

Multitask Matrix Completion for Learning Protein Interactions Across Diseases

MEGHANA KSHIRSAGAR,¹ KEERTHIRAM MURUGESAN,²
JAIME G. CARBONELL,² and JUDITH KLEIN-SEETHARAMAN³

ABSTRACT

Disease-causing pathogens such as viruses introduce their proteins into the host cells in which they interact with the host's proteins, enabling the virus to replicate inside the host. These interactions between pathogen and host proteins are key to understanding infectious diseases. Often multiple diseases involve phylogenetically related or biologically similar pathogens. Here we present a multitask learning method to jointly model interactions between human proteins and three different but related viruses: *Hepatitis C*, *Ebola* virus, and *Influenza A*. Our multitask matrix completion-based model uses a shared low-rank structure in addition to a task-specific sparse structure to incorporate the various interactions. We obtain between 7 and 39 percentage points improvement in predictive performance over prior state-of-the-art models. We show how our model's parameters can be interpreted to reveal both general and specific interaction-relevant characteristics of the viruses. Our code is available online.*

Keywords: host–pathogen protein–protein interaction, matrix completion, multitask learning, protein interaction prediction, viruses.

1. INTRODUCTION

INFECTIONOUS DISEASES such as H1N1 influenza, the recent Ebola outbreak, and bacterial infections such as the recurrent *Salmonella* and *Escherichia coli* outbreaks are a major health concern worldwide, causing millions of illnesses and many deaths each year. Key to the infection process are host–pathogen interactions at the molecular level, in which pathogen proteins physically bind to human proteins to manipulate important biological processes in the host cell, to evade the host's immune response, and to multiply within the host. Very little is known about these protein–protein interactions (PPIs) between pathogen and host proteins for any individual disease. However, such PPI data are widely available across several diseases, and the central question in this article is: *Can we model host–pathogen PPIs better by leveraging data across multiple diseases?* This is of particular interest for lesser known or recently evolved diseases, in which the data are particularly scarce. Furthermore, it allows us to learn models that generalize better across diseases by modeling global phenomena related to infection.

¹Memorial Sloan Kettering Cancer Center, New York, New York.

²Language Technologies Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania.

³Metabolic & Vascular Health, Warwick Medical School, University of Warwick, Coventry, United Kingdom.

*http://github.com/meghana-kshirsagar/bsl_mtl

An elegant way to formulate the interaction prediction problem is through a graph completion-based framework, in which we have several bipartite graphs over multiple hosts and pathogens. Nodes in the graphs represent host proteins (circles) and pathogen proteins (triangles), with edges between them representing interactions (host protein *interacts* pathogen protein). Given some observed edges (interactions obtained from laboratory-based experiments), we wish to predict the other edges in the graphs. Such bipartite graphs arise in a plethora of problems, including recommendation systems (user *prefers* movie), citation networks (author *cites* article), and disease–gene networks (gene *influences* disease). In our problem, each bipartite graph \mathcal{G} can be represented using a matrix M , in which the rows correspond to pathogen proteins and columns correspond to host proteins. The matrix entry M_{ij} encodes the edge between pathogen protein i and host protein j from the graph, with $M_{ij} = 1$ for the observed interactions. Thus, the graph completion problem can be mathematically modeled as a matrix completion problem (Candes and Recht, 2009).

Most of the prior work on host–pathogen PPI prediction has modeled each bipartite graph separately, and hence cannot exploit the similarities in the edges across the various graphs. Here we present a *multitask* matrix completion method that *jointly models* several bipartite graphs by sharing information across them. From the multitask perspective, a *task* is the graph between one host and one pathogen (can also be seen as interactions relevant to one disease). We focus on the setting in which we have a single host species (human) and several related viruses, in which we hope to gain from the fact that similar viruses will have similar strategies to infect and hijack biological processes in the human body. Such opportunities for sharing arise in other applications as well: for instance, predicting user preferences in movies may inform preferences in selection of books, or vice versa, as movies and books are semantically related. Multitask learning-based models that incorporate and exploit these correlations should benefit from the additional information.

Our multitask matrix completion-based model is motivated by the following biological intuition governing protein interactions across diseases.

1. An interaction depends on the structural properties of the proteins, which are conserved across similar viruses as they have evolved from common ancestors. Our model thus needs a component to capture these latent similarities, which is *shared* across tasks.
2. In addition to the shared properties already discussed, each pathogen has also evolved specialized mechanisms to target host proteins. These are unique to the pathogen and can be expressed using a *task-specific* parameter in the model.

This leads us to the following model that incorporates the mentioned ideas. The interactions matrix M_t of task t can be written as $M_t = \mu_t * (\text{shared component}) + (1 - \mu_t) * (\text{specific component})$ with hyperparameter μ_t , allowing each task to customize its amount of shared and specific components.

To incorporate the mentioned ideas, we assume that the interactions matrix M is generated from two components. The first component has low-rank latent factors over the human and virus proteins, with these latent factors jointly learned over all tasks. The second component involves a task-specific parameter, on which we additionally impose a sparsity constraint as we do not want this parameter to overfit the data. Section 3 discusses our model in detail. We trade-off the relative importance of the two components using task-specific hyperparameters. Our model can thus learn what is conserved and what is different across pathogens, rather than having to specify it manually.

The key challenges in inducing such a model are as follows. (1) In addition to the interactions from each graph, it should exploit information available in the form of features. (2) Exploiting features is particularly crucial because the graph \mathcal{G} is often extremely sparse, that is, there are a large number of nodes and very few edges are observed. There will be proteins (i.e., nodes) that are not involved in any known interactions—called the *cold start problem* in the recommendation systems community. The model should be able to predict the existence of links (or their absence) between such prior “unseen” node pairs. This is of particular significance in graphs that capture biological phenomena. For instance, the host–pathogen PPI network of human Ebola virus (column 3, Table 2) has ≈ 90 observed edges (equivalent to 0.06% of the possible edges), which involve only 2 distinct virus proteins. (3) A side effect of having scarce data is the availability of a large number of unlabeled examples, that is, pairs of nodes with no edge between them. These unlabeled examples can contain information about the graph as a whole, and a good model should be able to use them.

The main contributions of this work are as follows.

- We extend a prior matrix completion model (Abernethy et al., 2009) to the multitask setting. This extension is new.
- Unlike most prior approaches, our model exploits node-based features, which allows us to deal with the “cold start” problem (generating predictions on unseen nodes).
- We apply the model to an important, real-world problem—prediction of interactions in disease-relevant host–pathogen protein networks, for multiple related diseases. We demonstrate the superior performance of our model over prior state-of-the-art multitask models.
- We use unlabeled data to initialize the parameters of our model, which serves as a prior. This gives us a modest boost in prediction performance.

