Predictability of Genetic Interactions from Functional Gene Modules

Jonathan H. Young*,† and Edward M. Marcotte†,‡,1

*Institute for Computational Engineering and Sciences, †Center for Systems and Synthetic Biology, and ‡Department of Molecular Biosciences, The University of Texas at Austin, Texas 78712

ORCID ID: 0000-0003-3365-2702 (J.H.Y.)

ABSTRACT Characterizing genetic interactions is crucial to understanding cellular and organismal response to gene-level perturbations. Such knowledge can inform the selection of candidate disease therapy targets, yet experimentally determining whether genes interact is technically nontrivial and time-consuming. High-fidelity prediction of different classes of genetic interactions in multiple organisms would substantially alleviate this experimental burden. Under the hypothesis that functionally related genes tend to share common genetic interaction partners, we evaluate a computational approach to predict genetic interactions in *Homo sapiens*, *Drosophila melanogaster*, and *Saccharomyces cerevisiae*. By leveraging knowledge of functional relationships between genes, we cross-validate predictions on known genetic interactions and observe high predictive power of multiple classes of genetic interactions in all three organisms. Additionally, our method suggests high-confidence candidate interaction pairs that can be directly experimentally tested. A web application is provided for users to query genes for predicted novel genetic interaction partners. Finally, by subsampling the known yeast genetic interaction network, we found that novel genetic interactions are predictable even when knowledge of currently known interactions is minimal.

Determining the genetic interactions in an organism provides a basis for understanding how the role of a gene is influenced by the action of any other gene. By definition, two or more genes interact when combining variants of each gene produces a significantly pronounced phenotype when compared to the phenotypes of individual variants (Mani *et al.* 2008; Baryshnikova *et al.* 2013). The applications of exploiting such interactions extend to drug target discovery. Strategies such as targeting genes that interact with cancer-specific mutations have been proposed and reviewed extensively (Ashworth *et al.* 2011; Fece de la Cruz *et al.* 2015) and have led to clinical trials (Fong *et al.* 2009). Because experimental determination of genetic interactions involves examining all possible pairs from a group of genes, practical difficulties arise when a comprehensive interaction map of an entire organism is desired. Multicellular organisms present the challenge of various differentiated cell types, each having potentially differing genetic interactions. Moreover, there are different kinds of genetic interactions, ranging from those based on growth effects to other phenotypic effects. There exists a need to either reduce the search space for testing genetic interactions or to reliably predict them. Here, we evaluate a computational approach to predict and validate different types genetic interactions across multiple organisms.

Previous studies to predict genetic interactions leveraged existing sources of biological information. Integration of biological features in yeast (*i.e.*, gene coexpression, protein interaction and function) and their associated network topological properties guided the training of probabilistic decision trees to predict synthetic sick or lethal (SSL) interactions (Wong *et al.* 2004). In a similar vein, an ensemble classifier was trained on a set of 152 genetic interaction-independent features to predict SSL in yeast (Pandey *et al.* 2010). Compiling multiple biological features has also been extended to more than one organism. By considering the orthologous gene pairs among yeast, fly, and worm, features such as functional annotation were used to train a logistic regression model to predict a

genome-wide map of genetic interactions (Zhong and Sternberg 2006). Alternatively, studies have also explored network-based approaches for genetic interaction prediction. Novel SSL interactions were predicted by way of a diffusion kernel on a network of known SSL gene pairs (Qi *et al.* 2008). Interrogating functional gene networks that were constructed from integration of biological data from literature have proven useful in predicting modifier genes in yeast and worm (Lee *et al.* 2010). Many of these approaches have focused on a single genetic interaction type in a single organism.

Here, we examine an algorithm to predict multiple types of genetic interactions across diverse organisms based on the hypothesis that genes strongly participating in shared functions also share common genetic interaction partners. Our approach relies on a functional gene network for a given organism and knowledge of known genetic interactions of a particular type. We tested our approach on three organisms, human (*Homo sapiens*), fly (*Drosophila melanogaster*), and yeast (*Saccharomyces cerevisiae*), and found predictability across different types of genetic interactions. We also investigated how some interactions are enriched in yeast and human gene modules, specifically protein complexes, and the degree to which genetic interactions need to be experimentally determined before enrichment can be found.

MATERIALS AND METHODS

For various classes of genetic interactions in human, fly, and yeast, a list of genes and each of their known genetic interaction partners were assembled. A gene and its known interaction partners are collectively referred to as a "seed set." Seed sets with only a single interacting gene pair were filtered out. Receiver operating characteristic (ROC) analysis was performed to quantify whether the interaction partners of any given gene are clustered in the organism's functional gene network. Specifically, for every group of interaction partners of a gene, a score vector consists of entries that are sums of functional network edge weights between each gene in the network to the interaction partners. Because there are no self-edges in the network, leave-one-out cross-validation is carried out on the known interaction partners. An accompanying label vector indicates whether each gene in the network is indeed an interaction partner. The two vectors yield an ROC curve and the corresponding area under the curve (AUC). A seed set's AUC is the measure of how tightly connected the interaction partners are in the functional network, and therefore how predictive the seed set is for novel interactions (Lee *et al.* 2010). None of the known genetic interactions used for prediction were contained in the functional gene network. There were no fly or human genetic interactions incorporated into their respective functional gene networks. The yeast functional network did not include any genetic interactions discovered after 2006; every interacting gene pair before that year in BioGRID was excluded.

Enrichment of genetic interactions within yeast and human protein complexes was calculated with a binomial model defined as $P(X=k) = \binom{n}{k} p^k (1-p)^{n-k}$, where the background probability, p, equals the proportion of all possible gene pairs that are genetically interacting. The number of trials, n, is the number of possible gene pairs in the complex, and k equals the number of interacting pairs in the protein complex.

