medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Unsupervised representation learning improves genomic discovery and risk prediction for respiratory and circulatory functions and diseases

Taedong Yun[1,†], Justin Cosentino[2], Babak Behsaz[1], Zachary R. McCaw[2,◇], Davin Hill[3,4], Robert Luben[5,6], Dongbing Lai[7], John Bates[8], Howard Yang[2], Tae-Hwi Schwantes-An[7,9], Yuchen Zhou[1], Anthony P. Khawaja[5,6], Andrew Carroll[2], Brian D. Hobbs[4,10,11], Michael H. Cho[4,10,11], Cory Y. McLean[1,*,†], and Farhad Hormozdiari[1,*,†]

[*] Joint supervision.
[1] Google Research, Cambridge, MA 02142, USA.
[2] Google Research, Palo Alto, CA 94304, USA.
[3] Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 94304, USA.
[4] Channing Division of Network Medicine, Brigham and Women's Hospital, Boston, MA 02115, USA.
[5] NIHR Biomedical Research Centre at Moorfields Eye Hospital & UCL Institute of Ophthalmology, London EC1V 9EL, UK.
[6] MRC Epidemiology Unit, University of Cambridge, Cambridge CB2 0SL, UK.
[7] Department of Medical and Molecular Genetics, Indiana University School of Medicine, Indianapolis, IN 46202, USA.
[8] Verily Life Sciences, South San Francisco, CA 94080, USA.
[9] Division of Cardiology, Department of Medicine, Indiana University School of Medicine, Indianapolis, IN 46202, USA.
[10] Division of Pulmonary and Critical Care Medicine, Brigham and Women's Hospital, Boston, MA 02115, USA.
[11] Harvard Medical School, Boston, MA 02115, USA.
[◇] Present address: Insitro, South San Francisco, CA 94080, USA.
[†] Correspondence to: tedyun@google.com, cym@google.com, fhormoz@google.com

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Abstract

High-dimensional clinical data are becoming more accessible in biobank-scale datasets. However, effectively utilizing high-dimensional clinical data for genetic discovery remains challenging. Here we introduce a general deep learning-based framework, REpresentation learning for Genetic discovery on Low-dimensional Embeddings (REGLE), for discovering associations between genetic variants and high-dimensional clinical data. REGLE uses convolutional variational autoencoders to compute a *non-linear, low-dimensional, disentangled embedding* of the data with highly heritable individual components. REGLE can incorporate expert-defined or clinical features and provides a framework to create accurate disease-specific polygenic risk scores (PRS) in datasets which have minimal expert phenotyping. We apply REGLE to both respiratory and circulatory systems: spirograms which measure lung function and photoplethysmograms (PPG) which measure blood volume changes. Genome-wide association studies on REGLE embeddings identify more genome-wide significant loci than existing methods and replicate known loci for both spirograms and PPG, demonstrating the generality of the framework. Furthermore, these embeddings are associated with overall survival. Finally, we construct a set of PRSs that improve predictive performance of asthma, chronic obstructive pulmonary disease, hypertension, and systolic blood pressure in multiple biobanks. Thus, REGLE embeddings can quantify clinically relevant features that are not currently captured in a standardized or automated way.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

1 Introduction

High-dimensional clinical data (HDCD) provide a unique opportunity to reveal the genetic architecture of diseases and complex traits when coupled with biobank-scale genetic data [1, 2, 3, 4, 5, 6]. However, we lack statistical methods to fully utilize HDCD in genome-wide association studies (GWAS), as standard GWAS require the phenotype of interest to be encoded as a single scalar. Multiple methods have been developed to use HDCD in GWAS, but each has unique limitations.

A natural procedure to use HDCD in GWAS is to perform GWAS on every single data coordinate (e.g. time points or pixels). For example, prior work performed GWAS on each recorded point of electrocardiograms to identify its genetic architecture [7]. There are multiple shortcomings of this approach: 1) running GWAS on thousands of phenotypes can be prohibitively computationally expensive, 2) HDCD often have a correlation structure in which the actual number of degrees of freedom is much lower than the number of coordinates in the data, and 3) multiple hypothesis testing correction for highly correlated coordinates reduces statistical power [8, 9]. One popular approach to address these issues is to use principal component analysis (PCA) [10] on the HDCD and then perform GWAS on a subset of the principal components (PCs) [11]. However, PCA assumes a linear relationship between the raw HDCD and the underlying biological factors of interest, and does not explicitly model temporal or spatial structure of HDCD. This incomplete representation of the data can lead to suboptimal downstream genetic analysis.

Machine learning-based (ML-based) phenotyping uses HDCD as input to a supervised machine learning model to predict trait labels, and then performs GWAS using the model predictions as the target phenotype [3, 12, 6]. While ML-based phenotyping can augment standard GWAS on manually defined trait labels, the supervised model only learns signals related to the specific target trait. Additionally, for the common case in which the supervised model uses deep learning, many labeled examples may be required to achieve good performance.

The most common method for GWAS on HDCD uses a small number of expert-defined features (EDFs) extracted from the HDCD as the target phenotypes. For example, spirograms are a graphical representation of spirometry test results, a widely-used clinical test for lung function that measures airflow and volume over time [13, 14]. Spirograms can be summarized into EDFs including forced vital capacity (FVC), forced expiratory volume in the 1st second (FEV_1), FEV_1/FVC, peak expiratory flow (PEF) and forced mid-expiratory flow ($FEF_{25\text{-}75\%}$) [15]. Spirogram EDFs are used in clinical settings to diagnose diseases such as COPD [16, 17]. In another example, photoplethysmograms

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

(PPG) measure volumetric changes in peripheral blood circulation using infrared light. Previously studied EDFs of PPG include the presence (or absence) of a notch, the position of the notch, the position of the peak, the position of the shoulder, and the peak-to-peak time [18, 19, 20, 21, 22]. PPG EDFs are known to detect diseases such as coronary heart disease [18]. Spirograms and PPG EDFs are heritable, and GWAS on EDFs have helped identify the genetic architecture of lung [23, 24, 25] and circulatory function [26, 27, 28]. However, EDFs may not capture the entirety of biological factors encoded in spirograms or PPG, thus GWAS on these EDFs may not exploit the full potential of these HDCD.

To overcome these limitations we develop a principled method, REpresentation learning for Genetic discovery on Low-dimensional Embeddings (REGLE), that is computationally efficient, requires no labels, and can incorporate information from EDFs if they are available. We apply REGLE in two case studies to understand the genetic architecture of lung function from raw spirograms and circulatory function from PPG. Compared to GWAS on spirogram EDFs (e.g. FEV_1) and PPG EDFs (e.g., absence of notch), our GWAS on the learned encodings both recovers most known genetic loci linked to lung and circulatory function and also detects additional loci. We computed polygenic risk scores (PRSs) from GWAS on the learned encodings and evaluated their ability to discriminate disease risks. In the case of REGLE on spirograms, we observed strong enrichment for lung function, genetically causal links between the encodings and COPD and asthma, and significantly improved PRSs for asthma and COPD in multiple datasets and genetic ancestries. In the case of REGLE on PPGs, we observed that PRSs from GWAS on the learned encodings of PPGs can discriminate hypertension and systolic blood pressure in multiple datasets and produce stronger heart and blood vessel function enrichments. These results indicate that REGLE successfully extracts a meaningful low dimensional representation of lung function from spirograms and of circulatory function from PPG, which in turn improves genetic discovery for both organ systems.

2 Results

2.1 Overview of REGLE

REGLE consists of three main steps: 1) learning a non-linear, low-dimensional, disentangled representation (i.e. an encoding) of the HDCD, 2) performing GWAS on each encoding coordinate, and 3) using PRSs from the encoding coordinates as genetic scores of general biological functions, and potentially combining them to create a PRS for a disease or trait of interest (Figure 1).

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

In the first step, we train a variational autoencoder (VAE) [29] to compress and reconstruct HDCD (Figure 1 and Methods). Autoencoders consist of a pair of function approximators, typically called an encoder and a decoder, connected by a low-dimensional "bottleneck" layer. The encoder summarizes the input data efficiently into a small set of numbers represented at the bottleneck layer, and the decoder reconstructs the input data from the low-dimensional summary [30]. VAEs [29] are a special type of autoencoders that introduce stochasticity in the encoder. The VAE implicitly forces the learned encodings to be relatively disentangled [31], i.e. the encodings have relatively uncorrelated coordinates and separable biological factors can be better captured in each coordinate.

In addition, REGLE enables relevant EDFs to be optionally included in the input to the decoder of the model, so that the encoder is encouraged to learn only the residual signals not represented by the EDFs (Figure 1). This ability to incorporate prior knowledge of important data features (from users or clinicians) is one of the key advantages of REGLE.

In the second step, we perform GWAS independently on each learned encoding coordinate for all individuals (Methods). In the final step, we compute coordinate-specific PRSs that represent intermediate genetic scores of biological function, and linearly combine these coordinate-specific PRSs into a single disease-specific PRS by training on a small number of individuals with disease labels (Figure 1).

2.2 Overview of REGLE on spirograms

Spirograms are a graphical representation of clinical pulmonary function tests, typically represented by volume-time, flow-time, and flow-volume curves. Spirograms are used to diagnose respiratory diseases such as COPD and understand lung function [32, 24, 6]. To understand the genetic architecture of human lung function, we applied REGLE to obtain low-dimensional representations of spirogram curves, which we call spirogram encodings (SPINCs) (Figure 2). To construct SPINCs, we trained a convolutional VAE [29] to reconstruct spirograms (volume-time and flow-time; Figure 2a and Methods). In addition, we constructed another set of encodings we call *residual* spirogram encodings (RSPINCs) by injecting five EDFs (FEV_1, FVC, FEV_1/FVC, PEF, and $FEF_{25\text{-}75\%}$) into the input to the decoder of the VAE to reconstruct flow-volume (Figure 2a). We generated SPINCs and RSPINCs for all individuals (n=351,120) in UK Biobank [33] using their first visit spirogram, excluding individuals whose spirogram failed our quality control (QC) measures (Methods). We used 80% of the European-ancestry individuals (n=259,692) to train the (R)SPINCs models and 20% (n=65,266) to evaluate the reconstruction performance and choose hyperparam-

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

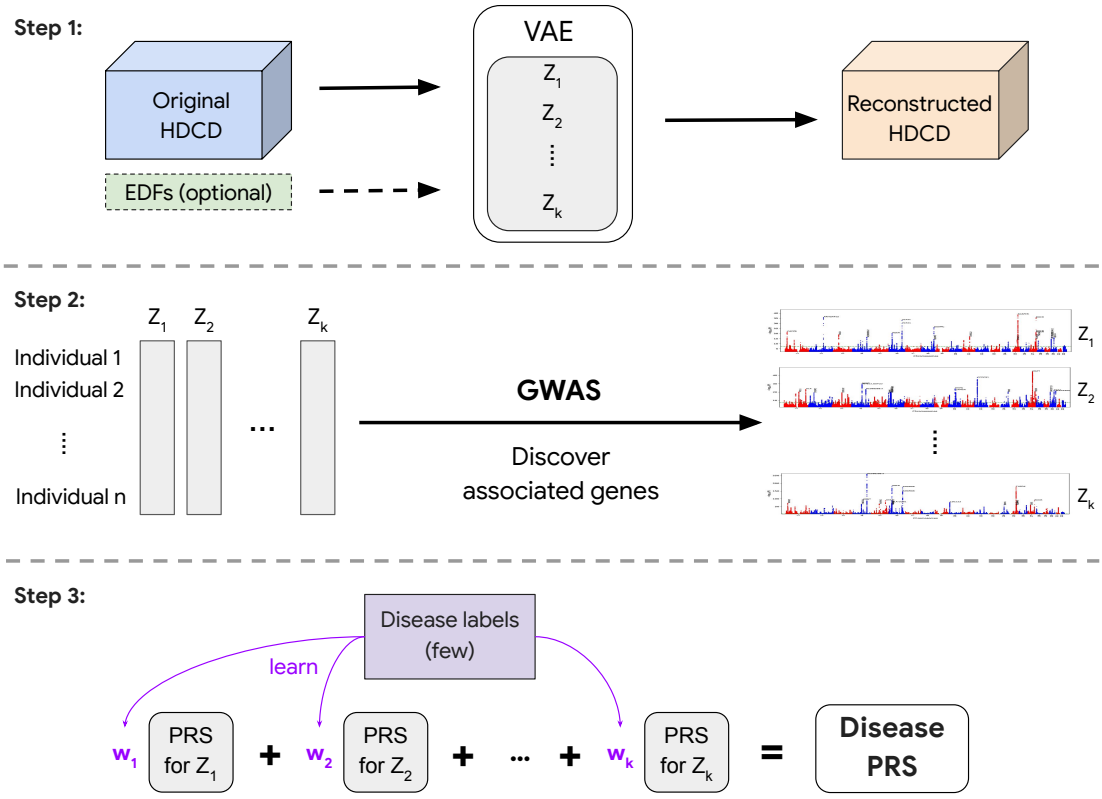

Figure 1: **Overview of representation learning for genetic discovery on low-dimensional embeddings (REGLE)**. In Step 1, we learn a low-dimensional embedding using a VAE where we optionally condition the decoder on EDFs. In Step 2, we perform GWAS on all learned variables (and EDFs if they are used). Finally in Step 3, we train a small linear model to learn weights for each latent variable PRS to obtain the final disease-specific PRS.

eters (Supplementary Figure 1, Supplementary Table 1, and Methods). Using just 5 SPINCs (the same as the number of common spirogram EDFs), we observed highly accurate reconstruction of the input spirograms (Figure 2b). SPINCs consistently outperformed an equivalent number of PCs in terms of reconstruction accuracy at small latent dimensions (Figure 2c, Supplementary Table 2, Supplementary Notes). We observed similarly accurate reconstructions using EDFs+RSPINCs and confirmed that the addition of RSPINCs improves the reconstruction quality significantly, compared to using a decoder-only model to reconstruct curves using EDFs only (Supplementary Figure 2); we used 2 RSPINCs to balance the number of additional coordinates and the reconstruction accuracy. Importantly, the learned representations are highly consistent when trained with multiple different initializations (Supplementary Figure 3, Supplementary Notes).

