

RESEARCH ARTICLE

Open Access



# VB-MK-LMF: fusion of drugs, targets and interactions using variational Bayesian multiple kernel logistic matrix factorization

Bence Bolgár\*  and Péter Antal

## Abstract

**Background:** Computational fusion approaches to drug-target interaction (DTI) prediction, capable of utilizing multiple sources of background knowledge, were reported to achieve superior predictive performance in multiple studies. Other studies showed that specificities of the DTI task, such as weighting the observations and focusing the side information are also vital for reaching top performance.

**Method:** We present Variational Bayesian Multiple Kernel Logistic Matrix Factorization (VB-MK-LMF), which unifies the advantages of (1) multiple kernel learning, (2) weighted observations, (3) graph Laplacian regularization, and (4) explicit modeling of probabilities of binary drug-target interactions.

**Results:** VB-MK-LMF achieves significantly better predictive performance in standard benchmarks compared to state-of-the-art methods, which can be traced back to multiple factors. The systematic evaluation of the effect of multiple kernels confirm their benefits, but also highlights the limitations of linear kernel combinations, already recognized in other fields. The analysis of the effect of prior kernels using varying sample sizes sheds light on the balance of data and knowledge in DTI tasks and on the rate at which the effect of priors vanishes. This also shows the existence of “small sample size” regions where using side information offers significant gains. Alongside favorable predictive performance, a notable property of MF methods is that they provide a unified space for drugs and targets using latent representations. Compared to earlier studies, the dimensionality of this space proved to be surprisingly low, which makes the latent representations constructed by VB-ML-LMF especially well-suited for visual analytics. The probabilistic nature of the predictions allows the calculation of the expected values of hits in functionally relevant sets, which we demonstrate by predicting drug promiscuity. The variational Bayesian approximation is also implemented for general purpose graphics processing units yielding significantly improved computational time.

**Conclusion:** In standard benchmarks, VB-MK-LMF shows significantly improved predictive performance in a wide range of settings. Beyond these benchmarks, another contribution of our work is highlighting and providing estimates for further pharmaceutically relevant quantities, such as promiscuity, druggability and total number of interactions.

**Keywords:** Drug-target interaction prediction, Matrix factorization, Multiple kernel learning, Variational Bayes, Probabilistic graphical models

\*Correspondence: bolgar@mit.bme.hu  
Department of Measurement and Information Systems, Budapest University of  
Technology and Economics, Magyar tudósok krt. 2., 1117 Budapest, Hungary

## Background

Drug-target interactions (DTI) or compound-protein interactions (CPIs) have become a focal point in chemo- and bioinformatics. There are many factors behind this trend, such as the direct, quantitative nature of bioactivity data [1], its unprecedented amount, public availability [2, 3], and variety including also phenotypic and content-rich assays and screenings [4]. Further factors are the semantic, linked open nature of the data [5, 6], collaborative initiatives in the pharmaceutical policy [1] and the construction of DTI benchmarks [7–13].

An additional factor is the varying granularity and multiple facets of the DTI task: it was already attacked in the 90's in single target scenarios, e.g. by using neural networks of that time [14] and subsequently by kernel methods [15, 16]. A series of similarity-based methods were also developed for virtual screening [17–19]; in the early 2000's molecular docking became popular [20, 21]; from the late 2000's matrix factorization methods were developed [7, 22, 23]. As the importance of data and knowledge integration in drug discovery was further emphasized [1, 24–26], the incorporation of prior knowledge in DTI became mainstream and indeed improved predictive performance [23, 27–29].

Computational data and knowledge fusion approaches in the DTI problem seem to be especially relevant, as the growth of DTI datasets is limited by experimental and publication time and cost, while the cross-linked repertoire of side information expands at an enormous rate. This grand pool of information complementing the DTI data and the full scope of the DTI fusion challenge is best illustrated by the drug repositioning problem [30, 31]. In repositioning, i.e. in the finding of a novel indication for an already marketed drug, extra information sources could also be used, such as off-label drug usage patterns, patient-reported adverse-effects and official side-effects [32]. Notably, this information pool can be linked back to early stage compound discovery [33].