Statistical analysis

If k is the number of genetic interactions within a protein complex, then the corresponding p-value is $P(X \geq k)$, according to a binomial model
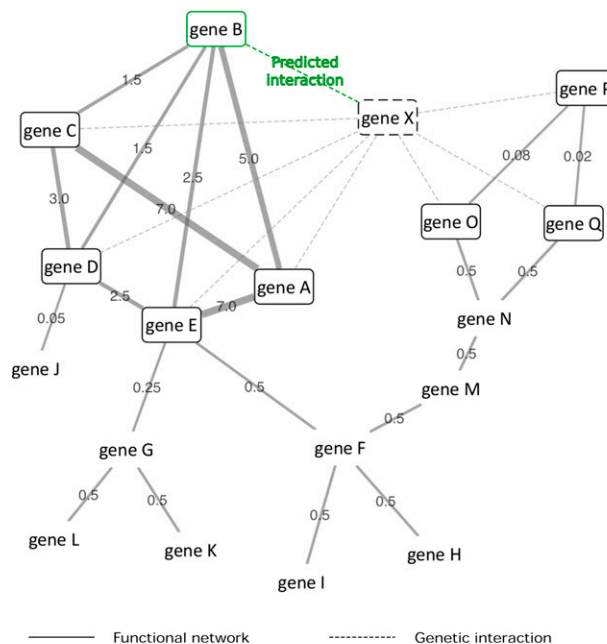
Denoting by Ω the set of genetic interaction partners of X:

$$\begin{bmatrix} \sum_{g \in \Omega} LLS_{g,A} & \sum_{g \in \Omega} LLS_{g,B} & \sum_{g \in \Omega} LLS_{g,C} & \cdots & \sum_{g \in \Omega} LLS_{g,O} & \cdots \\ 1 & 0 & 1 & \cdots & 1 & \cdots \end{bmatrix}$$
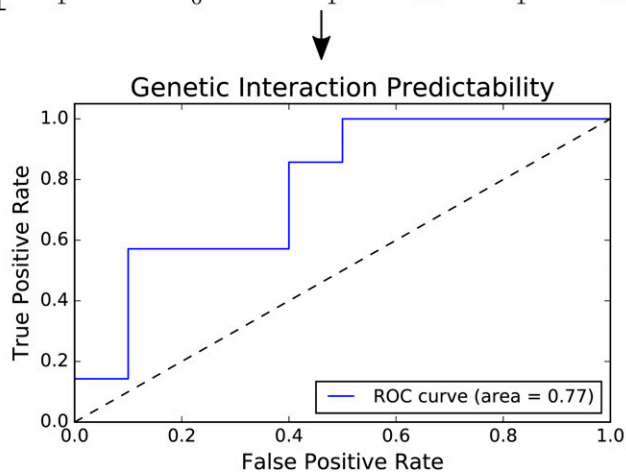
Figure 1 Genetic interaction prediction. Dashed edges indicate known genetic interactions. Solid edges connect genes that participate in the same biological process, with log likelihood scores (LLS) as edge weights reflecting the degree of confidence in the genes' shared functionality. Genes A, C–E are genetic interaction partners of gene X and members of a functional net cluster; then the remaining cluster member, gene B, is a predicted interaction partner of gene X as well. Candidate clusters are evaluated by first assigning scores to each gene in the network by summing the edge weights, as shown in the first row of the matrix. $LLS_{g,A}$ denotes