Bioinformatics
doi.10.1093/bioinformatics/xxxxxx
Advance Access Publication Date: Day Month Year
Manuscript Category

OXFORD

Subject Section

Phenotypic, functional and taxonomic features predict host-pathogen interactions

Wang Liu-Wei, Şenay Kafkas and Robert Hoehndorf*

Computer, Electrical and Mathematical Sciences and Engineering Division, Computational Bioscience Research Center, King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia.

*To whom correspondence should be addressed.

Associate Editor: XXXXXXX

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Abstract

Motivation: Identification of host-pathogen interactions (HPIs) can reveal mechanistic insights of infectious diseases for potential treatments and drug discoveries. Current computational methods for the prediction of HPIs often rely on our knowledge on the sequences and functions of pathogen proteins, which is limited for many species, especially for emerging pathogens. Matching the phenotypes elicited by pathogens with phenotypes associated with host proteins might improve the prediction of HPIs.
Results: We developed an ontology-based machine learning method that predicts potential interaction protein partners for pathogens. Our method exploits information about disease mechanisms through features learned from phenotypic, functional and taxonomic knowledge about pathogens and human proteins. Additionally, by embedding the phenotypic information of the pathogens within a formal representation of pathogen taxonomy, we demonstrate that our model can accurately predict interaction partners for pathogens without known phenotypes, using a combination of their taxonomic relationships with other pathogens and information from ontologies as background knowledge. Our results show that the integration of phenotypic, functional and taxonomic knowledge not only improves the prediction of HPIs, but also enables us to investigate novel pathogens in emerging infectious diseases.
Availability: `https://github.com/bio-ontology-research-group/hpi-predict`
Contact: robert.hoehndorf@kaust.edu.sa

1 Introduction

Infectious diseases from bacteria, viruses and fungi are one of the major causes of deaths around the globe, especially in low-income regions (World Health Organization, 2018). Pathogens disrupt host cell functions (Finlay and Cossart, 1997) and target immune pathways (Dyer *et al.*, 2010) through complex interactions among proteins (Dyer *et al.*, 2008), RNA (Fajardo *et al.*, 2015) and DNA (Weitzman *et al.*, 2004). The study of host-pathogen interactions (HPI) can provide insights into molecular mechanisms underlying infectious diseases and enable the development of novel therapeutic discoveries. For example, a previous study of many HPIs showed that pathogens typically interact with the protein hubs (those with many interaction partners) and bottlenecks (those of central location to important pathways) in human protein-protein interaction (PPI) networks (Dyer *et al.*, 2008). However, due to high costs and time constraints, experimentally validated pairs of interacting host-pathogen proteins are limited in number. Moreover, there exists a time delay for a validated HPI to be included in a database of HPIs, often requiring text-mining efforts (Thieu *et al.*, 2012). Therefore, the computational prediction of HPIs is useful in suggesting candidate interaction partners out of all the human proteins.

Existing HPI prediction methods typically utilizes features of interacting protein pairs, such as PPI network topology, structural and sequential homology, or functional profiling such as Gene Ontology similarity and KEGG pathway analysis (Nourani *et al.*, 2015). While such protein-specific features exist for some extensively studied pathogens, such as *Mycobacterium tuberculosis* (Huo *et al.*, 2015), human immunodeficiency virus (Mukhopadhyay *et al.*, 2014), *Salmonella* and *Yersinia* (Kshirsagar *et al.*, 2012), for most of the other pathogens, especially the newly emerging ones, these features are scarce (or non-existent) and expensive to obtain. As new virus species are discovered each year with potentially many more to come (Woolhouse *et al.*, 2012), a method is needed to aggregate our existing phenotypic, functional and

© The Author 2015.

taxonomic knowledge about HPIs and predict for pathogens that are less well studied or novel.

The phenotypes elicited by pathogens, i.e., the signs and symptoms observed in a patient, may provide information about molecular mechanisms (Gkoutos *et al.*, 2018). The information that phenotypes provide about molecular mechanisms is commonly exploited in computational studies of Mendelian disease mechanisms (Oellrich *et al.*, 2016; Schofield *et al.*, 2012), for example to suggest candidate genes (Hoehndorf *et al.*, 2011; Meehan *et al.*, 2017) or diagnose patients (Köhler *et al.*, 2009), but can also be used to identify drug targets (Hoehndorf *et al.*, 2013a) or gene functions (Hoehndorf *et al.*, 2013b). To the best of our knowledge, phenotypes and phenotype similarity have not yet been used to suggest host-pathogen interactions.

We use the hypothesis that phenotypes elicited by an infection with a pathogen are – among others – the result of molecular interactions, and that knowledge of the phenotypes in the host can be used to suggest the protein perturbations through which these phenotypes arise. While a large number of phenotypes resulting from infections are a consequence of immune system processes that are shared across a wide range of different types of pathogens, certain hallmark phenotypes – such as decreased CD4 count in infections with human immunodeficiency virus or microcephaly resulting from Zika virus infections – can be used to suggest candidate host proteins through which these phenotypes are elicited. Phenotypes shared across many pathogens could also contribute to the prediction of interacting host proteins as well as the absence of certain phenotypes.

