

1 **TITLE**

2 **Speos: An ensemble graph representation learning framework to predict core genes**
3 **for complex diseases**

4
5
6

7 **AUTHORS**

8 Florin Ratajczak¹, Mitchell Joblin², Marcel Hildebrandt³, Martin Ringsquandl³, Pascal Falter-
9 Braun^{1,4}✉, Matthias Heinig^{5,6,7}✉

10
11

12 **AFFILIATIONS**

13 ¹*Institute of Network Biology (INET), Molecular Targets and Therapeutics Center (MTTC), Helmholtz*
14 *Center Munich, German Research Center for Environmental Health, Munich-Neuherberg, Germany*

15 ²*Amazon, Seattle, USA*

16 ³*Siemens Technology, Siemens AG, Munich, Germany.*

17 ⁴*Microbe-Host Interactions, Faculty of Biology, Ludwig-Maximilians-Universität (LMU) München,*
18 *Planegg-Martinsried, Germany.*

19 ⁵*Institute of Computational Biology (ICB), Helmholtz Center Munich, German Research Center for*
20 *Environmental Health, Munich-Neuherberg, Germany*

21 ⁶*Department of Computer Science, TUM School of Computation, Information and Technology,*
22 *Technical University of Munich, Garching, Germany*

23 ⁷*German Centre for Cardiovascular Research (DZHK), Munich Heart Association, Partner Site*
24 *Munich, Berlin, Germany*

25

26 ✉email: matthias.heinig@helmholtz-muenchen.de; pascal.falter-braun@helmholtz-muenchen.de

27 **Correspondence** should be addressed to Matthias Heinig, Pascal-Falter-Braun

28
29

30 **KEYWORDS**

31 Graph representation-learning, interactome, regulatory networks, deep-learning, human disease
32 genes, GWAS, core genes, omnigenic model

33
34
35

36 **ABSTRACT**

37 Understanding phenotype-to-genotype relationships is a grand challenge of 21st century biology with
38 translational implications. The recently proposed “omnigenic” model postulates that effects of genetic
39 variation on traits are mediated by *core*-genes and -proteins whose activities mechanistically
40 influence the phenotype, whereas *peripheral* genes encode a regulatory network that indirectly
41 affects phenotypes via core gene products. We have developed a positive-unlabeled graph
42 representation-learning ensemble-approach to predict core genes for diverse diseases using
43 Mendelian disorder genes for training. Employing mouse knockout phenotypes for external
44 validation, we demonstrate that our most confident predictions validate at rates on par with the
45 Mendelian disorder genes, and all candidates exhibit core-gene properties like transcriptional
46 deregulation in diseases and loss-of-function intolerance. Predicted candidates are enriched for drug
47 targets and druggable proteins and, in contrast to Mendelian disorder genes, also for druggable but
48 yet untargeted gene products. Model interpretation suggests key molecular mechanisms and
49 physical interactions for core gene predictions. Our results demonstrate the potential of graph
50 representation learning and pave the way for studying core gene properties and future drug
51 development.

52

53 MAIN

54 Understanding phenotype-to-genotype relationships is one of the most fundamental problems of
55 current biological research with profound translational implications for questions ranging from human
56 healthcare to biotechnological crop improvement. Genome-wide association studies (GWAS)
57 statistically associate phenotypes with specific variants in genomic loci. This approach has been
58 immensely successful and led to the identification of thousands of variants affecting diverse
59 physiological, molecular, and even psychological phenotypes. The problem of identifying likely
60 causal variants within haplotype blocks is increasingly solved by advanced modeling approaches
61 that integrate GWAS and functional genomic data to identify the genetic variants and genes that are
62 likely causal for the observed phenotypic manifestation^{1,2}. However, after solving this issue recent
63 analyses still indicate that even simple traits can have thousands of causal variants³ distributed
64 uniformly across the genome, and many without obvious connection to the known molecular
65 mechanisms regulating the respective trait⁴⁻⁹. This insight raises the conceptual question which
66 molecular mechanisms could give rise to such a highly polygenic architecture and the practical
67 question about how to prioritize proteins as interventional and diagnostic targets. The recently
68 proposed “omnigenic” model postulates that the effects of genetic variation on a trait are mediated
69 by *core* genes, encoding *core* proteins (hereafter used interchangeably depending on context),
70 whose expression, and ultimately function, directly and mechanistically influences the phenotype,
71 whereas *peripheral* genes and proteins constitute a regulatory network that propagates the effects
72 of genetic variants on the phenotype by modulating core gene expression and function^{10,11}. The
73 model postulates that the effects of peripheral proteins converge on relatively few core proteins that
74 have a major influence on the trait¹²; consequently many functional mutations in core genes remain
75 at low frequency in the adult population⁸, making their detection in GWAS challenging. While the
76 original authors propose *trans*-eQTLs to infer the underlying genetic network for all diseases^{12,13},
77 they admit that the required cohort sizes make this approach impractical¹¹. Rare variant sequencing,
78 alternatively suggested to associate core genes to diseases¹⁰, similarly requires very large cohorts
79 and has been criticized as a suboptimal strategy¹⁴.

80 Conceptually, the impact of peripheral genes and proteins is transmitted to core proteins via
81 ‘regulatory networks’ that encompass all layers of biological regulation¹⁰, and which we more
82 generically refer to as *molecular networks* to include biochemical modes of regulation. Thus, to
83 identify core disease genes, here we propose an advanced machine learning (ML) approach that
84 utilizes physical and regulatory molecular network information to identify core genes using Mendelian
85 disorder genes as a positive training set. Mendelian disorder genes not only “clearly fulfill the core
86 gene definition”¹⁴, but are examples towards the extreme end of the distribution of core genes, as a
87 single Mendelian disorder gene can cause the disease¹⁴. Moreover, for nearly all modes of biological
88 regulation increasingly complete reference networks are available that describe biochemical
89 interactions and regulatory effects, e.g. protein-DNA contacts and transcriptional regulation¹⁵⁻¹⁷,
90 protein-protein interactions^{18,19} and signaling pathways²⁰, and human metabolism²¹. While similar

91 information is also available from aggregated small-scale studies and predictions^{20,22–24}, these are
92 affected by a heavy inspection bias of hypothesis driven approaches and therefore not ideal for
93 reliable identification of hitherto unknown core genes^{18,25,26} (**Extended Data Fig. 1, Supplementary**
94 **Note 1**).

95 With the uptake of graph representation learning in biomedicine²⁷, novel options exist to process
96 networks alongside the input features in a joint ML model, thus approaching an *in silico*
97 representation of biological regulation. First implementations based on random-walks^{28–37} or graph
98 neural networks (GNN)^{38–43} show promise in predicting ‘disease genes’, but are often disease
99 specific, depend on hard-coded and partially biased input data, and do not further explore the
100 properties of predicted (core) genes (**Extended Data Fig. 2**). Moreover, in many machine learning
101 applications ensemble approaches outperform individual models^{44,45}.

102 Here we present Speos, an extensible and generalizable positive-unlabeled (PU) ML framework that
103 integrates information from biological networks and multiple biological modalities including gene
104 expression and GWAS data to predict core gene candidates for five common complex diseases. For
105 this, we developed an ensemble-based machine learning classifier. Systematic evaluation of the
106 predicted candidate genes using six external datasets demonstrates that these exhibit key core gene
107 characteristics, impact phenotypic manifestations to a similar extent as Mendelian genes, and are
108 enriched for potential new drug targets.

109 As Mendelian genes display all characteristics of ‘strong’ core genes¹⁴ we use these as positive
110 labels for a positive unlabeled graph representation learning²⁷ ensemble. Tissue-specific gene
111 expression and gene-level GWAS summary statistics will be used as input features^{10,3}. To identify
112 suitable base classifiers, we first conducted a hyperparameter optimization on the full data set
113 assuming negative labels for unlabeled genes.

114

115 **Performance of Base Classifiers**

116 Although Speos uses an ensemble-approach to achieve a consensus, the performance of the base
117 classifier is expected to be indicative for the performance of the ensemble. We therefore explored
118 the performance of different commonly used base classifiers (**Fig. 1a, Extended Data Fig. 3**). Since
119 it is unknown by which regulatory modalities the effects of peripheral genes are transmitted to core
120 proteins and if these differ among diseases, we tested 35 biological networks (**Fig. 1b**) selected for
121 their unbiased, systematic construction or strict curation approach. Among several GNN base
122 classifiers, the widely used GCN⁴⁶ layer, which is limited to one network at a time, convolutes the
123 features of each gene with a nonlinear projection of its immediate (1-step) neighborhood in a given
124 network. The TAG⁴⁷ layer is similar to GCN but considers higher-order neighborhoods (3-steps) of
125 any node and can block out unhelpful information. RGCN⁴⁸, is the relational equivalent of GCN and
126 can consider multiple networks simultaneously. Lastly, FiLM⁴⁹ is similar to RGCN, but uses feature-
127 wise linear modulation⁵⁰ to exclude and even override unhelpful neighborhood features based on the
128 center node and the connecting edge. Additional GNN layers performed worse during

129 hyperparameter optimization and were not further included (**Extended Data Fig. 4, 5, 6,**
130 **Supplementary Note 2**). Lastly, we included Node2Vec⁵¹ (N2V), which uses random walks on the
131 network and techniques developed for natural language processing to embed the network topology
132 into vector space in an unsupervised setting. These N2V-generated vectors can then be used as
133 input features by methods that cannot ingest networks directly like multilayer perceptrons (MLP),
134 logistic regression (LR) and random forests (RF).

135 We compared the ability of these base classifiers to identify Mendelian disorder genes using a 4-fold
136 cross-validation analysis, and quantified performance on the holdout set using area under the
137 receiver operator curve (AUROC) (**Fig. 1a**). AUROC is suitable for model comparison in PU learning
138 as known positives receive higher predictions than the average unlabeled gene, even though the
139 unlabeled (actual) positives reduce the optimal AUROC score. While many classifiers perform
140 similarly, most methods strongly depend on the input features and the network used. In line with the
141 omnigenic model^{10,11}, removing tissue-specific gene expression from the input features reduces the
142 performance. The “direct” annotated interaction network from IntAct²² works well with single-network
143 layers, while the FiLM layer performs well using the union of all networks (**Fig. 1b**). However, not all
144 networks improve the performance compared to an MLP that does not use any network, likely
145 reflecting the different importance of biological modalities and tissues for different diseases. With
146 GCN, many networks have a detrimental effect on the performance; using TAG, this effect is less
147 pronounced. Equivalently, the FiLM layer improves the performance compared to the RGCN layer
148 when all networks are used simultaneously and tends to predict genes with very high GWAS Z-
149 scores as core genes of cardiovascular disease, consistent with the omnigenic model (**Extended**
150 **Data Fig. 7**). As TAG and FiLM, but not GCN or RGCN, can ignore unhelpful neighborhood
151 information, their increased performance could reflect the fact that not the complete reference
152 network is relevant for disease manifestation and prediction. Intriguingly, in this benchmarking the
153 best performing method (**Fig. 1a**), N2V+MLP, does not use graph convolutions but embeds all
154 networks simultaneously into vector space using Node2Vec⁵¹ and then feeds these vectors
155 alongside the GWAS and gene expression features into an MLP (**Supplementary Note 3**).

156 Based on these results we selected five methods as base classifiers for our nested cross-validation
157 ensemble: N2V+MLP, which had the best overall performance, FiLM trained on all networks, and
158 TAG trained on IntAct Direct Interaction as best performing GNN-based methods. Despite the lower
159 performance we decided to also include MLP as a baseline classifier that does not use relational
160 network information, and GCN⁴⁶, which is regularly used in graph-based problems to ensure
161 comparability with other studies.

162

163 **Ensemble training and external validation of candidate genes**

164 We used the selected base classifiers to train the ensembles, which takes the form of a nested cross
165 validation with $m = 11$ outer folds, each comprised of $n = 10$ (inner fold) models (**Fig. 2a**). Within
166 each outer fold we statistically assess the agreement of the 10 inner models. Using the outer fold

167 hold-out set we select an *inner threshold* at which the agreement among the 10 inner models on
168 held out Mendelian genes surpasses random expectation ($FDR < 0.05$; Student's t-test, **Fig. 2b**,
169 **Supplementary Table ST1**). All genes with higher agreement than this *inner threshold* are
170 considered candidate genes of this outer fold. Since each outer fold predicts one set of candidate
171 genes, the overlap among these sets can be used to assign confidence to each candidate gene
172 using a consensus score (CS) (**Fig. 2c**), which indicates the number of outer folds which predict a
173 given gene to be a candidate. Genes receiving a CS of 0 are non-candidates, while genes with the
174 highest CS of 11 are the most confident predictions. We aimed to validate the model and our
175 predictions using systematic, orthogonal functional data.

176

177 **Mouse knockout data**

178 Since core genes are defined as directly contributing to a disease phenotype¹⁰⁻¹², genetic deletion
179 of core genes in mice should cause mouse phenotypes related to the human disease. We therefore
180 investigated if genetic deletion of mouse orthologs of predicted core genes across the different
181 convergence scores led to relevant phenotypes more often than expected by chance (**Fig. 3a**,
182 **Supplementary Table ST2**). Mendelian genes of all five disorders show a significant enrichment,
183 serving as a positive control and benchmark for this validation. From the least conservative (CS ≥ 1)
184 to the most stringent bin (CS = 11) the odds ratio (OR) of mouse knockout genes among the
185 candidate genes increases for all five disease groups. This indicates that, indeed, Speos' CS
186 identifies gene sets of increasing biological relevance and thus can serve as a measure of the quality
187 of predictions. Overall, FiLM and TAG predicted gene sets show the highest enrichment and only
188 when all methods show a low performance, as for diabetes, the gap between the methods narrows.
189 For other diseases, represented by cardiovascular diseases and body mass disorder, FiLM and/or
190 TAG perform consistently better while GCN, N2V+MLP and MLP remain at the tail end of the
191 distribution. This contrasts with the initial benchmark (**Fig. 1a**), where N2V+MLP performed best.
192 This discrepancy is likely due to a distribution shift referred to as *probabilistic gap*⁵², which here is
193 the consequence of differences between strong Mendelian genes used for training and the additional
194 core genes we aim to predict, for which only genetic variants with weaker effects are commonly
195 observed in the population. Because of these differences, (**Fig. 3b**), the decision boundary that is
196 best suited to recover the 'extreme' Mendelian core genes, i.e. our labeled positives (**Fig. 3c**), is ill-
197 suited to discover the 'normal' core genes, i.e. unlabeled positives, we aim to discover (**Fig. 3d**).
198 Importantly, we noticed that the inspection bias of hypothesis-driven small-scale studies present in
199 the body of literature, and reflected in curated interaction datasets, is amplified in predictions relying
200 on these (**Extended Data Fig. 8**). Removing the affected networks resolves the bias in the results,
201 yet especially FiLM predictions still validate at similar rates even after removal of the IntAct datasets
202 (**Extended Data Fig. 9a, Supplementary Note 4, Supplementary Table ST3**). Furthermore, gene
203 set enrichment analysis for gene ontology (GO) biological processes highlights relevant terms, such

204 as muscle contraction for cardiovascular disease and immune response for immune dysregulation
205 (**Supplementary Table ST4**).

206 The strong performance in predicting genes with relevant mouse phenotypes clearly demonstrates
207 that Speos identifies disease relevant (core) genes. Importantly, at high CS scores, the orthogonal
208 KO validation rates for sets from all methods except MLP are statistically indistinguishable from the
209 positive control benchmarks for the majority of disease groups (FDR > 0.05, z-test, **Supplementary**
210 **Table ST5**). Thus, biologically, our predictions are on par with strong Mendelian core genes.
211 However, even in the lowest bin (CS ≥ 1) genes with disease-relevant mouse phenotypes are
212 enriched for every disorder and every method (FDR < 0.05, Fisher's exact test, **Supplementary**
213 **Table ST2**), indicating that these sets are meaningful and disease-specific. We therefore include all
214 genes with CS ≥ 1 as candidate genes for the remainder of this work.

215

216 **Differential Gene Expression**

217 Variation in gene expression can translate into altered enzyme activities and network dynamics and
218 is therefore an important mechanism by which core genes contribute to disease^{10,53}. Thus, in disease
219 conditions both Mendelian and predicted core genes are expected to be enriched among
220 differentially expressed genes. Indeed, for all disease groups Mendelian genes show a strong
221 enrichment among differentially expressed genes and, again, serve as the reference. The predicted
222 core genes are similarly enriched among differentially expressed genes, although the enrichment is
223 weaker for many diseases (**Fig. 4a, Supplementary Table ST6**). This difference is consistent with
224 the notion of 'extreme' and 'normal' core genes and reinforces the idea that both Mendelian genes
225 and core genes underlying the genetic architecture of complex traits can cause phenotypes by loss
226 of function or expression mediated change of activities.

227 FiLM and TAG predict gene sets with the strongest enrichment in differentially expressed genes with
228 average odds ratios (OR_{AV}) of 5.4 and 5.0, respectively. Although TAG shows a stronger enrichment
229 of cardiovascular disease subtypes and predicted differential expression-enriched gene sets for 21
230 out of 22 disease subtypes, FiLM shows the highest OR_{AV} of all methods with especially strong
231 performance in immune dysregulation. The candidate genes produced by GCN show a lower
232 enrichment (OR_{AV} of 4.0), reflecting its initial inclusion as the weakest of the selected GNNs. As
233 before, the performance of N2V+MLP in predicting unknown core genes is worse compared to TAG
234 and FiLM. While showing high ORs for some subtypes (OR_{AV} of 3.6), for 5 of 22 disease subtypes
235 the predicted candidate sets show no enrichment for differential expression in disease conditions
236 (FDR > 0.05). The MLP without the Node2Vec node embeddings shows a substantially weaker
237 performance (OR_{AV} = 2.1), indicating the importance of network information. Using hypothesis-driven
238 curated interaction datasets differentially impacts the enrichment of predictions for differentially
239 expressed genes for different diseases (**Extended Data Fig. 9b, Supplementary Table ST7,**
240 **Supplementary Note 4**).

241 Overall, these results indicate that the Mendelian genes tend to be differentially expressed in
242 disease, likely contributing to disease etiology, and that optimized graph convolutional methods such
243 as FiLM and TAG are best suited to generalize this pattern to identify non-Mendelian candidate core
244 genes.

245

246 **Loss of Function and Missense Intolerance**

247 Because core genes directly influence disease phenotypes, these are expected to accumulate
248 protein-function impairing mutations at a lower frequency than regulatory peripheral genes^{8,10}. Using
249 ExAc cohort Z-scores⁵⁴, we examined this conjecture for two types of functional mutations: loss-of-
250 function (LoF) and missense (**Fig. 4b, c, Supplementary Tables ST8, ST9**). Consistent with the
251 omnigenic model, Mendelian genes are enriched for LoF intolerant genes in four out of five disorders.
252 Similarly, candidate core genes identified by FiLM and TAG are significantly different from the non-
253 candidates in four out of five disorders. All significant candidate sets are enriched for LoF intolerance,
254 except FiLM predictions for cardiovascular disease genes. For missense mutation intolerant genes
255 we observed overall similar trends (**Fig. 4c**). Interestingly, the signal from the Mendelian genes is
256 less pronounced and does not reach significance in three of the five diseases. Correspondingly, for
257 four disease groups the signal of the FiLM and TAG predicted core genes exceed that of the
258 Mendelian genes and thus presenting the inverse picture than loss-of-function intolerance. Different
259 biological and clinical properties of LoF and non-LoF mutations are well recognized⁵⁵ and the
260 observed differences between Mendelian and predicted core genes demonstrate that Speos
261 identifies genes with different biological properties than the training set. For cardiovascular diseases,
262 the FiLM predictions again show a significant depletion indicating a potentially interesting, but at this
263 point unexplained phenomenon.

264 Taken together, all our analyses strengthen the view of Mendelian genes as ‘extreme’ core genes,
265 and demonstrate that Speos reliably identifies phenotypically relevant genes with key core genes
266 properties.

267

268 **Examples and Model interpretation**

269 After demonstrating that Speos predicts *bona fide* core genes, we were interested in exploring
270 specific predicted examples to assess plausibility and to understand which aspects of the known
271 biology reflected in the model were most relevant for their prediction as core genes (**Supplementary**
272 **Note 5**). We selected genes with high CS, which are differentially expressed in at least one disease
273 subtype (**Fig. 4a**). To explore translational potential, we filtered for genes encoding yet untargeted
274 but druggable⁵⁶ proteins and applied model interpretation techniques to investigate gene- and
275 network-level patterns underlying their prediction as candidates. Both TNFSF15 and IL18RAP are
276 predicted as candidate genes for immune dysregulation by FiLM (CS 11 & 9); the former also by
277 TAG (CS 5) (**Fig. 5**).

278 TNFSF15 is differentially expressed in Crohn's disease and ulcerative colitis and its protein product
279 TL1A is part of the tumor necrosis factor superfamily and a ligand for two receptors: DR3 encoded
280 by TNFRSF25, which activates pro-inflammatory signaling, and soluble TR6 encoded by
281 TNFRSF6B, which acts as a non-functional decoy-receptor^{57,58}. Increased binding of TL1A to DR3
282 results in gut inflammation^{59,60} and endothelial dysfunction⁶¹, while neutralization of TL1A by TR6
283 down-regulates apoptosis⁶². This ability of TL1A/TNFSF15 to tip the balance of inflammation is
284 mirrored in findings that different genetic variants in- or decrease the risk for Crohn's disease⁶³⁻⁶⁵,
285 ulcerative colitis⁶⁴ and inflammatory bowel disease⁶⁶. We investigated influential network-level
286 patterns by assigning importance values to edges using integrated gradients⁶⁷. Model interpretation
287 for TNFSF15 shows that the receptor-ligand relationships with the protein products of TNFRSF25
288 and TNFRSF6B are among the strongest influences (**Fig. 5a**) illustrating that Speos' predictions
289 point towards biologically relevant and actionable mechanisms. The model interpretation further
290 suggests that drugs mimicking TR6 can alleviate inflammation by competitively sequestering TL1A
291 and thereby reducing binding of TL1A to DR3. Indeed, monoclonal antibody treatments leveraging
292 this mechanism are in clinical testing and initial results demonstrate a reduction of free TL1A and
293 normalization of pathologically dysregulated gut mucosa⁶⁸.