1.1. Background: Host–pathogen PPIs

The experimental discovery of host–pathogen protein–protein interactions (HP PPIs) involves biochemical and biophysical methods such as coimmunoprecipitation, yeast two-hybrid assays, and cocrystallization. The HP PPIs from several small-scale and high-throughput experiments are aggregated by databases such as MINT (Chatranyamonti et al., 2009), HPIDB (Kumar and Nanduri, 2010), and PHISTO (Tekir et al., 2012) by literature curation. These databases are valuable sources of information to bioinformaticians for developing models.

1.1.1. Prediction of host–pathogen PPIs. The most reliable experimental methods for studying PPIs are often very time consuming and expensive, making it hard to investigate the prohibitively large set of possible host–pathogen interactions—for example, the bacterium *Bacillus anthracis* that causes anthrax has about 2321 proteins, which when coupled with the 100,000 or so human proteins gives ≈ 232 million protein pairs to test experimentally. Computational techniques complement laboratory-based methods by predicting highly probable PPIs. These techniques use the known interactions data from previous experiments and predict the most plausible new interactions. In particular, supervised machine learning-based methods use the few known interactions as training data and formulate the interaction prediction problem in a classification setting, with target classes “interacting” or “noninteracting.” Features are derived using various attributes of the two proteins such as protein sequences from UniProt (UniProt Consortium, 2011), protein structure from PDB, and gene ontology (GO) from GO database (Ashburner et al., 2000).

1.2. Prior work

Most of the prior work in PPI prediction has focused on building models separately for individual organisms (Chen and Liu, 2005; Qi et al., 2006; Singh et al., 2006; Wu et al., 2006) or on building a model specific to a disease in the case of HP PPI prediction (Dyer et al., 2007; Qi et al., 2009; Tastan et al., 2009; Kshirsagar et al., 2012). There has been little work on combining PPI data sets with the goal of improving prediction performance for multiple organisms. Qi et al. (2010) proposed a semisupervised multitask framework to predict PPIs from partially labeled reference sets. Kshirsagar et al. (2013) developed a task regularization-based framework called multitask pathway-based learning (MTPL) that incorporates the similarity in biological pathways targeted by various diseases to couple multiple tasks together. Matrix factorization-based PPI prediction has seen very little work, mainly because of the extremely sparse nature of these data sets that makes it very difficult to get reliable predictors. Xu et al. (2010) used a CMF-based approach in a multitask learning setting for within-species PPI prediction. The methods used in all prior work on PPI prediction do not explicitly model the features of the proteins and cannot be applied on proteins that have no known interactions available. Our work addresses both these issues.

A majority of the prior work in the relevant areas of collaborative filtering and link prediction includes single relation models that use neighborhood-based prediction (Sarwar et al., 2001), matrix factorization-based approaches (Koren et al., 2009; Menon and Elkan, 2011), and Bayesian approaches using graphical models (Jin et al., 2002; Phung et al., 2009). There have also been multitask approaches on link prediction (Singh and Gordon, 2008; Li et al., 2009; Cao et al., 2010; Zhang et al., 2012); however, many of those models do not incorporate features of the nodes in the graph. Menon and Elkan (2011) combine linear and bilinear features, latent parameters on the nodes, and several other parameters into a function that minimizes a ranking loss for link prediction in a single graph. Abernethy et al. (2009) cast the problem of matrix

completion in terms of the abstract problem of learning linear operators. Their framework allows the incorporation of features and kernels. We extend their bilinear model for the multitask setting. There has been a lot of work on other low-rank models for multitask learning (Ando and Zhang, 2005; Ji and Ye, 2009; Chen et al., 2012, 2013). Another related line of work is coembedding, in which the goal is to embed different types of objects in the same shared lower dimensional space (Mirzazadeh et al., 2015).

2. BILINEAR LOW-RANK MATRIX DECOMPOSITION

In this section, we present the matrix decomposition model that we extend for the multitask scenario. In the context of our problem, at a high level, this model states that protein interactions can be expressed as dot products of features in a lower dimensional subspace.

Let \mathcal{G}_t be a bipartite graph connecting nodes of type v with nodes of type ζ . Let there be m_t nodes of type v and n_t nodes of type ζ . We denote by $M \in \mathbb{R}^{m_t \times n_t}$ the matrix representing the edges in \mathcal{G}_t . Let the set of observed edges be Ω . Let \mathcal{X} and \mathcal{Y} be the feature spaces for the node types v and ζ , respectively. For the sake of notational convenience, we assume that the two feature spaces have the same dimension d_t^1 . Let $\mathbf{x}_i \in \mathcal{X}$ denote the feature vector for a node i of type v and $\mathbf{y}_j \in \mathcal{Y}$ be the feature vector for node j of type ζ . The goal of the general matrix completion problem is to learn a function $f: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ that also explains the observed entries in the matrix M . We assume that the function f is bilinear on $\mathcal{X} \times \mathcal{Y}$. This bilinear form was first introduced by Abernethy et al. (2009) and takes the following form:

$$f(\mathbf{x}_i, \mathbf{y}_j) = \mathbf{x}_i^T H \mathbf{y}_j = \mathbf{x}_i^T U V^T \mathbf{y}_j. \quad (1)$$

The factor $H \in \mathbb{R}^{d_t \times d_t}$ maps the two feature spaces \mathcal{X} and \mathcal{Y} . This model assumes that H has a low-rank factorization given by $H = UV^T$, where $U \in \mathbb{R}^{d_t \times k}$ and $V \in \mathbb{R}^{d_t \times k}$. The factors U and V project the two feature spaces to a common lower dimensional subspace of dimension k . Although the dimensionality of the feature spaces \mathcal{X} and \mathcal{Y} may be very large, the latent lower dimensional subspace is sufficient to capture all the information pertinent to interactions. To predict whether two new nodes (i.e., nodes with no observed edges) with features \mathbf{p}_i and \mathbf{q}_j interact, we simply need to compute the product: $\mathbf{p}_i U V^T \mathbf{q}_j$. This enables the model to avoid the cold start problem that many prior models suffer from. The objective function to learn the parameters of this model has two main terms: (1) a data-fitting term, which imposes a penalty for deviating from the observed entries in Ω and (2) a low-rank enforcing term on the matrix H .