We developed a method to predict candidate interacting proteins for pathogens by integrating phenotypic, functional and taxonomic information about pathogens and human proteins. We systematically compare phenotypes arising from infection with a pathogen to phenotypes associated with human genes; we further relate phenotypes to protein functions and compare pathogen-associated phenotypes with human protein functions. Relying on recent progress in deep learning with structured and semantic data (Smaili *et al.*, 2018), we construct a model that combines this information as well as structured information about taxonomic relations between pathogens to predict host-pathogen interactions. We demonstrate that our model can not only accurately predict the interacting protein partners for pathogens with phenotypic information, but also their relatives without phenotypes by aggregating knowledge through the taxonomy.

2 Results

2.1 Feature generation and representation learning for human proteins and pathogens

We first generated feature vectors for pathogens and human proteins based on different sets of features, and then trained an artificial neural network model (ANN) for predicting their interactions. We follow the workflow shown in Figure 1.

The first step in our method is the generation of features that represent a pathogen or a human protein. We represent human proteins through their associated phenotypes, the phenotypes associated with their mouse orthologs, and their functions. We obtained loss-of-function phenotypes of mouse genes from the Mouse Genome Informatics (MGI) database (Smith *et al.*, 2018) and phenotype annotations of human genes from Human Phenotype Ontology (HPO) database (Köhler *et al.*, 2018). We obtained protein function annotations from the Gene Ontology (GO) database (Ashburner *et al.*, 2000; The Gene Ontology Consortium, 2017). For pathogens, we use the phenotype annotations from the PathoPhenoDB (Kafkas *et al.*, 2018), a database of pathogen-phenotype associations, and taxonomic information from the NCBI Taxonomy (Sayers *et al.*, 2009).

We represent the phenotypes and functions through ontologies. To associate human and mouse phenotypes, we use the cross-species phenotype ontology PhenomeNET (Hoehndorf *et al.*, 2011; Rodríguez-García *et al.*, 2017), which combined the Human Phenotype Ontology (HP) (Köhler *et al.*, 2018) and the Mammalian Phenotype Ontology (MP) (Smith *et al.*, 2004). To incorporate knowledge of protein functions, we use the Gene Ontology (Ashburner *et al.*, 2000; The Gene Ontology Consortium, 2017). For the taxonomic relations between pathogens, we used the NCBI Taxonomy database (Sayers *et al.*, 2009). These ontologies contain formalized biological background knowledge (Hoehndorf *et al.*, 2015b). Using the information in ontologies as background knowledge during feature generation has the potential to significantly improve the performance of these features in machine learning and predictive analyses (Smaili *et al.*, 2018).

We use the OPA2Vec (Smaili *et al.*, 2018) method to generate features from annotations of pathogen taxa and human proteins while using the ontologies that are used to express them as background knowledge. To determine the effect of different annotations and ontologies on the performance of our models, we generated feature vectors of human proteins from human phenotype annotations, the phenotypes of their mouse orthologs, and functional annotations individually with their respective ontologies, as well as their union and intersection by combining the annotations and merging the ontologies.

2.2 Phenotypic and functional prediction of interaction partners

We obtained high-confidence interaction pairs of pathogens and human proteins from the Host-Pathogen Interaction Database (HPIDB) (Ammari *et al.*, 2016). We then trained an ANN to predict host-pathogen interactions using the feature vectors of pathogens and human proteins as input. We performed leave-one-taxon-out cross-validation (LOOCV), in which we leave out one pathogen taxon in our positive set for validation, and use all the other pathogens and their interactions as training data. We evaluated our model by ranking the true positive protein partners of the validation taxon (see Methods). We performed the training and evaluation separately based on taxonomic groups, i.e., viruses and bacteria. Table 1 provides an overview over the underlying data used for these predictions, and Figure 2 shows the prediction results.

Table 1. **Datasets used for phenotypic and functional predictions.** The features are human phenotypes (HP), mouse phenotypes (MP), gene ontology terms (GO), their union and their intersection. The coverage shows the total number of proteins annotated with the respective features. The number of unique interacting pairs, pathogen taxa and human proteins are listed for both viruses and bacteria.

Features	Coverage	Viruses			Bacteria		
		Interactions	Taxa	Proteins	Interactions	Taxa	Proteins
HP	3740	101	24	78	18	9	18
MP	9153	214	29	174	47	15	46
GO	18084	291	30	247	57	16	56
HP∪MP∪GO	18349	291	30	247	57	16	56
HP∩MP∩GO	2955	92	23	69	18	9	18

For viruses, the model using features that combine phenotypic and functional information has the best performance in prioritizing true protein partners. This model is able to prioritize 50% of the true positives within the first 10% of all proteins. For bacteria, the model that uses only functions of human proteins performs better than all other features. The TPR curve is less smooth compared to the viruses due to the limited number of exact matches of bacterial taxa between HPIDB and PathoPhenoDB.

