Age-dependent topic modelling of comorbidities in UK Biobank identifies disease subtypes with differential genetic risk

Xilin Jiang[1,2,3,4,5,6], Martin Jinye Zhang[4,7§], Yidong Zhang[1,8,9§], Michael Inouye[5,6,10,11,12,13], Chris Holmes[1,2,14], Alkes L. Price[4,7,15]*, Gil McVean[1]*

Affiliations

[1] Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, University of Oxford, Oxford OX3 7LF, UK

[2] Department of Statistics, University of Oxford, Oxford OX1 3LB, UK

[3] Wellcome Centre for Human Genetics, University of Oxford, Oxford OX3 7BN, UK

[4] Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, MA, USA

[5] British Heart Foundation Cardiovascular Epidemiology Unit, Department of Public Health and Primary Care, University of Cambridge, Cambridge UK

[6] Heart and Lung Research Institute, University of Cambridge, Cambridge UK

[7] Program in Medical and Population Genetics, Broad Institute of MIT and Harvard, Cambridge, MA, USA

[8] CAMS China Oxford Institute, Nuffield Department of Medicine, University of Oxford, Oxford OX3 7BN, UK

[9] Department of Radiation Oncology, Peking Union Medical College Hospital, Chinese Academy of Medical Sciences and Peking Union Medical College, Beijing, China

[10] Cambridge Baker Systems Genomics Initiative, Department of Public Health and Primary Care, University of Cambridge, Cambridge, UK.

[11] Health Data Research UK Cambridge, Wellcome Genome Campus and University of Cambridge, Cambridge, UK.

[12] British Heart Foundation Cambridge Centre of Research Excellence, Department of Clinical Medicine, University of Cambridge, Cambridge, UK.

[13] Cambridge Baker Systems Genomics Initiative, Baker Heart and Diabetes Institute, Melbourne, VIC, Australia.

[14] The Alan Turing Institute, London NW1 2DB, UK

[15] Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, MA, USA

[§] These authors contributed equally to this work

*These authors jointly supervised the work

Corresponding authors:

xilinjiang@hsph.harvard.edu

aprice@hsph.harvard.edu

gil.mcvean@bdi.ox.ac.uk

Abstract

Longitudinal data from electronic health records (EHR) has immense potential to improve clinical diagnoses and personalised medicine, motivating efforts to identify disease subtypes from age-dependent patient comorbidity information. We introduce an age-dependent topic modelling (ATM) method that provides a low-rank representation of longitudinal records of hundreds of distinct diseases in large EHR data sets. The model learns, and assigns to each individual, topic weights for several disease topics, each of which reflects a set of diseases that tend to co-occur as a function of age. Simulations show that ATM attains high accuracy in distinguishing distinct age-dependent comorbidity profiles. We applied ATM to 282,957 UK Biobank samples, analysing 1,726,144 disease diagnoses spanning 348 diseases with ≥1,000 incidences. We inferred 10 disease topics optimising model fit. We identified 52 diseases with heterogeneous comorbidity profiles (≥500 incidences assigned to each of ≥2 topics), including breast cancer, type 2 diabetes (T2D), hypertension, and hypercholesterolemia; for most of these diseases, topic assignments were highly age-dependent, suggesting differences in disease aetiology for early-onset vs. late-onset disease. We defined subtypes of the 52 heterogeneous diseases based on the topic assignments, and compared genetic risk across subtypes using polygenic risk scores (PRS). We identified 18 disease subtypes whose PRS differed significantly from other subtypes of the same disease, including a subtype of T2D characterised by cardiovascular comorbidities and a subtype of asthma characterised by dermatological comorbidities. We further identified specific SNPs underlying these differences. For example, the T2D-associated SNP rs1063192 in the *CDKN2B* locus has a higher odds ratio in the top quartile of cardiovascular topic weight (1.19±0.02) than in the bottom quartile (1.08±0.02) ($P=4 \times 10^{-5}$ for difference). In conclusion, ATM identifies disease subtypes with differential genome-wide and locus-specific genetic risk profiles.

Introduction

Longitudinal electronic health record (EHR) data, encompassing diagnoses across hundreds of distinct diseases, offers immense potential to improve clinical diagnoses and personalised medicine[1]. Despite intense interest in both the genetic relationships between distinct diseases[2–11] and the genetic relationships between biological subtypes of disease[12–15], there has been limited progress on classifying disease phenotypes into groups of diseases with frequent co-occurrences (comorbidities) and leveraging comorbidities to identify disease subtypes. Low-rank modelling has appealing theoretical properties[16,17] and has produced promising applications[18–24] to infer meaningful representations of high-dimensional data. In particular, low-rank representation is an appealing way to summarise data across hundreds of distinct diseases[25–27], providing the potential to identify patient-level comorbidity patterns and distinguish disease subtypes. Disease subtypes inferred from EHR data could be validated by comparing genetic profiles across subtypes, which is possible with emerging data sets that link genetic data with EHR data[28–31].

Previous studies have used low-rank representation to identify shared genetic components[25–27] across multiple distinct diseases, identifying relationships between diseases and generating valuable biological insights. However, age at diagnosis information in longitudinal EHR data has the potential to improve such efforts. For example, a recent study used longitudinal disease trajectories to identify disease pairs with statistically significant directionality[32], suggesting that age information could be leveraged to infer comorbidity profiles that capture temporal information. In addition, patient-level comorbidity information could potentially be leveraged to identify biological subtypes of disease, complementing its application to increase power for identifying genetic associations[12] and to cluster disease-associated variants into biological pathways[8]; disease subtypes are fundamental to disease aetiology[14,33–36].

Here, we propose an age-dependent topic modelling (ATM) method to provide a low-rank representation of longitudinal disease records. ATM learns, and assigns to each individual, topic weights for several disease topics, each of which reflects a set of diseases that tend to co-occur as a function of age. We applied ATM to 1.7 million disease diagnoses spanning 348 diseases in the UK Biobank, inferring 10 disease topics. We identified 52 diseases with heterogeneous comorbidity profiles that enabled us to define disease subtypes. We used genetic data to validate the disease subtypes, showing that they exhibit differential genome-wide and locus-specific genetic risk profiles.

Results

Overview of methods

We propose an age-dependent topic modelling (ATM) model, providing a low-rank representation of longitudinal records of hundreds of distinct diseases in large EHR data sets (Figure 1, Methods). The model assigns to each individual *topic weights* for several *disease topics*; each disease topic reflects a set of diseases that tend to co-occur as a function of age, quantified by age-dependent *topic loadings* for each disease. The model assumes that for each disease diagnosis, a topic is sampled based on the individual's topic weights (which sum to 1 across topics, for a given individual), and a disease is sampled based on the individual's age and the age-dependent topic loadings (which sum to 1 across diseases, for a given topic at a given age). The model generalises the latent dirichlet allocation (LDA) model[37,38] by allowing topic loadings for each topic to vary with age (Supplementary Note, Supplementary Figure 1).

We developed a method to fit this model that addresses several challenges inherent to large EHR data sets; the method estimates topic weights for each individual, topic loadings for each disease, and posterior diagnosis-specific topic probabilities for each disease diagnosis. First, we derived a scalable deterministic method that uses numerical approximation approaches to fit the parameters of the model, addressing the challenge of computational cost. Second, we used the prediction odds ratio[39] to compare model structures (e.g. number of topics and parametric form of topic loadings as a function of age), addressing the challenge of appropriate model selection; roughly, the prediction odds ratio quantifies the accuracy of correctly predicting disease diagnoses in held-out patients using comorbidity information, compared to a predictor based only on prevalence (see Methods and Supplementary Table 1). Third, we employed collapsed variational inference[40], addressing the challenge of sparsity in the data (e.g. in UK Biobank data that we analysed, the average patient has diagnoses for 6 of 348 diseases analysed); collapsed variational inference outperformed mean-field variational inference[37] in empirical data. Further details are provided in the Methods section and Supplementary Note; we have publicly released open-source software implementing the method (see Code Availability).

We applied ATM to longitudinal records of 282,957 individuals from the UK Biobank[29], containing a total of 1,726,144 disease diagnoses spanning 348 diseases (see Data Availability). Each disease diagnosis had an associated age at diagnosis, defined as the earliest age of reported diagnosis of the disease in that individual; we caution that age at diagnosis may differ from age at disease onset (see Discussion). ATM does not use genetic data, but we used genetic data to validate the inferred topics (Methods).

Simulations

We performed simulations to compare ATM with latent dirichlet allocation (LDA)[37,38], a simpler topic modelling approach that does not model age. We simulated 61,000 disease diagnoses spanning 20 diseases in 10,000 individuals, using the ATM generative model; the average

number of disease diagnoses per individual (6.1), ratio of #individuals/#diseases (500), topic loadings, and standard deviation in age at diagnosis (8.5 years for each disease) were chosen to match empirical UK Biobank data. We assigned each disease diagnosis to one of two subtypes for the underlying disease based on age and other subtype differences, considering high, medium, or low age-dependent effects by specifying an average difference of 20, 10, or 5 years respectively in age at diagnosis for the two subtypes. For each level of age-dependent effects, we varied the proportion of diagnoses belonging to the first subtype (*subtype sample size proportion*) from 10-50%. Further details of the simulation framework are provided in the Methods section. Our primary metric for evaluating the LDA and ATM methods was area under the precision-recall curve (AUPRC)[41], where precision is defined as the proportion of disease diagnoses that a given method assigned to the first subtype that were assigned correctly and recall is defined as the proportion of disease diagnoses truly belonging to the first subtype that were assigned correctly. We discretized the subtype assigned to each disease diagnosis by a given method by assigning the subtype with higher inferred probability. We note that AUPRC is larger when classifying the smaller subtype; results using the second subtype as the classification target are also provided. We used AUPRC (instead of prediction odds ratio) in our simulations because the underlying truth is known. Further details and justifications of metrics used in this study are provided in the Methods section and Supplementary Table 1.

In simulations with high age-dependent effects, ATM attained much higher AUPRC than LDA across all values of subtype sample size proportion (AUPRC difference: 24%-42%), with both methods performing better at more balanced ratios (Figure 2, Supplementary Table 2). Accordingly, ATM attained both higher precision and higher recall than LDA (Supplementary Figure 2). Results were qualitatively similar when using the second subtype as the classification target (Supplementary Figure 3). In simulations with medium or low age-dependent effects, ATM continued to outperform LDA but with smaller differences between the methods. In simulations without age-dependent effects, ATM slightly underperformed LDA (Supplementary Figure 4A).

We performed three secondary analyses. First, we varied the number of individuals, number of diseases, or number of disease diagnoses per individual. ATM continued to outperform LDA in each case, although increasing the number of individuals or the number of disease diagnoses per individual did not always increase AUPRC (Supplementary Figure 4B). Second, we performed simulations in which we increased the number of subtypes from two to five and changed the number of diseases to 50, and compared ATM models trained using different numbers of topics (in 80% training data) by computing the prediction odds ratio; we used the prediction odds ratio (instead of AUPRC) in this analysis both because it is a better metric to evaluate the overall model fit to the data, and because it is unclear how to compare AUPRC across scenarios of varying topic numbers (see Supplementary Table 1). We confirmed that the prediction odds ratio was maximised using five topics, validating the use of the prediction odds ratio for model

selection (Supplementary Figure 5A). Third, we computed the accuracy of inferred topic loadings, topic weights, and grouping accuracy (defined as proportion of pairs of diseases truly belonging to the same topic that ATM correctly assigned to the same topic), varying the number of individuals and number of diseases diagnoses per individual. We determined that ATM also performed well under these metrics (Supplementary Figure 5B-E).

We conclude that ATM (which models age) assigns disease diagnoses to subtypes with higher accuracy than LDA (which does not model age) in simulations with age-dependent effects. We caution that our simulations largely represent a best-case scenario for ATM given that the generative model and inference model are very similar (although there are some differences, e.g. topic loadings were generated using a model different from the inference model), thus it is important to analyse empirical data to validate the method.

Age-dependent disease topic loadings capture comorbidity profiles in the UK Biobank

We applied ATM to longitudinal records of 282,957 individuals from the UK Biobank[29]. We used Phecode[42] to define 1,726,144 disease diagnoses spanning 348 diseases with at least 1,000 diagnoses each; the average individual had 6.1 disease diagnoses, and the average disease had a standard deviation of 8.5 years in age at diagnosis. The optimal ATM model structure included 10 topics and modelled age-dependent topic loadings for each disease as a spline function with one knot (see below). We assigned names (and corresponding acronyms) to each of the 10 inferred topics based on the Phecode systems[42] assigned to diseases with high topic loadings (aggregated across ages) for that topic (Table 1, Supplementary Table 3).

Age-dependent topic loadings across all 10 topics and 348 diseases (stratified into Phecode systems), summarised as averages across age<60 and age≥60, are reported in Figure 3, Supplementary Figure 6, and Supplementary Table 4. Some topics such as NRI span diseases across the majority of Phecode systems, while other topics such as ARP are concentrated in a single Phecode system. Conversely, a single Phecode system may be split across multiple topics, e.g. the digestive system is split across UGI, LGI, and MDS. We note that topic loadings in diseases that span multiple topics are heavily age-dependent. For example, type 2 diabetes patients assigned to the CVD topic are associated with early onset of type 2 diabetes whereas type 2 diabetes patients assigned to MGND topic are associated with late onset of type 2 diabetes.