In this paper we investigate the multiple kernel-based fusion approach to the DTI task from a computational fusion perspective, by adopting widely used benchmark datasets, implementations and evaluation methodologies from Yamanishi et al. [7], Gönen [22], Pahikkala et al. [8] and Liu et al. [34]. Our contributions are as follows:

1. *VB-MK-LMF*: We present a Bayesian matrix factorization method with a novel variational Bayesian approximation, which unifies multiple kernel learning, importance weight for (positive) observations, network-based regularization and explicit modeling of probabilities of drug-target interactions.
2. *Effect of multiple kernels*: We report the results of a comparison against three leading solutions using two

benchmark datasets, in which VB-MK-LMF achieved significantly better performance in most settings. We systematically investigate factors behind its performance, such as the type of the kernels, the role of neighborhood restriction and Bayesian averaging. Finally, we evaluate the effect of priors using varying sample sizes highlighting the regions where using side-information improves predictive performance.

3. *Posteriors for promiscuity and druggability*: We show that probabilistic predictions from VB-MK-LMF can be used to quantify the expected values for promiscuity or the number of hits in a DTI task.
4. *Dimensionality of the unified “pharmacological” space*: We investigate the learned unified latent representations of drugs and targets, and contrary to many studies we argue that drastically smaller dimensions are sufficient. We discuss the possibility that this low dimension, around 10, could be utilized in visual analytics and exploratory data analysis.
5. *Accessibility*: We report the adaptation of the developed variational Bayesian approximation to general purpose graphics processing units (GP-GPU). Evaluations show that 30× speed-up can be achieved using a standard GP-GPU environment. To support the development of current DTI benchmarks towards “computational DTI fusion”, we release the applied kernels, code and parameter settings for academic use.

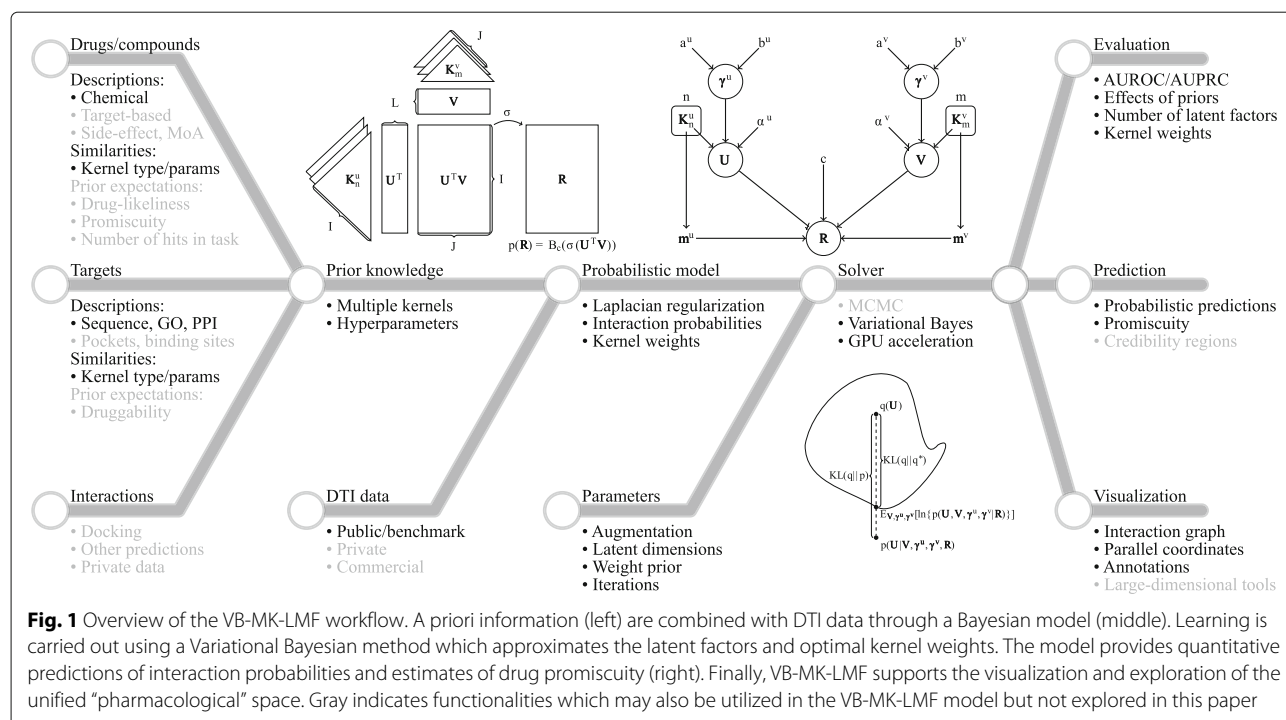
Figure 1 shows the overview of Variational Bayesian Multiple Kernel Logistic Matrix Factorization (VB-MK-LMF).

