

Medical history predicts phenome-wide disease onset and enables the rapid response to emerging health threats

Received: 17 November 2023

Accepted: 3 May 2024

Published online: 20 May 2024

 Check for updates

Jakob Steinfeldt ^{1,2,3,4,5,18}, Benjamin Wild ^{6,18}, Thore Buergele ^{5,6,18}, Maik Pietzner ^{3,7,8}, Julius Upmeier zu Belzen ⁶, Andre Vauvelle⁹, Stefan Hegselmann ^{10,11}, Spiros Denaxas^{9,12,13,14}, Harry Hemingway ^{9,13,14}, Claudia Langenberg ^{3,7,8}, Ulf Landmesser ^{1,2,4,15,16,19}, John Deanfield^{5,19} & Roland Eils ^{6,17,19} ✉

The COVID-19 pandemic exposed a global deficiency of systematic, data-driven guidance to identify high-risk individuals. Here, we illustrate the utility of routinely recorded medical history to predict the risk for 1883 diseases across clinical specialties and support the rapid response to emerging health threats such as COVID-19. We developed a neural network to learn from health records of 502,460 UK Biobank. Importantly, we observed discriminative improvements over basic demographic predictors for 1774 (94.3%) endpoints. After transferring the unmodified risk models to the All of US cohort, we replicated these improvements for 1347 (89.8%) of 1500 investigated endpoints, demonstrating generalizability across healthcare systems and historically underrepresented groups. Ultimately, we showed how this approach could have been used to identify individuals vulnerable to severe COVID-19. Our study demonstrates the potential of medical history to support guidance for emerging pandemics by systematically estimating risk for thousands of diseases at once at minimal cost.

The early phase of the COVID-19 pandemic exposed a global deficiency in delivering systematic, data-driven guidance for individual patients and healthcare providers with critical implications for pandemic preparedness. The assessment of an individual's risk for future disease is central to guiding preventive interventions, early detection of disease, and the initiation of treatments. However, bespoke risk scores are only available for a subset of common diseases^{1–4}, leaving healthcare providers and individuals with little to no guidance on most relevant diseases. Even for diseases with established risk scores, little consensus exists on which score to use and associated physical or laboratory measurements to obtain, leading to highly fragmented practice in routine care⁵. Importantly, in the early phases of emerging pandemics such as COVID-19, it is necessary to allocate sparse resources, but risk scores to identify vulnerable subpopulations are not available due to the lack of available data.

At the same time, most medical decisions on diagnosis, treatment, and prevention of diseases are fundamentally based on an individual's medical history⁶. With the widespread digitalization, this information is routinely collected by healthcare providers, insurance, and governmental organizations at a population scale in the form of electronic health records^{7–12}. These readily accessible records, which include diseases, medications, and procedures, are potentially informative about future risk trajectories, but their potential to improve medical decision-making is limited by the human ability to process and understand vast amounts of data¹³.

To date, routine health records have been used to guide clinical decision-making with etiological^{14–17}, diagnostic^{18,19}, and prognostic research^{15,16,20–22}. Existing efforts often extract and leverage known clinical predictors with new methodologies¹⁹, augment them with additionally extracted data modalities such as clinical notes²³, or aim to

Table 1 | The study population

	UK Biobank			All Of Us				
	Male, N = 229,107	Female, N = 273,353	Overall, N = 502,460	Male, N = 83,013	Female, N = 140,371	Diverse/ Unknown, N = 6446	Overall, N = 229,830	
Age (years)	58 (50, 64)	57 (50, 63)	58 (50, 63)	57 (42, 67)	52 (36, 64)	58 (44, 68)	54 (38, 65)	
	Unknown			1,804	3503	188	5495	
Ethnicity	Asian	5878 (2.6%)	5575 (2.0%)	11,453 (2.3%)	2201 (2.7%)	3742 (2.7%)	42 (0.7%)	5985 (2.6%)
	Black	3407 (1.5%)	4653 (1.7%)	8060 (1.6%)	17,487 (21%)	27,812 (20%)	779 (12%)	46,078 (20%)
	Mixed	1105 (0.5%)	1851 (0.7%)	2956 (0.6%)	2,852 (3.4%)	4883 (3.5%)	111 (1.7%)	7846 (3.4%)
	White	215,244 (95%)	257,413 (96%)	472,657 (95%)	47,467 (57%)	75,613 (54%)	1078 (17%)	124,158 (54%)
	Unknown	3473 (1.5%)	3861 (1.4%)	7334 (1.5%)	13,006 (16%)	28,321 (20%)	4436 (69%)	45,763 (20%)
Smoking status	Current	28,607 (13%)	24,364 (9.0%)	52,971 (11%)	12,553 (15%)	12,562 (8.9%)	668 (10%)	25,783 (11%)
	Previous	87,604 (38%)	85,440 (31%)	173,044 (35%)	21,789 (26%)	29,679 (21%)	1606 (25%)	53,074 (23%)
	Never	111,460 (49%)	162,037 (60%)	273,497 (55%)	39,621 (48%)	87,996 (63%)	3398 (53%)	130,985 (57%)
	Unknown	1436 (0.6%)	1512 (0.7%)	2948 (0.5%)	9050 (11%)	10,164 (7.2%)	774 (12%)	19,988 (8.7%)
Body Mass Index		27.3 (25.0, 30.1)	26.1 (23.5, 29.7)	26.7 (24.1, 29.9)	28 (25, 32)	29 (24, 35)	29 (25, 34)	29 (25, 34)
	Unknown	1646 (0.7%)	1458 (0.6%)	3104 (0.6%)	2802 (3.4%)	5711 (4.1%)	387 (6.0%)	8900 (3.9%)
Systolic Blood Pressure (mmHg)		139 (128, 152)	133 (121, 147)	136 (124, 150)	129 (119, 141)	123 (112, 136)	127 (115, 140)	125 (114, 138)
	Unknown	13,579 (5.9%)	16,536 (6.0%)	30,115 (6.0%)	1804 (2.2%)	3503 (2.5%)	188 (2.9%)	5495 (2.4%)
Median (IQR); n (%)								

identify novel predictors among the recorded concepts^{14–17}. Prior work on the prediction of disease onset has mainly focused on single diseases, including dementia^{15,24}, cardiovascular conditions^{23,25} such as heart failure²⁶ and atrial fibrillation^{27,28}. In contrast, phenome-wide association studies (PheWAS) quantifying the associations of genetic variants with comprehensive phenotypic traits are emerging in genetic epidemiology^{29,30}. While approaches have been developed for high-throughput phenotyping^{31,32} and to extract information from longitudinal health records^{33,34}, no studies have investigated the predictive potential and potential utility over the entire human phenome. Consequently, the predictive information in routinely collected health records and its potential to systematically guide medical decision-making is largely unexplored.

Here, we examined the predictive potential of an individual's entire medical history and propose a systematic approach for phenome-wide risk stratification. We developed, trained, and validated a neural network in the UK Biobank cohort³⁵ to estimate disease risk from routinely collected health records. Unlike alternative methods, such as linear models or survival trees, which require separate models for each disease, our approach employs a multi-layer perceptron that predicts multiple endpoints concurrently, resulting in a significantly simplified model architecture. These endpoints include preventable diseases (e.g., coronary heart disease), diseases that are not currently preventable, but the early diagnosis has been shown to substantially slow down the progression and development of complications (e.g., heart failure), and outcomes, which are currently neither entirely preventable nor treatable (e.g., death). They also include both diseases with risk prediction models recommended in guidelines and used in practice (e.g., cardiovascular diseases or breast cancer) as well as diseases without current risk prediction models (e.g., psoriasis and rheumatoid arthritis).

We evaluated our approach by integrating the endpoint-specific risk states estimated by the neural network in Cox Proportional Hazard models³⁶, investigating the phenome-wide predictive potential over basic demographic predictors, selected comorbidities, and established modifiable risk factors, and illustrating how phenome-wide risk stratification could benefit individuals by providing risk estimates, facilitating early disease diagnosis, and guiding preventive interventions. Furthermore, by externally validating in the All Of Us cohort³⁷, we show that our models can generalize across healthcare systems and

populations, including communities historically underrepresented in biomedical research.

Finally, we assessed the potential of our approach to aid risk stratification for the primary prevention of cardiovascular disease and to respond to emerging health threats at the example of COVID-19. We then show that the risk states of pneumonia, sepsis & all-cause death can be used to calculate a combined severity risk score using primary and secondary care records available before the global spread of the COVID-19 pandemic. Our results demonstrate the currently unused potential of routine health records to guide medical practice by providing comprehensive phenome-wide risk estimates.

Results

Characteristics of the study population and integration of routine health records

This study is based on the UK Biobank cohort^{35,38}, a longitudinal population cohort of 502,460 relatively healthy individuals of primarily British descent, with a median age of 58 (IQR 50, 63) years, 54.4% biological females, 11% current smokers, and a median BMI of 26.7 (IQR 24.1, 29.9) at recruitment (Table 1 for detailed information). Individuals recruited between 2006 and 2010 were followed for a median of 12.6 years, resulting in ~6.2 M overall person-years on 1883 phenome-wide endpoints³⁹ with ≥ 100 incident events ($>0.02\%$ of individuals having the event in the observation time). We externally validated our findings in individuals from the All of Us cohort, a longitudinal cohort of 229,830 individuals with linked health records recruited from all over the United States. Individuals in the All of Us cohort are of diverse descent, with 46% of reportedly non-white ethnicity and 78% of groups historically underrepresented in biomedical research^{37,40}, and have a median age of 54 (IQR 38, 65) years with 61.1% biological females (see Table 1 for detailed information). Individuals were recruited from 2019 on and followed for a median of 3.5 years, resulting in ~787,300 person-years on 1568 endpoints.

Central to this study is the prior medical history, defined as the entirety of routine health records before recruitment. Before further analysis, we mapped all health records to the OMOP vocabulary. While most records originate from primary care and, to a lesser extent, secondary care (Suppl. Figure 1a), the predominant record domains are drugs and observations, followed by conditions, procedures, and devices (Supplementary Fig. 1b). Interestingly, while rare medical concepts

(with a record in <1% of individuals in the study population) are not commonly included in prediction models²¹, they are often associated with high incident event rates (exemplified by the mortality rate in Supplementary Fig. 1c) compared to common concepts (a record present in $\geq 1\%$ of the study population). For example, the concept code for “portal hypertension” (OMOP 34742003) is only recorded in 0.04% (203) of individuals at recruitment, but 48.7% (99 individuals) will die over the course of the observation period. Importantly, there are many distinct rare concepts, and thus 91.7% of individuals have at least one rare record before recruitment, compared with 92.5% for common records. In addition, 60.7% of individuals have ≥ 10 rare records compared with 78.4% for common records, and individuals have only slightly fewer rare than common records (Supplementary Fig. 1d).

After excluding very rare concepts (<0.01%, less than 50 individuals with the record in this study), we integrated the remaining 15,595 unique concepts (Supplementary Data 2) with a multi-task multi-layer perceptron (with 88.4 M parameters) to predict the phenome-wide onset of 1883 endpoints (Supplementary Data 1) simultaneously (Fig. 1a). For comparison, we also include additional comparisons with a linear baseline (with 29.4 M parameters, Supplementary Fig. 2), demonstrating superior performance at a minimal increase of complexity.

To ensure that our findings are generalizable and transferable, we spatially validate our models in 22 recruitment centers (Fig. 1b) across England, Wales, and Scotland. We developed 22 models, each trained on individuals from 21 recruitment centers at recruitment, randomly split into training and validation sets (Fig. 1c). We subsequently tested the models on individuals from the additional recruitment center unseen for model development for internal spatial validation. After checkpoint selection on the validation data sets and obtaining the selected models' final predictions on the individual test sets, the test set predictions were aggregated for downstream analysis (Fig. 1d). Subsequently, disease-specific exclusions of prior events and sex-specificity were respected in all downstream analyses. After development, the models were externally validated in the All of Us cohort³⁷.

Routine health records stratify phenome-wide disease onset

Central to the utility of any predictor is its potential to stratify risk. The better the stratification of low and high-risk individuals, the more effective targeted interventions and disease diagnoses are.