The first term can be any loss function such as squared error, logistic loss, and hinge loss. We tried both squared error and logistic loss and found the performance to be similar. The squared error function has the advantage of being amenable to adaptive step-size-based optimization, which results in a much faster convergence. The low-rank constraint on H (mentioned in (2)) is NP-hard to solve and it is standard practice to replace it with either the trace norm or the nuclear norm. Minimizing the trace norm (i.e., sum of singular values) of $H = UV^T$ is equivalent to minimizing $\|U\|_F^2 + \|V\|_F^2$. This choice makes the overall function easier to optimize:

$$\mathcal{L}(U, V) = \sum_{(i,j) \in \Omega} c_{ij} \ell(M_{ij}, \mathbf{x}_i^T U V^T \mathbf{y}_j) + \lambda (\|U\|_F^2 + \|V\|_F^2) \quad (2)$$

where $\ell(a, b) = (a - b)^2$

The constant c_{ij} is the weight/cost associated with the edge (i, j) , which allows us to penalize the error on individual instances independently. The parameter λ controls the trade-off between the loss term and the regularizer.

3. THE BILINEAR SPARSE LOW-RANK MULTITASK LEARNING MODEL

In the previous section, we described the bilinear low-rank model for matrix completion. Note that to capture linear functions over the features, we introduce a constant feature for every protein (i.e., $[\mathbf{x}_i 1]$). We now discuss

¹The dimensions being different do not influence the method or the optimization in any way.

the multitask extensions that we propose. Let $\{\mathcal{G}_t\}$, where $t=1 \dots T$ be the set of T bipartite graphs and the corresponding matrices be $\{M_t\}$. Each matrix M_t has rows corresponding to node type v_t and columns corresponding to the node type ζ_t . The feature vectors for individual nodes of the two types be represented by \mathbf{x}_{ti} and \mathbf{y}_{tj} , respectively. Let Ω_t be the set of observed links (and nonlinks) in the graph \mathcal{G}_t . Our goal is to learn individual link prediction functions f_t for each graph. To exploit the relatedness of the T bipartite graphs, we make some assumptions on how they share information. We assume that each matrix M_t has a low-rank decomposition that is shared across all graphs and a sparse component that is specific to the task t . That is,

$$f_t(\mathbf{x}_{ti}, \mathbf{y}_{tj}) = \mathbf{x}_{ti}^T H_t \mathbf{y}_{tj}, \text{ where } H_t = \mu_t U V^T + (1 - \mu_t) S_t. \quad (3)$$

As before, the shared factors U and V are both $\mathbb{R}^{d_t \times k}$ (where the common dimensionality d_t of the two node types is assumed for convenience). The matrix $S_t \in \mathbb{R}^{d_t \times d_t}$ is a sparse matrix. The objective function for the multitask model is given by

$$\begin{aligned} \mathcal{L}(U, V, \{S_t\}) = & \frac{1}{N} \sum_{t=1}^T \sum_{(i,j) \in \Omega_t} c_{ij}^t \ell(M_{t_{ij}}, \mathbf{x}_{ti}^T H_t \mathbf{y}_{tj}) + \\ & \lambda (\|U\|_F^2 + \|V\|_F^2) + \sum_{t=1}^T \sigma_t \|S_t\|_1. \end{aligned} \quad (4)$$

Here $N = \sum_t |\Omega_t|$ is the total number of training examples (links and nonlinks included) from all tasks. To enforce the sparsity of S_t , we apply an ℓ_1 norm. In our experiments, we tried both ℓ_1 and ℓ_2 norms and found that the ℓ_1 norm works better.

Optimization: The function $\mathcal{L}(U, V, \{S_t\})$ is nonconvex. However, it is convex in every one of the parameters (i.e., when the other parameters are fixed), and a block coordinate descent method called alternating least squares (ALSs) is commonly used to optimize such functions. To speed up convergence, we use an adaptive step size. The detailed optimization procedure is shown in the Supplementary Data.

Convergence: The ALSs algorithm is guaranteed to converge only to a local minimum. There is work showing convergence guarantees to global optima for related simpler problems; however, the assumptions on the matrix and the parameter structure are not very practical and it is difficult to verify whether they hold for our setting.

Initialization of U and V : We tried random initialization (in which we randomly set the values to lie in the range [0 to 1]), and also the following strategies that initialize $U^0 \leftarrow$ top- k left singular vectors and $V^0 \leftarrow$ top- k right singular vectors from the singular value decomposition (SVD) of $\sum_{(i,j) \in \Gamma} m_{ij} \mathbf{x}_{ti} \mathbf{y}_{tj}^T$. We set Γ

to (a) training examples from all tasks or (b) a random sample of 10,000 unlabeled data from all tasks. We found that using the unlabeled data for initialization gives us a better performance.

3.1. Handling the ‘‘curse of missing negatives’’

For the matrix completion (MC) algorithm to work in practice, the matrix entries M_{ij} should represent interaction scores (range [0 to 1]) or take binary values (1 second for positives and 0 second for negatives). Our experiments with PPI probabilities (obtained using the MINT-scoring algorithm) gave bad models. The binary matrix setting requires some observed 0 second. However, noninteractions are not available as they cannot be verified experimentally for various reasons. Hence we derived a set of ‘‘probable negatives’’ using a heuristic often used in PPI prediction work (Qi et al., 2006; 2009; Dyer et al., 2011; Kshirsagar et al., 2013). We pair up all virus proteins with all human proteins and sample a random set to be negatives. This heuristic works in practice as the interaction ratio (i.e., number of positives in a large random set of protein pairs) is expected to be very low: $\approx 1/100$ to $1/500$. That is, the probability that our negatives contain true positives is negligible.

High class imbalance: we incorporate the prior on the interaction ratio by setting the size of our randomly sampled negatives set equal to 100 times the number of gold standard positives.

4. DATA SET

We use three human virus PPI data sets from the PHISTO (Tekir et al., 2012) database (version from 2014), the characteristics of which are summarized in Table 1. The *Influenza A* task includes various strains

TABLE 1. TASKS AND THEIR SIZES (A COLUMN CORRESPONDS TO ONE BIPARTITE GRAPH BETWEEN HUMAN PROTEINS AND THE VIRUS INDICATED IN THE COLUMN HEADER)

| Task → | <i>Influenza A</i> (28 strains) | <i>Hepatitis C</i> (5 strains, many isolates) | <i>Ebola</i> (3 strains) |
|--|------------------------------------|--|-----------------------------|
| No. of HP PPIs (positives) | 848 | 981 | 90 |
| No. of nonhomologous PPIs | 274 | 270 | 75 |
| No. of unique virus proteins in data set | 54 | 151 | 2 |
| No. of unique human proteins in data set | 362 | 385 | 88 |
| Total No. of virus proteins in UniProtKB across strains ^a | 542 | 163 | 38 |
| No. of negatives | 84800 | 98100 | 9000 |
| Density of observed graph ^b (as %) | 0.15 | 0.60 | 0.20 |

More importantly, each task represents several strains of one virus species. We pool together PPI from all strains as the number of interactions per strain is very small and not representative. PPI data from different strains often involve homologous proteins (virus proteins or/and human proteins) as these were curated from different publications/studies involving the virus. The second row gives a conservative estimate of the number of nonhomologous PPIs. These were obtained by computing BLAST sequence similarity (see Section 4.1 for details). All pathogens are single-stranded RNA viruses. The last row shows that each of our graphs is extremely sparse.