294 IL-18RAcP encoded by IL18RAP is an accessory protein for the receptor of the proinflammatory
295 interleukin 18 (IL-18) and greatly increases its affinity to its ligand⁶⁹. As such, it can increase the pro-
296 inflammatory effect of IL-18, exacerbating inflammation via the Interferon- γ pathway. IL18RAP is
297 differentially expressed in ulcerative colitis, its expression modulates treatment response in
298 rheumatoid arthritis⁷⁰ and it is considered a risk factor for celiac disease⁷¹ and autoimmune thyroid
299 diseases⁷². FiLM's prediction of IL18RAP is highly influenced by its connection to PIGH (**Fig. 5b**),
300 which is crucial for the first step of the glycosylphosphatidylinositol (GPI) biosynthesis⁷³. The GPI
301 glycan supports complex formation between IL-18RAcP and IL-18 receptor which increases
302 proinflammatory signaling⁷⁴. Thus, model interpretation suggests that interfering with the IL-18RAcP
303 - IL-18 receptor interaction reduces dysregulated inflammatory signaling. Indeed, it has recently been
304 demonstrated that cleaving IL-18RAcP using specific antibodies effectively reduces inflammation in
305 human blood cell cultures⁷⁵.

306 Both gene's predictions are strongly influenced by the GWAS input features for the complex forms
307 of the phenotype (**Fig. 5c, d**). For IL18RAP, high gene expression in whole blood, plasmacytoid
308 dendritic cells (DC) and the spleen are among the strongest influences, which is expected for a factor
309 contributing to autoimmunity^{76,77}. This combination of GWAS and disease-specific gene expression
310 are gene-level patterns expected for core genes by the omnigenic model¹⁰⁻¹². Beyond this, further
311 analyses and examples indicate that Speos finds core gene patterns along the entire continuum of
312 evidence combinations, from relying mostly on GWAS features (**Fig. 5c**), a combination of GWAS
313 and gene expression features (**Fig. 5d**) to almost exclusively utilizing gene expression features as
314 for obscurin and ITGA7 (**Extended Data Fig. 10, Supplementary Note 6**).

315 These examples showcase that Speos candidate genes constitute strong core gene hypotheses that
316 are consistent with the omnigenic model. Moreover, model interpretations suggest biochemically and
317 pharmaceutically plausible mechanisms for their impact on disease.

318

319 **Speos-candidates are potential drug targets**

320 Since core proteins are defined to directly and causally influence disease phenotypes, countering
321 the respective perturbations with pharmaceutical interventions should improve disease severity and
322 symptoms. To test this prediction systematically, we gathered drug-target-gene interactions from the
323 drug repurposing knowledge graph (DRKG)⁷⁸ and assessed the proportion of drug-target encoding
324 genes among the Mendelian and predicted core gene sets. Mendelian genes for all five disorders
325 are significantly enriched for genes encoding drug targets (DT), druggable proteins (Dr), and average
326 number of drugs targeting their products (xDC) (**Fig. 6**, FDR < 0.05, DT and Dr: Fisher's exact test,
327 xDC: U-test, **Supplementary Table ST7**). The enrichments of drug targets (DT) and of the average
328 number of drugs targeting the encoded proteins (xDC) both suggest that Mendelian genes have
329 been in the focus of drug development. The enrichment of druggable gene products (Dr) among
330 Mendelian genes and predicted candidates could be due to selection biases in the drug discovery
331 process, or may indicate that proteins with binding pockets for substrates or ligands are more likely
332 to be core disease genes that can directly cause disease phenotypes (**Fig. 6**). Crucially, Speos'
333 predicted candidate core genes are similarly enriched in all categories. In contrast to the analyses
334 of biological properties above, the observed enrichments are more varied among the methods with
335 each method predicting highly enriched subsets in one or two diseases, except for the network-
336 independent MLP.

337 In light of the retrospective confirmation of core gene products as suitable drug targets, core gene-
338 encoded proteins that are not drug-targets yet, are attractive candidates for future drug development.
339 However, proteins encoded by Mendelian genes are not enriched for druggable proteins once the
340 established drug targets have been removed (**Fig. 6**, Dr-), which indicates that the innovative
341 potential of Mendelian gene-products as drug targets has been largely exhausted. In contrast,
342 candidate genes produced by TAG and FiLM, as well as N2V+MLP jointly show a significant
343 enrichment for druggable proteins among the non-drug-targets in all five diseases (**Fig. 6**, FDR <
344 0.05, Fisher's exact test, **Supplementary Table ST10**) highlighting potential targets for development
345 of therapeutics for these epidemic disease groups. Removing the IntAct networks results in
346 prediction of significantly more not-targeted druggable (Dr-) genes for immune dysregulation
347 (**Extended Data Fig. 9c, ST11, Supplementary Note 4b**).

348

349 **Discussion**

350 Speos is a graph-representation learning framework that predicts novel core genes with high external
351 validation rates and properties predicted for core disease genes. In developing this framework, we
352 show that all investigated modalities of molecular networks carry relevant information to identify core

353 genes (**Fig. 5a, b**). At the same time, despite the strong GNN performance in the external biological
354 validation (**Fig. 3**), we were surprised by the moderate gain from including network information in the
355 initial prediction of held out Mendelian genes (**Fig. 1a**). This is mirrored by the finding that a
356 substantial part of the information that Speos extracts from molecular networks is encoded in the
357 topology and less so in features of neighboring genes (**Supplementary Note 3**). This is unexpected
358 as both the omnigenic model as well as the underlying biological thinking predict that the regulatory
359 and biochemical network surrounding a node modulates and impacts its function and activation. The
360 fact that the extensive network information we use does not result in an even greater gain in
361 performance may have a variety of possible reasons that could point towards future improvements.
362 Obvious shortfalls are imperfect SNP to gene mappings, and residual false-positives and the
363 incompleteness of all network maps^{18,79,80}. Similarly, models built on any single network are limited
364 by only accessing a small part and single modality of regulatory links explaining their weak
365 performance. Noteworthy, however, is the observation that learning methods that can selectively
366 ignore link information perform better than those that always consider the complete network
367 neighborhood. We also noticed that the average shortest path between all genes in the union of all
368 networks is close to 2 and many nodes have degrees exceeding 300 (**Extended Data Fig. 6c**)
369 indicating a very high network density. Likely, in any specific (patho-) physiological setting only few
370 of these interactions are responsible for dysregulated core protein activity, whereas others matter in
371 other conditions, other tissues, or for processes that do not influence disease etiology. We therefore
372 think that, in addition to a lack of relevant interactions, especially the abundance of disease-context-
373 irrelevant interactions constitute a challenge for learning algorithms and, in fact, for our
374 understanding of network function. For future implementations it may be helpful to include
375 directionality of signaling links for example based on systematic perturbation screens⁸¹⁻⁸⁴ and include
376 tissue specificity of edges as explicit features. Therefore, even in the absence of new systematic
377 experimental data, future iterations of this type of work are expected to jointly learn the network and
378 gene representations, thereby improving our understanding of network functioning.

379 In summary, we show that Speos is able to produce candidate core gene sets for different common
380 and complex diseases using Mendelian disorder genes as training examples (**Supplementary Table**
381 **ST12**). By building on properties predicted by the omnigenic model, we have further shown that
382 these candidate genes are enriched for mouse KO genes, differentially expressed genes, genes
383 intolerant of functional mutations and drug targets, all characteristics that are expected of core
384 genes. Furthermore, we show examples of candidate genes that have already been selected for
385 drug development and demonstrate that the model relies on similar data as domain experts. As such,
386 Speos is the first attempt at translating the omnigenic model into an ML framework for predicting and
387 prioritizing core genes across several disease areas. We anticipate that our results open the door
388 for a better understanding of core gene attributes and network functioning, and open possibilities for
389 future drug development.

390

391 **REFERENCES**

- 392 1. Ghoussaini, M. *et al.* Open Targets Genetics: systematic identification of trait-associated
393 genes using large-scale genetics and functional genomics. *Nucleic Acids Res.* **49**, D1311–
394 D1320 (2020).
- 395 2. Mountjoy, E. *et al.* An open approach to systematically prioritize causal variants and genes at
396 all published human GWAS trait-associated loci. *Nat. Genet.* **53**, 1527–1533 (2021).
- 397 3. Sinnott-Armstrong, N., Naqvi, S., Rivas, M. & Pritchard, J. K. GWAS of three molecular traits
398 highlights core genes and pathways alongside a highly polygenic background. *eLife* **10**,
399 e58615 (2021).
- 400 4. Loh, P.-R. *et al.* Contrasting genetic architectures of schizophrenia and other complex
401 diseases using fast variance-components analysis. *Nat. Genet.* **47**, 1385–1392 (2015).
- 402 5. Shi, H., Kichaev, G. & Pasaniuc, B. Contrasting the Genetic Architecture of 30 Complex Traits
403 from Summary Association Data. *Am. J. Hum. Genet.* **99**, 139–153 (2016).
- 404 6. Zhang, Y., Qi, G., Park, J.-H. & Chatterjee, N. Estimation of complex effect-size distributions
405 using summary-level statistics from genome-wide association studies across 32 complex
406 traits. *Nat. Genet.* **50**, 1318–1326 (2018).
- 407 7. Frei, O. *et al.* Bivariate causal mixture model quantifies polygenic overlap between complex
408 traits beyond genetic correlation. *Nat. Commun.* **10**, 2417 (2019).
- 409 8. O'Connor, L. J. *et al.* Extreme Polygenicity of Complex Traits Is Explained by Negative
410 Selection. *Am. J. Hum. Genet.* **105**, 456–476 (2019).
- 411 9. O'Connor, L. J. The distribution of common-variant effect sizes. *Nat. Genet.* **53**, 1243–1249
412 (2021).
- 413 10. Boyle, E. A., Li, Y. I. & Pritchard, J. K. An Expanded View of Complex Traits: From Polygenic
414 to Omnigenic. *Cell* **169**, 1177–1186 (2017).
- 415 11. Boyle, E. A., Li, Y. I. & Pritchard, J. K. The Omnigenic Model: Response from the Authors. *J.*
416 *Psychiatry Brain Sci.* **2**, (2017).
- 417 12. Liu, X., Li, Y. I. & Pritchard, J. K. Trans Effects on Gene Expression Can Drive Omnigenic
418 Inheritance. *Cell* **177**, 1022-1034.e6 (2019).
- 419 13. Vösa, U. *et al.* Large-scale cis- and trans-eQTL analyses identify thousands of genetic loci and
420 polygenic scores that regulate blood gene expression. *Nat. Genet.* **53**, 1300–1310 (2021).
- 421 14. Wray, N. R., Wijmenga, C., Sullivan, P. F., Yang, J. & Visscher, P. M. Common Disease Is
422 More Complex Than Implied by the Core Gene Omnigenic Model. *Cell* **173**, 1573–1580
423 (2018).
- 424 15. Gerstein, M. B. *et al.* Architecture of the human regulatory network derived from ENCODE
425 data. *Nature* **489**, 91–100 (2012).
- 426 16. Fuxman Bass, J. I. *et al.* Human gene-centered transcription factor networks for enhancers
427 and disease variants. *Cell* **161**, 661–673 (2015).
- 428 17. Zhu, X., Duren, Z. & Wong, W. H. Modeling regulatory network topology improves genome-

- 429 wide analyses of complex human traits. *Nat. Commun.* **12**, 2851 (2021).
- 430 18. Luck, K. *et al.* A reference map of the human binary protein interactome. *Nature* **580**, 402–408
431 (2020).
- 432 19. Huttlin, E. L. *et al.* Dual Proteome-scale Networks Reveal Cell-specific Remodeling of the
433 Human Interactome. *Cell* **184**, 3022–3040.e28 (2021).
- 434 20. Fabregat, A. *et al.* The Reactome Pathway Knowledgebase. *Nucleic Acids Res.* **46**, D649–
435 D655 (2018).
- 436 21. Brunk, E. *et al.* Recon3D enables a three-dimensional view of gene variation in human
437 metabolism. *Nat. Biotechnol.* **36**, 272–281 (2018).
- 438 22. Orchard, S. *et al.* The MIntAct project—IntAct as a common curation platform for 11 molecular
439 interaction databases. *Nucleic Acids Res.* **42**, D358–D363 (2014).
- 440 23. Szklarczyk, D. *et al.* The STRING database in 2021: customizable protein-protein networks,
441 and functional characterization of user-uploaded gene/measurement sets. *Nucleic Acids Res.*
442 **49**, D605–D612 (2021).
- 443 24. Oughtred, R. *et al.* The BioGRID database: A comprehensive biomedical resource of curated
444 protein, genetic, and chemical interactions. *Protein Sci. Publ. Protein Soc.* **30**, 187–200
445 (2021).
- 446 25. Yu, H. *et al.* High Quality Binary Protein Interaction Map of the Yeast Interactome Network.
447 *Science* **322**, 104–110 (2008).
- 448 26. Edwards, A. M. *et al.* Too many roads not taken. *Nature* **470**, 163–165 (2011).
- 449 27. Li, M. M., Huang, K. & Zitnik, M. Graph representation learning in biomedicine and healthcare.
450 *Nat. Biomed. Eng.* 1–17 (2022) doi:10.1038/s41551-022-00942-x.
- 451 28. Köhler, S., Bauer, S., Horn, D. & Robinson, P. N. Walking the Interactome for Prioritization of
452 Candidate Disease Genes. *Am. J. Hum. Genet.* **82**, 949–958 (2008).
- 453 29. Li, Y. & Patra, J. C. Genome-wide inferring gene-phenotype relationship by walking on the
454 heterogeneous network. *Bioinforma. Oxf. Engl.* **26**, 1219–1224 (2010).
- 455 30. Xie, M., Xu, Y., Zhang, Y., Hwang, T. & Kuang, R. Network-based Phenome-Genome
456 Association Prediction by Bi-Random Walk. *PLoS One* **10**, e0125138 (2015).
- 457 31. Zhao, Z.-Q., Han, G.-S., Yu, Z.-G. & Li, J. Laplacian normalization and random walk on
458 heterogeneous networks for disease-gene prioritization. *Comput. Biol. Chem.* **57**, 21–28
459 (2015).
- 460 32. Himmelstein, D. S. & Baranzini, S. E. Heterogeneous Network Edge Prediction: A Data
461 Integration Approach to Prioritize Disease-Associated Genes. *PLOS Comput. Biol.* **11**,
462 e1004259 (2015).
- 463 33. Huang, J. K. *et al.* A systematic evaluation of molecular networks for discovery of disease
464 genes. *Cell Syst.* **6**, 484–495.e5 (2018).
- 465 34. Valdeolivas, A. *et al.* Random walk with restart on multiplex and heterogeneous biological
466 networks. *Bioinformatics* **35**, 497–505 (2019).

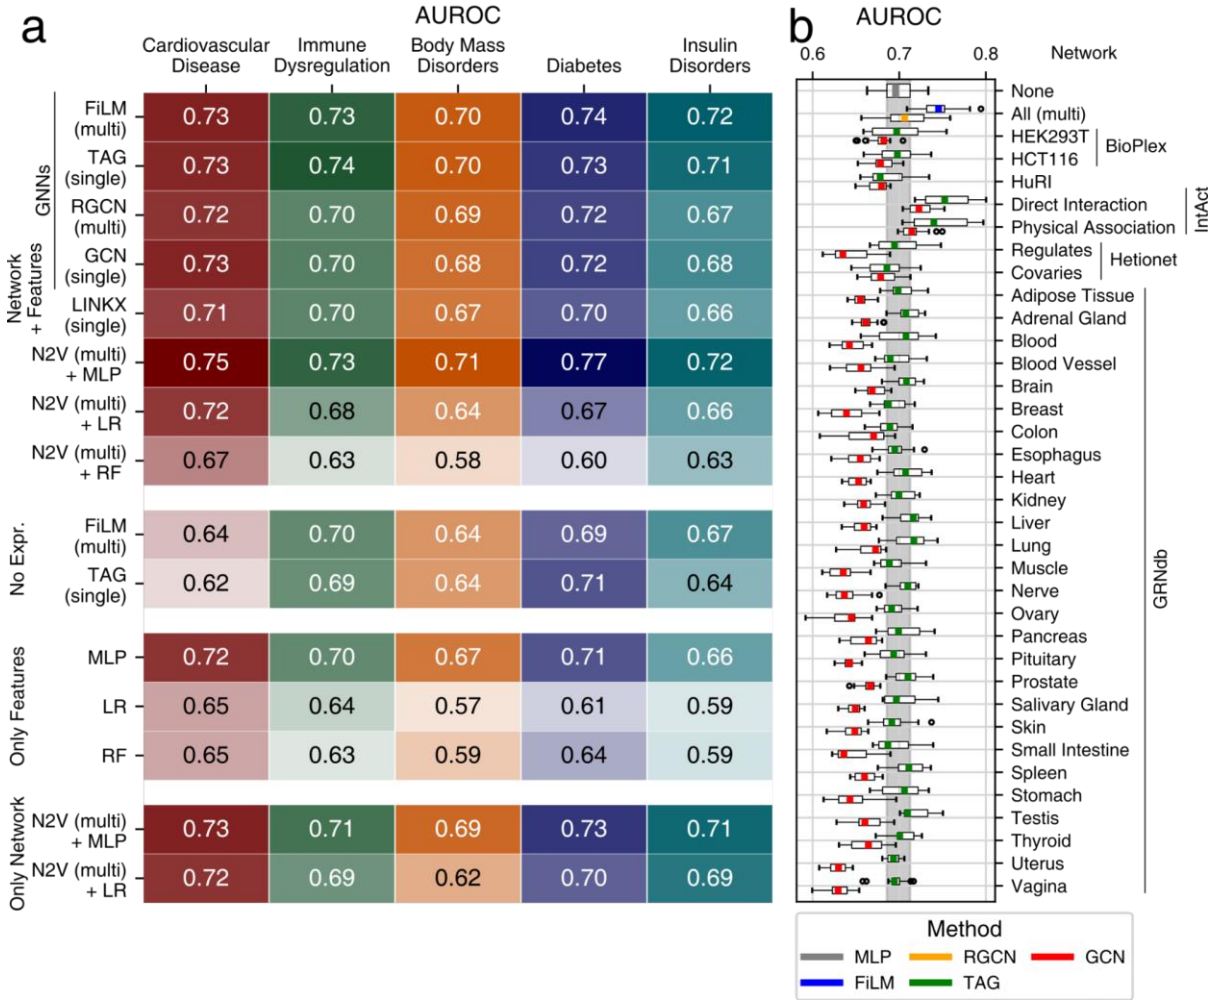
- 467 35. Joodaki, M., Ghadiri, N., Maleki, Z. & Lotfi Shahreza, M. A scalable random walk with restart
468 on heterogeneous networks with Apache Spark for ranking disease-related genes through
469 type-II fuzzy data fusion. *J. Biomed. Inform.* **115**, 103688 (2021).
- 470 36. Zhong, Y., Shen, C., Wu, H., Xu, T. & Luo, L. Improving the Prediction of Potential Kinase
471 Inhibitors with Feature Learning on Multisource Knowledge. *Interdiscip. Sci. Comput. Life Sci.*
472 **14**, 775–785 (2022).
- 473 37. Wang, L., Shang, M., Dai, Q. & He, P.-A. Prediction of lncRNA-disease association based on a
474 Laplace normalized random walk with restart algorithm on heterogeneous networks. *BMC*
475 *Bioinformatics* **23**, 5 (2022).
- 476 38. Vanunu, O., Magger, O., Ruppín, E., Shlomi, T. & Sharan, R. Associating Genes and Protein
477 Complexes with Disease via Network Propagation. *PLOS Comput. Biol.* **6**, e1000641 (2010).
- 478 39. Yang, P., Li, X., Chua, H.-N., Kwok, C.-K. & Ng, S.-K. Ensemble Positive Unlabeled Learning
479 for Disease Gene Identification. *PLOS ONE* **9**, e97079 (2014).
- 480 40. Schulte-Sasse, R., Budach, S., Hnisz, D. & Marsico, A. Integration of multiomics data with
481 graph convolutional networks to identify new cancer genes and their associated molecular
482 mechanisms. *Nat. Mach. Intell.* **3**, 513–526 (2021).
- 483 41. Choi, W. & Lee, H. Identifying disease-gene associations using a convolutional neural
484 network-based model by embedding a biological knowledge graph with entity descriptions.
485 *PLOS ONE* **16**, e0258626 (2021).
- 486 42. Du, J. *et al.* Graph Embedding Based Novel Gene Discovery Associated With Diabetes
487 Mellitus. *Front. Genet.* **12**, 779186 (2021).
- 488 43. Zhang, Y., Chen, L. & Li, S. CIPHER-SC: Disease-Gene Association Inference Using Graph
489 Convolution on a Context-Aware Network With Single-Cell Data. *IEEE/ACM Trans. Comput.*
490 *Biol. Bioinform.* **19**, 819–829 (2022).
- 491 44. Breiman, L. Bagging Predictors. *Mach. Learn.* **24**, 123–140 (1996).
- 492 45. Rokach, L. Ensemble-based classifiers. *Artif. Intell. Rev.* **33**, 1–39 (2010).
- 493 46. Kipf, T. N. & Welling, M. Semi-Supervised Classification with Graph Convolutional Networks.
494 *ArXiv160902907 Cs Stat* (2016).
- 495 47. Du, J., Zhang, S., Wu, G., Moura, J. M. F. & Kar, S. Topology Adaptive Graph Convolutional
496 Networks. Preprint at <http://arxiv.org/abs/1710.10370> (2018).
- 497 48. Schlichtkrull, M. *et al.* Modeling Relational Data with Graph Convolutional Networks.
498 *ArXiv170306103 Cs Stat* (2017).
- 499 49. Brockschmidt, M. GNN-FiLM: Graph Neural Networks with Feature-wise Linear Modulation.
500 Preprint at <http://arxiv.org/abs/1906.12192> (2020).
- 501 50. Perez, E., Strub, F., de Vries, H., Dumoulin, V. & Courville, A. FiLM: Visual Reasoning with a
502 General Conditioning Layer. Preprint at <http://arxiv.org/abs/1709.07871> (2017).
- 503 51. Grover, A. & Leskovec, J. node2vec: Scalable Feature Learning for Networks. in *Proceedings*
504 *of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*

- 505 - *KDD '16* 855–864 (ACM Press, 2016). doi:10.1145/2939672.2939754.
- 506 52. Bekker, J. & Davis, J. Learning from positive and unlabeled data: a survey. *Mach. Learn.* **109**,
507 719–760 (2020).
- 508 53. Botstein, D. & Risch, N. Discovering genotypes underlying human phenotypes: past
509 successes for mendelian disease, future approaches for complex disease. *Nat. Genet.* **33**,
510 228–237 (2003).
- 511 54. Lek, M. *et al.* Analysis of protein-coding genetic variation in 60,706 humans. *Nature* **536**, 285–
512 291 (2016).
- 513 55. Gerasimavicius, L., Livesey, B. J. & Marsh, J. A. Loss-of-function, gain-of-function and
514 dominant-negative mutations have profoundly different effects on protein structure. *Nat.*
515 *Commun.* **13**, 3895 (2022).
- 516 56. Finan, C. *et al.* The druggable genome and support for target identification and validation in
517 drug development. *Sci. Transl. Med.* (2017) doi:10.1126/scitranslmed.aag1166.
- 518 57. Migone, T. S. *et al.* TL1A is a TNF-like ligand for DR3 and TR6/DcR3 and functions as a T cell
519 costimulator. *Immunity* **16**, 479–492 (2002).
- 520 58. Zhan, C. *et al.* Decoy strategies: the structure of TL1A:DcR3 complex. *Struct. Lond. Engl.*
521 *1993* **19**, 162–171 (2011).
- 522 59. Furfaro, F. *et al.* TL1A: A New Potential Target in the Treatment of Inflammatory Bowel
523 Disease. *Curr. Drug Targets* **22**, 760–769 (2021).
- 524 60. Jin, S. *et al.* TL1A/TNFSF15 directly induces proinflammatory cytokines, including TNF α , from
525 CD3+CD161+ T cells to exacerbate gut inflammation. *Mucosal Immunol.* **6**, 886–899 (2013).
- 526 61. Della Bella, S. *et al.* Pathologic up-regulation of TNFSF15-TNFRSF25 axis sustains
527 endothelial dysfunction in unprovoked venous thromboembolism. *Cardiovasc. Res.* **116**, 698–
528 707 (2020).
- 529 62. Yang, C.-R. *et al.* Soluble decoy receptor 3 induces angiogenesis by neutralization of TL1A, a
530 cytokine belonging to tumor necrosis factor superfamily and exhibiting angiostatic action.
531 *Cancer Res.* **64**, 1122–1129 (2004).
- 532 63. Yang, D.-H. *et al.* TNFSF15 is an independent predictor for the development of Crohn's
533 disease-related complications in Koreans. *J. Crohns Colitis* **8**, 1315–1326 (2014).
- 534 64. He, L., Chen, J., Sun, J., Peng, J. & He, Q. Protective association of TNFSF15 polymorphisms
535 with Crohn's disease and ulcerative colitis: A meta-analysis. *Saudi J. Gastroenterol. Off. J.*
536 *Saudi Gastroenterol. Assoc.* **24**, 201–210 (2018).
- 537 65. Zhou, Y. *et al.* Polymorphism rs6478109 in the TNFSF15 gene contributes to the susceptibility
538 to Crohn's disease but not ulcerative colitis: a meta-analysis. *J. Int. Med. Res.* **48**,
539 300060520961675 (2020).
- 540 66. Richard, A. C. *et al.* Reduced monocyte and macrophage TNFSF15/TL1A expression is
541 associated with susceptibility to inflammatory bowel disease. *PLoS Genet.* **14**, e1007458
542 (2018).