To investigate whether health records can be used to identify high-risk individuals, we assessed the relationship between the risk states estimated by the neural network for each endpoint and the risk of future disease (Fig. 2). For illustration, we first aggregated the incident events over the percentiles of the risk states for each endpoint and subsequently calculated ratios between the top and bottom 10% of risk states over the entire phenome (Fig. 2a). We found that fewer than 10% of the individuals had an incident hypertension diagnosis in the observation window if they were estimated to be in the bottom risk percentile of the medical history, compared to more than 60% if they were estimated to be in the top risk percentile. Subsequently, the incident event ratio between the top and bottom deciles was -5.23. Importantly, we found differences in the event rates, reflecting a stratification of high and low-risk individuals for almost all endpoints covering a broad range of disease categories and etiologies: For 1341 of 1883 endpoints (71.2%), we observed >10-times as many events for individuals in the top 10% of the predicted risk states compared to the bottom 10%. For instance, these endpoints included rheumatoid arthritis (Ratio -11.3), ischemic heart disease (Ratio -23.5), or chronic obstructive pulmonary disease (Ratio -65.4). For 230 (12.2%) of the 1883 conditions, including abdominal aortic aneurysm (Ratio -163.4), more than 100 times the number of individuals in the top 10% of predicted risk states had incident events compared to the bottom 10%. For 542 (28.8%) endpoints, the separation between high and low-risk individuals was smaller (Ratio <10), which included hypertension (Ratio -5.2) and anemia (Ratio -6.7), often diagnosed earlier in life or

precursors for future comorbidities. Notably, the ratios were >1 for all but one of the 1883 investigated endpoints, even though all models were developed in spatially segregated assessment centers. To illustrate how high-risk individuals differ from the moderate cases, we also provide additional ratios comparing the top 10% to individuals in the median 20% of the population. The complete list of all endpoints and corresponding statistics can be found in Supplementary Data 4.

In addition to the phenome-wide analysis of 1883 endpoints, we also provide detailed associations between the risk percentiles and incident event ratios (Fig. 2b), as well as cumulative event rates for up to 15 years (Fig. 2c) of follow-up for the top, median, and bottom percentiles for a subset of 24 selected endpoints. This set was selected to comprise actionable endpoints and common diseases with significant societal burdens, specific cardiovascular conditions with pharmacological and surgical interventions, as well as endpoints without established tools to stratify risk to date. To exemplify the potential of our approach, among individuals in the top risk decile for heart failure, 8018 (16.06%) experienced an event, in contrast to 178 (0.35%) individuals in the bottom decile, resulting in a risk ratio of 46.35 (Fig. 2a, b, Supplementary Data 4). Consequently, those at high risk of heart failure could be prioritized for echocardiographic screening and, if necessary, prescribed effective guideline-directed medical therapy. Similarly, individuals with a high risk of developing COPD—where the top 10% face over 65 times the risk compared to the bottom 10%—may be considered for spirometry, an approach already established in the CAPTURE trial⁴¹. If confirmed, they could benefit from interventions such as long-acting bronchodilators. As a third example, a high-risk estimate for less common diseases, such as multiple sclerosis (risk ratio -8.3), could further support referring individuals to a specialist and potentially shorten the often extensive patient journey before a final diagnosis is reached.

In summary, the disease-specific states stratify the risk of onset for all 1883 investigated endpoints across clinical specialties. This indicates that routine health records provide a large and widely unused potential for the systematic risk estimation of disease onset in the general population.

Discriminative performance indicates potential utility

While routine health records can stratify incident event rates, this does not prove utility. To test whether the risk state derived from the routine health records could provide utility and information beyond ubiquitously available predictors, we investigated the predictive information over age and biological sex, selected comorbidities from the Charlson Comorbidity Index⁴², and established modifiable risk factors from the AHA ASCVD pooled cohort equation³. We modeled the risk of disease onset using Cox Proportional-Hazards (CPH) models for all 1883 endpoints, which allowed us to estimate adjusted hazard ratios (denoted as HR in Supplementary Data 6) and 10-year discriminative improvements (indicated as Delta C-index in Fig. 3a).

We found significant improvements over the baseline model (age and biological sex only) for 1774 (94.2%) of the 1883 investigated endpoints (Fig. 3, Supplementary Data 5). For many of these endpoints, the discriminative improvements were considerable (Delta C-Index Q25%: 0.094, Q50: 0.116, Q75: 0.141). We found significant improvements for 23 of the highlighted subset of 24 endpoints (indicated in Fig. 2a), with the largest increases for the prediction of back pain (Delta C-Index: +0.238 (CI 0.236, 0.241)), suicide attempts (Delta C-Index: +0.224 (CI 0.213, 0.235)), psoriasis (Delta C-Index: +0.171 (CI 0.161, 0.178)), all-cause mortality (Delta C-Index: +0.171 (0.169, 0.174)) and chronic obstructive pulmonary disease (Delta C-Index: +0.156 (0.151, 0.159)). In contrast, we did not find significant improvements in the prediction of 86 (4.6%) of the 1883 endpoints, including, e.g., Parkinson's disease (Delta C-Index: -0.006 (CI -0.013, 0)) or even deteriorations in the prediction of 23 (1.2%) of the endpoints, including neoplasm like cervical cancer (Delta C-Index: -0.025 (-0.059, -0.004))

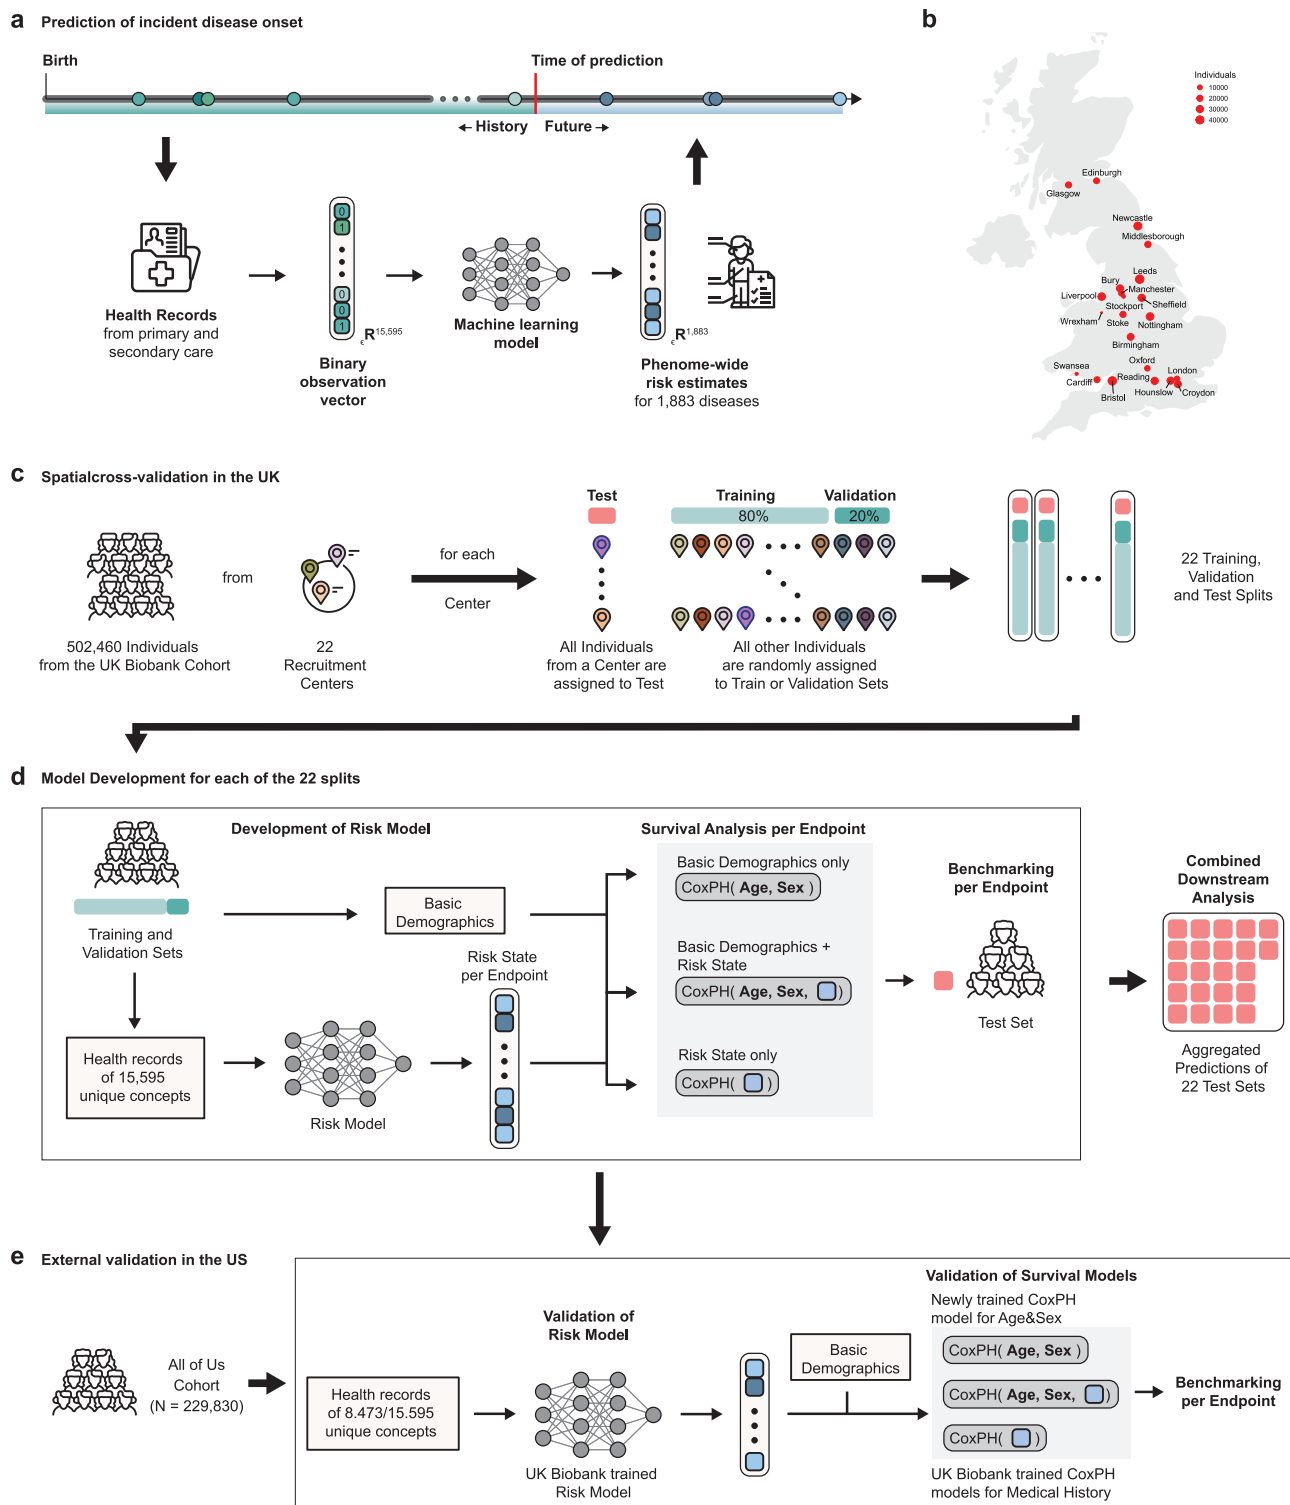


Fig. 1 | Overview of the study. **a** The medical history captures encounters with primary and secondary care, including diagnoses, medications, and procedures (ideally) from birth. Here we train a multi-layer perceptron on data before recruitment to predict phenome-wide incident disease onset for 1883 endpoints. **b** Location and size of the 22 assessment centers of the UK Biobank cohort across England, Wales, and Scotland. **c** To learn risk states from individual medical histories, the UK Biobank population was partitioned by their respective assessment center at recruitment. **d** For each of the 22 partitions, the Risk Model was trained to predict phenome-wide incident disease onset for 1883 endpoints. Subsequently,

for each endpoint, Cox proportional hazard (CPH) models were developed on the risk states in combination with sets of commonly available predictors to model disease risk. Predictions of the CPH model on the test set were aggregated for downstream analysis. **e** External validation in the All of US cohort. After mapping to the OMOP vocabulary, we transferred the trained risk model to the All of US cohort and calculated the risk state for all endpoints. To validate these risk states, we compared the unchanged CPH models developed in the UK Biobank with refitted CPH models for age and sex. Source data are provided. The Icons are made by Freepik from www.flaticon.com.

