

RESEARCH

Benchmarking network propagation methods for disease gene identification

Sergio Picart-Armada^{1,2*}, Steven J. Barrett⁴, David R. Willé⁴, Alexandre Perera-Lluna^{1,2}, Alex Gutteridge³ and Benoit H. Dessailly⁵

Abstract

Background: In-silico identification of potential disease genes has become an essential aspect of drug target discovery. Recent studies suggest that one powerful way to identify successful targets is through the use of genetic and genomic information. Given a known disease gene, leveraging intermolecular connections via networks and pathways seems a natural way to identify other genes and proteins that are involved in similar biological processes, and that can therefore be analysed as additional targets.

Results: Here, we systematically tested the ability of 12 varied network-based algorithms to identify target genes and cross-validated these using gene-disease data from Open Targets on 22 common diseases. We considered two biological networks, six performance metrics and compared two types of input gene-disease association scores. We also compared several cross-validation schemes and showed that different choices had a remarkable impact on the performance estimates. When seeding biological networks with known drug targets, we found that machine learning and diffusion-based methods are able to find novel targets, showing around 2-4 true hits in the top 20 suggestions. Seeding the networks with genes associated to disease by genetics resulted in poorer performance, below 1 true hit on average. We also observed that the use of a larger network, although noisier, improved overall performance.

Conclusions: We conclude that machine learning and diffusion-based prioritisers are suited for drug discovery in practice and improve over simpler neighbour-voting methods. We also demonstrate the large effect of several factors on prediction performance, especially the validation strategy, input biological network, and definition of seed disease genes.

Keywords: guilt by association; biological networks; network propagation; drug targets; genetic associations; cross validation; protein complexes

*Correspondence:

sergi.picart@upc.edu

¹B2SLab, Departament d'Enginyeria de Sistemes,

Automàtica i Informàtica

Industrial, Universitat Politècnica de Catalunya, CIBER-BBN,

08028, Barcelona, Spain

Full list of author information is available at the end of the article

†Equal contributor

¹Background

²The pharmaceutical industry faces considerable challenges in the efficiency of com-²
³mercial drug research and development [1] and in particular in improving its ability³
⁴to identify future successful drug targets.⁴

⁵ It has been suggested that using genetic association information is one of the best⁵
⁶ways to identify such drug targets [2]. In recent years, a large number of highly⁶
⁷powered GWAS studies have been published for numerous common traits (see for⁷
⁸example [3] or [4]) and have yielded many candidate genes. Further potential targets⁸
⁹can be identified by adding contextual data to the genetic associations, such as genes⁹
¹⁰involved in similar biological processes [5, 6]. Biological networks and biological¹⁰
¹¹pathways can be used as a source of contextual data.¹¹

¹² Biological networks are widely used in bioinformatics and can be constructed from¹²
¹³multiple data sources, ranging from macromolecular interaction data collected from¹³
¹⁴the literature [7] to correlation of expression in transcriptomics or proteomics sam-¹⁴
¹⁵ples of interest [8]. A large number of interaction network resources have been made¹⁵
¹⁶available over the years, many of which are now in the public domain, combining¹⁶
¹⁷thousands of interactions in a single location [9, 10]. They are based on three dif-¹⁷
¹⁸ferent fundamental types of data: (1) data-driven networks such as those built by¹⁸
¹⁹WGCNA [8] for co-expression; (2) interactions extracted from the literature using¹⁹
²⁰a human curation process as exemplified by IntAct [11] or BioGRID [12]; and (3)²⁰
²¹interactions extracted from the literature using text mining approaches [13].²¹

²² On the other hand, a plethora of network analysis algorithms are available for²²
²³extracting useful information from such large biological networks in a variety of²³
²⁴contexts. Algorithms range in complexity from simple first-neighbour approaches,²⁴
²⁵where the direct neighbours of a gene of interest are assumed to be implicated in²⁵
²⁶similar processes [14], to machine learning (ML) algorithms designed to learn from²⁶
²⁷the features of the network to make more useful biological predictions (e.g. [15]).²⁷

²⁸ One broad family of network analysis algorithms are the so-called Network Prop-²⁸
²⁹agation approaches [16], used in contexts such as protein function prediction [17],²⁹
³⁰disease gene identification [16] and cancer gene mutation identification [18]. In this³⁰
³¹paper, we perform a systematic review of the usefulness of network analysis methods³¹
³²for the purpose of identification of disease genes susceptible of being drug targets.³²
³³Claims that such methods are helpful in that context have been made on numerous³³

¹occasions but a comprehensive validation study is lacking. One major challenge in¹
²doing such a study is that it is not straightforward to define a list of known disease²
³genes to be used for this purpose.³

⁴ To address this, the Open Targets collaboration has been setup between phar-⁴
⁵maceutical companies and public institutions to collect information on known drug⁵
⁶targets and to help identify new ones [19]. A dedicated internet platform provides⁶
⁷a free-to-use accessible resource summarising known data on gene-disease relation-⁷
⁸ships from a number of data sources (e.g. known released drugs, genetic associations⁸
⁹from GWAS, etc) [19].⁹

¹⁰ The purpose of this work is to quantify the performance of network-based meth-¹⁰
¹¹ods to prioritise novel targets, using various networks and validation schemes, and¹¹
¹²aiming at a faithful reflection of a realistic drug development scenario. We select¹²
¹³a number of network approaches that are representative of several classes of algo-¹³
¹⁴rithms, and test their ability to recover known disease genes by cross-validation.¹⁴
¹⁵

¹⁶ We benchmark multiple definitions of disease genes, computational methods, bi-¹⁶
¹⁷ological networks, validation schemes and performance metrics. We account for all¹⁷
¹⁸possible combinations of such factors and derive guidelines for future disease gene¹⁸
¹⁹identification studies.¹⁹

²⁰

²¹**Results**

²²**Benchmark framework**

²³ Our general approach, summarised in figure 1, consisted in using a biological net-²³
²⁴work and a list of genes with prior disease-association scores as input to a network²⁴
²⁵propagation approach. We used three cross-validation schemes -two take into ac-²⁵
²⁶count protein complexes- in which some of the prior disease-association scores are²⁶
²⁷hidden. The desired output was a new ranking of genes in terms of their association²⁷
²⁸scores to the disease. Such ranking was compared to the known disease genes in²⁸
²⁹the validation fold using several performance metrics. Given the amount of design²⁹
³⁰factors and comparisons, the metrics were analysed through explorative additive³⁰
³¹models (see Methods). Alternatively, we provide plots on the raw metrics in the³¹
³²supplement, stratified by method in figures S10 and S11 or by disease in figures S12³²
³³and S13.³³

¹ We considered 2 metrics (AUROC and top 20 hits) and 2 input types (known drug¹

²target genes and genetically associated genes), resulting in a total of 4 combinations,²

³each described through an additive main effect model. Another 4 metrics were³

⁴explored and can be found in the supplement (figure S17 and tables S6, S7).⁴

⁵ Interactions were explored, but they did not provide any added value for the⁵

⁶extra complexity (see figure S18 from the supplement). The metrics used were the⁶

⁷dependent variables, while the regressors included the prediction method, the CV⁷

⁸scheme, the network and the disease.⁸

⁹

¹⁰Performance using known drug targets as input¹⁰

¹¹Figure 2 describes the additive models for AUROC and top 20 hits, and using known¹¹

¹²drug targets as input. Figure 3 contains their predictions for each method, network¹²

¹³and cross validation scheme with 95% confidence intervals, averaged over diseases.¹³

¹⁴The models are complex and we therefore review each main effect separately.¹⁴

¹⁵ AUPRC (quasi-binomial model) and top 20 hits (quasi-poisson) behave alike, as¹⁵

¹⁶can be observed by their similar ranking of model estimates in Figure 2. For inter-¹⁶

¹⁷pretability within real scenarios, only top 20 hits is shown in the main body. The¹⁷

¹⁸standard AUROC (quasi-binomial) clearly led to different conclusions and is kept¹⁸

¹⁹throughout the results section for comparison. The remaining metrics (AUPRC,¹⁹

²⁰pAUROC 5%, pAUROC 10% and top 100 hits) result in similar method priorita-²⁰

²¹tions as top 20 hits, see figure S17. Detailed models can be found in the supplement,²¹

²²indexed by tables S6 and S7.²²

²³Comparing cross-validation schemes²³

²⁴ Whether protein complexes were properly taken into account when performing the²⁴

²⁵cross-validation (see Methods) stood out as a key influence on the quality of predic-²⁵

²⁶tions: there was a dramatic reduction in performance for most methods when using²⁶

²⁷a complex-aware cross-validation strategy. For instance, method rf applied on the²⁷

²⁸STRING network dropped from almost 12 correct hits in the top 20 predicted dis-²⁸

²⁹ease genes when using our *classic* cross-validation scheme down to fewer than 4.5²⁹

³⁰when using either of our complex-aware cross-validation schemes. Likewise, table S5³⁰

³¹from the supplement ratifies that only the *classic* cross validation splits complexes.³¹

³² Our data suggests that the performance drop when choosing the appropriate³²

³³validation strategy is comparable to the performance gap of competitive methods³³

¹versus a simple neighbour-voting baseline (see figure 2). This highlights the im-¹
²portance of carefully controlling for this bias when estimating the performance of²
³network-based disease gene prediction methods. Overall, the *classic* cross-validation³
⁴scheme gave biased estimates in our dataset, whereas our *block* and *representative*⁴
⁵cross-validation schemes had similar effects on the prediction performance. Method⁵
⁶ranking was independent of the cross validation choice thanks to the use of an addi-⁶
⁷tive model. And since both the *block* and *representative* schemes make theoretical⁷
⁸sense, we chose to focus on results from the block scheme in the rest of this study.⁸
9

¹⁰*Comparing networks*

¹¹We found that using STRING as opposite to OmniPath improved overall perfor-¹¹
¹²mance of network-based disease gene prediction methods. Our models for top 20₁₂
¹³hits quantified this effect as noticeable although less important than that of the₁₃
¹⁴cross validation strategy. For reference, method **rf** obtains about 3 true hits under₁₄
¹⁵both complex-aware strategies in OmniPath. It has been previously shown that the₁₅
¹⁶positive effect on predictive power of having more interactions and coverage in a₁₆
¹⁷network can outweigh the negative effect of increased number of false positive in-₁₇
¹⁸teractions [20], which is in line with our findings. The authors also report STRING₁₈
¹⁹among the best resources to discover disease genes, which is a finding we reproduce₁₉
20 here.
20

²¹We focus on the STRING results in the rest of the text.
21

²²*Comparing methods*

²³Having identified the optimal cross-validation scheme and network for our bench-²³
²⁴mark in the previous sections, we quantitatively compared the performance of the²⁴
²⁵different methods.
25

²⁶First, network topology alone had a slight predictive power, as method **pr** (PageR-²⁶
²⁷ank approach that ignores the input gene scores) showed better performance than²⁷
²⁸the **random** baseline under all the metrics. The randomised diffusion **randomraw**²⁸
²⁹lied between **random** and **pr** in performance. Both facts support the existence of an²⁹
³⁰inherent network topology-related bias among the positives that benefits diffusion-³⁰
31 based methods.
31

³²Second, the basic GBA approach from EGAD had an advantage over the input-³²
³³naïve baselines **pr**, **randomraw** and **random**. It also outperformed prioritizing genes³³

¹using other Open Targets data stream scores such as genes associated to disease¹
²from pathways or from the literature (see supplement, table S19).²

³ Most diffusion-based and ML-based methods outperformed EGAD. Results from³
⁴top 20 hits suggest using **rf** for the best performance followed by, in order: **raw** and⁴
⁵**bagsvm**, **z** and **svm** (main models panel in Figure 6).⁵

⁶ To formally test the differences between methods, we carried a Tukey's multi-⁶
⁷ple comparison test on the model coefficients (Figure 5) as implemented in the R⁷
⁸package multcomp [21]. Although such differences were in most cases statistically⁸
⁹significant, even with such a strong multiplicity adjustment, their actual effect size⁹
¹⁰or magnitude can be modest in practice, see Figures 3 and 6.¹⁰

