transferGWAS: GWAS of images using deep transfer learning

Matthias Kirchler^{1,6*}, Stefan Konigorski^{1,2}, Matthias Norden^{3,4}, Christian Meltendorf⁵, Marius Kloft⁶, Claudia Schurmann^{2,3}, Christoph Lippert^{1,2*}

¹Digital Health - Machine Learning Research Group, Digital Health Center, Hasso Plattner Institute, University of Potsdam, Germany

²Hasso Plattner Institute for Digital Health at Mount Sinai, Icahn School of Medicine at Mount Sinai, New York, NY 10029, USA

³Digital Health & Personalized Medicine Research Group, Digital Health Center, Hasso Plattner Institute, University of Potsdam, Germany

⁴Department of Anesthesiology and Intensive Care Medicine, Charité – Universitätsmedizin Berlin, Berlin, Germany

⁵Department of Electrical Engineering - Mechatronics - Optometry, Beuth University of Applied Sciences Berlin, Berlin, Germany

⁶Department of Computer Science, TU Kaiserslautern, Germany

*Corresponding authors: matthias.kirchler@hpi.de, christoph.lippert@hpi.de

Abstract

Motivation:

Medical images can provide rich information about diseases and their biology. However, investigating their association with genetic variation requires non-standard methods. We propose *transferGWAS*, a novel approach to perform genome-wide association studies directly on full medical images. First, we learn semantically meaningful representations of the images based on a transfer learning task, during which a deep neural network is trained on independent but similar data. Then, we perform genetic association tests with these representations.

Results:

We validate the type I error rates and power of *transferGWAS* in simulation studies of synthetic images. Then we apply *transferGWAS* in a genome-wide association study of retinal fundus images from the UK Biobank. This first-of-a-kind GWAS of full imaging data yielded 60 genomic regions associated with retinal fundus images, of which 7 are novel candidate loci for eye-related traits and diseases.

Availability and Implementation:

Our method is implemented in Python and available at <https://github.com/mkirchler/transferGWAS/>

1 Introduction

The emergence of large-scale imaging-genetics datasets and deep learning provides us with the potential for powerful genome-wide association studies (GWAS) using image phenotypes. Current approaches to GWAS of imaging data have been largely focused on extracting fixed quantities of interest from the images, such as organ volumes, distances, and sizes. Such quantities have been obtained manually (e.g., segmenting organ volumes using graphics software), by automated software tools¹, or deep learning models trained to predict

segmentation masks or other biomarkers.²⁻⁷ However, such methods are limited in the scope of questions they can address and prohibit the detection of associations between unspecified variation of the studied organ and the genetic variation. Further challenges are that (i) deep learning methods require large amounts of labeled data, which can be very costly to produce, (ii) the extracted sets of biomarkers are generally study- and disease-specific^{8,9}, and (iii) deep learning-based feature extraction often leads to noisier estimates of the trait under investigation compared to hand-labeled and quality-controlled labels.¹⁰ Overall, it is unclear how to aggregate the complex geometric and topological structures present in images into a form beyond traditional biomarkers and useful in a GWAS setting.

While imaging GWASs based on specific extracted image features is valuable and important for scientific discovery, here we present a complementary approach called *transferGWAS*. Instead of extracting pre-specified phenotypes from images, *transferGWAS* treats the images themselves as phenotypes. This can be achieved with a version of transfer learning, which is well-established in deep learning.¹¹ By training a deep neural network on an independent dataset and extracting transformations of the original images at hidden layers of the network, we obtain low-dimensional, semantically meaningful representations.¹¹ These representations are continuous variables and can be analyzed using association tests for quantitative traits. A related line of work has sought to find general associations between histology images and genetic variants and gene expression using a deep-learning based version of canonical correlation analysis (CCA).^{12,13} However, the respective methods have not yet been adapted for a GWAS setting. We evaluate type I error rate and power of *transferGWAS* in Monte Carlo simulation studies, where, conditional on genetic effects, we generate images using generative adversarial networks (GAN). We apply *transferGWAS* to retinal fundus images from the UK Biobank and detail ways to understand and interpret the discovered associations.

2 System and Methods

transferGWAS is a multi-step method for using imaging data as phenotypes in GWAS. At its core, study images are condensed into lower-dimensional feature embeddings using a deep convolutional neural network (CNN) trained on an independent “transfer” dataset. CNNs can learn complex non-linear and layered continuous representations of the input data. Early layers usually extract simple information similar to edge detectors, while deeper layers learn more complex features.^{11,14} In each layer of the network, the representations become more sophisticated but also more specific to the learning task.¹⁴ Empirical evidence supports that deep neural networks learn useful representations of their input data that is transferable even to widely diverging tasks.¹¹ We show that the learned representations are to some degree task-agnostic, meaning that even though the CNN has to be trained on a specific task (such as predicting disease states from images), the representations capture considerably more information than what is necessary to perform the specific transfer task. In the context of two-sample testing, a recent study has shown both theoretically and empirically that a similar transfer-hypothesis testing approach yields valid p-values and exhibits high statistical power.¹⁵

transferGWAS consists of three main steps (Figure 1). First, a CNN is trained on an auxiliary dataset, called the “transfer dataset.” The transfer dataset can depict the same image modality as the study images in the next step but can also consist of non-medical images. At this stage, no genetic data is used. In the second step, the dataset of interest is analyzed, which must be independent of the transfer dataset regarding the observations. Using the CNN, the images (“study images”) available for the participants in the GWAS cohort are condensed into continuous low-dimensional phenotype representations. For this, the images are fed into the

network trained in the first step, after which the results are extracted at an intermediate, hidden layer of the network (instead of at the output layer) and then further reduced to ten dimensions via principal component analysis. Finally, genetic association testing is performed on the condensed phenotypes using linear mixed models. All steps can be adapted to fit the investigator's interests, and especially customization of the first two steps allows exploring different aspects of the data and indirect specification of prior knowledge.