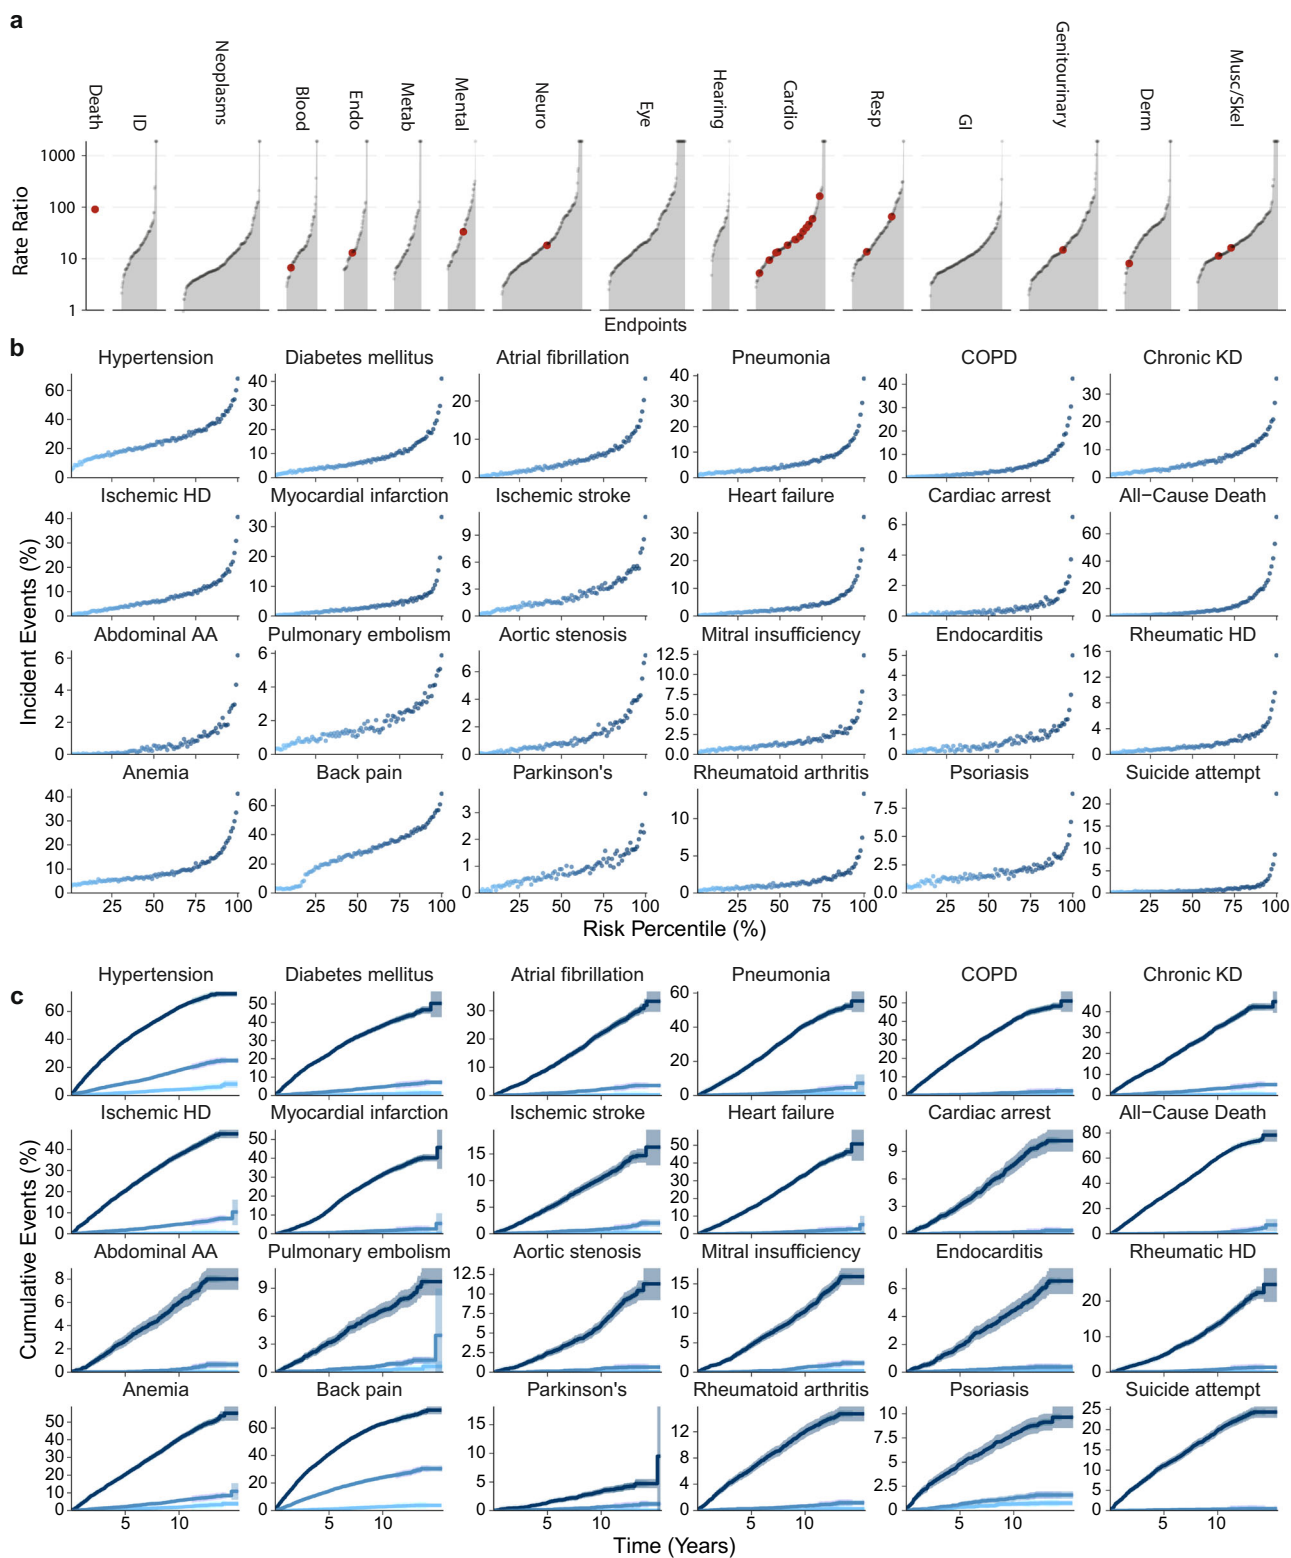


Fig. 2 | Routine health records stratify phenome-wide disease onset. a Ratio of incident events in the Top 10% compared with the Bottom 10% of the estimated risk states. Event rates in the Top 10% are higher than in the Bottom 10% for all but one of the 1883 investigated endpoints. Red dots indicate 24 selected endpoints detailed in Fig. 2b. To illustrate, 1198 (2.39%) individuals in the top risk decile for cardiac arrest experienced an event compared with only 30 (0.06%) in the bottom

decile, with a risk ratio of 39.93. **b** Incident event rates for each medical history risk percentile (if medical history was available) for a selection of 24 endpoints. **c** Cumulative event rates with 95% confidence intervals for the Top 1%, median, and Bottom 1% of risk percentiles in **b** over 15ys. Statistical measures were derived from 502,460 individuals. Individuals with prevalent diseases were excluded from the endpoints-specific analysis. Source data are provided.