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

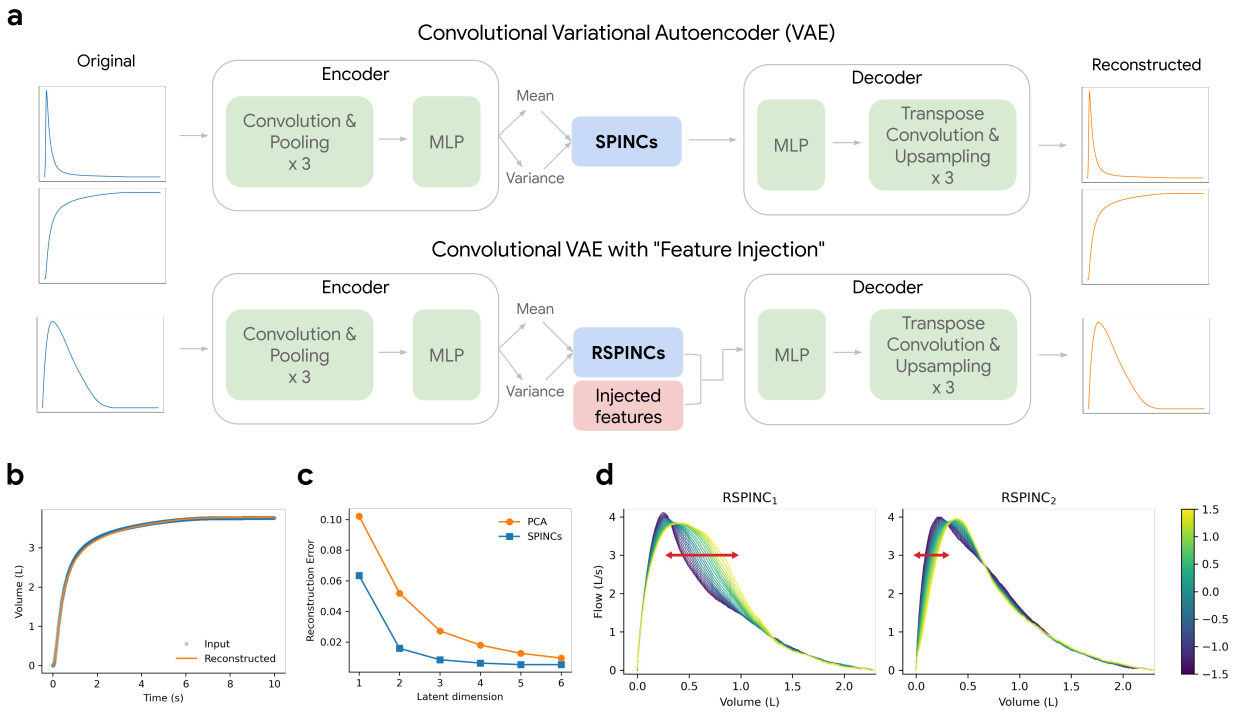

Figure 2: **Overview of REGLE on spirograms**. a) Learning spirogram encodings (SPINCs) using a convolutional variational autoencoder (VAE) and *residual* spirogram encodings (RSPINCs) using a convolutional VAE with "feature injection", using expert-defined features (EDFs) for example. b) Reconstructing a spirogram (volume-time curve) from SPINCs (dim=5). c) Reconstruction errors (mean squared error across time points) for reconstructed spirograms using the SPINCs model and PCA with a varying latent dimension. Both the SPINCs model and PCA are trained (or "fitted") on a training set and the reconstruction error was evaluated in a separate validation set. d) Spirograms created by RSPINCs (dim=2) decoder using a fixed set of injected features (i.e. EDFs) and varying one RSPINC coordinate while fixing the other one to be zero. Line color indicates the varying RSPINC coordinate value.

2.3 Overview of REGLE on PPGs

Photoplethysmograms (PPG) measure volumetric changes in blood in peripheral circulation, typically using infrared light. They are represented as pulse waveforms and used as an inexpensive and noninvasive way to measure arterial stiffness. PPGs can also be used to classify elevated blood pressure or cardiovascular disease (CVD), and to interrogate circulatory function [34, 35, 36]. To understand the genetic architecture of human circulatory function from these highly available data, we applied REGLE to obtain low-dimensional representations of PPG curves computed from a median single heartbeat (median is taken over multiple heartbeats as computed by the pulse waveform

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

device), which we call photoplethysmogram encodings (PLENCs) (Figure 3). To construct PLENCs, we trained a convolutional VAE [29] to reconstruct PPGs (Figure 3a and Methods). We generated PLENCs for all individuals (n=170,714) in UK Biobank [33] using their first-visit PPG, excluding individuals whose PPG failed our QC measures (Methods). We used 80% of the European-ancestry individuals (n=136,239) to train the PLENCs models and 20% (n=34,475) to evaluate the reconstruction performance and choose hyperparameters (Supplementary Figure 4, Supplementary Table 1, and Methods). With just 5 PLENCs (the same as the number of PPG EDFs), we observed highly accurate reconstruction of the input PPG (Figure 3b and Supplementary Table 3). PLENCs consistently outperformed an equivalent number of PCs in terms of reconstruction accuracy at small latent dimensions (Figure 3c, Supplementary Notes).

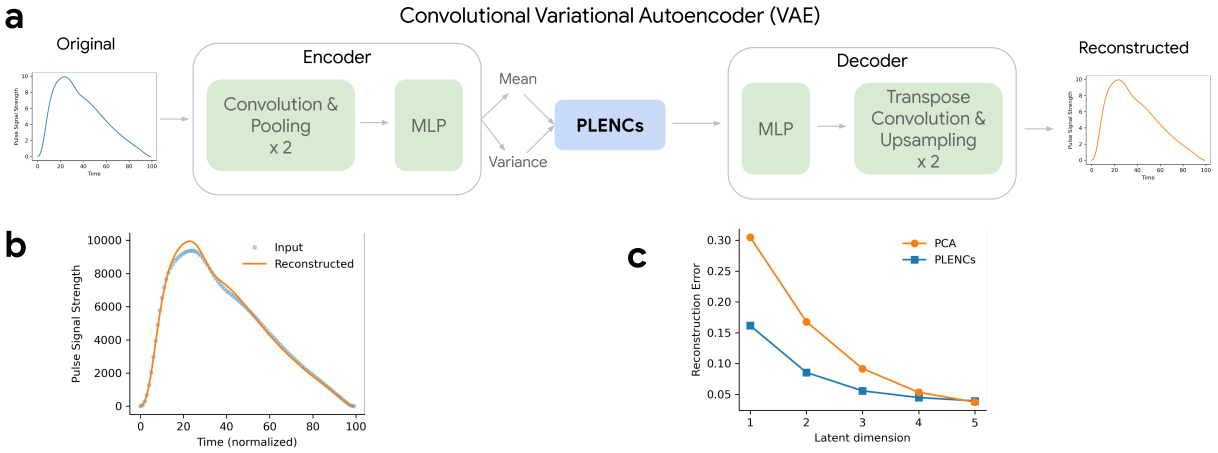

Figure 3: **Overview of REGLE on PPG**. a) Learning PPG encodings (PLENCs) using a convolutional variational autoencoder (VAE). b) Reconstructing PPG from PLENCs (dim=5). c) Reconstruction errors (mean squared error across time points) for reconstructed PPGs using the PLENCs model and PCA with a varying latent dimension. Both the PLENCs model and PCA are trained (or "fitted") on a training set and the reconstruction error was evaluated in a separate validation set.

2.4 (R)SPINCs are partially interpretable

To interpret the influence of RSPINC coordinates on spirogram shape, we fixed the values of EDFs (obtained from a randomly selected individual in the validation set) and varied one RSPINC coordinate while keeping the other one fixed at zero, and generated the corresponding flow-volume spirograms using only the decoder portion of the RSPINCs model (Figure 2d). A typical flow-volume spirogram consists of two distinct parts: the first part, a relatively brief part to reach peak flow where the flow increases monotonically as the volume increases, and the second part, the main

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

part of the spirogram where the flow decreases monotonically. In Figure 2d, we clearly observed that varying the first coordinate of RSPINCs amounts to widening or narrowing of the second part (negative slope) while keeping the first part relatively fixed. Similarly, varying the second coordinate of RSPINCs widens or narrows the first part (positive slope) while keeping the second part relatively fixed. Notably, when varying either coordinate of RSPINCs, the maximum flow value (PEF) and the final volume value (FVC) stay roughly the same, as expected since all EDFs were fixed.

2.5 (R)SPINCs and PLENCs encode information beyond EDFs

Some coordinates of SPINCs are highly correlated with known EDFs (Supplementary Figure 5). For example, the 3rd coordinate of SPINCs is 96% correlated with FVC and 94% correlated with FEV_1, while the 2nd coordinate is 73% correlated with FEV_1/FVC (after flipping the signs) (Supplementary Figure 5). Both RSPINCs coordinates have low correlation ($|R| < 0.3$) with all EDFs, which is expected since they were encouraged to learn only residual signals not captured by the EDFs (Supplementary Figure 5). (R)SPINCs are also correlated with other predictors of lung function ("covariates"), such as age, sex, height, body mass index, and smoking status (Supplementary Figure 5). To investigate if (R)SPINCs include information beyond EDFs and covariates, we residualized both the EDFs and the covariates from (R)SPINCs and computed correlation with tabular UK Biobank features. Multiple groups of fields strongly and significantly correlated with the (R)SPINCS even after residualizing the EDFs and the covariates, including asthma, breathing issues, cognitive function, and allergies (Supplementary Tables 4 to 6).

We observed similar qualitative results for PLENCs, which also show strong correlation with known PPG EDFs (Supplementary Figure 6). For example, the 3rd coordinate of PLENCs is 75% correlated with the position of the shoulder and 74% correlated with the position of the peak, while the first coordinate is 58% correlated with the pulse rate (after sign flip). PLENCs also show correlation with other predictors of circulatory function (Supplementary Figure 6). After residualizing EDFs and covariates from PLENCs, we still observed multiple groups of fields strongly and significantly correlated with the residualized PLENCs coordinates, including pulse rate, systolic and diastolic blood pressure, and position of shoulder (Supplementary Tables 7 and 8).

Finally, using the Cox proportional hazards model, we observed that both (R)SPINCs and PLENCs are associated with overall survival. For example, the third coordinate of SPINCs had a hazard ratio of 0.68 (95% CI, 0.65 to 0.71; P=1.6E-83) implying the hazard of death decreased by

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

32% per one standard deviation increase in the coordinate (Supplementary Notes, Supplementary Figures 7 to 9, Supplementary Table 9, Methods).

2.6 REGLE detects novel loci for lung and circulatory functions

We generated SPINCs (dim=5), RSPINCs (dim=2, in addition to 5 EDFs) for all individuals with valid first-visit spirograms in UK Biobank (Supplementary Figures 1, 10 and 11; Methods) and PLENCs (dim=5) for all individuals with valid first-visit PPGs in UK Biobank (Supplementary Figure 4; Methods). We then performed GWAS on all European-ancestry individuals across all encoding coordinates and 5 EDFs using BOLT-LMM [37, 38], adjusting for covariates (Supplementary Notes; Methods; Supplementary Figures 12 to 16, Supplementary Figures 17 and 18, and Supplementary Figures 19 to 23). S-LDSC analysis demonstrates that each of the SPINCs, RSPINCs, and PLENCs have significant SNP heritability. For example, the SNP heritability of SPINCs ranged from 4.3% for SPINC_5 to 24.8% for SPINC_3. Importantly, the RSPINCs, which are relatively uncorrelated with the EDFs, were also significantly heritable (4.5-16.2%), indicating the existence of additional heritable signal within the spirograms (Supplementary Table 10).