the LLS between genes g and A. The second row is populated with binary labels indicating whether the gene is a known interaction partner of X. In this fashion, an ROC curve is constructed to yield an AUC.

previously described, with control of FDR at 5% using the Benjamini–Hochberg procedure (Benjamini and Hochberg 1995). Seed sets with AUC ≥ 0.9 were considered highly predictive of novel genetic interactions.
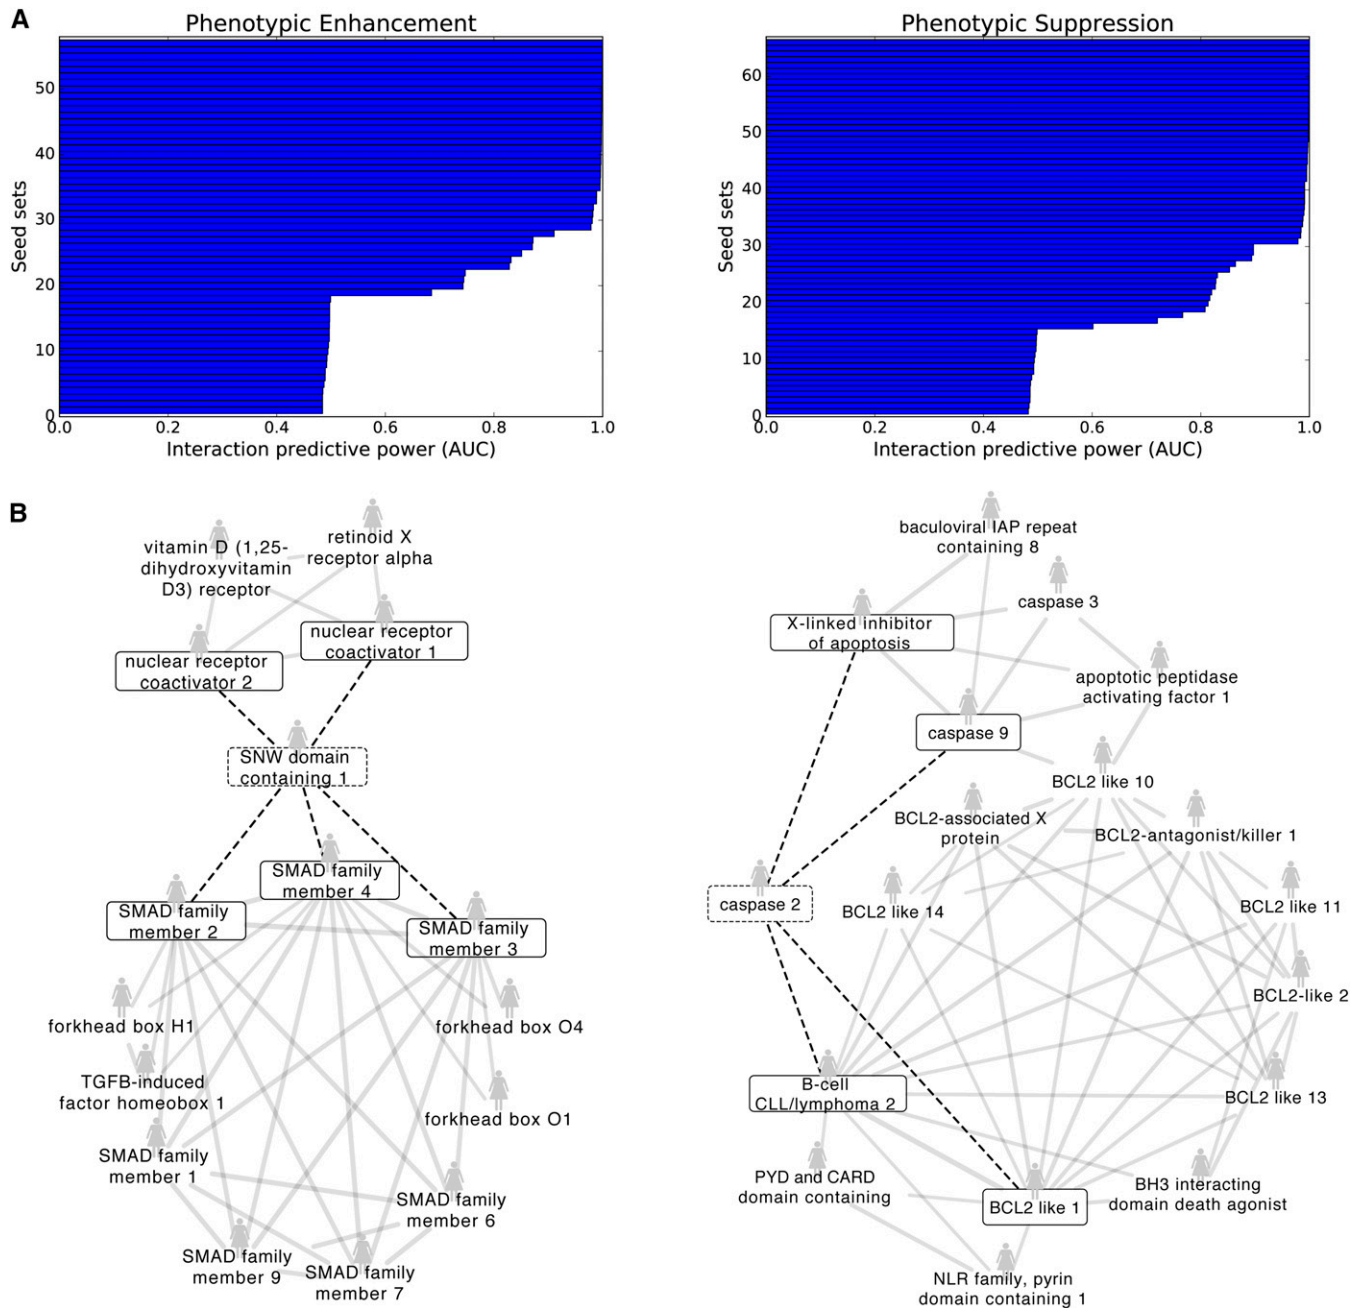
Figure 2 Predictive functional net clusters yield novel phenotypic enhancing and suppressing human gene pairs. (A) Each horizontal bar represents the set of known genetic interaction partners of a specific human gene; each of these sets is referred to as a "seed set." High AUC scores indicate that the interaction partners participate together in a cluster in HumanNet, the human functional gene network. Therefore, other members of the cluster are predicted as novel interaction partners. (B) Shown are two examples of well-defined HumanNet clusters that are highly predictive for phenotypic enhancement (left) and suppression (right), with the known interactions from the seed set denoted by the boxed genes and dashed edges.