¹¹ The ranking of methods was similar when using the metrics AUPRC, pAUROC¹¹
¹²and top k hits (see supplement, figure S17) and is only intended to be a general¹²
¹³reference, given the impact of the problem definition, cross validation scheme and¹³
¹⁴the network choice.¹⁴

¹⁵ With AUROC on the other hand, **rf** lost its edge whilst most diffusion-based¹⁵
¹⁶and ML-based methods appeared technically tied. Despite its theoretical basis, in-¹⁶
¹⁷terpretability and widespread use in similar benchmarks, these results support the¹⁷
¹⁸assertion that AUROC is a sub-optimal choice in drug discovery practical scenarios.¹⁸

¹⁹ Figure 4 further shows how the different methods compare with one another. Dis-¹⁹
²⁰tances between each pair of method in terms of their top 100 novel predictions were²⁰
²¹represented graphically. From this we observe that the supervised bagged SVM ap-²¹
²²proach (**bagsvm**) behaves similarly to the simple diffusion approach (**raw**), reflecting²²
²³the fact that they use the same kernel. We also observe that diffusion approaches do²³
²⁴not necessarily produce similar results (compare for example **raw** and **z**). And that²⁴
²⁵interestingly, methods EGAD (arguably one of the simplest) and COSNet (arguably²⁵
²⁶one of the most complex) seemed to result in similar predictions. Fully supervised²⁶
²⁷and semi-supervised approaches largely group in the top right hand quadrant of the²⁷
²⁸STRING plot away from diffusion methods, possibly showing some shared greater²⁸
²⁹potential for “learning effect” with the larger network.²⁹

³⁰ Interestingly, when comparing overall performances shown in figure 6 with the³⁰
³¹prediction differences from the MDS plot (figure 4), it appears that the better³¹
³²performing methods may be doing well for different reasons as they do not occur³²
³³within the same region of the plot (e.g. **rf** and **raw**). MDS plots on the eight possible³³

¹combinations of network, input type and inclusion of seed genes are displayed in¹
²the supplementary figures [S15](#) and [S16](#).

³ Regarding the STRING network and the block validation scheme, we fitted six³
⁴additive models (one per metric) to the known drug target data (see supplement,⁴
⁵table [S7](#)) and prioritised the methods (reduced models in figure [6](#)). These reduced⁵
⁶models better described this particular scenario, as they were not forced to fit the⁶
⁷trends in all networks and validation schemes in an additive way. Considering the⁷
⁸top 20 hits, **rf** and **svm** were the optimal choices, followed by **wsld** and **knn**.

⁹
¹⁰*Comparing diseases*

¹¹We next examine performance by disease. The top 20 hits model in figure [2](#) shows₁₁
₁₂that allergy (the figure's baseline reference), ulcerative colitis and rheumatoid arthri-₁₂
₁₃tis (group I) are the diseases for which prediction of disease genes was worst, whereas₁₃
₁₄cardiac arrhythmia, Parkinson's disease, stroke and multiple sclerosis (group II) are₁₄
₁₅those for which it was best. As shown in figure [7](#), group I diseases had fewer known₁₅
₁₆disease genes and lower modularity compared to group II diseases.

¹⁷Prediction methods worked better when more known disease genes were available₁₇
₁₈as input in the network, with two possible underlying reasons being the greater₁₈
₁₉data availability to train the methods, and the natural bias of top 20 hits towards₁₉
₂₀datasets with more positives. Likewise, a stronger modularity within disease genes₂₀
₂₁justifies the guilt-by-association principle and led to better performances. In turn,₂₁
₂₂the number of genes and the modularity were positively correlated, see supplement,₂₂
₂₃figure [S14](#).

²⁴*Performance using genetic associations as input*
²⁵Using genetically associated genes as input to a prediction approach to find known²⁵
²⁶drug targets mimicked a realistic scenario where novel genetic associations are²⁶
²⁷screened as potential targets. However, inferring known drug targets through the²⁷
²⁸indirect genetic evidence posed problems to prediction strategies, especially those²⁸
²⁹based on machine learning. Learning is done using one class of genes in order to pre-²⁹
³⁰dict genes that belong to another class, and the learning space suffers from intrinsic³⁰
³¹uncertainties in the genetic associations to disease.

³²Consequently, we observed a major performance drop on all the prioritisation³²
³³methods: using any network and cross-validation scheme, the predicted top 20 hits³³

¹were practically bounded by 1. This was more pronounced on supervised machine¹
²learning-focused strategies, as **rf** and **svm** lost their edge on diffusion-based strate²
³gies. The fact that the genetic associations of the validation fold were hidden further³
⁴hindered the predictions and can be a cause of our pessimistic performance esti⁴
⁵mates.

6

5

⁷*Comparing cross-validation schemes*

7

⁸For reference, we also ran all three cross validation schemes on the genetic data to⁸
⁹quantify and account for complex-related bias. The models confirm that, contrary⁹
¹⁰to the drugs-related input, the differences between the results for the different cross-¹⁰
¹¹validation schemes were rather modest. For example, method **raw** with the STRING¹¹
¹²network attains 0.59-0.64, 0.50-0.54 and 0.37-0.40 hits in the top 20 under the¹²
¹³classical, block and representative cross-validation strategies. The slightly larger¹³
¹⁴negative effect on top 20 hits observed with the representative scheme is expected¹⁴
¹⁵because the number of positives that act as validation decreased and this metric¹⁵
¹⁶is biased by the class imbalance. The agreement between method ranking using¹⁶
¹⁷AUPRC and top 20 hits was less consistent, possibly due to the performance drop,¹⁷
¹⁸whilst AUROC again yielded quite a different ranking. Further data can be found¹⁸
¹⁹in the supplement, tables **S15** and **S16**.

19

20

20

²¹*Comparing networks*

21

²²The change in performance for using the OmniPath network instead of the filtered²²
²³STRING network was also limited. For AUROC the effect was negative, whereas for²³
²⁴the top 20 hits metric the performance improved. Method **raw** changed from 0.50-²⁴
²⁵0.54 top 20 hits in STRING to 0.61-0.66 in OmniPath under the block validation²⁵
²⁶strategy.

26

²⁷*Comparing methods*

27

²⁸To be consistent with the drugs section, we take as reference the block cross-²⁸
²⁹validation strategy and the STRING network.

29

³⁰The baseline approach **pr** that effectively makes use of the network topology alone³⁰
³¹proved difficult to improve upon, with 0.43-0.47 expected true hits in the top 20.³¹
³²Methods **raw** and **rf** respectively achieved 0.50-0.54 and 0.23-0.26 – although sig-³²
³³nificant, the difference in practice would be minimal. The best performing method³³

¹was **mc** with 0.65-0.7 hits. All the performance estimates can be found in the sup-¹
²plement, table **S16**. To give an idea of the effort that would be required in a realistic²
³setting to find novel disease genes, the number of correct hits in the top 100 hits³
⁴was 3.29-3.45 with the best performing method (in this case, **ppr**), against 2.25-2.38⁴
⁵of **pr**.

⁶ Two main conclusions can be drawn from these results. First, the network topology⁶
⁷baseline retained some predictive power upon which most diffusion-based methods,⁷
⁸as well as machine-learning approaches **COSNet** and **bagsvm**, only managed to add⁸
⁹minor improvements, if any. Second, drug targets could still be found by combining⁹
¹⁰network analysis and genes with genetic associations to disease, but with a substan-¹⁰
¹¹tially lower performance and with a marginal gain compared to a baseline approach¹¹
¹²that would only use the network topology to find targets (e.g. by screening the most¹²
¹³connected genes in the network).

¹⁴ It is worth noting that gene-disease genetic association scores themselves have¹⁴
¹⁵drawbacks and that better prediction accuracy could result as genetic association¹⁵
¹⁶data improves.

¹⁷

¹⁸**Conclusions**

¹⁹We performed an extensive analysis of the ability of network-based approaches to¹⁹
²⁰identify novel disease genes. We exhaustively explored the effect of different factors²⁰
²¹including the biological network, the definition of disease genes, and the statistical²¹
²²framework being used to evaluate methods performance. We show that carefully²²
²³choosing an appropriate cross-validation framework and suitable performance met-²³
²⁴ric have an important effect in evaluating the utility of these methods.

²⁵ Our main conclusion is that network-based drug target discovery seems effective,²⁵
²⁶reflecting the fact that drug targets tend to cluster within the network. This in²⁶
²⁷turn may of course be due to the fact that the scientific community has so far²⁷
²⁸been focusing on testing the same proven mechanisms. In a strict cross-validation²⁸
²⁹setting, we found that even the most basic guilt-by-association method was useful,²⁹
³⁰with ~2 correct hits in its top 20 predictions, compared to ~0.1 when using a³⁰
³¹random ranking. The best diffusion based algorithm improved that figure to ~3.75,³¹
³²and the best overall performing method was a random forest classifier on network-³²
³³based features (~4.4 hits). Leading approaches can be notably different in terms³³

¹of their top predictions, suggesting potential complementarity. We found a better¹
²performance when using a network with more coverage at the expense of more false²
³positive interactions. In a more conservative network, random forest performance³
⁴dropped to ~3.1 hits. Comparing performance on different diseases shows that the⁴
⁵more known target genes, and the more clustered these are in the network, the⁵
⁶better the performance of network-based approaches for finding novel targets for it.⁶

⁷ We also explored the prediction of known drug target genes by seeding the network⁷
⁸with an indirect data stream, in particular, genetic association data. Here, the best⁸
⁹performing methods were diffusion-based and presented a statistically significant,⁹
¹⁰but marginal, improvement over approaches that only look at network connectivity.¹⁰

¹¹ We conclude that network propagation methods can help identify novel disease¹¹
¹²genes, but that the choice of the input network and the seed scores on the genes¹²
¹³needs careful consideration. Based on our approach and endorsed benchmarks, we¹³
¹⁴recommend the use of methods employing representations of diffusion-based in-¹⁴
¹⁵formation (the MashUp network-based features and the diffusion kernels), namely¹⁵
¹⁶random forest, the support vector machine variants, and raw diffusion algorithms¹⁶
¹⁷for optimal results.¹⁷

¹⁸

¹⁸

¹⁹**Methods**

¹⁹

²⁰**Selection of methods for investigation**

²⁰

²¹Algorithms were selected for validation based on the following criteria:²¹

²² 1 published in a peer-reviewed journal, with evidence of improved performance²²
²³ in gene disease prediction relative to contenders,²³
²⁴ 2 implemented as a well documented open source package, that is efficient, ro-²⁴
²⁵ bust and usable within a batch testing framework,²⁵
²⁶ 3 directly applicable for gene disease identification from a single gene/protein²⁶
²⁷ interaction network, without requiring fundamental changes to the approach²⁷
²⁸ or additional annotation information and²⁸
²⁹ 4 capable of outputting a ranked list of individual genes (as opposed to gene²⁹
³⁰ modules for example).³⁰

³¹ In addition, we selected methods that were representative of a diverse panel of³¹
³² algorithms, including diffusion-based approaches, supervised learning approaches,³²
³³ and a number of baseline approaches (see table 1).³³

¹Testing framework, algorithms and parameterisation

¹

²All tests and batch runs were set-up and conducted using the R statistical pro-

²

³gramming language [22]. When no R package was available, the methodology was

³

⁴re-implemented, building upon existing R packages whenever possible. Standard

⁴

⁵R machine learning libraries were used to train the support vector machine and

⁵

⁶random forest classifiers. Only the MashUp algorithm [23] required feature gener-

⁶

⁷ation outside of the R environment, using the Matlab code from their publication.

⁷

⁸Further details on the methods implementation can be found in the supplement,

⁸

⁹section “Method details”.

⁹

¹⁰ EGAD [24], a pure neighbour-voting approach, was used here as a baseline com-

¹⁰

¹¹parator.