^aThis is the number of “reviewed” proteins on UniProtKB from all strains in our data.

^bDensity=(No. of positives + No. of negatives)/total number of possible edges. The denominator was computed considering all proteins from all strains of the virus and $\approx 100,000$ human proteins.

HP PPIs, host–pathogen protein–protein interactions; PPI, protein–protein interaction.

of flu such as H1N1 (Puerto Rico, Texas strains, etc.), H3 N2 (England strain, etc.), H5 N1, and H7 N3. The *Hepatitis* task has several subtypes (1a and 1b), isolates such as H77 and HC-J6, and *Ebola* has the Zaire Ebola virus strain Mayinga and strain 1995. All three are single-strand RNA viruses, with *Hepatitis* being a positive-strand ssRNA, whereas *Influenza* and *Ebola* are negative-strand viruses. The density of the known interactions is quite small when considering the entire proteome (i.e., all known proteins) of the host and pathogen species (last row in Table 1).

4.1. Addressing the homologs in PHISTO

The interactions data reported in PHISTO database is aggregated from several host-pathogen interactions data sources such as MINT, IntAct, DIP, etc. We found that many of the interactions (within a task) involve homologous PPI—virus protein homologs from two strains interacting with either the same human protein or with homologous human proteins. Such PPIs usually come from different publications/studies, but they introduce a lot of redundancy in the data set. Removing such PPIs is challenging as there are several criteria to consider, the most important being *what threshold of sequence similarity constitutes a homolog?* The very existence of sequence similarity between proteins is what makes it possible for biologists to predict unknown interactions and most computational methods rely on it for PPI prediction.

In row 2 of Table 1, we give a conservative estimate of the number of nonhomologous PPIs. For this, we identify homologous PPI and retain only one of them in the data set. Homologs were obtained using BLAST sequence alignment (blastp) using an e-value cutoff threshold of $1e^{-5}$. This is a relaxed cutoff because there are cases in which two functionally different proteins (either virus–virus or human–human) with low query coverage and low identity end up as “homologs”. For example, human proteins Q8WV28 (gene BLNK, B cell linker protein) and O00459 (gene PIK3R2, phosphatidylinositol 3-kinase) are considered homologs by this cutoff. The number of nonhomologous PPIs should, therefore, be considered a lower limit.

In our experiments section, we thus evaluate our models in two settings—(1) on the original data set (2) by creating partitions: the first without any homologous PPI and the other with only the homologous PPI.

4.2. Features

As the sequence of a protein determines its structure and consequently its function, it may be possible to predict PPIs using the amino acid sequence of a protein pair. Shen et al. (2007) introduced the “conjoint triad model” for predicting PPIs using only amino acid sequences. They partitioned the 20 amino acids into

7 classes based on their electrostatic and water affinities.² A protein's amino acid sequence is first transformed to a class sequence (by replacing each amino acid by its class). For $k=3$, they count the number of times each distinct tri-mer (set of three consecutive amino acids) occurred in the sequence. As there are 343 (7^3) possible trimers (with an alphabet of size 7), the feature vector containing the tri-mer frequency counts will have 343 elements. To account for protein size, they normalized the counts by linearly transforming them to lie between 0 and 1. Thus the value of each feature in the feature vector is the normalized count for each of the possible amino acid tri-mers. We use di-, tri-, and four-mers thus leading to a total of 2793 features ($7^2 + 7^3 + 7^4$). Such features have been successfully applied in prior work (Dyer et al., 2007; Kshirsagar et al., 2013).

5. EXPERIMENTAL SETUP

5.1. Comparison with other machine learning approaches

Our baselines include recent low-rank and sparse models, conventional multitask methods, and prior work on HP PPI prediction. For a uniform comparison, we used least squared loss in all the methods. The MALSAR (Zhou et al., 2011) package was used to implement some of the models. For the baselines wherever appropriate, we concatenated the features of the two node types into a single feature vector. Let $W \in \mathbb{R}^{T \times d_t}$ be the matrix with the task-specific weight vectors \mathbf{w}_t .

Single task learning: We used ridge regression with ℓ_2 regularization (which performed better than ℓ_1).

MMTL: The mean regularized multitask learning model from Evgeniou and Pontil (2004).

Sparse + low-rank (Chen et al., 2012): W is assumed to have the decomposition: $W = P + Q$, where P is sparse and Q has a low-rank structure.

IMC (Jain and Dhillon, 2013; Natarajan and Dhillon, 2014): This is the link-prediction model from Section 2, in which data from all tasks are combined without incorporating any task relationships. U and V are shared by all tasks. We use the same initialization for this method as we do for our model. A comparison with this model tells us how much we gain from the task-specific sparsity component S_t .

MTPL (Kshirsagar et al., 2013): A biologically inspired regularizer is used to capture task similarity.

BSL-MTL: Bilinear sparse low-rank multitask learning, the method developed in this article.

5.2. Comparison with a homolog baseline

In addition to the machine learning methods presented in Section 5.1, we also tried a simple homology-based approach. As per this method, a virus protein v will interact with a human protein h if it interacts with any other human protein h' similar to h . The converse holds too. Such homology-motivated heuristic approaches have been traditionally popular in the interorganism PPI prediction literature (Matthews et al., 2001; Walhout and Vidal, 2001). Mika and Rost (2006) offer a commentary on the success of such approaches in prediction PPI within an organism and the transferability across organisms. Modifications of such approaches have been published recently (Chen et al., 2009; Lin et al., 2013; Murakami and Mizuguchi, 2014).