Data availability

All genetic interactions were downloaded from version 3.4.130 of BioGRID (Stark *et al.* 2006). Organism-specific functional gene networks were downloaded for human (Lee *et al.* 2011), fly (Shin *et al.* 2015), and yeast (Lee *et al.* 2007). Previous studies served as sources of protein complexes for yeast and human (Hart *et al.* 2007; Ruepp *et al.* 2010). Python code using the Matplotlib (Hunter *et al.* 2007), scikit-learn (Pedregosa *et al.*

2011), and *mygene* (Wu *et al.* 2012) libraries is available at https://bitbucket.org/youngjh/genetic_interact. All network visualizations were produced in Cytoscape (Shannon *et al.* 2003). A supplementary web page at http://marcottelab.org/Genetic_Interact allows users to query a gene of interest. If the gene has known genetic interaction partners that are predictive, then the functional network cluster is displayed. Raw data files listing the seed sets with AUC ≥ 0.7 are also available.
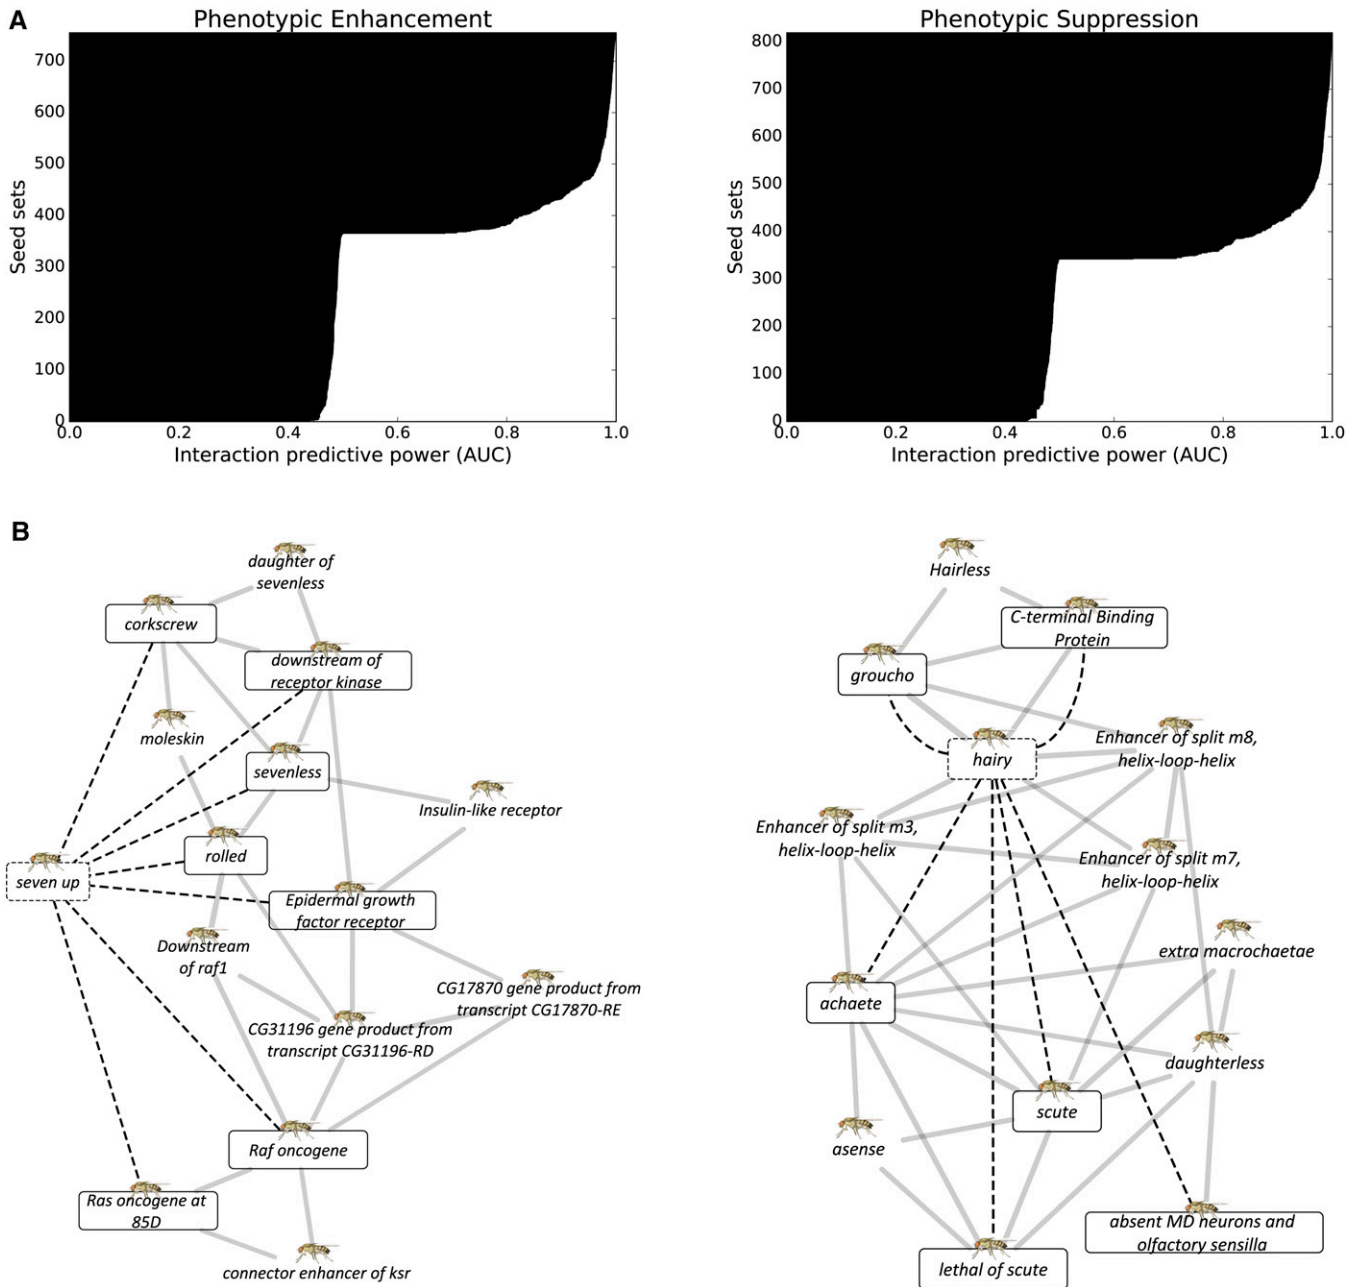
Figure 3 FlyNet predictability for phenotypic enhancing and suppressing genetic interactions. (A) Each horizontal bar represents a single fly gene that is known to interact with a number of other genes. (B) Predictive seed gene sets are shown for phenotypic enhancement (left) and suppression (right).

RESULTS

We sought to determine whether clusters of functionally related genes, for example, genes A–E in Figure 1, are predictive of genetic interactions. In this example, genes A and C–E are known to share genetic interactions with gene X, and our hypothesis would suggest gene B as a novel interaction partner of X. Our method identifies predictive clusters by leave-one-out cross-validation and ROC analysis; when applied to the network in Figure 1, each of genes A and C–E are individually withheld as known interaction partners one at a time, and predicted back with high recall. Subsequently, gene B is a novel high-confidence predicted interaction partner of X. The approach described here was evaluated for several classes of phenotypic and growth-based genetic interactions in human, fly, and yeast.

The human functional gene network is predictive for phenotypic genetic interactions

As shown in Figure 2A, our method demonstrated high performance in predicting phenotypic enhancing and suppressing human gene pairs. In these interactions, a double mutant has an enhanced or suppressed phenotype (other than growth) in comparison to either of the single mutants. The plots for phenotypic enhancement and suppression in Figure 2A display the performance of seed sets, each of which are defined as a group of known phenotypic enhancing or suppressing partners of a particular gene. There are 57 phenotypic enhancement seed sets, of which 30 have an AUC ≥ 0.9. Similarly, 36 of 66 phenotypic suppression seed sets have an AUC ≥ 0.9. The AUC is the area

	H. sapiens	D. melanogaster	S. cerevisiae
Dosage growth defect	Not tested	Not tested	176/488
Dosage lethality	2/3	Not tested	116/201
Dosage rescue	5/13	Not tested	203/638
Phenotypic enhancement	30/57	322/449	287/1175
Phenotypic suppression	36/66	398/518	223/1088
Synthetic growth defect	4/37	Not tested	576/2436
Synthetic rescue	2/3	5/6	218/972
Synthetic lethality	Not tested	Not tested	221/1469
Negative genetic	Not tested	Not tested	65/4352
Positive genetic	Not tested	Not tested	55/2826

For each fraction, the numerator indicates the number of seed sets with AUC ≥ 0.9 and the denominator equals the total number of seed sets tested. Definitions of each genetic interaction type are listed in Table S1.

under the ROC curve that measures how well the known interaction partners rank in our leave-one-out cross-validation scheme. Those that are not predictive are the ones with an AUC ≈ 0.5, indicating that their predictability is no better than random. For the most part, seed sets are either at least moderately predictive, or not at all. The sharp drop in AUC to 0.5 in Figure 2A does not appear to correspond with genes that simply have no edges above the background in the functional network. It does not appear that all random AUCs are the result of a small set size, as seed set sizes over two also produce random AUCs. As shown in Supplemental Material, Figure S1, low incident edge weight does not correlate with low AUC.