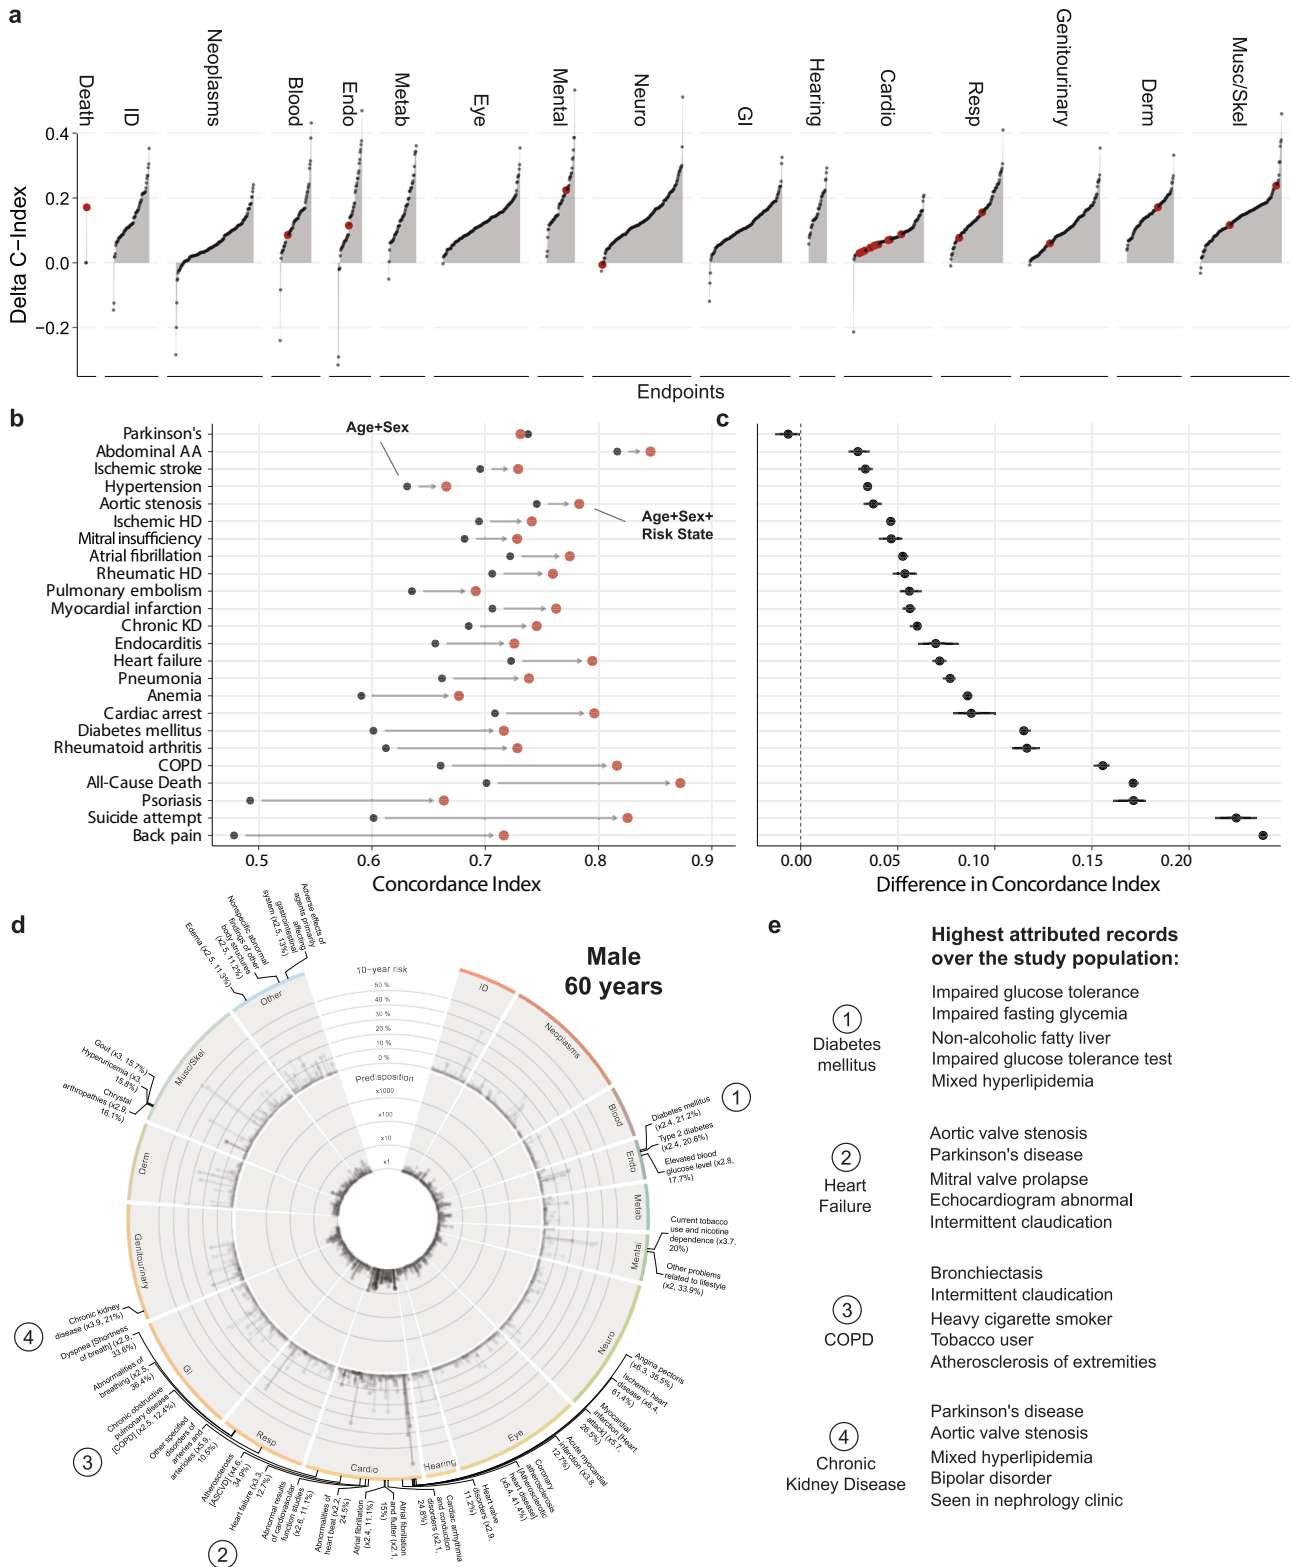


Fig. 3 | Discriminative performance indicates potential utility. a Differences in discriminatory performance quantified by the C-Index between CPH models trained on Age+Sex and Age+Sex+MedicalHistory for all 1883 endpoints. We found significant improvements over the baseline model (Age+Sex, age, and biological sex only) for 1774 (94.2%) of the 1883 investigated endpoints. Red dots indicate selected endpoints in Fig. 3b. **b** Absolute discriminatory performance in terms of C-Index comparing the baseline (Age+Sex, black point) with the added routine health records risk state (Age+Sex+RiskState, red point) for a selection of 24 endpoints. **c** The direct C-index differences for the same models. Dots indicate medians

and whiskers extend to the Bonferroni-corrected 95% confidence interval for a distribution bootstrapped over 100 iterations. **d** Example of individual predicted phenome-wide risk profile. Predisposition (10-year risk estimated by Age+Sex +RiskState compared to risk estimated by Age+Sex alone) is displayed in the inner circle, and absolute 10-year risk estimated by Age+Sex+RiskState can be found in the outer circle. Labels indicate endpoints with a high individual predisposition (>2 times higher than the Age+Sex-based reference estimate) and absolute 10-year risk > 10%. **e** Top 5 highest attributed records for selected endpoints. Statistical measures were derived from 502,460 individuals. Source data are provided.

and gastrointestinal diseases as chronic hepatitis (Delta C-Index: -0.032 (-0.064 , -0.007)).

We also present a comparison between our approach and the Charlson Comorbidity Index's⁴² predictive performance, both of which can be automated. Additionally, we compare our method to the well-established ASCVD predictors, which are widely accessible but require an additional blood draw. Notably, incorporating the comorbidities from the Charlson Comorbidity Index enhances the discriminative capacity beyond age and sex; however, adding medical history proves to be significantly more effective in improving performance (Supplementary Fig. 3, Supplementary Data 5). Likewise, while supplementing ASCVD predictors to age and sex augments the performance for most endpoints, it remains inferior to the combination of age, sex, and medical history alone. Incorporating the medical history alongside the comorbidities or ASCVD predictors further improves the predictive performance for the vast majority of endpoints (AgeSex+Comorbidities augmented by the MedicalHistory: $+1726/1883$ (91.7%), ASCVD +MedicalHistory: $+1727/1883$ (91.7%)), demonstrating complementary nature of these information sources.

For illustration, we also present individual phenome-wide risk profiles (Fig. 3c, Supplementary Fig. 4a–b and 5a+b). The risk profiles varied substantially in the predispositions relative to the age and sex reference (the inner circle, see methods for details) and the absolute 10-year risk estimates (the outer circle). The first individual (Fig. 3c), a 60-year-old man, is predicted to be at a particularly high 10-year risk of metabolic, cardiovascular, respiratory, and genitourinary conditions, including diabetes mellitus (19.4%), heart failure (22%), COPD (14.9%), and chronic kidney disease (16.8%). Increased risk of neoplastic, dermatological, and musculoskeletal conditions was not predicted by the prior health records of this individual. In contrast, another individual, a 48-year-old woman (Supplementary Fig. 5b), is not estimated at increased cardiovascular risk but conversely to have almost 10x the risk for suicide ideation and attempt or self-harm compared to the reference group.

Importantly, the model performance is robust to the removal of recent information, indicating that the model effectively incorporates both the individuals' long-term medical history and recent interactions with the healthcare system in order to predict future disease onset (Supplementary Fig. 6). We provide Shapley attributions⁴³ for the most important records (Fig. 3d, Suppl. Figure 4c, Suppl. Figure 5c) and all records for the 24 highlighted endpoints (Supplementary Data 9) in the study population, enhancing the interpretability of our findings.

These findings indicate that health records contain substantial predictive information over established predictors for the majority of disease endpoints from across clinical specialties.

Predictive models can generalize across healthcare systems and populations

While our findings indicate potential utility in the UK Biobank, health records vary substantially across healthcare systems and over time due to differences in medical and coding practices (“distribution shift”) and underlying differences in the populations. Thus, predictive models can fail to learn robust and generalizable information^{44–46}.

To better understand the generalisability across different healthcare systems, we predicted risk states and absolute risk estimates for all individuals in the All of Us cohort with linked medical records ($N=229,830$; see Table 1). Importantly, we found significant improvements over the baseline model (age and biological sex only) for 1347 (85.9%) of the 1568 investigated endpoints with at least 100 incident events (Fig. 4a, Supplementary Data 8), replicating 1347/1500 (89.8%) of all significant improvements in the UK Biobank (Fig. 4b, Supplementary Data 8). Generally, larger improvements in the UK Biobank were replicated in the All of Us cohort. It's noteworthy that smaller improvements in the UK Biobank often corresponded to proportionately larger improvements in All of Us, while larger improvements in the UK Biobank were attenuated in All of Us (Fig. 4c).

As the risk states were largely derived from white, middle-aged, and generally affluent and healthy individuals from the UK, it was critical to validate the discriminative performance in diverse and historically underserved and underrepresented groups and ethnicities. Generally, we found comparable discriminative performances (Fig. 4d) and substantial benefits over basic demographic predictors (example of cardiac arrest in Fig. 4e) across all investigated groups.

To illustrate these improvements further, we replicated significant improvements for all of the 24 a priori selected endpoints, with improvements ranging from modest for hypertension (Delta C-Index: $+0.021$ (0.016, 0.024)) and Parkinson's disease (Delta C-Index: $+0.035$ (0.021, 0.05)) to substantial for, e.g., All-Cause Death (Delta C-Index: $+0.116$ (0.104, 0.127)), Pulmonary embolism (Delta C-Index: $+0.125$ (0.112, 0.137)), and Cardiac arrest (Delta C-Index: $+0.176$ (0.146, 0.206)) (Fig. 4f, g and Supplementary Data 8). Only for a subset of 54 (3.44%) significantly improved endpoints in the UK Biobank, the discriminative performance in All Of Us deteriorated significantly upon transferring the pre-trained medical history risk model and integrating the information beyond age and biological sex alone, including hepatitis (Delta C-Index: -0.226 (-0.251 , -0.2)), substance abuse (Delta C-Index: -0.037 (-0.05 , -0.026)) and osteoporosis (Delta C-Index: -0.015 (-0.021 , -0.008)).

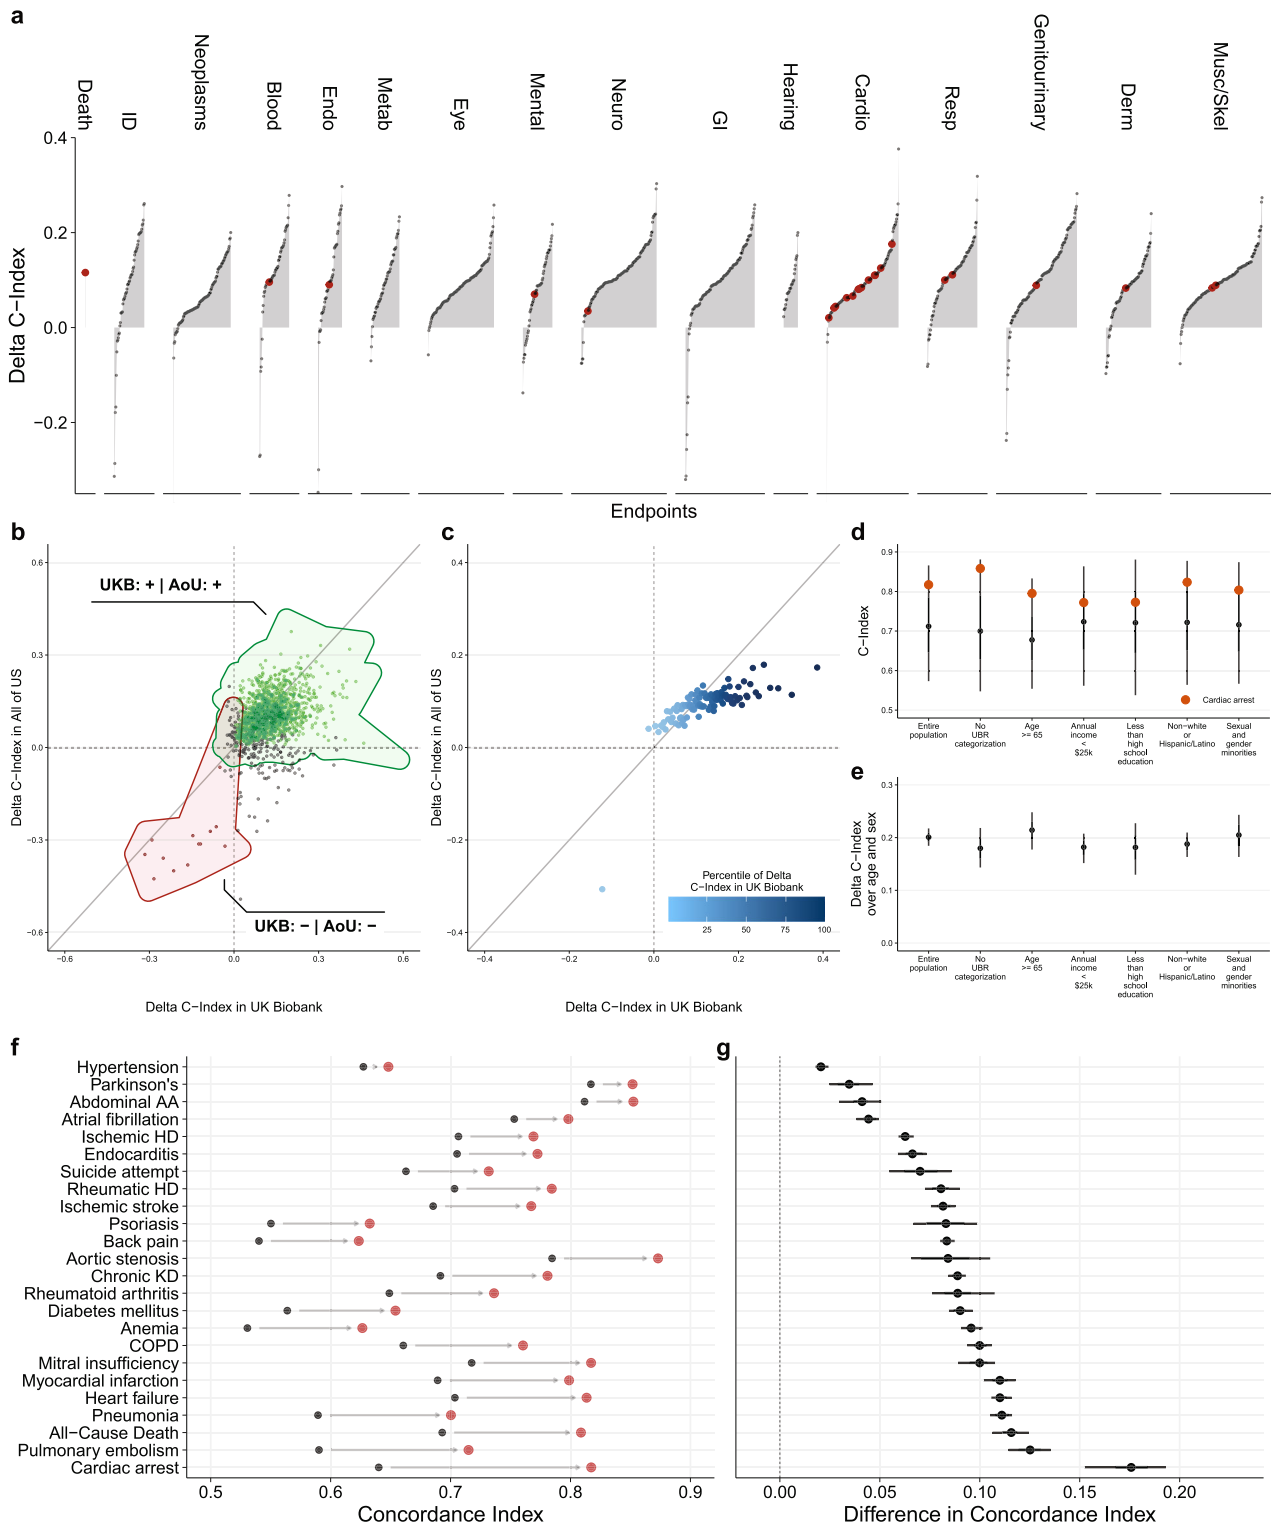
Taken together, our findings suggest that predictive models based on medical history can generalize across health systems and are robust to diverse populations.