¹¹

¹² Diffusion (propagation) methods are central in this study. We used the random

¹²

¹³walk-based personalised PageRank [25], previously used in similar tasks [26], as

¹³

¹⁴implemented in igraph [27]. The remaining diffusion-based methods were run on

¹⁴

¹⁵top of the regularised Laplacian kernel [28], computed through diffuStats [29]. We

¹⁵

¹⁶included the classical diffusion **raw**, a weighted approach version **gm** and two statis-

¹⁶

¹⁷tically normalised scores (**mc** and **z**), as implemented in diffuStats. In the scope of

¹⁷

¹⁸positive-unlabelled learning [30, 31], we included the kernelised scores **knn** and the

¹⁸

¹⁹linear decayed **wsld** from RANKS [32]. Closing this category, we implemented the

¹⁹

²⁰bagging Support Vector Machine approach from ProDiGe1 [33], here **bagsvm**.

²⁰

²¹ Purer ML-based methods were also included. On one hand, network-based features

²¹

²²were generated using MashUp [23] and two classical classifiers were fitted to them,

²²

²³based on caret [34] and mlr [35]. These are **svm**, the Support Vector Machine as

²³

²⁴implemented in kernlab [36], and **rf**, the Random Forest found in the randomForest

²⁴

²⁵package [37]. On the other hand, we tried the parametric Hopfield recurrent neural

²⁵

²⁶network classifier in the COSNet R package [38, 39].

²⁶

²⁷ Finally, we defined three naive baseline methods: (1) **pr**, a classic problem-naïve

²⁷

²⁸‘non-personalised’ PageRank implementation where input scores on the genes are

²⁸

²⁹ignored; (2) **randomraw**, which applies the **raw** diffusion approach to randomly per-

²⁹

³⁰muted input scores on the genes; and (3) **random**, a uniform re-ranking of input

³⁰

³¹genes without any network propagation. The inclusion of **pr** and **randomraw** al-

³¹

³²lowed us to quantify the predictive power of the network topology alone, without

³²

³³any consideration for the input scores on the genes.

³³

¹Biological networks

²The biological network used in the validation is of critical importance as current ²
³network resources contain both false positive and false negative interactions, and ³
⁴these will affect any subsequent predictions [20]. ⁴

⁵ Here, we used two human networks with different general properties, one more ⁵
⁶likely to contain false positive interactions (STRING [40]), and another more con- ⁶
⁷servative (OmniPath [41]), to test the effect of the network itself on network prop- ⁷
⁸agation performance. We further filtered STRING [40] to retain only a subset of ⁸
⁹interactions. Having tested several filters, we settled upon high-confidence inter- ⁹
¹⁰actions (combined score > 700) with some evidence from the “Experiments” or ¹⁰
¹¹“Databases” data sources (see supplement, table S2). Applying these filters and ¹¹
¹²taking the largest connected component resulted in a connected network of 11,748 ¹²
¹³nodes and 236,963 edges. Edges were assigned weights between 0 and 1 by rescaling ¹³
¹⁴the STRING combined score. ¹⁴

¹⁵ We did not filter the OmniPath network [41]. After removing duplicated edges ¹⁵
¹⁶and taking the largest connected component, the OmniPath network contained ¹⁶
¹⁷8,580 nodes and 42,145 unweighted edges. ¹⁷

¹⁸

¹⁹Disease gene data

²⁰We used the Open Targets platform [19] to select known disease-related genes. In ²⁰
²¹this analysis we define disease-related genes are those reported in Open Targets as ²¹
²²being the target of a known drug against the disease of interest, or as those with ²²
²³a genetic association of sufficient confidence with the disease. Associations were bi- ²³
²⁴narised: any non-zero drugs-related association was considered positive, implying ²⁴
²⁵that the methods would predict genes on which a drug has been essayed, regardless ²⁵
²⁶of whether the drug was eventually approved. Likewise, only genetic associations ²⁶
²⁷with an Open Targets score above 0.16 (see supplement, figure S1) were consid- ²⁷
²⁸ered positive. We considered exclusively common diseases with at least 1,000 Open ²⁸
²⁹Targets associations, of which a minimum of 50 could be based on known released ²⁹
³⁰drugs and 50 on genetic associations, in order to avoid empty folds in the nested ³⁰
³¹cross-validations. By applying these filters, we generated a list of phenotypes and ³¹
³²diseases which we then manually curated to remove cancers (where causal genetic ³²
³³mechanisms can differ from those of other common diseases), non-disease pheno- ³³

¹type terms (e.g. “body weight and measures”) as well as vague or broad terms¹
²(e.g. “cerebrovascular disorder” or “head disease”) and infectious diseases. This left²
³22 diseases considered in this study (table 2). Further descriptive material on the³
⁴role of disease genes within the STRING network can be found in the section “De-⁴
⁵scriptive disease statistics in the STRING network” from the supplement.

⁶ ⁶

⁷Validation strategies ⁷

⁸*Input Gene Scores* ⁸

⁹We used the binarised drug association scores and genetic association scores from⁹
¹⁰Open Targets as input gene-level scores to seed the network propagation analyses¹⁰
¹¹(figure 8) and test their ability to recover known drug targets. With the first ap-¹¹
¹²proach (subfigure (1) in figure 8), we tested the predictive power of current network¹²
¹³propagation methods for drug target identification using a direct source of evidence¹³
¹⁴(known drug targets). In the second approach (subfigure (2) in figure 8), we assessed¹⁴
¹⁵the ability of a reasonable but indirect source of evidence – genetic associations to¹⁵
¹⁶disease – in combination with network propagation to recover known drug targets.¹⁶

¹⁷ ¹⁷

¹⁸*Metrics* ¹⁸

¹⁹Methods were systematically compared using standard performance metrics. The¹⁹
²⁰Area under the Receiver Operating Characteristic curve (AUROC) is extensively²⁰
²¹used in the literature for binary classification of disease genes [42], but can be²¹
²²misleading in this context given the extent of the class imbalance between target and²²
²³non-target genes [43]. We however included it in our benchmark for comparison with²³
²⁴previous literature. More suitable measures of success in this case are Area under²⁴
²⁵the Precision-Recall curve (AUPRC) [43] and partial AUROC (pAUROC) [44].²⁵
²⁶AUROC, AUPRC and pAUROC were computed with the precrec R package [45].²⁶
²⁷We also included top 20 hits, defined as the number of true positives in the top²⁷
²⁸20 predicted genes (proportional to precision at 20). It is straightforward, intuitive²⁸
²⁹and most likely to be useful in practice, such as a screening experiment where only²⁹
³⁰a small number of predicted hits can be assayed.

³¹We considered another 3 metrics, reported only in Supplement, i.e. partial AU-³¹
³²ROC up to 5% FPR, partial AUROC up to 10% FPR, and number of hits within³²
³³the top 100 genes.

¹ *Cross validation schemes and protein complexes*

² Standard (stratified) and modified k-fold cross-validation were used to estimate²
³ the performance of network-based methods. Folds were based upon known drugs³
⁴ related genes, regardless of which type of input was used (see figure 8). Genes in⁴
⁵ the training fold were negatively or positively labelled according to the input type,⁵
⁶ whereas genes in the validation fold were left unlabelled.⁶

⁷ A fundamental challenge we faced when applying cross-validation to this prob-⁷
⁸ lem was that known drug targets often consist in protein complexes, e.g. multi-⁸
⁹ protein receptors. Drug-target associations typically have complex-level resolution.⁹
¹⁰ The drug target data from Open Targets comes from ChEmbl [46], in which all the¹⁰
¹¹ proteins in the targeted complex are labelled as targets.¹¹

¹² If left uncorrected, this could bias our cross-validation results: networks densely¹²
¹³ connect proteins within a complex, random folds would frequently split positively¹³
¹⁴ labelled complexes between train and validation, and therefore network-based meth-¹⁴
¹⁵ ods would have an unfair advantage at spotting positive hits in the training folds.¹⁵
¹⁶ In view of this, we benchmarked the methods under three cross validation strate-¹⁶
¹⁷ gies: a standard cross validation (A) in line with usual practice and two (B, C)¹⁷
¹⁸ complex-aware schemes (figure 9) addressing non-independence between folds when¹⁸
¹⁹ the known drug targets act as input.¹⁹

²⁰ Strategy (A), called **classic**, was a regular stratified *k*-fold repeated cross-²⁰
²¹ validation. We used *k* = 3 folds, averaging metrics over each set of folds, repeated²¹
²² 25 times (see also figure 1).²²

²³ Strategy (B), named **block**, performed a repeated cross validation while explicitly²³
²⁴ preventing any complexes that contain disease genes to be split across folds. The key²⁴
²⁵ point is that, where involved, shuffling was performed at the complex level instead²⁵
²⁶ of the gene level – overlapping complexes that shared at least one known drug target²⁶
²⁷ were merged into a larger pseudo-complex before shuffling. Fold boundaries were²⁷
²⁸ chosen so that no complex was divided into two folds, while keeping them as close²⁸
²⁹ as possible to those that would give a balanced partition, see figure 9. Nevertheless,²⁹
³⁰ a limitation of this scheme is that it can fail to balance fold sizes in the presence³⁰
³¹ of large complexes (see supplement, figure S9). For example, chronic obstructive³¹
³² pulmonary disease exhibited imbalanced folds, as 50 of the proteins involved belong³²
³³ to the Mitochondrial Complex I³³

¹ Strategy (C), referred to as **representative**, selected only a single representative¹
² or prototype gene for each complex to ensure that gene information in a complex²
³ was not mixed between training and validation folds. In each repetition of cross val-³
⁴ idation, after merging the overlapping complexes, a single gene from each complex⁴
⁵ was chosen uniformly at random and kept as positive. The remaining genes from the⁵
⁶ complexes involved in the disease were set aside from the training and validation⁶
⁷ sets, in order (1) not to mislead methods into assuming their labels were negative⁷
⁸ in the training phase, and (2) not to overestimate (if set as positives) or penalise⁸
⁹ (if set as negatives) methods that ranked them highly, as they were expected to do⁹
¹⁰ so. This strategy kept the folds balanced, but at the expense of a possible loss of¹⁰
¹¹ information by summarising each complex by a single gene at a time, reducing the¹¹
¹² number of positives for training and validation.¹²

¹³

¹³

¹⁴

¹⁴

Additive performance models

¹⁵

¹⁵

¹⁶ For a systematic comparison between diseases, methods, cross-validation schemes¹⁶
¹⁷ and input types, we fitted an additive regression model to the performance metrics of¹⁷
¹⁸ each (averaged) fold from the cross-validation. The use of main effect models eased¹⁸
¹⁹ the evaluation of each individual factor while correcting for the other covariates.¹⁹

²⁰ We modelled each metric f separately for each input type, not to mix problems of²⁰
²¹ different nature:²¹

²²

²²

²³

²³

$$f \sim \text{cv_scheme} + \text{network} + \text{method} + \text{disease}$$

²⁴

²⁴

²⁵

²⁵

We fitted dispersion-adjusted logistic-like *quasibinomial* distributions for the met-²⁶
rics AUROC, pAUROC and AUPRC and *quasipoisson* for top k hits. *The effect of*²⁶
*27 changing any of the four main effects is discussed in separate sub-sections in Re-*²⁷
28 sults, following the order from the formula above. After a data driven choice of²⁸
*29 cross-validation scheme and network we fitted reduced models within them for a*²⁹
*30 more accurate description:*³⁰

³¹

³¹

³²

³²

³³

³³

$$f \sim \text{method} + \text{disease}$$

¹*Qualitative methods comparison*

1

²The rankings produced by the different algorithms were qualitatively compared us-²
³ing Spearman's footrule [47]. Distances were computed between all method ranking³
⁴pairs for each individual combination of disease, input type, network and for the⁴
⁵top N predicted genes, excluding the original seed genes. This part does not involve⁵
⁶cross validation – all known disease-associated genes were used for gene prioriti-⁶
⁷sations. Pairs of rankings could include genes uniquely ranked highly by a single⁷
⁸algorithm from the comparison, so mismatch counts (i.e. percentage mismatches)⁸
⁹between these rankings were also taken into account. Mismatches occur when a⁹
¹⁰gene features in the top N predictions of one algorithm and is missing from the¹⁰
¹¹corresponding ranking by another algorithm. A compact visualisation of distance¹¹
¹²matrices was obtained using a multi-view extension of MDS [48, 49, 50]. For this¹²
¹³we used the R package *multiview* [51] that generates a single, low-dimensional pro-¹³
¹⁴jection of combined inputs (disease, input and network).¹⁴