We performed seven secondary analyses to validate the inferred comorbidity topics. First, we fit ATM models with different model structures using 80% training data, and computed their prediction odds ratios using 20% testing data. The ATM model structure with 10 topics and age-dependent topic loadings modelled as a spline function performed optimally (Supplementary Figure 7; see Methods). Second, we confirmed that ATM attained higher prediction odds ratios than LDA across different values of the number of topics (Supplementary Figure 8). Third, we

reached similar conclusions using evidence lower bounds[39] (ELBO; see Supplementary Table 1) when fitting the model without splitting training and testing data (Supplementary Figure 9). Fourth, we confirmed that collapsed variational inference [40] outperformed mean-field variational inference [37] (Supplementary Figure 10). Fifth, we computed a co-occurrence odds ratio evaluating whether diseases grouped into the same topic by ATM in the training data have higher than random probability of co-occurring in the testing data (Supplementary Table 1). The co-occurrence odds ratio is consistently above one and increases with the number of comorbid diseases, for each inferred topic (Supplementary Figure 11). Sixth, we compared the topic loadings by repeating the inference on female-only or male-only populations and observed no major discrepancies, except for genitourinary topics MGND and FGND (topic loading R^2 (female vs. all) = 0.788, topic loading R^2 (male vs. all) = 0.773, Supplementary Figure 12). Lastly, we verified that BMI, sex, Townsend deprivation index, and birth year explained very little of the information in the inferred topics (Supplementary Table 3).

Disease topics capture known biology as well as the age-dependency of comorbidities for the same diseases. For example, early onset of essential hypertension is associated with the CVD topic [43], which captures the established connection between lipid dysfunction ("hypercholesterolemia") and cardiovascular diseases[44], while later onset of essential hypertension is associated with the CER topic, which pertains to type 2 diabetes, obesity and COPD (Figure 4A). Continuously varying age-dependent topic loadings for all 10 topics, restricted to diseases with high topic loadings, are reported in Supplementary Figure 13 and Supplementary Table 5. We note that most diseases have their topic loadings concentrated into a single topic (Figure 4B, Supplementary Figure 14A and Supplementary Table 4), and that most individuals have their topic weights concentrated into 1-2 topics (Figure 4C and Supplementary Figure 14B). For diseases spanning multiple topics (Supplementary Figure 6 and Supplementary Table 4), the assignment of type 2 diabetes patients to the CVD topic is consistent with known pathophysiology and epidemiology [45,46] and has been shown in other comorbidity clustering studies, e.g. with the Beta Cell and Lipodystrophy subtypes described in ref. [35] and the severe insulin-deficient diabetes (SIDD) subtype described in ref.[14], which are characterised by early onset of type 2 diabetes and have multiple morbidities including hypercholesterolemia, hyperlipidemia, and cardiovascular diseases [47]. In addition, early-onset breast cancer and late-onset breast cancer are associated with different topics, e.g. NRI and FGND, consistent with known treatment effects for breast cancer patients which increase susceptibility to infections, especially bacterial pneumonias[48] and hypothyroidism [49]

We conclude that ATM identifies latent disease topics that robustly compress age-dependent comorbidity profiles and capture disease comorbidities both within and across Phecode systems.

Disease subtypes defined by distinct topics are genetically heterogeneous

We sought to define disease subtypes based on the diagnosis-specific topic probabilities of each disease diagnosis. We assigned a discrete topic assignment to each disease diagnosis based on its maximum diagnosis-specific topic probability, and defined the disease subtype of each disease diagnosis based on the topic assignment. We restricted our disease subtype analyses to 52 diseases with at least 500 diagnoses assigned to each of two distinct subtypes (Methods, Supplementary Figure 6, Supplementary Figure 15, and Supplementary Table 6).

Age-dependent distributions of subtypes (topics) for four diseases (type 2 diabetes, asthma, hypercholesterolemia, and essential hypertension) are reported in Figure 5A and Supplementary Table 7; results for all 52 diseases are reported in Supplementary Figure 16 and Supplementary Table 7. The number of subtypes can be large, e.g. six subtypes for essential hypertension. Subtypes are often age-dependent, e.g. for the CVD and MGND subtypes of type 2 diabetes[14,35] (discussed above).

ATM and the resulting subtype assignments do not make use of genetic data. However, we used genetic data to assess genetic heterogeneity across inferred subtypes of each disease. We first assessed whether polygenic risk scores (PRS) for overall disease risk varied across subtypes of each disease; PRS were computed using BOLT-LMM with five-fold cross validation [50,51] (see Methods and Code Availability). Results for four diseases (from Figure 5A) are reported in Figure 5B and Supplementary Table 8; results for all 10 well-powered diseases (10 of 52 diseases with highest z-scores for nonzero SNP-heritability) are reported in Supplementary Figure 17 and Supplementary Table 8. We identified 18 disease-topic pairs (of 100 disease-topic pairs analysed) for which PRS values in disease cases vary with patient topic weight. For example, for essential hypertension, hypercholesterolemia, and type 2 diabetes, patients assigned to the CVD subtype had significantly higher PRS values than patients assigned to other subtypes. For essential hypertension, patients assigned to the CER subtype had significantly higher PRS values; for type 2 diabetes, patients assigned to the CER subtype had lower PRS values than the CVD subtype, even though the majority of type 2 diabetes diagnoses are assigned to the CER subtype. We further verified that most of the variation in PRS values with disease subtype could not be explained by age[52] or differences in subtype sample size (Supplementary Figure 18). These associations between subtypes (defined using comorbidity data) and PRS (defined using genetic data) imply that disease subtypes identified through comorbidity are genetically heterogeneous, consistent with differences in disease aetiology.

We further investigated whether subtype assignments (defined using comorbidity data) revealed subtype-specific excess genetic correlations. We estimated excess genetic correlations between pairs of disease subtypes (relative to genetic correlations between the underlying diseases). Excess genetic correlations for 15 disease subtypes (spanning 11 diseases and 3 topics: CER, MGND and CVD) are reported in Figure 6A and Supplementary Table 9 (relative to genetic correlations between the underlying diseases; Figure 6B), and excess genetic correlations for all

306 89 well-powered disease subtypes (89 of 378 disease subtypes with z-score > 4 for nonzero SNP-
307 heritability) are reported in Supplementary Figure 19 and Supplementary Table 9. Genetic
308 correlations between pairs of subtypes involving the same disease were significantly less than 1
309 (FDR<0.1) for hypertension (CER vs. CVD: $\rho = 0.86 \pm 0.04$, P=0.0004; MGND vs. CVD: $\rho =$
310 0.74 ± 0.05, P=3×10^{-8}) and type 2 diabetes (CER vs. MGND: $\rho = 0.64 \pm 0.09$, P=8×10^{-5})
311 (Figure 6A; Supplementary Table 9). In addition, we observed significant excess genetic
312 correlations (FDR<0.1) for 8 pairs of disease subtypes involving different diseases (Figure 6A;
313 Supplementary Table 9). We verified that the excess genetic correlations could not be explained
314 by non-disease-specific differences in the underlying topics (which are weakly heritable;
315 Supplementary Table 3) by repeating the analysis using disease cases and controls with matched
316 topic weights (Methods, Supplementary Figure 20). We also estimated subtype-specific SNP-
317 heritability and identified some instances of differences between subtypes, albeit with limited
318 power (Supp Table 10).
319
320 Finally, we used the population genetic parameter F_{ST}[53,54] to quantify genome-wide differences
321 in allele frequency between two subtypes of the same disease; we used F_{ST} on control sets with
322 matched topic weights to assess statistical significance while accounting for non-disease-specific
323 differences in the underlying topics (excess F_{ST}; Methods). We determined that 63 of 104 pairs
324 of disease subtypes involving the same disease (spanning 29 of 49 diseases, excluding 3 diseases
325 that did not have enough controls with matched topic weights) had significant excess F_{ST}
326 estimates (FDR < 0.1) (Supplementary Figure 21, Supplementary Table 11). For example, the
327 CVD, CER, and MGND subtypes of type 2 diabetes had significant excess F_{ST} estimates (F-
328 statistic=0.0003, P=0.001 based on 1,000 matched control sets). This provides further evidence
329 that disease subtypes as determined by comorbidity have different molecular and physiological
330 aetiologies.
331
332 We conclude that disease subtypes defined by distinct topics are genetically heterogeneous.
333
334 *Disease-associated SNPs have subtype-dependent effects*
335 We hypothesised that disease genes and pathways might differentially impact the disease
336 subtypes identified by ATM. We investigated the genetic heterogeneity between disease
337 subtypes at the level of individual disease-associated variants. We employed a statistical test that
338 tests for SNP x topic interaction effects on disease phenotype in the presence of separate SNP
339 and topic effects (Methods). We verified via simulations that this statistical test is well-calibrated
340 under a broad range of scenarios with no true interaction, including direct effect of topic on
341 disease, direct effect of disease on topic, pleiotropic SNP effects on disease and topic, and
342 nonlinear effects (Supplementary Figure 22). We also assessed the power to detect true
343 interactions (Supplementary Figure 23). To limit the number of hypotheses tested, we applied
344 this test to independent SNPs with genome-wide significant main effects on disease (Methods).
345 We thus performed 2,530 statistical tests spanning 888 disease-associated SNPs, 14 diseases, and
346 35 disease subtypes (Supplementary Table 12). We assessed statistical significance using global

FDR<0.1 across the 2,530 statistical tests. We also computed main SNP effects specific to each quartile of topic weights across individuals, as an alternative way to represent SNP x topic interactions.

We identified 43 SNP x topic interactions at FDR<0.1 (Figure 7, Supplementary Figure 24, Supplementary Table 13 and Supplementary Table 14). Here, we highlight a series of examples. First, the type 2 diabetes-associated SNP rs1063192 in the *CDKN2B* locus has a higher odds ratio in the top quartile of CVD topic weight (1.19±0.02) than in the bottom quartile (1.08±0.02) (P=4×10^{-4} for difference). *CDKN2B* is associated with both coronary artery disease and type 2 diabetes[55–59], suggesting that shared pathways underlie the observed SNP x topic interaction. Second, the asthma-associated SNP rs1837253 in the *TSLP* locus has a higher odds ratio in the top quartile of SRD topic weight (1.17±0.02) than in the bottom quartile (1.05±0.02) (P=1×10^{-4} for difference). *TSLP* plays an important role in promoting Th2 cellular responses and is considered a potential therapeutic target, which is consistent with assignment of asthma and atopic/contact dermatitis[60] to the SRD topic (Supplementary Table 4). Third, the hypertension-associated SNP rs3735533 within the *HOTTIP* long non-coding RNA has a lower odds ratio in the top quartile of CVD topic weight (1.07±0.02) than in the bottom quartile (1.13±0.02). *HOTTIP* is associated with blood pressure[27,61] and conotruncal heart malformations[62]. Fourth, the hypothyroidism-associated SNP rs9404989 in the *HCG26* long non-coding RNA has a higher odds ratio in the top quartile of FGND topic weight (1.90±0.24) than in the bottom quartile (1.19±0.13) (P=3×10^{-3} for difference). Hypothyroidism associations have been reported in the HLA region[27], but not to our knowledge in relation to the *HCG26*. To verify correct calibration, we performed control SNP x topic interaction tests using the same 888 disease-associated SNPs together with random topics that did not correspond to disease subtypes, and confirmed that these control tests were well-calibrated (Supplementary Figure 24B).

We conclude that genetic heterogeneity between disease subtypes can be detected at the level of individual disease-associated variants.

Discussion

We have introduced an age-dependent topic modelling (ATM) method to provide a low-rank representation of longitudinal disease records, leveraging age-dependent comorbidity profiles to identify and validate biological subtypes of disease. Our study builds on previous studies on topic modelling[37,38,40,63], genetic subtype identification[13–15], and low-rank modelling of multiple diseases to identify shared genetic components[25–27]. We highlight three specific contributions of our study. First, we incorporated age at diagnosis information into our low-rank representation, complementing the use of age information in other contexts[32,52,64]; we showed that age information is highly informative for our inferred comorbidity profiles in both simulated and empirical data, emphasising the importance of accounting for age in efforts to classify disease diagnoses. Second, we identified 52 diseases with heterogeneous comorbidity profiles that we used to define disease subtypes, many of which had not previously been identified (Supplementary Table 15). Third, we used genetic data (including PRS, genetic correlation and F_{ST} analyses) to validate these disease subtypes, confirming that the inferred subtypes reflect true differences in disease aetiology.

We emphasise three downstream implications of our findings. First, it is of interest to perform disease subtype-specific GWAS on the disease subtypes that we have identified here, analogous to GWAS of previously identified disease subtypes[13–15]. Second, our findings motivate efforts to understand the functional biology underlying the disease subtypes that we identified; the recent availability of functional data that is linked to EHR is likely to aid this endeavor[29,65]. Third, it is of interest to apply ATM to identify age-dependent comorbidity profiles and disease subtypes in other EHR data sets[30,31]; establishing representations of disease topics that are transferable and robust across different healthcare systems and data sources represents a major future challenge.

Our findings reflect a growing understanding of the importance of context, such as age, sex, socioeconomic status and previous medical history, in genetic risk [52,66,67]. To maximise power and ensure accurate calibration, context information needs to be integrated into clinical risk prediction tools that combine genetic information (such as polygenic risk scores [1,68]) and non-genetic risk factors. Our work focuses on age, but motivates further investigation of other contexts. We note that aspects of context are themselves influenced by genetic risk factors, hence there is an open and important challenge in determining how best to combine medical history and/or causal biomarker measurements with genetic risk to predict future events[69].

We note several limitations of our work. First, age at diagnosis information in EHR data may be an imperfect proxy for true age at onset, particularly for less severe diseases that may be detected as secondary diagnoses; although perfectly accurate age at onset information would be ideal, our study shows that that imperfect age at diagnosis information is sufficient to draw meaningful conclusions. Second, raw EHR data may be inaccurate and/or difficult to parse[1]; again, although perfectly accurate EHR data would be ideal, our study shows that imperfect EHR data is

sufficient to draw meaningful conclusions. Third, our ATM approach incurs substantial computational cost (Supplementary Table 16); however, analyses of biobank-scale data sets are computationally tractable, with our main analysis requiring only 4.7 hours of running time. Finally, we have applied ATM to a UK population of predominantly European ancestry; it is of interest to apply ATM to diverse populations[30,31]. Despite these limitations, ATM is a powerful approach for identifying age-dependent comorbidity profiles and disease subtypes.