## Related works

To give an overview about related, earlier works [7, 27–29, 35–54], we summarize the main properties of their applied datasets, side information, methods and evaluation methodologies in Additional file 1).

## DTI data

Drug-target interaction data has become a fundamental resource in pharmaceutical research, which can be attributed to its public availability in an open linked format, see e.g. [1, 5, 6, 55–58]. The relative objectivity of interaction activities and the side information about drugs and targets renders a unique status to the comprehensive tabular DTI data, even compared to media and e-commerce data [59], despite the issues of quality [60, 61], duality of commercial and public repositories [62–64] and selection bias related to the lack of negative samples [12] and promiscuity [65]. However, at present the heterogeneous, real-valued activity data are usually treated as binary relations, even though the use of raw data together with information about the measurement



context is expected in more realistic DTI prediction scenarios [8, 46, 52]. Another largely overlooked property of the binary drug-target interaction data is its possibly indirect nature, which influences the applicable target-target similarities, e.g. in the indirect case protein-protein networks may have relevance (for the explicit treatment of direct and indirect relations, see e.g. RBM [45]).

### DTI prior knowledge

The molecular similarity property principle [66, 67], the drug-likeness of a compound [68, 69] and druggability of proteins [70] are essential concepts in the broader drug discovery context, together with molecular docking [20, 21] and binding site, pocket predictors [71], if structure information is available. However, their use as priors in the computational DTI task is still largely unexplored. If the goal is the discovery of indirect drug-target interactions, possibly including multiple paths, which are especially relevant in polypharmacology [72], then the use of molecular interaction and regulatory networks alongside protein-protein similarities is another open issue.

Chemical similarity, the most widespread source of prior knowledge in DTI, was the basis of many “guilt-by-association” approaches in chemo- and bioinformatics. Earlier investigations helped to understand the use of multiple, heterogeneous representations, similarity measures and introduced the concept of fusion methods in ligand-based virtual screening [17, 18, 73–75]. Beyond chemical similarities, target-based similarities can also be used to exceed activity cliffs [32]; moreover, side-effect based and

off-label usage based similarities can be constructed for compounds using FDA-approved drugs as canonical bases in a group-representation [33].

Target-target similarities are another diverse and voluminous of prior information, which can be defined using sequence similarities, common motifs and domains, phylogenetic relations or shared binding sites and pockets [71]. In case of indirect drug-target interactions, a broader set of target-target similarities could be based on relatedness in pathways, protein-protein networks and functional annotations, e.g. from Gene Ontology [76].

We concentrate on predicting presumably direct activities in this paper, thus we demonstrate the capability of the developed method and the effect multiple information sources using multiple chemical similarities, although the method can incorporate symmetrically multiple target-target similarities. Furthermore, the method can also incorporate separate prior expectations about the success rates of drugs in a given DTI, which could be combined with drug-likeness [77], promiscuity prediction [78] and decoy prediction in case of their use [79]. Symmetrically, it can also incorporate separate prior expectations about the success rates of targets in a given DTI, which could be combined with druggability predictions [70, 80, 81] and the presence of pockets [82]. For an overview of available resources relevant for the DTI task, see e.g. [83, 84].