15

15

16Competing interests

16

SB, DW, AG and BD are paid employees and shareholders of GlaxoSmithKline PLC. The commercial affiliation of
17SB, DW, AG and BD does not alter our adherence to BioMed Central policies.¹⁷

18Availability of data and materials

18

The datasets supporting the conclusions of this article are available in <https://github.com/b2sLab/genedise>.¹⁹

19

20Author's contributions

20

SP, SB, DW, and BD analysed and interpreted the data. AP, AG and BD helped supervise the project. All authors
21provided critical feedback and helped shape the research, analysis and manuscript. All authors approved the final
version of this manuscript for publication.²¹

22

22Acknowledgements

22

AG would like to acknowledge Philippe Sanseau, Matt Nelson and John Whittaker for critical feedback on the
23applicability of this research to drug discovery.²⁴

24

This work was supported by the Spanish Ministry of Economy and Competitiveness (MINECO) [TEC2014-60337-R
25and DPI2017-89827-R to AP].²⁵

AP and SP thank for funding the Spanish Networking Biomedical Research Centre in the subject area of
26Bioengineering, Biomaterials and Nanomedicine (CIBER-BBN), initiative of Instituto de Investigación Carlos III
27(ISCIII). SP thanks the AGAUR FI-scholarship programme.²⁷

28Author details

28

¹B2SLab, Departament d'Enginyeria de Sistemes, Automàtica i Informàtica Industrial, Universitat Politècnica de
29Catalunya, CIBER-BBN, 08028, Barcelona, Spain. ²Institut de Recerca Pediàtrica Hospital Sant Joan de Déu,
3008950, Esplugues de Llobregat, Barcelona, Spain. ³Computational Biology and Statistics, GSK, Gunnels Wood

30

Road, SG1 2NY Stevenage, United Kingdom. ⁴Research Statistics, GSK, Gunnels Wood Road, SG1 2NY

31

³¹Stevenage, United Kingdom. ⁵GSK Vaccines, Rue de l'Institut 89, 1330 Rixensart, Belgium.³¹

32References

32

1. Scannell, J.W., Blanckley, A., Boldon, H., Warrington, B.: Diagnosing the decline in pharmaceutical r&d
33 efficiency. *Nature reviews Drug discovery* 11(3), 191–200 (2012)

33

1 2. Nelson, M.R., Tipney, H., Painter, J.L., Shen, J., Nicoletti, P., Shen, Y., Floratos, A., Sham, P.C., Li, M.J., 1
2 Wang, J., Cardon, L.R., Whittaker, J.C., Sanseau, P.: The support of human genetic evidence for approved 2
drug indications. *Nature Genetics* **47**(8), 856–860 (2015). doi:[10.1038/ng.3314](https://doi.org/10.1038/ng.3314)

3 3. Verstocket, B., Smith, K.G., Lee, J.C.: Genome-wide association studies in crohn's disease: Past, present and 3
future. *Clinical & Translational Immunology* **7**(1), 1001 (2018). doi:[10.1002/cti2.1001](https://doi.org/10.1002/cti2.1001)

4 4. of the Psychiatric Genomics Consortium, S.W.G.: Biological insights from 108 schizophrenia-associated genetic 4
loci. *Nature* **511**(7510), 421–427 (2014). doi:[10.1038/nature13595](https://doi.org/10.1038/nature13595)

5 5. Jia, P., Zhao, Z.: Network-assisted analysis to prioritize GWAS results: principles, methods and perspectives. 5
Human Genetics **133**(2), 125–138 (2013). doi:[10.1007/s00439-013-1377-1](https://doi.org/10.1007/s00439-013-1377-1)

6 6. Boyle, E.A., Li, Y.I., Pritchard, J.K.: An expanded view of complex traits: From polygenic to omnigenic. *Cell* 6
169(7), 1177–1186 (2017). doi:[10.1016/j.cell.2017.05.038](https://doi.org/10.1016/j.cell.2017.05.038)

7 7. Orchard, S., Kerrien, S., Abbani, S., Aranda, B., Bhat, J., Bidwell, S., Bridge, A., Brigandt, L., Brinkman, 7
F.S., Brinkman, F., Cesareni, G., Chatr-aryamontri, A., Chautard, E., Chen, C., Dumousseau, M., Goll, J., 8
Hancock, R.E., Hancock, R., Hannick, L.I., Jurisica, I., Khadake, J., Lynn, D.J., Mahadevan, U., Perfetto, L., 9
Raghunath, A., Ricard-Blum, S., Roechert, B., Salwinski, L., Stumpflen, V., Tyers, M., Uetz, P., Xenarios, I., 10
Hermjakob, H.: Protein interaction data curation: the International Molecular Exchange (IMEx) consortium. 11
Nat. Methods **9**(4), 345–350 (2012)

8 8. Langfelder, P., Horvath, S.: WGCNA: an r package for weighted correlation network analysis. *BMC 12
Bioinformatics* **9**(1), 559 (2008). doi:[10.1186/1471-2105-9-559](https://doi.org/10.1186/1471-2105-9-559)

9 9. Razick, S., Magklaras, G., Donaldson, I.M.: iRefIndex: A consolidated protein interaction database with 13
provenance. *BMC Bioinformatics* **9**(1), 405 (2008). doi:[10.1186/1471-2105-9-405](https://doi.org/10.1186/1471-2105-9-405)

10 10. Szklarczyk, D., Morris, J.H., Cook, H., Kuhn, M., Wyder, S., Simonovic, M., Santos, A., Doncheva, N.T., 14
Roth, A., Bork, P., Jensen, L.J., von Mering, C.: The STRING database in 2017: quality-controlled 15
protein–protein association networks, made broadly accessible. *Nucleic Acids Research* **45**(D1), 362–368 (2016). 16
doi:[10.1093/nar/gkw937](https://doi.org/10.1093/nar/gkw937)

11 11. Kerrien, S., Aranda, B., Breuza, L., Bridge, A., Broackes-Carter, F., Chen, C., Duesbury, M., Dumousseau, M., 17
Feuermann, M., Hinz, U., *et al.*: The intact molecular interaction database in 2012. *Nucleic acids research* 18
40(D1), 841–846 (2011)

12 12. Chatr-Aryamontri, A., Oughtred, R., Boucher, L., Rust, J., Chang, C., Kolas, N.K., O'Donnell, L., Oster, S., 19
Theesfeld, C., Sellam, A., *et al.*: The biogrid interaction database: 2017 update. *Nucleic acids research* **45**(D1), 20
369–379 (2017)

13 13. Al-Aamri, A., Taha, K., Al-Hammadi, Y., Maalouf, M., Homouz, D.: Constructing genetic networks using 21
biomedical literature and rare event classification. *Scientific reports* **7**(1), 15784 (2017)

14 14. Piovesan, D., Giollo, M., Ferrari, C., Tosatto, S.C.E.: Protein function prediction using guilty by association 22
from interaction networks. *Amino Acids* **47**(12), 2583–2592 (2015). doi:[10.1007/s00726-015-2049-3](https://doi.org/10.1007/s00726-015-2049-3)

15 15. Re, M., Mesiti, M., Valentini, G.: A fast ranking algorithm for predicting gene functions in biomolecular 23
networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* **9**(6), 1812–1818 (2012)

16 16. Cowen, L., Ideker, T., Raphael, B.J., Sharan, R.: Network propagation: a universal amplifier of genetic 24
associations. *Nature Reviews Genetics* **18**(9), 551–562 (2017). doi:[10.1038/nrg.2017.38](https://doi.org/10.1038/nrg.2017.38)

17 17. Sharan, R., Ulitsky, I., Shamir, R.: Network-based prediction of protein function. *Molecular Systems Biology* **3** 26
(2007). doi:[10.1038/msb4100129](https://doi.org/10.1038/msb4100129)

18 18. Leiserson, M.D.M., Vandin, F., Wu, H.-T., Dobson, J.R., Eldridge, J.V., Thomas, J.L., Papoutsaki, A., Kim, 27
Y., Niu, B., McLellan, M., Lawrence, M.S., Gonzalez-Perez, A., Tamborero, D., Cheng, Y., Ryslik, G.A., 28
Lopez-Bigas, N., Getz, G., Ding, L., Raphael, B.J.: Pan-cancer network analysis identifies combinations of rare 29
somatic mutations across pathways and protein complexes. *Nature Genetics* **47**(2), 106–114 (2014). 29
doi:[10.1038/ng.3168](https://doi.org/10.1038/ng.3168)

30 19. Koscielny, G., An, P., Carvalho-Silva, D., Cham, J.A., Fumis, L., Gasparyan, R., Hasan, S., Karamanis, N., 30
Maguire, M., Papa, E., *et al.*: Open targets: a platform for therapeutic target identification and validation. 31
Nucleic acids research **45**(D1), 985–994 (2016)

32 20. Huang, J.K., Carlin, D.E., Yu, M.K., Zhang, W., Kreisberg, J.F., Tamayo, P., Ideker, T.: Systematic evaluation 32
of molecular networks for discovery of disease genes. *Cell systems* **6**(4), 484–495 (2018) 33

¹21. Hothorn, T., Bretz, F., Hothorn, M.T.: The multcomp package (2009) 1

²22. R Core Team: R: A language and environment for statistical computing (2016). R Foundation for Statistical 2 Computing

³23. Cho, H., Berger, B., Peng, J.: Compact integration of multi-network topology for functional analysis of genes. 3 Cell systems 3(6), 540–548 (2016)

⁴24. Ballouz, S., Weber, M., Pavlidis, P., Gillis, J.: Egad: ultra-fast functional analysis of gene networks. 4 Bioinformatics 33(4), 612–614 (2017). doi:[10.18129/B9.bioc.EGAD](https://doi.org/10.18129/B9.bioc.EGAD)

⁵25. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. 5 Technical report, Stanford InfoLab (1999)

⁶26. Jiang, B., Kloster, K., Gleich, D.F., Gribkov, M.: Aptrank: an adaptive pagerank model for protein function 6 prediction on bi-relational graphs. Bioinformatics 33(12), 1829–1836 (2017)

⁷27. Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal Complex 7 Systems, 1695 (2006)

⁸28. Smola, A.J., Kondor, R.: Kernels and regularization on graphs. In: Learning Theory and Kernel Machines, pp. 8 144–158. Springer, ??? (2003)

⁹29. Picart-Armada, S., Thompson, W.K., Buil, A., Perera-Lluna, A.: An r package to compute diffusion-based 9 scores on biological networks: diffustats. Bioinformatics 32(18), 2872–2874 (2016).