- 543 67. Sundararajan, M., Taly, A. & Yan, Q. Axiomatic Attribution for Deep Networks. in *Proceedings*
544 *of the 34th International Conference on Machine Learning* 3319–3328 (PMLR, 2017).
- 545 68. Kokkotis, G. & Bamias, G. TL1A as a therapeutic target in inflammatory bowel disease. *Expert*
546 *Rev. Clin. Immunol.* **18**, 551–555 (2022).
- 547 69. Wu, C. *et al.* IL-18 Receptor β -Induced Changes in the Presentation of IL-18 Binding Sites
548 Affect Ligand Binding and Signal Transduction. *J. Immunol.* **170**, 5571–5577 (2003).
- 549 70. Cherlin, S. *et al.* Investigation of genetically regulated gene expression and response to
550 treatment in rheumatoid arthritis highlights an association between IL18RAP expression and
551 treatment response. *Ann. Rheum. Dis.* **79**, 1446–1452 (2020).
- 552 71. Hunt, K. A. *et al.* Newly identified genetic risk variants for celiac disease related to the immune
553 response. *Nat. Genet.* **40**, 395–402 (2008).
- 554 72. Wang, X. *et al.* Polymorphisms of ST2-IL18R1-IL18RAP gene cluster: a new risk for
555 autoimmune thyroid diseases. *Int. J. Immunogenet.* **43**, 18–24 (2016).
- 556 73. Watanabe, R. *et al.* The first step of glycosylphosphatidylinositol biosynthesis is mediated by a
557 complex of PIG-A, PIG-H, PIG-C and GPI1. *EMBO J.* **17**, 877–885 (1998).
- 558 74. Fukushima, K., Ikehara, Y. & Yamashita, K. Functional Role Played by the
559 Glycosylphosphatidylinositol Anchor Glycan of CD48 in Interleukin-18-induced Interferon- γ
560 Production*. *J. Biol. Chem.* **280**, 18056–18062 (2005).
- 561 75. Li, S. *et al.* A novel anti-human IL-1R7 antibody reduces IL-18-mediated inflammatory
562 signaling. *J. Biol. Chem.* **296**, 100630 (2021).
- 563 76. Galicia, G. & Gomerman, J. L. Plasmacytoid dendritic cells and autoimmune inflammation.
564 *Biol. Chem.* **395**, 335–346 (2014).
- 565 77. Jang, E., Cho, S., Pyo, S., Nam, J.-W. & Youn, J. An Inflammatory Loop Between Spleen-
566 Derived Myeloid Cells and CD4+ T Cells Leads to Accumulation of Long-Lived Plasma Cells
567 That Exacerbates Lupus Autoimmunity. *Front. Immunol.* **12**, (2021).
- 568 78. Ioannidis, V. N. *et al.* DRKG - drug repurposing knowledge graph for covid-19. (2020).
- 569 79. Venkatesan, K. *et al.* An empirical framework for binary interactome mapping. *Nat. Methods* **6**,
570 83–90 (2009).
- 571 80. Peel, L., Peixoto, T. P. & De Domenico, M. Statistical inference links data and theory in
572 network science. *Nat. Commun.* **13**, 6794 (2022).
- 573 81. Dixit, A. *et al.* Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA
574 Profiling of Pooled Genetic Screens. *Cell* **167**, 1853-1866.e17 (2016).
- 575 82. Datlinger, P. *et al.* Pooled CRISPR screening with single-cell transcriptome readout. *Nat.*
576 *Methods* **14**, 297–301 (2017).
- 577 83. Gasperini, M. *et al.* A Genome-wide Framework for Mapping Gene Regulation via Cellular
578 Genetic Screens. *Cell* **176**, 377-390.e19 (2019).
- 579 84. Bock, C. *et al.* High-content CRISPR screening. *Nat. Rev. Methods Primer* **2**, 1–23 (2022).
- 580

581 **FIGURES**

582



583

584

585

586

587

588

589

590

591

592

593

594

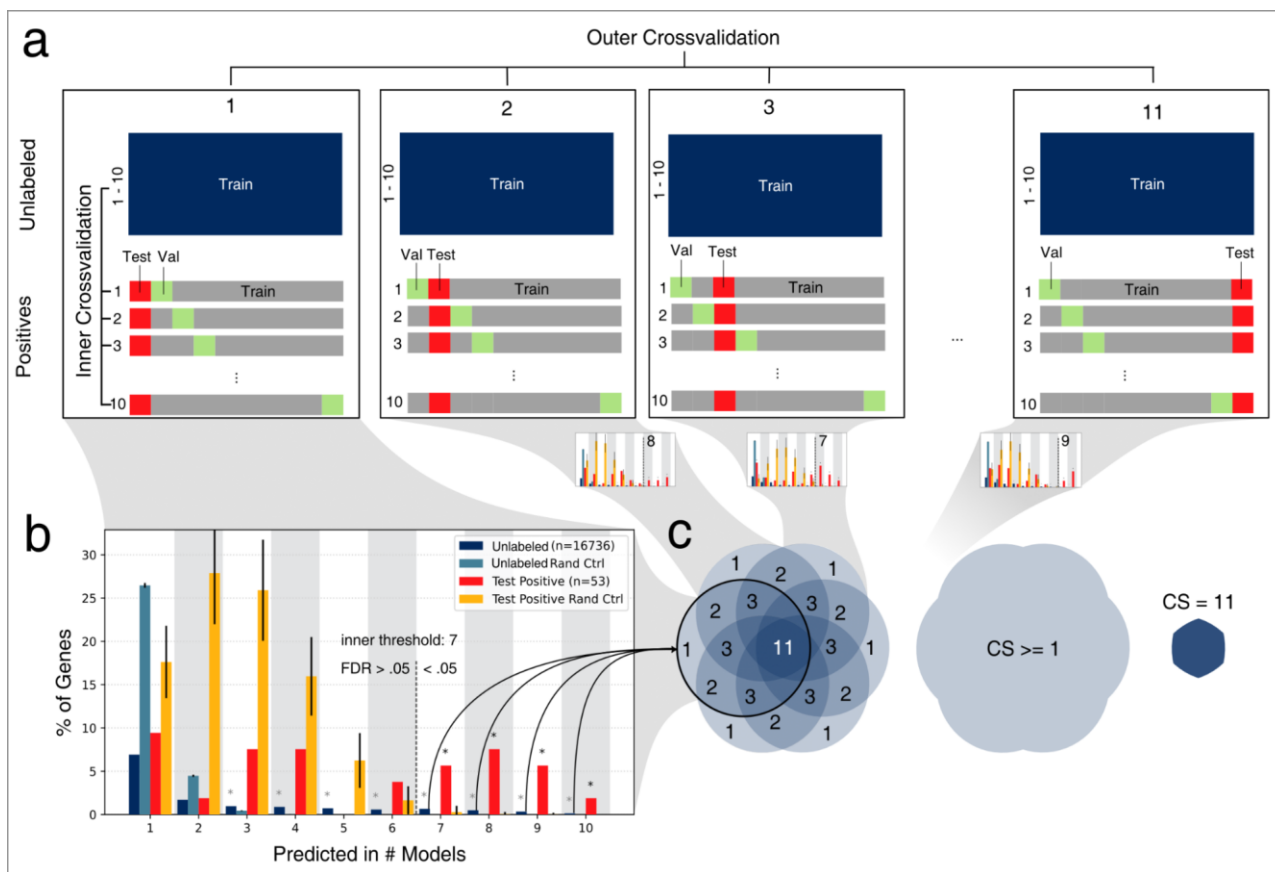
595

596

597

598

Fig. 1 | Performance AUROC. **a**, the mean area under the receiver operator characteristic curve (AUROC) metric (higher is better) over $n = 16$ models for different base classifiers, dataset variants and phenotypes. For AUPRC and mean rank and see **Extended Data Fig. 3**. Combinations of methods and input data are indicated along the y-axis. The blocks group models using common input data as indicated: Only Network: adjacency matrix/matrices; Only Features: gene expression and GWAS input features but no adjacency matrices; No Expression: GWAS input features and adjacency of individual (single) or multiple (multi) networks; Network + Features: adjacency of individual (single) or multiple (multi) networks, GWAS and gene expression. **b**, AUROC of 4 repetitions of a 4-fold cross validation for the indicated individual networks, all networks simultaneously (multi) using the classifier methods indicated by color. The vertical gray area indicates the interquartile range of the MLP, which does not use any network information (uppermost boxplot). Each boxplot is based on $n = 16$ values. Boxes represent the interquartile range, colored bars are medians, whiskers extend at most 1.5 times the interquartile range, and outliers are shown individually.



599

600

601

602

603

604

605

606

607

608

609

610

611

612

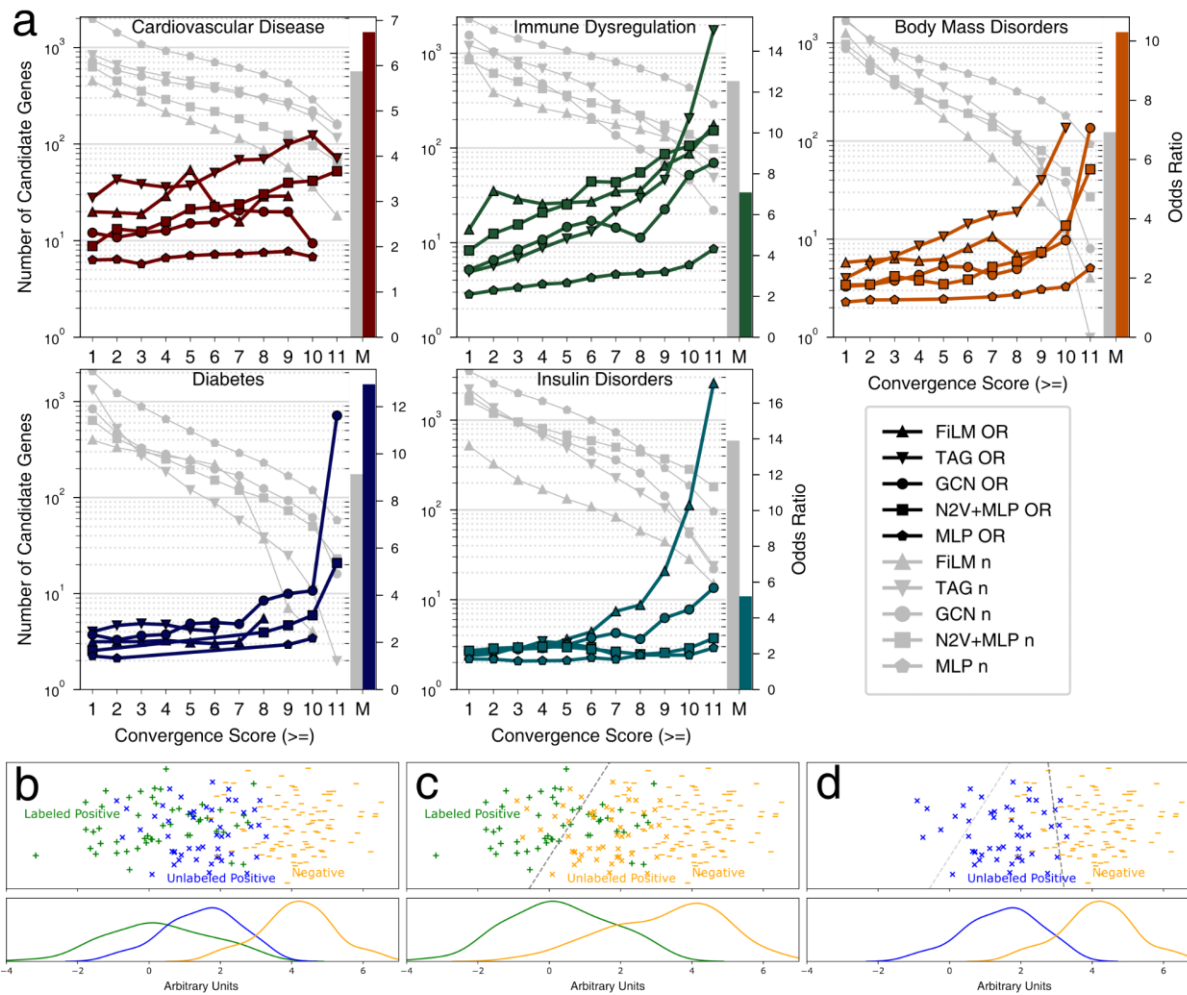
613

614

615

616

Fig. 2 | Cross Validation Ensemble. **a**, The 11 folds of the outer cross-validation, each with 10 inner cross-validation folds. Each inner cross validation fold corresponds to one ML model. The positives correspond to the Mendelian disorder genes for the given phenotype. Every model within one outer fold has the same positive test set (red square), but different positive validation sets (green squares) used for early stopping. All unlabeled genes are used for training for every model of every fold. **b**, For each outer fold, the overlap of candidate-predicted unknowns (dark blue bars) and correct predictions of the positive test set (red bars) of the 10 models are compared to random sets of the same size. Mean and standard deviation of the random sets are shown colored according to the legend (light blue and orange bars, error bars denote one standard deviation). If the observed overlap of correctly classified held out positives is significantly higher than expected by chance (FDR < 0.05, one-sided t-tests, **Supplementary Table ST1**, marked with black asterisk), the predicted unlabeled genes of these overlap bins (*inner threshold*) are considered candidate core genes for this outer fold. **c**, the candidate genes of each outer fold are aggregated. The Consensus Score (CS) of candidate genes ranges from 1 to 11 and indicates by how many outer folds a given gene is selected as candidate core gene. Genes with CS of 0 are considered non-candidate genes.