### DTI methods

The rapid growth, especially the public availability of tabular (dyadic) DTI data in the last decade caused a dramatic

shift of the applied statistical methods. For an overview of classical single prediction oriented machine learning and data mining in drug discovery, especially in DTI and ADME predictions, see e.g. [85], for large-scale, comprehensive applications of DTI data, see e.g. [86]. The tabular nature of the DTI data called for new methods not only handling this type of data natively, but also capable of using side information. Transfer learning and multitask learning paradigms addressed this challenge [8, 87, 88], but in the DTI context, two groups of methods, the pairwise conditional methods and the matrix factorization based generative methods proved to be particularly successful.

Pairwise conditional approaches or pairwise kernel methods flatten the dyadic structure of the DTI data and use drug and target descriptors, optionally even explanatory descriptors about the drug-target relations to predict interaction properties of drug-target pairs (for the assumptions behind the conditional approach, see e.g. [89], for its early DTI application, see e.g. [90]). Classification and regression methods, such as MLPs, decision trees and SVMs remain directly applicable in this conditional approach (not modeling the distribution of the drug-target pairs), however, the high number of drug-target pairs is challenging for kernel based methods [51, 91], but recent developments in deep learning show promising results [92]. Using multiple representations for drugs and targets is directly possible in this pairwise approach, but the construction of an aggregate pair-pair (interaction-interaction) similarity or an efficient set of pair-pair similarities from drug-drug and target-target similarities is an open problem. In the case of single drug-drug and target-target similarities, the Kroneckerian combination was proposed in the work of van Laarhoven [91] with corresponding computational simplifications to maintain scalability. Additionally, kernel techniques were extended to use multiple kernels, which are potentially derived from heterogeneous representations and similarities [51]. Recent extensions include non-linear kernel fusion in the RLS-KF system [50] and using boosting to learn from unscreened controls [54].

Matrix factorization (MF) methods differ from pairwise approaches in multiple properties crucial in the DTI task. The central operation of these methods is the construction of a joint space with latent factors for drugs and targets and modeling their interactions based on the inner product of the respective vectorial representations. Contrary, pairwise approaches, such as kernel methods or deep learning cannot directly exploit the tabular prior constraint of the data. The MF approach also allows the direct incorporation of drug-drug similarities and target-target similarities. Additionally, the low dimensionality of the latent space supports data visualization, although its interpretation is still in its infancy.

Finally, probabilistic MF methods construct a distribution over the latent representations of drugs and targets, which in fact means that they are full-fledged generative models.

Matrix factorization methods were adopted early in gene expression data analysis [93, 94]. They were used for dimensionality reduction and the construction of a unified space for ligands and receptors [95], applied in biomedical text-mining and [96] and chemogenomics [97]. Later in the 2000's media and e-commerce recommendation applications dominated the research of matrix factorization methods [98] and many developments were motivated and reported in these contexts, such as solutions for new items without interactions, selection bias, model regularization, automated parameter selection and incorporation of side information from multiple sources. An early work from Srebro et al. addressed the problems of using weights to represent importance or trust in the observations and the use of logistic regression as a non-linear transformation to predict probabilities of binary observations [99]. A special weighting of observations compared to unknowns were investigated in [100]. Salakhutdinov introduced Bayesian matrix factorization, which addressed regularization and automated parameter selection by Bayesian model averaging, also indicating the principled and flexible options for prior incorporation [101]. Severinski demonstrated the advantages of the full Bayesian approach versus a Maximum a Posteriori based alternative in this context [102]. Zhou introduced Gaussian process priors over the latent dimensions to enforce two kernels over row and column items [103]. Lobato et al. reported a variational Bayesian approach for logistic matrix factorization [104].