Acknowledgements

This research has been conducted using the UK Biobank Resource; application number 12788.

Funded by Wellcome (BST00080- H503.01 to XJ, 100956/Z/13/Z to GM, https://wellcome.org); the Li Ka Shing Foundation (to GM, https://www.lksf.org); NIH grants R01 HG006399, R01 MH101244, and R37 MH107649 (to ALP); The Alan Turing Institute (https://www.turing.ac.uk), Health Data Research UK (https://www.hdruk.ac.uk), the Medical Research Council UK (https://mrc.ukri.org), the Engineering and Physical Sciences Research Council (EPSRChttps://epsrc.ukri.org) through the Bayes4Health programme Grant EP/R018561/1, and AI for Science and Government UK Research and Innovation (UKRI, https://www.turing.ac.uk/ research/asg) (to CH); BHF Chair award CH/12/2/29428 (to XJ). This work was supported by core funding from the: British Heart Foundation (RG/13/13/30194; RG/18/13/33946), BHF Cambridge Centre of Research Excellence (RE/13/6/30180) and NIHR Cambridge Biomedical Research Centre (BRC-1215-20014). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

This work uses data provided by patients and collected by the NHS as part of their care and Support. Computation used the Oxford Biomedical Research Computing (BMRC) facility, a joint development between the Wellcome Centre for Human Genetics and the Big Data Institute supported by Health Data Research UK and the NIHR Oxford Biomedical Research Centre. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health. We thank Kushal Dey, Luke Kelly and Yunlong Jiao for the discussion.

Data availability

UK Biobank data is publicly available at https://www.ukbiobank.ac.uk/.

Code availability

Open-source software implementing the ATM method is available at https://github.com/Xilin-Jiang/ATM. BOLT-LMM is available at https://alkesgroup.broadinstitute.org/BOLT-LMM/. Heritability and genetic correlation analysis were performed using LDSC, which is available at

455 https://github.com/bulik/ldsc. PLINK v1.9, which was used for F_{ST} and association tests, is
456 available at https://www.cog-genomics.org/plink/.

457

458

Methods

Age-dependent topic model (ATM)

Our Age-dependent topic model (ATM) is a Bayesian hierarchical model to infer latent risk profiles for common diseases. The model assumes that each individual possesses several age-evolving disease profiles (topic loadings), which summarise the risk over age for multiple diseases that tend to co-occur within an individual's lifetime, namely the age specific multi-morbidity profiles. At each disease diagnosis, one of the disease profiles is first chosen based on individual weights of profile composition (topic weights), the disease is then sampled from this profile conditional on the age of the incidence.

We constructed a Bayesian hierarchical model to infer *K* latent risk profiles for *D* distinct common diseases. Each latent risk profile (comorbidity topics) is age-evolving and contains risk trajectories for all *D* diseases considered. Each individual might have a different number of diseases, while the disease risk is determined by the weighted combination of latent risk topics. The indices in this note are as follows:

- $s = 1, \ldots, M$;
- $n = 1, \ldots, N_s$;
- $i = 1, \ldots, K$;
- $j = 1, \ldots, D$;

where *M* is the number of subjects, N_s is the number of records within s^{th} subject, *K* is the number of topics, and *D* is the total number of diseases we are interested in. The plate notation of the generative model is summarised in Supplementary Figure 1:

- $\theta \in R^{M \times K}$ is the topic weight for all individuals (referred to as patient topic weights), each row of which ($\in R^K$) is assumed to be sampled from a Dirichlet distribution with parameter α. α is set as a hyper parameter: $\theta_s \sim Dir(\alpha)$.
- $z \in \{1,2,\ldots,K\}^{\sum_S N_S}$ (referred to as diagnosis-specific topic probability) is the topic assignment for each diagnosis $w \in \{1,2,\ldots,D\}^{\sum_S N_S}$. Note the total number of diagnoses across all patients are $\sum_S N_S$. The topic assignment for each diagnosis is generated from a categorical distribution with parameters equal to s^{th} individual topic weight: $z_{sn} \sim Multi(\theta_s)$.
- $\beta(t) \in F(t)^{K \times D}$ is the topic loading which is $K \times D$ functions of age t. $F(t)$ is the class of functions of t. At each plausible t, the following is satisfied: $\sum_j \beta_{ij}(t) = 1$. In practice we use softmax function to ensure above is true and add smoothness by constrain $F(t)$ to be spline or polynomial functions: $\beta_{ij}(t) = \frac{exp(p_{ij}{}^T \phi(t))}{\{\sum_{j=1}^{D} exp(p_{ij}{}^T \phi(t))}$, where $p_{ij} = \{p_{ijd}\}$; $d = 1,2,\ldots,P$; P is the degree of freedom that controls the smoothness; $\phi(t)$ is polynomial and spline basis for age t.

- $w \in \{1,2,\ldots,D\}^{\sum_s N_s}$ are observed diagnoses. The nth diagnosis of sth individual w_{sn} is sampled from the topic $\beta_{z_{sn}}(t)$ chosen by z_{sn}: $w_{sn} \sim Multi(\beta_{z_{sn}}(t_{sn}))$, here t_{sn} is the age of the observed age at diagnosis of the observed diagnosis w_{sn}.

The values of interest in this model are global topic parameter β, individual (patient) level topic weight θ, and diagnosis-specific topic probability z. Based on the generative process above, we notice that each patient is independent conditional on α. Therefore, the inference of θ and z(discussed below) could be performed by looping each individual in turn.

The key element in our model is age-evolving risk profiles, which is achieved by model the comorbidity trajectories $\beta(t) \in F(t)^{K \times D}$ as functions of age. The functionals $F(t)$ considered are linear, quadratic, cubic polynomials, and cubic splines with one, two and three knots.

Inference of ATM

The variables of interest are global topic parameter $\beta(t)$, individual (patient) level topic weight θ, and diagnosis-specific topic probability z of each diagnosis. We could adopt an EM strategy, where in the E-step we first estimate posterior distribution of θ and z, then in the M-step we estimate β which maximises the evidence lower bound (ELBO).

The details of the inference is explained in Supplementary Note. In summary, in a Bayesian setting, the model could be evaluated by the evidence function $p(w|\alpha,\beta)$. The best $\beta(t)$ is found by maximise the evidence function, while for θ and z we aim to find or approximate their posterior distribution $p(z,\theta \mid w,\alpha,\beta)$. Given that the posterior distribution is intractable, we use variational distribution $q(z,\theta)$ to approximate them. Now we could write the evidence function as:

$$p(w|\,\alpha,\beta) = L(z,\theta,\beta,\alpha) + KL(q||p),$$

here $KL(q||p) = -\int_{z,\theta} q(z,\theta)\ ln\frac{p(z,\theta \mid w,\alpha,\beta)}{q(z,\theta)}$ is the KL divergence. Since KL divergence is always positive, $L(z,\theta,\beta,\alpha)$ is a lower bound of the evidence function:

$$L(z,\theta,\beta,\alpha) = E_q\{\, ln\, p(w,z,\theta \mid \alpha,\beta) - ln\, q(z,\theta)\}.$$

When finding the posterior of θ and z, we want $ln\, q(z,\theta)$ to be as close to the posterior $p(z,\theta \mid w,\alpha,\beta)$ as possible. Since $KL(q||p) = 0$ when $q(z,\theta) = p(z,\theta \mid w,\alpha,\beta)$, this could be achieved by minimising $KL(q||p)$ or maximise $L(z,\theta,\beta,\alpha)$. The most commonly used form of $q(z,\theta)$ assume the distribution is factorised, which might cause instability when signal-to-noise ratio is low[70]. Therefore, more accurate inference methods such as collapsed variational inference is considered[40]. Comparison of the evidence lower bound $L(z,\theta,\beta,\alpha)$ shows collapsed variational inference is consistently more accurate than LDA (Supplementary Figure 8). Therefore we choose the collapsed variational inference[40]. The collapsed variational inference is achieved by integrate out θ from the likelihood function $p(w,z,\theta \mid \alpha,\beta)$ and find the

approximated posterior distribution $q(z)$. For detailed derivation, the comparison between collapsed variational inference and mean-field variational inference, and update algorithms, see Supplementary Note.

When finding the $\beta(t)$ that maximises the evidence function, we again maximise $L(z, \theta, \beta, \alpha)$. Maximising $L(z, \theta, \beta, \alpha)$ with respect to $\beta(t)$ does not have an analytical solution due to its softmax structure. We use local variational methods and numeric optimisation to find the distribution of $\beta(t)$. Details are provided in Supplementary Note.

We extract topic weights at patient-level and diagnosis-level from the posterior distribution inferred from the data. Our model has the desired property that each patient and patient-diagnosis are assigned to comorbidity topics. The model estimates the posterior distribution $q(z)$, which is a categorical distribution (Supplementary Note.) Several metrics related to topic assignments could be derived from the $q(z)$:

- Each patient-diagnosis (incident disease) has a diagnosis-specific topic probability, which is computed as $E_q\{z_n\}$.
- Each patient has a posterior topic assignment θ_s , which is a dirichlet distribution $\theta_s \sim Dir(\alpha + \sum_{n=1}^{N_S} E_q\{z_n\})$. The topic weights of each patient is the mode of this Dirichlet distribution $\frac{\sum_{n=1}^{N_S} E_q\{z_n\}}{\sum_{i=1}^{K} \sum_{n=1}^{N_S} E_q\{z_{ni}\}}$ (we used $\alpha = 1$). The value is used as the patient low-rank representation of disease history, for analysis including PRS association with comorbidity within cases and G x Topic interaction analysis.
- The average topic assignments of disease j is the mean over all incidences $\overline{E_q\{ z_{sn \in \{w_{sn}=J\}}\}}$. This metric is used to measure which comorbidity topic a disease is associated with (Figure 4B), and it is equivalent to a weighted average of topic loadings (for the specific weighted average expression, see equation 5 of Supplementary Note). A disease assigned to multiple topics is considered to have comorbidity subtypes.
- A hard assignment of a patient-diagnosis to a subtype is based on the max value of the vector $E_q\{z_n\}$. The incident disease is assigned to topic $argmax_i\ (E_q\{z_{ni}\})$.

Metrics for evaluating ATM

ATM is evaluated for different purposes, which requires different metrics (Supplementary Table 1). Here we list the details of the four metrics considered: *Prediction odds ratio, Evidence Lower Bound (ELBO), AURPC, and Co-occurrence odds ratio.*

Prediction odds ratio: To compare models of different topic numbers and configuration of age profiles, we compare the prediction odds ratio of each model. Briefly, prediction odds ratio is defined on 20% held-out test data as the odds that the true diseases are within the top 1% diseases predicted by ATM (trained on 80% of the training set and uses earlier diagnoses as

input), divided by the odds that the true diseases are within the top 1% of diseases ranked by prevalence.

Specifically, we separate UK Biobank patients into a training set (80%) and a testing set (20%). On the training set, we estimate the comorbidity topic loadings. On the testing set, we fix the topic loadings and infer the patient topic weights to predict the next disease in chronological order. The topic loadings are estimated using the n diseases and compute the risk rank of diseases at the age of the n+1 disease. The odds ratio is computed by the odds of the n+1 disease being in the top 1% of diseases versus being in the top 1% most prevalent diseases. We use the top 1% most prevalent diseases instead of randomly chosen diseases as it represents a naive prediction model that predicts disease based on prevalence. The patient topic weights computation is in section Inference of ATM and the risk is computed as the linear combination of topics using topic weights as coefficients. We also compute the prediction odds ratio using the LDA model. We repeat the procedure for 10 times for each model configuration.

We compared the prediction odds ratio for topic number between 5 to 20, with linear, quadratic polynomial, cubic polynomial, and splines with one, two and three knots. We also compare the ATM model with the LDA model of topic number between 5 to 20.

Evidence Lower Bound (ELBO): ELBO evaluated the accuracy of the variational inference method on a specific data set. Mathematical expression of ELBO for ATM is presented in equation 9 in Supplementary Note. To find the best model that fit to the entire dataset, we evaluate the ELBO for models with topic numbers between 5 to 20, 25, 30, and 50 topics and age profiles configured by linear, quadratic polynomial, cubic polynomial, and splines with one, two and three knots. Each model is run for 10 times with random initialisations. We choose the model that has the highest ELBO after converging.

AURPC: To evaluate whether a model could capture the comorbidity subtypes in simulation analysis, we compute the precision, recall, and area under precision-recall curve (AUPRC) to correctly classify disease diagnosis to be from the topic that it is generated from. The topic of each diagnosis is determined by diagnosis-specific topic probability. Note we could only evaluate AUPRC in simulations where the truth is known.

Co-occurrence odds ratio: To verify that the comorbidity profiles that the model captured are capturing diseases that are more likely to present within the same individual, we estimate the odds ratio of the disease duo, trio, quartet, and quintet that are captured by the topic versus that of random combinations. We divide the population into an 80% training set and a 20% testing set. We trained the ATM model with five random initialisations and kept the inference with the highest ELBO. Each disease is assigned to a topic by the highest average topic assignments. (section Inference of ATM) We focus on the top 100 diseases ranked by prevalence to avoid the

combination being too rare to appear in the population. In the testing set, we computed the odds of individuals who have all diseases in the comorbidities versus the odds implied if all diseases are independent (computed as the product of disease prevalence). The odds ratio is computed for all combinations of duo, trio, quartet, and quintet that are assigned to the same topics. We perform the same analysis using PCA for comparison.

Simulations of ATM method.