Shown in Figure 2B are illustrative seed sets with high predictability that form well-defined clusters in the human functional gene network, HumanNet. For clarity, only functional network edges with log likelihood scores (LLS) above 3.0 are shown. Furthermore, HumanNet genes are shown only if they connect to at least two of the known genetic interaction partners. The seed set consisting of the SNW domain containing one gene (*SNW1*) that phenotypically enhances members of the SMAD family and nuclear receptor coactivators yielded an AUC of 0.91. The prediction is that *SNW1* also phenotypically enhances with other members of the SMAD family along with members of the forkhead box. In the phenotypic suppression case, we find that known phenotypic suppressors of caspase 2 are tightly functionally linked with members of the BCL2-like family, among other genes. With a resulting AUC of 0.90, these BCL2-like genes are expected to participate in phenotypic suppression with caspase 2.

Fly phenotypic enhancement and suppression interactions are predicted from functional net clusters

Similar to the human case, the fly functional network FlyNet is particularly predictive of phenotypic enhancement and suppression, as shown in Figure 3. A larger proportion of the seed sets are predictive than in the human case. For phenotypic enhancement, 322 out of 449 seed sets had AUC ≥ 0.9 and 398 phenotypic suppression seed sets (out of 518) met the same threshold. Figure 3B shows a well-defined gene cluster (AUC = 0.94) containing phenotypic enhancement interaction partners of *seven up*. From this cluster, genes involved in the *sevenless* signaling and the *Drosophila epidermal growth factor receptor* signal transduction pathways achieved high recall, and neighbor genes also involved in the same signaling pathways are expected to phenotypically enhance *seven up*.
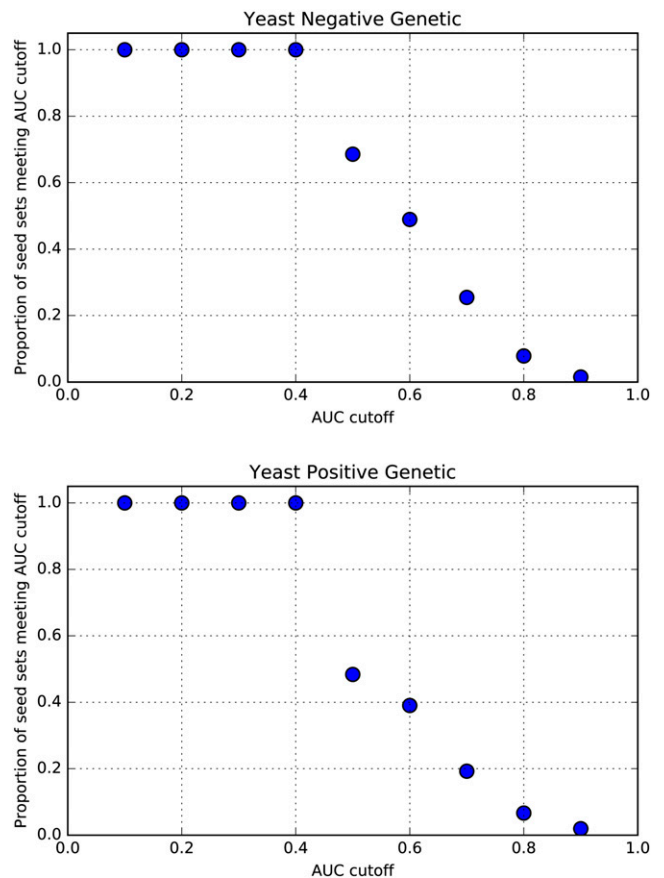

Figure 4 Performance in unbiased yeast screens as a function of AUC cutoff. For both negative and positive genetic interactions in *S. cerevisiae*, the proportion of seed gene sets exceeding a set AUC threshold ranging from 0.1 to 0.9 is shown. Nonrandom performance is observed and increases as the AUC threshold is relaxed.

Turning to phenotypic suppression, several *Enhancer of split* genes are tightly clustered (AUC = 0.98) with known phenotypic suppressors of *hairy* that include the *achaete-scute* complex, thereby implicating them as additional, novel phenotypic suppressors of *hairy*.

High-confidence predictability is found in human, fly, and yeast

The full range of various genetic interaction classes that were analyzed from BioGRID are listed in Table 1. Genetic interactions were generally based on phenotypic effects or growth and lethality measurements. Each entry in Table 1 lists the number of predictive seed sets having AUC ≥ 0.9 of out the total examined. For human, our method performed well primarily for phenotypic enhancement and suppression, as described above, but did not offer predictability for the dosage lethality and synthetic growth defect and rescue interactions determined to date. For fly, most of the known interactions fall into the phenotypic enhancement and suppression categories, for which high predictability was observed. In both human and fly, several classes of interactions have not been extensively determined and thus were untested in our prediction scheme.

Our method also performed well in most of the interaction categories for *S. cerevisiae* (Figure S2 and Table 1). Notably, negative and positive genetic interactions fared poorly as few predictive seed sets were identified, even though most of the experimentally determined interactions in yeast fall into these categories.
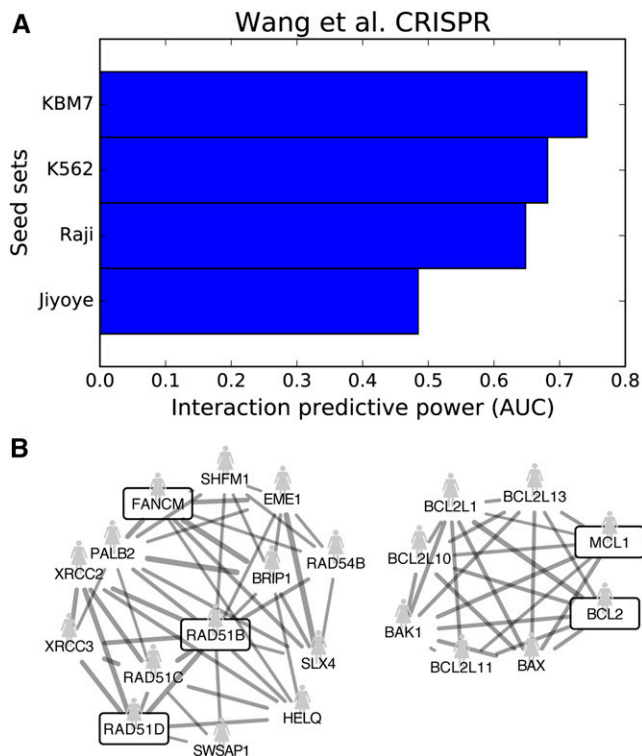
Figure 5 Predictability observed in unbiased screens of higher eukaryotes. A CRISPR-Cas9 screen against four human cancer cell lines uncovered cell line-specific essential genes (Wang *et al.* 2014). (A) The specific hits for each cell line were used as the seed sets, and we found moderate predictability in three of the four cell lines, especially for KBM7. (B) Shown is the predictive functional net cluster induced by KBM7-specific hits, which are the boxed genes. Solid edges indicate connections in the human functional network, HumanNet. For clarity, only edges above a LLS weight cutoff of 3.0 and genes with more than one interaction to the hits are shown.