¹⁰30. doi:[10.18129/B9.bioc.diffuStats](https://doi.org/10.18129/B9.bioc.diffuStats)

¹¹31. Elkan, C., Noto, K.: Learning classifiers from only positive and unlabeled data, 213–220 (2008). ACM 11

¹²32. Yang, P., Li, X.-L., Mei, J.-P., Kwoh, C.-K., Ng, S.-K.: Positive-unlabeled learning for disease gene 12 identification. Bioinformatics 28(20), 2640–2647 (2012)

¹³33. Valentini, G., Paccanaro, A., Caniza, H., Romero, A.E., Re, M.: Ranks: a flexible tool for node label ranking 13 and classification in biological networks. Artificial Intelligence in Medicine 61(2), 63–78 (2014)

¹⁴34. Mordelet, F., Vert, J.-P.: ProDiGe: Prioritization Of Disease Genes with multitask machine learning from 14 positive and unlabeled examples. BMC bioinformatics 12(1), 389 (2011)

¹⁵35. Kunn, M.: Building predictive models in r using the caret package. Journal of Statistical Software 28(5), 1–26 15 (2008)

¹⁶36. Bischl, B., Lang, M., Kotthoff, L., Schiffner, J., Richter, J., Stederus, E., Casalicchio, G., Jones, Z.M.: mlr: 16 Machine learning in r. Journal of Machine Learning Research 17(170), 1–5 (2016)

¹⁷37. Karatzoglou, A., Smola, A., Hornik, K., Zeileis, A.: kernlab – an S4 package for kernel methods in R. Journal of 17 Statistical Software 11(9), 1–20 (2004)

¹⁸38. Liaw, A., Wiener, M.: Classification and regression by randomforest. R News 2(3), 18–22 (2002) 18

¹⁹39. Frasca, M., Bertoni, A., Re, M., Valentini, G.: A neural network algorithm for semi-supervised node label 19 learning from unbalanced data. Bioinformatics 43(C), 84–98 (2013)

²⁰40. Bertoni, A., Frasca, M., Valentini, G.: Cosnet: a cost sensitive neural network for semi-supervised learning in 20 graphs. Lecture Notes in Computer Science 6911(4), 219–234 (2011). doi:[10.18129/B9.bioc.COSNet](https://doi.org/10.18129/B9.bioc.COSNet)

²¹41. Szklarczyk, D., Franceschini, A., Wyder, S., Forsslund, K., Heller, D., Huerta-Cepas, J., Simonovic, M., Roth, 21 A., Santos, A., Tsafou, K.P., *et al.*: String v10: protein–protein interaction networks, integrated over the tree of 22 life. Nucleic acids research 43(D1), 447–452 (2014)

²³42. Türei, D., Korcsmáros, T., Saez-Rodríguez, J.: Omnipath: guidelines and gateway for literature-curated 23 signaling pathway resources. Nature methods 13(12), 966 (2016)

²⁴43. Lee, I., Blom, U.M., Wang, P.I., Shim, J.E., Marcotte, E.M.: Prioritizing candidate disease genes by 24 network-based boosting of genome-wide association data. Genome Research 21(7), 1109–1121 (2011). 25 doi:[10.1101/gr.118992.110](https://doi.org/10.1101/gr.118992.110)

²⁵44. Saito, T., Rehmsmeier, M.: The precision-recall plot is more informative than the ROC plot when evaluating 25 binary classifiers on imbalanced datasets. PLOS ONE 10(3), 0118432 (2015). 26 doi:[10.1371/journal.pone.0118432](https://doi.org/10.1371/journal.pone.0118432)

²⁶45. Jiang, Y., Metz, C.E., Nishikawa, R.M.: A receiver operating characteristic partial area index for highly sensitive 26 diagnostic tests. Radiology 201(3), 745–750 (1996)

²⁷46. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 27 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

²⁸47. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 28 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

²⁹48. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 29 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

³⁰49. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 30 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

³¹50. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 31 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

³²51. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 32 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

³³52. Saito, T., Rehmsmeier, M.: Precrec: fast and accurate precision-recall and roc curve calculations in r. 33 Bioinformatics 33 (1), 145–147 (2017). doi:[10.1093/bioinformatics/btw570](https://doi.org/10.1093/bioinformatics/btw570)

1 46. Bento, A.P., Gaulton, A., Hersey, A., Bellis, L.J., Chambers, J., Davies, M., Krüger, F.A., Light, Y., Mak, L., 1
2 McGlinchey, S., *et al.*: The chembl bioactivity database: an update. Nucleic acids research **42**(D1), 1083–1090 2
3 (2014)

4 47. Spearman, C.: footrulefor measuring correlation. British Journal of Psychology **2**(1), 89–108 (1906) 3
4 48. Mardia, K.V.: Some properties of classical multi-dimensional scaling. Communications in Statistics-Theory and 4
4 Methods **7**(13), 1233–1241 (1978)

5 549. Gower, J.C.: Some distance properties of latent root and vector methods used in multivariate analysis. 5
6 Biometrika **53**(3-4), 325–338 (1966) 6

7 50. Kanaan-Izquierdo, S., Ziyatdinov, A., Perera-Lluna, A.: Multiview and multifeature spectral clustering using 7
7 common eigenvectors. Pattern Recognition Letters **102**, 30–36 (2018) 7

8 51. Kanaan-Izquierdo, S.: multiview: A package with multiview clustering and dimensionality reduction methods. 8
8 <https://github.com/b2slab/multiview>

9 52. Team, B.C.: Bioconductor software project. Bioconductor Open Source Software For Bioinformatics 9
10 53. Gentleman, R.C., Carey, V.J., Bates, D.M., Bolstad, B., Dettling, M., Dudoit, S., Ellis, B., Gautier, L., Ge, Y., 10
11 Gentry, J., Hornik, K., Hothorn, T., Huber, W., Iacus, S., Irizarry, R., Leisch, F., Li, C., Maechler, M., Rossini, 10
11 A.J., Sawitzki, G., Smith, C., Smyth, G., Tierney, L., Yang, J.Y., Zhang, J.: Bioconductor: open software 11
11 development for computational biology and bioinformatics. Genome Biology **5**(R80) (2004) 11

12 1254. Vandin, F., Upfal, E., Raphael, B.J.: Algorithms for detecting significantly mutated pathways in cancer. Journal 12
13 of Computational Biology **18**(3), 507–22 (2011) 13

14 55. Mostafavi, S., Ray, D., Warde-Farley, D., Grouios, C., Morris, Q.: Genemania: a real-time multiple association 14
14 network integration algorithm for predicting gene function. Genome Biology **9**(S4), 1–15 (2008) 14

15 56. Bersanelli, M., Mosca, E., Remondini, D., Castellani, G., Milanesi, L.: Network diffusion-based analysis of 15
15 high-throughput data for the detection of differentially enriched modules. Scientific Reports **6**(34841), 1–12 15
16 (2016) 16

17 57. Vandin, F., Upfal, E., Raphael, B.J.: Algorithms for detecting significantly mutated pathways in cancer. IEEE 17
17 Signal Processing Magazine **30**(4), 87–97 (2013) 17

18 58. Valentini, G., Armano, G., Frasca, M., Lin, J., Mesiti, M., Re, M.: Ranks: a flexible tool for node label ranking 18
18 and classification in biological networks. Bioinformatics **32**(18), 2872–2874 (2016)

19 1959. Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple 19
20 testing. Journal of the royal statistical society. Series B (Methodological), 289–300 (1995) 20

21 **Figures** 21

22

23 **Tables** 23

24

25 **Table 1 List of methods included in this benchmark.** Method identifiers are shortened method 25
25 names used throughout the text. Other columns are self-explanatory.

26	Method Identifier	Method Name	Method Class	Implementation	Reference	26
27	pr	PageRank with a uniform prior	Baseline	igraph (Bioconductor [52, 53] package)	[25]	
random	Random	Baseline	R	(see text)		27
randomraw	Random Raw	Baseline	R	(see text)		
28	EGAD	Extending Guilt by Association' by Degree	Baseline	EGAD (Bioconductor package)	[24]	
ppr	Personalized PageRank	Diffusion	igraph (R package)	[26]		28
raw	Raw Diffusion	Diffusion	diffuStats (Bioconductor package)	[54]		
29	gm	GeneMania-based weights	Diffusion	diffuStats (Bioconductor package)	[55]	
mc	Monte Carlo normalised scores	Diffusion	diffuStats (Bioconductor package)	[56]		29
30	z	Z-scores	Diffusion	diffuStats (Bioconductor package)	[57]	
knn	K nearest neighbours	Semi-supervised learning	RANKS (R package)	[58]		30
wld	Weighted Sum with Linear Decay	Semi-supervised learning	RANKS (R package)	[58]		
31	COSNet	COst Sensitive neural Network	Supervised learning	COSNet (R package)	[38]	
bagsvm	Bagging SVM (based on ProDiGeI)	Supervised learning	kernlab (R package)	[33]		31
32	rf	Random Forest	Supervised learning	randomForest (R package) + Matlab (features)	[23]	
	svm	Support Vector Machine	Supervised learning	kernlab (R package) + Matlab (features)	[23]	32

33

33

Table 2 List of diseases included in this study. Diseases included in this study, with a minimum of 1
50 associated genes both in the known drug targets and the genetic categories (see text). The overlap 1
2 between these two lists of genes showed a degree of dependence between these two Open Targets 2
3 data streams for some of the diseases. P-values were calculated using Fisher's exact test and are 3
reported without and with correction for false discovery rate (Benjamini and Hochberg [59]).

Disease	N(genetic)	N(drugs)	Overlap	P-value	FDR
allergy	112	57	1	4.22e-01	4.42e-01
Alzheimers disease	208	103	4	1.10e-01	1.42e-01
arthritis	174	188	6	6.08e-02	1.03e-01
asthma	105	80	6	7.77e-05	5.70e-04
bipolar disorder	117	148	3	1.83e-01	2.12e-01
cardiac arrhythmia	75	177	6	9.15e-04	3.36e-03
chronic obstructive pulmonary disease (COPD)	154	116	6	4.18e-03	1.31e-02
coronary heart disease	111	171	4	7.86e-02	1.24e-01
drug dependence	75	143	6	2.96e-04	1.30e-03
hypertension	66	188	2	2.85e-01	3.14e-01
multiple sclerosis	71	167	4	1.83e-02	4.03e-02
obesity	69	194	3	1.06e-01	1.42e-01
Parkinson's disease	55	145	0	1	1
psoriasis	131	105	7	1.68e-04	9.23e-04
rheumatoid arthritis	138	95	5	5.18e-03	1.42e-02
schizophrenia	410	163	17	5.44e-05	5.70e-04
stroke	90	156	3	1.18e-01	1.44e-01
systemic lupus erythematosus (lupus)	126	109	5	6.30e-03	1.54e-02
type I diabetes mellitus	87	106	3	4.39e-02	8.04e-02
type II diabetes mellitus	130	154	4	9.14e-02	1.34e-01
ulcerative colitis	136	51	7	1.81e-06	3.98e-05
unipolar depression	123	121	4	3.81e-02	7.63e-02

18

18

19 Additional Files

19

Additional file 1 — Supplement

20

This document contains complementary material that supports our claims in the main body. It includes topics such 20
21 as descriptive statistics, topological properties of disease genes, raw metrics plots, method details, MDS plots, 21
alternative performance metrics and further explicative models.

22

23 Additional file 2 — MDS_plots

23

Complementary single-disease MDS plots and distance matrices.

24

25 Additional file 3 — Interactions.html.viewer

25

Stand-alone viewer to explore models with interaction terms.