To test whether the algorithm could assign disease diagnosis to correct comorbidity profiles, we simulated disease from two disease topics within a population of 10,000, using following parameters:

- $M = 10{,}000$;
- $\overline{N_S} = 6.1$;
- $N_S \sim exp\{\overline{N_S}\}$;
- $D = 20$;
- $K = 2$;

Here M is the number of individuals in the population, $\overline{N_S}$ is the average number of diseases for each individual, D is the total number of diseases, K is the number of comorbidity topics. The distribution of disease number per-individual N_S is sampled from an exponential distribution, which matches those from UK Biobank data (Supplementary Figure 26). According to equation 3.1 in Ghorbani et al.[70], whether the topic model could capture the true latent structure is determined by the information signal-to-noise ratio and could be evaluated with limits $M \rightarrow \infty$; $D \rightarrow \infty$; $\frac{D}{M} \rightarrow \delta$, where δ is a constant. Therefore we choose D and M at scales that make $\frac{D}{M}$ approximately similar to those of the UK Biobank dataset (Samples size = 282,957; distinct disease number = 349).

The simulated topics loadings are constructed as follows:

- All but K diseases are simulated to be associated with comorbidity profiles. Each of them has a risk period of 30 years and overlaps for 10 years with the next disease. For example, if disease 1 has a risk period from 30 to 59 years of age, disease 2 will have a risk period between 50 to 79 years of age. When the risk period reaches the maximal age, the truncated part will be carried to the next disease to create diseases with shorter risk period. All risk periods are assigned a value 1.
- K diseases that are not associated with comorbidity are simulated to span all topics. The values of these diseases are sampled from $Unif(0, \frac{0.1}{K})$ for each topic. Here K is the number of topics.
- The age profiles are then normalised at each age point to ensure $\sum_{j=1}^{D} \beta_j(t) = 1$ for all t. With this constraint we could sample a disease at each age t using a multinomial probability with the topic loading as the parameter. The age range of the simulated topics

is 30 to 81 years of age, which is the minimal and maximal age at diagnosis of incident disease in the UK Biobank population. An example of a simulated topic is shown in Supplementary Figure 27.

For each individual, we sampled the Dirichlet parameter α from a gamma distribution (shape = 50, rate = 50). Topic loadings are sampled from the Dirichlet distribution for each patient as the generative process. For each patient, we first sample the number of diseases N_S. For each incident disease, we sample the disease age from uniform distribution between age 30 to 81 and a topic from the topic loading. We then choose the incident disease based on the age at diagnosis from the chosen topic. The procedure follows the generative process described in Supplementary Note.

Since in real data we only use the first age at diagnosis for diseases that are recorded repeatedly within the same patient, we filter the simulated diseases accordingly. The filtered data are fed into the inference functions to infer the latent topics and disease assignments. The inferred topics resemble the true topics used to simulate diseases as shown in Supplementary Figure 27. For the initialisation of each inference, we first sample β and θ from the Dirichlet distribution of non-informative hyperparameters, then initialise other variables parameters following the generative process. The variational inference converged where the relative increase of ELBO is below 10^{-6}.

To simulate disease having distinct comorbidity subtypes, we first simulate diseases using the procedure described above. We consider two scenarios: (1) the subtype of diseases have the same age at diagnosis distribution. (2) the subtypes of disease have distinct age at diagnosis distribution.

We create diseases with distinct comorbidity profiles by combining diseases that are sampled from distinct topics and labelling them as a single disease. We first chose one disease (**disease A**) then sampled a proportion of a second disease (**disease B**) to label as **disease A**. The proportion is varied to create a different sample size ratio of the two subtypes. In scenario one, **disease B** is a disease that has the exact same age distribution as **disease A** but from the other topic. In scenario two, **disease B** is from the other topic and has a different age distribution (age at diagnosis moves up for 20 years, 10 years, or 5 years, respectively) than **disease A**. After changing the labels of **disease B** to be the same as **disease A**, we used the inference procedure described as above to get the posterior distribution.

To evaluate whether a model could capture the comorbidity subtypes, we compute the precision, recall, and area under precision-recall curve (AUPRC) to correctly classify incident **disease B** to be from the topic that it is generated from. The topic of each diagnosis is determined by diagnosis-specific topic probability. We use other diseases from the topic of **disease B** to benchmark the topic label. Topic modelling on the simulated data is performed with both ATM

and LDA (both implemented using collapsed variational inference for fair comparison) to compare the performances.

We evaluate the subtype classification with varying values for three simulation parameters:

- ratio of sample sizes between the two subtypes. We change the ratio of the two subtypes by a grid between 0 to 0.9 with a step size 0.1. The default value of sample size ratio is set as 0.1 in other simulations to test for other parameters that have impacts on the precision and recall.
- Simulated population size. We simulated population sizes equal to 200, 500, 1000, 2000, 5000, and 10,000. The default population size is 10,000 in other simulations.
- Number of distinct diseases. We simulated datasets with 20, 30, 40, and 50 distinct diseases, with 2, 3, 4 and 5 underlying disease topics respectively. The default number of distinct diseases is 20 in other simulations.
- Difference of age distribution. We considered three scenarios of subtype age distribution, with 0, 10, and 20 years of difference in the average age at diagnosis.

UK Biobank comorbidity data.

We analysed comorbidity data from 282,957 UK Biobank samples with diagnoses for at least two of the 348 focal diseases that we studied (see below). We use the hospital episode statistics (HES) data within the UK Biobank dataset, which records diseases using the ICD-10/ICD-10CM coding system. Codes started with letters from A to N are kept as they correspond to disease code (opposed to procedure codes). The disease records were mapped from ICD-10/ICD-10CM codes to PheCodes using a three-step procedure: Firstly, we map the first four letters of each ICD-10 records to the phecodes, using the map file downloaded from phewascatalog.org; Secondly, we map the remaining records using ICD-10CM map file downloaded from phewascatalog.org; Lastly, we map remaining records to a collapsed ICD-10CM mapping system which only use the first four character of ICD-10CM codes. We also noticed an ICD-10/ICD-10CM code could map to multiple PheCodes. When a single ICD-10/ICD-10CM code s mapped to more than one PheCodes, we only kept the Phecode that are mapped to the most ICD-10 codes (i.e. PheCode is constructed by combining ICD-10 that represent similar diseases. The Phecode that represent a larger number of ICD-10 codes are more likely to be a well defined disease, which we chose to keep.), which ensure that one ICD-10(CM) code only maps to one PheCode. Using the procedure above, we mapped 99.7% ICD-10/ICD-10CM code to PheCodes, with 4,637,127 records in total.

The mapped Phecodes are filtered to keep only the first age at diagnosis for the same diseases within a patient. The age at diagnosis for each record is computed as the difference between month of birth to the episode starting date. We then computed the occurrence of each disease in the UK Biobank and kept 348 that have more than 1,000 occurrences (Supplementary Table 4).

Starting with all 488,377 UK Biobank patients (including both European and non-European ancestries), we filtered the patients to keep only those who have at least two distinct diseases from the 348 focal diseases, as we are most interested in the comorbidity information. We treated the death as an additional disease (8,666 records) to evaluate if certain comorbidities are more likely to lead to fatal events. After these procedures, there are in total 1,726,144 distinct records across 282,957 patients.

To name the topics inferred from the UK Biobank, we take the sum of average topic assignments (section Inference of ATM) over diseases that are within each phecode system and extract the top 3 systems. Most of comorbidity topics are named using the first three topics (e.g. CER: cardiovascular, endocrine/metabolic, respiratory), except for topics that are predominantly associated with one system (LGI: lower gastrointestinal; UGI: upper gastrointestinal; CVD: cardiovascular).

We present focal diseases for each topic in two ways. Firstly, we filter each topic using the profile mean value between age 30 to 81 to keep the top seven diseases. We chose seven for visualisation, as we found more diseases would be harder to read on a plot. Secondly, we also show seven diseases that have the highest average assignment to each topic. This will give a picture of diseases that are not the most prevalent in the population but are predominantly associated with the target topic.

To compare the comorbidity heterogeneity between age groups, we group the incidences for each disease to two age groups: young group (<60 years of age) and old group (≥60 years of age). We compute the average topic assignment of each group as described in section Inference of ATM. Additionally, we inferred topics for male (984,554 records in 156,366 individuals) and female (741,590 records in 126,591 individuals) populations respectively using a model with 10 topics and spline function with one knot. We extract the average topic assignment for each disease, and use Pearson's correlation to match the topics for both sexes to the topics inferred on the entire population.

Each diagnosis could be assigned to a specific topic using max diagnosis-specific topic probability. We focus our disease heterogeneity analysis on 52 diseases that have at least 500 incidences assigned to a secondary topic.

UK Biobank genotype data.

For all analyses except BOLT-LMM we use 488,377 UK Biobank participants. For BOLT-LMM analyses, we constrain our analysis to 409,694 British Isle ancestry individuals to remove the possibility that topics are capturing population structure. For F_{ST} analysis with PLINK we used 805,426 genotyped SNPs; for BOLT-LMM PRS analysis we used 727,882 genotyped SNP with MAF>0.1%; for genetic correlation analysis using LDSC, we used 157,756 Genotyped SNPs

mapped to HapMap3 SNPs; for BOLT-LMM and subsequent LDSC analysis that use imputed SNPs, we used 1,201,838 imputed SNPs mapped to HapMap3 SNPs SNPs.

Polygenic risk scores (PRS) analysis.

To exclude the possibility of population stratification, we compute PRS using mixed-effect association on the British Isle ancestry group (N = 409,694) for the 10 heritable diseases that have the highest heritability z-scores. We used a mixed model to estimate effect size implemented by BOLT-LMM and constructed genome-wide PRS [50]. For the computation of PRS, we randomly sampled half of the British isle ancestry population (N = 204,847) for computation efficiency (essential hypertension, arthropathy, asthma, and hypercholesterolemia) or sampled 9 controls for each case to ensure case proportion at or above 10% as recommended by BOLT-LMM (type 2 diabetes, varicose veins of lower extremity, hypothyroidism, other peripheral nerve disorders, major depressive disorder, and GRED). We used PLINK to select genotyped SNPs with MAF > 0.1% as recommended in BOLT-LMM. For each disease, we used 5-fold cross validation to estimate effect sizes using BOLT-LMM and computed the PRS on the held-out testing set. The predictive PRS are then used to compute the excess PRS over different topic loadings, by a linear regression where PRS is the response variable and topic weights is the predictor.

We compute the relative risk for each percentile of PRS using the following formula:

$$RR_{pt,s} = \frac{n_{pt,s} \times 100}{n_S},$$

where $RR_{pt,s}$ is the relative risk of s subtype for the pt^{th} PRS percentile (computed for the entire population); $n_{pt,s}$ is the number of cases in s subtype that has PRS within the pt^{th} percentile; n_S is the number of cases in the s subtype.

Genetic correlation analysis.

For each disease and disease subtype, we use a case-control matching strategy to construct data to estimate coefficients for genetic correlation analysis. For each case in the disease group, we pick four nearest neighbors (without replacement) from the control group, matching sex, BMI, year of birth and 40 genetic principal components. The covariates are available within the UK Biobank data set, over which we computed the principal components. We then compute the Euclidean distance of the principal components to find the nearest neighbours in the population. All cases are matched with four controls except for 401.1 essential hypertension which has a sample size larger than 20% of the population. We match only one control for each hypertension case.

We perform logistic regression with sex and top 10 principal components as covariates to estimate the main variant effect of the 805,426 variants that are genotyped. We used PLINK 1.9 for association analysis[71]. With the summary statistics from the association analysis, we use

LDSC to map the summary statistics to HapMap3 SNPs and match the effect and non-effect alleles[2,72]. Since UK Biobank is mostly of British Isle ancestry, we use the pre-computed LD score from the LDSC website. We estimated the heritability for each disease or disease subtype which has more than 1000 incidences (378 diseases subtypes and diseases). We use 1000 incidence threshold as LDSC are more accurate with larger sample size. We focus on 71 disease and 18 disease subtypes that have heritability z-score above 4 for genetic correlation analysis.

The genetic correlation is computed for each pair of disease/subtypes using the same summary statistics and LD score regression. We report the estimate of genetic correlation and z-scores. Additionally, for pairs that involve subtypes (disease-subtype or subtype-subtype), we compute the excess genetic correlation, defined as the difference between the genetic correlation involving subtypes and the genetic correlation involving all disease diagnoses. For example, the genetic correlation between T2D-CER and hypertension-CVD is compared to the genetic correlation between all T2D and all hypertension. The z-score and p-value of the genetic correlation differences are reported. We note that genetic correlations between subtypes of the same disease are compared to 1. We only reported p-values of excess genetic correlation when both genetic correlation estimation has standard error <0.1 and at least one of the genetic correlation has |z-score|>4.

To avoid potential collider effects where subtypes are defined by topic components that are independent of the diseases, we further match cases in each subtype with controls that match the topic loadings. We computed PCs from 23 variables (10 topic loadings, 10 PCs, year of birth, sex, and BMI) and use the nearest neighbour procedure (by Euclidean Distance) to find controls for each case. Here controls are chosen from individuals without the targeting disease, i.e. an individual with one subtype of the target disease could not be a control for the other subtypes. We performed the same analysis using this case-control matching procedure and compared the genetic correlation with the case-control procedure described above. We perform the analysis for four diseases that have evidence for genetic subtypes: asthma, type 2 diabetes, hypercholesterolemia, and hypertension. For one subtype (hypertension-CVD), the heritability (0.0313, s.e. = 0.0289) is below threshold after matching the topic, which was excluded in genetic correlation analysis.

F_{ST} analysis.

To evaluate the genetic heterogeneity between disease subtypes, we estimated the F_{ST} for 52 diseases that have at least 500 incidences assigned to a secondary topic. To test the statistical significance of Fst, we adopted a permutation strategy and sampled the same number of controls of similar topic weights distribution for each subtype. The topic weights are matched by sampling (without replacement) the same number of controls for each dominant topic weight quartile of the cases (i.e. matching the topic that defines the subtype), which ensures the controls have the same topic weight stratification as the disease subtypes. We then compute the F_{ST} across

the control groups matched for subtypes. We excluded three diseases, “hypertension”, “hypercholesterolemia”, and “arthropathy”, from F_{ST} analysis as we do not have enough controls that match topic weight distribution. The F_{ST}s are computed using PLINK 1.9's weighted mean across all genotyped SNPs, which report F statistics across all subtypes.