Predictability observed in unbiased screens of eukaryotes

The negative and positive genetic interaction screens for yeast are nearly exhaustive and therefore largely unbiased. We investigated the performance on the yeast data using different AUC cutoffs and found that predictability increased as the AUC threshold was relaxed (Figure 4). This suggests that our method predicts positive and negative interactions even in yeast from unbiased data, albeit to more modest levels.

Given the sparsity of currently known genetic interactions in higher eukaryotes, we analyzed a genome-scale CRISPR screen against four human cancer cell lines that uncovered essential genes specific to each cell line (Wang *et al.* 2014). The screen was not previously included among the BioGRID datasets in our study and, unlike other data on higher eukaryotes, is unbiased in the sense that there was no selection of candidate genetic interaction partners. Rather, the characteristic mutations of each cell line determined the specific hits. As shown in Figure 5, our method found moderate predictability from three of the four cell lines, especially for KBM7, with an AUC of 0.74.

Protein complexes inform trends of genetic interaction predictability

With genetic interactions predicted across multiple organisms, it was natural to investigate their evolutionary conservation. In particular, if a protein complex were enriched in genetic interactions then perhaps a homologous protein complex would also exhibit similar enrichment. We found enrichment of various types of interactions within yeast protein complexes, but none thus far for human. Therefore, instead the problem shifted to identifying the degree to which genetic interactions must be determined in order to find enrichment, and therefore predictability. Using yeast as a test case, simulations successively withheld increasing proportions of genetic interactions, with enrichment within yeast protein complexes computed at each point. The interaction types considered were negative and positive genetic, and synthetic growth defect and lethality. We note that according to BioGRID, negative and positive genetic interactions are reserved for high-throughput assays with growth scores. As shown in Figure 6, when withholding genes with a genetic interaction degree (the number of interacting partners of a certain gene) of >5, corresponding to withholding >90% of synthetic growth defect and >80% of synthetic lethality pairs, then an immediate drop-off in enrichment resulted. No such behavior was observed for negative and positive genetic interactions, for which enrichment linearly decreased as a function of the withheld proportion. Similarly, when removing interacting pairs at random, there was a steady decrease in the number of significantly enriched complexes among all types. Finally, when withholding pairs under a degree cutoff, there was also no point beyond which enrichment failed to be found (Figure S3).

DISCUSSION

Our results demonstrate that various classes of genetic interactions in different organisms can be successfully predicted based on the hypothesis that functional gene clusters tend to share genetic interaction partners. For *S. cerevisiae* in particular, predictability was obtained whether the genetic interaction type was based on growth effects or nongrowth phenotype-based measurements (*i.e.*, phenotypic suppression). Interestingly, our method did not yield predictability for negative and positive genetic interactions, which happen to be the interaction types for which most of the pairs have been tested (Costanzo *et al.* 2010). As BioGRID reserves negative and positive genetic interactions for high-throughput assays with growth scores, a threshold is applied to determine the presence of an interaction. The choice of this threshold may affect the predictability of the algorithm we have evaluated here.

In many of the genetic interaction datasets analyzed in our study, a subset of the organism's genome was selected for screening as candidate genetic interaction partners of a chosen gene. In contrast, the yeast negative and positive genetic interactions and the human CRISPR dataset we analyzed are unbiased in that nearly the entire genome was screened, and therefore no decision was taken to include or exclude genes as candidate genetic interaction partners. Using the unbiased data, our method found predictability of genetic interactions, albeit to a lesser degree than in other datasets. We note that this result does not confirm or rule out potential bias in analyzing data from other organisms, given that the genetic interactions are sparser and hence have the potential for ascertainment bias elevating the prediction scores. Moreover, our method would likely benefit from increased coverage of the functional gene networks.

While the range of predictable genetic interaction classes for human and fly were limited to phenotypic enhancement and suppression, we believe that this is probably due to the sparsity of known genetic interactions for these organisms. In this study, the source of known genetic interactions, BioGRID, had over 150,000 yeast gene pairs but only ~2800 pairs for fly and ~1500 for human. As shown in Table 1, many types of genetic interactions could simply not be tested for fly and human.