617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

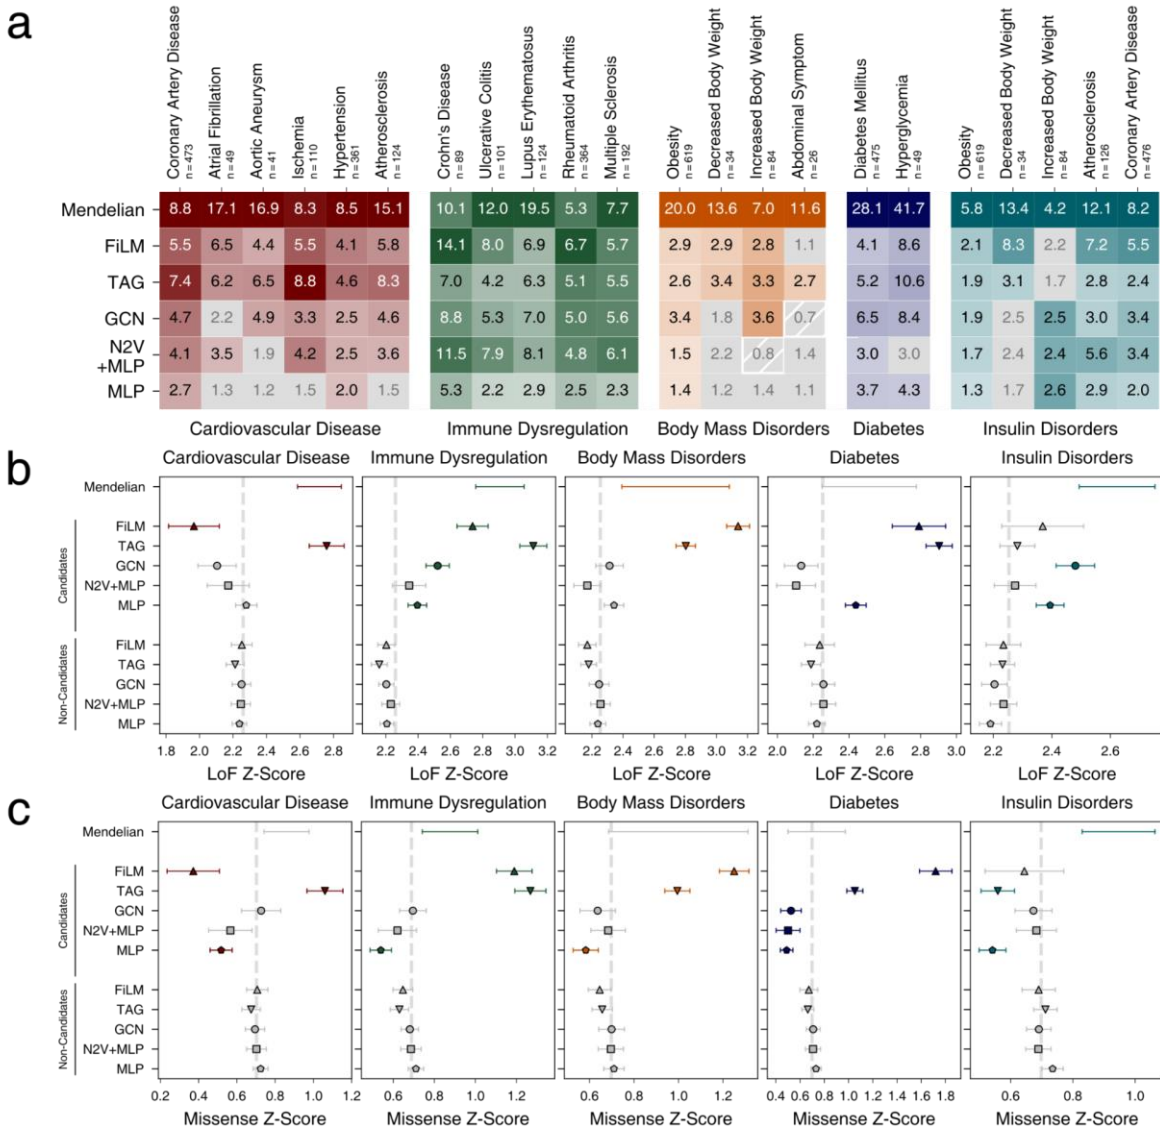
633

634

635

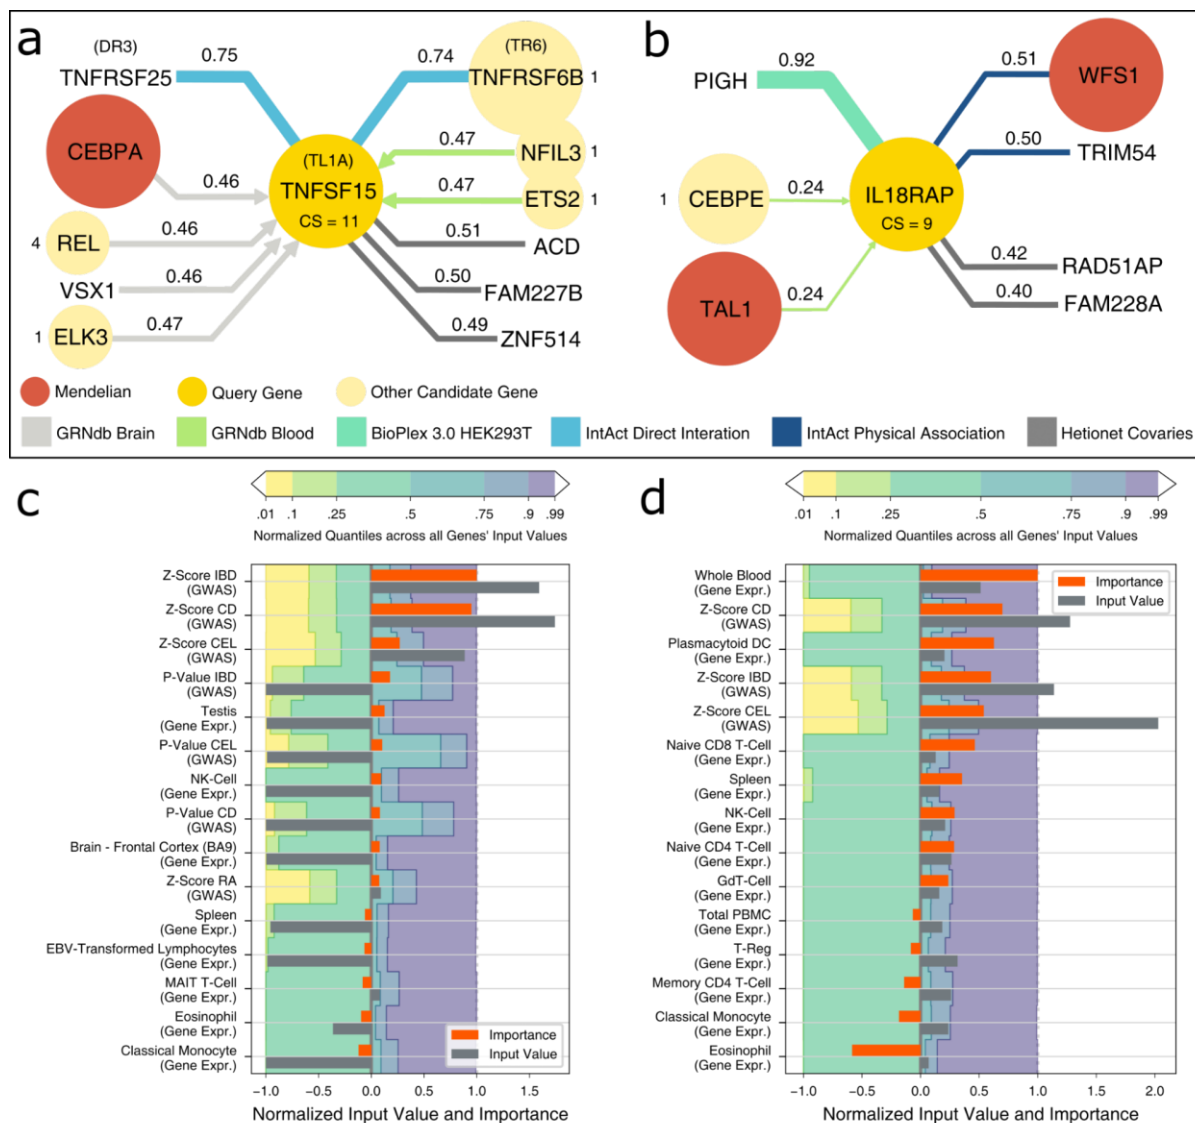
636

Fig. 3 | Mouse Knockout Validation. **a**, Odds ratio (OR) (right y-axis) for observing disease relevant phenotypes in mice with knock-outs of orthologs of candidate core genes in the indicated convergence score bins (x-axis) of the five classifier methods (colored lines). Gray lines indicate strength of candidate gene sets (left y-axis) in the corresponding bin for the phenotypes as indicated in the panel. Only ORs with an FDR < 0.05 (Fisher's exact test) are shown. Bars to the right of each plot (M) indicate set strength (gray) and OR (colored) of Mendelian genes for each phenotype. Precise *P*-values, FDR, and *n* for each test are shown in **Supplementary Table ST2**. **b**, Illustration of the probabilistic gap according to the "sampled at random with probabilistic gap positive unlabeled" (SAR-PGPU) case from ref. ⁵². Labeled and unlabeled positives are drawn from the same underlying distribution, however the label frequency increases towards the more extreme end of the positive distribution, e.g. due to detection bias. We assume this scenario to be true for Mendelian genes as "extreme" core genes¹⁴. **c**, For the internal cross validation on a holdout set (as in **Fig. 1a**) all unlabeled genes are considered negatives. Consequently, models with the indicated decision boundary (gray dashed line) will perform well. **d**, For prediction and subsequent validation of less 'extreme' true, but unknown, core genes indicated by blue labels (**Fig. 3a**), a model with a decision boundary near the dark gray dashed line is expected to perform well, while the decision boundary from panel **b** (light gray dashed line) is not optimal anymore.



637
638
639
640
641
642
643
644
645
646
647
648
649

Fig. 4 | External Validation. **a**, Odds ratios (ORs) of Mendelian genes (first row) and of candidate genes of the five selected methods (rows) for common complex subtypes of the five Mendelian disorder groups. ORs with FDR > 0.05 (Fisher's exact test) in gray. **b**, **c**, LoF intolerance and missense mutation intolerance Z-scores of Mendelian genes, and the indicated candidate and non-candidate sets generated by the five methods. Shown are group means and 95% confidence intervals of Tukey's HSD test. Colored symbols and error bars indicate P < 0.05 in comparison with respective non-candidate sets; not significant sets in gray. Dashed line indicates the mean across all genes. Precise P-values, FDR, and n for each test in each panel are shown in **Supplementary Tables ST6, ST8, and ST9**, respectively.



650
651

652 **Fig. 5 | Model Interpretation.** **a.** Most important edges for FiLM's prediction of TNFSF15 as
653 candidate gene for immune dysregulation. Shown are HGNC gene symbols, protein symbols are
654 added in parenthesis where necessary. The query gene node is shown in the center, with adjacent
655 relevant nodes in the periphery. Candidate genes are signed with their Consensus Score (CS). The
656 color of the edges denotes the network and the strength of the edge shows the relative importance
657 for the prediction of the query gene which is also written at the edge. Arrowheads indicate direction
658 of edges, undirected edges have no arrowheads. A value of 1 means that it is the most important
659 edge for all models of the ensemble, while a value of 0 indicates that it is the least important
660 edge for every model. Shown are 11 out of 4.3 million edges, 301 of which are in the direct neighborhood
661 of the query gene. **b.** Most important edges for FiLM's prediction of IL18RAP as candidate gene for
662 immune dysregulation. Shown are 7 out of 4.3 million edges, 431 of which are in the direct
663 neighborhood of the query gene. **c, d:** Input feature importance for TNFSF15 and IL18RAP alongside
664 the respective feature's input value, compared to the input values of other genes by the quantile
665 borders in the background. Shown are the 10 features with the strongest positive influence and the
666 5 features with the strongest negative influence. Negative input values are normalized to the interval
667 [-1; 0] and positive input values to [0; 1] for visualization. Gray bars exceeding the colored areas are
668 either below the 1% quantile or above the 99% quantile of that input feature. Importance values are
669 obtained by integrated gradients and normalized to the interval [-1; 1]. Positive importance values
670 are in favor of the prediction as candidate genes, negative importance values are attributed to

671 features that contradict the prediction. For the input feature importance of surrounding nodes see
672 **Supplementary Note 5.**
673

674
675

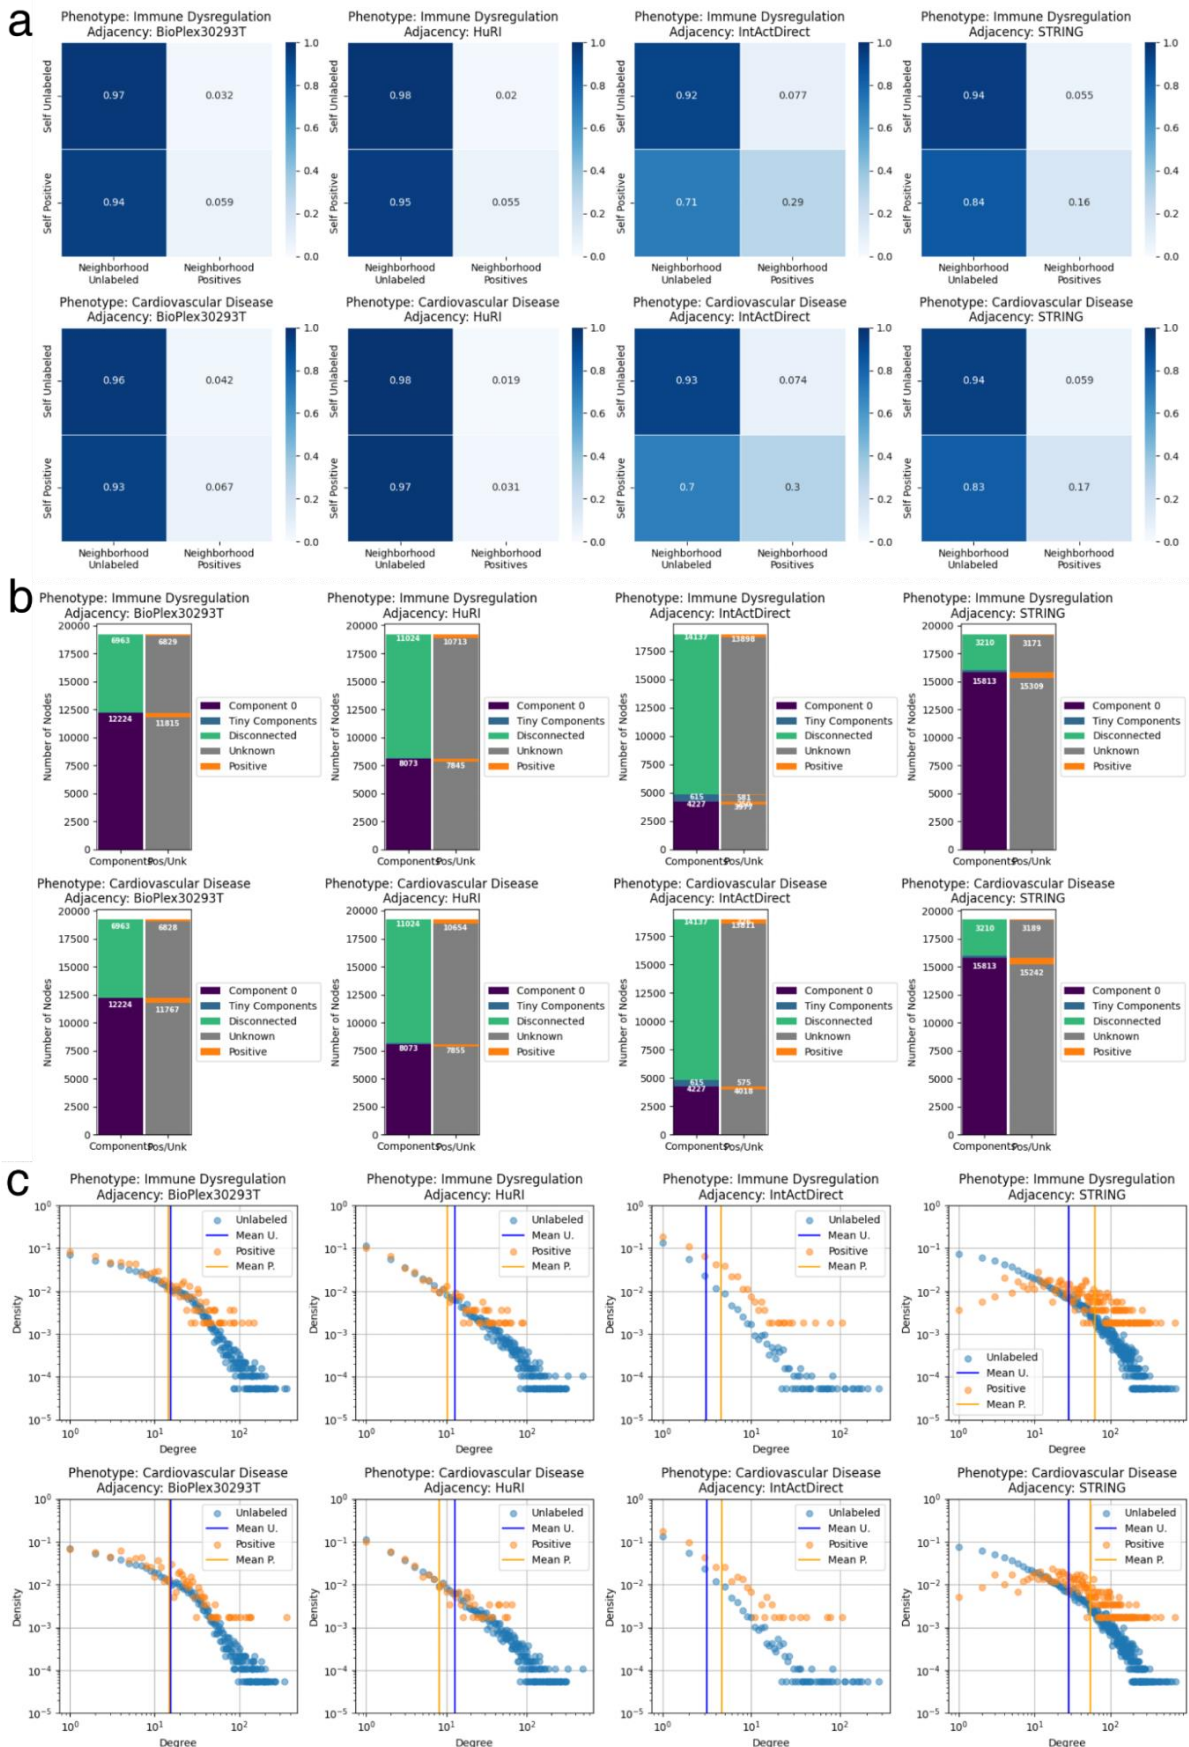
	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-
Mendelian	4.7	2.5	2.0	0.8	5.3	3.0	3.1	1.2	8.7	3.5	3.4	0.7	6.1	3.8	3.3	1.2	5.8	3.0	2.4	1.0
FILM	3.5	2.0	2.1	1.1	3.1	2.2	1.7	0.7	2.8	2.0	1.5	0.8	2.5	2.2	0.8	0.2	6.3	2.2	2.7	1.9
TAG	5.9	2.2	2.4	1.6	5.1	2.2	2.3	1.6	2.9	2.0	2.2	1.4	3.4	2.0	2.1	1.3	2.2	1.5	1.6	1.2
GCN	4.4	1.5	1.8	1.2	2.4	1.8	2.6	1.8	2.3	2.0	2.4	1.6	3.9	1.8	3.0	2.2	2.4	1.5	1.6	1.2
N2V+MLP	2.7	2.0	1.9	1.3	2.7	1.8	2.7	1.7	1.4	1.2	1.8	1.4	1.8	1.8	1.8	1.5	3.6	1.8	2.3	1.5
MLP	1.9	1.2	1.3	0.8	1.4	1.2	1.4	1.3	1.0	1.2	1.4	1.3	1.5	1.2	1.7	1.2	1.8	1.2	1.3	1.0
	Cardiovascular Disease				Immune Dysregulation				Body Mass Disorders				Diabetes				Insulin Disorders			

676
677

678 **Fig. 6 | Drug Target Analysis.** Enrichment of drug targets and druggability in Mendelian disorder
679 genes and indicated candidate gene sets. DT: OR of known drug targets. xDC: Ratio of median
680 number of drug-gene interactions per candidate gene to the median of non-candidates, only genes
681 with drug-gene interactions are considered. Ratios with FDR > 0.05 (U-test) are grayed out. Dr: OR
682 of druggable genes. Dr-: OR of druggable genes, after all drug targets have been removed. Odds
683 Ratios with FDR > 0.05 (Fisher's exact test) are grayed out. For all panels, precise *P*-values, FDR,
684 and n for each test are shown in **Supplementary Table ST10**.

685 EXTENDED DATA FIGURES

686



687

688

689

Extended Data Fig. 1 | Properties of biased and unbiased networks. **a**, the top and bottom rows

690 show the dependency of the neighborhood labels on the label of the center node of unlabeled genes
691 and Mendelian disorder genes for immune dysregulation and cardiovascular disease, respectively.
692 The two left columns show results for systematically generated Bioplex 3.0 HEK293T and HuRI
693 adjacency matrices. The two right columns show the adjacency matrices IntAct Direct Interaction
694 and STRING (confidence > 0.7), which are largely assembled from hypothesis-driven small-scale
695 data. Connectivity of Mendelian gene encoded proteins in the systematic networks is similar to that
696 of unlabeled nodes. In the collated networks, proteins encoded by Mendelian disorder genes show
697 higher assortativity, i.e. tendency to interact with each other, for both phenotypes. **b**, in each panel
698 the left bar shows the fraction of nodes in the largest connected component (component 0) versus
699 isolated small components and disconnected nodes. The right bar shows how the positive and
700 unlabeled nodes are distributed among these components. **c**, the top and bottom rows show the
701 degree distributions of Mendelian disorder genes and unlabeled genes for immune dysregulation
702 and cardiovascular disease, respectively. The two left columns show the adjacencies Bioplex 3.0
703 293T and HuRI, which are unbiased, systematically generated networks. The two right columns
704 show the adjacencies IntAct Direct and STRING (confidence > 0.7), which are not systematically
705 generated. The bias towards known disease genes in the two right networks can be seen for both
706 phenotypes. First, the average degree of Mendelian disorder genes is higher than the average
707 degree of unlabeled genes. Second, the degree distribution of the Mendelian disorder genes in
708 STRING does not follow a scale-free degree distribution. On the contrary, nodes with a medium
709 degree are the most abundant, while nodes of low and very high degree are rare.

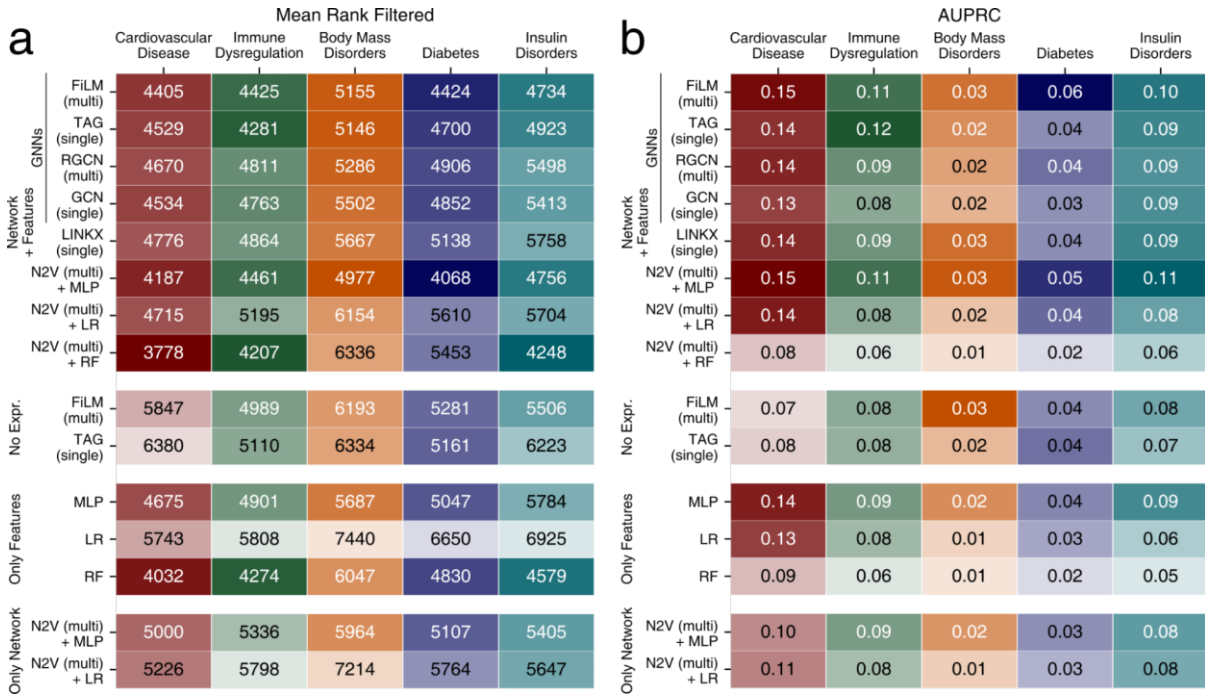
710
711

Method Reference	class prior assumption-free*	ensemble	unbiased networks**	networks extensible***	systematic external validation****	uses random walks	uses graph convolutions	uses input features	input extensible***
PRINCE ³⁸	✗	✗	✗	✗	✗	✗	✓ [†]	✗	n.a.
Yang, 2014 ³⁹	✗	✓	✗	✗	✗	✗	✓ [†]	✓	✗
DWPC ³²	✗	✗	✗	✗	✗	✓	✗	✗	✓
Huang, 2018 ³³	✗	✗	✓	✗	✓	✓	✗	✗	n.a.
RWR-MH ³⁴	✗	✗	✗	✗	✗	✓	✗	✗	n.a.
RWRHN-FF ³⁵	✗	✗	✗	✗	✗	✓	✗	✗	n.a.
EMOGI ⁴⁰	✗	✗	✗	✓	✓	✗	✓	✓	✓
KGED ⁴¹	✗	✗	✗	✗	✗	✗	✓ [°]	✓	✗
Du, 2021 ⁴²	✗	✗	✗	✗	✗	✓	✗	✗	n.a.
CIPHER-SC ⁴³	✗	✗	✗	✗	✗	✗	✓	✓	✗
Speos (Ours)	✓	✓	✓	✓	✓	✓	✓	✓	✓