26

27

27

28

28

29

29

30

30

31

31

32

32

33

33

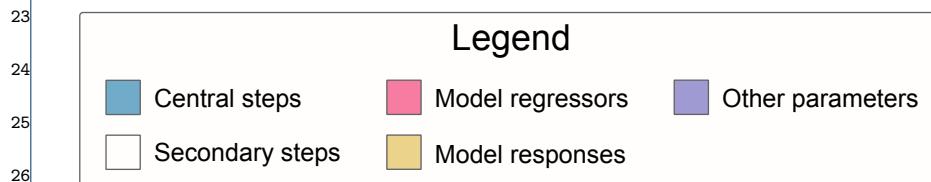
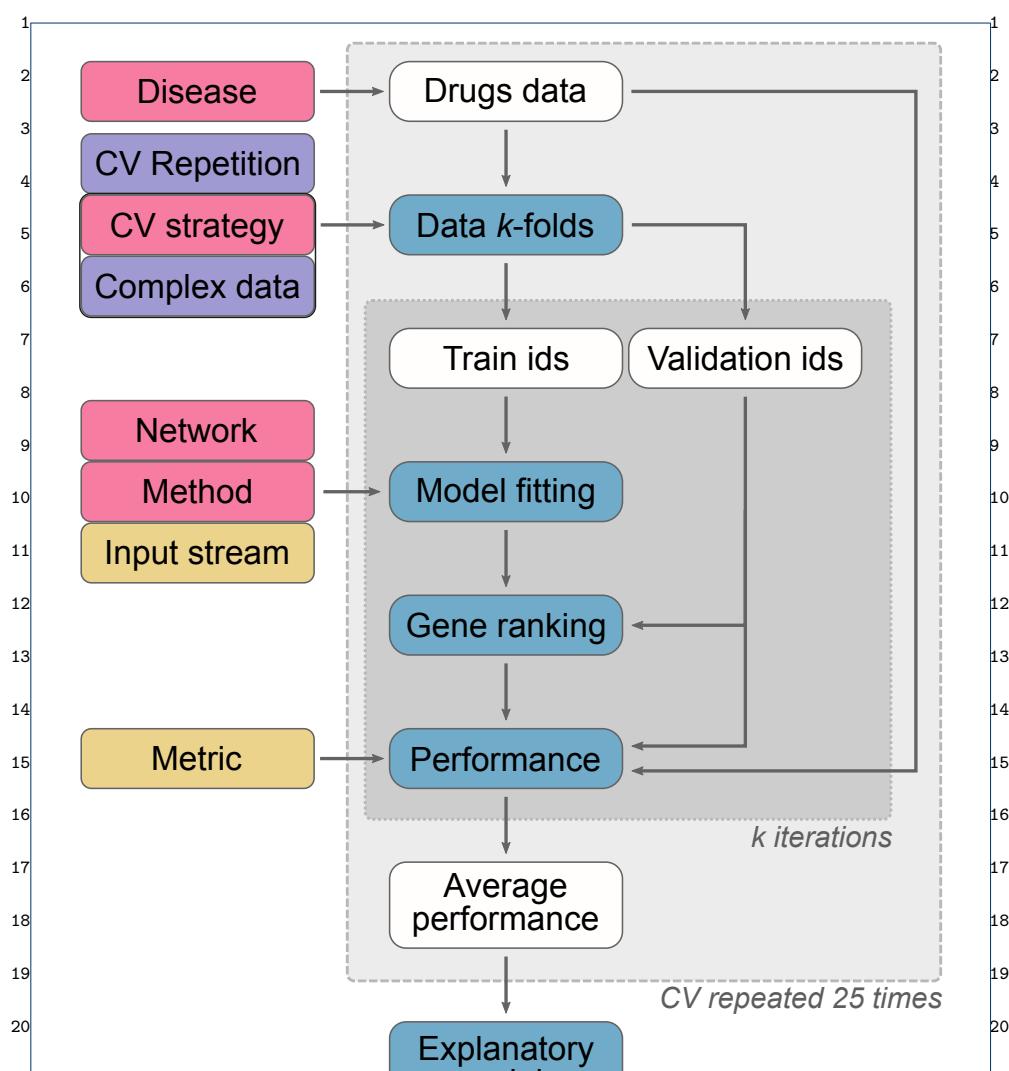


Figure 1 Benchmark overview. This work describes six performance metrics using two input streams (genetic association and drug-based genes) to predict drug-based genes for 22 common diseases. 3-fold cross validation (CV), repeated 25 times, was run under three CV strategies. The gene identifiers in each fold are determined using only the drugs data, regardless of the input. Two validation strategies are complex-aware and therefore needed this data to define the splits. 15 network-based methods (including 4 baselines) were evaluated, using two networks with different properties, by modelling their performance, averaged on every CV round. The explanatory models allowed hypothesis testing and a direct comparison between diseases, CV strategies, networks and methods.

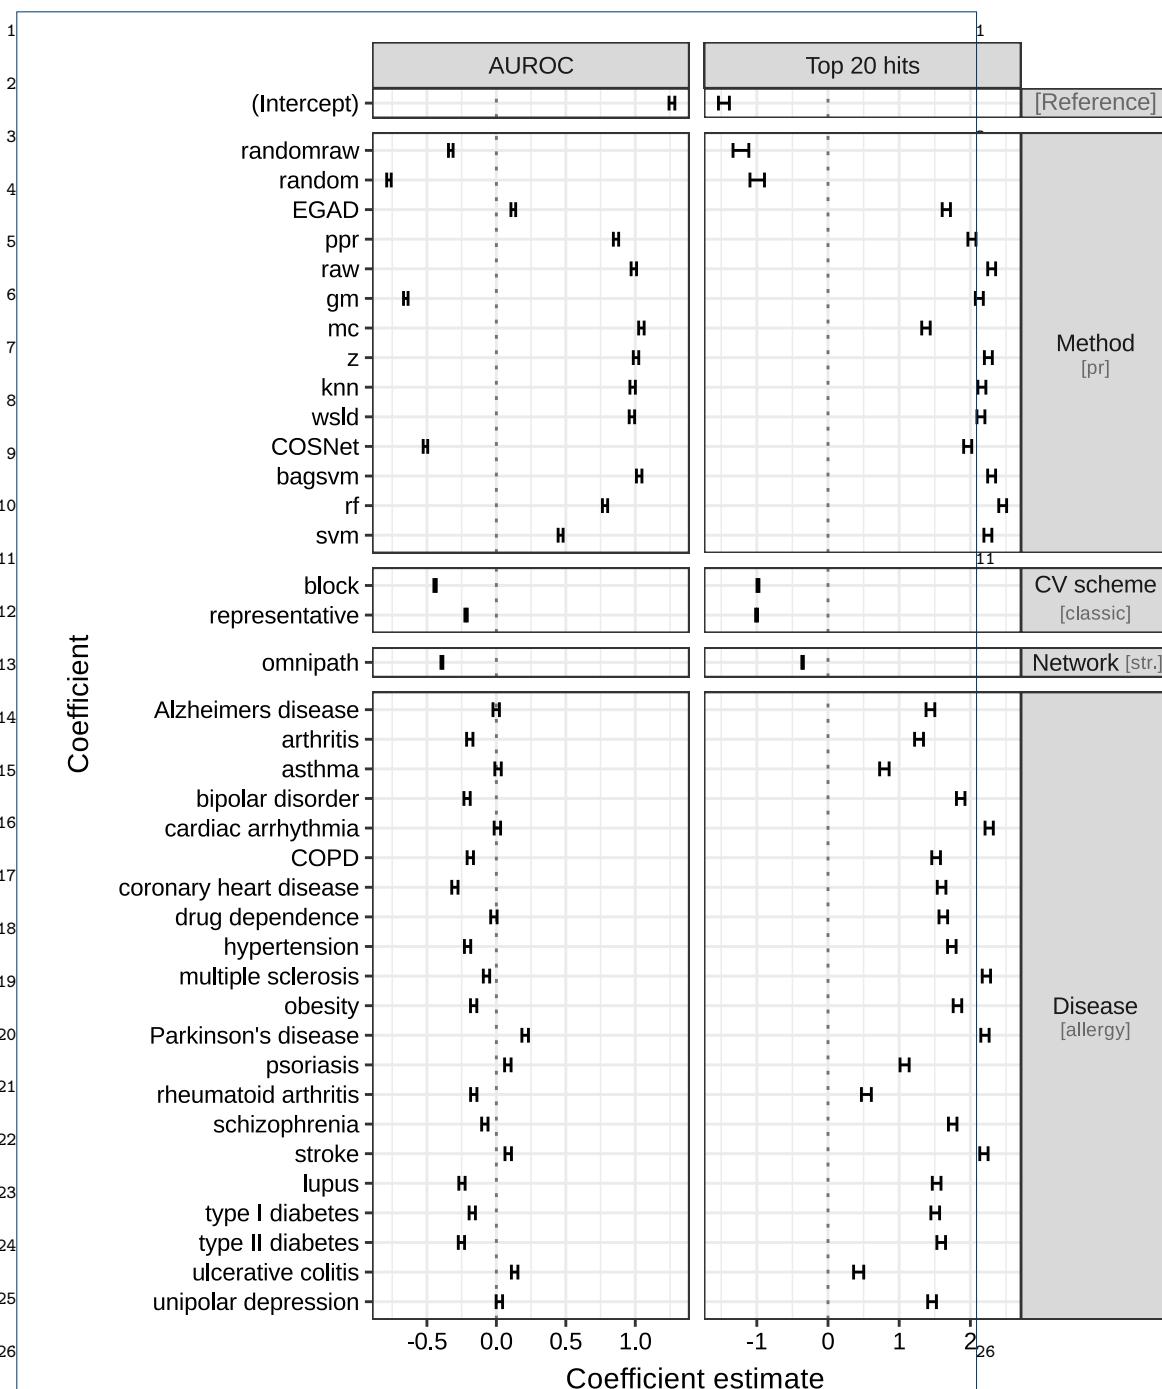


Figure 2 Additive models for AUROC and top 20 hits. Each column corresponds to a different model, whereas each row depicts the 95% confidence interval for each model coefficient. Rows are grouped by the categorical variable they belong to: method, cv scheme, network and disease. Each variable has a **reference level**, implicit in the intercept and specified in brackets: pr method, classic validation scheme, STRING network and allergy. Positive estimates improve performance over the reference levels, whereas negative ones reduce it. For example, the data suggest that method rf performs better than the baseline using both metrics, and is the preferred method using the top 20 hits. Switching from STRING to the OmniPath network, or from classic to block or representative cross-validation, has a negative effect on both performance metrics. Specific model estimates and confidence intervals can be found in the supplement, see tables S8 and S9.