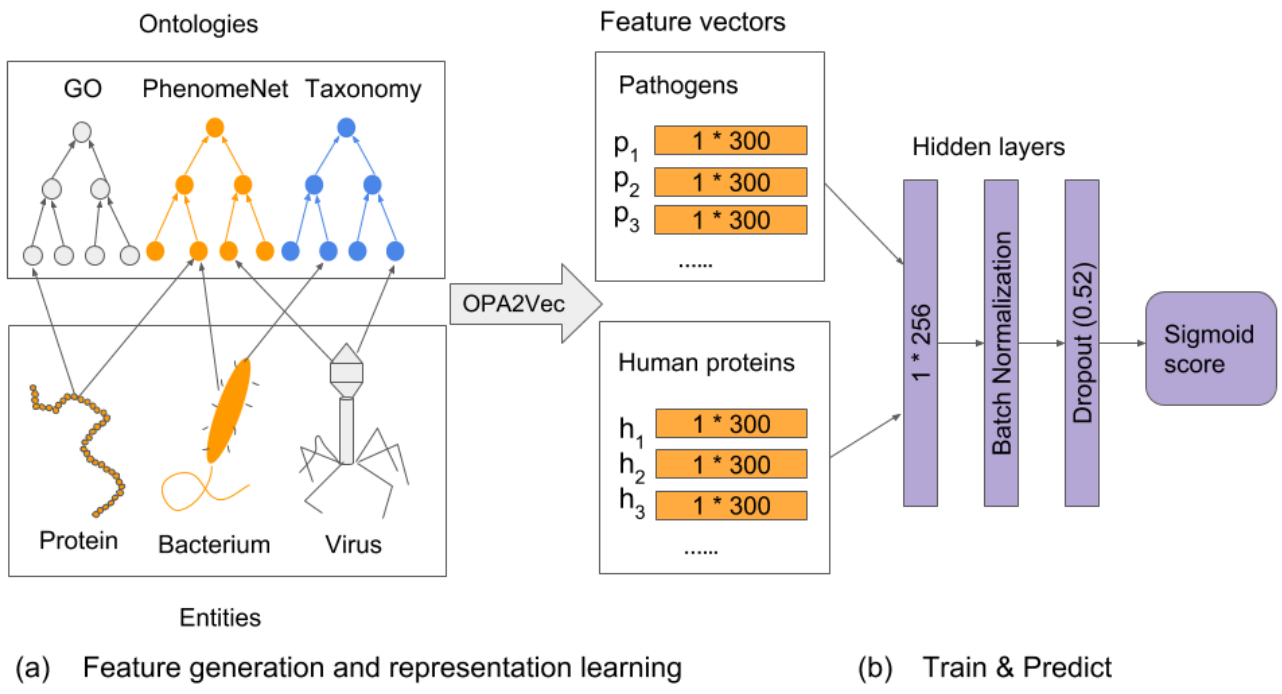

Fig. 1: An illustration of our overall workflow.

2.3 Taxonomic prediction of interaction partners

The phenotypic and functional prediction of HPIs has several drawbacks: firstly, there are many pathogen taxa in HPIDB that cannot be matched with those in PathoPhenoDB at the exact taxon level; secondly, there is considerably less training data for bacteria compared to viruses, leading to the non-smoothness of the TPR curve. To address the taxonomic inconsistencies between HPIDB and PathoPhenoDB, we generated feature vectors for all the pathogens that can be represented as a class in NCBI Taxonomy by merging the NCBI Taxonomy Ontology and the PhenomeNet Ontology during the generation of features. Our aim is not only to increase the taxonomic coverage but also to encode phenotypic and taxonomic information for pathogens and their relatives by exploiting the underlying ontological structures. Table 2 shows that almost all of the pathogens we consider have some relatives within the same family that have some phenotypic annotations in PathoPhenoDB, and over two thirds within the same species (i.e., species, subspecies, or strain).

Table 2. **Coverage of phenotypic features through taxonomy.** The percentages stand for the proportion of pathogens that have relatives within the same family, genus or species with phenotypic annotations.

Features	Viruses			Bacteria		
	Family	Genus	Species	Family	Genus	Species
HP	100%	90%	77%	97%	97%	79%
MP	100%	90%	76%	98%	98%	83%
GO & HP∪MP∪GO	100%	89%	75%	98%	98%	83%
HP∩MP∩GO	100%	90%	75%	97%	97%	79%

We retrained the models with the extended set of pathogens and evaluated their performance using LOOCV. Table 3 shows that there is more than 50% increase in the number of interactions and pathogen taxa after the inclusion of taxonomic features. Figure 3 shows the model performance after retraining with the extended set of pathogens. For both viruses and bacteria, the prediction model shows an increase in AUC for every feature of human proteins. Again, the features representing functions of gene products perform better than the phenotype features individually. However, collectively the features that combine phenotypic and functional information perform best and prioritize around 70% of the true positive protein partners within the first 10% of all proteins. Table 4 shows a summary of the performances of the different models.

Table 3. **Datasets used for taxonomic predictions** The numbers in brackets indicate the increase in data after adding taxonomic features.

