1:500) for 24 hours and washed five times in PBS. Retinas were then incubated with streptavidin conjugated to Alexa Fluor™ 488 (Thermo Fisher Scientific™; NC0186832; 1:2000) for 24 hours to detect the NHS-Biotin tracerand washed three times in PBS. Retinas were flat-mounted with Fluoromount-G® (SouthernBiotech; 0100-01) and imaged by confocal microscopy.

Acknowledgements

This work was supported by grants from the National Institute of Aging (AG070112-01A1) to Z.X., W.Z., L.G., H.C., and D.Z.. In addition, this work was supported by American Diabetes Association (1-16-INI-16 to S.W. , C.L., A.D. and M. W.), NIH 1R01EY03258501 (to S.W. , C.L., A.D. and M.W.), P30 to Stanford Ophthalmology (to S.W. , C.L., A.D. and M.W.). This work was also supported by grants from the National Eye Institute (EY022356, EY018571, EY002520), Retinal Research Foundation, and NIH shared instrument grant S10OD023469 to R.Chen. R Channa is supported by the grant from the National Eye Institute (1K23EY030911-01). This work was supported in part by an Unrestricted Grant from Research to Prevent Blindness to the UW-Madison Department of Ophthalmology and Visual Sciences. L.G. was also supported by the Translational Research Institute through NASA Cooperative Agreement NNX16AO69A, NIH grants UL1TR003167 and R01NS121154, and a Cancer

Prevention and Research Institute of Texas grant (RP 170668). We thank Kaitlyn Xiong for critical reading of the manuscript.

Author contributions

R.Chen and D.Z. conceptualized the concept of iGWAS and initiated the project. Z.X., D.Z., and L.G. designed the SSuPER Model. Z.X. implemented the model, designed and conducted image analysis. R.Channa performed image grading. T.Z., Z.X. and H.C. conducted GWAS analyses. Z.X., S.K., and W.Z. conducted post-GWAS analyses. J.L., R.Chen., S.W. , C.L., A.D. and M. W. conducted experiment validation. D.Z., Z.X., R.Chen. and S.W. wrote the manuscript. All authors discussed the results and commented on the manuscript.

References:

1. Pirruccello, J. P. *et al.* Deep learning enables genetic analysis of the human thoracic aorta. *Bioxiv* https://www.biorxiv.org/content/10.1101/2020.05.12.091934v1 (2020).
2. Marvao, A. de *et al.* Genome wide association analysis of the heart using high-resolution 3D cardiac MRI identifies new genetic loci underlying cardiac structure and function. *Journal of Cardiovascular Magnetic Resonance* vol. 18 (2016).
3. Zhao, B. *et al.* Common genetic variation influencing human white matter microstructure. *Science* **372**, eabf3736 (2021).
4. Taghavi Namin, S., Esmaeilzadeh, M., Najafi, M., Brown, T. B. & Borevitz, J. O. Deep phenotyping: deep learning for temporal phenotype/genotype classification. *Plant Methods* **14**, 66 (2018).
5. Wang, X. *et al.* High-throughput phenotyping with deep learning gives insight

into the genetic architecture of flowering time in wheat. *Gigascience* **8**, (2019).

6. Sommer, C. & Gerlich, D. W. Machine learning in cell biology - teaching computers to recognize phenotypes. *J. Cell Sci.* **126**, 5529–5539 (2013).
7. Pirruccello, J. P. *et al.* Deep learning enables genetic analysis of the human thoracic aorta. doi:10.1101/2020.05.12.091934.
8. Alipanahi, B. *et al.* Large-scale machine learning-based phenotyping significantly improves genomic discovery for optic nerve head morphology. *Am. J. Hum. Genet.* (2021) doi:10.1016/j.ajhg.2021.05.004.
9. Poplin, R. *et al.* Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat Biomed Eng* **2**, 158–164 (2018).
10. Mitani, A. *et al.* Detection of anaemia from retinal fundus images via deep learning. *Nat Biomed Eng* **4**, 18–27 (2020).
11. Varadarajan, A. V. *et al.* Deep Learning for Predicting Refractive Error From Retinal Fundus Images. *Invest. Ophthalmol. Vis. Sci.* **59**, 2861–2868 (2018).
12. van den Oord, A., Li, Y. & Vinyals, O. Representation Learning with Contrastive Predictive Coding. *arXiv [cs.LG]* (2018).
13. He, K., Fan, H., Wu, Y., Xie, S. & Girshick, R. Momentum contrast for unsupervised visual representation learning. in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 9729–9738 (2020).
14. Chen, T., Kornblith, S., Norouzi, M. & Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. in *Proceedings of the 37th*

International Conference on Machine Learning (eds. Iii, H. D. & Singh, A.) vol. 119 1597–1607 (PMLR, 2020).