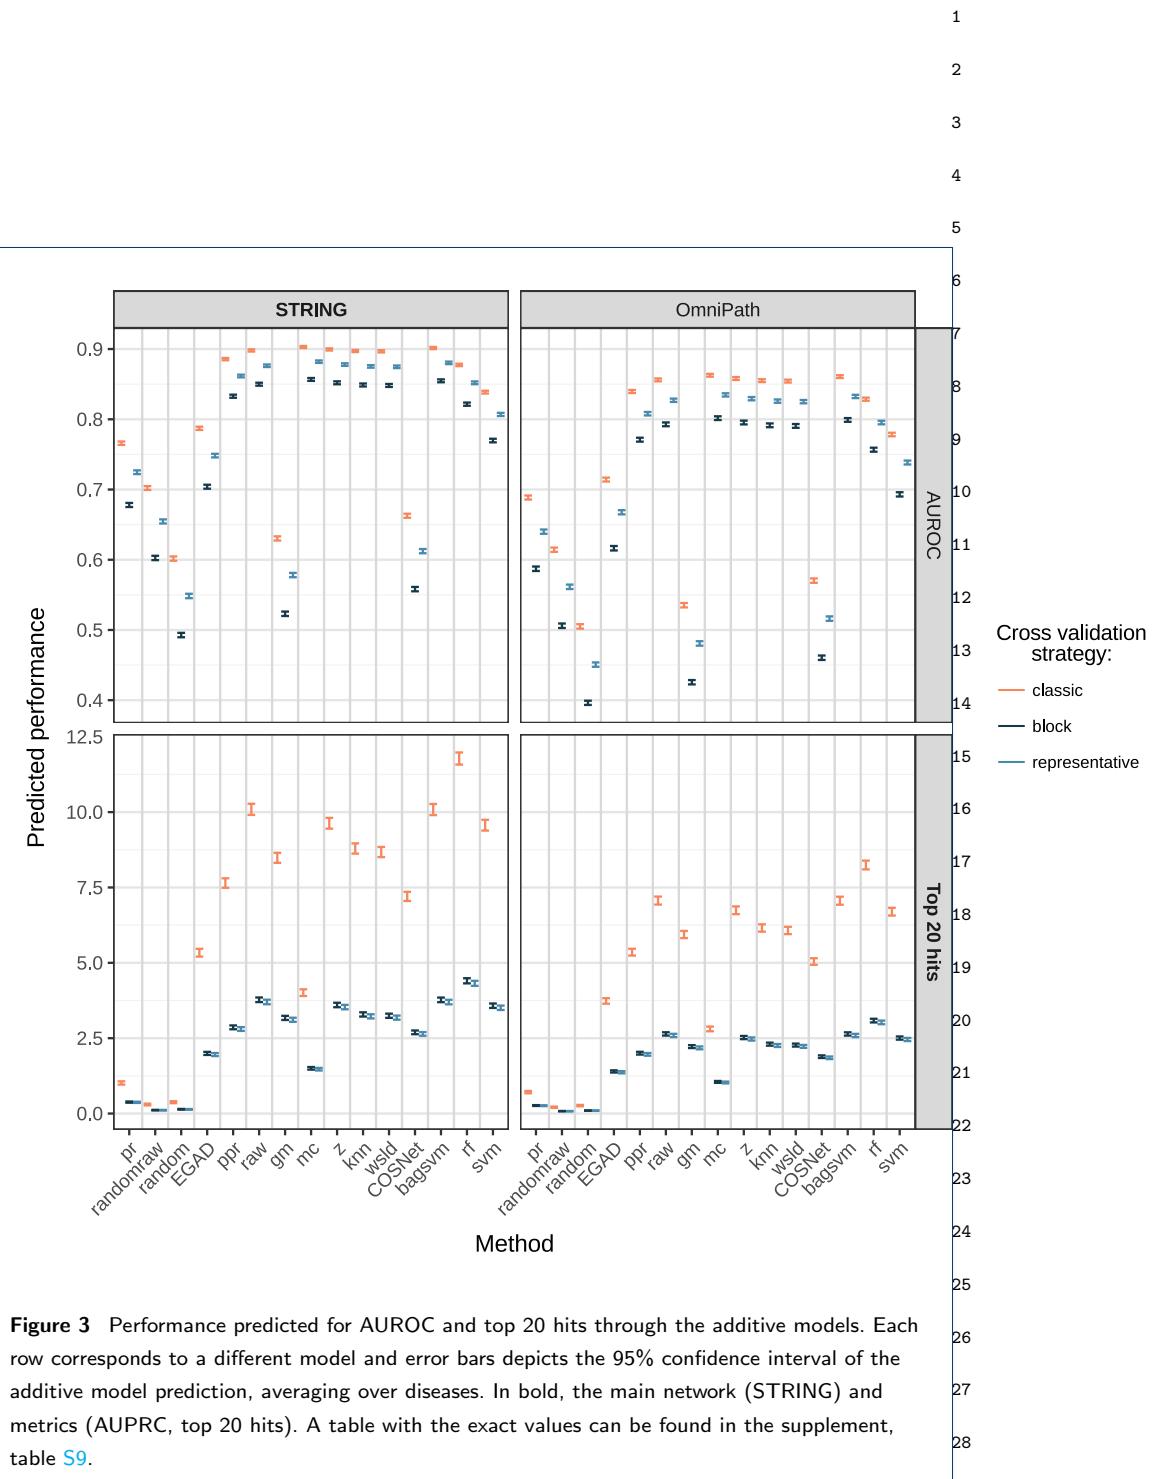
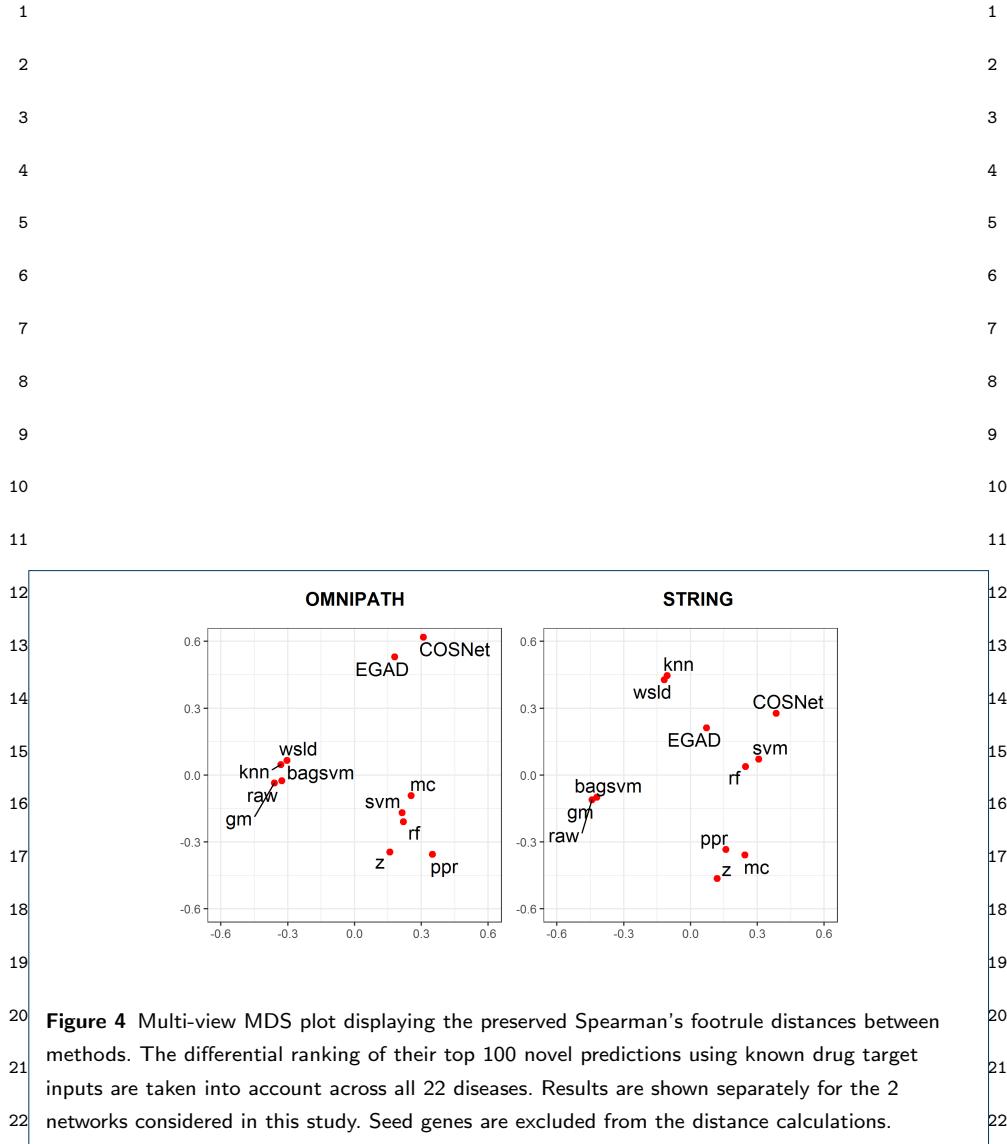


Figure 3 Performance predicted for AUROC and top 20 hits through the additive models. Each row corresponds to a different model and error bars depicts the 95% confidence interval of the additive model prediction, averaging over diseases. In bold, the main network (STRING) and metrics (AUPRC, top 20 hits). A table with the exact values can be found in the supplement, table S9.