2.1 Pre-training

We evaluate two different pre-training strategies. In both cases, we use the same convolutional neural network architecture, a ResNet50¹⁶. First, we use a ResNet50 trained on the ImageNet dataset of everyday objects (provided by the PyTorch library¹⁷). Second, we train another ResNet50 with randomly initialized weights on the EyePACS dataset of retinal fundus images.^{18,19} We perform this training in a multi-task manner to predict different stages of diabetic retinopathy (DR) and reconstruct the original image from the bottleneck layer simultaneously. Both tasks share a common CNN body, but while for the DR prediction task the head of the network is a simple linear layer, for the image reconstruction task the head is a convolutional decoder network as in an autoencoder (Supplementary Figure 6). This allows the network to be primed to disease-related aspects while capturing as much information about the original image as possible. Although the network is explicitly trained to detect DR, the hidden layer embeddings are also expected to capture very different characteristics. This is partly because the network needs to distinguish DR from other eye diseases efficiently, and hence implicitly learns versatile eye attributes. For training the DR model, we use the Adam optimizer²⁰ with a learning rate set to 10^{-3} , batch size 25, and the mean squared error loss function for both reconstruction and DR prediction (weighed equally). We train the network for 100 epochs, using the OneCycle learning rate schedule (using the weights from the epoch with the lowest validation loss). To avoid overfitting, images are randomly rotated, color-jittered, and flipped during each training step ("data augmentation").

2.2 Feature Condensation

The ResNet50 architecture is organized into one simple convolutional layer at the beginning and four subsequent multi-layer blocks, usually also called layers. We extract image representations from the fourth (i.e., last) and second of the larger layer blocks for the ImageNet and EyePACS-trained networks, respectively. In both cases, we use the representation at the end of the layer block directly before applying the last batch normalization and rectified linear unit (ReLU) activation. We perform eight deterministic image transformations (rotation by 0°, 90°, 180°, and 270°, each both with and without a horizontal mirroring), feed all images individually into the network, and stack the outputs into a single feature vector. This is a version of test-time augmentation (TTA) shown to improve the robustness of deep learning models on prediction tasks.²¹ For each individual, we stack the corresponding feature vectors of both left and right eye. Since we extract features at convolutional layers, the corresponding embeddings contain spatial information, which we average out ("global average pooling"). At this stage, the dimensionality of the feature representations is $n_channels * 8 * 2$, where $n_channels$ is 512 for the second layer block and 2048 for the last layer block of a ResNet50 architecture. We further reduce the dimensionality via a principal component analysis to the first ten PCs, which accounted for 52.5% of the explained variance on the ImageNet-trained network. For the EyePACS-trained network, the first ten PCs accounted for about 98.6% of the explained variance of the full embedding phenotype; see Supplementary Figure 7a and Section 2.4.

2.3 Genetic Association Testing

We perform genetic association tests in a linear mixed model (LMM) framework. There are different possible approaches to handle multivariate phenotypic data in an LMM framework, such as multivariate LMMs^{22,23} or other approaches²⁴. In this paper, we perform a univariate association test for each of the ten feature dimensions from Section 2.2 and then compute the minimum p-value of all ten dimensions, Bonferroni-corrected by a factor of 10 to adjust for multiple testing of the ten feature dimensions. This procedure controls the family-wise error rate and is conservative since it does not account for the relationships between different feature dimensions. We use the infinitesimal model with leave-one-chromosome-out scheme in the BOLT-LMM software package^{25–27}. To adjust for confounding, we include age, sex, indicator variables for genotyping microarray and assessment center, and the first ten genetic PCs as provided by the UK Biobank as covariates. Most feature dimensions were approximately normally distributed, but a small number of dimensions were skewed in one direction, violating the Gaussianity assumption of BOLT-LMM. To counteract this effect, we apply a rank-based inverse normal transformation (INT)²⁸ for all feature dimensions before fitting the LMM. We first compute residuals from fitting only the fixed-effect covariates (age, sex, genotyping array, assessment center, and first ten genetic PCs) in a linear regression model, adjust those residuals with the INT and finally use BOLT-LMM on those adjusted residuals while again correcting for the covariates, as proposed by Sofer et al.²⁹. Notably, the association results were virtually identical with and without INT adjustment on the first ten feature dimensions both for ImageNet- and EyePACS-pre-training. However, for a small number of higher-order feature dimensions, the INT suppressed many likely false-positive association results (Section 2.4).

We perform the GWAS on the 16,255,985 SNPs from the imputed genotype data provided by the UK Biobank (Supplementary Section 1.3). To estimate the genetic relationship matrix, we only use a subset of the directly genotyped variants, which accounted for 557,891 SNPs (BOLT-LMM recommends using ~500k genotypes) after filtering for $MAF \geq 0.1\%$, Hardy-Weinberg equilibrium (significance level: 0.001), pairwise LD pruning ($R^2 = 0.8$) and maximum missingness of 10% per SNP and individual. For LD Score regression to calibrate the BOLT-LMM test statistics, we used the European-ancestry scores supplied by the library. For interpretation, we group the associated variants in regions, which were defined as 500kb genomic segments carrying at least one associated SNP with $p\text{-value} < 1 \cdot 10^{-9}$.