We obtained 1,000 permutation samples and reported the permutation p-value. Under the assumption that causal and non-causal variants have similar allele frequency differences across the subtypes, F_{ST} could be a measure of causal genetic effect heterogeneity across subtypes.

SNP x topic interaction test.

For the diseases that have heritability z-score above 4 in the UK Biobank, we further investigated whether there are interactions between genetic risk factors with the topic loadings. We used a fit a logistic regression model using following model:

$$logit(p) = \beta_0 + \beta_1 * T + \beta_2 * T^2 + \beta_3 * G + \beta_4 * G * T,$$

where T is individual topic weights for a specified topic, G is the genotype, and p is the probability of getting the disease. We computed the test statistics under the null that $\beta_4 = 0$. We used QQ plots to check that the test statistics are well calibrated for each disease-topic pair.

Since the simulation shows the interaction test is underpowered when the variant effects are small, we focus on the set of SNP that reaches genome-wide significance level to increase power to detect interaction effects. We performed LD-clumping using $r^2 > 0.6$ to remove variants that are in strong LD with the lead variants. We computed the test statistics using the model above (for testing $\beta_4 = 0$) and computed study-wise FDR across disease-topic pairs.

To verify the significant interactions, we divided cases into quartiles based on topic loading for each disease-topic pair, and randomly sampled two controls that match the topic loading for each case. We estimated the main effect sizes for all GWAS-SNP within each quartile of topic loadings to capture effects that are modulated by topic weights. We focus on the SNPs that have significant interaction test statistics computed in the previous step and compare it with background SNPs that have genome-wide significant main effects but no interaction effect (P>0.05).

Simulations of SNP x topic interaction

We simulate comorbidity with genetics to test interaction between genetic and comorbidity topics. We simulated 100 independent variants with MAF randomly sampled from $Unif(0, 0.5)$. We assumed an additive model and simulated genotypes for the population using Hardy-Weinberg equilibrium. We simulated three types of genetic effects on topic and diseases on topic of the simulation framework described in Simulations of ATM method section:

- Genetics-topic effect: each variant is simulated to have an linear effect of 0.04 on the topic loading. We choose this value as after normalising the topic, a regression of causal

variant to topic would have an effect size approximately 0.01 which is similar to our observation in the UK Biobank. The number of variants that are causal to the topic varies between 2 to 20. We simulated the effect on one topic by adding additive SNP effects and normalise the topic loadings of each patient. The topic-disease causality is a natural consequence following the generative process of sampling data.

- Genetic-disease-topic effect: we simulated a heritable disease that is causal to the topic. The disease is simulated with 20 causal variants each of effect size 0.15. We vary the disease-to-topic causal effect from 0.05 to 0.5, with a default value of 0.1 in other analyses (similar to the correlation we found in UK Biobank analysis). We simulated the effect on one topic by adding additive causal disease effects and normalise the topic loadings of each patient.
- The genetic effect could interact with the topic when contributing to disease risk. We simulated four additional diseases to represent different structures (Supplementary Figure 22).
    - Genetic effects interact with topic loading on altering disease risk. The interaction term is added to the mean of disease liability, which is sampled from a Gaussian distribution. The disease is then sampled by a threshold on the liability, where the incidence rate is by default 0.5. The interaction effect is varied from 0.4 to 4, with default value equal to 2.
    - Pleiotropy effects are simulated with a variant that have both genetic-disease and genetic-topic-disease effects. Both genetic and topic effects are added to the mean of disease liability. A disease is sampled by a threshold with default incidence rate equal to 0.5. The topic-disease effect is varied from 0.4 to 4, with default value equal to 2.
    - Pleiotropy effect with nonlinear topic-disease effect. A quadratic term of topic-disease effect added to the second model.
    - Pleiotropy effect with nonlinear genetic-disease effect. A quadratic term of genetic-disease effect added to the second model.

For disease-topic or topic-disease causal effects, we simulated 50 repetition at each causal effect size. For interaction analysis, we repeated 10 times at each parameter value, as there are more SNPs for uncertainty estimation. The simulated disease sets are fed into the inference procedure to infer the patient topic weights.

927 # Tables

928

Acronym	Disease systems	Representative diseases	Number of associated diseases
NRI	Neoplasms, respiratory, infectious diseases	Secondary malignancy of lymph nodes; Pneumococcal pneumonia; Bacterial infection NOS	53
CER	Circulatory system, endocrine/metabolic, respiratory	Type 2 diabetes; Obesity; Chronic airway obstruction	41
SRD	Sense organs, respiratory, dermatologic	Cataract; Septal Deviations/Turbinate Hypertrophy; Benign neoplasm of skin	38
CVD	Cardiovascular disease	Hypercholesterolemia; Coronary atherosclerosis; Myocardial infarction	27
UGI	Upper gastrointestinal disease	Diaphragmatic hernia; Benign neoplasm of other parts of digestive system; Gastritis and duodenitis;	22
LGI	Lower gastrointestinal disease	Irritable Bowel Syndrome; Benign neoplasm of colon; Anal and rectal polyp;	13
FGND	Female genitourinary, neoplasms, digestive	Uterine leiomyoma; Malignant neoplasm of female breast; Hypothyroidism NOS	34
MGND	Male genitourinary, neoplasms, digestive	Urinary tract infection; Cancer of prostate; Other disorders of bladder	33
MDS	Musculoskeletal, digestive, symptoms	Back pain; Cholelithiasis; Other disorders of soft tissues	29
ARP	Arthropathy-related disease	Arthropathy NOS; Rheumatoid arthritis; Enthesopathy	26

929

930
931
932 **Table 1. Summary of 10 inferred disease topics in the UK Biobank.** For each topic, we list its
933 3-letter acronym, disease systems, representative diseases, and number of associated diseases
934 (defined as diseases with average diagnosis-specific topic probability >50% for that topic).
935 Topics are ordered by the Phecode system (see Figure 3). 316 of 348 diseases analysed are
936 associated with a topic; the remaining 32 diseases do not have a topic with average diagnosis-
937 specific topic probability >50%.
938
939

Figures

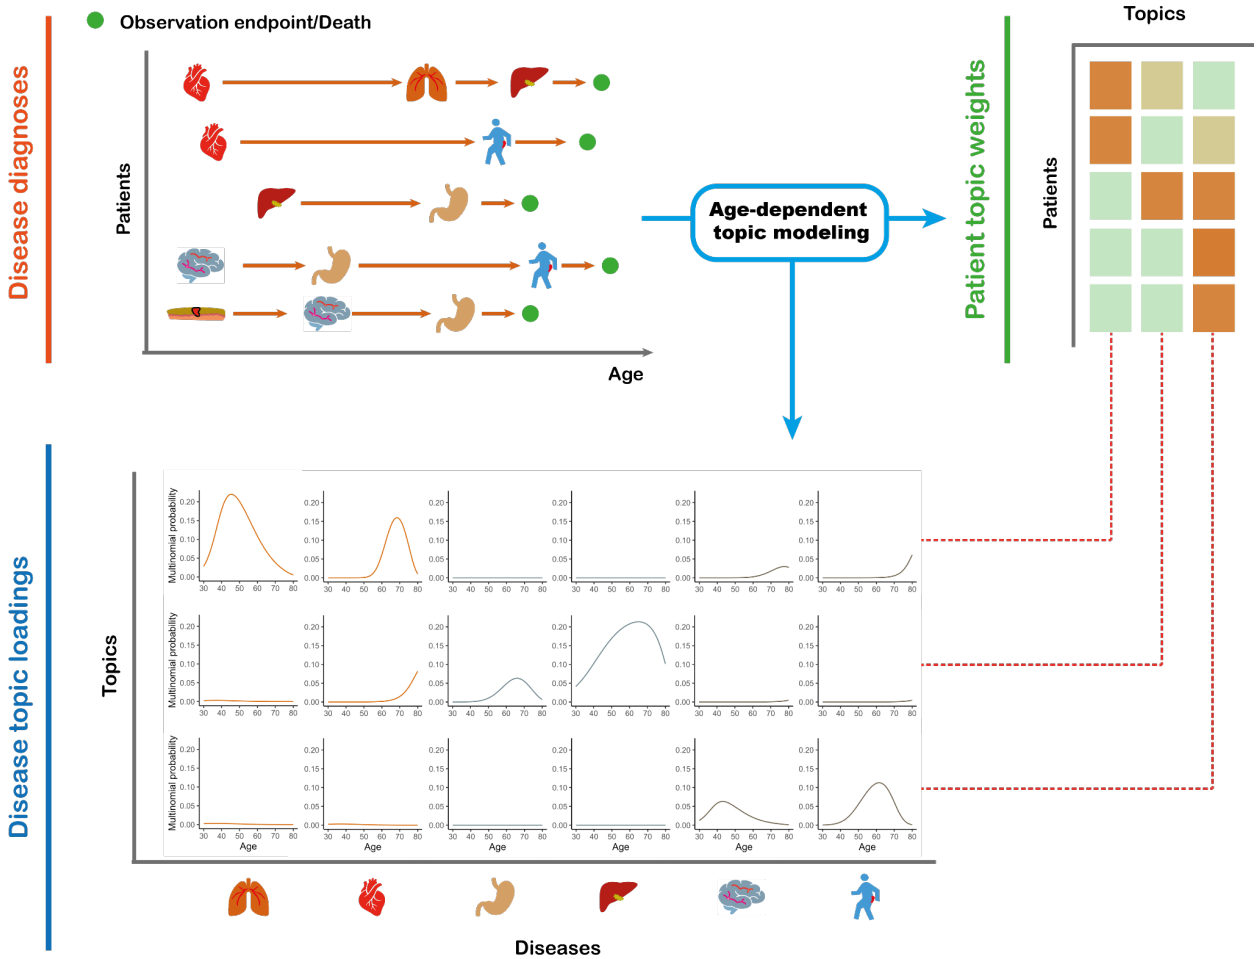

Figure 1: ATM provides an efficient way to represent longitudinal comorbidity data. Top left: input consists of disease diagnoses as a function of age. Top right: ATM assigns a topic weight to each patient. Bottom: ATM infers age-dependent topic loadings.

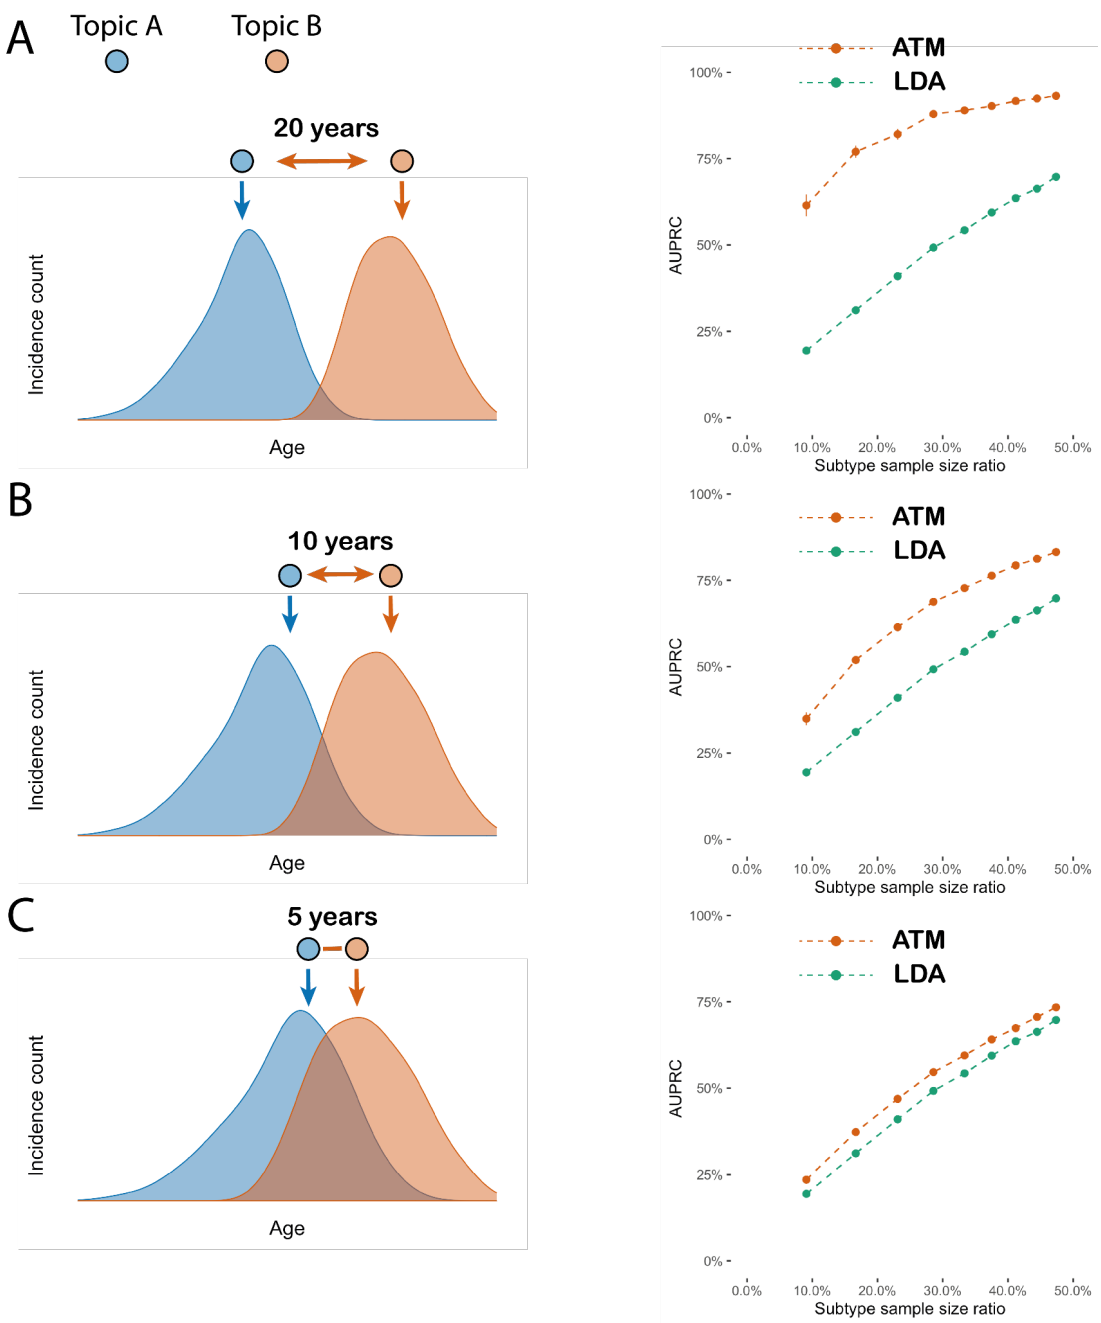

Figure 2: ATM outperforms LDA in simulations with age-dependent effects. In simulations at different levels of age-dependent effects (left panels), we report the area under the precision and recall curve (AUPRC) for ATM vs. LDA as a function of subtype sample size proportion (the proportion of diagnoses belonging to the smaller subtype) (right panels). Each dot represents the mean of 100 simulations of 10,000 individuals. Error bars denote 95% confidence intervals. (A) 20-year difference in age at diagnosis for the two subtypes. (B) 10-year difference in age at diagnosis for the two subtypes. (C) 5-year difference in age at diagnosis for the two subtypes. Numerical results are reported in Supplementary Table 2.