15. Grill, J.-B. *et al.* Bootstrap your own latent: A new approach to self-supervised Learning. *arXiv [cs.LG]* (2020).
16. Giancardo, L., Roberts, K. & Zhao, Z. Representation Learning for Retinal Vasculature Embeddings". in 243–250 (2017).
17. J Yan, S Salazar Marioni, A Jagolino-Cole, R Channa, S Pachade, S A Sheth, L Giancardo. "Towards Stroke Biomarkers on Fundus Retinal Imaging: A Comparison Between Vasculature Embeddings and General Purpose Convolutional Neural Network." in *Conf Proc IEEE Eng Med Biol Soc.*
18. Tian, Y., Krishnan, D. & Isola, P. Contrastive Multiview Coding. *arXiv [cs.CV]* (2019).
19. Kramer, M. A. Nonlinear principal component analysis using autoassociative neural networks. *AIChE J.* **37**, 233–243 (1991).
20. Carneiro, G. *et al. Deep Learning and Data Labeling for Medical Applications: First International Workshop, LABELS 2016, and Second International Workshop, DLMIA 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 21, 2016, Proceedings*. (Springer, 2016).
21. Shin, S. Y., Lee, S., Yun, I. D. & Lee, K. M. Deep vessel segmentation by learning graphical connectivity. *Med. Image Anal.* **58**, 101556 (2019).
22. Fraz, M. M. *et al.* Blood vessel segmentation methodologies in retinal images--a survey. *Comput. Methods Programs Biomed.* **108**, 407–433

(2012).

23. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* 234–241 (Springer International Publishing, 2015).
24. Deng, J., Guo, J., Xue, N. & Zafeiriou, S. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) doi:10.1109/cvpr.2019.00482.
25. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the inception architecture for computer vision. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 2818–2826 (2016).
26. Schroff, F., Kalenichenko, D. & Philbin, J. Facenet: A unified embedding for face recognition and clustering. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 815–823 (2015).
27. Oh Song, H., Xiang, Y., Jegelka, S. & Savarese, S. Deep metric learning via lifted structured feature embedding. in *Proceedings of the IEEE conference on computer vision and pattern recognition* 4004–4012 (2016).
28. Loh, P.-R. *et al.* Efficient Bayesian mixed-model analysis increases association power in large cohorts. *Nat. Genet.* **47**, 284–290 (2015).
29. Tomasoni, M. *et al.* Genome-Wide Association Studies of retinal vessel tortuosity identify 173 novel loci, capturing genes and pathways associated with disease and vascular tissue pathomechanics. *bioRxiv* (2020)

doi:10.1101/2020.06.25.20139725.

30. Chappell, J. C., Darden, J., Payne, L. B., Fink, K. & Bautch, V. L. Blood Vessel Patterning on Retinal Astrocytes Requires Endothelial Flt-1 (VEGFR-1). *J Dev Biol* **7**, (2019).
31. Wang, H. U., Chen, Z. F. & Anderson, D. J. Molecular distinction and angiogenic interaction between embryonic arteries and veins revealed by ephrin-B2 and its receptor Eph-B4. *Cell* **93**, 741–753 (1998).
32. Erber, R. *et al.* EphB4 controls blood vascular morphogenesis during postnatal angiogenesis. *EMBO J.* **25**, 628–641 (2006).
33. Zamora, D. O., Davies, M. H., Planck, S. R., Rosenbaum, J. T. & Powers, M. R. Soluble forms of EphrinB2 and EphB4 reduce retinal neovascularization in a model of proliferative retinopathy. *Invest. Ophthalmol. Vis. Sci.* **46**, 2175–2182 (2005).
34. Daneman, R. *et al.* Wnt/ -catenin signaling is required for CNS, but not non-CNS, angiogenesis. *Proceedings of the National Academy of Sciences* vol. 106 641–646 (2009).
35. Shu, W., Jiang, Y. Q., Lu, M. M. & Morrisey, E. E. Wnt7b regulates mesenchymal proliferation and vascular development in the lung. *Development* **129**, 4831–4842 (2002).
36. Saleem, M. *et al.* Sox6 as a new modulator of renin expression in the kidney. *Am. J. Physiol. Renal Physiol.* **318**, F285–F297 (2020).
37. El-Darzi, N. *et al.* Studies of ApoD-/- and ApoD-/-ApoE-/- mice uncover the APOD significance for retinal metabolism, function, and status of

chorioretinal blood vessels. *Cell. Mol. Life Sci.* **78**, 963–983 (2021).

38. Jensen, R. A. *et al.* Novel Genetic Loci Associated With Retinal Microvascular Diameter. *Circ. Cardiovasc. Genet.* **9**, 45–54 (2016).
39. Veluchamy, A. *et al.* Novel genetic locus influencing retinal venular tortuosity is also associated with risk of coronary artery disease. *Arterioscler. Thromb. Vasc. Biol.* **39**, 2542–2552 (2019).
40. Simcoe, M. *et al.* Genome-wide association study in almost 195,000 individuals identifies 50 previously unidentified genetic loci for eye color. *Sci Adv* **7**, (2021).
41. Morgan, M. D. *et al.* Genome-wide study of hair colour in UK Biobank explains most of the SNP heritability. *Nat. Commun.* **9**, 5271 (2018).
42. Landi, M. T. *et al.* Genome-wide association meta-analyses combining multiple risk phenotypes provide insights into the genetic architecture of cutaneous melanoma susceptibility. *Nat. Genet.* **52**, 494–504 (2020).
43. Yasuda, M. *et al.* Familial acanthosis nigricans with the FGFR3 mutation: Differences of pigmentation between male and female patients. *J. Dermatol.* **45**, 1357–1361 (2018).
44. Cheverud, J. M. A COMPARISON OF GENETIC AND PHENOTYPIC CORRELATIONS. *Evolution* **42**, 958–968 (1988).
45. Sodini, S. M., Kemper, K. E., Wray, N. R. & Trzaskowski, M. Comparison of Genotypic and Phenotypic Correlations: Cheverud's Conjecture in Humans. *Genetics* **209**, 941–948 (2018).
46. Matsuda, T. & Cepko, C. L. Electroporation and RNA interference in the