712
713

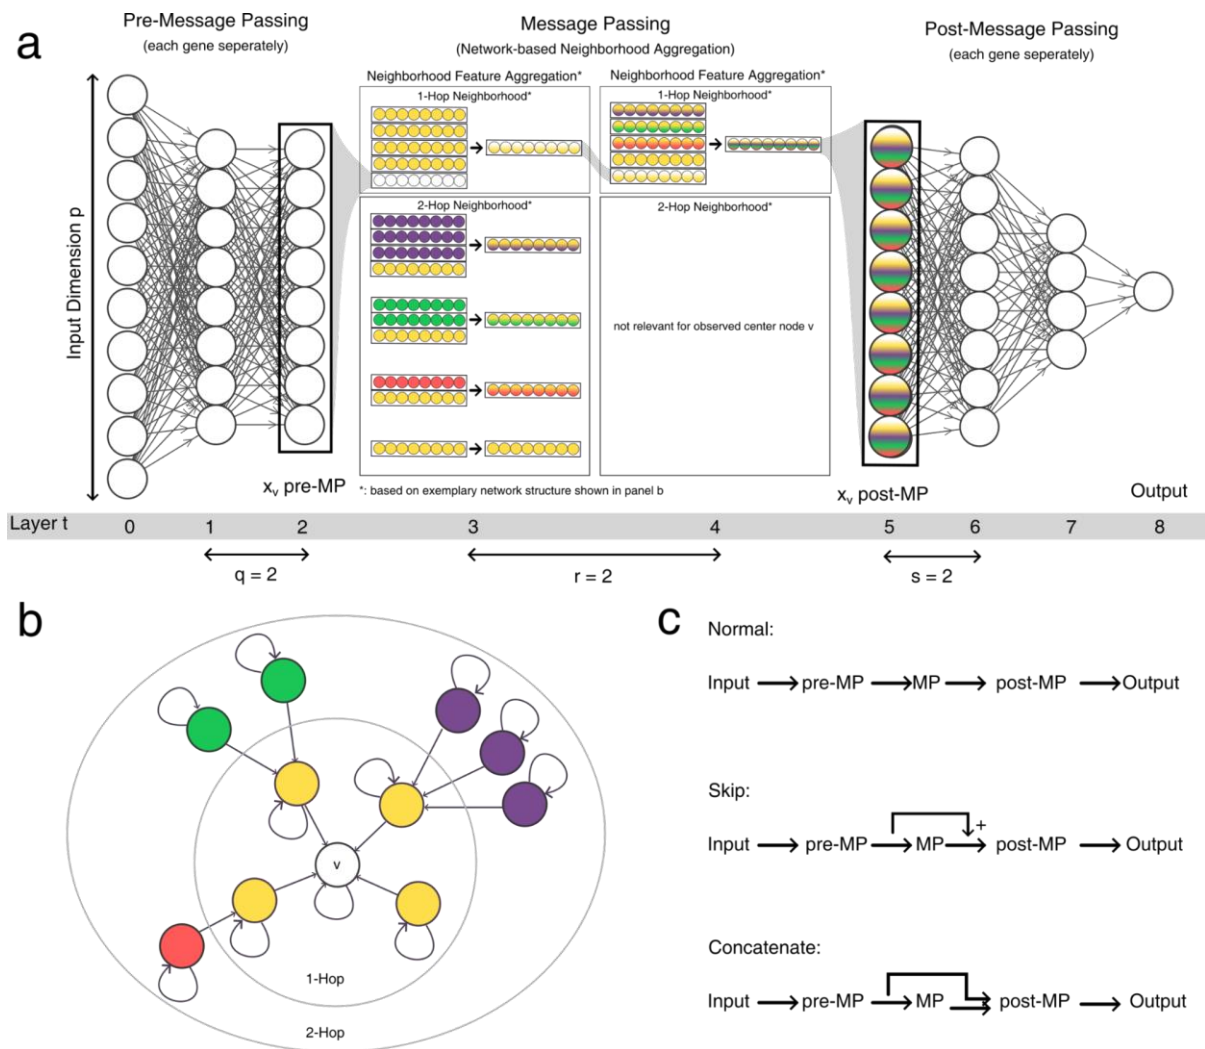
714 **Extended Data Fig. 2 | Methods Comparison.** *) assumptions about the prior class distribution of
715 the PU learning problem in form of arbitrary cut-off that are imposed on rank distributions to divide
716 candidate and non-candidate genes, in form of predetermined proportion of unlabeled positives. **)
717 positive if the article also reports results using only unbiased networks or if the method can be
718 reduced to only use unbiased networks without having to re-implement it. P-“Unbiased” networks
719 represent large-scale, systematic experiments whereas “biased” networks stem from aggregating
720 small-scale literature or using disease or gene ontologies. ***) within reason, i.e. without having to
721 re-implement the method. ****) negative if the labels for validation are sourced from the same
722 database as the input data, networks or training labels. †) uses label propagation instead of GNNs.
723 °) Uses knowledge graph embedding models instead of GNNs.
724

725
726




727
728

Extended Data Fig. 3 | Additional performance metrics. **a**, the mean rank of all held out positives, ranked individually (i.e. filtered) against all unlabeled genes (lower is better), for different base classifiers, dataset variants and phenotypes. Combinations of methods and input data are indicated along the y-axis. The blocks group models using common input data as indicated: Only Network: adjacency matrix/matrices; Only Features: gene expression and GWAS input features but no adjacency matrices; No Expression: GWAS input features and adjacency of individual (single) or multiple (multi) networks; Network + Features: adjacency of individual (single) or multiple (multi) networks, GWAS and gene expression. **b**, the mean area under the precision recall curve (AUPRC) metric (higher is better) for different base classifiers, dataset variants and phenotypes. Combinations of methods and input data are indicated along the y-axis. The blocks group models using common input data as indicated: Only Network: adjacency matrix/matrices; Only Features: gene expression and GWAS input features but no adjacency matrices; No Expression: GWAS input features and adjacency of individual (single) or multiple (multi) networks; Network + Features: adjacency of individual (single) or multiple (multi) networks, GWAS and gene expression.



743
744

745 **Extended Data Fig. 4 | GNN Model Architecture.** **a**, the general model architecture of all GNN
 746 models used in the experiments. The input features of node v are transformed into latent space by
 747 the pre-message passing module, which produces the latent vector x_v pre-MP. This latent vector is
 748 fed into the message-passing module, where the neighborhood feature aggregation takes place
 749 according to the graph shown in panel **b**. Each layer aggregates one hop in the network. Arrows
 750 denote the aggregation operators of the respective GNN layers described in the Methods section.
 751 After message passing, the latent vector x_v post-MP contains information of its n -hop neighborhood
 752 and is fed into the post-message passing module, which predicts the class of node v . The
 753 hyperparameters q , r and s control the number of layers per module. Not shown are nonlinearity
 754 functions and normalization layers. **b**, the simplified graph structure for the message passing shown
 755 in **a** with the observed node v in the center. Arrowheads denote the direction of the message passing;
 756 circles denote the respective n -Hop neighborhoods. **c**, normal versus alternative information flow
 757 through the network. Most commonly, all modules are chained consecutively, each feeding its output
 758 to the next. In the 'Skip' setting, the output vectors of the pre-MP and of the MP are summed up
 759 before being fed into the post-MP module. In the 'concatenate' setting, the output vectors of the pre-
 760 MP and of the MP are concatenated before being jointly fed into the post-message passing module.
 761 In this setting, the first layer of the post-message passing module has twice the number of
 762 dimensions.



763

764

765

766

767

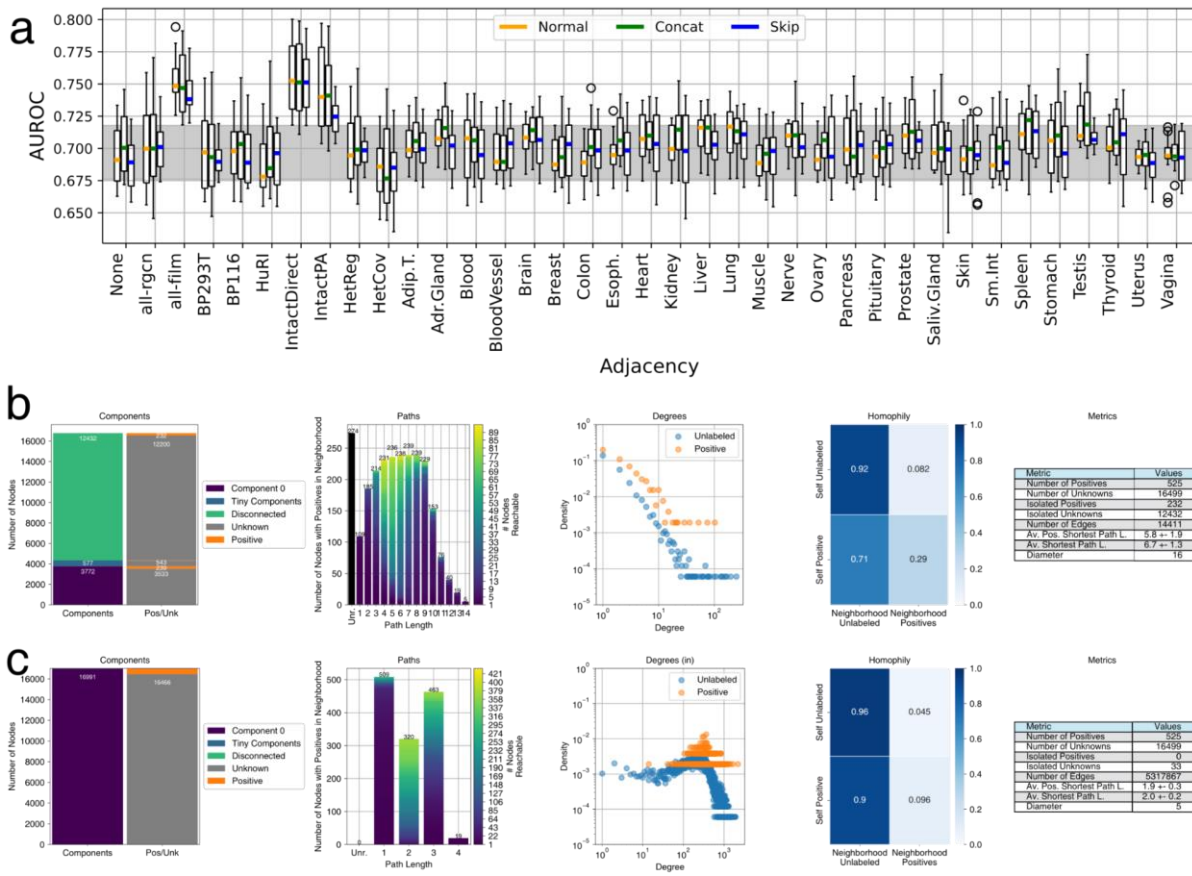
768

769

770

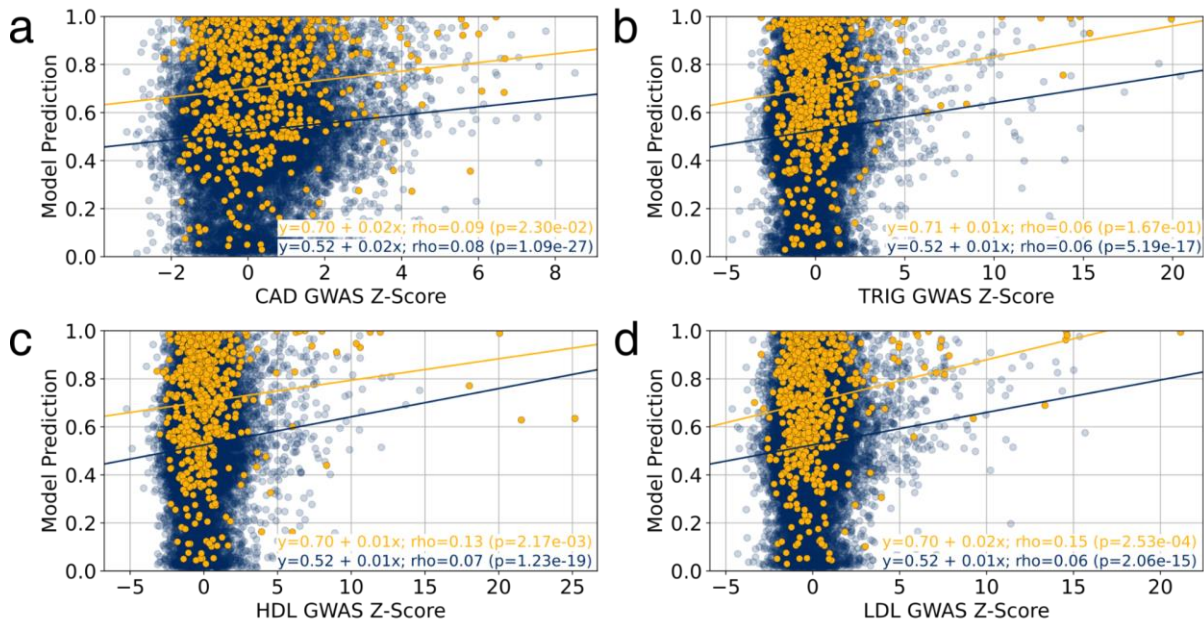
771

Extended Data Fig. 5 | GNN Depth. Influence of the number of graph convolution layers r on model performance using the adjacency matrices **a**, BioPlex 3.0 HEK293T, **b**, GRNdb Adipose Tissue, **c**, IntAct Direct Interaction. Zero convolutional layers correspond to an MLP. The gray bar in the background denotes the interquartile range of all MLP-runs. Each boxplot is based on $n = 16$ values. Boxes represent the interquartile range, colored bars are medians, whiskers extend at most 1.5 times the interquartile range, and outliers are shown individual. The color coding indicates the type of GNN (see **Supplementary Note 2** for more details).



772
773
774

775 **Extended Data Fig. 6 | Network Performance and Properties.** **a**, Boxplots of model performance
 776 (y-axis) for different adjacency matrices (x-axis). Adjacency “None” refers to an MLP that does not
 777 use any graph information. Boxes represent the interquartile range, colored bars are medians,
 778 whiskers extend at most 1.5 times the interquartile range, and outliers are shown individual. The
 779 gray bar in the background denotes the interquartile range of all MLP-runs. “Normal” indicates the
 780 normal information flow from pre-MP to MP to post-MP (**Extended Data Fig. 4c**). “Concat” indicates
 781 that the output of pre-MP is concatenated to the output of MP before being passed into post-MP.
 782 “Skip” indicates that the output of pre-MP is added to the output of MP using a sum operation before
 783 being passed into post-MP. **b**, Network properties of IntAct Direct with the label set for immune
 784 dysregulation. **c**, Network properties of all networks merged together with the label set for immune
 785 dysregulation. Components: The left bar shows the fraction of the network that is either in the largest
 786 connected component (component 0), in microcomponents (smaller than 1% of all nodes), or
 787 isolated nodes which have no incident edge, right bar shows the distribution of labeled and unlabeled
 788 nodes. Paths: Each bar shows the number of positives which have other positives in the
 789 neighborhood of the indicated size. Color indicates the number of positives in the neighborhood for
 790 each node according to scale on the right. The black bar on the left indicates the number of isolated
 791 positives. Degrees: Degree distributions of positives and unlabeled nodes. Homophily: Plot shows
 792 the percentage of nodes in the neighborhood of a node that either share the same label or have the
 793 opposite label. Metrics: additional metrics of the graphs.



794
795

796 **Extended Data Fig. 7 | Relation of gene-level GWAS association and core gene prediction.**
797 Scatterplots show the relation between the model prediction of a single FiLM prediction model trained
798 for cardiovascular disease (y-axis) and the gene-level association z-scores computed with MAGMA⁸⁵
799 (x-axis) from different GWAS studies: yellow points represent Mendelian disorder genes; blue points
800 represent unlabeled genes. Predictions are obtained on the holdout set. **a**, coronary artery disease
801 (CAD), **b**, triglyceride levels (TRIG), **c**, high density lipoprotein levels (HDL) and **d**, low density
802 lipoprotein levels (LDL).

a

Method	Degree	CS ≥ 1			CS ≥ 8		
		Candidates	Non-Candidates	<i>p</i> , OR	Candidates	Non-Candidates	<i>p</i> , OR
FiLM	$d > 0$	186	4366	$p = 6.19e-08$	41	4511	$p = 3.24e-04$
	$d = 0$	299	11885	OR = 1.69	50	12134	OR = 2.21
TAG	$d > 0$	715	3837	$p = 8.99e-251$	276	4276	$p = 6.82e-122$
	$d = 0$	182	12002	OR = 12.3	30	12154	OR = 26.1
N2V+MLP	$d > 0$	218	4334	$p = 7.52e-03$	43	4509	$p = 5.48e-01$
	$d = 0$	469	11715	OR = 1.26	129	12055	OR = 0.89
FiLM	$d > 0$	168	4384	$p = 4.98e-02$	21	4531	$p = 9.00e-01$
Unbiased	$d = 0$	375	11809	OR = 1.21	59	12125	OR = 0.95

b

Method	Degree	CS ≥ 1			CS ≥ 8		
		Candidates	Non-Candidates	<i>p</i> , OR	Candidates	Non-Candidates	<i>p</i> , OR
FiLM	$d > 0$	665	3771	$p = 1.58e-160$	143	4293	$p = 2.47e-60$
	$d = 0$	340	11723	OR = 6.08	20	12043	OR = 20.1
TAG	$d > 0$	1277	3159	$p = 0.00$	212	4224	$p = 2.67e-123$
	$d = 0$	0	12063	OR = <i>inf</i>	0	12063	OR = <i>inf</i>
N2V+MLP	$d > 0$	348	4064	$p = 2.52e-16$	103	4309	$p = 3.22e-14$
	$d = 0$	536	11381	OR = 1.82	93	11824	OR = 3.04
FiLM	$d > 0$	238	4198	$p = 5.76e-14$	72	4364	$p = 3.87e-10$
Unbiased	$d = 0$	341	11722	OR = 1.95	66	11997	OR = 3.0

803
804

805

806

807

808

809

810

811

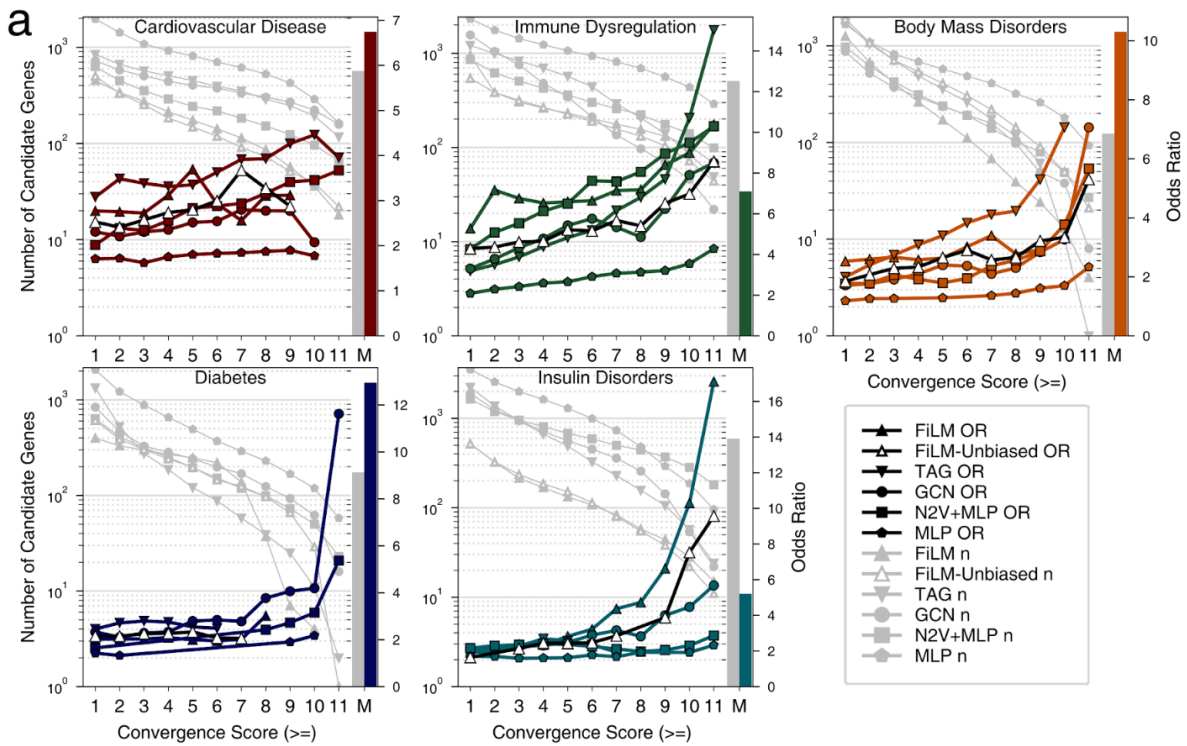
812

813

814

815

Extended Data Fig. 8 | Biased results from aggregating small-scale Literature. **a**, shown are the 2 x 2 contingency tables of candidates and noncandidates with two different Convergence Score (CS) cutoffs for cardiovascular disease and their involvement in the IntAct Direct Interaction network. Degree $d > 0$ denotes genes that have at least one incident edge in IntAct Direct Interaction, while $d = 0$ denotes genes that are isolated in IntAct Direct Interaction. Odds Ratios (OR) and *P*-values obtained via Fisher's exact test, unadjusted. **b**, shown are the 2 x 2 contingency tables of candidates and noncandidates with two different Convergence Score (CS) cutoffs for immune dysregulation and their involvement in the IntAct Direct Interaction network. Degree $d > 0$ denotes genes that have at least one incident edge in IntAct Direct Interaction, while $d = 0$ denotes genes that are isolated in IntAct Direct Interaction. Odds Ratios (OR) and *P*-values obtained via Fisher's exact test, unadjusted.



b

	Cardiovascular Disease					Immune Dysregulation					Body Mass Disorders				Diabetes		Insulin Disorders					
	Coronary Artery Disease n=473	Atrial Fibrillation n=49	Aortic Aneurysm n=41	Ischemia n=110	Hypertension n=361	Atherosclerosis n=124	Crohn's Disease n=89	Ulcerative Colitis n=101	Lupus Erythematosus n=124	Rheumatoid Arthritis n=364	Multiple Sclerosis n=192	Obesity n=619	Decreased Body Weight n=34	Increased Body Weight n=84	Abdominal Symptom n=26	Diabetes Mellitus n=475	Hyperglycemia n=49	Obesity n=619	Decreased Body Weight n=34	Increased Body Weight n=84	Atherosclerosis n=126	Coronary Artery Disease n=476
Mendelian (Ground Truth)	8.8	17.1	16.9	8.3	8.5	15.1	10.1	12.0	19.5	5.3	7.7	20.0	13.6	7.0	11.6	28.1	41.7	5.8	13.4	4.2	12.1	8.2
FiLM	5.5	6.5	4.4	5.5	4.1	5.8	14.1	8.0	6.9	6.7	5.7	2.9	2.9	2.8	1.1	4.1	8.6	2.1	8.3	2.2	7.2	5.5
FiLM-Unbiased	5.5	2.1	3.9	4.0	3.5	5.1	18.9	10.5	7.2	8.3	6.3	2.9	0.8	3.4	1.1	7.1	7.9	1.8	4.3	1.7	7.4	4.7
TAG	7.4	6.2	6.5	8.8	4.6	8.3	7.0	4.2	6.3	5.1	5.5	2.6	3.4	3.3	2.7	5.2	10.6	1.9	3.1	1.7	2.8	2.4
GCN	4.7	2.2	4.9	3.3	2.5	4.6	8.8	5.3	7.0	5.0	5.6	3.4	1.8	3.6	0.7	6.5	8.4	1.9	2.5	2.5	3.0	3.4
N2V+MLP	4.1	3.5	1.9	4.2	2.5	3.6	11.5	7.9	8.1	4.8	6.1	1.5	2.2	0.8	1.4	3.0	3.0	1.7	2.4	2.4	5.6	3.4
MLP	2.7	1.3	1.2	1.5	2.0	1.5	5.3	2.2	2.9	2.5	2.3	1.4	1.2	1.4	1.1	3.7	4.3	1.3	1.7	2.6	2.9	2.0