The virus-virus and human-human protein similarity graphs used in this baseline were computed using protein sequence similarity. We ran BLAST sequence alignment (blastp) with an e-value cutoff of 0.1, and the negative logarithm of the e-value was used as a measure of similarity (we use a somewhat relaxed cutoff as we are comparing across organisms). We use the entire human-virus interactome from PHISTO as input to this method. Given a human protein h , we rank all virus proteins v based on the following simple scheme. Let $V(h)$ be a set of all virus proteins interacting with the human protein h and let $H(v)$ be all human proteins interacting with the virus protein v . We score each pair (h, v) from the bipartite graph as follows: $\max\{\max \text{similarity between } h \text{ and all human proteins in } H(v); \max \text{similarity between } v \text{ and all virus proteins in } V(h)\}$. This baseline is shown in the first row as *Homolog* in Table 2.

5.3. Evaluation setup

We compare all the methods in two settings, in which a small proportion of the available labeled data are randomly sampled and used to train a model, which is then evaluated on the remaining data. For the first

²For details of these classes, please refer to the supplementary or the original article.

TABLE 2. AREA UNDER THE PRECISION-RECALL CURVE ON THE PHISTO DATA SET

| | 10% training, 90% test | | | 30% training, 70% test | | |
|------------------|------------------------|--------------------|--------------------|------------------------|--------------------|--------------------|
| | <i>Ebola</i> | <i>Hep-C</i> | <i>Influenza</i> | <i>Ebola</i> | <i>Hep-C</i> | <i>Influenza</i> |
| Homolog | 0.230 ±0.06 | 0.178±0.01 | 0.158±0.01 | 00.311 ±.03 | 0.180±0.01 | 0.198±0.01 |
| STL (Ridge Reg.) | 0.189±0.09 | 0.702±0.08 | 0.286±0.02 | 00.130±.03 | 0.802±0.03 | 0.428±0.03 |
| MMTL | 0.113±0.04 | 0.767±0.03 | 0.321±0.02 | 00.129±.02 | 0.802±0.04 | 0.430±0.03 |
| Sparse+low-rank | 0.144±0.07 | 0.767±0.02 | 0.318±0.02 | 00.153±.02 | 0.814 ±0.01 | 0.414±0.03 |
| MTPL | 0.217±0.08 | 0.695±0.02 | 0.345±0.02 | 00.260±.05 | 0.713±0.01 | 0.496 ±0.03 |
| IMC | 0.087±0.04 | 0.779 ±0.02 | 0.362 ±0.01 | 00.122±.02 | 0.801±0.01 | 0.410±0.03 |
| BSL-MTL | 0.233 ±0.10 | 0.807 ±0.02 | 0.486 ±0.02 | 00.361 ±.03 | 0.842 ±0.01 | 0.560 ±0.02 |

The column header “X% training” indicates the fraction of the labeled data used for training and tuning the model with the rest (100 – X)% used as test data. We report the average AUC-PR over 10 random train–test splits (stratified splits that maintain the class-skew of 1:100). The standard deviation is also shown. The performance of the best baseline and the overall best method (BSL-MTL) is highlighted in bold. The Homolog baseline performs comparably to our BSL-MTL method on the *Ebola* task. This is explained by the virus–human PPI graph, from several other viruses, that the Homolog baseline has access to.

AUC-PR, area under the precision-recall curve; BSL-MTL, bilinear sparse low-rank multitask learning; MTPL, multitask pathway-based learning; STL, single task learning.

setting, we randomly split the labeled data from each task into 10% training and 90% test, such that the class-skew of 1:100 is maintained in both splits (i.e., stratified splits). The second setting uses a 30% training, 70% test split. In each setting, we generate 10 random splits and average the performance over the 10 runs (parameter tuning details are given in the Supplementary Data).

We report the area under the precision-recall curve (AUC-PR) along with the standard deviation. AUC-PR has been shown to give a more informative picture of an algorithm’s performance than receiver operating characteristic (ROC) curves in high-class imbalance data sets (Davis and Goadrich, 2006) such as ours.

6. RESULTS

The PPI data we obtain from PHISTO have homologous PPI within a task. Note that the presence of sequence similarity across virus proteins is a strong criterion for computational methods to work; however, orthologs or paralogs of proteins within a task lead to redundancy in the data, and some readers may find it harder to judge the contribution of the methods. We, therefore, present results on three sets:

1. The original PPIs from PHISTO. AUC-PR is shown in Table 2.
2. Nonhomologous partition: after removal of all homologous PPIs from the test data. Here we ensure that the test data do not contain homologs of PPI observed in the training data (this ensures there are

TABLE 3. AREA UNDER THE PRECISION-RECALL CURVE ON THE NONHOMOLOGOUS PROTEIN–PROTEIN INTERACTION WITHIN EACH TASK, IN THE 10% SETTING

| 10% training, 90% test on nonhomologous PPI | | | |
|---|--------------------|--------------------|--------------------|
| | <i>Ebola</i> | <i>Hep-C</i> | <i>Influenza</i> |
| Homolog | 0.207±0.07 | 0.113±0.01 | 0.109±0.01 |
| STL | 0.210±0.14 | 0.474±0.06 | 0.302±0.03 |
| Sparse+LR | 0.216±0.13 | 0.482±0.05 | 0.312±0.04 |
| BSL-MTL | 0.295 ±0.10 | 0.544 ±0.06 | 0.505 ±0.05 |

To get an estimate of the relative size of the nonhomologous partition for each task, see Table 1. Only the most representative baselines are retained for this experiment. Our model performs better than the others on the nonhomologous part of the test data, as it captures nonlinear similarities among the interactions within a task and across tasks very well.

Highest AUC is highlighted in bold.

TABLE 4. AREA UNDER THE PRECISION-RECALL CURVE ON THE HOMOLOGOUS PROTEIN-PROTEIN INTERACTION WITHIN EACH TASK, IN THE 10% SETTING

| <i>10% training, test only homologous PPI</i> | | | |
|---|--------------------|--------------------|--------------------|
| | <i>Ebola</i> | <i>Hep-C</i> | <i>Influenza</i> |
| Homolog | 0.399 ±0.07 | 0.161±0.01 | 0.180±0.01 |
| STL | 0.192±0.15 | 0.504±0.10 | 0.305±0.03 |
| Sparse+LR | 0.184±0.13 | 0.512±0.10 | 0.315±0.04 |
| BSL-MTL | 0.130±0.10 | 0.815 ±0.02 | 0.475 ±0.02 |

The Homolog baseline outperforms all on the *Ebola* task and BSL-MTL beats all other methods on the other tasks. The Homolog method has access to a wider virus-human PPI graph involving several other viruses and is able to infer homologous PPI for *Ebola* better (note that some of the homologous PPIs are not redundant examples but involve proteins with very small regions of sequence similarity).