GWAS on 5 SPINCs detected 575 independent genome-wide significant (GWS) loci ($R^2 \leq 0.1$ and $P \leq 5 \times 10^{-8}$) after merging hits within 250kb together (Table 1, Methods). To compare our results to known lung function loci from previous literature, we combined the largest published GWAS on lung function (using FEV_1, FVC, PEF, and FEV_1/FVC) from Shrine et al. [25] (580,869 individuals, compared to our 324,702 individuals in UK Biobank) with all lung function-related loci in the NHGRI-EBI GWAS Catalog [39] (Methods). This resulted in 1104 independent loci after merging loci by distance (250kb), hereafter referred to as "previously known loci". Most GWS loci from SPINCs and EDFs+RSPINCs recover previously known loci (89% for SPINCs, 90% for EDFs+RSPINCs). Out of 575 GWS SPINCs loci, 65 (11%) were not previously known, compared to 32 from EDFs and 15 from PCA. Of 659 EDFs+RSPINCs GWS loci, 63 (10%) were not previously known (Table 1). Functional enrichment analysis with GARFIELD [40] shows that these loci are enriched for lung tissue DNase I hypersensitive sites (Supplementary Figures 24 to 28; Supplementary Figures 29 and 30; Supplementary Notes) and the EDFs+RSPINCs loci show stronger ontology term enrichments than EDFs loci alone (Supplementary Figure 31) using GREAT [41]. Notably, we found a strong enrichment for RSPINC_2 in blood (Supplementary Figure 30). Finally, we performed multiple analyses to ensure these novel loci were not detected by EDFs or previous works (Supplementary Notes; Supplementary Tables 11 and 12).

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

GWAS on 5 PLENCs detected 90 independent GWS loci (Table 1, Methods). To compare our results to known cardiovascular function loci from previous literature, we utilize all cardiovascular function-related loci in the NHGRI-EBI GWAS Catalog [39] (Methods; 520 independent known loci after merging by distance). For cardiovascular function and PPG, EDFs were the absence of notch, the position of notch, the position of peak, the position of shoulder, and the peak-to-peak time. Out of 90 GWS PLENCs loci, 50 (56%) were not previously known (Table 1 and Supplementary Table 13). Functional enrichment analysis with GARFIELD [40] shows that PLENCs GWS loci are enriched for fetal heart, heart, and blood vessel tissue DNase I hypersensitive sites (Supplementary Figures 32 to 36; Supplementary Notes).

System	Method (# traits)	Sample size	Total	Known (%)	Novel (%)
Lung	GWAS Catalog + Shrine 2023	-	1104	–	–
	Shrine 2023	581K	754	–	–
	Spirogram EDFs (5)	325K	613	581 (95%)	32 (5%)
	Spirogram PCA (5)	325K	412	397 (96%)	15 (4%)
	SPINCs (5)	325K	575	510 (89%)	65 (11%)
	EDFs+RSPINCs (7)	325K	659	596 (90%)	63 (10%)
Cardio-vascular	GWAS Catalog	-	520	-	-
	PPG EDFs (5)	141K	62	24 (39%)	38 (61%)
	PPG PCA (5)	141K	43	20 (47%)	23 (53%)
	PLENCs (5)	141K	90	40 (44%)	50 (56%)

Table 1: **Comparison of GWAS significant loci.** For lung function and spirograms, expert-defined features (EDFs) are FEV_1, FVC, FEV_1/FVC, PEF, and $FEF_{25\text{-}75\%}$, and "known" and "novel" is in reference to lung function loci in GWAS Catalog and Shrine et al. [25]. For cardiovascular function and PPG, EDFs are the absence of notch, the position of notch, the position of peak, the position of shoulder, and the peak-to-peak time, and "known" and "novel" is in reference to cardiovascular disease loci in GWAS Catalog.

2.7 (R)SPINCs improves asthma and COPD PRS over EDFs in UK Biobank

We computed PRSs using BOLT-LMM [37, 38] effect sizes for 5 SPINCs and 2 RSPINCs coordinates, in addition to 5 spirogram EDFs. We treat these sets of PRSs as intermediate genetic scores for lung function. Given a specific trait, a set of such intermediate PRSs, and a (small) set of individuals for whom the trait status is available, one can combine the intermediate PRSs into a single trait-specific PRS via a weighted linear sum of the intermediate PRSs. The weights of the linear score can be learned in any dataset composed of individuals with available genetics and trait labels (Figure 1). We created disease-specific PRSs for asthma and COPD from three sets of

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

intermediate PRSs: 1) 5 EDFs, 2) 5 SPINCs, and 3) 5 EDFs plus 2 RSPINCs. We learned the disease-specific PRS weights within the modeling set (n=324,958) of European ancestry individuals in UK Biobank using medical-record-based asthma and COPD statuses. To evaluate the performance of each disease-specific PRS, we computed the accuracy of the PRS in a completely separate set of individuals of European ancestry (n=110,722) not previously used for model training or GWAS.

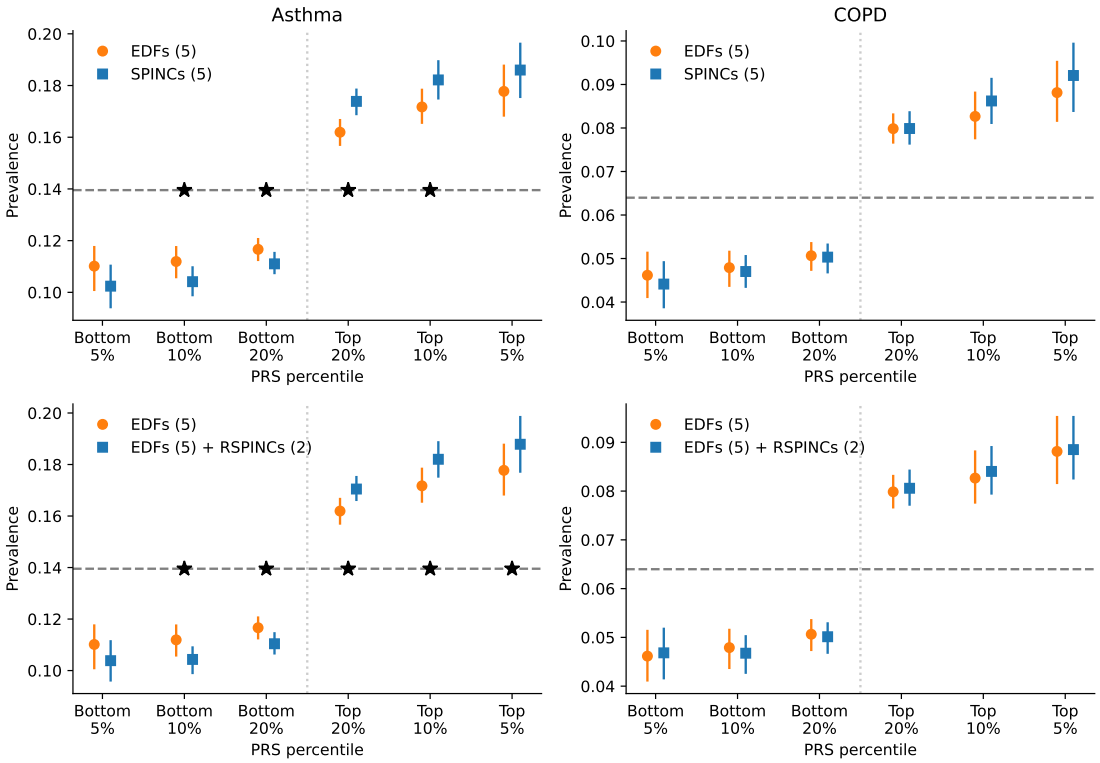

Figure 4: **PRS using SPINCs and RSPINCs in UK Biobank.** Combined PRS for medical-record-based asthma and COPD using three sets of intermediate PRS: 5 EDFs, 5 SPINCs, and 5 EDFs + 2 RSPINCs. Each set of PRS is combined by a linear model trained using the target phenotype labels and the prevalence of the phenotypes in the top and bottom 5%, 10%, and 20% PRS individuals is evaluated in a separate evaluation set. Vertical line segments indicate 95% confidence interval generated by bootstrapping (300 repetitions). The horizontal dashed line shows the total prevalence. Star (*) sign indicates a statistically significant difference between the two methods using *paired* bootstrapping (300 repetitions) with 95% confidence. Lower is better for the bottom percentiles; higher is better for the top percentiles.

We observed that the top-decile high-risk individuals based on the SPINCs asthma PRS have significantly higher asthma prevalence than individuals in the top-decile of the EDF PRS (Figure 4, Supplementary Table 14). The SPINCs asthma PRS more effectively stratifies the risk groups than EDFs on both ends of the risk spectrum. In addition, we observed statistically significant improvements in AUC-ROC, AUC-PR, and Pearson correlation by using the SPINCs asthma PRS (Supple-

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

mentary Table 14). We observed the same trend for COPD (Figure 4, Supplementary Table 15). Furthermore, we observed that the EDFs+RSPINCs PRS outperforms the EDFs PRS on all metrics for both asthma and COPD; The differences in AUC-ROC, AUC-PR, and Pearson correlation are statistically significant for both diseases, and the difference in the top decile prevalence is significant for asthma (Figure 4, Supplementary Tables 14 and 15). We observed that the SPINCs COPD PRS outperforms the FEV_1/FVC PRS (Supplementary Table 15) for predicting medical-record-based COPD, despite FEV_1/FVC having previously been shown to be one of the best phenotypes for generating a COPD PRS, even outperforming the binary COPD PRS [6]. These results provide further evidence that SPINCs capture more genetic determinants of lung function related to asthma and COPD than the same number of EDFs, and that RSPINCs capture additional genetic factors not captured by the EDFs.

Finally, since each disease-specific PRS is constructed from intermediate (R)SPINCs coordinate PRSs using just 5–7 learned weights, we explored whether these weights could be learned from a subset of the training data. For both asthma and COPD, (R)SPINCs-based PRS fit with as few as 100 disease cases performed indistinguishably from those trained on the full training data of 45,609 and 12,468 cases, respectively (Figure 1 "Step 3", Supplementary Figure 37). We also evaluated PRSs generated by GWAS with a cohort-level phenotype adjustment using inverse-normal transformation. We observed fewer significant differences with this adjustment, though SPINCs and EDFs+RSPINCs still maintain some statistically significant differences for asthma PRS (Supplementary Figure 38).

2.8 (R)SPINCs PRS transferred to independent datasets and ancestry groups

The performance of PRS computed in one dataset can degrade significantly when transferring the variant weights directly to another dataset or a different ancestry group [42]. In addition, the quality of COPD and asthma labels in different datasets can vary widely (e.g. using physician diagnoses, self report, or medical records). To test the generalizability of our (R)SPINCs PRSs to individuals outside of UK Biobank and non-European ancestry individuals, we transferred the variant weights from the asthma and COPD PRSs in UK Biobank directly to the Genetic Epidemiology of COPD (COPDGene) dataset [43], the eMERGE III dataset, the EPIC-Norfolk dataset [44], and Indiana Biobank dataset [45] (Supplementary Table 16).

For the COPDGene dataset, we computed PRS of all individuals using the same variant effect sizes obtained by BOLT-LMM in UK Biobank and the same linear weights to combine the EDFs,

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

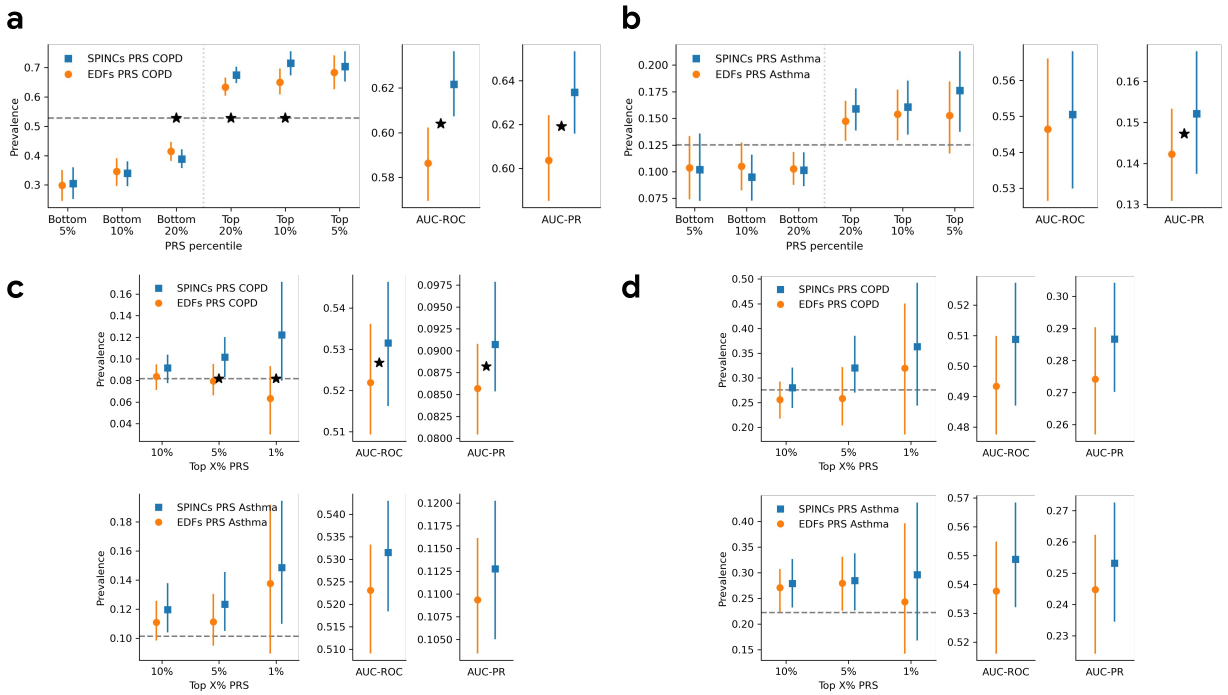

Figure 5: **SPINCs PRS transferred to multiple independent datasets**. SPINCs PRSs for COPD and asthma generated on UK Biobank are transferred to four independent datasets: COPDGene, eMERGE III, EPIC Norfolk, and Indiana Biobank. a) PRS evaluation in COPDGene dataset on COPD. b) PRS evaluation in eMERGE III dataset on asthma. c) PRS evaluation in EPIC Norfolk study on COPD and asthma. d) PRS evaluation in Indiana Biobank on COPD and asthma. In all figures, solid vertical intervals represent 95% confidence interval generated by bootstrapping (300 repetitions). The horizontal dashed lines show the total prevalence in the evaluation set. Star symbols indicate a statistically significant difference between the two methods using paired bootstrapping (300 repetitions) with 95% confidence.