Predictions can support cardiovascular disease prevention and the response to emerging health threats

While comprehensive phenome-wide risk profiles provide opportunities to guide medical decision-making, not all of the predictions are actionable. To illustrate the potential clinical utility, we focused on the primary prevention of cardiovascular disease and the response to newly emerging health threats at the example of COVID-19.

Risk scores are well established in the primary prevention of cardiovascular events and have been recommended to guide preventive lipid-lowering interventions⁴⁷. While cardiovascular predictors are accessible at a low cost, dedicated visits and resources from healthcare providers for physical and laboratory measurements are required. Therefore, we compared our phenome-wide risk score, based only on age, sex, and routine health records, to models based on established cardiovascular risk scores, the SCORE2⁴⁸, the ASCVD³, and the British QRISK3⁴ score. Interestingly, the discriminative performance of our phenome-wide model is competitive with the established cardiovascular risk scores for all investigated cardiovascular endpoints (Fig. 5a, Supplementary Data 7): we found comparable C-Indices with differences $+0.001$ (CI -0.002 , 0.005) for ischemic stroke, $+0.002$ (CI 0.002, 0.005) for ischemic heart disease and $+0.006$ (CI 0.003, 0.009) for myocardial infarction compared with the comprehensive QRISK3 score. It is noteworthy that these discriminative improvements are substantially better for later-stage diseases, including heart failure ($+0.018$ (CI 0.015, 0.021)), cardiac arrest ($+0.05$ (CI 0.042, 0.059)), and all-cause mortality ($+0.13$ (CI 0.128, 0.132)) when prior health records are considered.

To further illustrate potential utility, we look at newly emerging pathogenic health threats, where rapid and reliable risk stratification is required to protect high-risk groups and prioritize preventive interventions. We investigated how our phenome-wide risk states could have been used in the context of COVID-19, a respiratory infection with pneumonia and sepsis as common, life-threatening complications of severe cases. We repurposed the risk states for pneumonia, sepsis, and all-cause mortality to calculate a combined COVID-19 severity risk score using information available at the end of 2019 before the global spread of the COVID-19 pandemic (see Methods for details). The COVID-19 severity risk score resembles the risk for developing severe or fatal COVID-19 and illustrates how health records could have helped to identify individuals at high risk and to prioritize individuals in initial



vaccination campaigns better. Augmenting age with the COVID-19 severity risk score, we found substantially improved discriminative performance for both severe and fatal COVID-19 outcomes (Severe: C-Index (age) 0.597 (CI 0.591, 0.604) → C-Index (age + COVID-19 severity risk score) 0.647 (CI 0.641, 0.654); Fatal: C-Index (age) 0.720 (CI 0.710, 0.731) → C-Index (age + COVID-19 severity risk score) 0.780 (CI 0.772, 0.789). These discriminative improvements translate into higher cumulative incidence in the Top 5% population compared to age alone (Suppl. Figure 6C, age (left), COVID-19 severity score (right),

severe COVID-19 (top), fatal COVID-19 (bottom)): In the top 5% of the age-based risk group (-79 (IQR 77, 81) years old), 0.42% (CI 0.34%, 0.5%, $n = 105$) have been hospitalized, and 0.26% (CI 0.2%, 0.33%, $n = 66$) had died by the end of the first wave. By the end of the second wave, around 0.96% (CI 0.83%, 1.08%, $n = 240$) had been hospitalized, and 0.44% (0.36%, 0.52%, $n = 111$) had died. In contrast, for individuals in the top 5% of the COVID-19 severity risk score, by the end of the first wave, around 0.64% (CI 0.54%, 0.74%, $n = 160$) had been hospitalized, and 0.32% (0.25%, 0.39%, $n = 80$) had died, while by the end of the second

Fig. 4 | Predictive models can generalize across healthcare systems and populations. **a** External validation of the differences in discriminatory performance quantified by the C-Index between CPH models trained on age and biological sex and age, biological sex, and the risk state for 1,568 endpoints in the All of Us cohort. We find significant improvements over the baseline model (age and biological sex only) for 1,347 (85.9%) of the 1,568 investigated endpoints. **b** Direct comparison of the absolute C-Index in the UK Biobank (x-axis) and the All Of Us cohort (y-axis). Significant improvements can be replicated for 1,347 (89.8%, green points) of 1,500 endpoints in the All Of Us cohort. **c** Comparison of mean delta C-Index per delta percentile (derived from the UK Biobank from the 1,568 endpoints available in All Of Us). Improvements in the All Of Us cohort are consistent with the UK Biobank cohort: Small improvements in the UK Biobank tend to be larger in All Of Us, while large improvements in the UK Biobank tend to be attenuated in All Of Us. **d** Distribution of C-Indices for the 1,568 investigated endpoints

stratified by communities historically underrepresented in biomedical research (UPD)⁷³. Dots indicate medians and whiskers extend to the Bonferroni-corrected 95% confidence interval for a distribution bootstrapped over 100 iterations. **e** For the same groups, confidence intervals for the additive performance as measured by the C-Index compared to the baseline model. Dots indicate medians and whiskers extend to the Bonferroni-corrected 95% confidence interval for a distribution bootstrapped over 100 iterations. **f** Absolute discriminatory performance in terms of C-Index comparing the baseline (age and biological sex, black point) with the added routine health records risk state (red points) for a selection of 24 endpoints. **g** The differences in C-index for the same models. Statistical measures for UKB (in **b** and **c**) were derived from 502,460 individuals and for AoU (in **a–g**) were derived from 229,830 individuals. Dots indicate medians and whiskers extend to the Bonferroni-corrected 95% confidence interval for a distribution bootstrapped over 100 iterations. Source data are provided.

wave, 1.74% (CI 1.57%, 1.9%, $n = 436$) had been hospitalized and 0.68% (0.58%, 0.79%, $n = 172$) had died.

In summary, our findings illustrate the clinical utility of medical history for primary prevention of cardiovascular diseases and the rapid response to emerging health threats.

Discussion

Current clinical practice lacks systematic, data-driven guidance for individuals and care providers. Our study demonstrated that medical history can systematically inform on phenome-wide risk across clinical specialties, as shown in the British UK Biobank cohort. Subsequently, we show that these risk states can be repurposed to identify individuals vulnerable to severe COVID-19 and mortality. Importantly, we found significant improvements in the discriminated performance for the vast majority of disease endpoints, of which almost 90% could be replicated in the US All of US cohort. Our results indicated utility beyond age, sex, selected comorbidities, and established cardiovascular risk factors commonly considered in clinical practice for preventable diseases, treatable diseases, and diseases without existing risk stratification tools. We anticipate that our approach has the potential to facilitate population health at scale.

Designed for outpatient settings and focused on patients without acute complaints, our approach identifies incident disease onset from early (e.g., hypertension) and later (e.g., bypass surgery) health system contacts. We identified three primary scenarios of potential utility: Firstly, medical history can be exploited in diseases that are preventable with effective interventions, such as the prescription of lipid-lowering medication for primary prevention of coronary heart disease⁴⁷. Lowering LDL cholesterol in 10,000 individuals at increased risk by 2 mmol/L with atorvastatin 40 mg daily (-2€ per month) for 5 years is estimated to prevent 500 vascular events, reducing the individual relative risk by more than a third^{49,50}. Secondly, in conditions that are not preventable anymore individuals can benefit from early detection and treatment, like in type 2 diabetes or systolic heart failure. In individuals with heart failure with reduced ejection fraction, a comprehensive treatment regime (including ARNI, beta-blockers, MRA, and SGLT2 inhibitors) compared to a conventional regime (ACEi or ARB and beta blockers) reduced the hospital admissions for heart failure by more than two thirds, all-cause mortality by almost half⁵¹. For a 55-year-old male, this translated into an estimated 8.3 additional years free from cardiovascular death or readmission for heart failure. Lastly, in cases where outcomes are neither preventable nor treatable, estimates of prospective individual risk may be of high importance for personal decisions or the planning of advanced care, e.g., a high short-term mortality could identify patients needing to transition from curative to palliative strategies for optimal care^{52,53}. Multiple studies have shown that palliative care services can improve patients' symptoms and life quality and may even increase survival⁵⁴. Overall, our approach could facilitate the identification of high-risk populations for specific screening programs, potentially improving the value of national health programs.