This sparseness of experimentally determined genetic interactions, especially in human, led to the lack of enrichment in gene modules such
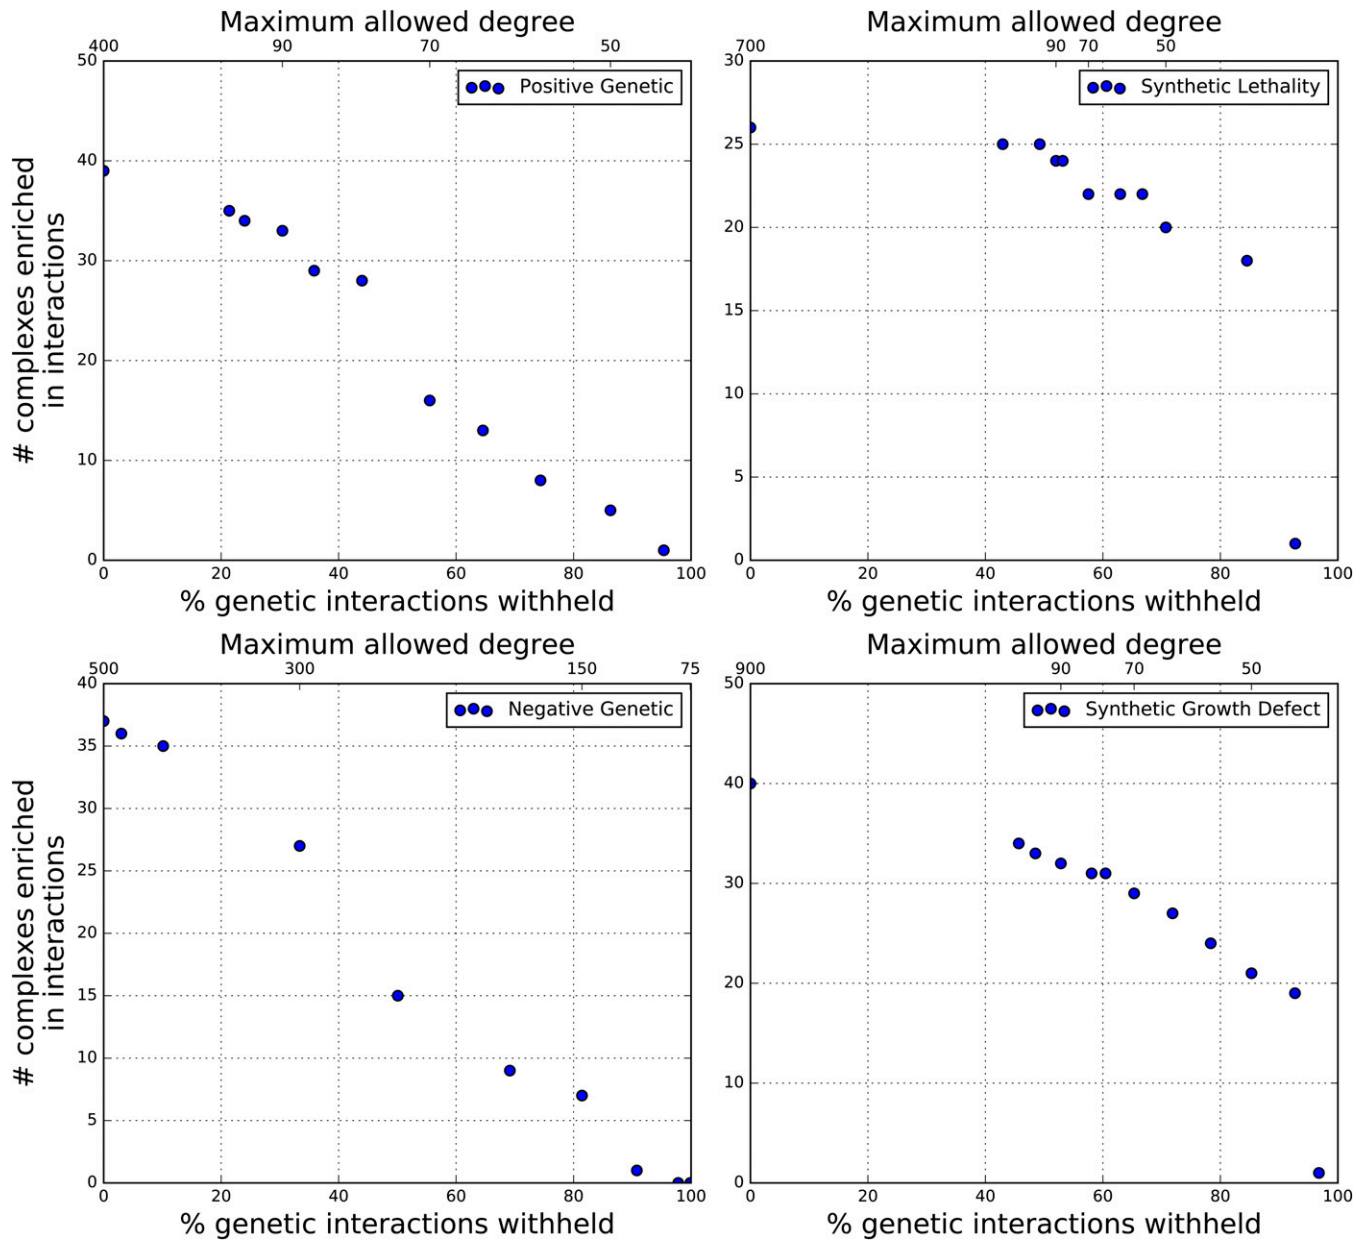
Figure 6 Predictability of genetic interactions can be found even when known interactions are sparse. By successively withholding known yeast genetic interactions according to each gene's interaction degree (*e.g.*, number of interaction partners), enrichment and therefore predictability is still detectable when information of known interactions is minimal. This effect is especially pronounced for synthetic growth defect and lethality, provided genes possess sufficiently high interaction degree.

as protein complexes. In our simulations of withholding genetic interacting pairs, we expected that regardless of the interaction type, there would be a point after which no enrichment would be found. Thus, it was surprising that negative and positive genetic interactions exhibited a linear decrease in enrichment, regardless of how the pairs were withheld (by degree or at random). On the other hand, the enrichment signal in synthetic growth defect and lethality is sensitive to the interaction degree, as there was a steep drop-off when most of the interaction pairs were withheld. In the negative and positive genetic networks, there appears to sufficient genetic interaction density such that even when high numbers of interacting pairs are withheld, enrichment under a binomial model can still be found. By extrapolating to the human case, a modest increase in the number of screened human gene pairs is likely to dramatically

increase the ability to predict additional genetic interactions, especially for synthetic growth defect and lethality, where the genes have multiple interaction partners.

While we did not investigate which classes or properties of functional interactions are predictive of genetic interactions, it would be interesting to further explore such effects. The functional gene network corresponding to a certain organism is an integration of multiple networks from various biological evidence, such as coexpression or protein–protein interactions. It may be informative to hold out certain networks for given evidence classes to determine the resulting effect on the predictability of genetic interactions.

Similar to previous genetic interaction prediction approaches (Qi *et al.* 2008; Zhong and Sternberg 2006), our algorithm requires knowledge

of known experimentally determined genetic interactions. While other studies proceed without such requirements, the assimilation of a host of biologically annotated features are still necessary for their prediction method (Pandey *et al.* 2010; Wong *et al.* 2004). In contrast to the aforementioned studies, our methodology systematically examined more than one class of genetic interaction and was successfully applied to multiple eukaryotic organisms, thereby generalizing results from a previous study by Lee *et al.* (2010). We did not explicitly carry out a comparative evaluation of our method against other approaches, however, and cannot quantitatively comment on the relative performance of different methods. Since the detection of tightly connected sets of nodes in a network is central to our method, further avenues for exploration perhaps include investigating methods, such as graph clustering (Enright *et al.* 2002) or community detection algorithms (Fortunato 2010), though these algorithms lack built-in validation. It would also be interesting to explore using tissue-specific gene networks instead of a single integrated functional gene network for more targeted predictions (Greene *et al.* 2015).