Features	Viruses			Bacteria		
	Interactions	Taxa	Proteins	Interactions	Taxa	Proteins
HP	363 (+262)	98 (+74)	211 (+133)	47 (+29)	29 (+20)	34 (+16)
MP	829 (+615)	144 (+115)	510 (+336)	98 (+51)	40 (+25)	80 (+34)
GO	1125 (+834)	155 (+125)	746 (+499)	121 (+64)	42 (+26)	100 (+44)
HP∪MP∪GO	1126 (+835)	155 (+125)	747 (+500)	121 (+64)	42 (+26)	100 (+44)
HP∩MP∩GO	331 (+239)	96 (+73)	184 (+115)	47 (+29)	29 (+20)	34 (+16)

3 Discussion

We developed an ontology-based method for predicting HPIs using the phenotypic, functional and taxonomic features of pathogens and human proteins. Unlike previous methods that target protein-protein interactions and utilize features from protein structures and sequences (Nourani *et al.*, 2015), we focused on predicting pathogen-protein interactions and use features that are relatively easy and fast to obtain from patients' symptoms and taxonomic identification of pathogens.

Phenotype-based approaches have been successful in predicting disease-gene associations for Mendelian diseases (Hoehndorf *et al.*, 2011)

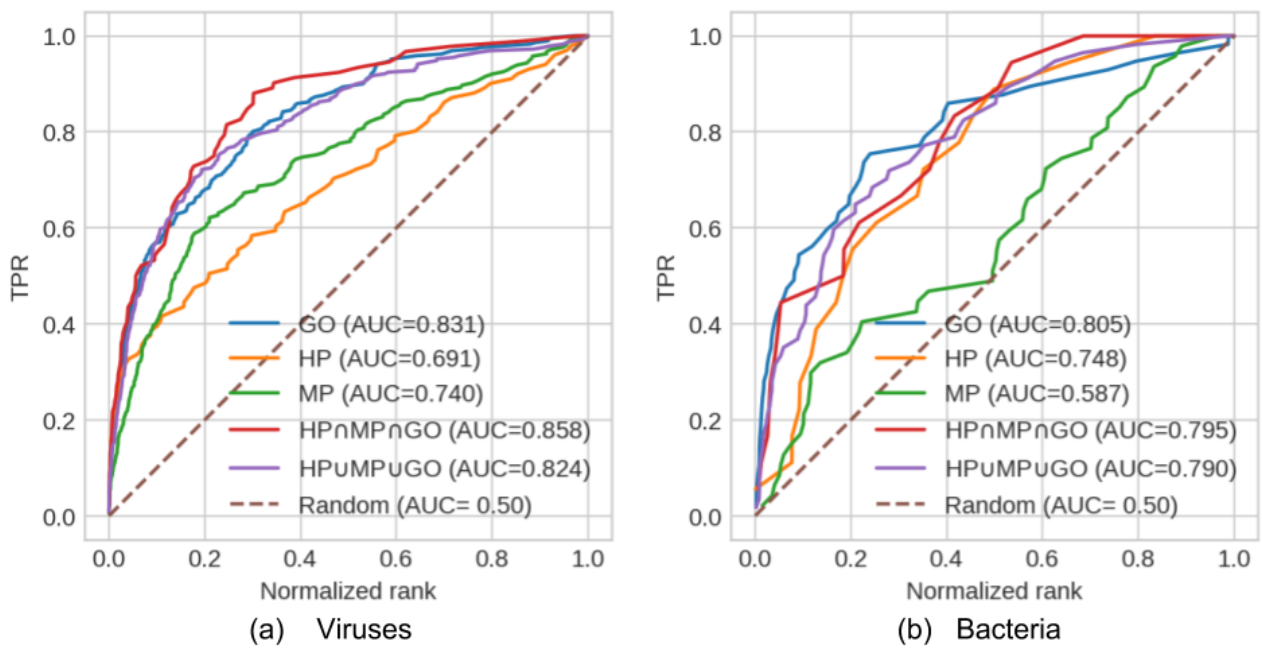

Fig. 2: **The TPR curves by normalized rank of proteins with different phenotypic and functional features** The solid curves represent the different features used to generate the feature vectors of human proteins. The dashed line represents a random classifier.

Table 4. **A summary of the model performances.** The bold AUC values represent the best features of human proteins for each prediction model. The underlined AUC values shows the best performing features of pathogen taxa after the addition of the taxonomic knowledge.

Features	Phenotypic & functional		Taxonomic	
	Virus	Bacteria	Virus	Bacteria
HP	0.691	0.748	0.776	0.793
MP	0.740	0.587	0.764	0.700
GO	0.831	**0.805**	0.841	0.870
HP∪MP∪GO	0.824	0.790	0.842	0.871
HP∩MP∩GO	**0.858**	0.795	**0.863**	**0.899**

but have not yet been used for prediction of host-pathogen interactions for infectious diseases. Our results demonstrate that the phenotypes associated with human proteins and protein functions (i.e., GO classes), can generate accurate predictions. While GO classes have been shown to perform well in HPI predictions previously (Mei, 2013; Mei and Zhu, 2014), they still rely on sequence homology of pathogen proteins. We only target interactions between pathogen taxa and host proteins instead of direct protein-protein interactions, and therefore avoid the common bottleneck of identifying the functions of pathogen proteins. Instead, we introduce a novel and – in the context of infectious diseases – rarely explored type of feature, the phenotypes elicited by pathogens in their hosts, as a "proxy" for the molecular mechanisms that result in the phenotypes.