2.4 Hyperparameter Selection in *transferGWAS*

For the ImageNet-trained network, we select the penultimate layer (after the last convolution operation, before batch normalization and ReLU activation), as is common in many transfer settings. Since ImageNet contains 1,000 different classes, the resulting image representation is sufficiently diverse and general to yield useful results. For the CNN trained on the EyePACS dataset, the training task is very narrow. To find more than only diabetic-retinopathy related traits, we choose the end of the second layer (of the four convolutional layer-blocks in the ResNet50 architecture).

We select other hyperparameters based on a parameter sweep in a simulation study, designed analogously to the validation study for empirical type I error rate and power (Section 3). The number of principal components had relatively little influence on the power in the simulation study (as long as it was at least 5 PCs), and we fix it to ten. We validate this behavior on the real datasets (Supplementary Figure 7b). Without correcting skewed distributions via INT (Section 2.3), we found (in data not shown here) that a single feature dimension (#47) in

EyePACS pre-training had a large number of likely false-positive SNP hits. Those were mostly one- or two-SNP loci with $MAF < 0.5\%$ and relatively low imputation quality, spread over 19 different chromosomes. These clusters of likely false-positive hits disappear when correcting the feature dimensions with INT while giving similar association results for other feature dimensions. The remaining hyperparameters we select to maximize empirical power in the simulation hyperparameter sweep. They were: image size (set to 448); whether or not to use test-time augmentation (set to use test-time augmentation); and whether to pool the spatial output at the convolutional layers via maximum or mean pooling (set to mean). We also considered using the hyperparameter sweep to select the optimal layer. However, the best layer for simulated images might very well not coincide with the best layer in the realistic scenarios of interest. This is because image features generated by the StyleGAN2 might operate on different spatial scales and can not be expected to capture significant pathologic patterns.

2.5 Interpretation of black-box results

Association results returned by *transferGWAS* cannot directly be traced back to specific phenotypic abstractions. To overcome this drawback, we apply several explanatory approaches to interpret the association results. First, we visually inspect the individual feature dimensions by contrasting images that strongly activate the respective dimensions in either direction. Next, we cross-correlate each individual feature dimension with all phenotypic traits collected in the UKBiobank in a phenome-wide association study (PheWAS, Section 2.6) and compare association results to previous findings in the GWAS catalog³⁰. Finally, we can also explore individual association results by training a CNN to predict the genotype of a specific SNP and then contrasting high- and low-activating images for this network. Namely, we train a ResNet18 architecture to predict the genotype from the left retinal fundus image. Images are resized to 448x448 pixels and randomly rotated by up to 20° during training. Assuming an additive genetic model, we code labels continuously as 0, 1, and 2. We use the Pearson correlation coefficient between predictions and true genotypes in each mini-batch as loss function. Using ImageNet pre-training, we further train the network for 10 epochs (with early stopping) with Adam optimizer and OneCycle learning rate scheduler, learning rate of 0.0001, and a batch-size of 64. We visualize the images in the validation dataset that most activate in negative or positive direction after training.

2.6 Phenome-Wide Association Study

For the PheWAS, we use the software tool PHESANT³¹ to perform automated phenome-wide association tests between all traits available in UK Biobank and the 10 extracted feature dimensions (after INT correction, Section 2.3). Association test results were available for 11,644 tested traits, including personal characteristics, lifestyle factors, body and organ measurements, and ICD10 codes from hospital visits. The association tests are based on linear, ordered logistic, multinomial, or logistic regression models depending on the scale of the outcome variable, with covariates age, sex, and assessment center. We use a Bonferroni correction to account for the $11,644 \cdot 10$ tests. For the characterization of the features and to compare the two different pre-trainings, we highlighted and manually explored unique and shared traits that showed significant associations. Due to the many different data categories (e.g., questionnaires, verbal interviews, physical measures, biological samples) some of the traits are overlapping and provide similar information. Therefore, we focus on two specific eye-related questionnaires in the dataset and repeat the association tests. We explicitly select the questionnaires on “Eye problems/disorders” and “Reasons for wearing glasses” to evaluate and compare the plausibility of the information captured by the different features.

3 Method validation in simulation study

To investigate the empirical type I error rate and power of *transferGWAS*, we perform a novel Monte Carlo simulation study. We generate synthetic images conditioned on 512-dimensional latent “code vectors.” These code vectors, in turn, are simulated from a small set of “causal” SNPs. We choose simulated causal SNPs uniformly at random from the first half of each chromosome (to estimate power) and count SNPs on the second half of each chromosome as null SNPs to evaluate type I error rate. To counteract proximal contamination³² of the null SNPs, we additionally leave a gap of 2 Mb before and after the center of each chromosome. The effect size of each causal SNP is drawn from a standard normal distribution and then normalized such that all causal SNPs together account for the specified variation in every single dimension of the code vectors. We fix the number of causal SNPs to 1,250 across all chromosomes (out of 275,142 SNPs in total on the first halves of all chromosomes), vary the explained variation of latent embeddings between 0.5 and 0.7, and also vary sample size between 6,000 and 46,731. We replicate each setting 10 times with different random seeds for a total of 80 genome-wide simulation runs. See Supplementary Table 1 for an overview of the scenarios. For the generation of synthetic images, we use the state-of-the-art StyleGAN2³³, a deep learning architecture for synthesizing high-resolution images from low-dimensional latent codes. We train StyleGAN2 on a subset of healthy retinal fundus images from the EyePACS dataset, an independent dataset of retinal fundus images used to detect DR¹⁸, using normally-distributed latent variables conditional on genetic variants as input. Then we perform association tests between the genetic variants and the synthetic images using *transferGWAS* with a ResNet50 architecture trained on ImageNet.