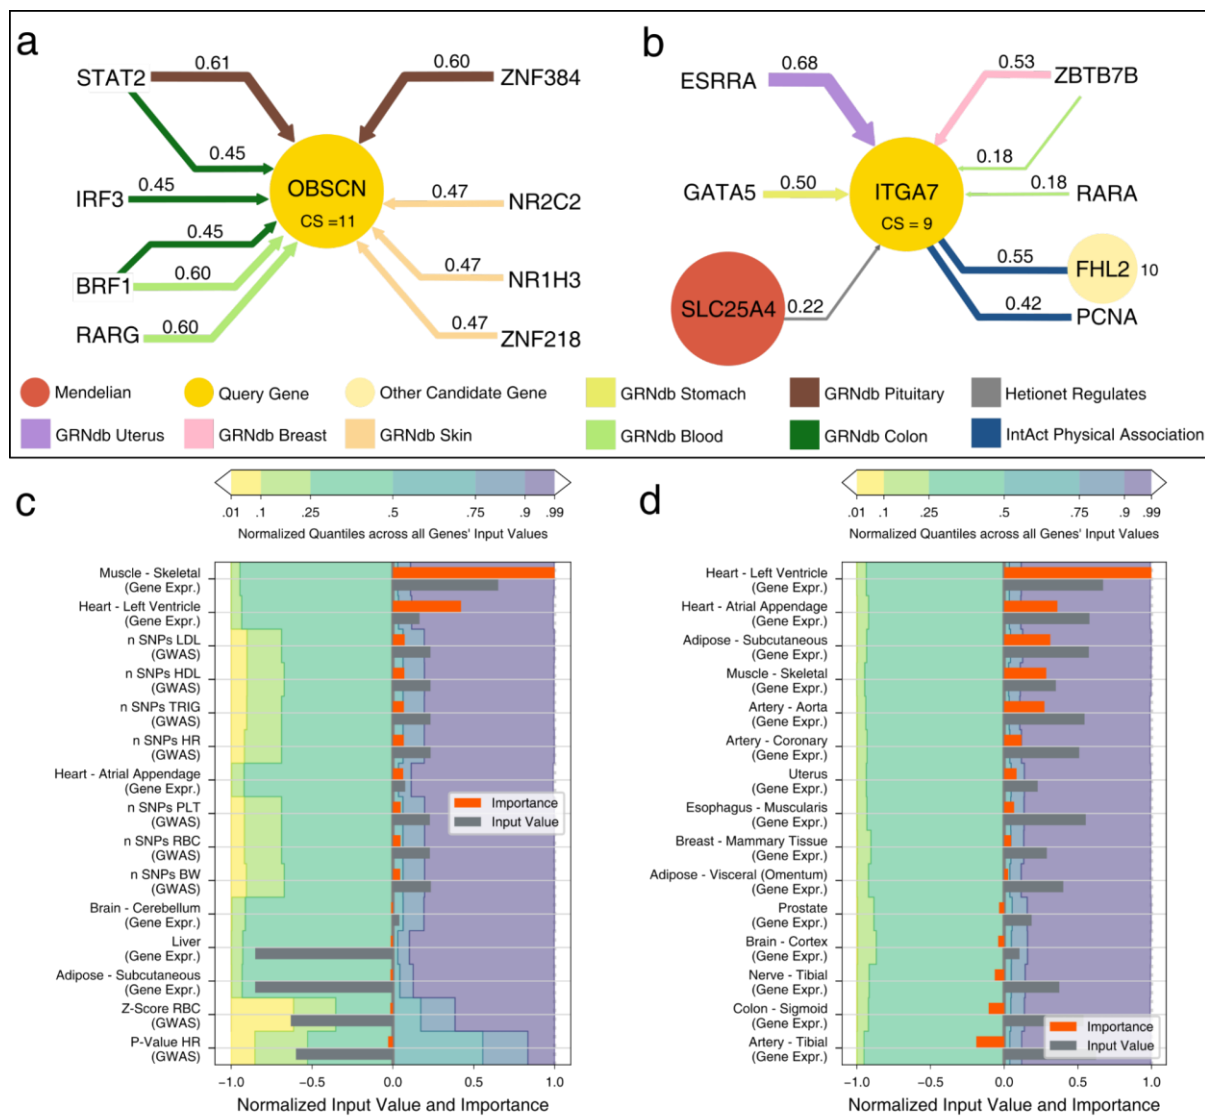
c

	Cardiovascular Disease				Immune Dysregulation				Body Mass Disorders				Diabetes				Insulin Disorders			
	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-	DT	xDC	Dr	Dr-
Mendelian (Ground Truth)	4.7	2.5	2.0	0.8	5.3	3.0	3.1	1.2	8.7	3.5	3.4	0.7	6.1	3.8	3.3	1.2	5.8	3.0	2.4	1.0
FiLM	3.5	2.0	2.1	1.1	3.1	2.2	1.7	0.7	2.8	2.0	1.5	0.8	2.5	2.2	0.8	0.2	6.3	2.2	2.7	1.9
FiLM-Unbiased	3.9	2.0	2.3	1.3	3.2	2.2	2.9	1.7	2.7	1.8	1.3	0.7	3.2	2.5	1.4	0.6	4.8	2.2	2.9	2.6
TAG	5.9	2.2	2.4	1.6	5.1	2.2	2.3	1.6	2.9	2.0	2.2	1.4	3.4	2.0	2.1	1.3	2.2	1.5	1.6	1.2
GCN	4.4	1.5	1.8	1.2	2.4	1.8	2.6	1.8	2.3	2.0	2.4	1.6	3.9	1.8	3.0	2.2	2.4	1.5	1.6	1.2
N2V+MLP	2.7	2.0	1.9	1.3	2.7	1.8	2.7	1.7	1.4	1.2	1.8	1.4	1.8	1.8	1.8	1.5	3.6	1.8	2.3	1.5
MLP	1.9	1.2	1.3	0.8	1.4	1.2	1.4	1.3	1.0	1.2	1.4	1.3	1.5	1.2	1.7	1.2	1.8	1.2	1.3	1.0

816
817

818 **Extended Data Figure 9 | Unbiased external validation.** a, Odds ratio (OR) (right y-axis) for
819 observing disease relevant phenotypes in mice with knock-outs of orthologs of candidate core genes
820 in the indicated convergence score bins (x-axis) of the five classifier methods (colored lines). Gray
821 lines indicate strength of candidate gene sets (left y-axis) in the corresponding bin for the phenotypes
822 as indicated in the panel. Only ORs with an FDR < 0.05 are shown. Bars to the right of each plot (M)

823 indicate set strength (gray) and OR (colored) of Mendelian genes for each phenotype. Precise P -
824 values, FDR, and n for each test are shown in **Supplementary Table ST3. b**, Odds ratios (ORs) of
825 Mendelian genes (first row) and of candidate genes of the five selected methods (rows) for common
826 complex subtypes of the five Mendelian disorder groups. ORs with a FDR > 0.05 in gray. **c**,
827 Enrichment of drug targets and druggability in Mendelian disorder genes and indicated candidate
828 gene sets. DT: OR of known drug targets. xDC: Ratio of median number of drug-gene interactions
829 per candidate gene to the median of non-candidates, only genes with drug-gene interactions are
830 considered. Dr: OR of druggable genes. Dr-: OR of druggable genes, after all drug targets have been
831 removed. Odds Ratios with FDR > 0.05 are grayed out. For all panels, precise P -values, FDR, and
832 n for each test are shown in **Supplementary Tables ST7, ST11**.



833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

Extended Data Fig. 10 | Model Interpretation. **a.** Most important edges for FiLM's prediction of OBSCN as candidate gene for cardiovascular disease. Shown are HGNC gene symbols. The query gene node is shown in the center, with adjacent nodes connected by relevant edges in the periphery. Candidate genes are signed with their Consensus Score (CS). The color of the edges denotes the network and the strength of the edge shows the relative importance for the prediction of the query gene which is also written at the edge. Arrowheads indicate direction of edges, undirected edges have no arrowheads. A value of 1 means that it is the most important edge for all models of the ensemble, while a value of 0 indicates that it is the least important edge for every model. Shown are 10 out of 4.5 million edges, 312 of which are in the direct neighborhood of the query gene. **b.** Most important edges for FiLM's prediction of ITGA7 as candidate gene for cardiovascular disease. Shown are 8 out of 4.5 million edges, 316 of which are in the direct neighborhood of the query gene. **c,d:** Input feature importance for OBSCN and ITGA7 alongside the respective feature's input value, compared to the input values of other genes by the quantile borders in the background. Shown are the 10 features with the strongest positive influence and the 5 features with the strongest negative influence. Negative input values are normalized to the interval [-1; 0] and positive input values to (0 ; 1] for visualization. Gray bars exceeding the colored areas are either below the 1% quantile or above the 99% quantile of that input feature. Importance values are obtained by integrated gradients and normalized to the interval [-1; 1]. Positive importance values are in favor of the prediction as candidate genes, negative importance values are attributed to features that contradict the prediction.

854 **METHODS**

855

856 **Speos: An ensemble-based PU learning framework**

857 Speos is a fully equipped Python framework which manages the modeling of networks and input
858 modalities as well as the training, evaluation and validation of several machine learning (ML)
859 methods for the prediction of novel core gene candidates. It is available at
860 <https://github.com/fratajcz/speos> and allows the configuration of experiments, including those
861 reported in this article, without the necessity to write or read any code, facilitating the uptake of
862 computational methods. For our experiments we used the Python version 3.7.12. Furthermore, it is
863 fully extensible, as input data, networks, label and validation sets as well as ML methods can be
864 chosen and added by the user. The following sections describe the modeling, training and validation
865 approaches of Speos as they are used in the experiments shown in the manuscript.

866 To fully exploit all available data for training and to avoid overestimating the performance of the
867 model, we first conduct a hyperparameter optimization on the full data set assuming negative labels
868 for unlabeled genes to find promising base classifiers and then proceed with an ensemble approach,
869 which we evaluate on independent data sources.

870 **Selection of base classifiers by cross validation**

871 We first optimize hyperparameters of base classifiers to identify the best setting of the model
872 architecture based on the performance on immune dysregulation and cardiovascular disease (see
873 **Hyperparameter optimization** and **Supplementary Note 2**). Next, we apply these optimal
874 hyperparameters to all diseases and select the most promising base classifiers for the ensemble
875 approach. We performed a fourfold cross validation with four repeats per fold, each holding out 25%
876 of positive and unlabeled genes. We assume negative labels for unlabeled genes and compare
877 performance by mean AUROC on the holdout set.

878 **The Ensemble Approach**

879 PU learning describes a subdomain of ML approaches for problems where a small set of data points
880 (in our case genes) is labeled positive and the rest of the data points are unlabeled. An intrinsic
881 challenge for this class of problems is that the number of true positives, i.e. the prior class
882 distribution⁸⁶⁻⁹¹, is unknown and most classifiers require labels for training. Motivated by the
883 robustness and the performance of ensemble approaches such as bagging in PU learning^{39,86,87}, we
884 develop a statistical approach to separate candidate genes from non-candidate genes using an
885 ensemble approach^{87,88,92} which eliminates the need to predefine³⁹ or estimate⁸⁸ a prior class
886 distribution or to choose an arbitrary cut-off^{40,42} on predicted rank distributions. At the heart of Speos
887 is the cross validation ensemble consisting of m outer folds, each containing n models. It is an
888 approach to maximize the utilization of scarce, strong labels and simultaneously exploit the
889 constraints of the genetic domain while satisfying the assumptions of the positive-unlabeled training
890 regimen. In addition to the two-step approach and the ensemble learning, we introduce the following
891 measures to improve PU learning: we designed a specific loss function that upweights positives and

892 we employed a variant of the stochastic gradient descent algorithm that downsamples negatives
893 inspired by refs ^{93,94,86}.

894 **Nested cross-validation**

895 In each outer fold $j \in \{1, 2, \dots, m\}$ the positive labels are split up into the training set $train_j$ and the
896 hold-out test set $test_j$. All positives in $test_j$ are treated as unknown and consequently labeled as
897 negatives (class $y = 0$) during training. The remaining positives are labeled as class $y = 1$ during the
898 training. There are two options to define the hold out sets: 1) define hold out sets containing positives
899 and negatives (i.e. unlabeled examples) or 2) define the hold out sets to only contain positives. In
900 option 1) the held out negatives do not contribute to the loss function during training, whereas in
901 option 2) all negatives contribute to the loss function during training. Therefore, in option 2) the
902 additional negatives increase the loss if they are unknown positives while they would not contribute
903 to the loss in option 1). In general, the model will only produce supposed false-positive predictions if
904 alternative parameters increase false negative predictions, i.e. decrease sensitivity. Thus, by design
905 of the loss function, such a change of parameters results in a greater loss than “admitting” the “false-
906 positive” predictions of the unlabeled. However, only in option 2 this trade-off is reflected in the overall
907 loss across all negatives used for training. In option 1, the prediction of the held out negatives would
908 have no implication on model parameters, thus failing to induce a loss-guided trade-off between false
909 positives and false negatives. The penalty of making these errors is stronger in option 2 because it
910 was applied to the positively predicted candidate genes that are selected from the training set.
911 Therefore, this leads to even fewer positive predictions overall, i.e. more stringent predictions and
912 thus a more conservative choice than option 1.

913 Each model $i \in \{1, 2, \dots, n\}$ of the inner cross validation is trained on the entirety of unlabeled genes
914 treated as negatives ($y = 0$) and the subset $train_j$ of positives (**Fig 3a**). The set of positives $train_j$
915 is used for all models in cross validation run j , but each inner model $model_{j,i}$ divides it further into
916 $train_{j,i}$ and $val_{j,i}$. Since our holdout $val_{j,i}$ set contains only positives, we quantify overfitting by
917 measuring precision pr on the training data and recall rec on the holdout set for the performance
918 measure $f_1 = 2(pr \cdot rec)$. We train for a maximum of 1000 epochs and always retain model
919 parameters corresponding to the maximum f_1 , but we stop training if f_1 did not improve during 100
920 epochs from the maximum. Within each outer fold j , each model i produces a prediction $\hat{y}_{i,j}^g = 1$ for
921 every gene g if the model prediction is greater than 0.7 and 0 otherwise. The global holdout set of
922 $test_j$ is not accessible for any model in outer fold j . We compute the number of concordant
923 predictions for each gene $c_j^g = \sum_{i=1}^n (\hat{y}_{i,j}^g)$ for this given cross validation run j . Each gene is
924 considered a candidate gene if $c_j^g \geq c^*$ and forwarded to the outer cross validation. The inner
925 threshold c^* is introduced in the next section.

926 **Calculation of the inner threshold**

927 To assess if predictions at any threshold have higher concordance than expected by chance, and
928 hence are potentially meaningful, we set aside a global holdout set $test_j$ for every outer fold j (**Fig.**

929 **2b).** We then quantify the overlap of the held-out positive genes g in $test_j$ with concordant predictions
930 of c^* models as $C_j = |\{g | c_j^g = c^* \wedge g \in test_j\}|$. To obtain a background model for the distribution of
931 model overlaps c_j^g we setup n random classifiers \hat{i} that produce the same number of positive
932 predictions $\hat{y}_{i,j} = 1$ as the original models and analogously count the overlaps of their ‘predictions’
933 as \bar{c}_j^g . We quantify the overlap of the held-out positive genes g in $test_j$ with concordant predictions
934 of c^* random models as $\bar{C}_j =$. This is repeated $B = 1000$ times to obtain an empirical background
935 distribution. Finally, we compare C_j against \bar{C}_j using a one sample t-test for each value of $c^* \in$
936 $\{1, 2, \dots, n\}$ and apply FDR across the n tests. We choose the inner threshold as the minimal c^* where
937 C_j is significantly greater than the random expectation \bar{C}_j (FDR < 0.05, Student’s t-test,
938 **Supplementary Table ST1**) or if \bar{C}_j is smaller than 0.1. All positively predicted unlabeled genes
939 which reach at least the inner threshold are considered candidate genes for this outer fold j .

940 **Ensemble prediction**

941 The m candidate gene sets, one from every fold of the outer cross validation, are assessed for
942 overlapping genes to arrive at a Consensus Score (CS) for every gene. The CS reflects the number
943 of outer folds, which has predicted the gene as a candidate gene. Thus, the CS ranges from 0 to m ,
944 with 0 indicating that the gene has never been chosen as a candidate and thus is termed a non-
945 candidate gene. Candidate genes have a CS of 1 to m , where 1 indicates the least and m the most
946 stringent cutoff.

947

948 **Input Data**

949 **Labels**

950 Freund *et al.* have recently defined 20 classes of Mendelian disorders which resemble common
951 complex diseases. This defines sets of Mendelian disorder genes for each of the disease groups⁹⁵.
952 Importantly, Mendelian genes “clearly fulfill the core gene definition”¹⁴. Thus, we have chosen the
953 Mendelian gene sets proposed by Freund *et al.* as reliable “known positives” for each disease group.
954 The Mendelian disorder gene sets are originally obtained by querying the Online Mendelian
955 Inheritance in Men (OMIM)⁹⁶ database for symptoms and phenotypes that relate to the diseases⁹⁵.
956 In total we use 598 Mendelian disorder genes for cardiovascular disease, 550 for immune
957 dysregulation, 128 for body mass disorders, 182 for diabetes, and 623 for insulin disorders.

958 Disease gene prediction is inherently a positive unlabeled (PU) learning problem. Despite this, it is
959 a common approach to compose a supposedly “reliable” negative training set to transform the
960 problem from a PU learning task into semi-supervised classification^{39,40}. Precise negative training
961 sets are inherently difficult, if not impossible, to obtain as this requires a positive demonstration that
962 a given gene has no function in a specific, or even a panel of diverse diseases. In light of the
963 modification of genetic risk by genetic variation and environmental factors it requires immense
964 resources to demonstrate the lack of involvement, which renders this approach essentially
965 impossible, if a statistically meaningful negative training set is required. Alternative approaches make

966 assumptions about the nature of disease genes and then define negatives that contrast these
967 assumptions. In different contexts this has shown to lead to very strong biases^{97,98}, since even
968 inconspicuous household genes host a higher-than-average rate of disease genes⁹⁹. Moreover,
969 using negatives that are most dissimilar to the positives in the input space encourages ML algorithms
970 to find trivial solutions, artificially inflating performance metrics while leading to suboptimal results.
971 In light of these substantial challenges, we decided to use a PU learning approach for core disease
972 gene identification and rely on an internal threshold and external validation to assess precision of
973 the results.

974 **Nodes and node features**

975 In the following we define an input matrix $X^{(0)} \in \mathbb{R}^{n \times p}$, which is used for all experiments where the
976 number of nodes n and the number of features per node p depend on the disease and data
977 availability. The full list of nodes contains $n_{full} = 19220$ protein coding genes. We use tissue-wise
978 median gene expression and GWAS summary statistics as input features, which have to be available
979 for every gene. For the gene expression we use GTEx v7 data which has been obtained via RNASeq
980 across 44 human tissues encoded as median transcript per million (TPM)¹⁰⁰ across all GTEx samples
981 of one tissue. Additionally, we use normalized gene expression levels for 18 different blood cells and
982 total peripheral mononuclear blood cells (PBMC) from the human protein atlas¹⁰¹. As GWAS traits
983 we used those that are mapped to the relevant Mendelian disorder by ref. ⁹⁵ (**Supplementary Table**
984 **ST13**). We gathered genome-wide summary statistics from a collection of GWAS of 114 traits that
985 were assembled by the GTEx consortium¹⁰² and are available on zenodo¹⁰³. We converted our
986 protein names/gene symbols to Entrez gene ids and mapped them to the corresponding annotations
987 on the human genome assembly 38. Next we aggregated the GWAS summary statistics on the gene-
988 level using MAGMA⁸⁵ based on the positional overlap of SNPs and gene annotations with an
989 upstream/downstream window of 10 kb. Finally we used 8 GWAS traits for cardiovascular disease
990 (BW, CAD, HDL, HR, LDL, RBC, PLT, TRIG), 7 for immune dysregulation (CD, CEL, IBD, MS, RA,
991 SLE, UC), 7 for body mass disorders (BMI, BW, HDL, FAT, T2D, TRIG, WHR), 6 for (monogenic)
992 diabetes (BW, HDL, FAT, T1D, T2D, TRIG) and 4 for insulin disorders (BMI, CAD, FG, WHR). We
993 used the total number of SNPs per gene, the gene-level P -value and the gene-level Z -score for every
994 trait as input node features for Speos. Genes for which at least one of the mentioned input features
995 could not be gathered are excluded from the analysis. This leaves $n = 17320$ out of $n_{full} = 19220$ for
996 cardiovascular disease, $n = 17042$ for immune dysregulation, $n = 17398$ for body mass disorder, n
997 $= 17460$ for diabetes and $n = 17401$ for insulin disorders (see **Supplementary Table ST13**).
998 Finally, all input features were scaled by quantiles using scikit-learn's (*v1.0.2*) RobustScaler¹⁰⁴ to
999 facilitate the processing in neural networks. Unlike gaussian normalization, this method is more
1000 robust to outliers and extreme skewness of input features.

1001 It is important to point out that Speos is a fully extensible framework, which allows the user to add
1002 more features by adding a minimal description and a preprocessing function as outlined here:

1003 <https://speos.readthedocs.io/en/latest/extension.html#additonal-node-features>

1004 **Edges and types of networks**

1005 Network maps have been generated for different modalities of biological regulation or tissue-specific
1006 manifestations. In total, we use 33 different networks in our model. Protein-protein interaction
1007 networks (PPI), have been widely used for the analysis of the genetic background of diseases and
1008 can be obtained using a variety of methods¹⁰⁵. Affinity-purification mass spectrometry based maps
1009 predominantly identify stable complexes and contain a mix of direct and indirect associations¹⁰⁵. For
1010 this, we use the systematically collected BioPlex 3.0 HEK293T and BioPlex 3.0 HCT116¹⁹ (accessed
1011 17.3.22). Additionally, we use the Human Reference Interactome (HuRI) (accessed 17.3.22), which
1012 has been obtained using a binary multi-assay mapping pipeline, which identifies predominantly
1013 directly contacting proteins¹⁸. Both BioPlex and HuRI were generated in systematic experimental
1014 approaches. Additional PPI network data are derived from the IntAct database²² (accessed 11.5.22),
1015 which is a manually curated and annotated source of protein-protein interactions. For our analysis,
1016 we use only human interactions and further filter them into two subsets. The first contains all
1017 interactions that have been labeled as “Physical Association” including its subcategories, and
1018 includes, analogous to AP-MS-based data, direct and indirect protein associations e.g. in large
1019 complexes or mediated by rRNAs in the ribosome. The second category “direct interactions” is a
1020 strict subset of IntAct Physical Association and requires unambiguous evidence for direct interactions
1021 using biochemically purified proteins. In contrast to the systematically collected HuRI and BioPlex
1022 datasets, IntAct contains interactions sourced from hypothesis-driven small-scale studies and thus
1023 represents the biases inherent to this research approach^{18,25,26}.