The focus of our method on utilizing features generated based on endo-phenotypes observed in humans and mice (Schofield *et al.*, 2016) has the crucial advantage that we can identify host-pathogen interactions that may contribute to particular signs and symptoms. For example, our model predicts an interaction between Zika virus (`NCBITaxon:64320`) and DDX3X (`UniProt:O00571`). Infections with Zika virus have the potential to result in abnormal embryogenesis and, specifically, microcephaly (Wang *et al.*, 2017). Phenotypes associated with DDX3X in the mouse ortholog include abnormal embryogenesis, microcephaly, and abnormal neural tube closure (Chen *et al.*, 2016). DDX3X mutations in humans have been found to result in intellectual disability, specifically in females and affecting individuals in dose-dependent manner (Blok *et al.*, 2015). While DDX3X has previously been linked to the infectivity of Zika virus (Doñate-Macián *et al.*, 2018), our model further suggests a role of DDX3X in the development of the embryogenesis phenotypes resulting from Zika virus infections.

In addition to identifying the role of host proteins in the pathogenesis of the disease, we can also use our models to suggest candidate drug targets that inhibit these interactions and therefore prevent disease progression or decrease the severity of symptoms. Proteins that are known drug targets for diseases caused by each pathogen are ranked higher by our models than proteins that are not drug targets. For example, the taxonomic prediction model ranks known drug targets ($n = 163$) significantly higher than non-drug targets as potential interaction partners (Mann-Whitney U test, $p = 0.98 \times 10^{-25}$ for bacteria; $p = 0.02$ for viruses).

The lack of phenotypic annotations is often a challenge for such phenotype-based predictions, which is also the case for many pathogens in our study. We demonstrated that by integrating a formalized taxonomy in our feature generation, not only did we greatly increase the number of taxa but also improved the prediction performance. This suggests that our model was able to learn high level representations of the phenotypes caused by pathogens across their taxonomic hierarchy. This is the first attempt, to our knowledge, for a taxonomy ontology to be used in a machine learning method to transfer and distribute information across different taxonomic levels. This method can be potentially extended to many applications where data are only present for some levels of taxonomy, such as the modeling of some other ecological systems where many species interact.

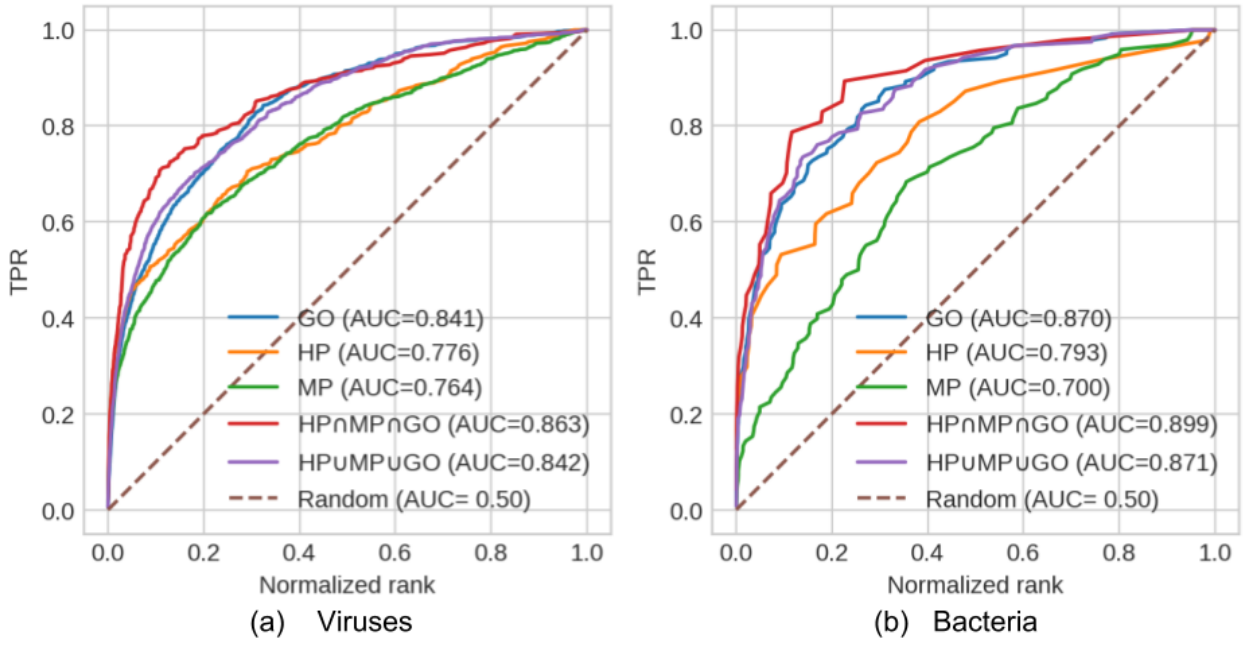

Fig. 3: **The TPR curves by normalized rank of proteins after adding taxonomic features.**

4 Materials and methods

4.1 Data sources

We downloaded interactions between hosts and pathogens from the Host Pathogen Interaction Database (HPIDB) (Ammari *et al.*, 2016) on July 16, 2018. The dataset contains 27,533 distinct interaction pairs of human proteins to pathogen species with the IntAct and Virhostnet MI scores provided. We further filtered the interactions by the MI scores with a threshold of 0.5, reducing the number of interaction pairs to 1,682. We obtained phenotypes associated with pathogens from the PathoPhenoDB (Kafkas *et al.*, 2018), a database of manually curated and text-mined associations of pathogens, diseases and phenotypes. We downloaded the version 1 of the PathoPhenoDB database (`http://patho.phenomebrowser.net/`) on Oct 11, 2018. Similarly to the implementation at the PathoPhenoDB web interface, species level taxa are also annotated with the phenotypes of all their subclasses in the taxonomy of pathogens.