In the DTI context, an early kernel regression-based method (KRM) was reported in [7], and emphasized the advantages of a unified "pharmacological space". Gönen introduced a kernelized Bayesian matrix factorization (KBMF) [22], which applies kernel-based averaging over the latent vectorial representations of rows and columns. The paper also introduced an efficient variational Bayesian approximation and indicated the interpretability of the latent space. Zheng et al. proposed a non-probabilistic multiple kernel learning approach, which achieved superior performance [23]. Multiple kernel learning was also realized in KBMF [27] and was also extended towards regression [105]. Special non-missing-at-random DTI data models were proposed in [52], which applied Gaussian priors to incorporate multiple kernels and used Gibbs sampling to approximate the posteriors. In an integrative work, Liu et al. proposed the combination of special neighborhood restricted kernels, network-based regularization, importance weights for the observations and logistic link functions in a non-Bayesian framework [48]. A recent extension applied a

nonlinear kernel diffusion technique to boost relevant, complementary information in similarity matrices [49].

### DTI benchmarks

The most widely used DTI benchmark from Yamanishi et al. [7] defined DTI prediction as a binary prediction problem with a single source of drug-drug and a target-target similarity, which induced the development of variety of methods and datasets (see Additional file 1). These datasets are still in the range of  $1000 \times 1000$  and contain  $10k$  interactions, but they inherit the problem of the selection bias present in the DTI repositories [11, 12, 65, 83, 106, 107]. Pahikkala et al. stressed the importance of fully observed bioactivity values in benchmarks [8], such as from Davis [9], to avoid misleading results because of selection bias, indirect interactions and the binary nature of the interactions. Liu et al. [48] reported a comprehensive evaluation of methods and released a corresponding benchmark implementation, the *pyDTI* package. For real, experimental evaluation of DTI methods, see e.g. [108, 109].

### Methods

Our work directly builds upon Gönen's work on kernel-based matrix factorization using twin kernels (KBMF-MKL), which applied variational Bayesian approximations [27]. Another direct predecessor of our work is Liu et al's neighborhood regularized logistic matrix factorization [48].

### Materials

To maintain consistency with earlier works, we evaluated the methods on the data sets provided by Yamanishi et al. [7] and Pahikkala et al. [8]. While the latter comes with multiple similarity matrices based on various molecular fingerprints, the former is one-kernel and therefore needed to be extended to properly test the MKL performance. We used the RDKit package [110] to compute additional MACCS and Morgan fingerprints for the molecules and used these in conjunction with the Tanimoto and Gaussian RBF similarity measures. Target similarities were obtained from Nascimento et al. [51] which utilized sequential, GO- and PPI-based similarities.

### Probabilistic model

Let  $\mathbf{R} \in \{0, 1\}^{I \times J}$  denote the matrix of the interactions, where  $\mathbf{R}_{ij} = 1$  indicates a known interaction between the  $i$ th drug and  $j$ th target. In order to formulate a Bayesian model, we put a Bernoulli distribution on each  $\mathbf{R}_{ij}$  with parameter  $\sigma(\mathbf{u}_i^T \mathbf{v}_j)$  where  $\sigma$  is the logistic sigmoid function and  $\mathbf{u}_i, \mathbf{v}_j$  are the  $i$ th and  $j$ th columns of the respective factor matrices  $\mathbf{U} \in \mathbb{R}^{L \times I}$  and  $\mathbf{V} \in \mathbb{R}^{L \times J}$ . One can think of  $\mathbf{u}_i$  and  $\mathbf{v}_j$  as  $L$ -dimensional latent representations of the

$i$ th drug and  $j$ th target, and the *a posteriori* probability of an interaction between them is modeled by  $\sigma(\mathbf{u}_i^T \mathbf{v}_j)$ .