Highest AUC is highlighted in bold.

no redundant examples). Furthermore, every set of homologous PPIs in the test data is replaced with a single PPI. A BLAST sequence similarity cutoff of $1e^{-5}$ was used to find homologs. The results are given in Table 3.

- Homologous PPI partition: only homologous PPI part of the test data. This is essentially the set of PPI removed in the set-(b). Table 4 shows these results.

Note that the partitions in (b) and (c) together comprise our test data. Also, the negatives in the test data are randomly split into the “no-homolog” and “homolog” partitions while maintaining the number of negatives in each partition to be 100 times the number of positives.

Overall performance: For reference, the AUC-PR of a random classifier model is ≈ 0.01 because of the class imbalanced nature of the data. In general, we notice that multitask learning benefits all tasks. The first three columns show the results in the 10% setting. Our model (last row) has large gains for influenza (1.4 times better than the next best) and modest improvements for the other tasks. The variance in the AUC is high for the Ebola task (column 1) owing to the small number of positives in the training splits. The most benefits for our model are seen in the 30% setting for all tasks. Notably, Homolog performs comparably to our model. In addition to the training data seen by all other methods, this baseline also uses additional data in the form of the virus-virus, human-human sequence similarities and a virus-human interactome from several other viruses (besides the three tasks). Results on the partitions shed further light on this.

Performance on the nonhomologous PPI partition: Table 3 shows these results. Overall, our method performs best on this partition over all tasks. Comparison with the AUC in Table 2 shows that AUC for the *Ebola* task is higher on all baselines, except Homolog, indicating that a bigger fraction of the overall performance comes from the nonhomologous partition (this is also marginally the case for *Influenza*). This effect is the most pronounced for our method (BSL-MTL) and demonstrates that we are able to predict such PPI by sharing information with other tasks. There are thus two components to the prediction performance—(1) the presence of sequence-similar proteins within a task and (2) nonlinear similarities between the interactions within a task and across tasks. We are able to capture the latter well through the shared subspace structure of our model.

Performance on the homologous PPI partition (Table 4): Except the *Ebola* task, BSL-MTL has the highest AUC. The Homolog baseline best captures the homologous PPI in *Ebola* (with an AUC twice the next best). However, as there are far fewer homologous PPI in Ebola than in the other tasks (row 2 of Table 1), this contributes less to the overall performance of the Homolog baseline. *Hepatitis-C* has the highest fraction of homologous PPI and BSL-MTL does significantly better on these than the other baselines, which also leads to a higher overall AUC. The same can be said of the *Influenza* task.

Disadvantages of the Homolog method: In Figure 3, we show for two tasks the protein interactions in the test data (for one train-test split) that were not found by the Homolog model that relies purely on homology-based sequence similarity.




FIG. 1. Principal component analysis (PCA) of virus proteins in the original feature space. The first two principal components are shown. Shape and color of the points indicate which virus that protein comes from.

6.1. Biological significance of the model

The model parameters U , V , and S are a source of rich information that can be used to further understand host–pathogen interactions. Note that our features are derived from the amino acid sequences of the proteins that provide opportunities to interpret the parameters.

1. Clustering proteins based on interaction propensities: We analyze the proteins by projecting them using the model parameters U and V into a lower dimensional subspace (i.e., computing XU^T and YV^T to get projections of the virus and human proteins, respectively). The principal component analysis (PCA) of this lower dimensional representation is compared with PCA in the original feature space (protein sequence features) in Figures 1 and 2. First, the projected data have a much better separation than the original data. Second, Figure 2 shows that Hepatitis-C and Influenza have many proteins with similar binding tendencies, and that these behave differently than most Ebola virus proteins. This observation is not obvious in the PCA of the original feature space (Fig. 1), in which proteins with similar sequences group together. We analyze the projected data further by looking at the clusters of




FIG. 2. PCA of virus proteins in the projected subspace. The first two principal components are shown. Shape of the points indicates which virus that protein comes from. We have highlighted three clusters for which we show gene ontology term enrichment analysis results in Table 5.