3.1 Results of simulation study

Supplementary Figure 1 shows a random selection of images generated by the StyleGAN2. The images generally appear reasonably realistic. There is a small number of synthetic retinal images with no or two optic disks. The GAN also exhibits some amount of “mode collapse”: several images look very similar, most images are of the left eye, and a central artifact can be seen in most images.

In the explored simulation settings under the null hypothesis of no association between genetic variants and image phenotypes, the type I error rate was properly controlled: only two null SNPs reached genome-wide significance at $\alpha = 5 \cdot 10^{-8}$ (i.e., a type I error rate of 2.5%). Supplementary Figure 2a shows a QQ-plot of null SNPs for one exemplary simulation run. For analysis of the empirical power, we selected 1,250 causal SNPs, which accounted for 50% explained variance of the latent code. 12% of the 1,250 causal SNPs could be recovered with a sample size of 46,731 individuals. With 70% explained variance, the number of discovered causal SNPs increased to 24% (Supplementary Figures 2b and Supplementary Table 1).

4 Results

We present results from applying *transferGWAS* on retinal fundus images of 46,731 participants of European ancestry from the UK Biobank, see Supplementary Table 2 for descriptive statistics and Supplementary Section 1.3 for details on the dataset and quality control. To characterize the information captured by the different feature dimensions, an experienced ophthalmologist visually explored those retinal fundus images which “activated” the extracted feature dimensions the most in positive and negative directions.

4.1 Feature Characterization

For the visual characterization of the features, Figure 2 and Supplementary Figures 3 and 4 show those images that activate each of the 10 feature dimensions the most in positive and negative direction, based on the ImageNet pre-training and the EyePACS pre-training, respectively. The most apparent characteristics captured by the feature dimensions based on ImageNet pre-training are brightness and hue, and structural attributes such as local smoothness. Besides this, some clinically relevant disease states are visible. For example, in dimension 6, the images activating in negative direction indicate severe myopic fundus, while in dimension 9, the images activating in positive direction show typical signs of an atrophy of the retinal pigment epithelial layer. The visual observation of the images activating the features based on the EyePACS pre-training show considerably more disease-related characteristics. The positive-activating images of dimension 1 show signs of pigment deposits of the retinal pigment epithelial layer, which can be caused by different eye pathologies such as retinitis pigmentosa. In the positive-activating images for dimensions 3 and 4 and negative-activating images for dimension 8, again, typical signs of a myopic fundus are visible, while the negative-activating images in dimension 4 also exhibit emphasized choroidal vessels. The positive-activating images for dimension 7 show strong indications of peripheral retinal drusen. For dimension 10, opacities in the vitreous humor as seen in pathologies like asteroid hyalosis can be discerned in both eyes. The remaining dimensions (2, 5, 6, and 9) do not show clear pathological patterns. Dimensions 2 and 6 seem to capture low-quality and noise signals, such as opacities and overexposure. Note that the images were recorded without dilation, which may explain the large number of exposure and shadow artifacts.

The PheWAS results are shown in Supplementary File 1. For ImageNet pre-training and EyePACS pre-training, there were 205 and 266 significantly associated traits, respectively (Bonferroni-corrected). Each feature dimension was associated with between 19 and 124 traits (ImageNet pre-training) and between 1 and 139 traits (EyePACS pre-training). 179 of these traits were associated in both pre-trainings. Shared traits were related to ophthalmic information, containing specific eye diseases, history of eye surgery, results of different eye measurement procedures, and reasons for wearing glasses. In addition, associations with pigmentation traits, cardiometabolic traits, and lifestyle factors but also with technical parameters, blood cell counts, or lung function measurements can be observed. Of the 26 associated traits unique to ImageNet pre-training, two were eye disease/eye procedure traits, and three were other disease/medication-related traits. On the other hand, the 87 traits uniquely identified after the EyePACS pre-training included 21 eye disease/eye procedure traits and 12 other disease/medication-related traits. For a more specific comparison and accounting for various overlaps of the phenotypic traits, we focused on the reasons for wearing glasses and doctor-diagnosed eye problems and diseases as recorded by touchscreen questionnaires (Table 1). All 14 traits were associated with at least one feature dimension based on the EyePACS pre-training, and 9 traits were associated with at least one feature dimension based on the ImageNet pre-training.

4.2 GWAS

There was no evidence of unaccounted stratification in the GWASs (genomic control parameter λ_{gc} ranged between 0.999-1.087). We identified 48 and 31 regions associated with at least one feature dimension in the ImageNet and EyePACS pre-training, respectively. Using PLINK's clumping method, we identified 95 and 63 independent loci located in the 48 and 31 regions for ImageNet and EyePACS, respectively. They account for 117 signals that can be mapped to 60

genomic regions. Manhattan plots are shown in Figures 3 and 4, QQ-plots in Supplementary Figure 5; for the full list of associated loci, see Supplementary Files 2 and 3. We identified 19 regions that were significantly associated in both analyses based on the ImageNet and EyePACS pre-training. Furthermore, 23 regions were significant in one analysis and showed sub-threshold significance ($p\text{-value} < 10^{-4}$) in the other analysis. Additionally, we found 16 and 2 regions exclusively associated with at least one feature dimension in the ImageNet and EyePACS pre-training, respectively.