Importantly, our approach, based on routine health records, shows large discriminative improvements for the majority of diseases compared with conventionally tested biomarkers^{55–57} and can generalize across diverse health systems, populations, and ethnicities. However, we also see that including the medical history over age and sex deteriorated the performance for a subset of 1.2% (UK Biobank) and 4.9% (All Of Us cohort), respectively. Three central challenges remain: First, health records, being products of interactions with the medical system, are subject to biological, procedural, and socio-economic biases⁵⁸, as well as being dependent on the evolving nature of medical knowledge and policies. Furthermore, certain measurements and laboratory values are often inaccessible at the point of care, and harmonization in and across health systems presents a significant barrier to implementation⁵⁹. Integrating these measures into the model holds considerable promise to improve the predictive performance further. While our approach is based on the standardized OMOP vocabulary, implementation requires a robust harmonization infrastructure, and data drift might necessitate model updates. Second, research cohorts often comprise healthier individuals with lower disease prevalence than the general population⁶⁰, potentially leading to underestimating absolute risks. While discriminative improvements provide evidence of the potential clinical utility, they are insufficient to prove it, as it is highly context-dependent on the population, the disease, and the interventions available. This is particularly relevant for very rare diseases, where screening the general population poses the risk of false positive findings. Future randomized implementation studies must investigate how this discriminatory information can translate into improved clinical outcomes in the respective target populations. The third challenge concerns ensuring the interpretability of our approach on such complex data. Our approach provided unique insights into how the model used patients' medical history to make risk predictions. The Shapley value attributions highlighted features the model found most informative for inference on both individual and population levels. These attributions are reflective of the model's decision-making process, and while they aligned with our clinical understanding, they should not replace clinician judgment or other forms of evidence. As we refine and deploy this approach, we must remain vigilant in evaluating its performance and understanding the interpretational limitations. Interestingly, the attributions also expose the challenges of implementing predictive models across primary care and clinical specialties. For example, statins and chest pain are among the most highly attributed records for a high future likelihood of developing heart disease, indicating that in some cases, prior healthcare providers have already considered or even acted upon a high suspected risk of the disease, without entering the actual diagnosis into the records. Consequently, employing the model for such patients, when low-density lipoprotein (LDL) cholesterol levels are already managed, may not lead to further preventive actions if the patient's care aligns with established standards. Importantly, we find that such cases do not drive the model's predictive performance by assessing the robustness of the model performance to the removal of recent information

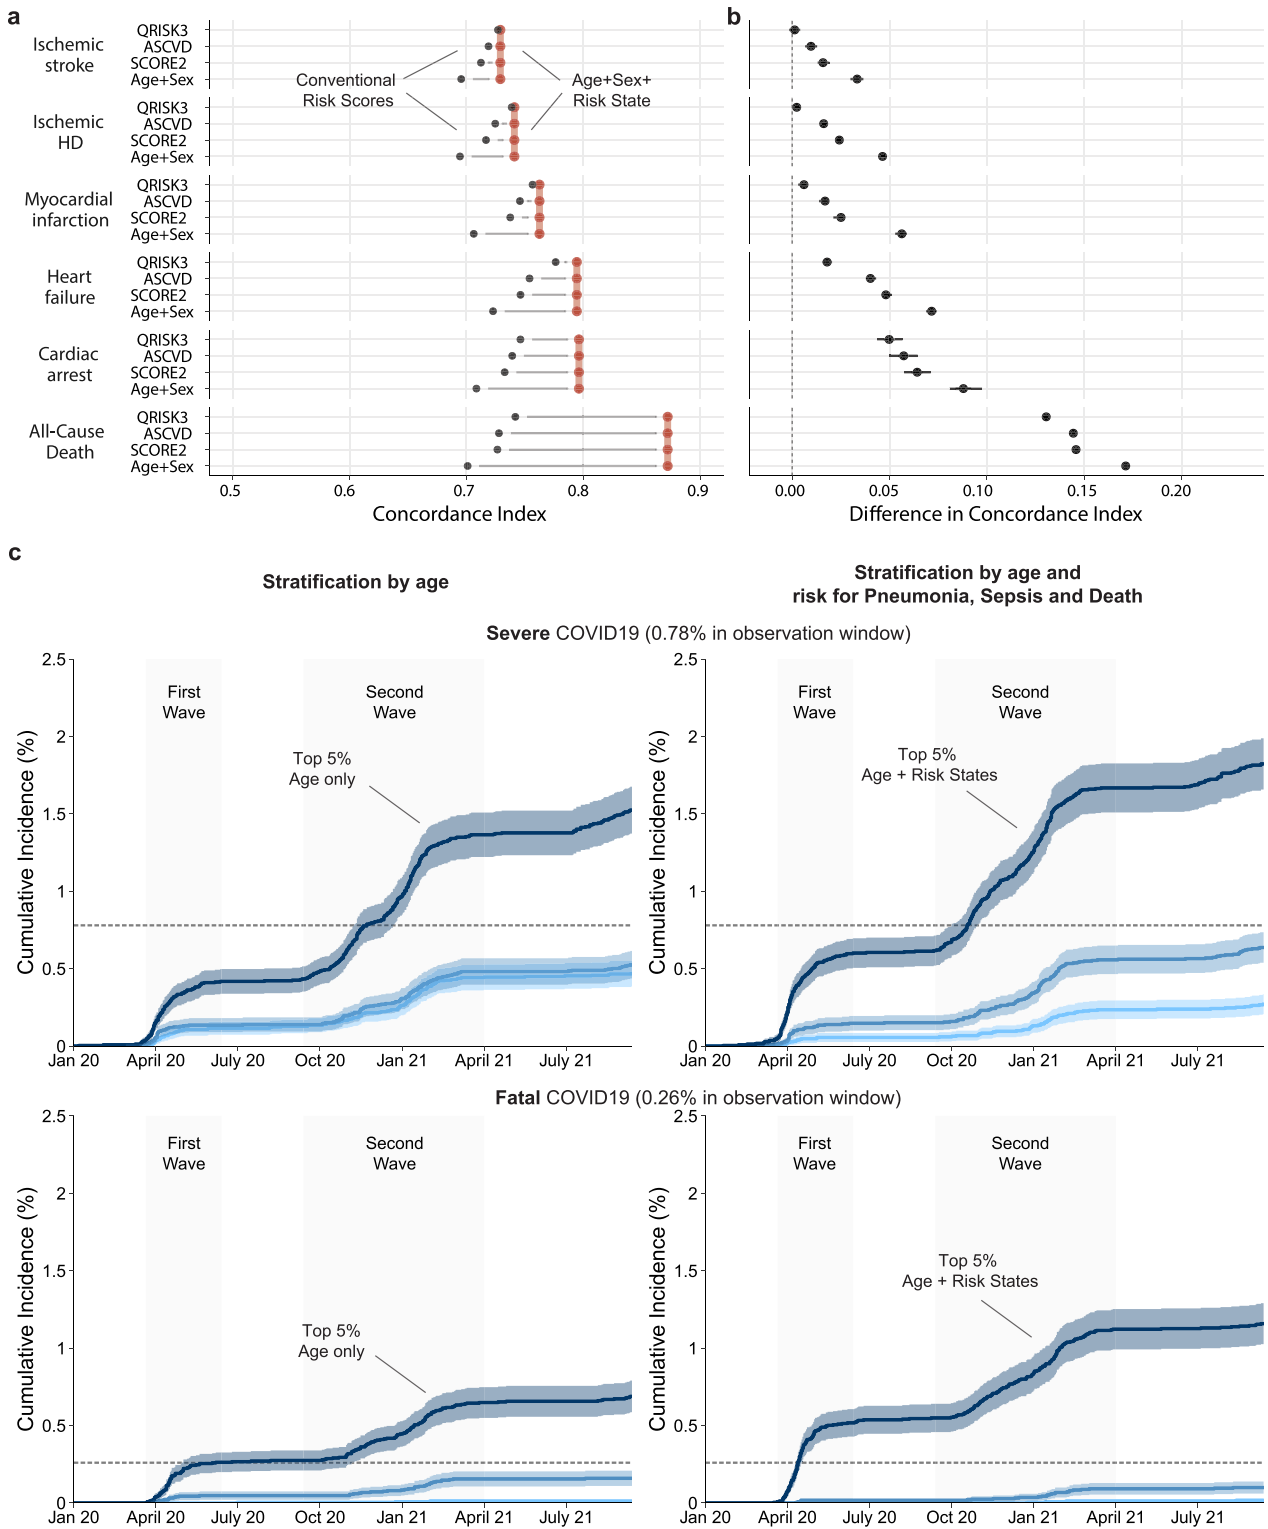


Fig. 5 | Predictions can support cardiovascular disease prevention and the response to emerging health threats. a Discriminatory performances in terms of absolute C-Indices comparing risk scores (Age+Sex, SCORE2, ASCVD, and QRISK as indicated, black point) with the risk model based on Age+Sex+RiskState (red segment). **b** Direct differences between risk scores (Age+Sex, SCORE2, ASCVD, and QRISK as indicated) and the risk model based on Age+Sex+RiskState in terms of C-index. Dots indicate medians and whiskers extend to the Bonferroni-corrected

95% confidence interval for a distribution bootstrapped over 100 iterations. **c** Estimated cumulative event trajectories, including 95% confidence intervals of severe (with hospitalization) and fatal (death registry) COVID-19 outcomes stratified by the Top, Median, and Bottom 5% based on age (left) or risk states of pneumonia, sepsis, and all-cause mortality as estimated by Kaplan-Meier analysis. Statistical measures were derived from 502,460 individuals. Source data are provided.

(Supplementary Fig. 6). Ultimately, if routine health records are to be used for risk prediction, robust governance rules to protect individuals, such as opt-out and usage reports, need to be implemented. With many national initiatives emerging to curate routine health records for millions of individuals in the general population, future studies will allow us to better understand how to overcome these challenges.

Our study presents a systematic approach to simultaneous risk stratification for thousands of diseases across clinical specialties based on readily available medical history. These risk states can then be used to rapidly respond to emerging health threats such as COVID-19. Our findings demonstrate the potential to link clinical practice with already collected data to inform and guide preventive interventions, early diagnosis, and treatment of disease.

Methods

Data source and definitions of predictors and endpoints

To derive risk states, we analyzed data from the UK Biobank cohort. Participants were enrolled from 2006 to 2010 in 22 recruitment centers across England, Scotland, and Wales; the follow-up is ongoing, and records until the 24th of September 2021 are included in this analysis. The UK Biobank cohort comprises 273,353 women and 229,107 men aged between 37–73 years at the time of their assessment visit. Participants are linked to routinely collected records from primary care (GP), hospital records (HES, PEDW, and SMR), and death registries (ONS), providing longitudinal information on diagnosis, procedures, and prescriptions for the entire cohort from Scotland, Wales, and England. Routine health records were mapped to the OMOP CDM and represented as a 71,036-dimensional binary vector, indicating whether a concept has been recorded at least once in an individual prior to recruitment. A subset of 15,595 unique concepts, all found in at least 50 individuals, was chosen for model development. Endpoints were defined as the set of PheCodes X^{39,61}, and after the exclusion of very rare endpoints (recorded in <100 individuals), 1883 PheCodes X endpoints were included in the development of the models. Due to the adult population, congenital, developmental, and neonatal endpoints were excluded. For each endpoint, subsequently, time-to-event outcomes were extracted, defined by the first occurrence after recruitment in primary care, hospital, or death records. Detailed information on the predictors and endpoints is provided in Supplementary Data 1–2.

While all individuals in the UK Biobank were used to integrate the routine health records, develop the model, and estimate phenome-wide log partial hazards, individuals were excluded from endpoint-specific downstream analysis if they were already diagnosed with a disease (defined by a prior record of the respective endpoint) or are generally not eligible for the specific endpoint (females were excluded from the risk estimation for prostate cancer).

To externally validate our risk states, we investigate individuals from the All of Us cohort³⁷, containing information on 229,830 individuals of diverse descent and from minorities historically underrepresented in biomedical research⁴⁰. Because we only use the All of Us cohort for validation, we evaluate the predictive performance for the subset of 1568 endpoints with at least 100 incident events in the All of Us cohort.

The study adhered to the TRIPOD (Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis) statement for reporting⁶². The completed checklist can be found in the Supplementary Information.

Extraction and preparation of the routine health records

To extract the routine health records of each individual, we first aggregated the linked primary care, hospital records, and mortality records and mapped the aggregated records to the OMOP CDM (mostly SNOMED and RxNorm). Specifically, we used mapping tables provided by the UK Biobank, the OHDSI community, and SNOMED

International to map concepts from the provider and country-specific non-standard vocabularies to OMOP standard vocabularies.

We restricted the analysis to the domains “Observation”, “Condition”, “Procedure”, “Drug” and “Device”. To reduce the complexity, we did not include any laboratory measures. The PheCode X endpoints^{39,61} were derived from either mapping directly from ICD-10 (hospital and death records) or mapping from SNOMED to ICD-10 (using the official mapping table) and subsequently to Phecodes X.