We downloaded phenotypes of human genes on August 26, 2018 from the Human Phenotype Ontology (HPO) database. The dataset contains 3,777 human genes and 131,045 associations between human genes and HPO classes. We downloaded phenotypes associated with mouse genes from the Mouse Genome Informatics (MGI) database (Smith *et al.*, 2018) on September 17, 2018, using the `MGI_GenePheno.rpt` file with 12,182 mouse genes and 165,892 associations of mouse genes and Mammalian Phenotype ontology (MP) classes. We then associated MP classes with human genes by mapping mouse genes to human genes based on the information about orthologous genes provided in `HMD_HumanPhenotype.rpt` downloaded from MGI on July 16, 2018. We remove all "no abnormal phenotype detected" (`MP:0002169`) annotations.

We mapped all the genes to Uniprot proteins using the Uniprot Retrieve/ID mapping tool (`https://www.uniprot.org/uploadlists`), and used only reviewed Swissprot mappings. We downloaded Gene Ontology (GO) annotations of human proteins from the Gene Ontology website (Ashburner *et al.*, 2000; The Gene Ontology Consortium, 2017) on September 27, 2018. To identify broad taxonomic groups, i.e., bacteria and viruses, we used the NCBI Taxonomy Categories file (`ftp://ftp.ncbi.nih.gov/pub/taxonomy`) downloaded on September 3, 2018.

To obtain formal representations of phenotypes and GO classes, we downloaded the cross-species PhenomeNET Ontology (Hoehndorf *et al.*, 2011; Rodríguez-García *et al.*, 2017) from the AberOWL ontology repository (Hoehndorf *et al.*, 2015a) on September 13, 2018, and we downloaded the Gene Ontology (Ashburner *et al.*, 2000; The Gene Ontology Consortium, 2017) from the Gene Ontology website on September 27, 2018. For taxonomy features, we downloaded the NCBI Taxonomy in OWL format from the OBO Foundry website (`http://www.obofoundry.org/ontology/ncbitaxon.html`) on October 8th, 2018.

4.2 Generation of feature embeddings

For feature generation, we used OPA2Vec (Smaili *et al.*, 2018), a tool that generates feature vectors from entities and their annotation from ontologies. OPA2vec first generates a corpus consisting of relevant annotations of pathogens and proteins (i.e., HP, MP, GO, and NCBI Taxonomy classes), the complete set of class axioms and the metadata of the ontologies.

For OPA2Vec parameters, we used the default skipgram model with a window size of 5 and a minimum count parameter of 2. We tested different embedding sizes (i.e., 50, 100, 150, 200 and 300) and fixed the embedding size to 300 throughout our experiments in this paper.

4.3 Supervised prediction models

We trained an ANN as a binary classifier that predicts a sigmoid score of an interaction between a host protein and a pathogen taxon, and we used