Inspecting the known associations for the 60 regions in the GWAS catalog, we found associations with eye diseases and related traits (including myopia, age-related macular degeneration, diabetic retinopathy, glaucoma, cataract, macular thickness, refractive error, intraocular pressure, corneal and optic nerve measurements) for 36/60 regions. For example, 8 of the regions had previously been associated with myopia. Of them, 7 were associated with ImageNet feature dimensions 2, 4, 6 or 10, and 6 of those were associated with EyePACS feature dimensions 4, 7 or 8, thereby supporting our PheWAS findings where the ImageNet and EyePACS feature dimensions 6 and 8, respectively, showed the strongest correlation with myopia. To further evaluate the myopia top hits, we performed a separate GWAS for myopia in UK Biobank using the same set of samples as for the *transferGWAS*. Out of 70 known myopia loci³⁴, we could replicate 46% in the myopia GWAS and 32% and 23% in the ImageNet and EyePACS GWAS, respectively, using Bonferroni correction ($p\text{-value} < 0.05/70$). Of the 12 genome-wide significant regions from our myopia GWAS, we found that 75% of them showed significant associations in both *transferGWAS* after Bonferroni correction ($p\text{-value} < 0.05/12$) and 17-25% reached genome-wide significance.

Inspecting the known associations for the 60 regions in the GWAS catalog further, 30 regions have previously been associated with anthropometric traits (e.g., body weight, height, waist-hip ratio, or waist circumference), 29 with cardiovascular and metabolic traits (e.g., blood pressure, lipid levels, hypertension, and coronary artery disease), 9 with diabetes and related traits, 7 with smoking status, and 5 with Alzheimer's disease or dementia. Beyond that, a large proportion of the regions ($N=17/60$) has previously been reported to be associated with eye and hair color and skin pigmentation, for example, the *OCA2-HERC2* region. Overall, 17 regions did not have previously reported associations with eye-related traits and diseases or pigmentation traits, and 7 regions did not have any previously reported associations with traits listed above.

The *GOLGA8M-APBA2* region (significantly associated in EyePACS and ImageNet analysis) is located close to the *OCA2-HERC2* region and has previously been shown to be associated with hair color. Using stepwise conditional analyses, we identified 3 independent signals within that region (Supplementary Table 3). Interestingly, the *APBA2* gene encodes for a neuronal adapter protein expressed in the brain that interacts with the Alzheimer's disease amyloid precursor protein (APP). Low levels of amyloid-beta proteins are biomarkers of Alzheimer's disease and can be detected in the eye.³⁵ Alternatively, the detection of loci associated with Alzheimer's disease may be due to artifacts because of failure to comply with instructions during image capture. Another interesting locus that has not previously been reported is the *FOXF2* locus (EyePacs, lead SNP rs9405472, intronic variant, MAF = 0.4279, $p\text{-value} 2.0 \cdot 10^{-11}$). *FOXF2* is mainly expressed in the lung, colon, and esophagus and has been implicated in anterior segment dysgenesis³⁶. Further genes in that region include additional genes from the forkhead family of transcription factors, including *FOXC1*, which plays a role in ocular development and has been implicated in glaucoma³⁷.

To better understand the found associations, we trained separate CNNs to predict the respective genetic variant from the fundus images (Figure 5). First, we investigated the *GNB3* locus with lead SNP rs5442. rs5442 is a missense variant (MAF = 0.070) in *GNB3* whose minor

allele A has been shown to be associated with a 30% increased risk of high myopia³⁸. Images activating in direction of the risk allele A show myopic fundus, which was not present in G-allele direction. Second, we inspected the *ACTN4* locus with lead SNP rs10415219 (MAF = 0.486). In the A-allele activating images, retinal venular tortuosity is clearly visible, a potential marker for risk factors for cardiovascular diseases³⁹. rs10415219 is in LD with rs1808382 ($R^2=1$), whose major allele has previously been reported to be associated with increased retinal venular tortuosity⁴⁰ which is in line with our findings. For other selected SNPs, the visual inspection was not always conclusive.

5 Discussion

In the application of *transferGWAS* to the UK Biobank, we identified 60 genetic regions associated with retina scans, including 7 novel candidate loci. The remaining loci contain signals for eye-related traits and diseases, anthropometric, cardiovascular and metabolic traits, smoking status, and Alzheimer's disease. One intriguing example of the power of *transferGWAS* is the identification of loci on chromosome 15 in the *OCA2-HERC2* locus, for which there is evidence of their implication in oculocutaneous albinism.^{41,42} Oculocutaneous albinism exhibits general symptoms which provides a straightforward explanation of why *transferGWAS* can identify it in retina scans by linking pigmentation and eye diseases. Other novel loci proposed by *transferGWAS* have also been linked to disease-relevant processes and are interesting candidates for follow-up studies, such as the *APBA2* locus on chromosome 15. The signals within that locus might represent novel independent loci. One possible explanation is the interaction of the translated protein of *APBA2* with the amyloid precursor protein APP, which is involved in Alzheimer's disease and related to eye diseases³⁵. Importantly, we demonstrated how combining PheWAS results, re-training of CNN classifiers for discovered SNPs, and subsequent visual inspection can reveal clinically relevant features such as myopic fundus and retinal venular tortuosity. Our results also confirm previous findings that a wide variety of physical signals and functions can be detected from retinal fundus images.⁴³

Comparing the two different pre-training approaches, the specialized pre-training on EyePACS yielded feature dimensions that capture more pronounced eye-specific and eye-disease specific information compared to ImageNet pre-training. In the GWAS, however, ImageNet pre-training resulted in considerably more association results. This apparent contradiction is probably due to the more general training task for ImageNet pre-training. Although EyePACS pre-training captures more information about eye diseases, ImageNet pre-training identifies more non-specific information related to traits such as pigmentation and cardiometabolics. In the UK Biobank disease-specific phenotypes are overrepresented compared to non-pathological traits; e.g., "eye color" is not recorded at all, while most diseases are redundantly recorded in multiple fields.