1024 The next type of network that is usually imposed on genes is gene regulatory networks (GRN). Gene
1025 regulatory networks are usually directed. Edges run from a transcriptional regulator (the transcription
1026 factor - TF) to its target gene. We use 27 tissue-specific GRNs obtained from GRNdb¹⁰⁶ (accessed
1027 29.3.22). These networks have been inferred using enriched TF motifs and RNA-seq expression
1028 data of healthy human tissues from GTEX^{100,106}. Finally, we use two types of relations from Hetionet
1029 (accessed 18.3.22) to define edges^{107,108}. The relation “Gene→regulates→Gene” is a non-tissue-
1030 specific GRN that has been established from RNA-seq data by the original authors of Hetionet. The
1031 relation “Gene–covaries–Gene” captures coevolutionary patterns of two genes which has been
1032 shown to aid in disease gene prioritization¹⁰⁹. We do not include the third relation that runs between
1033 genes, “Gene–interacts–Gene”, since we already include several prime candidates for PPIs.

1034 At this point we would like to emphasize that Speos is a fully extensible framework, which allows
1035 users to add more adjacency matrices by adding a minimal description as outlined in the
1036 documentation: <https://speos.readthedocs.io/en/latest/extension.html#adding-a-network>

1037

1038 **Modeling Networks for Machine Learning**

1039 All nodes in the used networks represent genes or their encoded protein products, thus the networks
1040 represent homogeneous graphs. For our machine learning approach we model each network as a
1041 directed graph. In case of PPI networks, which are inherently undirected, we introduce two edges
1042 between two connected genes *gene_a* and *gene_b*, each going in a different direction, so that the there

1043 exist both edges $gene_a \rightarrow gene_b$ and $gene_b \rightarrow gene_a$. In case of gene regulatory networks which are
1044 inherently directed, we only model the edges running from the transcription factor (TF) to the
1045 modulated genes $gene_{TF} \rightarrow gene_b$, but not vice versa. In the experiments where multiple networks
1046 are used simultaneously, each edge is also given a type $r \in R$, which indicates the network the edge
1047 is sourced from. This means that two connected genes $gene_a$ and $gene_b$ can, but don't have to be
1048 connected by more than one edge of different edge types r_1 and r_2 : $gene_a \xrightarrow{r_1} gene_b$ and $gene_a \xrightarrow{r_2}$
1049 $gene_b$.

1050

1051 **Model Architecture**

1052 Our general model architecture for most of our base classifiers consists of three consecutively
1053 arranged modules: pre message passing, message passing and post message passing (**Extended**
1054 **Data Fig. 4a**). The pre and post message passing consist of fully connected linear layers,
1055 interspersed with exponential linear unit (ELU) activation functions¹¹⁰. Their task is to transform the
1056 dimensionality from the input dimension to the desired hidden or output layer's dimension.
1057 Additionally, they serve as feature extraction layers, where pre message passing extracts and
1058 transforms the features so that the message passing can be most efficient, and the post message
1059 passing transforms the result of the message passing into a prediction for every gene. Based on
1060 hyperparameter optimization, we have chosen two hidden layers plus the input/output layer for both
1061 pre and post message passing with a hidden dimension of 50 (see **Supplementary Note 2**). The
1062 message passing is where the actual graph convolutions take place using graph neural network
1063 (GNN) layers. Based on hyperparameter optimization (**Extended Data Fig. 5**) we have chosen two
1064 GNN layers, each followed by ELU nonlinearity and instance normalization layers¹¹¹.

1065 **GNN-based methods**

1066 GNNs have recently seen a rapid development since Kipf and Welling have proposed their seminal
1067 GCN layer⁴⁶. Since then, numerous adaptations of the GCN layer have been proposed, focusing on
1068 different weaknesses of the original formulation. We have explored 11 different types of GNNs
1069 implemented in PyTorch Geometric¹¹² (v2.0.4) and assessed their suitability for our task. Speos
1070 allows the user to choose any of these convolution layers, as well as the number of hidden layers
1071 and hidden dimensions of the network. For a detailed account of the graph convolutions we
1072 examined alongside with the resulting change in performance, see **Supplementary Note 2**. Here
1073 we introduce layers that are used throughout our work.

1074 **Graph Convolutional Network Layer (GCN)**

1075 The GCN layer is defined as follows:

$$1076 \quad X^{(t+1)} = D^{-1/2}(A + I)D^{-1/2}X^{(t)}W_t$$

1077 where t corresponds to the t -th layer of the network. Usually, self-loops are added by adding the
1078 identity matrix I to the adjacency matrix A which is then normalized by the node degree matrix D .
1079 The resulting normalized adjacency matrix is then multiplied with the node feature matrix $X^{(t)}$ and a
1080 trainable weight matrix W_t . The node-specific update rule following this layer definition, also called

1081 message passing, is defined as follows

$$1082 \quad x_v^{(t+1)} = W_t^\top \sum_{u \in N(v)} \frac{a_{v,u}}{\sqrt{d_v d_u}} x_u^{(t)} .$$

1083 where $x_v^{(t+1)}$ is the latent representation of node v at layer $t+1$, which is composed of a linear
 1084 combination of the latent representations $x_u^{(t)}$ of nodes at layer t in the neighborhood of v , $N(v)$,
 1085 weighted by an optional weight $a_{v,u}$ of the edge between u and v and the degree of the nodes, d_v and
 1086 d_u . In our experiments, all edges are weighted identically with $a_{v,u} = 1$.

1087 **Topology Adaptive Graph Convolution (TAG)**

1088 TAG⁴⁷ has been proposed to address the limitation of GCN layers to the 1-hop neighborhood of each
 1089 node, which implies that the receptive field of GCNs in the graph is directly dependent on the number
 1090 of layers. TAG contains a hyperparameter K which manages the depth (number of hops) that each
 1091 TAG layer can reach within the graph. It achieves this by using powers of the adjacency matrix

$$1092 \quad X^{(t+1)} = \sum_{k=0}^K (D^{-1/2}(A + I)D^{-1/2})^k X^{(t)} W_{t,k} .$$

1093 We use two layers of TAG with a K of 3, which means that each node's representation can be
 1094 influenced by nodes 3 hops away for each TAG layer used. It furthermore employs skip-connections
 1095 between layers so that unhelpful information can be blocked. These skip connections are encoded
 1096 in the weight matrix $W_{t,k}$ for $k = 0$, as $(D^{-1/2}(A + I)D^{-1/2})^0 = I$. Like GCN, TAG is not aware of edge
 1097 types, so it is only applied on individual networks.

1098 **Relational Graph Convolution (RGCN)**

1099 RGCN extends the idea of GCN to be aware of multiple types R of edges between nodes, denoted
 1100 as $r \in \{0, 1, \dots, |R| - 1\}$. Every layer t therefore learns separate weights $W_r^{(t)}$ of node v 's
 1101 neighborhood for each type of edge r and then sums these up

$$1102 \quad x_v^{(t+1)} = W_{root}^{(t)} x_v^{(t)} + \left(\sum_{r \in R} \sum_{u \in N_r(v)} \frac{1}{|N_r(v)|} W_r^{(t)} x_u^{(t)} \right) .$$

1103 It furthermore learns edge-independent weights $W_{root}^{(t)}$ that are multiplied with v 's node features and
 1104 added to the neighborhood representation.

1105 **Feature-wise Linear Modulation Convolution (FiLM)**

1106 The FiLM⁴⁹ GNN layer has been proposed as a generalization of several relational GCN
 1107 architectures such as relational graph convolution (RGCN)⁴⁸ or relational graph attention (RGAT)¹¹³
 1108 and is based on the idea of feature-wise linear modulation which has recently been proposed for
 1109 visual reasoning⁵⁰. As such, it introduces an offset beta and a linear coefficient gamma for every
 1110 feature of an incoming message $x_u^{(t)}$ based on the edge type r and the receiver node v

$$1111 \quad x_v^{(t+1)} = \sum_{r \in R} \sum_{u \in N_r(v)} \sigma(\gamma_{r,v}^{(t)} \odot W_r x_u^{(t)} + \beta_{r,v}^{(t)}) .$$

1112 Where σ is a nonlinearity function (Rectified Linear Unit: ReLU) and \odot is the element-wise or
 1113 Hadamard product. The coefficients $\gamma_{r,v}^{(t)}$ and offsets $\beta_{r,v}^{(t)}$ applied to every message $x_u^{(t)}$ from node u
 1114 in the neighborhood of v for each edge type r , $N_r(v)$, are obtained by training a hypernetwork g

$$1115 \quad \beta_{r,v}^{(t)}, \gamma_{r,v}^{(t)} = g(x_v^{(t)}, W_{g,r}^{(t)}) .$$

1116 so that both $W_{g,r}$ and W_r contain trainable parameters. Hypernetworks are neural networks that learn
1117 parameters of other neural networks in an attempt to increase weight-sharing and reduce model
1118 complexity and memory requirements^{114,115}. In FiLM, g is implemented as a single linear layer. In
1119 other words, FiLM modifies the message that a node u passes to a node v conditioned on the relation
1120 r and the latent representation of the receiving node v .

1121 **Node2Vec**

1122 Methods like Node2Vec⁵¹ can bridge the gap between graph-native and non-graph methods by first
1123 preprocessing the graph, embedding each node into vector space in an unsupervised setting using
1124 random walks. These embeddings can then be used by MLPs or regressions as regular input
1125 features. We used the *fastnode2vec*¹¹⁶ (*v0.0.5*) command line interface of *gensim*'s¹¹⁷ (*v4.1.2*)
1126 implementation of Node2Vec with context 5, 100 dimensions, walk length 100 and 500 training
1127 epochs on all networks simultaneously. Because Node2Vec does not use edge types, using all input
1128 networks is effectively equivalent to using a single network.

1129 **Non-GNN Methods**

1130 LINKX¹¹⁸ is an MLP-based method that first trains MLPs on the input features and adjacency matrix
1131 separately and then a third MLP that joins the information of the previous two. It has been proposed
1132 to address the shortcomings of GNNs when the first order neighborhood is heterophilous, i.e. the
1133 connected nodes do not tend to have the same label. To do so, it trains multiple MLPs: MLP_A is
1134 trained directly on the adjacency matrix, using each row of the matrix as feature vectors for the
1135 respective nodes. MLP_X is trained on the feature matrix $X^{(0)}$. Finally, MLP_f uses the concatenated
1136 latent representations produced by MLP_A and MLP_X as input and predicts the class label \hat{y} . We
1137 implemented LINKX in PyTorch¹¹⁹ (*v1.8.0*) and found that it is prone to overfitting due to the large
1138 weight matrix of the first layer of MLP_A . We have therefore placed an L1 regularization term on this
1139 matrix which we multiply with a factor α and add it to the task-specific loss. We have searched α in
1140 powers of ten from 10^0 to 10^{-5} and found the best performance with $\alpha = 10^{-2}$.

1141 The MLP used as a base-classifier resembles the general model architecture outlined above with
1142 the number of message passing layers set to 0, only leaving fully connected layers interspersed with
1143 ELU nonlinearity. Logistic regression and random forests are implemented using scikit-learn's¹⁰⁴
1144 (*v1.0.2*) `LogisticRegression` and `RandomForestClassifier` classes with balanced class weights and
1145 sample weights 2 for positives and 1 for unlabeled genes. As they are not able to directly use graph-
1146 structured data, they either only use the feature matrix $X^{(0)}$ (Only Features) or use a concatenation
1147 of $X^{(0)}$ and the latent node features obtained via Node2Vec (Network + Features).

1148 **Hyperparameter Optimization (HPO)**

1149 A systematic HPO is crucial for most machine learning purposes. We utilize a 4-fold cross validation
1150 for HPO and report the performance in recovering held out known positives, considering all unlabeled
1151 genes as negatives. We assess the area under the receiver operator characteristic curve (AUROC)
1152 as performance metric since we expect an ideal classifier to rank the known positives higher than
1153 the average unlabeled gene. To avoid a bias towards a small holdout set given our already very

1154 small set of reliable positives, each fold trains on 75% of all genes and assess holdout performance
1155 on the remaining 25%. Using the same data for HPO and the validation of the final ensembles would
1156 be considered an information leak, resulting in overestimation of model performance. This is why we
1157 evaluate the final performance of the ensembles exclusively on additional independent label sets
1158 (external validation) which are not present during the HPO. Therefore, the integrity of the training
1159 regimen is not compromised. For HPO, we train four models on each fold and report the mean of all
1160 16 resulting models. We first searched for optimal hidden dimension (data not shown), number of
1161 hidden layers and type of single-network GNN layer using a selection of networks (**Extended Data**
1162 **Fig. 5**). Then we searched for the optimal network using all 35 networks and for the optimal multi-
1163 network GNN layer using the union of all networks (**Extended Data Fig. 6a**). See **Supplementary**
1164 **Note 2** for detailed results.

1165 **Loss Function**

1166 The loss or risk function L measures the goodness of fit of the model and provides the error term
1167 from which the gradient is calculated which directly influences the tuning of model parameters via
1168 backpropagation. We use class-label 0 for unlabeled genes and class-label 1 for labeled genes and
1169 use binary cross entropy, also called logistic loss, as loss function. To reflect the uncertainty of the
1170 true label of class 0 and the strength of evidence for our label class 1, we have implemented two
1171 mechanisms for loss tuning which we refer to as *dilution* and *amplification*, inspired by ref. ^{93,94,86}.
1172 *Dilution* is a downsampling process where, for each training epoch, we gather a different random
1173 subset $U_{sampled}$ sampled uniformly with replacement from all unlabeled genes U so that $|U_{sampled}| =$
1174 $|P_{train}| \cdot d = u^*$ where P_{train} is the set of all positives in $train_{j,i}$ and d is the dilution hyperparameter.
1175 This has the advantage that not every unlabeled gene contributes to the loss term in every epoch,
1176 allowing unlabeled genes that resemble positive genes to receive a higher prediction, and balancing
1177 the contribution of unlabeled and positives to the loss term, eradicating the influence of class
1178 imbalance.

1179 The final loss function is composed as follows:

$$1180 \quad L = \sum_{u=0}^{u^*} \frac{BCE(y_u, \hat{y}_u)}{d} + a \cdot \sum_{u=0}^{|P_{train}|} BCE(y_p, \hat{y}_p)$$

1181 Where BCE stands for binary cross entropy or logistic loss, a is our amplification hyperparameter,
1182 $y_u = 0$ and $y_p = 1$. We use $d = 10$ and $a = 2$ in our experiments. For *amplification*, we sum the
1183 individual loss terms of positives used for training and multiply it with the amplification factor a . This
1184 has the effect that false-negative predictions become a times more costly than false-positive
1185 predictions. If there exists an unlabeled gene, which is indistinguishable from a known positive, both
1186 *dilution* and *amplification* result in a loss that encourages the model to predict both genes as positive
1187 (class 1) rather than both as negative (class 0). Although this might lead individual models to overfit
1188 to their positive examples in training, ensembles are expected to thrive under these
1189 circumstances¹²⁰. We optimize L via gradient descent using an Adam¹²¹ optimizer with learning rate

1190 10^{-3} .

1191

1192 **Model Interpretation**

1193 As candidate genes are predicted by an ensemble, we provide model interpretations based on the
1194 average importance of an edge or input feature across the whole ensemble. A related idea of model
1195 interpretation has recently been formulated as *model class reliance*¹²². To assess the reliance of the
1196 ensemble on certain edges and node features, we gather the respective edges' and nodes'
1197 importance using integrated gradients⁶⁷ from every model of the ensemble for a query gene c .
1198 Broadly speaking, integrated gradients assign importance values based on the change in gradients
1199 when input features are substituted with a contrast, usually a vector containing only zeros. For edge
1200 importance, this means that we introduce edge weights of 1 for every edge which are then substituted
1201 with edge weights of 0. An edge weight of 1 does not alter the message passing and an edge weight
1202 of 0 means removing the edge, while gradients backpropagated to the respective edge weights can be
1203 used for inspection. As we predict the importance based on the gradients backpropagated from gene
1204 c , the obtained importance values are valid only for the interpretation of the prediction for gene c .
1205 Each individual model's absolute integrated gradients are minmax scaled to the interval $[0, 1]$ across
1206 all nodes and edges in the graph. Minmax scaling has the advantage of a comparable output space,
1207 but has the tendency to over-emphasize negligible differences in the input space. To alleviate this
1208 problematic tendency, we use the mean value of all models' minmax scaled importance values,
1209 assuming that an important edge or input feature will repeatedly be close to 1 and an unimportant
1210 edge or input feature close to 0, leaving the intermediate values to edges and features that are of
1211 ambivalent importance.

1212 Edge Importance $I_{v,e} \in \mathbb{R}$ of edges $e \in E$ for candidate gene v over all models $i \in \{1, 2, \dots, n \cdot m\}$ from
1213 all inner and outer folds against a contrast edge weight of 0:

$$1214 \quad I_{v,e} = \frac{1}{n \cdot m} \sum_i \minmax_{v \in E} (|IntegratedGradients_i(e, 0)|)$$

1215 Note that minmax operates across the set of all edges E (union of all edges across networks in the
1216 case of FiLM). Node input feature Importance $I_{v,n} \in \mathbb{R}^p$ of for input features $f \in \mathbb{R}^p$ nodes $n' \in N$ for
1217 candidate gene v over all models m from ensemble M against a contrast vector containing only zeros
1218 ($0^{\rightarrow} \in \mathbb{R}^p$):

$$1219 \quad I_{v,n'} = \frac{1}{n \cdot m} \sum_i \minmax_{v \in N} (|IntegratedGradients_i(n', 0^{\rightarrow})|)$$

1220 To get a more detailed interpretation of node v 's own input features, we also obtain the importance
1221 $I_{v,v}$ without removing the sign of the output of integrated gradients and minmax scale it across its
1222 own dimensions:

$$1223 \quad I_{v,v} = \frac{1}{n \cdot m} \sum_i \minmax_v (IntegratedGradients_i(v, 0^{\rightarrow}))$$

1224 This way, the most important feature across all models will receive an importance score close to 1
1225 or -1, depending on the direction of its influence, and the least important feature will receive an

1226 importance score close to 0.

1227 For implementation of GNN interpretations we use the interface of PyTorch Geometric¹¹² (v2.0.4)
1228 with the PyTorch¹¹⁹-based model interpretation library Captum¹²³ (v0.4.1).

1229

1230 **External validation and core gene properties**

1231 As outlined above, we use all available positive labels for training due to their scarcity. To avoid an
1232 information leak between training and validation, we base the validation of our candidate genes on
1233 labels sourced from external datasets which are not present during training and hyperparameter
1234 optimization but reflect several characteristics of core genes.

1235 **Mouse KO Experiments**

1236 We assume that if a gene plays a pivotal role in a disease, severely disrupting the gene's function
1237 will result in a phenotype that resembles the disease. To assess this hypothesis, we gathered the
1238 same phenotypical queries that ref.⁹⁵ used to obtain the labels for the Mendelian genes
1239 (**Supplementary Table ST14**). We then used these queries to retrieve a set of genes that, if
1240 deliberately knocked out in mice, produce phenotypes that match the queries using the Mouse
1241 Genome Database (MGD)^{124,125} (<http://www.informatics.jax.org/allele>, accessed 17.3.22). We used
1242 an empty query to get a background set of all available mouse knockout genes. We then translated
1243 the mouse genes to their human orthologs using the official MGD mouse-human homolog table
1244 (<http://www.informatics.jax.org/homology.shtml>, accessed 28.11.22), entries without a human
1245 ortholog were discarded, resulting in 16370 genes. For the assessment of candidates, we removed
1246 Mendelian genes from the background sets and those genes that were excluded from the predictions
1247 due to missing input features, such that the respective intersections of 14116; 13936; 14586; 14541;
1248 14123 (**Supplementary Table ST12**) formed the background sets for the following analysis
1249 (**Supplementary Table ST2 & ST3**). Next we tested the Mendelian genes of each disease for an
1250 enrichment in mouse KO genes against all non-Mendelian genes in the background set, and the
1251 candidate genes against all non-Mendelian non-candidate genes in the background set using
1252 Fisher's exact test (**Supplementary Table ST2 & ST3**). We further tested if restricting the candidate
1253 genes to a higher consensus score increases their enrichment. To do so, we tested each CS bin for
1254 enrichment against all protein coding non-Mendelian genes with a lower CS. We adjusted the *P*-
1255 values of the multiple Fisher's exact tests by FDR correction.

1256 **Differential Gene Expression**

1257 We gathered differentially expressed genes for subcategories of cardiovascular disease and immune
1258 dysregulation by individually querying the following disease subtype in the GEMMA database¹²⁶:
1259 coronary artery disease (DOID_3393), Atrial Fibrillation (HP_0005110), aortic aneurysm
1260 (DOID_3627), ischemia (DOID_326), hypertension (DOID_10763), atherosclerosis (DOID_1936),
1261 Crohn's disease (DOID_8778), ulcerative colitis (DOID_8577), lupus erythematosus (DOID_8857),
1262 rheumatoid arthritis (DOID_7148), multiple sclerosis (DOID_2377), obesity (DOID_9970),
1263 Decreased body weight (HP_0004325), Increased body weight (HP_0004324), Abdominal symptom

1264 (HP_0011458), diabetes mellitus (DOID_9351), hyperglycemia (DOID_4195). Non-human entries
1265 were removed. We applied Fisher's exact tests (**Supplementary Table ST6 & ST7**) to look for an
1266 enrichment of differentially expressed genes in the respective gene sets.

1267 **Gene Set Enrichment Analysis**

1268 We applied gene set enrichment analysis (GSEA) to our candidate gene sets using all using the
1269 respective list of 'considered genes' as background. Gene Ontology (GO) Enrichment Analysis
1270 performs GSEA based on the GO ontology *biological process*^{127,128} (**Supplementary Table ST4**).
1271 We obtained the GO annotations through the tool GeneSCF¹²⁹.