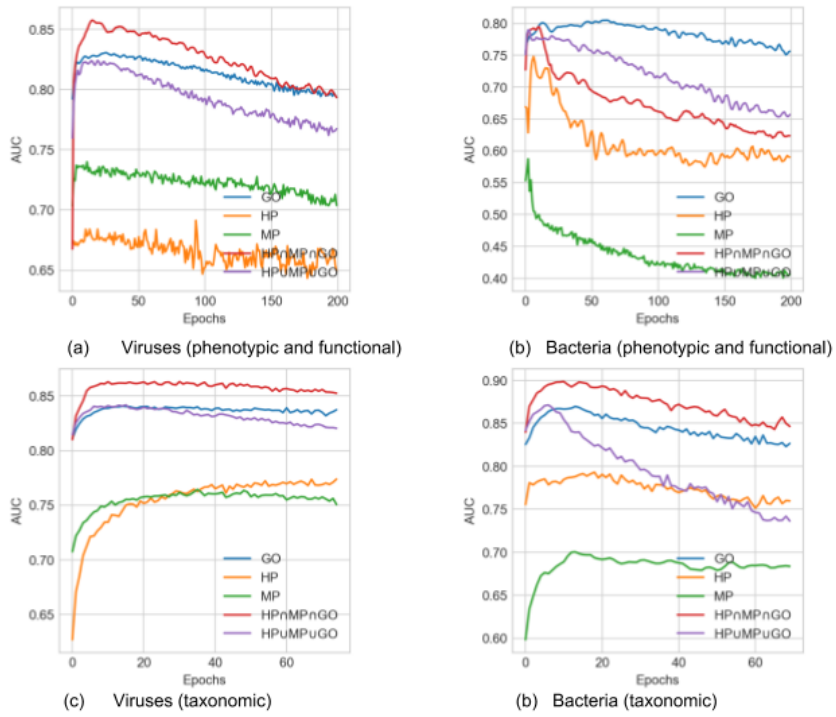

Fig. 4: **The AUC to epoch relationship**

leave-one-taxon-out cross-validation to report performance results. During each taxon-fold, we use the trained ANN to predict possible interacting proteins for the validation taxon. We evaluate the prediction results by ranking the sigmoid scores of its true protein partners among those of all the proteins. By aggregating the ranks of these true positive at each taxon fold, we plotted a true positive rate (TPR) curve by the normalized ranks of proteins and calculated an Area Under Curve (AUC) metric by trapezoidal rule for integration. The AUC metric shows our model's ability to prioritize the true protein partners under the assumption that a perfectly random model will have an AUC of 0.5. Due to the large number of negatives in relation to the positives, normalized ranks approximate the true negative rate (TNR) and the AUC measure approximates the area under the receiver operating characteristic (ROC) curve (Fawcett, 2006).

We obtained positives by extracting interaction pairs of human proteins and pathogen taxa from HPIDB. To generate the negatives, we first generated all possible combinations between all the proteins with a feature vector and pathogens that are in the positives, and then randomly sampled within the ones that did not occur in the positives. We treat all "unknown" interactions as negatives.

We searched for an optimal set of hyperparameters using the Hyperas library (`https://github.com/maxpumperla/hyperas`) based on the TPE algorithm (Bergstra *et al.*, 2011) for one of our models (i.e., the model using only the HP features). After searching, we fixed our parameters to the following throughout all the experiments: one hidden layer of 256 units with Rectified Linear Unit (ReLU) (Nair and Hinton, 2010) as activation function; one batch normalization layer; one dropout layer with a dropout rate of 0.52; one output layer with a sigmoid activation function; the Adam optimizer with cross entropy as loss function. We provide the complete list of hyperparameter search space in Table 5 and the AUC by epoch relationship is plotted in Figure 4. The preprocessing and training pipeline is provided as Jupyter Notebooks, which are available at `https://github.com/bio-ontology-research-group/hpi-predict`.

Table 5. **Hyperparameter search space for Hyperas Optimizer.**

Dimension	First layer	Second layer	Third layer	Activation	Dropout	Optimizer
50	64, 32, 16, 8, 4	16, 8, 4	-			
100	128, 64, 32, 16, 8	32, 16, 8	8, 4, 2			
150	256, 128, 64, 32, 16	32, 16, 8, 4	8, 4, 2	Relu, Sigmoid	(0, 1)	Rmsprop, Adam
200	256, 128, 64, 32, 16	32, 16, 8, 4	8, 4, 2			
300	512, 256, 128, 64, 32	64, 32, 16, 8	8, 4, 2			

4.4 Drug target analysis

We downloaded the drug target information from DrugBank (Wishart *et al.*, 2018) and obtained pathogen-disease and disease-drug associations from PathoPhenoDB (Kafkas *et al.*, 2018). We ranked the prediction scores of the drug target proteins among all proteins for each pathogen using the taxonomic prediction model and the union of HP, MP and GO features for human proteins, and performed a Mann-Whitney U test using the SciPy library version 1.0.1 (Jones *et al.*, 2001).

Acknowledgements

We would like to thank Maxat Kulmanov and Mona Alshahrani for their advice on using machine learning libraries.

Funding

This work was supported by funding from King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No URF/1/3454-01-01.

References

Ammari, M. G. *et al.* (2016). Hpidb 2.0: a curated database for host-pathogen interactions. *Database*, **2016**, baw103.

Ashburner, M. *et al.* (2000). Gene ontology: tool for the unification of biology. *Nature genetics*, **25**(1), 25.

Bergstra, J. S. *et al.* (2011). Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems*, pages 2546–2554.

Blok, L. S. *et al.* (2015). Mutations in ddx3x are a common cause of unexplained intellectual disability with gender-specific effects on wnt signaling. *The American Journal of Human Genetics*, **97**(2), 343 – 352.

Chen, C.-Y. *et al.* (2016). Targeted inactivation of murine ddx3x: essential roles of ddx3x in placentation and embryogenesis. *Human Molecular Genetics*, **25**(14), 2905–2922.

Doñate-Macián, P. *et al.* (2018). The trpv4 channel links calcium influx to ddx3x activity and viral infectivity. *Nature Communications*, **9**, 2307.

Dyer, M. D. *et al.* (2008). The landscape of human proteins interacting with viruses and other pathogens. *PLOS Pathogens*, **4**(2), 1–14.

Dyer, M. D. *et al.* (2010). The human-bacterial pathogen protein interaction networks of bacillus anthracis, francisella tularensis, and yersinia pestis. *PLOS ONE*, **5**(8), 1–12.

Fajardo, Jr., T. *et al.* (2015). Disruption of specific rna-rna interactions in a double-stranded rna virus inhibits genome packaging and virus infectivity. *PLOS Pathogens*, **11**(12), 1–22.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recogn Lett*, **27**(8), 861 – 874.

Finlay, B. B. and Cossart, P. (1997). Exploitation of mammalian host cell functions by bacterial pathogens. *Science*, **276**(5313), 718–725.

Gkoutos, G. V. *et al.* (2018). The anatomy of phenotype ontologies: principles, properties and applications. *Briefings in Bioinformatics*, **19**(5), 1008–1021.