We also presented a novel way to validate imaging-based GWASs using synthetically generated images. In Monte Carlo simulation studies, we generated images conditional on genetic variants to investigate the empirical type I error rate and power of *transferGWAS*. This evaluation helped to validate *transferGWAS* and indicated that hypothesis testing within *transferGWAS* can be made with valid type I error control.

5.1 Conclusions

We presented *transferGWAS*, a novel approach to utilize full images in GWAS, yielding interesting insights that might be easily overlooked in GWAS of pre-selected biomarkers extracted from images. We showed how *transferGWAS* can incorporate prior expert knowledge and illustrated its use for the analysis of retinal fundus images. Besides this use case, *transferGWAS* can be extended to any imaging modality. It can be used on 2D image (e.g., histological tissue images), 3D volumetric image (e.g., MR images of the brain), or CINE data (e.g., MR videos of the cardiac cycle). Even more generally, our method can be adapted to work on other complex biomarkers, such as graphs (e.g., brain connectomes) and time-series (e.g., ECG or EEG).

Data availability

The individual-level genotype, covariates, and image data are available by application from the UK Biobank (<https://www.ukbiobank.ac.uk/register-apply/>). The images used for training are publicly available from <http://image-net.org/index> (ImageNet dataset) and <https://www.kaggle.com/c/diabetic-retinopathy-detection> (EyePACS diabetic retinopathy dataset).

Code availability

Code and trained neural networks underlying all analyses, including the simulation study, are available at <https://github.com/mkirchler/transferGWAS>. For the GWAS we used the BOLT-LMM library^{25,26}, and for the phenome-wide association scans we used the PHESANT software library³¹, both of which are publicly available.

Acknowledgements

We thank the participants of the UK Biobank study. This research has been conducted using the UK Biobank resource. This work was supported by the German Ministry of Research and Education (BMBF) [project numbers 01IS21069A, 01IS18051A, 031B0770E, and 01IS21010C]; the German Research Foundation (DFG) [project numbers KL 2698/2-1 and KL 2698/5-1]; and by the Carl-Zeiss Foundation to MKloft.

Competing interests

The authors declare no competing interests.

References

1. Fischl, B. FreeSurfer. *NeuroImage* **62**, 774–781 (2012).
2. Ji, Y. *et al.* Genome-wide and abdominal MRI data provide evidence that a genetically determined favorable adiposity phenotype is characterized by lower ectopic liver fat and lower risk of type 2 diabetes, heart disease, and hypertension. *Diabetes* **68**, 207–219 (2019).

3. Elliott, L. T. *et al.* Genome-wide association studies of brain imaging phenotypes in UK Biobank. *Nature* **562**, 210–216 (2018).
4. Satizabal, C. L. *et al.* Genetic architecture of subcortical brain structures in 38,851 individuals. *Nat. Genet.* **51**, 1624–1636 (2019).
5. Wen, C., Mehta, C. M., Tan, H. & Zhang, H. Whole genome association study of brain-wide imaging phenotypes: a study of the PING cohort. *Genet. Epidemiol.* **42**, 265–275 (2018).
6. Barry, J. D. *et al.* Histopathological image QTL discovery of immune infiltration variants. *iScience* **5**, 80–89 (2018).
7. Glastonbury, C. A. *et al.* Machine Learning based histology phenotyping to investigate the epidemiologic and genetic basis of adipocyte morphology and cardiometabolic traits. *PLOS Comput. Biol.* **16**, e1008044 (2020).
8. Chen, C. *et al.* Improving the generalizability of convolutional neural network-based segmentation on CMR images. *Front. Cardiovasc. Med.* **7**, (2020).
9. Schreier, J., Attanasi, F. & Laaksonen, H. Generalization vs. specificity: in which cases should a clinic train its own segmentation models? *Front. Oncol.* **10**, (2020).
10. Guenther, F. *et al.* Chances and challenges of machine learning-based disease classification in genetic association studies illustrated on age-related macular degeneration. *Genet. Epidemiol.* (2020) doi:10.1002/gepi.22336.
11. Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. How transferable are features in deep neural networks? in *Advances in Neural Information Processing Systems 27* (eds. Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D. & Weinberger, K. Q.) 3320–3328 (2014).
12. Ash, J. T., Darnell, G., Munro, D. & Engelhardt, B. E. Joint analysis of gene expression levels and histological images identifies genes associated with tissue morphology. *bioRxiv* 458711 (2018) doi:10.1101/458711.
13. Gundersen, G., Dumitrescu, B., Ash, J. T. & Engelhardt, B. E. End-to-end training of deep probabilistic CCA on paired biomedical observations. in *Uncertainty in Artificial Intelligence* 945–955 (PMLR, 2020).
14. Zeiler, M. D. & Fergus, R. Visualizing and understanding convolutional networks. in *Computer Vision – ECCV 2014* (eds. Fleet, D., Pajdla, T., Schiele, B. & Tuytelaars, T.) 818–833 (Springer International Publishing, 2014).
15. Kirchler, M., Khorasani, S., Kloft, M. & Lippert, C. Two-sample testing using deep learning. *AISTATS* (2020).
16. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. *2016 IEEE Conf. Comput. Vis. Pattern Recognit. CVPR* (2016).
17. Paszke, A. *et al.* PyTorch: an imperative style, high-performance deep learning library. in *Advances in Neural Information Processing Systems 32* (eds. Wallach, H. *et al.*) 8024–8035 (2019).
18. Cuadros, J. & Bresnick, G. EyePACS: an adaptable telemedicine system for diabetic retinopathy screening. *J. Diabetes Sci. Technol. Online* **3**, 509–516 (2009).
19. Diabetic retinopathy detection. <https://kaggle.com/c/diabetic-retinopathy-detection>.
20. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *ICLR Int. Conf. Learn. Represent.* (2015).
21. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. in *Advances in Neural Information Processing Systems 25* (eds. Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q.) 1097–1105 (2012).
22. Zhou, X. & Stephens, M. Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nat. Methods* **11**, 407–409 (2014).
23. Lippert, C., Casale, F. P., Rakitsch, B. & Stegle, O. LIMIX: genetic analysis of multiple traits. *bioRxiv* 003905 (2014).
24. Aschard, H. *et al.* Maximizing the power of principal-component analysis of correlated phenotypes in genome-wide association studies. *Am. J. Hum. Genet.* **94**, 662–676 (2014).