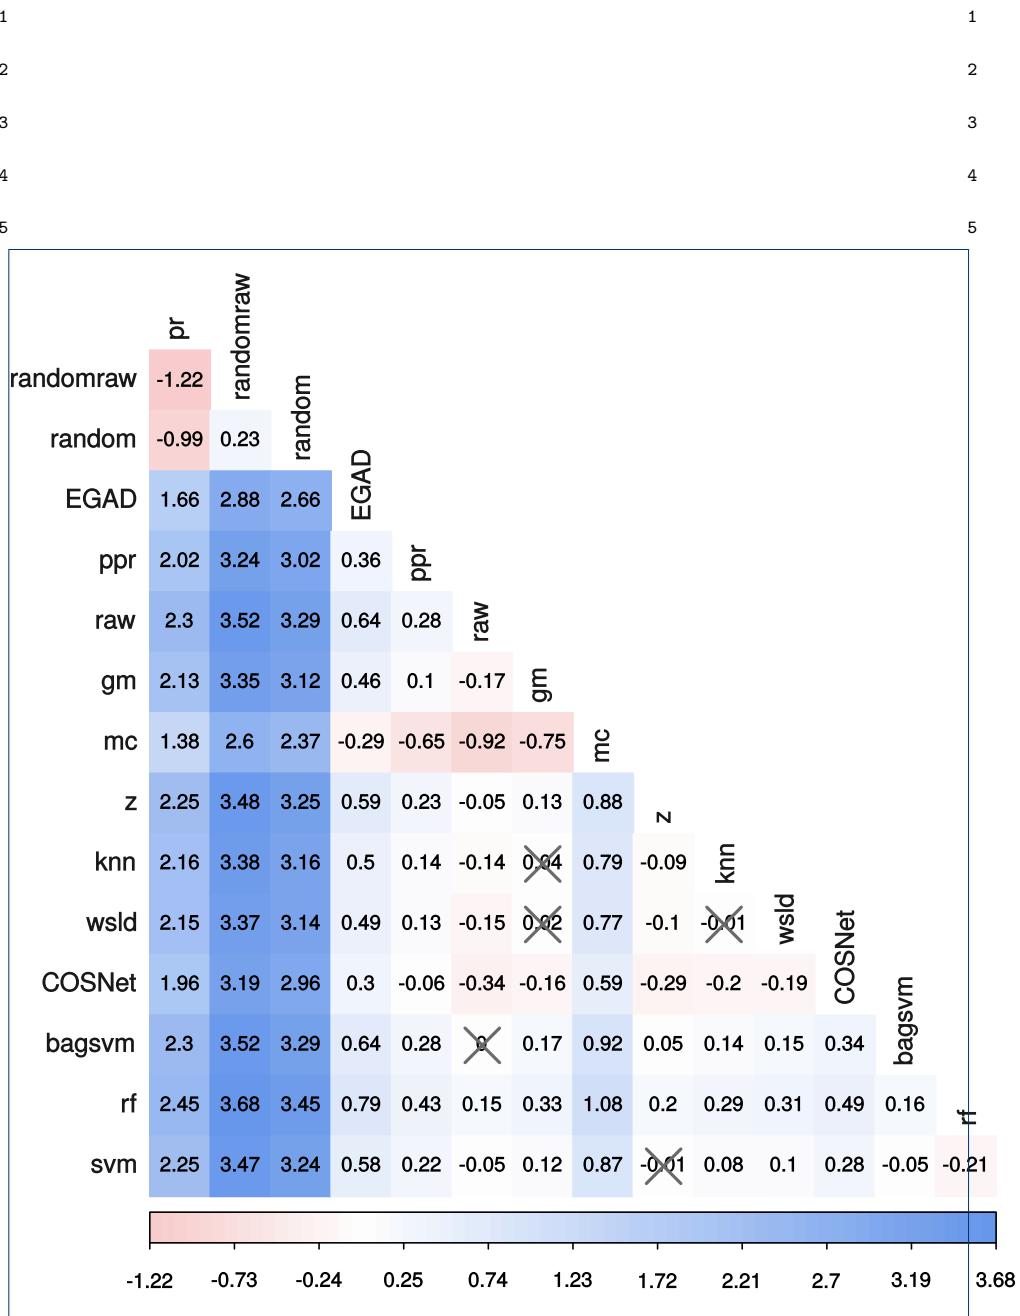
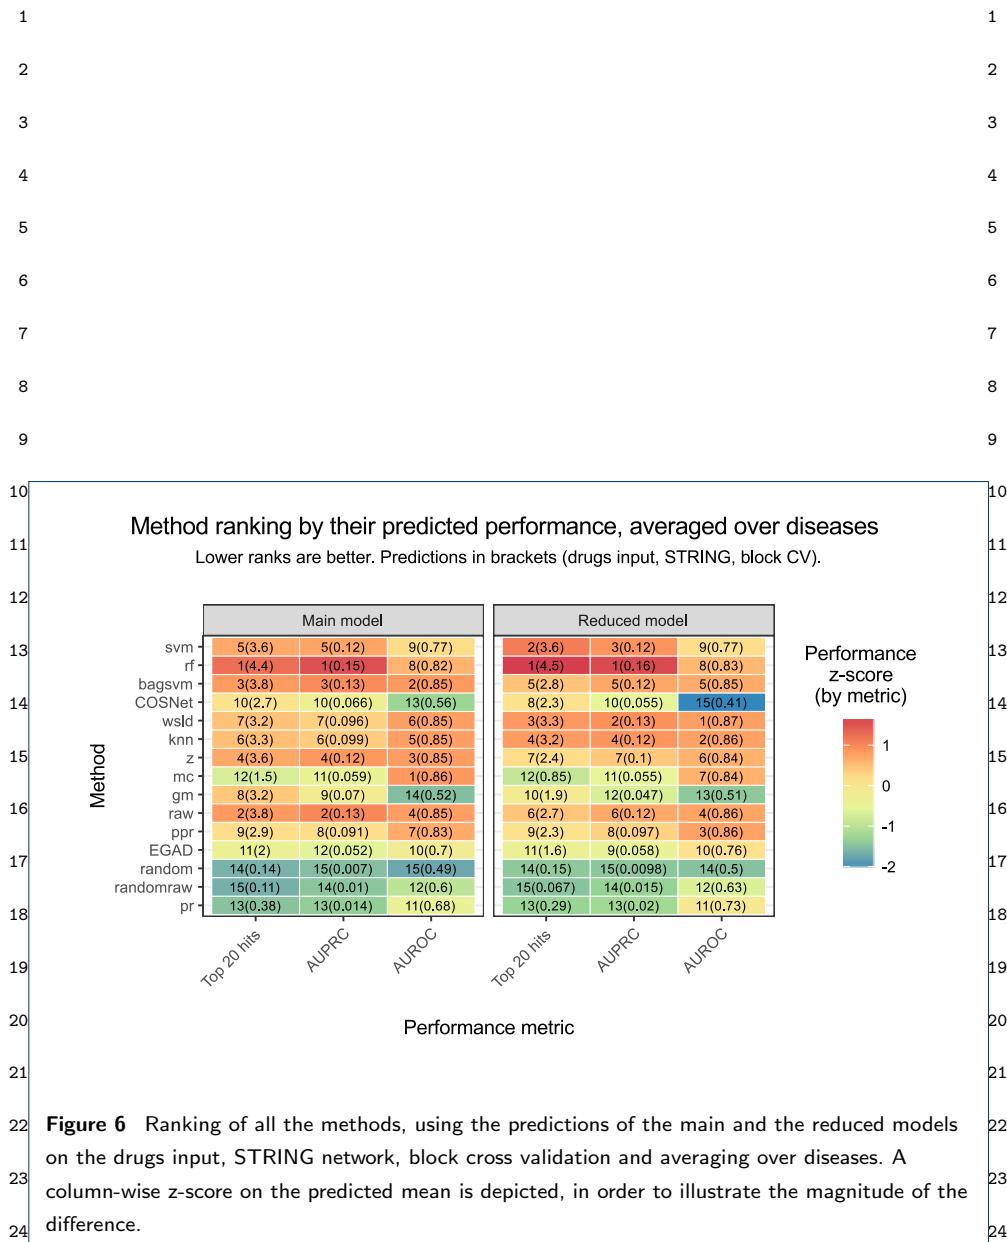
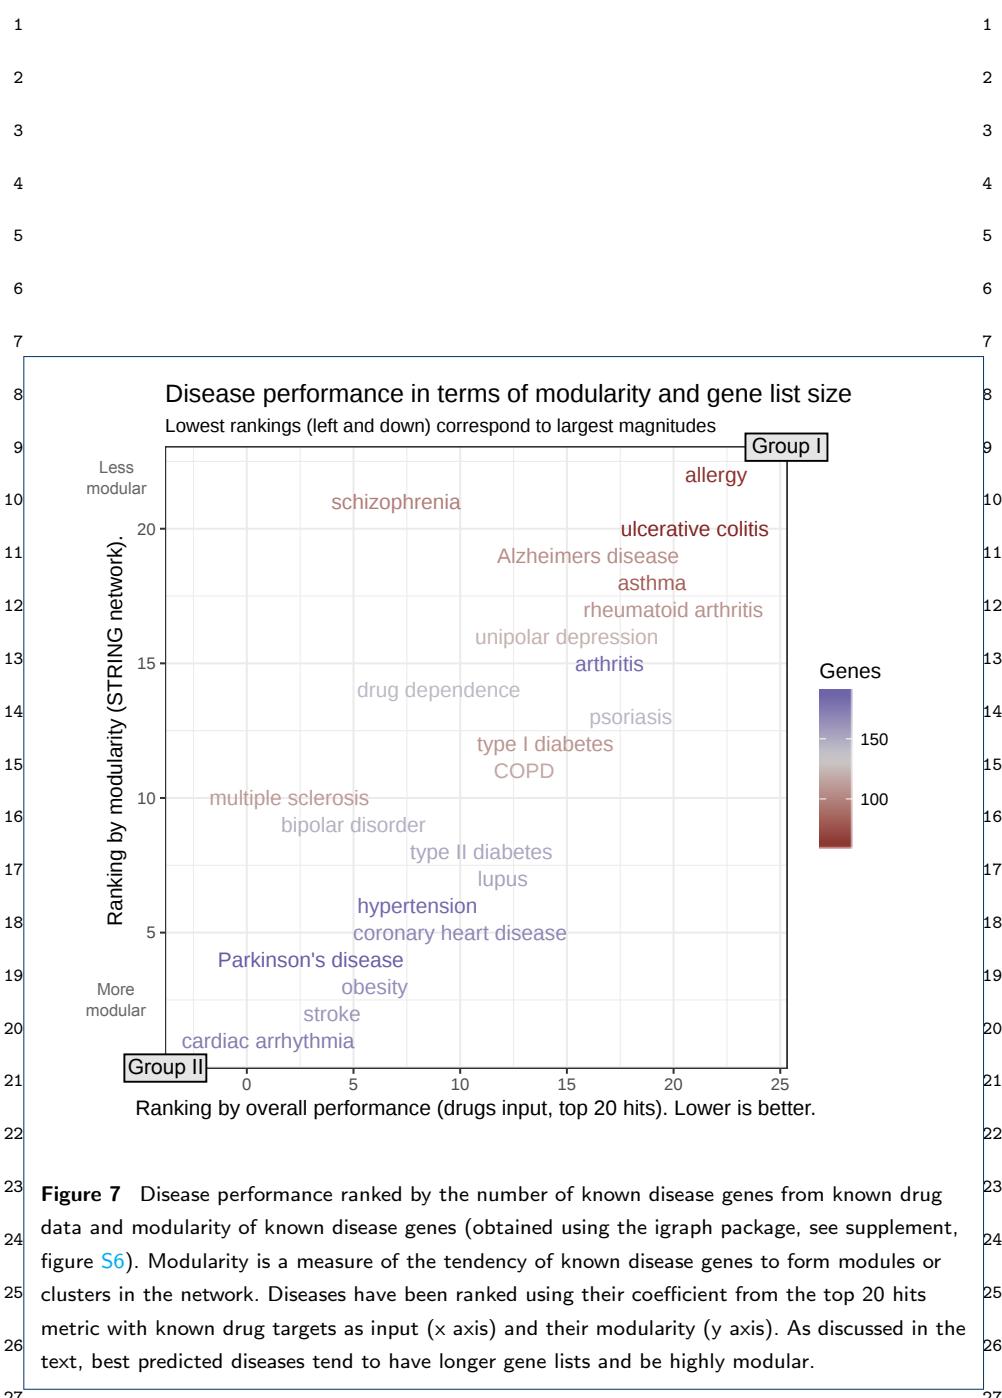


Figure 5 Pairwise contrasts on top 20 hits predicted by the quasipoisson model. Differences are expressed in the model space. Most of the pairwise differences are significant (Tukey's test, $p < 0.05$) – non-significant differences have been crossed out.





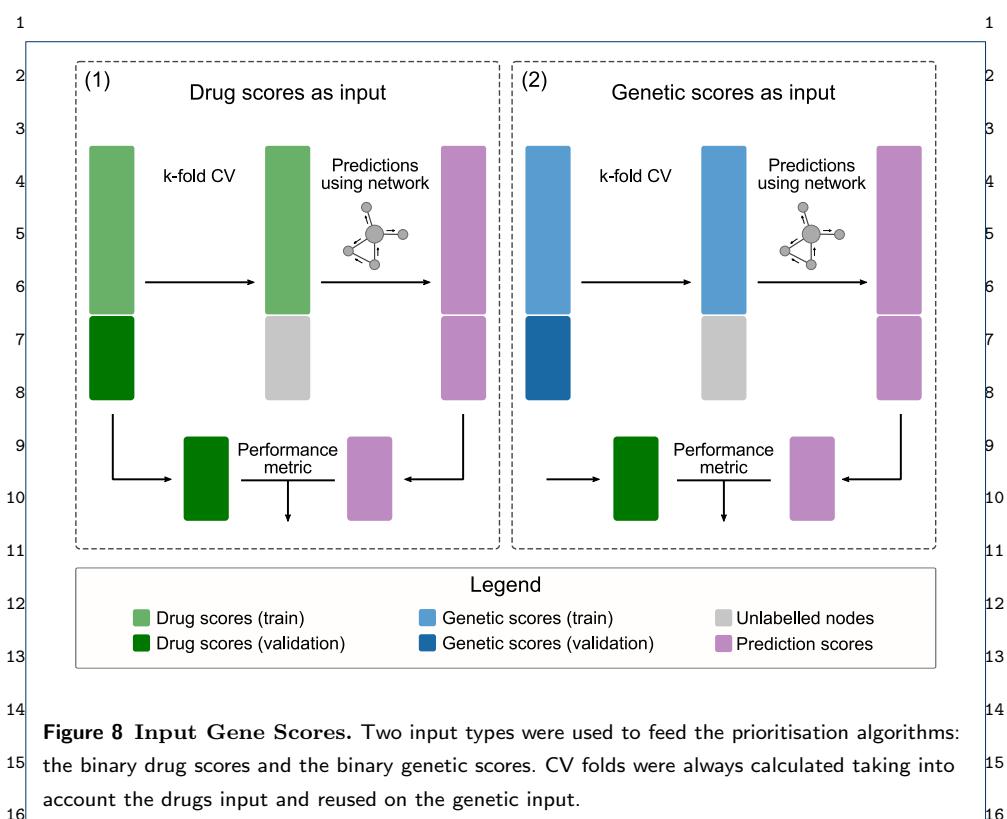


Figure 8 Input Gene Scores. Two input types were used to feed the prioritisation algorithms: the binary drug scores and the binary genetic scores. CV folds were always calculated taking into account the drugs input and reused on the genetic input.

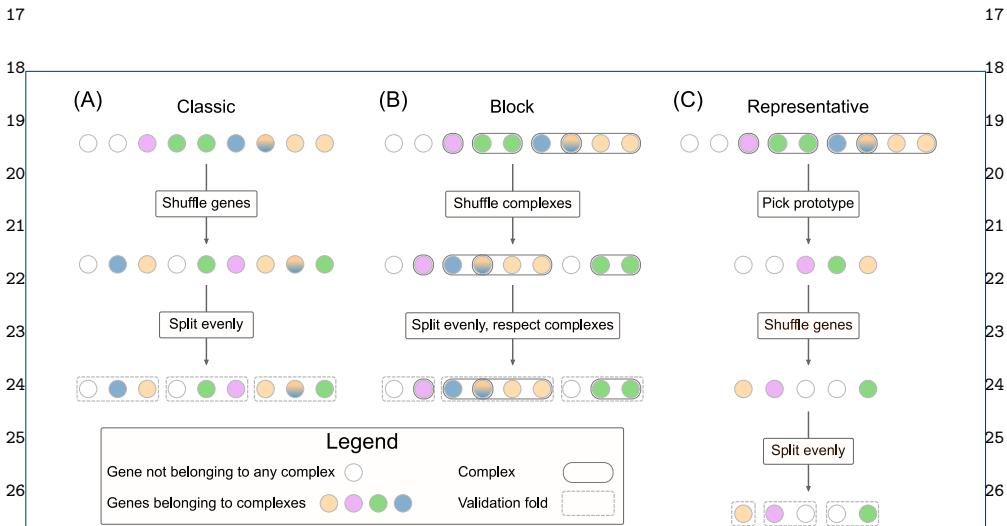


Figure 9 Cross-validation schemes. Three cross-validation schemes were tested. **(A):** standard k-fold stratified cross-validation that ignored the complex structure. **(B):** block k-fold cross-validation. Overlapping complexes were merged and the resulting complexes were shuffled. The folds were computed as evenly as possible without breaking any complex. **(C):** representative k-fold cross validation. Overlapping complexes were merged and the resulting complexes from which unique representatives were chosen uniformly at random. Then a standard k-fold cross-validation was run on the representatives, but excluding the non-representatives from train and validation.