As one major goal of any genetic interaction prediction is to at least narrow down the search space for experimentally testing genetically interacting pairs, our predictions are specifically testable experimentally, perhaps through CRISPR-Cas9 for human cells (Wong *et al.* 2016). We also contribute to available prediction methodologies for suggesting genetic interactions as candidate therapeutic targets. Ultimately, we demonstrate the power of leveraging knowledge of known genetic interactions and integrated biological information in functional gene networks to predict novel genetic interactions from single-cell to multicellular organisms.

ACKNOWLEDGMENTS

LITERATURE CITED

Ashworth, A., C. J. Lord, and J. S. Reis-Filho, 2011   Genetic interactions in cancer progression and treatment. Cell 145: 30–38.

Baryshnikova, A., M. Costanzo, C. L. Myers, B. Andrews, and C. Boone, 2013   Genetic interaction networks: toward an understanding of heritability. Annu. Rev. Genomics Hum. Genet. 14: 111–133.

Benjamini, Y., and Y. Hochberg, 1995   Controlling the false discovery rate: a practical and powerful approach to multiple testing. J. R. Stat. Soc. B. 57: 289–300.

Costanzo, M., A. Baryshnikova, J. Bellay, Y. Kim, E. D. Spear *et al.*, 2010   The genetic landscape of a cell. Science 327: 425–431.

Enright, A. J., S. Van Dongen, and C. A. Ouzounis, 2002   An efficient algorithm for large-scale detection of protein families. Nucleic Acids Res. 30: 1575–1584.

Fece de la Cruz, F., B. V. Gapp, and S. M. Nijman, 2015   Synthetic lethal vulnerabilities of cancer. Annu. Rev. Pharmacol. Toxicol. 55: 513–531.

Fong, P. C., D. S. Boss, T. A. Yap, A. Tutt, P. Wu *et al.*, 2009   Inhibition of poly (ADP-ribose) polymerase in tumors from BRCA mutation carriers. N. Engl. J. Med. 361: 123–134.

Fortunato, S., 2010   Community detection in graphs. Phys. Rep. 486: 75–174.

Greene, C. S., A. Krishnan, A. K. Wong, E. Ricciotti, R. A. Zelaya *et al.*, 2015   Understanding multicellular function and disease with human tissue-specific networks. Nat. Genet. 47: 569–576.

Hart, G. T., I. Lee, and E. M. Marcotte, 2007   A high-accuracy consensus map of yeast protein complexes reveals modular nature of gene essentiality. BMC Bioinformatics 8: 236.

Hunter, J. D., 2007   Matplotlib: a 2D graphics environment. Comput. Sci. Eng. 9: 90–95.

Lee, I., Z. Li, and E. M. Marcotte, 2007   An improved, bias-reduced probabilistic functional gene network of baker's yeast, *Saccharomyces cerevisiae.* PLoS One 2: e988.

Lee, I., B. Lehner, T. Vavouri, J. Shin, A. G. Fraser *et al.*, 2010   Predicting genetic modifier loci using functional gene networks. Genome Res. 20: 1143–1153.

Lee, I., U. M. Blom, P. I. Wang, J. E. Shim, and E. M. Marcotte, 2011   Prioritizing candidate disease genes by network-based boosting of genome-wide association data. Genome Res. 21: 1109–1121.

Mani, R., R. P. S. Onge, J. L. Hartman, G. Giaever, and F. P. Roth, 2008   Defining genetic interaction. Proc. Natl. Acad. Sci. USA 105: 3461–3466.

Pandey, G., B. Zhang, A. N. Chang, C. L. Myers, J. Zhu *et al.*, 2010   An integrative multi-network and multi-classifier approach to predict genetic interactions. PLOS Comput. Biol. 6: e1000928.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion *et al.*, 2011   Scikit-learn: machine learning in Python. J. Mach. Learn. Res. 12: 2825–2830.

Qi, Y., Y. Suhail, Y. -Y. Lin, J. D. Boeke, and J. S. Bader, 2008   Finding friends and enemies in an enemies-only network: a graph diffusion kernel for predicting novel genetic interactions and co-complex membership from yeast genetic interactions. Genome Res. 18: 1991–2004.

Ruepp, A., B. Waegele, M. Lechner, B. Brauner, I. Dunger-Kaltenbach *et al.*, 2010   CORUM: the comprehensive resource of mammalian protein complexes–2009. Nucleic Acids Res. 38: D497–D501.

Shannon, P., A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang *et al.*, 2003   Cytoscape: a software environment for integrated models of biomolecular interaction networks. Genome Res. 13: 2498–2504.

Shin, J., S. Yang, E. Kim, C. Y. Kim, H. Shim *et al.*, 2015   FlyNet: a versatile network prioritization server for the *Drosophila* community. Nucleic Acids Res. 43: W91–W97.

Stark, C., B.-J. Breitkreutz, T. Reguly, L. Boucher, A. Breitkreutz *et al.*, 2006   BioGRID: a general repository for interaction datasets. Nucleic Acids Res. 34: D535–D539.

Wang, T., J. J. Wei, D. M. Sabatini, and E. S. Lander, 2014   Genetic screens in human cells using the CRISPR-Cas9 system. Science 343: 80–84.

Wong, A. S., G. C. Choi, C. H. Cui, G. Pregernig, P. Milani *et al.*, 2016   Multiplexed barcoded CRISPR-Cas9 screening enabled by CombiGEM. Proc. Natl. Acad. Sci. USA 113: 2544–2549.

Wong, S. L., L. V. Zhang, A. H. Tong, Z. Li, D. S. Goldberg *et al.*, 2004   Combining biological networks to predict genetic interactions. Proc. Natl. Acad. Sci. USA 101: 15682–15687.

Wu, C., I. MacLeod, and A. I. Su, 2012   BioGPS and MyGene.info: organizing online, gene-centric information. Nucleic Acids Res. 41: D561–D565.

Zhong, W., and P. W. Sternberg, 2006   Genome-wide prediction of *C. elegans* genetic interactions. Science 311: 1481–1484.

Communicating editor: M. Walhout