1272 **LoF and Missense Intolerance**

1273 We gathered gene-level LoF and Missense Intolerance Z-scores from the ExAc Cohort⁵⁴ where a
1274 high value indicates a high intolerance for LoF or missense mutations, respectively. In total we
1275 obtained Z-scores for 16834 of our n_{full} of 19220 genes, which correspond to 15709 for
1276 cardiovascular disease, 15450 for immune dysregulation, 15781 for body mass disorders, 15787 for
1277 diabetes and 15784 for insulin disorders. We conducted a Tukey's Honestly Significant Difference
1278 test (**Supplementary Table ST8, ST9**) between Mendelian disorder genes, candidate genes and
1279 non-candidate genes.

1280 **Drug targets and druggability**

1281 We obtained drug-gene interactions from the Drug Repurposing Knowledge Graph⁷⁸, which has
1282 been gathered from a large compendium of databases relating genes, diseases, drugs and several
1283 other biomedical domains. We extracted only edges linking drugs and genes and removed edges
1284 that have been automatically mined from preprint servers. We considered as drug targets (DT) genes
1285 that have at least one edge to any compound and applied Fisher's exact tests (**Supplementary**
1286 **Table ST10**) to look for enrichment of drug targets in our gene sets. To analyze the drug-targeting
1287 degree we counted for all drug targets the number of drug-gene interactions. We then applied
1288 pairwise Wilcoxon rank sum tests between the counts of Mendelian disorder genes, candidates and
1289 non-candidates and adjust the P -values using FDR (**Supplementary Table ST10**). We report the
1290 fold increase of the median drug-targeting degree compared to non-candidate genes (xDC). Genes
1291 encoding druggable proteins were obtained from DGldb¹³⁰. Enrichment for "druggable genes" (Dr)
1292 in any set was assessed using Fisher's exact test. To evaluate not-targeted but druggable genes
1293 (Dr-), genes encoding products that are already targeted by a drug from the respective gene sets
1294 were removed and the remaining druggable proteins tested for enrichment using a Fisher's exact
1295 test.

1296 **REFERENCES FOR FIGURES, EXTENDED DATA FIGURES, METHODS, AND**

1297 **SUPPLEMENTARY INFORMATION**

1298

- 1299 85. Leeuw, C. A. de, Mooij, J. M., Heskes, T. & Posthuma, D. MAGMA: Generalized Gene-Set
1300 Analysis of GWAS Data. *PLOS Comput. Biol.* **11**, e1004219 (2015).
- 1301 86. Mordelet, F. & Vert, J.-P. A bagging SVM to learn from positive and unlabeled examples.
1302 *Pattern Recognit. Lett.* **37**, 201–209 (2014).
- 1303 87. Claesen, M., De Smet, F., Suykens, J. A. K. & De Moor, B. A robust ensemble approach to
1304 learn from positive and unlabeled data using SVM base models. *Neurocomputing* **160**, 73–84
1305 (2015).
- 1306 88. Yang, P., Liu, W. & Yang, J. Positive unlabeled learning via wrapper-based adaptive
1307 sampling. in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial*
1308 *Intelligence* 3273–3279 (International Joint Conferences on Artificial Intelligence
1309 Organization, 2017). doi:10.24963/ijcai.2017/457.
- 1310 89. Yao, Y. *et al.* Rethinking Class-Prior Estimation For Positive-Unlabeled Learning. in *ICLR 12*
1311 (2022).
- 1312 90. Hu, W. *et al.* Predictive Adversarial Learning from Positive and Unlabeled Data. in 9 (2021).
- 1313 91. Chen, H., Liu, F., Wang, Y., Zhao, L. & Wu, H. A Variational Approach for Learning from
1314 Positive and Unlabeled Data. in *Advances in Neural Information Processing Systems* vol. 33
1315 14844–14854 (Curran Associates, Inc., 2020).
- 1316 92. Navlakha, S. & Kingsford, C. The power of protein interaction networks for associating genes
1317 with diseases. *Bioinformatics* **26**, 1057–1063 (2010).
- 1318 93. Liu, B., Lee, W. S., Yu, P. S. & Li, X. Partially Supervised Classification of Text Documents.
1319 in *Proceedings of the Nineteenth International Conference on Machine Learning* 387–394
1320 (Morgan Kaufmann Publishers Inc., 2002).
- 1321 94. Liu, B., Dai, Y., Li, X., Lee, W. S. & Yu, P. S. Building text classifiers using positive and
1322 unlabeled examples. in *Third IEEE International Conference on Data Mining* 179–186 (2003).
1323 doi:10.1109/ICDM.2003.1250918.
- 1324 95. Freund, M. K. *et al.* Phenotype-Specific Enrichment of Mendelian Disorder Genes near
1325 GWAS Regions across 62 Complex Traits. *Am. J. Hum. Genet.* **103**, 535–552 (2018).
- 1326 96. Amberger, J. S., Bocchini, C. A., Schiettecatte, F., Scott, A. F. & Hamosh, A. OMIM.org:
1327 Online Mendelian Inheritance in Man (OMIM®), an online catalog of human genes and
1328 genetic disorders. *Nucleic Acids Res.* **43**, D789–D798 (2015).
- 1329 97. Ben-Hur, A. & Noble, W. S. Choosing negative examples for the prediction of protein-protein
1330 interactions. *BMC Bioinformatics* **7**, S2 (2006).
- 1331 98. Qi, Y., Bar-Joseph, Z. & Klein-Seetharaman, J. Evaluation of Different Biological Data and
1332 Computational Classification Methods for Use in Protein Interaction Prediction. *Proteins* **63**,
1333 490–500 (2006).

- 1334 99. Cacheiro, P. *et al.* Human and mouse essentiality screens as a resource for disease gene
1335 discovery. *Nat. Commun.* **11**, 655 (2020).
- 1336 100. GTEx Consortium *et al.* Genetic effects on gene expression across human tissues. *Nature*
1337 **550**, 204–213 (2017).
- 1338 101. Uhlen, M. *et al.* A genome-wide transcriptomic analysis of protein-coding genes in human
1339 blood cells. *Science* **366**, eaax9198 (2019).
- 1340 102. Barbeira, A. N. *et al.* Exploiting the GTEx resources to decipher the mechanisms at GWAS
1341 loci. *Genome Biol.* **22**, 49 (2021).
- 1342 103. Barbeira, A. N. *et al.* Publicly available GWAS summary statistics, harmonized and imputed
1343 to GTEx v8' variant reference. (2020) doi:10.5281/zenodo.3629742.
- 1344 104. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**,
1345 2825–2830 (2011).
- 1346 105. Braun, P. Interactome mapping for analysis of complex phenotypes: insights from
1347 benchmarking binary interaction assays. *Proteomics* **12**, 1499–1518 (2012).
- 1348 106. Fang, L. *et al.* GRNdb: decoding the gene regulatory networks in diverse human and mouse
1349 conditions. *Nucleic Acids Res.* **49**, D97–D103 (2021).
- 1350 107. Himmelstein, D. S. *et al.* Systematic integration of biomedical knowledge prioritizes drugs for
1351 repurposing. *eLife* **6**, e26726 (2017).
- 1352 108. Hetionet - An integrative network of biomedical knowledge. <https://het.io>.
- 1353 109. Priedigkeit, N., Wolfe, N. & Clark, N. L. Evolutionary Signatures amongst Disease Genes
1354 Permit Novel Methods for Gene Prioritization and Construction of Informative Gene-Based
1355 Networks. *PLOS Genet.* **11**, e1004967 (2015).
- 1356 110. Clevert, D.-A., Unterthiner, T. & Hochreiter, S. Fast and Accurate Deep Network Learning by
1357 Exponential Linear Units (ELUs). Preprint at <https://doi.org/10.48550/arXiv.1511.07289>
1358 (2016).
- 1359 111. Ulyanov, D., Vedaldi, A. & Lempitsky, V. Instance Normalization: The Missing Ingredient for
1360 Fast Stylization. Preprint at <https://doi.org/10.48550/arXiv.1607.08022> (2017).
- 1361 112. Fey, M. & Lenssen, J. E. Fast Graph Representation Learning with PyTorch Geometric.
1362 *ArXiv190302428 Cs Stat* (2019).
- 1363 113. Busbridge, D., Sherburn, D., Cavallo, P. & Hammerla, N. Y. Relational Graph Attention
1364 Networks. *ArXiv190405811 Cs Stat* (2019).
- 1365 114. Ha, D., Dai, A. & Le, Q. V. HyperNetworks. Preprint at <http://arxiv.org/abs/1609.09106>
1366 (2016).
- 1367 115. von Oswald, J., Henning, C., Grewe, B. F. & Sacramento, J. Continual learning with
1368 hypernetworks. Preprint at <http://arxiv.org/abs/1906.00695> (2022).
- 1369 116. Abraham, L. louisabraham/fastnode2vec version-0.0.5. (2020) doi:10.5281/zenodo.3902942.
- 1370 117. Řehůřek, R. & Sojka, P. *Software Framework for Topic Modelling with Large Corpora*.
1371 (University of Malta, 2010).

- 1372 118. Zhu, J. *et al.* Beyond Homophily in Graph Neural Networks: Current Limitations and Effective
1373 Designs. *ArXiv200611468 Cs Stat* (2020).
- 1374 119. Paszke, A. *et al.* PyTorch: An Imperative Style, High-Performance Deep Learning Library. in
1375 *Advances in Neural Information Processing Systems* vol. 32 (Curran Associates, Inc., 2019).
- 1376 120. Sollich, P. & Krogh, A. Learning with ensembles: How overfitting can be useful. in *Advances*
1377 *in Neural Information Processing Systems* vol. 8 190–196 (1996).
- 1378 121. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Preprint at
1379 <https://doi.org/10.48550/arXiv.1412.6980> (2017).
- 1380 122. Fisher, A., Rudin, C. & Dominici, F. All Models are Wrong, but Many are Useful: Learning a
1381 Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *J*
1382 *Mach Learn Res* **20**, 81 (2019).
- 1383 123. Kokhlikyan, N. *et al.* Captum: A unified and generic model interpretability library for PyTorch.
1384 Preprint at <https://doi.org/10.48550/arXiv.2009.07896> (2020).
- 1385 124. Bult, C. J. *et al.* Mouse Genome Database (MGD) 2019. *Nucleic Acids Res.* **47**, D801–D806
1386 (2019).
- 1387 125. Mouse genome database (MGD), MGI, The Jackson Laboratory, Bar Harbor, Maine.
1388 Accessed May 2022. <http://www.informatics.jax.org/>.
- 1389 126. Lim, N. *et al.* Curation of over 10 000 transcriptomic studies to enable data reuse. *Database*
1390 **2021**, baab006 (2021).
- 1391 127. Ashburner, M. *et al.* Gene Ontology: tool for the unification of biology. *Nat. Genet.* **25**, 25–29
1392 (2000).
- 1393 128. The Gene Ontology Consortium. The Gene Ontology resource: enriching a GOld mine.
1394 *Nucleic Acids Res.* **49**, D325–D334 (2021).
- 1395 129. Subhash, S. & Kanduri, C. GeneSCF: a real-time based functional enrichment tool with
1396 support for multiple organisms. *BMC Bioinformatics* **17**, 365 (2016).
- 1397 130. Cotto, K. C. *et al.* DGldb 3.0: a redesign and expansion of the drug–gene interaction
1398 database. *Nucleic Acids Res.* **46**, D1068–D1073 (2018).
- 1399 131. Hamilton, W. L., Ying, R. & Leskovec, J. Inductive Representation Learning on Large
1400 Graphs. *ArXiv170602216 Cs Stat* (2018).
- 1401 132. Shi, Y. *et al.* Masked Label Prediction: Unified Message Passing Model for Semi-Supervised
1402 Classification. Preprint at <https://doi.org/10.48550/arXiv.2009.03509> (2021).
- 1403 133. Defferrard, M., Bresson, X. & Vandergheynst, P. Convolutional Neural Networks on Graphs
1404 with Fast Localized Spectral Filtering. Preprint at <https://doi.org/10.48550/arXiv.1606.09375>
1405 (2017).
- 1406 134. Wu, F. *et al.* Simplifying Graph Convolutional Networks. Preprint at
1407 <https://doi.org/10.48550/arXiv.1902.07153> (2019).
- 1408 135. Veličković, P. *et al.* Graph Attention Networks. *ArXiv171010903 Cs Stat* (2017).
- 1409 136. Xu, K., Hu, W., Leskovec, J. & Jegelka, S. How Powerful are Graph Neural Networks?

- 1410 *ArXiv181000826 Cs Stat* (2019).
- 1411 137. Zhang, S.-W., Xu, J.-Y. & Zhang, T. DGMP: Identifying Cancer Driver Genes by Jointing
1412 DGCN and MLP from Multi-Omics Genomic Data. 2022.02.16.480791 Preprint at
1413 <https://doi.org/10.1101/2022.02.16.480791> (2022).
- 1414 138. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural message passing
1415 for Quantum chemistry. in *Proceedings of the 34th International Conference on Machine*
1416 *Learning - Volume 70* 1263–1272 (JMLR.org, 2017).
- 1417 139. Yang, Y. & Li, D. NENN: Incorporate Node and Edge Features in Graph Neural Networks. in
1418 *Proceedings of The 12th Asian Conference on Machine Learning* 593–608 (PMLR, 2020).
- 1419 140. Chamberlain, B. P. *et al.* GRAND: Graph Neural Diffusion. Preprint at
1420 <https://doi.org/10.48550/arXiv.2106.10934> (2021).
- 1421 141. Horn, M. *et al.* Topological Graph Neural Networks. Preprint at
1422 <https://doi.org/10.48550/arXiv.2102.07835> (2022).
- 1423 142. Cunningham, F. *et al.* Ensembl 2022. *Nucleic Acids Res.* **50**, D988–D995 (2022).
- 1424 143. Randazzo, D., Pierantozzi, E., Rossi, D. & Sorrentino, V. The potential of obscurin as a
1425 therapeutic target in muscle disorders. *Expert Opin. Ther. Targets* **21**, 897–910 (2017).
- 1426 144. Arimura, T. *et al.* Structural analysis of obscurin gene in hypertrophic cardiomyopathy.
1427 *Biochem. Biophys. Res. Commun.* **362**, 281–287 (2007).
- 1428 145. Grogan, A. *et al.* Deletion of obscurin immunoglobulin domains Ig58/59 leads to age-
1429 dependent cardiac remodeling and arrhythmia. *Basic Res. Cardiol.* **115**, 60 (2020).
- 1430 146. Lagor, W. R. *et al.* Genetic manipulation of the ApoF/Stat2 locus supports an important role
1431 for Type I Interferon signaling in atherosclerosis. *Atherosclerosis* **233**, 234–241 (2014).
- 1432 147. Lee, C.-J. *et al.* Stat2 stability regulation: an intersection between immunity and
1433 carcinogenesis. *Exp. Mol. Med.* **52**, 1526–1536 (2020).
- 1434 148. Liu, S. *et al.* ZNF384: A Potential Therapeutic Target for Psoriasis and Alzheimer’s Disease
1435 Through Inflammation and Metabolism. *Front. Immunol.* **13**, 892368 (2022).
- 1436 149. Desroches-Castan, A., Cherradi, N., Feige, J.-J. & Ciaïis, D. A novel function of Tis11b/BRF1
1437 as a regulator of Dll4 mRNA 3'-end processing. *Mol. Biol. Cell* **22**, 3625–3633 (2011).
- 1438 150. King, K. R. *et al.* IRF3 and type I interferons fuel a fatal response to myocardial infarction.
1439 *Nat. Med.* **23**, 1481–1487 (2017).
- 1440 151. Bhandari, A. *et al.* ITGA7 functions as a tumor suppressor and regulates migration and
1441 invasion in breast cancer. *Cancer Manag. Res.* **10**, 969–976 (2018).
- 1442 152. Zhang, X., Ke, S., Lu, Y. & An, H. ITGA7 relates to disease risk, pathological feature,
1443 treatment response and survival in Ph- acute lymphoblastic leukemia. *Biomark. Med.* **15**,
1444 1589–1597 (2021).
- 1445 153. Xia, W. *et al.* Case Report: A Boy From a Consanguineous Family Diagnosed With
1446 Congenital Muscular Dystrophy Caused by Integrin Alpha 7 (ITGA7) Mutation. *Front. Genet.*
1447 **12**, 706823 (2021).

- 1448 154. Bugiardini, E. *et al.* Integrin $\alpha 7$ Mutations Are Associated With Adult-Onset Cardiac
1449 Dysfunction in Humans and Mice. *J. Am. Heart Assoc.* **11**, e026494 (2022).
- 1450 155. Burkin, D. J. *et al.* Transgenic expression of $\{\alpha\}7\{\beta\}1$ integrin maintains muscle
1451 integrity, increases regenerative capacity, promotes hypertrophy, and reduces
1452 cardiomyopathy in dystrophic mice. *Am. J. Pathol.* **166**, 253–263 (2005).
- 1453 156. Esposito, T. *et al.* Digenic mutational inheritance of the integrin alpha 7 and the myosin
1454 heavy chain 7B genes causes congenital myopathy with left ventricular non-compact
1455 cardiomyopathy. *Orphanet J. Rare Dis.* **8**, 91 (2013).
- 1456 157. Argirò, A. *et al.* Sex-Related Differences in Genetic Cardiomyopathies. *J. Am. Heart Assoc.*
1457 **11**, e024947 (2022).
- 1458 158. Rowland, T. J. *et al.* Obscurin Variants in Patients With Left Ventricular Noncompaction. *J.*
1459 *Am. Coll. Cardiol.* **68**, 2237–2238 (2016).
- 1460 159. Marston, S. *et al.* OBSCN Mutations Associated with Dilated Cardiomyopathy and
1461 Haploinsufficiency. *PLOS ONE* **10**, e0138568 (2015).
- 1462 160. Ramlawi, B., Abu Saleh, W. K. & Edgerton, J. The Left Atrial Appendage: Target for Stroke
1463 Reduction in Atrial Fibrillation. *Methodist DeBakey Cardiovasc. J.* **11**, 100–103 (2015).

1464 **AUTHOR CONTRIBUTIONS**

1465 Conceptualization: M.H., P.F.-B.; Model development and validation analyses: F.R., M.H., P.F.-B.;
1466 Method development: F.R., M.J., M.Hi., M.R., P.F.-B; M.H., Implementation: F.R.; paper writing:
1467 F.R., M.H., P.F.-B.

1468

1469 **ACKNOWLEDGEMENTS**

1470 F.R. is supported by the Helmholtz Association under the joint research school "Munich School for
1471 Data Science - MUDS".

1472 **DATA AVAILABILITY**

1473 All datasets used in this study are already published and were obtained from public data
1474 repositories. Edges and Networks: BioPlex 3.0 edgelists were downloaded from
1475 <https://bioplex.hms.harvard.edu/interactions.php> . HuRI edgelist was downloaded from
1476 <http://www.interactome-atlas.org/download> . Intact edgelist was downloaded from
1477 <ftp://ftp.ebi.ac.uk/pub/databases/intact/current/psimitab/intact.txt> . GRNdb edgelists were
1478 downloaded from <http://grndb.com/download/> . Hetionet edgelist was downloaded from
1479 <https://github.com/hetio/hetionet/tree/master/hetnet/tsv> . Nodes and Features: Full list of human
1480 protein-coding genes was downloaded from [https://www.genenames.org/download/statistics-and-](https://www.genenames.org/download/statistics-and-files/)
1481 [files/](https://www.genenames.org/download/statistics-and-files/) , accessed 18.3.22. Positive labels were downloaded from
1482 https://github.com/bogdanlab/gene_sets/tree/master/mendelian_gene_sets , accessed 17.3.22.
1483 GWAS summary statistics were downloaded from <https://doi.org/10.5281/zenodo.3629742> .
1484 Tissue-specific median gene expression values were downloaded from
1485 [https://storage.googleapis.com/gtex_analysis_v7/rna_seq_data/GTEx_Analysis_2016-01-](https://storage.googleapis.com/gtex_analysis_v7/rna_seq_data/GTEx_Analysis_2016-01-15_v7_RNASeQCv1.1.8_gene_median_tpm.gct.gz)
1486 [15_v7_RNASeQCv1.1.8_gene_median_tpm.gct.gz](https://storage.googleapis.com/gtex_analysis_v7/rna_seq_data/GTEx_Analysis_2016-01-15_v7_RNASeQCv1.1.8_gene_median_tpm.gct.gz) . Median gene expression in blood cells was
1487 downloaded from https://v19.proteinatlas.org/download/rna_blood_cell.tsv.zip , accessed 17.3.22.
1488 External validation: Mouse knockout genes were downloaded from
1489 <http://www.informatics.jax.org/allele> , accessed 17.3.22. Lists of differentially expressed genes
1490 were downloaded from <https://gemma.msl.ubc.ca/phenotypes.html> , accessed 2.8.22. LoF and
1491 Missense Mutation intolerance Z-scores were downloaded from
1492 [ftp://ftp.broadinstitute.org/pub/ExAC_release/release1/manuscript_data/forweb_cleaned_exac_r03](ftp://ftp.broadinstitute.org/pub/ExAC_release/release1/manuscript_data/forweb_cleaned_exac_r03_march16_z_data_pLI.txt.gz)
1493 [_march16_z_data_pLI.txt.gz](ftp://ftp.broadinstitute.org/pub/ExAC_release/release1/manuscript_data/forweb_cleaned_exac_r03_march16_z_data_pLI.txt.gz) . List of drug targets was downloaded from [https://dgl-data.s3-us-](https://dgl-data.s3-us-west-2.amazonaws.com/dataset/DRKG/drkg.tar.gz)
1494 [west-2.amazonaws.com/dataset/DRKG/drkg.tar.gz](https://dgl-data.s3-us-west-2.amazonaws.com/dataset/DRKG/drkg.tar.gz) . Lists of druggable genes were downloaded
1495 from <https://www.dgidb.org/downloads> , accessed 24.3.22.

1496
1497 For reproducibility, the data can be jointly obtained via Speos' repository:
1498 <https://github.com/fratajcz/speos> or in its processed form from
1499 <https://doi.org/10.5281/zenodo.7468127> .

1500

1501 **CODE AVAILABILITY**

1502 Speos is open source, implemented in python and available at <https://github.com/fratajcz/speos> .
1503 Config files to reproduce the benchmarks and experiments are also available in that repository.

1504

1505 **COMPETING INTERESTS**

1506 The authors declare no competing interests.

1507

1508 **REPORTING SUMMARY**

1509 Further information on research design is available in the Nature Research Reporting Summary
1510 linked to this paper.

1511

1512 **ADDITIONAL INFORMATION**

1513

1514 **Supplementary Information** is available for this paper.