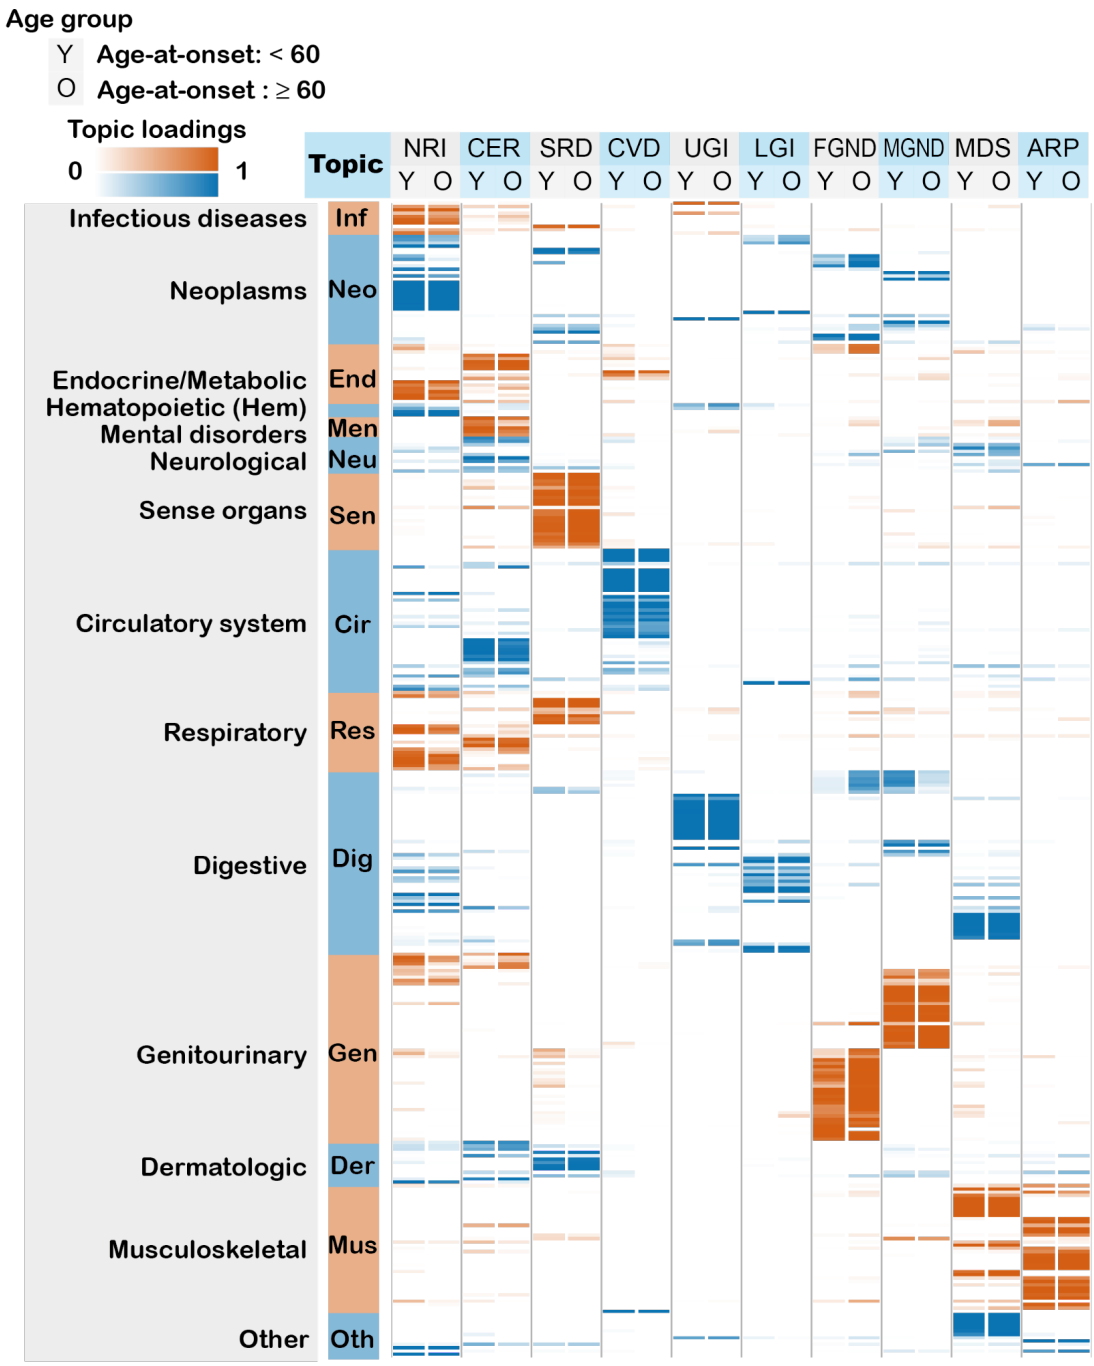

Figure 3. Age-dependent topic loadings of 10 inferred disease topics across 348 diseases in the UK Biobank. We report topic loadings averaged across younger ages (age at diagnosis < 60) and older ages (age at diagnosis > 60). Row labels denote disease categories ordered by Phecode systems, with alternating blue and red color for visualisation purposes; "Other" is a merge of five Phecode systems: "congenital anomalies", "symptoms", "injuries & poisoning", "other tests", and "death" (which is treated as an additional disease, see Methods). Topics are ordered by the corresponding Phecode system. Further details on the 10 topics are provided in Table 1. Further

details on the diseases discussed in the text (type 2 diabetes and breast cancer) are provided in Supplementary Figure 6. Numerical results are reported in Supplementary Table 4.

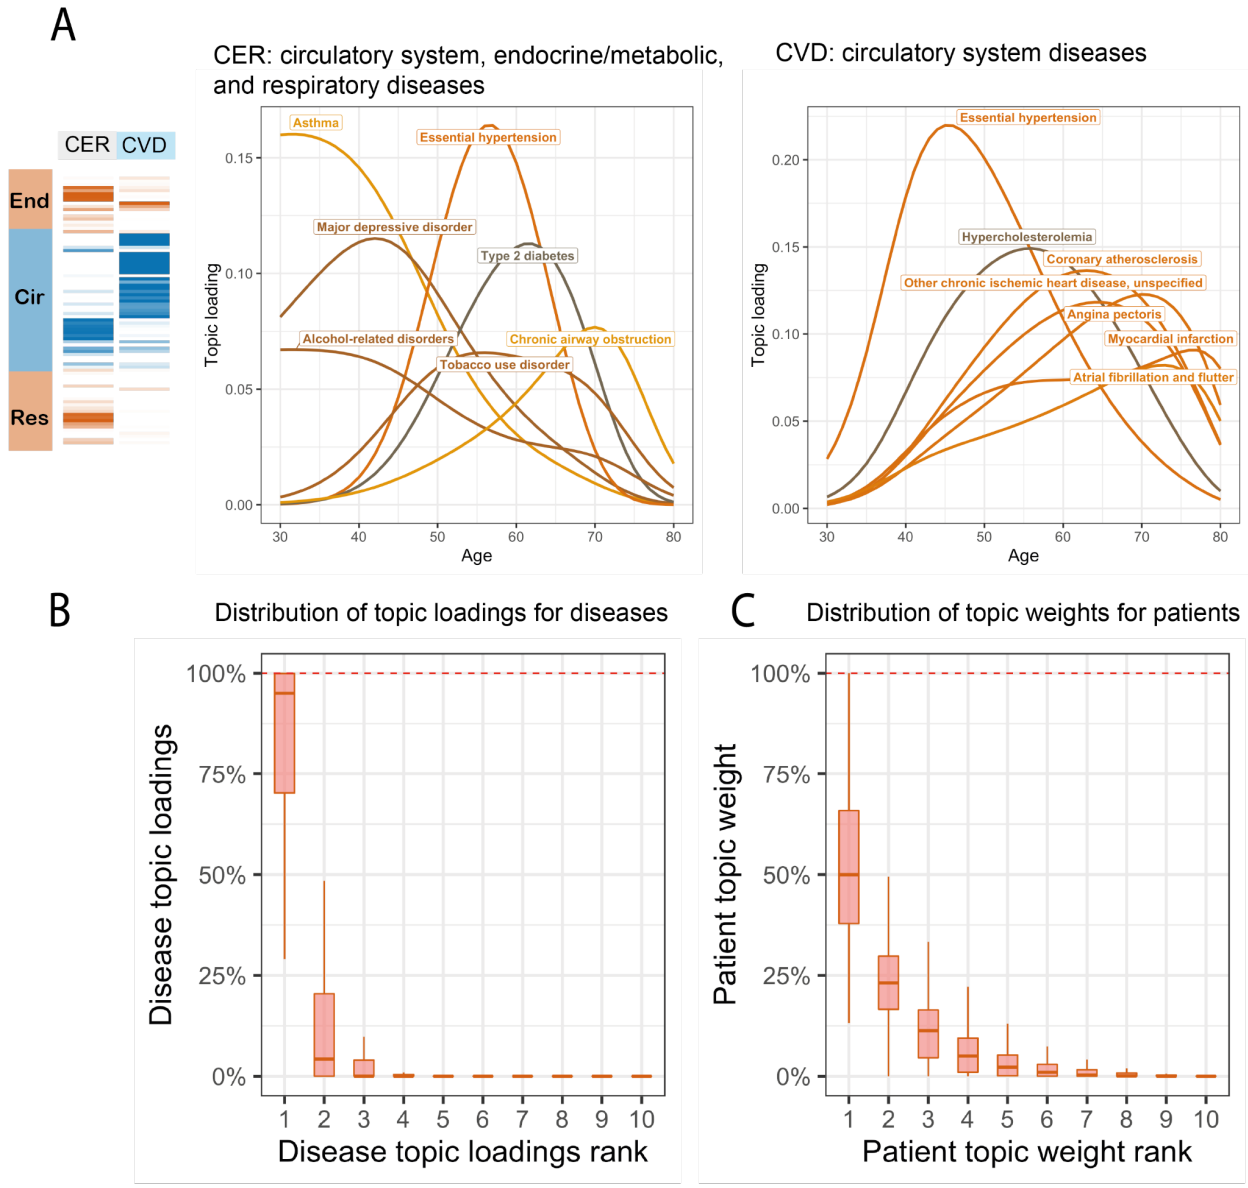

Figure 4. Topic loadings capture age-dependent comorbidities. (A) Age-dependent topic loadings for two representative topics, CER and CVD; for each topic, we include the top seven diseases with highest topic loadings. Results for all 10 topics are reported in Supplementary Figure 13. (B) Box plot of disease topic loading as a function of rank; disease topic loadings are computed as a weighted average across all values of age at diagnosis. (C) Box plot of patient topic weight as a function of rank. Numerical results are reported in Supplementary Table 5.