25. Loh, P.-R. *et al.* Efficient Bayesian mixed-model analysis increases association power in large cohorts. *Nat. Genet.* **47**, 284–290 (2015).
26. Loh, P.-R., Kichaev, G., Gazal, S., Schoech, A. P. & Price, A. L. Mixed-model association for biobank-scale datasets. *Nat. Genet.* **50**, 906–908 (2018).
27. Lippert, C. *et al.* FaST linear mixed models for genome-wide association studies. *Nat. Methods* **8**, 833–835 (2011).
28. Beasley, T. M., Erickson, S. & Allison, D. B. Rank-based inverse normal transformations are increasingly used, but are they merited? *Behav. Genet.* **39**, 580–595 (2009).
29. Sofer, T. *et al.* A fully adjusted two-stage procedure for rank-normalization in genetic association studies. *Genet. Epidemiol.* **43**, 263–275 (2019).
30. Buniello, A. *et al.* The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Res.* **47**, D1005–D1012 (2019).
31. Millard, L. A. C., Davies, N. M., Gaunt, T. R., Davey Smith, G. & Tilling, K. Software application profile: PHESANT: a tool for performing automated phenome scans in UK Biobank. *Int. J. Epidemiol.* **47**, 29–35 (2018).
32. Listgarten, J. *et al.* Improved linear mixed models for genome-wide association studies. *Nat. Methods* **9**, 525–526 (2012).
33. Karras, T. *et al.* Analyzing and improving the image quality of StyleGAN. *2020 IEEE Conf. Comput. Vis. Pattern Recognit. CVPR* (2020).
34. Tedja, M. S. *et al.* Genome-wide association meta-analysis highlights light-induced signaling as a driver for refractive error. *Nat. Genet.* **50**, 834–848 (2018).
35. Wright, L. M. *et al.* Association of cognitive function with amyloid- β and tau proteins in the vitreous humor. *J. Alzheimers Dis. JAD* **68**, 1429–1438 (2019).
36. McKeone, R., Vieira, H., Gregory-Evans, K., Gregory-Evans, C. Y. & Denny, P. Foxf2: a novel locus for anterior segment dysgenesis adjacent to the FOXC1 gene. *PLoS ONE* **6**, (2011).
37. Siggs, O. M. *et al.* Prevalence of FOXC1 variants in individuals with a suspected diagnosis of primary congenital glaucoma. *JAMA Ophthalmol.* **137**, 348–355 (2019).
38. Boutin, T. S. *et al.* Insights into the genetic basis of retinal detachment. *Hum. Mol. Genet.* **29**, 689–702 (2020).
39. Cheung, C. Y. *et al.* Retinal vascular tortuosity, blood pressure, and cardiovascular risk factors. *Ophthalmology* **118**, 812–818 (2011).
40. 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).
41. Dolinska, M. B. *et al.* Oculocutaneous Albinism Type 1: link between mutations, tyrosinase conformational stability, and enzymatic activity. *Pigment Cell Melanoma Res.* **30**, 41–52 (2017).
42. Donnelly, M. P. *et al.* A global view of the OCA2-HERC2 region and pigmentation. *Hum. Genet.* **131**, 683 (2012).
43. Poplin, R. *et al.* Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat. Biomed. Eng.* **2**, 158–164 (2018).

Tables

Table 1. Comparison of the results of association tests between eye-related traits and diseases and the feature dimensions obtained from the different pre-trainings.