rodent retina in vivo and in vitro. *Proc. Natl. Acad. Sci. U. S. A.* **101**, 16–22 (2004).

47. Chow, B. W. & Gu, C. Gradual Suppression of Transcytosis Governs Functional Blood-Retinal Barrier Formation. *Neuron* **93**, 1325-1333.e3 (2017).
48. Jonsson, B. A. *et al.* Brain age prediction using deep learning uncovers associated sequence variants. *Nature Communications* vol. 10 (2019).
49. Program Planner. https://www.abstractsonline.com/pp8/#!/9070/presentation/2029.
50. Han, X. *et al.* Automated AI labeling of optic nerve head enables insights into cross-ancestry glaucoma risk and genetic discovery in >280,000 images from UKB and CLSA. *Am. J. Hum. Genet.* (2021) doi:10.1016/j.ajhg.2021.05.005.
51. Higgins, I. *et al.* beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. (2016).
52. Higgins, I. *et al.* Towards a Definition of Disentangled Representations. *arXiv [cs.LG]* (2018).
53. Sevik, U., Kose, C., Berber, T. & Erdol, H. Identification of suitable fundus images using automated quality assessment methods. *JBO* **19**, 046006 (2014).
54. Bankhead, P., Scholfield, C. N., McGeown, J. G. & Curtis, T. M. Fast retinal vessel detection and measurement using wavelets and edge location refinement. *PLoS One* **7**, e32435 (2012).

55. Fraz, M. M. *et al.* An ensemble classification-based approach applied to retinal blood vessel segmentation. *IEEE Trans. Biomed. Eng.* **59**, 2538–2548 (2012).

56. Staal, J., Abràmoff, M. D., Niemeijer, M., Viergever, M. A. & van Ginneken, B. Ridge-based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imaging* **23**, 501–509 (2004).

57. Budai, A., Bock, R., Maier, A., Hornegger, J. & Michelson, G. Robust vessel segmentation in fundus images. *Int. J. Biomed. Imaging* **2013**, 154860 (2013).

58. Zhang, J. *et al.* Robust Retinal Vessel Segmentation via Locally Adaptive Derivative Frames in Orientation Scores. *IEEE Trans. Med. Imaging* **35**, 2631–2644 (2016).

59. Hoover, A. & Goldbaum, M. Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels. *IEEE Trans. Med. Imaging* **22**, 951–958 (2003).

60. Hoover, A., Kouznetsova, V. & Goldbaum, M. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Trans. Med. Imaging* **19**, 203–210 (2000).

61. Decencière, E. *et al.* Feedback on a publicly distributed image database: The Messidor database. *Image Anal. Stereol.* **33**, 231 (2014).

62. Bycroft, C. *et al.* The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**, 203–209 (2018).

63. Eastwood, S. V. *et al.* Algorithms for the capture and adjudication of

prevalent and incident diabetes in UK biobank. *PLoS One* **11**, e0162388 (2016).

64. Greig, D. M., Porteous, B. T. & Seheult, A. H. Exact MaximumA PosterioriEstimation for Binary Images. *Journal of the Royal Statistical Society: Series B (Methodological)* vol. 51 271–279 (1989).
65. BRADSKI & G. The OpenCV library. *Dr Dobb's J. Software Tools* **25**, 120–125 (2000).
66. Mehta, S., Rastegari, M., Caspi, A., Shapiro, L. & Hajishirzi, H. Espnet: Efficient spatial pyramid of dilated convolutions for semantic segmentation. in *Proceedings of the european conference on computer vision (ECCV)* 552–568 (2018).
67. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. *arXiv [cs.LG]* (2014).
68. Sørenson, T. *A Method of Establishing Groups of Equal Amplitude in Plant Sociology Based on Similarity of Species Content and Its Application to Analyses of the Vegetation on Danish Commons*. (I kommission hos E. Munksgaard, 1948).
69. Pruim, R. J. *et al.* LocusZoom: regional visualization of genome-wide association scan results. *Bioinformatics* **26**, 2336–2337 (2010).
70. Bulik-Sullivan, B. *et al.* An atlas of genetic correlations across human diseases and traits. *Nat. Genet.* **47**, 1236–1241 (2015).
71. Wang, S., Sengel, C., Emerson, M. M. & Cepko, C. L. A gene regulatory network controls the binary fate decision of rod and bipolar cells in the

vertebrate retina. *Dev. Cell* **30**, 513–527 (2014).