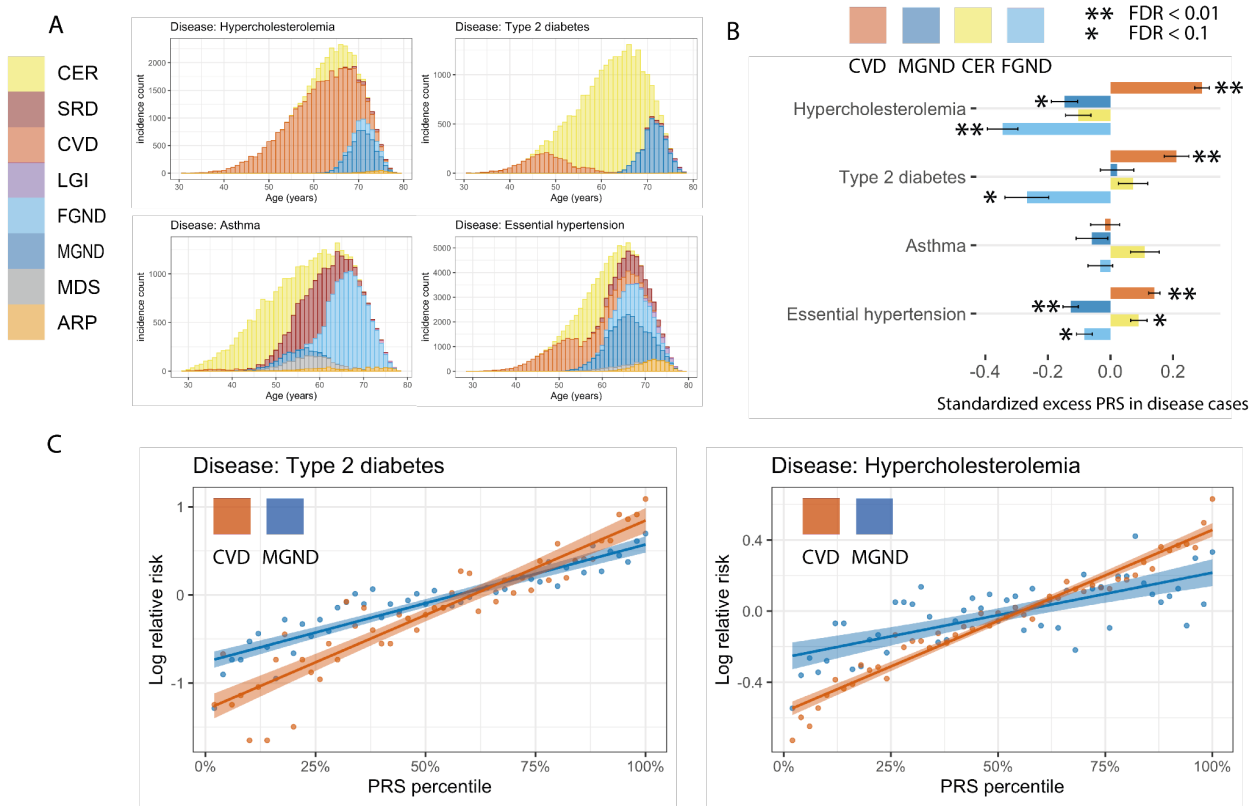

Figure 5. Polygenic risk scores vary across disease subtypes defined by distinct topics. (A) Stacked barplots of age-dependent subtypes (defined by topics) for 4 representative diseases (type 2 diabetes, asthma, hypercholesterolemia, and essential hypertension); for each disease, we include all subtypes with at least one diagnosis. Results for all 52 diseases are reported in Supplementary Figure 16. (B) Standardised excess PRS values in disease cases (s.d. increase in PRS per unit increase in patient topic weight) for 4 representative diseases and 4 corresponding topics. (C) Relative risk for cases of type 2 diabetes and hypercholesterolemia of CVD and MGND subtypes (vs. controls) across PRS percentiles. Each point spans 2 PRS percentiles. Lines denote regression on log scale. Error bars denote 95% confidence intervals. Numerical results are reported in Supplementary Table 6.

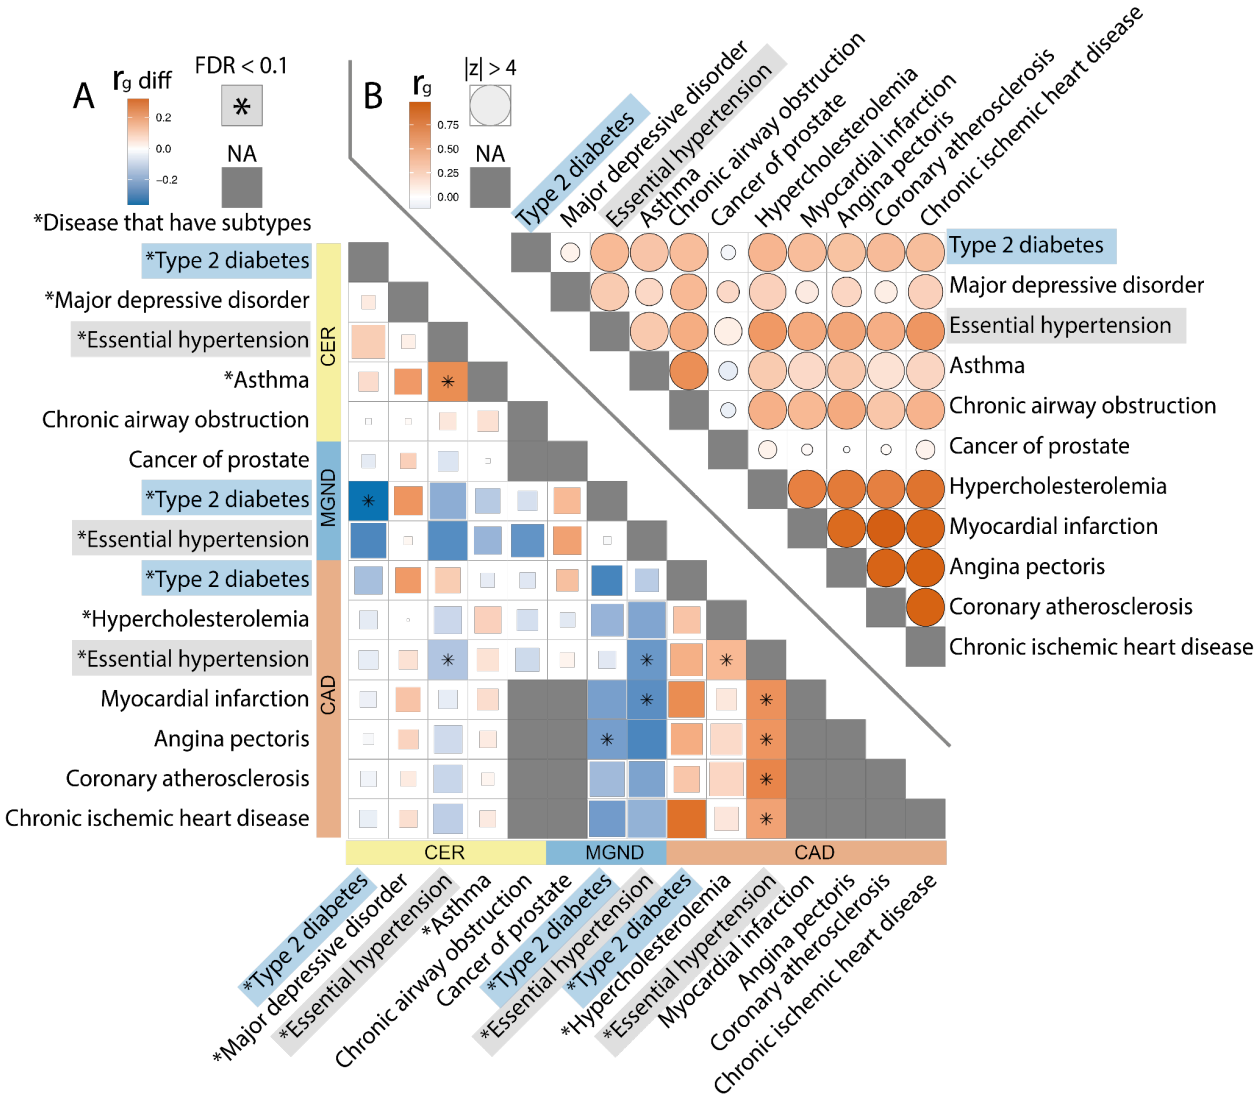

992
993 **Figure 6. Genetic correlations vary across disease subtypes defined by distinct topics.** (a)
994 Excess genetic correlations for pairs of 15 disease subtypes or diseases (9 disease subtypes
995 (denoted with asterisks) + 6 diseases without subtypes), relative to genetic correlations between
996 the underlying diseases. Full square with asterisk denotes FDR < 0.1; less than full squares have
997 area proportional to z-scores for difference. Grey squares denote NA (pair of diseases without
998 subtypes or pair of same disease subtype or disease). (b) Genetic correlations between the
999 underlying diseases. Full circle denotes |z-score| > 4 for nonzero genetic correlation; less than
1000 full circles have area proportional to |z-score|. Numerical results are reported in Supplementary
1001 Table 9.
1002

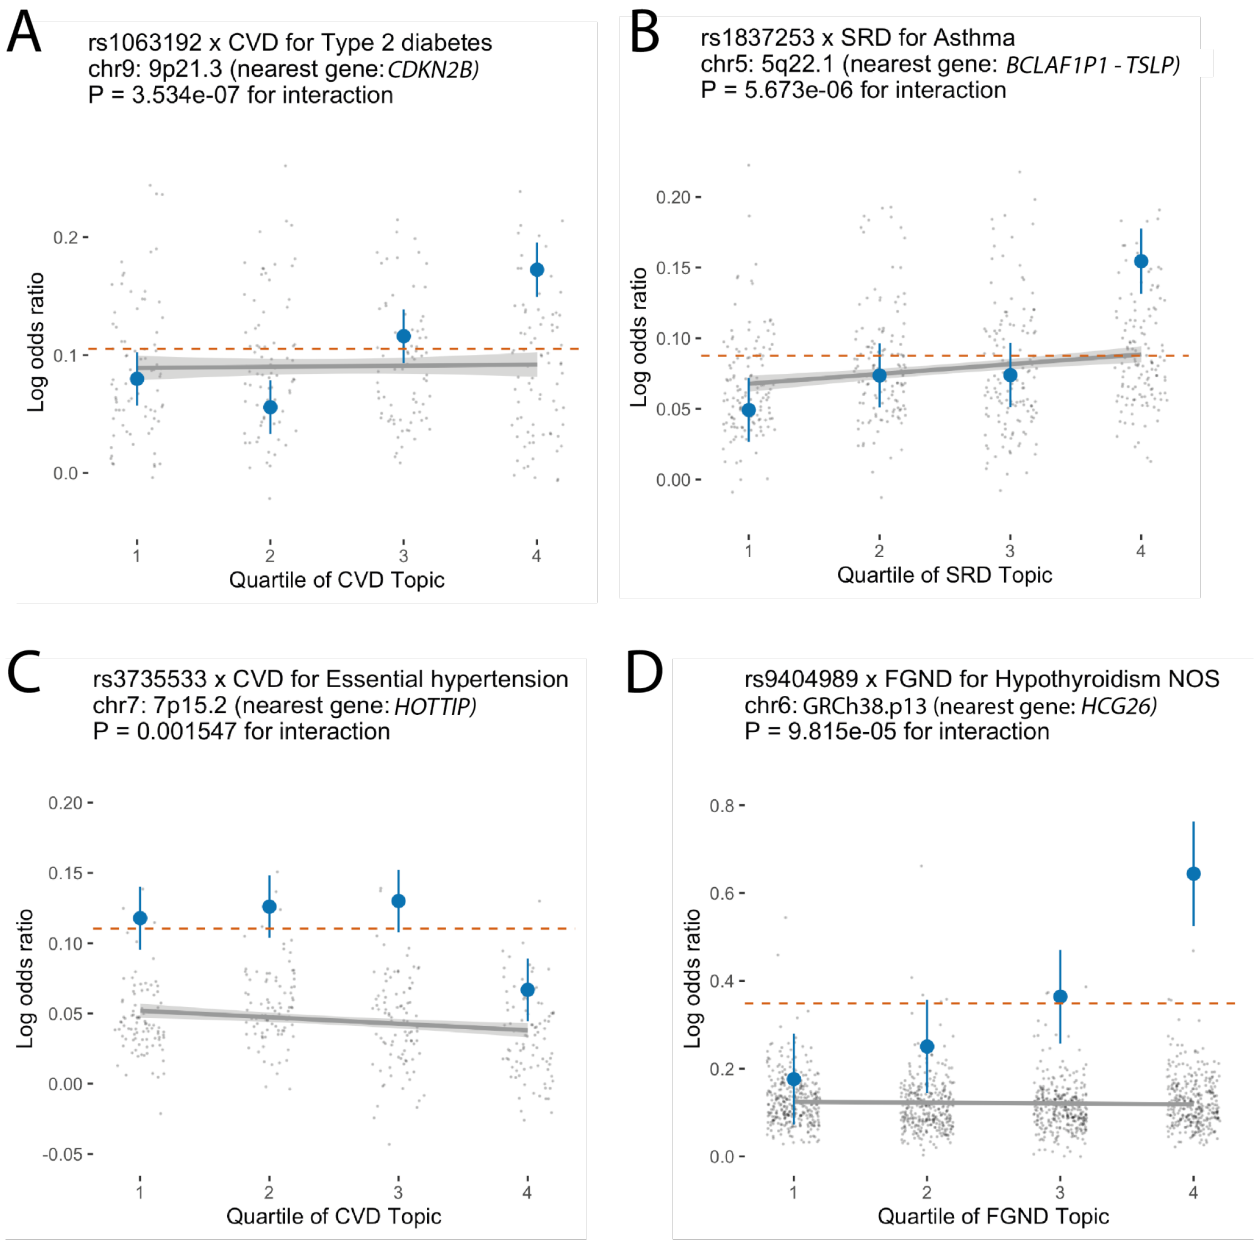

Figure 7. Examples of SNP x topic interaction effects on disease phenotypes. For each example, we report main SNP effects (log odds ratios) specific to each quartile of topic weights across individuals, for both the focal SNP (blue dots) and background SNPs for that disease and topic (genome-wide significant main effect ($P < 5 \times 10^{-8}$) but non-significant SNP x topic interaction effect ($P > 0.05$); grey dots). Dashed red lines denote aggregate main SNP effects for each focal SNP. Error bars denote 95% confidence intervals. Grey lines denote linear regression of grey dots, with grey shading denoting corresponding 95% confidence intervals. Numerical results are reported in Supplementary Table 14.