Eye-related trait	ImageNet	EyePACS		Cases	
	Significant dimensions	P-value	Significant dimensions		P-value
Myopia*	1, 2, 3, 4, 5, 6', 7, 8, 9, 10	$7.9 \cdot 10^{-296+}$	1, 3, 4, 5, 7, 8'	$7.7 \cdot 10^{-293+}$	30.6%
Hypermetropia*	1, 2, 3, 4, 5, 6', 7, 8, 9, 10	$1.4 \cdot 10^{-58+}$	3, 4, 5, 7, 8'	$3.7 \cdot 10^{-57+}$	14.8%
Presbyopia*	1, 2, 4', 5, 6, 10	$3.9 \cdot 10^{-67+}$	1', 3, 4, 5, 8	$1.0 \cdot 10^{-63+}$	32.4%
Astigmatism*	1, 2, 4', 6	$1.9 \cdot 10^{-23+}$	1', 3, 8	$4.2 \cdot 10^{-38+}$	10.0%
Strabismus*	2, 3', 9	$8.1 \cdot 10^{-15+}$	7, 8'	$5.3 \cdot 10^{-14+}$	1.4%
Amblyopia*	2, 3', 7, 9	$2.4 \cdot 10^{-20+}$	1, 7, 8'	$1.4 \cdot 10^{-16+}$	2.7%
Other eye condition*		$2.4 \cdot 10^{-4}$	4'	$1.3 \cdot 10^{-6+}$	1.0%
Diabetes related eye disease**		$8.9 \cdot 10^{-6}$	1'	$1.3 \cdot 10^{-7+}$	0.8%
Glaucoma**		$7.3 \cdot 10^{-6}$	1'	$3.7 \cdot 10^{-7+}$	1.7%
Injury or trauma**		$1.0 \cdot 10^{-2}$	4'	$7.0 \cdot 10^{-8+}$	0.5%
Cataract**	1, 2, 4, 8, 10'	$2.0 \cdot 10^{-24+}$	1', 4	$4.2 \cdot 10^{-44+}$	4.6%
Macular degeneration**		$1.4 \cdot 10^{-5}$	4', 7	$3.5 \cdot 10^{-27+}$	1.1%
Other serious eye condition**	4'	$9.1 \cdot 10^{-11+}$	1', 4	$9.9 \cdot 10^{-16+}$	1.8%
No eye disease**	1, 2, 4', 8, 10	$2.0 \cdot 10^{-24+}$	1', 4	$2.6 \cdot 10^{-67+}$	90.4%

Shown are association testing results between specific eye-related traits and diseases and the 10 feature dimensions after ImageNet and EyePACS pre-training, respectively. For each trait, smallest p-values (Bonferroni-adjusted for 10 feature dimensions) and all significant feature dimension are shown.

*Traits are subcategories of the touchscreen question of reasons for wearing glasses/contact lenses. Available nonmissing values: n= 46,140. The subcategory "Don't know" (7% of answers) was not analyzed in the PheWAS.

**Traits are subcategories of the touchscreen question of eye problems/disorders. Available nonmissing values: n= 39,308.

+Significant Bonferroni-adjusted significance level $4.3 \cdot 10^{-6}$.

'Dimension with strongest association.

Figures

Step 1: Train deep CNN on independent transfer task

Step 2: Condense study images into low-dimensional representations

Step 3: Perform multivariate GWAS on condensed image representations

Figure 1. Overview of transferGWAS.

Overview of the transferGWAS approach which consists of three steps. First, a convolutional neural network is trained on an independent transfer task to prime the network. Second, transferGWAS uses the trained network and a principal component analysis to condense study images into low-dimensional embeddings. Last, a linear mixed model association analysis is performed on the image representations.

Figure 2. Illustration of information in the retinal fundus images captured by the feature dimensions. Shown are example images activating 2 selected feature dimensions: dimension 6 from ImageNet pre-training and dimension 8 from EyePACS pre-training. The remaining feature dimensions are shown in Supplementary Figures 3 and 4. Retinal fundus images are displayed for the left eye (upper row) and the right eye (lower row); left column (“most negative”) shows the six individuals activating the corresponding feature dimension the most in negative direction; right column (“most positive”) shows the six individuals activating the corresponding feature dimension the most in positive direction. Reproduced by kind permission of UK Biobank ©.

Figure 3. GWAS results of transferGWAS based on ImageNet pre-training.

Manhattan plot for the genome-wide association study based on ImageNet pre-training ($n=46,731$ individuals, 16,255,985 SNPs). Shown are $-\log_{10}$ -transformed p-values across chromosomes 1 through 22. For each SNP, the minimum p-value over the ten feature dimensions multiplied by 10 (Bonferroni correction across the feature dimensions) is plotted. For SNPs exceeding the significance threshold ($5 \cdot 10^{-8}$), the highest activating dimension is indicated by the corresponding marker. P-values below 10^{-100} (there are 225 SNPs with p-values $< 10^{-100}$ on chromosomes 5 and 15) are truncated and indicated with a red circle for visualization purposes.

Figure 4. GWAS results of transferGWAS based on EyePACS pre-training.

Manhattan plot for the genome-wide association study based on pre-training on the Diabetic Retinopathy task using the EyePACS dataset ($n=46,731$ individuals, 16,255,985 SNPs). Shown are $-\log_{10}$ -transformed p-values across chromosomes 1 through 22. For each SNP, the minimum p-value over the ten feature dimensions multiplied by 10 (Bonferroni correction across the feature dimensions) is plotted. For SNPs exceeding the significance threshold ($5 \cdot 10^{-8}$) the highest activating dimension is indicated by the corresponding marker. P-values below 10^{-100} (there are 292 SNPs with p-values $< 10^{-100}$ on chromosome 15) are truncated and indicated with a red circle for visualization purposes.

Figure 5. Illustration of associations found by transferGWAS.

Visualization of two exemplary lead SNPs. Only left eyes are shown and were used for training. rs5442 shows strong signs of myopic fundus in the A direction, while rs10415219 shows retinal venular tortuosity in A direction. Images are plotted according to the predictions of a specifically trained CNN, but true genotypes may differ.