EDFs+RSPINCs, and SPINCs PRS as before after matching variants. We used the "race" field in COPDGene as a proxy for genetic ancestry and computed PRS performance in the two available subsets separately: "Non-Hispanic White" and "African American". We observed that for COPD, the SPINCs COPD PRS outperforms the EDFs COPD PRS for both subset of individuals in COPDGene for all four evaluation metrics (AUC-ROC, AUC-PR, top-decile prevalence, and Pearson correlation). In the "Non-Hispanic White" subset (n=6,576), which matches the UK Biobank ancestry group on which the PRS was trained, all four metrics are statistically significant (paired bootstrapping; Figure 5a, Supplementary Table 17). In the "African American" subset (n=3,140), differences were statistically significant for AUC-ROC and Pearson correlation (Supplementary Table 17). The EDFs+RSPINCs COPD PRS significantly outperformed the EDFs COPD PRS in "Non-Hispanic White" in AUC-ROC and Pearson correlation, but did not in the "African American" subset (Supple-

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

mentary Table 17).

We also transferred the UK Biobank PRSs to eMERGE III ("White" subset, n=8,288), EPIC Norfolk (self-reported "White", n=21,010), and the Indiana Biobank (mostly European, see Methods, n=5,254), to evaluate asthma, asthma and COPD, and asthma and COPD, respectively. We observed consistent improvement from using SPINCs PRSs over EDFs PRSs for both COPD and asthma phenotypes for top-percentile prevalences, AUC-ROC, and AUC-PR. The improvement was statistically significant for AUC-PR and the top 1% and 5% prevalence in eMERGE III, and for AUC-ROC and AUC-PR in EPIC Norfolk (Figure 5b-d).

2.9 PLENCs improves hypertension and systolic blood pressure PRS over EDFs

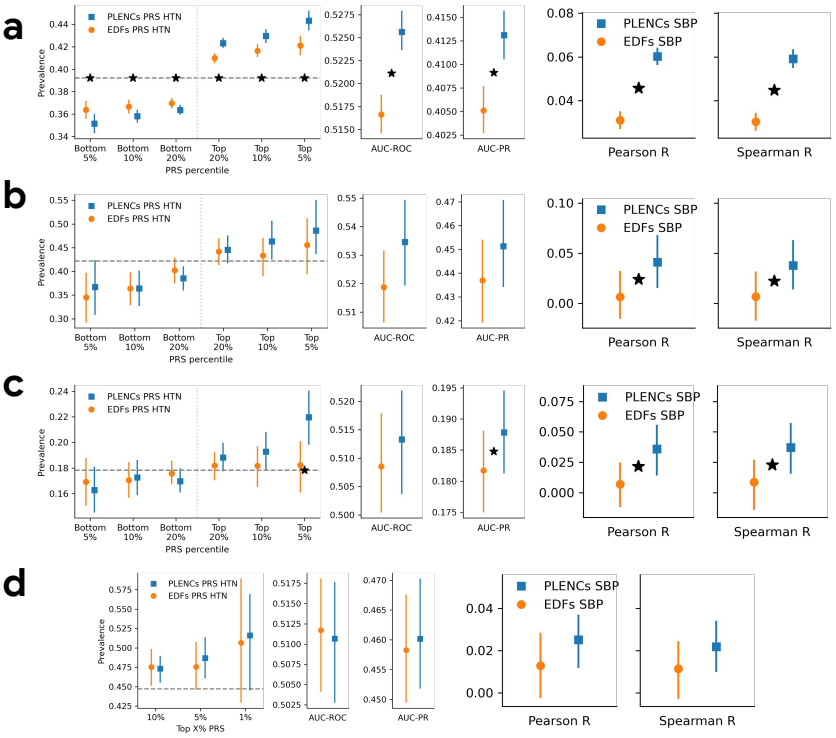

Figure 6: **PLENCs PRS generated in UK Biobank evaluated in multiple independent datasets**. PLENCs PRS for hypertension (HTN) and systolic blood pressure (SBP) generated on UK Biobank are evaluated in four independent datasets: UK Biobank, COPDGene, eMERGE III, and EPIC Norfolk. a) PRS evaluation in UK Biobank, evaluated in a separate test set not used for GWAS. b) PRS evaluation in COPDGene dataset. c) PRS evaluation in eMERGE III. d) PRS evaluation in EPIC Norfolk. In all figures, solid vertical intervals represent 95% confidence interval generated by bootstrapping (300 repetitions). The horizontal dashed lines show the total prevalence in the evaluation set. Star symbols indicate a statistically significant difference between the two methods using paired bootstrapping (300 repetitions) with 95% confidence.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

We computed PRSs using BOLT-LMM for 5 PLENC coordinates in addition to 5 PPG EDFs, and utilized these sets of PRSs as intermediate genetic scores for cardiovascular function. Since PPG is known to carry signals about blood pressure [46], we created trait-specific PRSs for hypertension (HTN; a binary trait) and systolic blood pressure (SBP; a continuous trait) using the REGLE framework (Figure 1, Supplementary Table 18). We evaluated HTN and SBP PRS generated by PLENCs and PPG EDFs in independent datasets (COPDGene, eMERGE III, and EPIC-Norfolk) in addition to the held-out test set in UK Biobank. For HTN we measured the top and bottom PRS group prevalences, AUC-ROC, and AUC-PR, and for SBP we measured the Pearson correlation and the Spearman rank correlation. Note that we did not perform HTN and SBP PRS evaluation in Indiana Biobank due to unusually high prevalence of hypertension (more than 80%) and the usage of blood pressure medication by a majority of its population.

We observed consistent trend of improvement from using PLENCs PRSs over EDFs PRSs for both HTN and SBP, with the exception of AUC-ROC of HTN in EPIC Norfolk (Figure 6). Notably, PLENCs PRS for SBP outperformed EDFs PRS for all datasets for both correlation metrics, e.g. $2\times$ higher Pearson correlation (6% vs. 3%) in the UK Biobank test set, and the differences were statistically significant in three out of four datasets.

2.10 High association between REGLE encodings and UKB phenotypes PRSs

To assess the influence of (R)SPINCs and PLENCs on traits and health outcomes, we performed phenome-wide association studies (PheWAS). We compared pruning+thresholding PRSs of all (R)SPINCs and and PLENCs coordinates to PRSs of 7,145 phenotypes computed by the Pan-UK Biobank consortium (Methods, URLs). For SPINCs, significant associations with strong correlation magnitude are driven by the 2nd and 3rd coordinates, which as mentioned above are strongly correlated with EDFs, and include expected phenotypes such as FEV_1, FEV_1/FVC, and PEF (Supplementary Tables 19 and 20). These coordinates also show strong correlation to other diseases and traits associated with alterations in lung function, identifying relationships with systemic lupus erythematosus [47, 48], thyroid dysfunction [49], and gluten-free diet [50]. For RSPINCs, significant correlations with strong magnitude are nearly all driven by $RSPINC_1$ and also include the same diseases (Supplementary Tables 21 and 22). For PLENCs, we observed all coordinates show significant correlation with different traits including blood traits (red blood cell, eosinophil count, age high blood pressure diagnosed, and haemoglobin concentration), PPG traits (pulse rate and pulse wave reflection index), ECG traits (QRS duration and P duration), blood pressure, and

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

vascular/heart problems (Supplementary Tables 23 and 24). The strongest correlation was obtained from 1st PLENCs coordinate with pulse rate (R=-0.67; P $\leq$ 1.00E-300) and ECG heart rate (R=-0.52; P $\leq$ 1.00E-300).

Finally, we performed latent causal variable (LCV) analysis [51] to identify potentially causal relationships between (R)SPINCs and EDFs with COPD, asthma, and other lung-related phenotypes (sarcoidosis, systemic lupus erythematosus, thyroid dysfunction, and gluten-free diet) where we observed significant correlation between their PRSs in PheWAS (Supplementary Tables 25 and 26). We observed a significantly positive genetic causality proportion (GCP) from the 2nd SPINCs coordinate to COPD, suggesting a genetic causal link from SPINCs to COPD. We also observed a significant negative GCP from the 5th SPINCs coordinate to asthma, suggesting a genetic causal link from asthma to SPINCs (Supplementary Tables 25 and 26; Supplementary Notes). Interestingly, all significant GCPs were negative from SPINCs to other lung-related phenotypes we considered, implying that the genetic causal directions are from those phenotypes to SPINCs. In the case of PLENCs, we observed that 1st PLENC coordinate has a significant negative GCP with hypertension (Supplementary Tables 25 and 26), implying a genetic causal link from hypertension to PLENCs.

3 Discussion

Large biobanks provide unique opportunities to identify the genetic factors underlying complex traits and diseases, but accurately phenotyping [52] the cohorts remains a core challenge. We proposed a general unsupervised deep representation learning method, REGLE (representation learning for genetic discovery on low-dimensional embeddings), to improve discovery of genetic components of high-dimensional clinical data (HDCD). We showcased the effectiveness of REGLE in spirograms and photoplethysmograms, where it produced latent variables ("encodings") that are both partially interpretable and effective for identifying genetic variants associated with lung and circulatory functions. Multiple lines of evidence show the relevance of the model representations for quantifying general lung and circulatory functions.

Unsupervised representation learning of HDCD for genomic discovery is attractive owing to the difficulty of acquiring (or defining) manual phenotypes at scale. Prior work has explored applying transfer learning [53] and contrastive learning [54] to retinal fundus images, or multimodal autoencoders to cardiac data modalities [55]. A key strength of our method is the use of a VAE to generate the non-linear low-dimensional representations of HDCD. First, by construction, the

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

coordinates of the latent representation are minimally correlated, which strengthens the combined power of the downstream GWASs. As a result, the PRSs of the learned encodings are also minimally correlated and contain relatively orthogonal genetic signals, which may contribute to the superior accuracy of the disease/trait specific PRS created by the REGLE framework. Second, the learned representations are stable up to changes in sign or order as we observed empirically (Section 2.4), potentially due to a grounding effect of a VAE prior in the probabilistic model. As changes in sign or order do not affect genetic discovery–though they may affect the sign of the effect size–the results of the REGLE framework are stable and replicable. Regular autoencoders without a prior do not have this stability property as they can learn any invertible linear transformation of a specific learned representation.

The architecture modification we introduce in REGLE to support expert-defined features (EDFs) enables a principled use of expert human knowledge and encourages the remaining latent coordinates to encode biological function explicitly not captured by the EDFs. This provides the opportunity to build upon and improve existing clinical practices with the extra information provided from the residual features. For example, clinical review of pulmonary function tests includes visual inspection of the curves for variation in shape. Coving of the flow volume loop is an indicator of obstruction. The ability of (R)SPINCs to identify these differences while holding EDFs constant suggest that (R)SPINCs are quantifying clinically relevant parameters that are not currently captured in a standardized or automated way. While we demonstrated the value of the method on spirograms to model lung function, it can be generalized to other HDCD modalities such as images.

The improved performance of SPINCs, EDFs+RSPINCs, and PLENCs PRS over EDFs PRS provides evidence for the existence of such residual genetic information. Importantly, the (R)SPINCs PRS for respiratory diseases (asthma and COPD) and PLENCs PRS for cardiovascular functions (hypertension and SBP) combine each latent coordinate GWAS into a PRS based solely on the effect size estimates from the GWAS on the learned coordinates, and the disease-specific PRS is simply a learned weighted sum of the handful (i.e., five or seven) coordinate PRSs. This has important implications for disease risk prediction: it demonstrates that given a dataset with widespread unlabeled quantification of lung system function (i.e. spirograms) and/or circulatory system function (i.e. photoplethysmograms), genetic predictors for specific lung or cardiovascular diseases can be accurately learned with very few disease labels (enough to learn 5-7 features). We hypothesize that unsupervised quantification of other organ systems may be similarly beneficial for improving polygenic prediction across a wealth of diseases. Finally, we note that the PRS performance re-

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

ported here represents a lower bound achievable by the method; jointly estimating disease-specific variant effect size estimates on the set of variants identified by the (R)SPINCs and PLENCs GWAS may further improve prediction performance.