1013 # References

1014 1. Abul-Husn, N. S. & Kenny, E. E. Personalized Medicine and the Power of Electronic Health
1015 Records. *Cell* **177**, 58–69 (2019).
1016 2. Bulik-Sullivan, B. *et al.* An atlas of genetic correlations across human diseases and traits.
1017 *Nat. Genet.* **47**, 1236–1241 (2015).
1018 3. Wang, K., Gaitsch, H., Poon, H., Cox, N. J. & Rzhetsky, A. Classification of common
1019 human diseases derived from shared genetic and environmental determinants. *Nat. Genet.*
1020 **49**, 1319–1325 (2017).
1021 4. Zhao, W. *et al.* Identification of new susceptibility loci for type 2 diabetes and shared
1022 etiological pathways with coronary heart disease. *Nat. Genet.* **49**, 1450–1457 (2017).
1023 5. Zhu, Z. *et al.* A genome-wide cross-trait analysis from UK Biobank highlights the shared
1024 genetic architecture of asthma and allergic diseases. *Nat. Genet.* **50**, 857–864 (2018).
1025 6. Turley, P. *et al.* Multi-trait analysis of genome-wide association summary statistics using
1026 MTAG. *Nat. Genet.* **50**, 229–237 (2018).
1027 7. O'Connor, L. J. & Price, A. L. Distinguishing genetic correlation from causation across 52
1028 diseases and complex traits. *Nat. Genet.* **50**, 1728–1734 (2018).
1029 8. Cortes, A., Albers, P. K., Dendrou, C. A., Fugger, L. & McVean, G. Identifying cross-
1030 disease components of genetic risk across hospital data in the UK Biobank. *Nat. Genet.* **52**,
1031 126–134 (2019).
1032 9. Morrison, J., Knoblauch, N., Marcus, J. H., Stephens, M. & He, X. Mendelian randomization
1033 accounting for correlated and uncorrelated pleiotropic effects using genome-wide summary
1034 statistics. *Nat. Genet.* **52**, 740–747 (2020).
1035 10. Peyrot, W. J. & Price, A. L. Identifying loci with different allele frequencies among cases of
1036 eight psychiatric disorders using CC-GWAS. *Nat. Genet.* **53**, 445–454 (2021).
1037 11. Mattheisen, M. *et al.* Identification of shared and differentiating genetic architecture for

1038 autism spectrum disorder, attention-deficit hyperactivity disorder and case subgroups. *Nat.*
1039 *Genet.* **54**, 1470–1478 (2022).
1040 12. Cortes, A. *et al.* Bayesian analysis of genetic association across tree-structured routine
1041 healthcare data in the UK Biobank. *Nat. Genet.* **49**, 1311–1318 (2017).
1042 13. Zhang, H. *et al.* Genome-wide association study identifies 32 novel breast cancer
1043 susceptibility loci from overall and subtype-specific analyses. *Nat. Genet.* **52**, 572–581
1044 (2020).
1045 14. Mansour Aly, D. *et al.* Genome-wide association analyses highlight etiological differences
1046 underlying newly defined subtypes of diabetes. *Nat. Genet.* **53**, 1534–1542 (2021).
1047 15. Hautakangas, H. *et al.* Genome-wide analysis of 102,084 migraine cases identifies 123 risk
1048 loci and subtype-specific risk alleles. *Nat. Genet.* **54**, 152–160 (2022).
1049 16. Srebro, N. & Shraibman, A. Rank, Trace-Norm and Max-Norm. in *Learning Theory* 545–
1050 560 (Springer Berlin Heidelberg, 2005).
1051 17. Candès, E. & Recht, B. Exact matrix completion via convex optimization. *Commun. ACM*
1052 **55**, 111–119 (2012).
1053 18. Yan, J. & Pollefeys, M. A General Framework for Motion Segmentation: Independent,
1054 Articulated, Rigid, Non-rigid, Degenerate and Non-degenerate. in *Computer Vision – ECCV*
1055 *2006* 94–106 (Springer Berlin Heidelberg, 2006).
1056 19. Ma, Y., Derksen, H. & Hong, W. Segmentation of multivariate mixed data via Lossy data
1057 coding and compression. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**, 1546–1562 (2007).
1058 20. Rao, S., Tron, R., Vidal, R. & Ma, Y. Motion segmentation in the presence of outlying,
1059 incomplete, or corrupted trajectories. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**, 1832–
1060 1845 (2010).
1061 21. Liu, G. & Yan, S. Latent Low-Rank Representation for subspace segmentation and feature
1062 extraction. in *2011 International Conference on Computer Vision* 1615–1622
1063 (ieeexplore.ieee.org, 2011).

22. Liu, Z. *et al.* Efficient Low-rank Multimodal Fusion with Modality-Specific Factors. *arXiv [cs.AI]* (2018).
23. Chen, Y. & Chi, Y. Harnessing Structures in Big Data via Guaranteed Low-Rank Matrix Estimation: Recent Theory and Fast Algorithms via Convex and Nonconvex Optimization. *IEEE Signal Process. Mag.* **35**, 14–31 (2018).
24. Chen, Y. & Candès, E. J. The projected power method: An efficient algorithm for joint alignment from pairwise differences. *Commun. Pure Appl. Math.* **71**, 1648–1714 (2018).
25. Jia, G. *et al.* Estimating heritability and genetic correlations from large health datasets in the absence of genetic data. *Nat. Commun.* **10**, 1–11 (2019).
26. Tanigawa, Y. *et al.* Components of genetic associations across 2,138 phenotypes in the UK Biobank highlight adipocyte biology. *Nat. Commun.* **10**, 4064 (2019).
27. Sakaue, S. *et al.* A cross-population atlas of genetic associations for 220 human phenotypes. *Nat. Genet.* **53**, 1415–1424 (2021).
28. Jia, G. *et al.* Discerning asthma endotypes through comorbidity mapping. *Nat. Commun.* **13**, 1–19 (2022).
29. Bycroft, C. *et al.* The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**, 203–209 (2018).
30. of Us Research Program, A. The 'All of Us' Research Program. *N. Engl. J. Med.* **381**, 668–676 (2019).
31. Ishigaki, K. *et al.* Large-scale genome-wide association study in a Japanese population identifies novel susceptibility loci across different diseases. *Nat. Genet.* **52**, 669–679 (2020).
32. Siggaard, T. *et al.* Disease trajectory browser for exploring temporal, population-wide disease progression patterns in 7.2 million Danish patients. *Nat. Commun.* **11**, 1–10 (2020).
33. Posey, J. E. *et al.* Resolution of Disease Phenotypes Resulting from Multilocus Genomic Variation. *N. Engl. J. Med.* **376**, 21–31 (2017).

1090 34. Cook, E. K. *et al.* Comorbid and inflammatory characteristics of genetic subtypes of clonal
1091 hematopoiesis. *Blood Adv* **3**, 2482–2486 (2019).
1092 35. Udler, M. S. *et al.* Type 2 diabetes genetic loci informed by multi-trait associations point to
1093 disease mechanisms and subtypes: A soft clustering analysis. *PLoS Med.* **15**, e1002654
1094 (2018).
1095 36. Wani, B., Aziz, S. A., Ganaie, M. A. & Mir, M. H. Metabolic Syndrome and Breast Cancer
1096 Risk. *Indian J. Med. Paediatr. Oncol.* **38**, 434–439 (2017).
1097 37. Blei, Ng & Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.* (2003).
1098 38. Pritchard, J. K., Stephens, M. & Donnelly, P. Inference of population structure using
1099 multilocus genotype data. *Genetics* **155**, 945–959 (2000).
1100 39. Bishop, C. M. *Pattern Recognition and Machine Learning*. (Springer New York, 2006).
1101 40. Teh, Y., Newman, D. & Welling, M. A collapsed variational Bayesian inference algorithm for
1102 latent Dirichlet allocation. *Adv. Neural Inf. Process. Syst.* **19**, (2006).
1103 41. Grau, J., Grosse, I. & Keilwagen, J. PRROC: computing and visualizing precision-recall and
1104 receiver operating characteristic curves in R. *Bioinformatics* **31**, 2595–2597 (2015).
1105 42. Wu, P. *et al.* Mapping ICD-10 and ICD-10-CM Codes to Phecodes: Workflow Development
1106 and Initial Evaluation. *JMIR Med Inform* **7**, e14325 (2019).
1107 43. Suvila, K. *et al.* Early Onset Hypertension Is Associated With Hypertensive End-Organ
1108 Damage Already by MidLife. *Hypertension* HYPERTENSIONAHA11913069 (2019).
1109 44. Wong, B. *et al.* Cardiovascular Disease Risk Associated With Familial
1110 Hypercholesterolemia: A Systematic Review of the Literature. *Clin. Ther.* **38**, 1696–1709
1111 (2016).
1112 45. Shah, M. S. & Brownlee, M. Molecular and Cellular Mechanisms of Cardiovascular
1113 Disorders in Diabetes. *Circ. Res.* **118**, 1808–1829 (2016).
1114 46. Shah, A. D. *et al.* Type 2 diabetes and incidence of cardiovascular diseases: a cohort study
1115 in 1·9 million people. *The Lancet Diabetes & Endocrinology* **3**, 105–113 (2015).

1116 47. Dabelea, D. & Hamman, R. F. Elevated Cardiometabolic Risk Profile Among Young Adults
1117 With Diabetes: Need for Action. *Diabetes care* vol. 42 1845–1846 (2019).
1118 48. Wong, J. L. & Evans, S. E. Bacterial Pneumonia in Patients with Cancer: Novel Risk
1119 Factors and Management. *Clin. Chest Med.* **38**, 263–277 (2017).
1120 49. Falstie-Jensen, A. M. *et al.* Incidence of hypothyroidism after treatment for breast cancer—
1121 a Danish matched cohort study. *Breast Cancer Res.* **22**, 1–10 (2020).
1122 50. Loh, P.-R., Kichaev, G., Gazal, S., Schoech, A. P. & Price, A. L. Mixed-model association
1123 for biobank-scale datasets. *Nat. Genet.* **50**, 906–908 (2018).
1124 51. Loh, P.-R. *et al.* Efficient Bayesian mixed-model analysis increases association power in
1125 large cohorts. *Nat. Genet.* **47**, 284–290 (2015).
1126 52. Jiang, X., Holmes, C. & McVean, G. The impact of age on genetic risk for common
1127 diseases. *PLoS Genet.* **17**, e1009723 (2021).
1128 53. Weir, B. S. & Cockerham, C. C. ESTIMATING F-STATISTICS FOR THE ANALYSIS OF
1129 POPULATION STRUCTURE. *Evolution* **38**, 1358–1370 (1984).
1130 54. Bhatia, G., Patterson, N., Sankararaman, S. & Price, A. L. Estimating and interpreting FST:
1131 the impact of rare variants. *Genome Res.* **23**, 1514–1521 (2013).
1132 55. Scott, L. J. *et al.* A genome-wide association study of type 2 diabetes in Finns detects
1133 multiple susceptibility variants. *Science* **316**, 1341–1345 (2007).
1134 56. Maller, J. B. *et al.* Bayesian refinement of association signals for 14 loci in 3 common
1135 diseases. *Nat. Genet.* **44**, 1294–1301 (2012).
1136 57. Cheng, M., An, S. & Li, J. CDKN2B-AS may indirectly regulate coronary artery disease-
1137 associated genes via targeting miR-92a. *Gene* **629**, 101–107 (2017).
1138 58. Scott, R. A. *et al.* An Expanded Genome-Wide Association Study of Type 2 Diabetes in
1139 Europeans. *Diabetes* **66**, 2888–2902 (2017).
1140 59. Aragam, K. G. *et al.* Discovery and systematic characterization of risk variants and genes
1141 for coronary artery disease in over a million participants. *bioRxiv* (2021)

1142 doi:10.1101/2021.05.24.21257377.
1143 60. Indra, A. K. Epidermal TSLP: a trigger factor for pathogenesis of atopic dermatitis. *Expert*
1144 *Rev. Proteomics* **10**, 309–311 (2013).
1145 61. Kanai, M. *et al.* Genetic analysis of quantitative traits in the Japanese population links cell
1146 types to complex human diseases. *Nat. Genet.* **50**, 390–400 (2018).
1147 62. Oluwafemi, O. O. *et al.* Genome-Wide Association Studies of Conotruncal Heart Defects
1148 with Normally Related Great Vessels in the United States. *Genes* **12**, (2021).
1149 63. Blei, D. M. & Lafferty, J. D. A correlated topic model of Science. *Ann. Appl. Stat.* **1**, 17–35
1150 (2007).
1151 64. Zaitlen, N. *et al.* Informed conditioning on clinical covariates increases power in case-
1152 control association studies. *PLoS Genet.* **8**, e1003032 (2012).
1153 65. Sun, B. B. *et al.* Genetic regulation of the human plasma proteome in 54,306 UK Biobank
1154 participants. *bioRxiv* 2022.06.17.496443 (2022) doi:10.1101/2022.06.17.496443.
1155 66. Mostafavi, H. *et al.* Variable prediction accuracy of polygenic scores within an ancestry
1156 group. *Elife* **9**, e48376 (2020).
1157 67. Dumitrescu, L. *et al.* Evidence for age as a modifier of genetic associations for lipid levels.
1158 *Ann. Hum. Genet.* **75**, 589–597 (2011).
1159 68. Chatterjee, N., Shi, J. & García-Closas, M. Developing and evaluating polygenic risk
1160 prediction models for stratified disease prevention. *Nat. Rev. Genet.* **17**, 392–406 (2016).
1161 69. Lin, J. *et al.* Integration of biomarker polygenic risk score improves prediction of coronary
1162 heart disease in UK Biobank and FinnGen. *bioRxiv* (2022)
1163 doi:10.1101/2022.08.22.22279057.
1164 70. Ghorbani, B., Javadi, H. & Montanari, A. An Instability in Variational Inference for Topic
1165 Models. in *Proceedings of the 36th International Conference on Machine Learning* (eds.
1166 Chaudhuri, K. & Salakhutdinov, R.) vol. 97 2221–2231 (PMLR, 09--15 Jun 2019).
1167 71. Chang, C. C. *et al.* Second-generation PLINK: rising to the challenge of larger and richer

1168 datasets. *Gigascience* **4**, 7 (2015).

1169 72. International HapMap 3 Consortium *et al.* Integrating common and rare genetic variation in

1170 diverse human populations. *Nature* **467**, 52–58 (2010).

1171