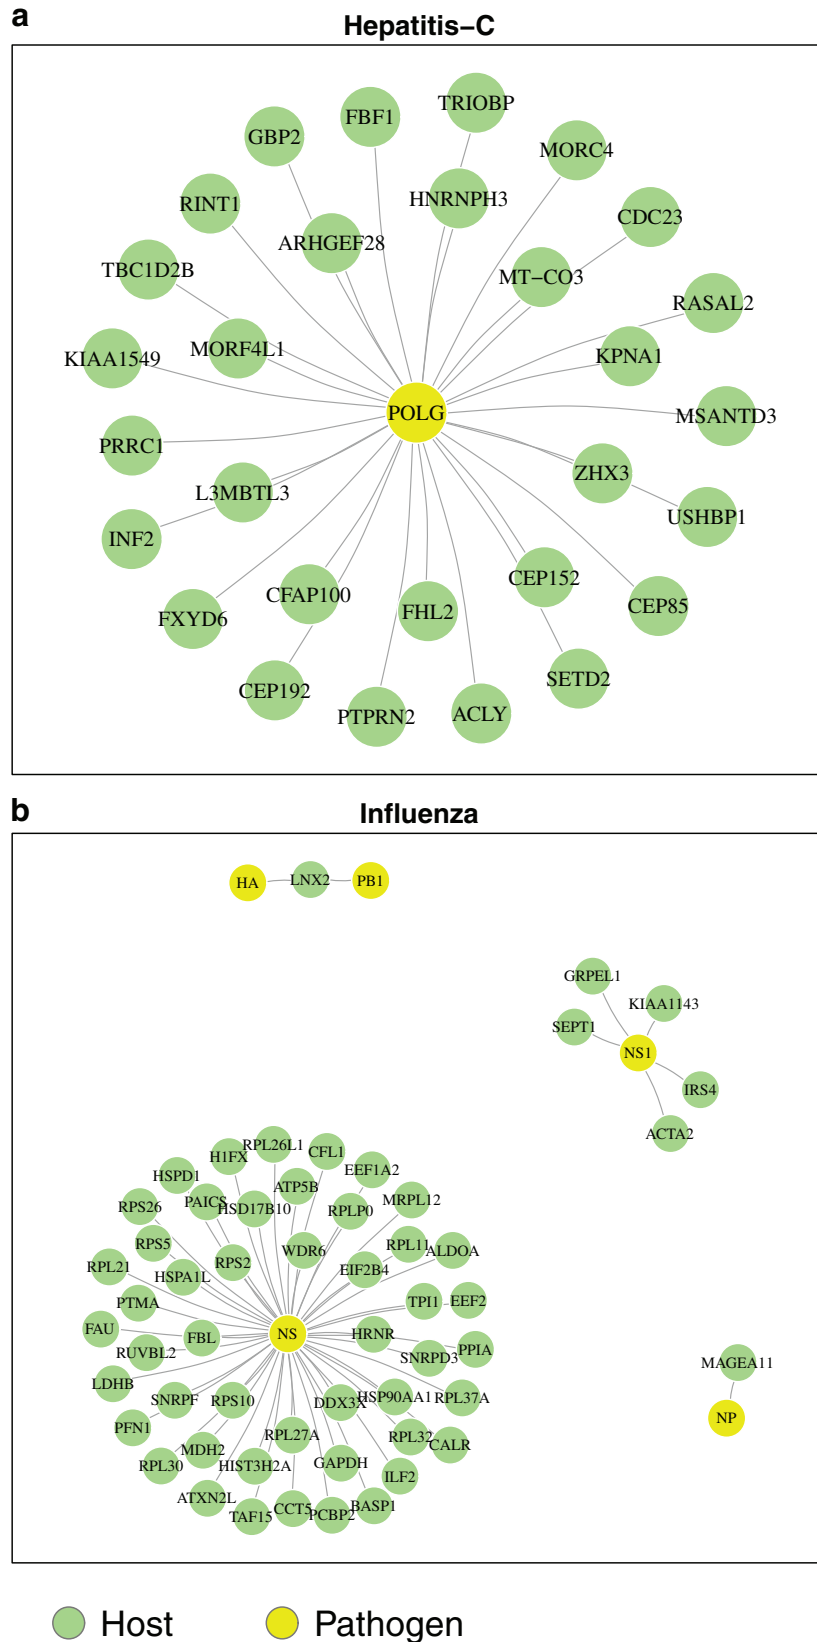


FIG. 3. PPIs from each task that are found by our method but are missed by the homolog method. These are those for which sequence-similar proteins were not found. (a) *Hepatitis-C*-human PPI and (b) *Influenza*-human PPI. PPI, protein-protein interaction.

TABLE 5. GENE ONTOLOGY TERMS (FOR PROCESS, FUNCTION, COMPONENT) ENRICHED IN THE HIGHLIGHTED CLUSTERS FROM FIGURE 2

| <i>GO terms for clusters from Figure 2</i> | |
|--|---|
| Cluster 1 (Ebola only) | mRNA (guanine-N7-)-methyltransferase activity; RNA-directed RNA polymerase activity; ATP binding |
| Cluster 2 (hepatitis and influenza proteins) | Host cell endoplasmic reticulum membrane; host cell lipid particle; induction by virus of host autophagy; host cell mitochondrial membrane; suppression by virus of host STAT1 activity |
| Cluster 3 (all) | Fusion of virus membrane with host plasma membrane |

FuncAssociate GO enrichment tool⁴ was used and terms with $p < 0.001$ are shown. The first column also indicates the composition of each cluster (i.e., which task's proteins appear in the cluster).

proteins for enrichment of GO³ (Ashburner et al., 2000) annotations (proteins were first clustered in this lower dimensional space using k-means and setting the number of clusters to 6). Of particular interest are the highlighted three clusters that contain either proteins projected far from others (such as cluster 1) or proteins from different viruses projected close together (cluster 2 and cluster 3). For enrichment analysis, we use the FuncAssociate 3.0⁴ (Berriz et al., 2003) tool and GO annotations for the three viruses from UniProtKB.

2. Novel interactions with Ebola proteins: The top four Ebola–human PPIs are all predictions for the Ebola envelope glycoprotein (GP) with four different human proteins (Note: GP is not in the gold standard PPIs). There is abundant evidence in the published literature (Nanbo et al., 2010) for the critical role played by GP in virus docking and fusion with the host cell.
3. Sequence motifs from virus proteins: We analyze sequence motifs derived from the top k -mers that contribute to interactions. The significant entries of the model parameters U , V , and $\{S_t\}$ were used to compute these motifs. The top positive-valued entries from the product UV^T indicate which pairs of features $((f_v, f_h)$: virus protein feature, human protein feature) are important for interactions across all the virus–human PPI tasks. Analogously, the entries from S_t give us pairs of features important to a particular virus–human task “ t ”. These results are given in the Supplementary Data.

7. CONCLUSIONS AND FUTURE EXTENSIONS

This work developed and tested a new multitask learning method for HP PPI prediction based on low-rank matrix completion for sharing information across tasks. Our method, BSL-MTL, exhibited large increases in prediction accuracy compared with a variety of baselines. The model parameters provide several avenues for further studying HP PPI that can lead to interesting observations and insights. Finally, the model we present is general enough to be applicable on other problems such as gene–disease relevance prediction across organisms or disease conditions.

AUTHOR DISCLOSURE STATEMENT

No conflicting financial interests exist.

REFERENCES

- Abernethy, J., Bach, F., Evgeniou, T., and Vert, J.P. (2009). A new approach to collaborative filtering: Operator estimation with spectral regularization. *J. Mach. Learn. Res.* 10, 803–826.