There are several limitations to this work. First, while unsupervised representation learning of HDCD for biological function is likely beneficial across data modalities, the generalizability of the specific VAE method introduced here to even higher dimensional modalities like imaging and video has not been assessed. In particular, VAEs tend to produce blurry image reconstructions [56, 57]; whether and how that affects the ability of a VAE to encode representations meaningful for genomic discovery is important future work. Second, we did not directly optimize multiple GWASs for novel genomic discovery, but used a straightforward (conservative) method to define and merge independent associated loci. A possible extension would be to combine the signals from multiple (R)SPINC and PLENC coordinate GWAS [11] or apply methods that maximize heritability (e.g. [58]). Third, we did not fully explore model architecture and training strategies specifically for genomic discovery. Some ideas which may warrant further investigation include: 1) using previously proposed modifications to the VAE loss function and the training procedure to maximize the degree of disentanglement of coordinates while balancing the reconstruction error [59, 60, 61], 2) incorporating an additional loss term to explicitly discourage correlation between RSPINCs and EDFs, and 3) introducing (semi-) supervision in model training to overcome the limitations of purely unsupervised training [62]. Fourth, in the case spirograms, none of the spirograms in UK Biobank were obtained after bronchodilation, and the asthma and COPD labels defined using clinical records are known to be noisy [6]. Fifth, we generated individual-level spirogram representations from the first measurement, despite some individuals having up to three acceptable blows in UK Biobank. Integrating all acceptable blows from an individual could produce a more comprehensive representation of their lung function [63]. Sixth, model training was performed exclusively in individuals of European ancestry. While PRS evaluation was performed in multiple datasets and ancestries, the impact of ancestry-specific model training was not explored.

Despite these limitations, REGLE provides a mechanism for identifying genetic influences on organ function in the absence of labeled data, and naturally admits incorporating expert features in the model. It also provides a method to create disease/trait specific PRS with very few labels (i.e. in the order of hundreds). As biobanks with rich imaging, activity monitoring, medical records, and paired genetic data continue to grow, we anticipate that this or similar methods will be increasingly used to further elucidate the genetic underpinnings of human traits and diseases.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

4 Methods

4.1 UK Biobank data preparation for spirograms

Spirograms from UK Biobank were sourced from the data field 3066, which contains the volume in milliliters of exhalation at 10 millisecond intervals (volume-time curve) and were preprocessed closely following the procedures in Cosentino et al. [6]. To generate flow-time curves, we approximated the first derivative of volume with respect to time by taking a finite difference in the volume-time curves. We normalized the volume-time and flow-time curves to 1000 time points, by either truncating longer curves or by right-padding shorter curves with zero (for flow-time curves) or the final value (for volume-time curves), and removed FEV_1, FVC, PEF values in the extreme tail (top/bottom 0.5%) of the observed values and all blows that fail the acceptability provided by UK Biobank data field 3061. We used the first acceptable blow of an individual when there is more than one. In addition, we dropped all flow curves whose values don't fall in [-10, 20], all volume curves whose values are not in [-5, 10], and all flow curves in which the proportion of nonzero values are less than 20%. Finally, we generated flow-volume curves from volume-time and flow-time curves by interpolating 1000 evenly spaced volume values between 0 and 6.58 liters (the maximum observed volume in the dataset).

We then subdivide all European ancestry individuals processed this way into a 80% training set and a 20% validation set. After additionally removing related individuals, there are 259,692 individuals in the training set and 65,266 individuals in the validation set (Supplementary Figure 1).

Asthma and COPD status were determined by medical records using self report, ICD9, and ICD10 codes as defined in Cosentino et al. [6].

4.2 UK Biobank data preparation for photoplethysmograms

Photoplethysmograms (PPG) from UK Biobank were sourced from the data field 4205, which contains the arterial stiffness pressure curve. Each waveform is actually a single pulse with 100 points. Then we computed the minimum, maximum, mean, and median distribution value of PPG. We keep the PPG when all four statistics fall in 0.1 and 99.9 percentiles of the related statistics values of all PPGs. We then subdivide all European ancestry individuals processed this way into a 80% training set and a 20% validation set. After additionally removing related individuals, there are 112,730 individuals in the training set and 28,545 individuals in the validation set (Supplementary Figure 4).

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Hypertension status was determined by medical records using self report, ICD9 codes (401.* and 405.*), and ICD10 codes (I10 and I15.*). Systolic blood pressure (SBP) was determined by automated reading and Data-Field 4080 was used in UK Biobank.

4.3 Convolutional VAE model architecture and training

To generate SPINCs, we encode the flow-time and volume-time curves. In our VAE, we use one-dimensional convolutional layers to utilize the temporal context of this time series, encoding the two curves in two channels. In the encoder, we first apply three 1D convolutional layers, each followed by max pooling, and use three fully-connected layers to generate the mean and variance of the bottleneck layer. We use 5 latent dimensions, identical to the number of EDFs, and each latent coordinate is sampled from the Gaussian distribution with the learned means and variances. The decoder architecture is a mirror image of the encoder. We start with three fully-connected layers followed by transpose convolutions layers, each prepended by an upsampling layer. See "SPINCs model architecture" in Supplementary Notes for full details.

For RSPINCs, we encode the flow-volume curve alone, and we apply the same sequences of convolutional and fully-connected layers as we did for SPINCs, while using only 2 latent dimensions in this case. Importantly, we use a novel VAE architecture to concatenate the 5 EDFs directly to the sampled output of the bottleneck layer (the layer right before the decoder) to learn only the residual signals not represented by the EDFs (Figure 2a). As a result, the encoder output dimension is 2 while the decoder input has dimension $5+2=7$. See "RSPINCs model architecture" in Supplementary Notes for full details.

For PLENCs, we encode the PPG curves. In our VAE, we use one-dimensional convolutional layers to utilize the temporal context of this time series. In the encoder, similar to SPINCs, we first apply three 1D convolutional layers, each followed by max pooling, and use three fully-connected layers to generate the mean and variance of the bottleneck layer. We use 5 latent dimensions, identical to the number of EDFs, and each latent coordinate is sampled from the Gaussian distribution with the learned means and variances. The decoder architecture is a mirror image of the encoder. We start with three fully-connected layers followed by transpose convolutions layers, each prepended by an upsampling layer. See "PLENCs model architecture" in Supplementary Notes for full details.

All models are trained using the standard VAE loss function consisting of the reconstruction loss and the (rescaled) Kullback–Leibler (KL) divergence loss. For RSPINCs the KL divergence loss

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

is only applied to the learned encodings, not to the injected EDFs. For optimization, the Adam optimizer [64] is used with varying learning rates and batch sizes. The final learning rate and batch size values ("hyperparameters") for (R)SPINCs and PLENCs were chosen to minimize the VAE loss in the validation set (Supplementary Notes, Supplementary Table 1).

After training SPINCs, RSPINCs, and PLENCs models, we use the encoders of the trained models to generate the encodings for each individuals, using the mean value of the learned Gaussian distribution of the encodings.

All models were implemented in TensorFlow V2 [65] and XManager (URLs) was used to manage multiple machine learning experiments.

4.4 Phenotypic Correlation Analysis

To residualize EDFs and/or covariates from (R)SPINCs and PLENCs, we used ordinary least squares linear regression. To compute the correlation of the EDFs-and-covariates-residualized (R)SPINCs and PLENCs with the tabular fields in UK Biobank, we first preprocessed the tabular fields to remove special codes, normalize, impute and aggregate the values, and finally transformed the categorical fields into one-hot encodings. For each individual correlation analysis between a feature and one of the (R)SPINCs and PLENCs, we computed the Pearson correlation and the P-value with two-sided alternative hypothesis.

4.5 Survival analysis

We performed analysis of overall survival for European individuals in the spirometry (n=65,266) and PPG (n=28,545) validation sets using the time from first assessment (field 53) to death from any cause (field 40000). Subjects who were not known to have died were right-censored at the date of UKB data ingestion (Dec 18, 2020). We quantified the association between overall survival and each SPINC, RSPINC, PLENC, and EDF per standard deviation using the hazard ratio, which was estimated from a Cox proportional hazards model adjusting for age and sex. The proportional hazards assumption, with respect to each SPINC, RSPINC, PLENC, and EDF, was assessed using the Schoenfeld residual test. After stratifying patients into quartiles using each SPINC, RSPINC, PLENC, or EDF coordinate, the overall survival curves in Supplementary Figures 7 to 9 were constructed using the standard Kaplan-Meier estimator with bootstrapped 95% confidence intervals.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

4.6 Genome-wide association studies and polygenic risk scores

GWAS on all spirograms EDFs, SPINCs, and RSPINCs were performed using BOLT-LMM v2.3.6 [37, 38], adjusting for age, sex, age^2, age $\times$ sex, height, height^2, body mass index, smoking status, the number of packs of cigarettes smoked per year, the type of genotyping array, and the top 15 genetic principal components as covariates. GWAS on all PPG EDFs and PLENCs were performed using BOLT-LMM v2.3.6 [37, 38], adjusting for age, sex, age^2, age $\times$ sex, height, height^2, body mass index, the type of genotyping array, and the top 15 genetic principal components as covariates.

All GWAS were restricted to European ancestry individuals to minimize confounding. For quality control we kept variants with minor allele frequency ≥ 0.001, imputation INFO score ≥ 0.8, missing call fraction ≤ 0.05, and Hardy-Weinberg equilibrium P-value $\geq 10^{-10}$, among all genotyped and imputed variants provided by UK Biobank. After GWAS, we performed S-LDSC [66] to estimate SNP-heritability and detect potential confounding. Genome-wide significant "hits" were defined as the most significant variants with $p \leq 5 \times 10^{-8}$ and independent at $R^2 < 0.1$ using the PLINK `-clump` command. A reference panel for linkage disequilibrium (LD) calculation contained 10,000 unrelated European samples from the UKB. Significant "loci" were created based on the span of reference panel SNPs in LD ($R^2 \geq 0.1$) with the hits. Loci separated by fewer than 250 kb were subsequently merged.

While performing GWAS, PRSs for all traits were computed using the `-predBetasFile` option of BOLT-LMM, which generates PRS coefficients using a Bayesian linear mixed model. While GWAS was performed on individuals with valid spirometry measurements, we evaluated the performance of the PRS in a separate set of individuals not used for GWAS.

To determine the known lung function loci from previous literature, we extracted all significant loci from Shrine et al. [25] and searched for lung function-related keywords in the NHGRI-EBI GWAS Catalog (version `v1.0.2-associations_e106_r2022-07-09`) [39]. We used the following keywords for the Catalog search (case insensitive): "asthma", "chronic obstructive pulmonary disease", "copd", "expiratory flow", "fev1", "forced expiratory", "forced vital capacity", and "lung function". To determine the known cardiovascular function loci from previous literature, we used the following keywords for the NHGRI-EBI GWAS Catalog search (case insensitive): "arrhythmia", "afib", "atrial fibrillation", "coronary artery disease", "stroke", "heart attack", "myocardial infarction", "heart failure", "mace", and "rheumatic heart disease".

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

4.7 Summary statistics conditional and joint analysis

We applied conditional and joint analysis (COJO) on set of previously known loci using GCTA (genome-wide complex trait analysis) software (version `1.93.3beta`) and we set `-cojo-cond` to the set of known loci. We provided 10,000 unrelated European samples randomly chosen from UKB as the reference samples, which is the same reference used to perform LD clumping to define hits, as part of GCTA-COJO input parameters.

4.8 PRS evaluation on respiratory diseases and cardiovascular traits on multiple datasets

COPDGene dataset: COPDGene is a study of 10,300 current and former smokers with and without COPD, self-reported non-Hispanic White and African-American, without known lung diseases other than COPD and asthma (dbGaP accession phs000179.v6.p2). Additional study details, the study protocol and details of genotyping have been previously published [43, 67], and additionally detailed at copdgene.org. We used the provided variant calls in VCF files and imputed the variants to the Haplotype Reference Consortium (HRC) reference panel using Michigan Imputation Server [68], resulting in 39,127,678 total variants. COPD cases were determined using the GOLD (Global Initiative for Chronic Obstructive Lung Disease) criteria, where GOLD stage 2 or higher was considered as cases. Among 6,576 non-Hispanic White individuals, we had access to 1,131 (17%) asthma cases and 2,781 (42%) COPD cases, and the rest were used as controls. Meanwhile, among 3,140 African-American individuals, 760 (24%) were asthma cases and 802 (26%) were COPD cases. We used the blood pressure measurements and the "high blood pressure" variable included in the dataset to define systolic blood pressure and hypertension traits.