To ensure the accuracy and integrity of our data, we implemented multiple validation steps. After each stage in the extraction and mapping process, we conducted plausibility and sanity checks on the distribution of the mapped records, along with spot checks of individual records. This approach was critical in verifying the validity of the data. Additionally, post-model training, the data underwent further verification. This included analyzing the calculated record attributions and removing recent records, as detailed in Supplementary Fig. 6. These steps were essential to identify and mitigate any potential issues of record leakage. In the accompanying code release, we have provided the exact code used to extract and prepare the health records.

Spatial validation and data preprocessing

For model development and testing, we split the data set into 22 spatially separated partitions based on the location of the assessment center at recruitment. We analyzed the data in 22-fold nested cross-validation, setting aside one of the spatially separated partitions as a test set, aggregating the remaining partitions, and randomly selecting 10% of the aggregated data for the validation set. Within each of the 22 cross-validation loops, the individual test set (i.e., the spatially separated partition) remained untouched throughout model development, and the validation set was used to validate the fitting progress and checkpoint selection. All 22 obtained models were then evaluated on their respective test sets. We assumed missing data occurred randomly and performed multiple imputations using chained equations with gradient boosting machines^{63,64}. Imputation models were fitted on the training sets and applied to the respective validation and test sets. Continuous variables were standardized; Categorical variables were one-hot encoded.

Development of the phenome-wide risk model

The risk model is a multi-task neural network that uses the binary representations of an individual's prior health records before recruitment to simultaneously predict log partial hazards⁶⁵ for a set of 1883 endpoints. The model consists of three fully connected linear layers with 4096 hidden units, each with layer normalisation⁶⁶, dropout⁶⁷, and leaky ReLU activations. The last latent representation serves as a regulariser as it incentivises the extraction of robust features for multiple diseases. For comparison, we also benchmarked the linear version of our model with 29.4 M instead of 88.4 M parameters (see Suppl. Figure 2). The model subsequently computes the log partial hazard (the risk state) for each endpoint with an adapted proportional hazard loss⁶⁵, resulting in a 1,883-dimensional output representation. The individual losses are averaged and then summed to derive the final loss of the model. We subsequently tuned hyperparameters (via Bayesian Optimization) on train and validation splits over a constrained parameter space, tuning batch size, learning rate, weight decay, number of nodes in the layers of the endpoint heads, number of hidden layers, dropout rates, and size of the output vector of the shared network. The final models were trained with batch size 512 using the Adam optimiser⁶⁸ with a learning rate of 0.0006 and weight decay of 0.3, and early stopping tracking of the performance on the validation set. We implemented the model in Python 3.9 using PyTorch 1.11⁶⁹ and PyTorch-lightning 1.5.5 (for code availability, see below). The training of a single model on an NVIDIA A100 GPU node for 18 epochs required approximately 11 hours, equating to the emission of approximately 1.08 kg CO₂ eq, 4.36 km driven by an average ICE car or 0.54 kgs of coal

burned as calculated by the mlco2 calculator⁷⁰. The external validation of these models, conducted within the All of Us cloud computing environment and including data preprocessing, inference, and evaluation, incurred a total compute cost of approximately 150 USD.

Downstream analysis and performance comparisons

We fitted Cox proportional hazards models³⁶ (CPH) to derive absolute risk predictions from the endpoint-specific risk states for the individual endpoints. For each endpoint, we developed models with distinct covariate sets: for all endpoints, we investigated age, biological sex, and the risk states from the health records. For cardiovascular endpoints, we additionally investigated predictors from established and guideline-recommended scores for the primary prevention of cardiovascular diseases, the SCORE2, ASCVD, and QRISK3. Model development was repeated independently for each assessment center thus, for each cross-validation split, models were trained on the respective train set, and checkpoints were selected on the respective validation set. For the final evaluation, test set predictions from the spatially separate recruitment centers were aggregated. Event risk rates were calculated over the full observation period. Harrell's C-Index⁷¹ was calculated with the lifelines package⁷² by bootstrapping both the aggregated test set and individual assessment centers within ten years after recruitment to control for right-censoring. The C-Index is a measure of rank correlation that quantifies the agreement between predicted and observed outcomes. It ranges between 0.5 (no better than random prediction) to 1 (perfect prediction). Statistical inferences about model differences were based on the distribution of bootstrapped differences in the C-Index; models were considered different whenever the Bonferroni-corrected 95% CI of the difference did not overlap cross zero, to account for multiple testing. CPH models were fitted with the CoxPHFitter from the Python package lifelines⁷² with default parameters and a step size of 0.5, 0.1, or 0.01 to facilitate model convergence. Confidence intervals for all statistical analyses were calculated over 1000 bootstrapping iterations.

Response to emerging health threats

We retrained our models using data limited to records until the end of December 2019, keeping the setting (in particular time zero for training) unchanged. Using these updated models, we then predicted the risk states using all data available at the end of 2019, just as the first cases of COVID-19 were reported. We then manually selected specific risk states associated with pneumonia, sepsis, and all-cause mortality to create an unweighted COVID-19 severity risk score. This risk score was subsequently tested against age for the identification of incident severe and fatal COVID-19 cases.

Independent validation in the All Of Us cohort

After mapping the linked health records from All Of Us to the OMOP vocabulary, we transferred the neural networks developed in the UK Biobank to the All Of Us research environment. We then used the models to predict the disease-specific risk states for all individuals. Subsequently, we predicted absolute risks with the CPH models developed in the UK Biobank. Finally, we calculated the mean of the predictions from the models for each individual and disease. For baseline comparison with Age and Sex, we fitted new CPH models in the All Of Us cohort.

Calculation of record attributions

To determine which records are most important on an individual level, we calculated attributions for the selection of 24 endpoints based on Shapley values. For computational efficiency, we approximated Shapley values via sampling for only 17,236 individuals unseen to the model during development⁴³. Please refer to Supplementary Data 9 for the aggregated attributions from individuals without prior events. Shapley values in the table are provided in two forms: averaged (so-

called local attributions to quantify importance for affected individuals) and summed (global attributions to quantify importance for population ranking). The average Shapley attributions, presented in the main text and figures, closely reflect our understanding of importance for affected individuals.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

UK Biobank data, including all linked routine health records, are publicly available to bona fide researchers upon application at <http://www.ukbiobank.ac.uk/using-the-resource/>. In this study, primary care data was used following the COPI regulations. The All Of Us cohort data were provided by the All Of Us Research Program by permission that can be sought by scientists and the public alike. Currently, however, data access requires affiliation with a US institution. All patient data used throughout this study has been subject to patient consent as covered by the UK Biobank and All Of Us. Detailed information on the predictors and endpoints is presented in Supplementary Data 1-3. Source data are provided with this paper.

Code availability

All code developed and used throughout this study has been made open source and is available on GitHub. The code to train the medical history model can be found here: github.com/nebw/medhist, while the code to run analysis on trained models can be found here: github.com/JakobSteinfeldt/MedicalHistoryPhenomeWide.

References

1. Sindi, S. et al. The CAIDE Dementia Risk Score App: The development of an evidence-based mobile application to predict the risk of dementia. *Alzheimers Dement* **1**, 328–333 (2015).
2. Lindström, J. & Tuomilehto, J. The diabetes risk score: a practical tool to predict type 2 diabetes risk. *Diab. Care* **26**, 725–731 (2003).
3. Goff, D. C. et al. 2013 ACC/AHA Guideline on the Assessment of Cardiovascular Risk. *Circulation* **129**, S49–S73 (2014).
4. Hippisley-Cox, J., Coupland, C. & Brindle, P. Development and validation of QRISK3 risk prediction algorithms to estimate future risk of cardiovascular disease: prospective cohort study. *BMJ* **357**, j2099 (2017).
5. Steyerberg, E. W. et al. Prognosis Research Strategy (PROGRESS) 3: prognostic model research. *PLoS Med* **10**, e1001381 (2013).
6. Hampton, J. R., Harrison, M. J., Mitchell, J. R., Prichard, J. S. & Seymour, C. Relative contributions of history-taking, physical examination, and laboratory investigation to diagnosis and management of medical outpatients. *Br. Med J.* **2**, 486–489 (1975).
7. Danish eHealth Portal. Danish eHealth Portal. Danish eHealth Portal. 2001. <https://www.sundhed.dk/borger/service/om-sundheddk/om-organisationen/ehealth-in-denmark/background/>.
8. e-Health Record. e-Health Record. 2005. <https://e-estonia.com/solutions/healthcare/e-health-records/>.
9. Clalit Research Institute. Clalit Health Services. Clalit Health Services. 2010. <http://clalitresearch.org/about-us/our-data/> (accessed 2010).
10. National Electronic Health Record. National Electronic Health Record. National Electronic Health Record. 2011. <https://www.ihis.com.sg/nehr/about-nehr>.
11. My Health Record. My Health Record. My Health Record. 2016. <https://www.myhealthrecord.gov.au/>.
12. Wood, A. et al. Linked electronic health records for research on a nationwide cohort of more than 54 million people in England: data resource. *BMJ* **373**, n826 (2021).