Hoehndorf, R. *et al.* (2011). Phenomenet: a whole-phenome approach to disease gene discovery. *Nucleic acids research*, **39**(18), e119–e119.

Hoehndorf, R. *et al.* (2013a). Mouse model phenotypes provide information about human drug targets. *Bioinformatics*.

Hoehndorf, R. *et al.* (2013b). Systematic analysis of experimental phenotype data reveals gene functions. *PLoS ONE*, **8**(4), e60847.

Hoehndorf, R. *et al.* (2015a). Aber-owl: a framework for ontology-based data access in biology. *BMC bioinformatics*, **16**(1), 26.

Hoehndorf, R. *et al.* (2015b). The role of ontologies in biological and biomedical research: a functional perspective. *Briefings in Bioinformatics*.

Huo, T. *et al.* (2015). Prediction of host - pathogen protein interactions between mycobacterium tuberculosis and homo sapiens using sequence motifs. *BMC Bioinformatics*, **16**(1), 100.

Jones, E. *et al.* (2001). SciPy: Open source scientific tools for Python.

Kafkas, S. *et al.* (2018). Pathophenodb: linking human pathogens to their disease phenotypes in support of infectious disease research. *bioRxiv*.

Köhler, S. *et al.* (2018). Expansion of the human phenotype ontology (hpo) knowledge base and resources. *Nucleic Acids Research*, page gky1105.

Kshirsagar, M. *et al.* (2012). Techniques to cope with missing data in host–pathogen protein interaction prediction. *Bioinformatics*, **28**(18), i466–i472.

Köhler, S. *et al.* (2009). Clinical diagnostics in human genetics with semantic similarity searches in ontologies. *The American Journal of Human Genetics*, **85**(4), 457 – 464.

Meehan, T. F. *et al.* (2017). Disease model discovery from 3,328 gene knockouts by the international mouse phenotyping consortium. *Nature genetics*, **49**(8), 1231–1238.

Mei, S. (2013). Probability weighted ensemble transfer learning for predicting interactions between hiv-1 and human proteins. *PLOS ONE*, **8**(11), 1–14.

Mei, S. and Zhu, H. (2014). Computational reconstruction of proteome-wide protein interaction networks between htlv retroviruses and homo sapiens. *BMC Bioinformatics*, **15**(1), 245.

Mukhopadhyay, A. *et al.* (2014). Incorporating the type and direction information in predicting novel regulatory interactions between hiv-1 and human proteins using a biclustering approach. *BMC Bioinformatics*, **15**(1), 26.

Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.

Nourani, E. *et al.* (2015). Computational approaches for prediction of pathogen-host protein-protein interactions. *Frontiers in Microbiology*, **6**, 94.

Oellrich, A. *et al.* (2016). The digital revolution in phenotyping. *Briefings in Bioinformatics*, **17**(5), 819–830.

Rodríguez-García, M. Á. *et al.* (2017). Integrating phenotype ontologies with phenomenet. *Journal of biomedical semantics*, **8**(1), 58.

Sayers, E. W. *et al.* (2009). Database resources of the national center for biotechnology information. *Nucleic Acids Research*, **37**(suppl_1), D5–D15.

Schofield, P. N. *et al.* (2012). Mouse genetic and phenotypic resources for human genetics. *Human Mutation*.

Schofield, P. N. *et al.* (2016). 25 - the informatics of developmental phenotypes. In R. B. B. R. D. Morriss-Kay, editor, *Kaufman's Atlas of Mouse Development Supplement*, pages 307 – 318. Academic Press, Boston.

Smaili, F. Z. *et al.* (2018). Opa2vec: combining formal and informal content of biomedical ontologies to improve similarity-based prediction. *Bioinformatics*, page bty933.

Smith, C. L. *et al.* (2004). The mammalian phenotype ontology as a tool for annotating, analyzing and comparing phenotypic information. *Genome Biology*, **6**(1), R7.

Smith, C. L. *et al.* (2018). Mouse genome database (mgd)-2018: knowledgebase for the laboratory mouse. *Nucleic Acids Research*, **46**(D1), D836–D842.

The Gene Ontology Consortium (2017). Expansion of the gene ontology knowledgebase and resources. *Nucleic Acids Research*, **45**(D1), D331–D338.

Thieu, T. *et al.* (2012). Literature mining of host–pathogen interactions: comparing feature-based supervised learning and language-based approaches. *Bioinformatics*, **28**(6), 867–875.

Wang, A. *et al.* (2017). Zika virus genome biology and molecular pathogenesis. *Emerging Microbes & Infections*, **6**(3), e13.

Weitzman, M. D. *et al.* (2004). Interactions of viruses with the cellular dna repair machinery. *DNA Repair*, **3**(8), 1165 – 1173. BRIDGE OVER BROKEN ENDS - The Cellular Response to DNA Breaks in Health and Disease.

Wishart, D. S. *et al.* (2018). Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic Acids Research*, **46**(D1), D1074–D1082.

Woolhouse, M. *et al.* (2012). Human viruses: discovery and emergence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **367**(1604), 2864.

World Health Organization (2018). The top 10 causes of death. Available at `www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death` [Accessed 15 November 2018].