EPIC-Norfolk dataset: The European Prospective Investigation into Cancer in Norfolk (EPIC-Norfolk) is a general population-based cohort study of men and women aged 40–79 years living in Norfolk, UK and recruited from general practices between 1993 to 1997. EPIC-Norfolk cohort participants were linked annually to nationally held hospital records and death certificates from 1999 to 2019 using UK National Health Service numbers. COPD was defined as any hospital admission or cause of death coded 490–492, 494–496 (ICD-9) or J40–J44, J47 (ICD-10). Asthma was similarly defined using codes 493 (ICD-9) or J45, J46 (ICD-10). Hypertension was defined using hospital records and death certificates for ICD codes 401.*, 405.* (ICD-9) and I10, I15.* (ICD-10). The systolic blood pressure was determined using the continuous systolic blood pressure from EPIC-Norfolk

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

health examination at baseline, which is the time-point with the highest number of individuals. In a small set of participants who do not have a baseline blood pressure measurement, we used blood pressure measured at the earliest subsequent health examination. Blood pressure was measured at two time points during the examination, with the mean used for analysis.

eMERGE III dataset: We utilize five consent groups that does not require IRB approval: General Research Use (GRU), Health/Medical/Biomedical-Genetic Studies, (HBM-GSO), Health/Medical/Biomedical (HMB), Health/Medical/Biomedical (MDS) HMB-MDS, and Health/Medical/Biomedical (PUB, GSO) (HMB-PUB-GSO) (dbGaP accession phs001584.v2.p2). In eMERGE III, we have access to 1,038 asthma cases and 7,250 controls for European ancestry while in the case of African ancestry we have access to 649 asthma cases and 1,367 controls. We used the 39,131,578 variants that are imputed to the HRC reference provided by dbGaP [69]. Asthma and systolic blood pressure traits were defined using the corresponding variables in the dataset. Hypertension was defined using two variables, "CASE_CONTROL_CKD_T2D_HTN" and "CASE_CONTROL_CKD_T2D", where the individuals in the former group but not in the latter group were defined as the hypertension cases. Note that this limited our analysis to hypertensive individuals without chronic kidney disease or type 2 diabetes.

Indiana Biobank dataset: The Indiana Biobank is a state-wide collaboration that provides centralized processing and storage of specimens that are linked to participants' electronic medical information via Regenstrief Institute at Indiana University. COPD was diagnosed by using ICD9: 491, 492, and 496, and ICD10: J41, J42, J43, and J44. Asthma was diagnosed by using ICD9: 493, and ICD10: J45 and J46. Cases were defined as having at least one in-patient diagnosis or two out-patient diagnoses. Those participants not having any diagnosis were defined as controls. Thus, we have 1,445 COPD cases and 3,808 controls while we have 1,171 asthma cases and 4,083 controls. Among 5,253 individuals for COPD evaluation, 3797 were of European ancestry, 1,371 of African ancestry, and 85 Hispanic ancestry. Among 5,254 individuals for asthma evaluation, 3805 of European ancestry, 1,362 of African ancestry, and 87 of Hispanic ancestry. Indiana Biobank samples used in this study were genotyped using Illumina Infinium Global Screening Array (GSA, Illumina, San Diego, CA) by Regeneron (Tarrytown, NY). SNPs with missing rate $> 5\%$, MAF $\leq 1\%$, HWE P-value $< 1 \times 10^{-10}$ among cases and $< 1 \times 10^{-6}$ in controls were excluded as reported previously [45]. Genotyping data were imputed to 1000 Genomes using the Michigan Imputation Server [68]. Imputed variants with $R^2 < 0.30$ and MAF$< 1\%$ were excluded. PLINK [70, 71] was used to calculate PRS by using imputation dosages.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

4.9 Functional significance of discovered loci

We ran GREAT v4.0.4 [41] on the human GRCh37 assembly to perform functional enrichment analysis of SPINCs, RSPINCs, PLENCs, and EDFs loci. We used the default "basal+extension" region-gene association rule with 5 kb upstream, 1 kb downstream, 1000 kb extension, and curated regulatory domains included. Furthermore, we ran GARFIELD [40] with default parameters to perform tissue-specific analysis where we utilized 424 DNase I hypersensitive site hotspot annotations provided by the GARFIELD authors [40].

4.10 Latent causal variable analysis

We applied LCV [51] (URLs) on genome-wide summary statistics for each pair of phenotypes. To create the right input for LCV, we used the munge script provided by S-LDSC software (URLs) to restrict the variants to HapMap3 SNPs with MAF > 0.05 and outside the MHC region. We utilized the baseline LD scores for HapMap3 SNPs. A two-sided test was used for the estimated GCP and to compute the significant level.

4.11 Genetic phenome-wide association study

To compute PheWAS, we downloaded GWAS summary statistics for 7,221 phenotypes from the Pan-UKB consortium 20200615 release (URLs). After restricting to phenotypes that contained European statistics and did not persistently fail in LD clumping, we were left with 7,145 pruning+thresholding PRSs generated by PLINK (URLs) using the `-clump` command with an index variant significance threshold of 5×10^{-8} and LD threshold of 0.1, with LD computed from a random subset of 10,000 European individuals in UK Biobank.

SPINCs, RSPINCs, PLENCs pruning+thresholding PRSs were computed analogously to the Pan-UKB PRSs. We computed the Pearson correlations between the PRSs derived from latent dimensions with the PRSs derived from Pan-UKB phenotypes and the P-values with two-sided alternative hypothesis.

URLs

Baseline and BaselineLD annotations: `https://data.broadinstitute.org/alkesgroup/ldscore`

BOLT-LMM software: `https://data.broadinstitute.org/alkesgroup/bolt-lmm`

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

COPDGene study: https://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000179.v6.p2

GCTA software: https://github.com/jianyangqt/gcta

GREAT software: http://great.stanford.edu

GWAS Catalog: https://www.ebi.ac.uk/gwas/

Indiana Biobank study: https://indianabiobank.org/

Pan-UK Biobank GWAS: https://pan.ukbb.broadinstitute.org

PLINK software: https://www.cog-genomics.org/plink1.9

TensorFlow: https://www.tensorflow.org

UCSC LiftOver: https://genome.ucsc.edu/cgi-bin/hgLiftOver

UK Biobank study: https://www.ukbiobank.ac.uk

XManager: https://github.com/deepmind/xmanager

Michigan Imputation Server https://imputationserver.sph.umich.edu/index.html#!pages/home

Data & code availability

Open-source code and trained model weights are available at https://github.com/Google-Health/genomics-research under the regle directory. SPINCs, RSPINCs, and PLENCs values of UK Biobank individuals will be returned to UK Biobank and will be made available by UK Biobank. GWAS summary statistics for SPINCs, RSPINCs, PLENCs, and EDFs are freely available on Google Cloud Storage at https://console.cloud.google.com/storage/browser/brain-genomics-public/research/regle.

Acknowledgments

We thank all participants, dataset creators, and maintainers of UK Biobank, COPDGene, eMERGE III, EPIC Norfolk, and Indiana Biobank. Full acknowledgement for each dataset can be found in "Dataset acknowledgment" in Supplementary Notes. We also thank Nick Furlotte for helpful discussions.

M.H.C. was supported by NHLBI R01HL153248, R01HL149861, and R01HL147148. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. The funding body has no role in the design of the study and collection, analysis, and interpretation of data and in writing the manuscript.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

References

[1] Lloyd T Elliott, Kevin Sharp, Fidel Alfaro-Almagro, Sinan Shi, Karla L Miller, Gwenaëlle Douaud, Jonathan Marchini, and Stephen M Smith. Genome-wide association studies of brain imaging phenotypes in UK biobank. *Nature*, 562(7726):210–216, October 2018.

[2] Wenjia Bai, Hideaki Suzuki, Jian Huang, Catherine Francis, Shuo Wang, Giacomo Tarroni, Florian Guitton, Nay Aung, Kenneth Fung, Steffen E Petersen, Stefan K Piechnik, Stefan Neubauer, Evangelos Evangelou, Abbas Dehghan, Declan P O'Regan, Martin R Wilkins, Yike Guo, Paul M Matthews, and Daniel Rueckert. A population-based phenome-wide association study of cardiac and aortic structure and function. *Nat. Med.*, 26(10):1654–1662, October 2020.

[3] Babak Alipanahi, Farhad Hormozdiari, Babak Behsaz, Justin Cosentino, Zachary R McCaw, Emanuel Schorsch, D Sculley, Elizabeth H Dorfman, Paul J Foster, Lily H Peng, Sonia Phene, Naama Hammel, Andrew Carroll, Anthony P Khawaja, and Cory Y McLean. Large-scale machine-learning-based phenotyping significantly improves genomic discovery for optic nerve head morphology. *Am. J. Hum. Genet.*, 108(7):1217–1230, July 2021.

[4] Nay Aung, Jose D Vargas, Chaojie Yang, Kenneth Fung, Mihir M Sanghvi, Stefan K Piechnik, Stefan Neubauer, Ani Manichaikul, Jerome I Rotter, Kent D Taylor, Joao A C Lima, David A Bluemke, Steven M Kawut, Steffen E Petersen, and Patricia B Munroe. Genome-wide association analysis reveals insights into the genetic architecture of right ventricular structure and function. *Nat. Genet.*, pages 1–9, June 2022.

[5] James P Pirruccello, Paolo Di Achille, Victor Nauffal, Mahan Nekoui, Samuel F Friedman, Marcus D R Klarqvist, Mark D Chaffin, Lu-Chen Weng, Jonathan W Cunningham, Shaan Khurshid, Carolina Roselli, Honghuang Lin, Satoshi Koyama, Kaoru Ito, Yoichiro Kamatani, Issei Komuro, Sean J Jurgens, Emelia J Benjamin, Puneet Batra, Pradeep Natarajan, Kenney Ng, Udo Hoffmann, Steven A Lubitz, Jennifer E Ho, Mark E Lindsay, Anthony A Philippakis, and Patrick T Ellinor. Genetic analysis of right heart structure and function in 40,000 people. *Nat. Genet.*, 54(6):792–803, June 2022.

[6] Justin Cosentino, Babak Behsaz, Babak Alipanahi, Zachary R. McCaw, Davin Hill, Tae-Hwi Schwantes-An, Dongbing Lai, Andrew Carroll, Brian D. Hobbs, Michael H. Cho, Cory Y.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

McLean, and Farhad Hormozdiari. Inference of chronic obstructive pulmonary disease with deep learning on raw spirograms identifies new genetic loci and improves risk models. *Nature Genetics*, 55(5):787–795, April 2023. doi: 10.1038/s41588-023-01372-4. URL https://doi.org/10.1038/s41588-023-01372-4.

[7] Niek Verweij, Jan-Walter Benjamins, Michael P Morley, Yordi J van de Vegte, Alexander Teumer, Teresa Trenkwalder, Wibke Reinhard, Thomas P Cappola, and Pim van der Harst. The genetic makeup of the electrocardiogram. *Cell Syst*, 11(3):229–238.e5, September 2020.

[8] Yoav Benjamini, Dan Drai, Greg Elmer, Neri Kafkafi, and Ilan Golani. Controlling the false discovery rate in behavior genetics research. *Behavioural Brain Research*, 125(1-2):279–284, nov 2001. doi: 10.1016/s0166-4328(01)00297-2. URL https://doi.org/10.1016%2Fs0166-4328%2801%2900297-2.

[9] David L. Streiner and Geoffrey R. Norman. Correction for multiple testing. *Chest*, 140(1):16–18, jul 2011. doi: 10.1378/chest.11-0523. URL https://doi.org/10.1378%2Fchest.11-0523.

[10] Karl Pearson. LIII. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, November 1901.

[11] Hugues Aschard, Bjarni J Vilhjálmsson, Nicolas Greliche, Pierre-Emmanuel Morange, David-Alexandre Trégouët, and Peter Kraft. Maximizing the power of principal-component analysis of correlated phenotypes in genome-wide association studies. *Am. J. Hum. Genet.*, 94(5):662–676, May 2014.

[12] Xikun Han, Kaiah Steven, Ayub Qassim, Henry N. Marshall, Cameron Bean, Michael Tremeer, Jiyuan An, Owen M. Siggs, Puya Gharahkhani, Jamie E. Craig, Alex W. Hewitt, Maciej Trzaskowski, and Stuart MacGregor. 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. *The American Journal of Human Genetics*, 108(7):1204–1216, July 2021. doi: 10.1016/j.ajhg.2021.05.005. URL https://doi.org/10.1016/j.ajhg.2021.05.005.

[13] David P. Johns, Julia A.E. Walters, and E. Haydn Walters. Diagnosis and early detection of

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

copd using spirometry. *Journal of Thoracic Disease*, 6(11), 2014. ISSN 2077-6624. URL https://jtd.amegroups.com/article/view/3088.

[14] Bartolome R. Celli. The importance of spirometry in copd and asthma: Effect on approach to management. *Chest*, 117(2, Supplement):15S–19S, 2000. ISSN 0012-3692. doi: https://doi.org/10.1378/chest.117.2_suppl.15S. URL https://www.sciencedirect.com/science/article/pii/S0012369215527505.

[15] M R Miller, J Hankinson, V Brusasco, F Burgos, R Casaburi, A Coates, R Crapo, P Enright, C P M van der Grinten, P Gustafsson, R Jensen, D C Johnson, N MacIntyre, R McKay, D Navajas, O F Pedersen, R Pellegrino, G Viegi, J Wanger, and ATS/ERS Task Force. Standardisation of spirometry. *Eur. Respir. J.*, 26(2):319–338, August 2005.

[16] David M Mannino and A Sonia Buist. Global burden of COPD: risk factors, prevalence, and future trends. *The Lancet*, 370(9589):765–773, September 2007. doi: 10.1016/s0140-6736(07)61380-4. URL https://doi.org/10.1016/s0140-6736(07)61380-4.