13. Rush, R. Taking Note. *N. Engl. J. Med* **381**, 9 (2019).
14. Tsang, G., Zhou, S.-M. & Xie, X. Modeling Large Sparse Data for Feature Selection: Hospital Admission Predictions of the Dementia Patients Using Primary Care Electronic Health Records. *IEEE J. Transl. Eng. Health Med* **9**, 3000113 (2021).
15. Langham J. et al. Predicting risk of dementia with machine learning and survival models using routine primary care records. In: 2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). 2021: 3036–3042.
16. Rajkomar, A. et al. Scalable and accurate deep learning with electronic health records. *npj Digital Med.* **1**, 1–10 (2018).
17. Appelbaum, L. et al. Development and validation of a pancreatic cancer risk model for the general population using electronic health records: An observational study. *Eur. J. Cancer* **143**, 19–30 (2021).
18. Kronzer, V. L. et al. Investigating the impact of disease and health record duration on the eMERGE algorithm for rheumatoid arthritis. *J. Am. Med. Assoc.* **327**, 601–605 (2020).
19. Sekelj, S. et al. Detecting undiagnosed atrial fibrillation in UK primary care: Validation of a machine learning prediction algorithm in a retrospective cohort study. *Eur. J. Prev. Cardiol.* **28**, 598–605 (2021).
20. Miotto, R., Li, L., Kidd, B. A. & Dudley, J. T. Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. *Sci. Rep.* **6**, 26094 (2016).
21. Estiri, H. et al. Predicting COVID-19 mortality with electronic medical records. *NPJ Digit Med* **4**, 15 (2021).
22. Wu J., Nadarajah R., Raveendra K., Cowan J. C., & Gale C. P. FIND-AF: a widely applicable artificial intelligence algorithm to target systematic screening for atrial fibrillation in older individuals through primary care electronic health records. *Europace* 2022; **24**. <https://doi.org/10.1093/europace/euac053.565>.
23. Bagheri A. et al. Multimodal Learning for Cardiovascular Risk Prediction using EHR Data. arXiv [cs.LG]. 2020; published online Aug 27. <http://arxiv.org/abs/2008.11979>.
24. Ben Miled, Z. et al. Predicting dementia with routine care EMR data. *Artif. Intell. Med* **102**, 101771 (2020).
25. Zhao, J. et al. Learning from Longitudinal Data in Electronic Health Record and Genetic Data to Improve Cardiovascular Event Prediction. *Sci. Rep.* **9**, 717 (2019).
26. Jin, B. et al. Predicting the Risk of Heart Failure With EHR Sequential Data Modeling. *IEEE Access* **6**, 9256–9261 (2018).
27. Hill, N. R. et al. Predicting atrial fibrillation in primary care using machine learning. *PLoS One* **14**, e0224582 (2019).
28. Tiwari, P. et al. Assessment of a Machine Learning Model Applied to Harmonized Electronic Health Record Data for the Prediction of Incident Atrial Fibrillation. *JAMA Netw. Open* **3**, e1919396 (2020).
29. Denny, J. C. et al. Systematic comparison of phenome-wide association study of electronic medical record data and genome-wide association study data. *Nat. Biotechnol.* **31**, 1102–1110 (2013).
30. Bush, W. S., Oetjens, M. T. & Crawford, D. C. Unravelling the human genome-phenome relationship using phenome-wide association studies. *Nat. Rev. Genet* **17**, 129–145 (2016).
31. Zhang, Y. et al. High-throughput phenotyping with electronic medical record data using a common semi-supervised approach (PheCAP). *Nat. Protoc.* **14**, 3426–3444 (2019).
32. Zheng, N. S. et al. PheMap: a multi-resource knowledge base for high-throughput phenotyping within electronic health records. *J. Am. Med. Assoc.* **327**, 1675–1687 (2020).
33. Rasmy, L., Xiang, Y., Xie, Z., Tao, C. & Zhi, D. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *npj Digital Med.* **4**, 1–13 (2021).
34. Li, Y. et al. BEHRT: Transformer for Electronic Health Records. *Sci. Rep.* **10**, 1–12 (2020).
35. Bycroft, C. et al. The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**, 203–209 (2018).
36. Cox, D. R. Regression Models and Life-Tables. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **34**, 187–202 (1972).
37. All of Us Research Program Investigators, Denny, J. C. et al. The ‘All of Us’ Research Program. *N. Engl. J. Med* **381**, 668–676 (2019).
38. Sudlow, C. et al. 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**, e1001779 (2015).
39. Wu, P. et al. Mapping ICD-10 and ICD-10-CM Codes to Phecodes: Workflow Development and Initial Evaluation. *JMIR Med Inf.* **7**, e14325 (2019).
40. Ramirez, A. H. et al. The All of Us Research Program: Data quality, utility, and diversity. *Patterns (N. Y)* **3**, 100570 (2022).
41. Martinez, F. J. et al. A New Approach for Identifying Patients with Undiagnosed Chronic Obstructive Pulmonary Disease. *Am. J. Respir. Crit. Care Med* **195**, 748–756 (2017).
42. Charlson, M. E., Pompei, P., Ales, K. L. & MacKenzie, C. R. A new method of classifying prognostic comorbidity in longitudinal studies: development and validation. *J. Chronic Dis.* **40**, 373–383 (1987).
43. Castro, J., Gómez, D. & Tejada, J. Polynomial calculation of the Shapley value based on sampling. *Comput Oper. Res* **36**, 1726–1730 (2009).
44. Finlayson, S. G. et al. The Clinician and Dataset Shift in Artificial Intelligence. *N. Engl. J. Med* **385**, 283–286 (2021).
45. Wong, A. et al. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Intern Med* **181**, 1065–1070 (2021).
46. Guo, L. L. et al. Evaluation of domain generalization and adaptation on improving model robustness to temporal dataset shift in clinical medicine. *Sci. Rep.* **12**, 2726 (2022).
47. National Institute for health and Care Excellence (NICE). Cardiovascular disease: risk assessment and reduction, including lipid modification. 2014; published online July 18. <https://www.nice.org.uk/guidance/cg181> (accessed Sept 16, 2022).
48. SCORE2 working group and ESC Cardiovascular risk collaboration. SCORE2 risk prediction algorithms: new models to estimate 10-year risk of cardiovascular disease in Europe. *Eur. Heart J.* **42**, 2439–2454 (2021).
49. Collins, R. et al. Interpretation of the evidence for the efficacy and safety of statin therapy. *Lancet* **388**, 2532–2561 (2016).
50. Ference, B. A. et al. Low-density lipoproteins cause atherosclerotic cardiovascular disease. 1. Evidence from genetic, epidemiologic, and clinical studies. A consensus statement from the European Atherosclerosis Society Consensus Panel. *Eur. Heart J.* **38**, 2459–2472 (2017).
51. Vaduganathan, M. et al. Estimating lifetime benefits of comprehensive disease-modifying pharmacological therapies in patients with heart failure with reduced ejection fraction: a comparative analysis of three randomised controlled trials. *Lancet* **396**, 121–128 (2020).
52. Adelson, K. et al. Standardized Criteria for Palliative Care Consultation on a Solid Tumor Oncology Service Reduces Downstream Health Care Use. *J. Oncol. Pr.* **13**, e431–e440 (2017).
53. Weissman, D. E. & Meier, D. E. Identifying patients in need of a palliative care assessment in the hospital setting: a consensus report from the Center to Advance Palliative Care. *J. Palliat. Med* **14**, 17–23 (2011).
54. Centeno, C. & Arias-Casais, N. Global palliative care: from need to action. *Lancet Glob. Health* **7**, e815–e816 (2019).
55. de Lemos, J. A. et al. Multimodality Strategy for Cardiovascular Risk Assessment: Performance in 2 Population-Based Cohorts. *Circulation* **135**, 2119–2132 (2017).
56. Steinfeldt, J. et al. Neural network-based integration of polygenic and clinical information: development and validation of a

- prediction model for 10-year risk of major adverse cardiac events in the UK Biobank cohort. *Lancet Digit Health* **4**, e84–e94 (2022).
57. Buerger T., et al. Metabolomic profiles predict individual multi-disease outcomes. *Nat. Med.* <https://doi.org/10.1038/s41591-022-01980-3> 2022.
 58. Vayena, E. Value from health data: European opportunity to catalyse progress in digital health. *Lancet* **397**, 652–653 (2021).
 59. Denaxas, S. et al. A semi-supervised approach for rapidly creating clinical biomarker phenotypes in the UK Biobank using different primary care EHR and clinical terminology systems. *JAMIA Open* **3**, 545–556 (2020).
 60. Fry, A. et al. Comparison of Sociodemographic and Health-Related Characteristics of UK Biobank Participants With Those of the General Population. *Am. J. Epidemiol.* **186**, 1026–1034 (2017).
 61. Wei, W.-Q. et al. Evaluating phecodes, clinical classification software, and ICD-9-CM codes for phenome-wide association studies in the electronic health record. *PLoS One* **12**, e0175508 (2017).
 62. Moons, K. G. M. et al. Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis (TRIPOD): explanation and elaboration. *Ann. Intern. Med.* **162**, W1–W73 (2015).
 63. Stekhoven, D. J. & Bühlmann, P. MissForest-non-parametric missing value imputation for mixed-type data. *Bioinformatics* **28**, 112–118 (2012).
 64. miceforest. PyPI. <https://pypi.org/project/miceforest/> (accessed July 6, 2022).
 65. Katzman J. L. et al. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC. Med. Res. Methodol.* **18**, 24 (2018).
 66. Ba J. L., Kiros J. R. & Hinton G. E. Layer Normalization. arXiv [stat.ML]. 2016; published online July 21. <http://arxiv.org/abs/1607.06450>.
 67. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn Res* **15**, 1929–1958 (2014).
 68. Kingma D. P. & Ba J. L. Adam: a Method for stochastic optimization. In *International Conference on Learning Representations 2015* (ICLR, 2015).
 69. Paszke, A. et al. Automatic differentiation in PyTorch. *Adv. Neural Inf. Process. Syst.* **30**, 1–4 (2017).
 70. Machine Learning CO2 impact calculator. <https://mlco2.github.io/impact/> (accessed May 10, 2023).
 71. Harrell, F. E. et al. Evaluating the yield of medical tests. *JAMA* **247**, 2543–2546 (1982).
 72. lifelines 0.25.8. 2021. <https://lifelines.readthedocs.io/en/latest/> (accessed Feb 3, 2021).
 73. How does All of Us assess diversity? What communities does All of Us consider ‘underrepresented in biomedical research?’ <https://www.researchallofus.org/faq/how-does-all-of-us-assess-diversity-what-communities-does-all-of-us-consider-underrepresented-in-biomedical-research/> (accessed May 5, 2023).
- Data Science Centre led by HDR UK (grant SP/19/3/34678), b) BigData@Heart Consortium, funded by the Innovative Medicines Initiative-2 Joint Undertaking under grant agreement 116074, c) the NIHR Biomedical Research Centre at University College London Hospital NHS Trust (UCLH BRC), d) a BHF Accelerator Award (AA/18/6/24223), e) the CVD-COVID-UK/COVID-IMPACT consortium and f) the Multimorbidity Mechanism and Therapeutic Research Collaborative (MMTRC, grant number MR/V033867/1). HH is supported by Health Data Research UK and the National Institute for Health Research, Biomedical Research Centre at University College London Hospitals.

Author contributions

J.S., B.W., T.B., M.P., H.H., C.L., U.L., J.D., and R.E. conceived and designed the project. J.S., B.W., and T.B. implemented models, conducted experiments, and performed data analysis. J.U. and A.V. supported the analysis. S.H. performed the external validation. M.P., S.D., H.H., and C.L. provided methodological support and contributed to the discussion of the results. J.S., B.W., T.B., U.L., J.D., and R.E. wrote and prepared the manuscript. All authors read, revised, and approved the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

U.L. received research grants to the institution from Abbott, Amgen, Bayer and Novartis. J.D. received honoraria from Amgen, Boehringer Ingelheim, Merck, Pfizer, Aegerion, Novartis, Sanofi, Takeda, Novo Nordisk, Bayer, and is a Trustee of Our Future Health. R.E. received honoraria from Sanofi and consulting fees from Boehringer Ingelheim. All other authors do declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-024-48568-8>.

Correspondence and requests for materials should be addressed to Roland Eils.

Peer review information *Nature Communications* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024

Acknowledgements

We would like to acknowledge the support of the UK Biobank and the All of Us Research Program in providing access to their respective datasets. This research has been conducted using data from the UK Biobank (application number 51157) and the All of Us Research Program (by S.H. UserID 5703). Both studies have received ethical approval from their respective institutional review boards and have obtained informed consent from participants. We are grateful to the participants who generously contributed their time and data to make this research possible. This project has been funded by the Charité - Universitätsmedizin Berlin and the Einstein Foundation Berlin through the Einstein BIH Visiting Fellowship awarded to J.D. The study has been supported by the BMBF-funded Medical Informatics Initiative (HiGHmed, 01ZZ1802A – 01ZZ1802Z) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 437531118 – SFB 1470. SD is supported by a) the BHF

¹Department of Cardiology, Angiology and Intensive Care Medicine, Deutsches Herzzentrum der Charité (DHZC), Berlin, Germany. ²Charité – Universitätsmedizin Berlin, corporate member of Freie Universität Berlin and Humboldt Universität zu Berlin, Klinik/Centrum, Charitéplatz 1, 10117 Berlin, Germany. ³Computational Medicine, Berlin Institute of Health (BIH), Charite - University Medicine Berlin, Berlin, Germany. ⁴Friede Springer Cardiovascular Prevention Center@Charite, Charite - University Medicine Berlin, Berlin, Germany. ⁵Institute of Cardiovascular Sciences, University College London, London, UK. ⁶Center for Digital Health, Berlin Institute of Health (BIH), Charite - University Medicine Berlin, Berlin, Germany. ⁷MRC Epidemiology Unit, Institute of Metabolic Science, University of Cambridge, Cambridge, UK. ⁸Precision Health University Research Institute, Queen Mary University of London and Barts NHS Trust, London, UK. ⁹Institute of Health Informatics, University College London, London, UK. ¹⁰Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Massachusetts, USA. ¹¹Pattern Recognition and Image Analysis Lab, University of Münster, Münster, Germany. ¹²British Heart Foundation Data Science Centre, London, UK. ¹³Health Data Research UK, London, UK. ¹⁴National Institute for Health Research, Biomedical Research Centre at University College London Hospitals National Institute for Health Research, Biomedical Research Centre, London, UK. ¹⁵Berlin Institute of Health (BIH), Charite - University Medicine Berlin, Berlin, Germany. ¹⁶DZHK (German Centre for Cardiovascular Research), Partner Site Berlin, Berlin Berlin, Germany. ¹⁷Health Data Science Unit, Heidelberg University Hospital and BioQuant, Heidelberg, Germany. ¹⁸These authors contributed equally: Jakob Steinfeldt, Benjamin Wild, Thore Buerger. ¹⁹These authors jointly supervised this work: Ulf Landmesser, John Deanfield, Roland Eils. ✉ e-mail: roland.eils@charite.de