³GO: gene ontology. URL: www.geneontology.org

⁴<http://llama.mshri.on.ca/funcassociate>

- Ando, R.K., and Zhang, T. (2005). A framework for learning predictive structures from multiple tasks and unlabeled data. *J. Mach. Learn. Res.* 6, 1817–1853.
- Ashburner, M., Ball, C.A., Blake, J.A., et al. (2000). Gene ontology: Tool for the unification of biology. *Nat. Gene.* 25, 25–29.
- Berriz, G., King, O., Bryant, B., et al. (2003). Characterizing gene sets with funcassociate. *Bioinformatics* 19, 2502–2504.
- Candes, E., and Recht, B. (2009). Exact matrix completion via convex optimization. *Found. Comput. Math.* 9, 717.
- Cao, B., Liu, N.N., and Yang, Q. (2010). Transfer learning for collective link prediction in multiple heterogenous domains. In *International Conference on Machine Learning*, 159–166.
- Chatraramontri, A., Ceol, A., Peluso, D., et al. (2009). Virusmint: A viral protein interaction database. *Nucleic Acids Res.* 37, D669–D673.
- Chen, C.-C., Lin, C.-Y., Lo, Y.-S., and Yang, J.-M. (2009). Ppsearch: A web server for searching homologous protein–protein interactions across multiple species. *Nucleic Acids Res.* 37(suppl 2), W369–W375.
- Chen, J., Liu, J., and Ye, J. (2012). Learning incoherent sparse and low-rank patterns from multiple tasks. *ACM Trans. Knowl. Discov. Data* 5, 22.
- Chen, J., Tang, L., Liu, J., and Ye, J. (2013). A convex formulation for learning a shared predictive structure from multiple tasks. *IEEE Trans. Pattern Anal. Mach. Intell.* 35, 1025–1038.
- Chen, X., and Liu, M. (2005). Prediction of protein–protein interactions using random decision forest framework. *Bioinformatics* 21, 4394–4400.
- Davis, J., and Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the 23rd International Conference on Machine Learning*, 233–240.
- Dyer, M., Murali, T., and Sobral, B. (2007). Computational prediction of host–pathogen protein–protein interactions. *Bioinformatics* 23, i159–66.
- Dyer, M., Murali, T.M., Sobral, B.W. (2011). Supervised learning and prediction of physical interactions between human and hiv proteins. *Infect. Genet. Evol.* 11, 917–923.
- Evgeniou, T., and Pontil, M. (2004). Regularized multi-task learning. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 109–117.
- Jain, P., and Dhillon, I.S. (2013). Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*.
- Ji, S., and Ye, J. (2009). An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 457–464. ACM.
- Jin, R., Si, L., and Zhai, C. (2002). Preference-based graphic models for collaborative filtering. 329–336, In *Proceedings of the 19th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc.
- Koren, Y., Bell, R., and Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *IEEE Comput.* 42, 30–37.
- Kshirsagar, M., Carbonell, J.G., and Klein-Seetharaman, J. (2012). Techniques to cope with missing data in host–pathogen protein interaction prediction. *Bioinformatics* 28, i466–i472.
- Kshirsagar, M., Carbonell, J.G., and Klein-Seetharaman, J. (2013). Multi-task learning for host–pathogen protein interactions. *Bioinformatics* 29, i217–i226.
- Kumar, R., and Nanduri, B. (2010). Hpidb—A unified resource for host–pathogen interactions. *BMC Bioinformatics* 11(Suppl 6), S16.
- Li, B., Yang, Q., and Xue, X. (2009). Transfer learning for collaborative filtering via a rating-matrix generative model. In *International Conference on Machine Learning*, 617–624. ACM.
- Lin, C.-Y., Chen, Y.-C., Lo, Y.-S., and Yang, J.-M. (2013). Inferring homologous protein–protein interactions through pair position specific scoring matrix. *BMC Bioinformatics* 14, 1.
- Matthews, L.R., Vaglio, P., Reboul, J., Ge, H., Davis, B.P., Garrels, J., Vincent, S., and Vidal, M. (2001). Identification of potential interaction networks using sequence-based searches for conserved protein–protein interactions or interologs. *Genome Res.* 11, 2120–2126.
- Menon, A.K., and Elkan, C. (2011). Link prediction via matrix factorization. *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 437–452.
- Mika, S., and Rost, B. (2006). Protein–protein interactions more conserved within species than across species. *PLoS Comput. Biol.* 2, e79.
- Mirzazadeh, F., White, M., György, A., and Schuurmans, D. (2015). Scalable metric learning for co-embedding. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-KDD)*, 625–642. Springer.
- Murakami, Y. and Mizuguchi, K. (2014). Homology-based prediction of interactions between proteins using averaged one-dependence estimators. *BMC Bioinformatics* 15, 213.
- Nanbo, A., Imai, M., Watanabe, S., et al. (2010). Ebola virus is internalized into host cells via macropinocytosis in a viral glycoprotein-dependent manner. *PLoS Pathog.* 6, e1001121.
- Natarajan, N., and Dhillon, I.S. (2014). Inductive matrix completion for predicting gene–disease associations. *Bioinformatics* 30, i60–i68.

- Phung, D.Q., Venkatesh, S., et al. (2009). Ordinal boltzmann machines for collaborative filtering. In *UAI*, 548–556. AUAI Press.
- Qi, Y., Bar-Joseph, Z., Klein-Seetharaman, J. (2006). Evaluation of different biological data and computational classification methods for use in protein interaction prediction. *Proteins* 63, 490–500.
- Qi, Y., Dhiman, H., Bar-Joseph, Z., et al. (2009). Systematic prediction of human membrane receptor interactions. *Proteomics* 23, 5243–55.
- Qi, Y., Tastan, O., Carbonell, J.G., Klein-Seetharaman, J., and Weston, J. (2010). Semi-supervised multi-task learning for predicting interactions between HIV-1 and human proteins. *Bioinformatics* 26, i645–i652.
- Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web*, 285–295. ACM.
- Shen, J., Zhang, J., Luo, X., et al. (2007). Predicting protein-protein interactions based only on sequences information. *Proc. Natl. Acad. Sci. U. S. A.* 104, 4337–4341.
- Singh, A.P., and Gordon, G.J. (2008). Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 650–658.
- Singh, R., Xu, J., and Berger, B. (2006). Struct2net: Integrating structure into protein-protein interaction prediction. *Pac. Symp. Biocomput.* 11, 403–411.
- Tastan, O., Qi, Y., Carbonell, J.G., et al. (2009). Prediction of interactions between hiv-1 and human proteins by information integration. *Pac. Symp. Biocomput.* 516–527.
- Tekir, S.D., Cakir, T., and Ulgen, K.Ö. (2012). Infection strategies of bacterial and viral pathogens through pathogen-human protein protein interactions. *Syst. Biol. Microb. Infect.* 3, 74.
- UniProt Consortium (2011). Ongoing and future developments at the universal protein resource. *Nucleic Acids Res.* 39, D214–D219.
- Walhout, A.J., and Vidal, M. (2001). Protein interaction maps for model organisms. *Nat. Rev. Mol. Cell Biol.* 2, 55–63.
- Wu, X., Zhu, L., Guo, J., et al. (2006). Prediction of yeast protein-protein interaction network: Insights from the gene ontology and annotations. *Nucleic Acids Res.* 34, 2137–2150.
- Xu, Q., Xiang, E.W., and Yang, Q. (2010). Protein-protein interaction prediction via collective matrix factorization. In *International Conference on Bioinformatics and Biomedicine*.
- Zhang, Y., Cao, B., and Yeung, D.-Y. (2012). Multi-domain collaborative filtering. *arXiv preprint arXiv:1203.3535*.
- Zhou, J., Chen, J., and Ye, J. (2011). *MALSAR: Multi-task Learning via Structural Regularization*. Arizona State University.

Address correspondence to:

Dr. Meghana Kshirsagar
Memorial Sloan Kettering Cancer Center
1275 York Avenue
New York, NY 10021

E-mail: meghanak@ymail.com