[17] Jørgen Vestbo, Suzanne S. Hurd, Alvar G. Agustí, Paul W. Jones, Claus Vogelmeier, Antonio Anzueto, Peter J. Barnes, Leonardo M. Fabbri, Fernando J. Martinez, Masaharu Nishimura, Robert A. Stockley, Don D. Sin, and Roberto Rodriguez-Roisin. Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease. *American Journal of Respiratory and Critical Care Medicine*, 187(4):347–365, feb 2013. doi: 10.1164/rccm.201204-0596pp.

[18] Thomas R. Dawber, H. Emerson Thomas, and Patricia M. McNamara. Characteristics of the dicrotic notch of the arterial pulse wave in coronary heart disease. *Angiology*, 24(4):244–255, April 1973. doi: 10.1177/000331977302400407. URL https://doi.org/10.1177/000331977302400407.

[19] John Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1–R39, February 2007. doi: 10.1088/0967-3334/28/3/r01. URL https://doi.org/10.1088/0967-3334/28/3/r01.

[20] N. Selvaraj, A. Jaryal, J. Santhosh, K. K. Deepak, and S. Anand. Assessment of heart rate variability derived from finger-tip photoplethysmography as compared to electrocardiography.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Journal of Medical Engineering &amp Technology, 32(6):479–484, January 2008. doi: 10.1080/03091900701781317. URL https://doi.org/10.1080/03091900701781317.

[21] Xiaochuan He, Rafik A. Goubran, and Xiaoping P. Liu. Secondary peak detection of PPG signal for continuous cuffless arterial blood pressure measurement. *IEEE Transactions on Instrumentation and Measurement*, 63(6):1431–1439, June 2014. doi: 10.1109/tim.2014.2299524. URL https://doi.org/10.1109/tim.2014.2299524.

[22] Farhan Riaz, Muhammad Ajmal Azad, Junaid Arshad, Muhammad Imran, Ali Hassan, and Saad Rehman. Pervasive blood pressure monitoring using photoplethysmogram (PPG) sensor. *Future Generation Computer Systems*, 98:120–130, September 2019. doi: 10.1016/j.future.2019.02.032. URL https://doi.org/10.1016/j.future.2019.02.032.

[23] Edwin Silverman, Scott Weiss, Steven Shapiro, and David Lomas. *Respiratory genetics*. CRC Press, 2005.

[24] Nick Shrine, Anna L Guyatt, A Mesut Erzurumluoglu, Victoria E Jackson, Brian D Hobbs, Carl A Melbourne, Chiara Batini, Katherine A Fawcett, Kijoung Song, Phuwanat Sakornsakolpat, Xingnan Li, Ruth Boxall, Nicola F Reeve, Ma'en Obeidat, Jing Hua Zhao, Matthias Wielscher, Stefan Weiss, Katherine A Kentistou, James P Cook, Benjamin B Sun, Jian Zhou, Jennie Hui, Stefan Karrasch, Medea Imboden, Sarah E Harris, Jonathan Marten, Stefan Enroth, Shona M Kerr, Ida Surakka, Veronique Vitart, Terho Lehtimäki, Richard J Allen, Per S Bakke, Terri H Beaty, Eugene R Bleecker, Yohan Bossé, Corry-Anke Brandsma, Zhengming Chen, James D Crapo, John Danesh, Dawn L DeMeo, Frank Dudbridge, Ralf Ewert, Christian Gieger, Amund Gulsvik, Anna L Hansell, Ke Hao, Joshua D Hoffman, John E Hokanson, Georg Homuth, Peter K Joshi, Philippe Joubert, Claudia Langenberg, Xuan Li, Liming Li, Kuang Lin, Lars Lind, Nicholas Locantore, Jian'an Luan, Anubha Mahajan, Joseph C Maranville, Alison Murray, David C Nickle, Richard Packer, Margaret M Parker, Megan L Paynton, David J Porteous, Dmitry Prokopenko, Dandi Qiao, Rajesh Rawal, Heiko Runz, Ian Sayers, Don D Sin, Blair H Smith, María Soler Artigas, David Sparrow, Ruth Tal-Singer, Paul R H J Timmers, Maarten Van den Berge, John C Whittaker, Prescott G Woodruff, Laura M Yerges-Armstrong, Olga G Troyanskaya, Olli T Raitakari, Mika Kähönen, Ozren Polašek, Ulf Gyllensten, Igor Rudan, Ian J Deary, Nicole M Probst-Hensch, Holger Schulz, Alan L James, James F Wilson, Beate Stubbe, Eleftheria Zeggini, Marjo-Riitta Jarvelin, Nick Wareham, Edwin K Silverman,

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Caroline Hayward, Andrew P Morris, Adam S Butterworth, Robert A Scott, Robin G Walters, Deborah A Meyers, Michael H Cho, David P Strachan, Ian P Hall, Martin D Tobin, Louise V Wain, and Understanding Society Scientific Group. New genetic signals for lung function highlight pathways and chronic obstructive pulmonary disease associations across multiple ancestries. *Nat. Genet.*, 51(3):481–493, March 2019.

[25] Nick Shrine, Abril G Izquierdo, Jing Chen, Richard Packer, Robert J Hall, Anna L Guyatt, Chiara Batini, Rebecca J Thompson, Chandan Pavuluri, Vidhi Malik, Brian D Hobbs, Matthew Moll, Wonji Kim, Ruth Tal-Singer, Per Bakke, Katherine A Fawcett, Catherine John, Kayesha Coley, Noemi Nicole Piga, Alfred Pozarickij, Kuang Lin, Iona Y Millwood, Zhengming Chen, Liming Li, China Kadoorie Biobank Collaborative Group, Sara R A Wijnant, Lies Lahousse, Guy Brusselle, Andre G Uitterlinden, Ani Manichaikul, Elizabeth C Oelsner, Stephen S Rich, R Graham Barr, Shona M Kerr, Veronique Vitart, Michael R Brown, Matthias Wielscher, Medea Imboden, Ayoung Jeong, Traci M Bartz, Sina A Gharib, Claudia Flexeder, Stefan Karrasch, Christian Gieger, Annette Peters, Beate Stubbe, Xiaowei Hu, Victor E Ortega, Deborah A Meyers, Eugene R Bleecker, Stacey B Gabriel, Namrata Gupta, Albert Vernon Smith, Jian'an Luan, Jing-Hua Zhao, Ailin F Hansen, Arnulf Langhammer, Cristen Willer, Laxmi Bhatta, David Porteous, Blair H Smith, Archie Campbell, Tamar Sofer, Jiwon Lee, Martha L Daviglus, Bing Yu, Elise Lim, Hanfei Xu, George T O'Connor, Gaurav Thareja, Omar M E Albagha, Qatar Genome Program Research (QGPR) Consortium, Karsten Suhre, Raquel Granell, Tariq O Faquih, Pieter S Hiemstra, Annelies M Slats, Benjamin H Mullin, Jennie Hui, Alan James, John Beilby, Karina Patasova, Pirro Hysi, Jukka T Koskela, Annah B Wyss, Jianping Jin, Sinjini Sikdar, Mikyeong Lee, Sebastian May-Wilson, Nicola Pirastu, Katherine A Kentistou, Peter K Joshi, Paul R H J Timmers, Alexander T Williams, Robert C Free, Xueyang Wang, John L Morrison, Frank D Gilliland, Zhanghua Chen, Carol A Wang, Rachel E Foong, Sarah E Harris, Adele Taylor, Paul Redmond, James P Cook, Anubha Mahajan, Lars Lind, Teemu Palviainen, Terho Lehtimäki, Olli T Raitakari, Jaakko Kaprio, Taina Rantanen, Kirsi H Pietiläinen, Simon R Cox, Craig E Pennell, Graham L Hall, W James Gauderman, Chris Brightling, James F Wilson, Tuula Vasankari, Tarja Laitinen, Veikko Salomaa, Dennis O Mook-Kanamori, Nicholas J Timpson, Eleftheria Zeggini, Josée Dupuis, Caroline Hayward, Ben Brumpton, Claudia Langenberg, Stefan Weiss, Georg Homuth, Carsten Oliver Schmidt, Nicole Probst-Hensch, Marjo-Riitta Jarvelin, Alanna C Morrison, Ozren Polasek, Igor Rudan, Joo-Hyeon Lee, Ian Sayers, Emma L Rawlins, Frank Dudbridge, Edwin K Silverman, David P Strachan, Robin G Walters,

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Andrew P Morris, Stephanie J London, Michael H Cho, Louise V Wain, Ian P Hall, and Martin D Tobin. Multi-ancestry genome-wide association analyses improve resolution of genes and pathways influencing lung function and chronic obstructive pulmonary disease risk. *Nat. Genet.*, 55(3):410–422, March 2023.

[26] Kenneth Fung, Julia Ramírez, Helen R. Warren, Nay Aung, Aaron M. Lee, Evan Tzanis, Steffen E. Petersen, and Patricia B. Munroe. Genome-wide association study identifies loci for arterial stiffness index in 127, 121 UK biobank participants. *Scientific Reports*, 9(1), June 2019. doi: 10.1038/s41598-019-45703-0. URL `https://doi.org/10.1038/s41598-019-45703-0`.

[27] Seyedeh M. Zekavat, Krishna Aragam, Connor Emdin, Amit V. Khera, Derek Klarin, Hongyu Zhao, and Pradeep Natarajan. Genetic association of finger photoplethysmography-derived arterial stiffness index with blood pressure and coronary artery disease. *Arteriosclerosis, Thrombosis, and Vascular Biology*, 39(6):1253–1261, June 2019. doi: 10.1161/atvbaha.119.312626. URL `https://doi.org/10.1161/atvbaha.119.312626`.

[28] Jonathan W. Cunningham, Paolo Di Achille, Valerie N. Morrill, Lu-Chen Weng, Seung Hoan Choi, Shaan Khurshid, Victor Nauffal, James P. Pirruccello, Scott D. Solomon, Puneet Batra, Jennifer E. Ho, Anthony A. Philippakis, Patrick T. Ellinor, and Steven A. Lubitz. Machine learning to understand genetic and clinical factors associated with the pulse waveform dicrotic notch. *Circulation: Genomic and Precision Medicine*, 16(1), February 2023. doi: 10.1161/circgen.121.003676. URL `https://doi.org/10.1161/circgen.121.003676`.

[29] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv*, 2013. doi: 10.48550/ARXIV.1312.6114. URL `https://arxiv.org/abs/1312.6114`.

[30] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. In *Parallel distributed processing: explorations in the microstructure of cognition, vol. 1: foundations*, pages 318–362. MIT Press, Cambridge, MA, USA, January 1986.

[31] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

[32] Brian D Hobbs, Kim de Jong, Maxime Lamontagne, Yohan Bossé, Nick Shrine, María Soler

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

Artigas, Louise V Wain, Ian P Hall, Victoria E Jackson, Annah B Wyss, Stephanie J London, Kari E North, Nora Franceschini, David P Strachan, Terri H Beaty, John E Hokanson, James D Crapo, Peter J Castaldi, Robert P Chase, Traci M Bartz, Susan R Heckbert, Bruce M Psaty, Sina A Gharib, Pieter Zanen, Jan W Lammers, Matthijs Oudkerk, H J Groen, Nicholas Locantore, Ruth Tal-Singer, Stephen I Rennard, Jørgen Vestbo, Wim Timens, Peter D Paré, Jeanne C Latourelle, Josée Dupuis, George T O'Connor, Jemma B Wilk, Woo Jin Kim, Mi Kyeong Lee, Yeon-Mok Oh, Judith M Vonk, Harry J de Koning, Shuguang Leng, Steven A Belinsky, Yohannes Tesfaigzi, Ani Manichaikul, Xin-Qun Wang, Stephen S Rich, R Graham Barr, David Sparrow, Augusto A Litonjua, Per Bakke, Amund Gulsvik, Lies Lahousse, Guy G Brusselle, Bruno H Stricker, André G Uitterlinden, Elizabeth J Ampleford, Eugene R Bleecker, Prescott G Woodruff, Deborah A Meyers, Dandi Qiao, David A Lomas, Jae-Joon Yim, Deog Kyeom Kim, Iwona Hawrylkiewicz, Pawel Sliwinski, Megan Hardin, Tasha E Fingerlin, David A Schwartz, Dirkje S Postma, William MacNee, Martin D Tobin, Edwin K Silverman, H Marike Boezen, Michael H Cho, COPDGene Investigators, ECLIPSE Investigators, LifeLines Investigators, SPIROMICS Research Group, International COPD Genetics Network Investigators, UK BiLEVE Investigators, and International COPD Genetics Consortium. Genetic loci associated with chronic obstructive pulmonary disease overlap with loci for lung function and pulmonary fibrosis. *Nat. Genet.*, 49(3):426–432, March 2017.

[33] Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins. UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.*, 12(3):e1001779, March 2015.

[34] Sandrine C Millasseau, James M Ritter, Kenji Takazawa, and Philip J Chowienczyk. Contour analysis of the photoplethysmographic pulse measured at the finger. *Journal of Hypertension*, 24(8):1449–1456, August 2006. doi: 10.1097/01.hjh.0000239277.05068.87. URL `https://doi.org/10.1097/01.hjh.0000239277.05068.87`.

[35] Ramakrishna Mukkamala, Jin-Oh Hahn, Omer T. Inan, Lalit K. Mestha, Chang-Sei Kim, Hakan Toreyin, and Survi Kyal. Toward ubiquitous blood pressure monitoring via pulse transit time: Theory and practice. *IEEE Transactions on Biomedical Engineering*, 62(8):1879–1901, Au-

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

gust 2015. doi: 10.1109/tbme.2015.2441951. URL https://doi.org/10.1109/tbme.2015.2441951.

[36] M. Abdullah Said, Ruben N. Eppinga, Erik Lipsic, Niek Verweij, and Pim van der Harst. Relationship of arterial stiffness index and pulse pressure with cardiovascular disease and mortality. *Journal of the American Heart Association*, 7(2), January 2018. doi: 10.1161/jaha.117.007621. URL https://doi.org/10.1161/jaha.117.007621.

[37] Po-Ru Loh, George Tucker, Brendan K Bulik-Sullivan, Bjarni J Vilhjálmsson, Hilary K Finucane, Rany M Salem, Daniel I Chasman, Paul M Ridker, Benjamin M Neale, Bonnie Berger, Nick Patterson, and Alkes L Price. Efficient bayesian mixed-model analysis increases association power in large cohorts. *Nat. Genet.*, 47(3):284–290, March 2015.

[38] Po-Ru Loh, Gleb Kichaev, Steven Gazal, Armin P Schoech, and Alkes L Price. Mixed-model association for biobank-scale datasets. *Nat. Genet.*, 50(7):906–908, July 2018.

[39] Elliot Sollis, Abayomi Mosaku, Ala Abid, Annalisa Buniello, Maria Cerezo, Laurent Gil, Tudor Groza, Osman Güneş, Peggy Hall, James Hayhurst, Arwa Ibrahim, Yue Ji, Sajo John, Elizabeth Lewis, Jacqueline A L MacArthur, Aoife McMahon, David Osumi-Sutherland, Kalliope Panoutsopoulou, Zoë Pendlington, Santhi Ramachandran, Ray Stefancsik, Jonathan Stewart, Patricia Whetzel, Robert Wilson, Lucia Hindorff, Fiona Cunningham, Samuel A Lambert, Michael Inouye, Helen Parkinson, and Laura W Harris. The NHGRI-EBI GWAS catalog: knowledgebase and deposition resource. *Nucleic Acids Res.*, 51(D1):D977–D985, January 2023.

[40] Valentina Iotchkova, Graham R S Ritchie, Matthias Geihs, Sandro Morganella, Josine L Min, Klaudia Walter, Nicholas John Timpson, UK10K Consortium, Ian Dunham, Ewan Birney, and Nicole Soranzo. GARFIELD classifies disease-relevant genomic features through integration of functional annotations with association signals. *Nat. Genet.*, 51(2):343–353, February 2019.

[41] Cory Y McLean, Dave Bristor, Michael Hiller, Shoa L Clarke, Bruce T Schaar, Craig B Lowe, Aaron M Wenger, and Gill Bejerano. GREAT improves functional interpretation of cis-regulatory regions. *Nat. Biotechnol.*, 28(5):495–501, May 2010.

[42] Alicia R Martin, Masahiro Kanai, Yoichiro Kamatani, Yukinori Okada, Benjamin M Neale, and Mark J Daly. Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat. Genet.*, 51(4):584–591, April 2019.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

[43] Elizabeth A Regan, John E Hokanson, James R Murphy, Barry Make, David A Lynch, Terri H Beaty, Douglas Curran-Everett, Edwin K Silverman, and James D Crapo. Genetic epidemiology of COPD (COPDGene) study design. *COPD*, 7(1):32–43, February 2010.

[44] N Day, S Oakes, R Luben, K T Khaw, S Bingham, A Welch, and N Wareham. EPIC-Norfolk: study design and characteristics of the cohort. european prospective investigation of cancer. *Br. J. Cancer*, 80 Suppl 1:95–103, July 1999.

[45] Dongbing Lai, Tae-Hwi Schwantes-An, Marco Abreu, Grace Chan, Victor Hesselbrock, Chella Kamarajan, Yunlong Liu, Jacquelyn L. Meyers, John I. Nurnberger, Martin H. Plawecki, Leah Wetherill, Marc Schuckit, Pengyue Zhang, Howard J. Edenberg, Bernice Porjesz, Arpana Agrawal, and Tatiana Foroud. Gene-based polygenic risk scores analysis of alcohol use disorder in african americans. *Translational Psychiatry*, 12(1), July 2022. doi: 10.1038/s41398-022-02029-2. URL https://doi.org/10.1038/s41398-022-02029-2.

[46] Xiaoman Xing, Zhimin Ma, Mingyou Zhang, Ying Zhou, Wenfei Dong, and Mingxuan Song. An unobtrusive and calibration-free blood pressure estimation method using photoplethysmography and biometrics. *Sci. Rep.*, 9(1):8611, June 2019.

[47] Kari Hemminki, Xiangdong Liu, Jianguang Ji, Kristina Sundquist, and Jan Sundquist. Subsequent COPD and lung cancer in patients with autoimmune disease. *European Respiratory Journal*, 37(2):463–465, 2011.

[48] Te-Chun Shen, Cheng-Li Lin, Chia-Hung Chen, Chih-Yen Tu, Te-Chun Hsia, Chuen-Ming Shih, Wu-Huei Hsu, and Yen-Jung Chang. Increased risk of chronic obstructive pulmonary disease in patients with systemic lupus erythematosus: a population-based cohort study. *PLOS ONE*, 9(3):e91821, 2014.

[49] Nikolaos M. Siafakas, Ira Milona, Vasiliki Salesiotou, Vasiliki Filaditaki, Nikolaos Tzanakis, and Demosthenes Bouros. Respiratory muscle strength in hyperthyroidism before and after treatment. *American Review of Respiratory Disease*, 146(4):1025–1029, October 1992. doi: 10.1164/ajrccm/146.4.1025. URL https://doi.org/10.1164/ajrccm/146.4.1025.

[50] Jonas F Ludvigsson, Malin Inghammar, Marie Ekberg, and Arne Egesten. A nationwide cohort study of the risk of chronic obstructive pulmonary disease in coeliac disease. *Journal of Internal Medicine*, 271(5):481–489, 2012.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

[51] Luke J O'Connor and Alkes L Price. Distinguishing genetic correlation from causation across 52 diseases and complex traits. *Nat. Genet.*, 50(12):1728–1734, December 2018.

[52] Cathryn M Delude. Deep phenotyping: the details of disease. *Nature*, 527(7576):S14–S15, 2015.

[53] Matthias Kirchler, Stefan Konigorski, Matthias Norden, Christian Meltendorf, Marius Kloft, Claudia Schurmann, and Christoph Lippert. transferGWAS: GWAS of images using deep transfer learning. *Bioinformatics*, 38(14):3621–3628, July 2022.

[54] Ziqian Xie, Tao Zhang, Sangbae Kim, Jiaxiong Lu, Wanheng Zhang, Cheng-Hui Lin, Man-Ru Wu, Alexander Davis, Roomasa Channa, Luca Giancardo, Han Chen, Sui Wang, Rui Chen, and Degui Zhi. iGWAS: image-based genome-wide association of self-supervised deep phenotyping of human medical images. *medRxiv*, 2022. doi: 10.1101/2022.05.26.22275626. URL `https://doi.org/10.1101/2022.05.26.22275626`.

[55] Adityanarayanan Radhakrishnan, Sam Freesun Friedman, Shaan Khurshid, Kenney Ng, Puneet Batra, Steven Lubitz, Anthony Philippakis, and Caroline Uhler. A cross-modal autoencoder framework learns holistic representations of cardiovascular state. *bioRxiv*, 2022. doi: 10.1101/2022.05.26.493497. URL `https://doi.org/10.1101/2022.05.26.493497`.

[56] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[57] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1558–1566, New York, New York, USA, 20–22 Jun 2016. PMLR. URL `https://proceedings.mlr.press/v48/larsen16.html`.

[58] Jin J Zhou, Michael H Cho, Christoph Lange, Sharon Lutz, Edwin K Silverman, and Nan M Laird. Integrating multiple correlated phenotypes for genetic association analysis by maximizing heritability. *Hum. Hered.*, 79(2):93–104, June 2015.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

[59] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/forum?id=Sy2fzU9gl`.

[60] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2649–2658. PMLR, 10–15 Jul 2018. URL `https://proceedings.mlr.press/v80/kim18b.html`.

[61] Abhishek Kumar, Prasanna Sattigeri, and Avinash Balakrishnan. Variational inference of disentangled latent concepts from unlabeled observations. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=H1kG7GZAW`.

[62] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR, 2019.

[63] Davin Hill, Max Torop, Aria Masoomi, Peter J Castaldi, Edwin K Silverman, Sandeep Bodduluri, Surya P Bhatt, Taedong Yun, Farhad Hormozdiari, Cory Y McLean, Jennifer Dy, Michael H Cho, and Brian D Hobbs. Deep learning utilizing suboptimal spirometry data to improve lung function and mortality prediction in the UK Biobank. *Preprint*, 2023.

[64] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. December 2014.

[65] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale machine learning on heterogeneous distributed systems. March 2016.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

[66] Brendan K Bulik-Sullivan, Po-Ru Loh, Hilary K Finucane, Stephan Ripke, Jian Yang, Schizophrenia Working Group of the Psychiatric Genomics Consortium, Nick Patterson, Mark J Daly, Alkes L Price, and Benjamin M Neale. LD score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.*, 47(3):291–295, March 2015.

[67] Michael H Cho, Merry-Lynn N McDonald, Xiaobo Zhou, Manuel Mattheisen, Peter J Castaldi, Craig P Hersh, Dawn L Demeo, Jody S Sylvia, John Ziniti, Nan M Laird, Christoph Lange, Augusto A Litonjua, David Sparrow, Richard Casaburi, R Graham Barr, Elizabeth A Regan, Barry J Make, John E Hokanson, Sharon Lutz, Tanda Murray Dudenkov, Homayoon Farzadegan, Jacqueline B Hetmanski, Ruth Tal-Singer, David A Lomas, Per Bakke, Amund Gulsvik, James D Crapo, Edwin K Silverman, Terri H Beaty, and NETT Genetics, ICGN, ECLIPSE and COPDGene Investigators. Risk loci for chronic obstructive pulmonary disease: a genome-wide association study and meta-analysis. *Lancet Respir Med*, 2(3):214–225, March 2014.

[68] Sayantan Das, Lukas Forer, Sebastian Schönherr, Carlo Sidore, Adam E Locke, Alan Kwong, Scott I Vrieze, Emily Y Chew, Shawn Levy, Matt McGue, David Schlessinger, Dwight Stambolian, Po-Ru Loh, William G Iacono, Anand Swaroop, Laura J Scott, Francesco Cucca, Florian Kronenberg, Michael Boehnke, Gonçalo R Abecasis, and Christian Fuchsberger. Next-generation genotype imputation service and methods. *Nature Genetics*, 48(10):1284–1287, August 2016. doi: 10.1038/ng.3656. URL `https://doi.org/10.1038/ng.3656`.

[69] Ian B Stanaway, Taryn O Hall, Elisabeth A Rosenthal, Melody Palmer, Vivek Naranbhai, Rachel Knevel, Bahram Namjou-Khales, Robert J Carroll, Krzysztof Kiryluk, Adam S Gordon, Jodell Linder, Kayla Marie Howell, Brandy M Mapes, Frederick T J Lin, Yoonjung Yoonie Joo, M Geoffrey Hayes, Ali G Gharavi, Sarah A Pendergrass, Marylyn D Ritchie, Mariza de Andrade, Damien C Croteau-Chonka, Soumya Raychaudhuri, Scott T Weiss, Matt Lebo, Sami S Amr, David Carrell, Eric B Larson, Christopher G Chute, Laura Jarmila Rasmussen-Torvik, Megan J Roy-Puckelwartz, Patrick Sleiman, Hakon Hakonarson, Rongling Li, Elizabeth W Karlson, Josh F Peterson, Iftikhar J Kullo, Rex Chisholm, Joshua Charles Denny, Gail P Jarvik, eMERGE Network, and David R Crosslin. The eMERGE genotype set of 83,717 subjects imputed to 40 million variants genome wide and association with the herpes zoster medical record phenotype. *Genet. Epidemiol.*, 43(1):63–81, February 2019.

medRxiv preprint doi: https://doi.org/10.1101/2023.04.28.23289285; this version posted August 29, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license .

[70] Shaun Purcell, Benjamin Neale, Kathe Todd-Brown, Lori Thomas, Manuel A.R. Ferreira, David Bender, Julian Maller, Pamela Sklar, Paul I.W. de Bakker, Mark J. Daly, and Pak C. Sham. PLINK: A tool set for whole-genome association and population-based linkage analyses. *The American Journal of Human Genetics*, 81(3):559–575, September 2007. doi: 10.1086/519795. URL https://doi.org/10.1086/519795.

[71] Christopher C Chang, Carson C Chow, Laurent CAM Tellier, Shashaank Vattikuti, Shaun M Purcell, and James J Lee. Second-generation PLINK: rising to the challenge of larger and richer datasets. *GigaScience*, 4(1), February 2015. doi: 10.1186/s13742-015-0047-8. URL https://doi.org/10.1186/s13742-015-0047-8.