Similarly to NRLMF, we utilize an augmented version of the Bernoulli distribution parameterized by  $c \geq 1$  which assigns higher importance to observations (positive examples). NRLMF also uses a post-training weighted average to infer interactions corresponding to empty rows and columns in  $\mathbf{R}$  (i.e. these would have to be estimated without using any corresponding observations). We account for them by introducing variables  $\mathbf{m}^u, \mathbf{m}^v \in \{0, 1\}$  indicating whether the row or column is empty. In these cases, only the side information will be used in the prediction. The conditional on the interactions can be written as

$$p(\mathbf{R} | \mathbf{U}, \mathbf{V}, c, \mathbf{m}^u, \mathbf{m}^v) \propto \prod_i \prod_j \left[ \left( \sigma(\mathbf{u}_i^T \mathbf{v}_j) \right)^{c \mathbf{R}_{ij}} \left( 1 - \sigma(\mathbf{u}_i^T \mathbf{v}_j) \right)^{1 - \mathbf{R}_{ij}} \right]^{\mathbf{m}_i^u \mathbf{m}_j^v} \quad (1)$$

Specifying priors on  $\mathbf{U}$  and  $\mathbf{V}$  presents an opportunity to incorporate multiple sources of side information. In particular, we can use a Gaussian distribution with a weighted linear combination of kernel matrices  $\mathbf{K}_n, n = 1, 2, \dots$  in the precision matrix, which corresponds to a combined  $L_2$ -Laplacian regularization scheme [36]

$$p(\mathbf{U} | \alpha^u, \gamma^u, \mathbf{K}^u) \propto \prod_i \prod_k \exp \left\{ -\frac{1}{2} \sum_n \gamma_n^u \mathbf{K}_{n,ik}^u \|\mathbf{u}_i - \mathbf{u}_k\|^2 \right\} \cdot \prod_i \exp \left\{ -\frac{\alpha^u}{2} \|\mathbf{u}_i\|^2 \right\} \quad (2)$$

The prior on  $\mathbf{V}$  can be written similarly. To automate the learning of the optimal value of kernel weights  $\gamma_n^u$ , we introduce another level of uncertainty using Gamma priors:

$$p(\gamma_n^u | a, b) = \frac{b^a (\gamma_n^u)^{a-1} e^{-b\gamma_n^u}}{\Gamma(a)} \quad (3)$$

### Variational approximation

In the Bayesian approach, the combination of the data  $\mathbf{R}$  and prior knowledge through kernel matrices  $\mathbf{K}_n$  and hyperparameters defines the posterior

$$p(\mathbf{U}, \mathbf{V}, \gamma^u, \gamma^v | \mathbf{R}, \mathbf{K}_n^u, a^u, b^u, \mathbf{K}_n^v, a^v, b^v, \alpha^u, \alpha^v, c)$$

In the variational setting [111], we approximate the posterior with a variational distribution  $q(\mathbf{U}, \mathbf{V}, \gamma^u, \gamma^v)$ . Suppressing the hyperparameters for notational simplicity, the expectation

$$p(\mathbf{R}) = \int p(\mathbf{R} | \mathbf{U}, \mathbf{V}) p(\mathbf{U} | \gamma^u) p(\mathbf{V} | \gamma^v) p(\gamma^u) p(\gamma^v) d\mathbf{U} d\mathbf{V} d\gamma^u d\gamma^v,$$

can be decomposed as

$$\ln p(\mathbf{R}) = \mathcal{L}(q) + KL(q \parallel p),$$

and, since the left hand side is constant with respect to  $q$ , maximizing the evidence lower bound  $\mathcal{L}(q)$  with respect to  $q$  is equivalent to minimizing the Kullback–Leibler divergence  $KL(q \parallel p)$  between the variational distribution and the true posterior. In the mean field variational approach, maximization of  $\mathcal{L}(q)$  is achieved by using a factorized variational distribution

$$q(\mathbf{U}, \mathbf{V}, \gamma^u, \gamma^v) = q(\mathbf{U})q(\mathbf{V})q(\gamma^u)q(\gamma^v).$$

In particular, the evidence lower bound takes the form [112]

$$\mathcal{L}(q) = \int q(\mathbf{U})q(\mathbf{V})q(\gamma^u)q(\gamma^v) \ln \left\{ \frac{p(\mathbf{R}, \mathbf{U}, \mathbf{V}, \gamma^u, \gamma^v)}{q(\mathbf{U})q(\mathbf{V})q(\gamma^u)q(\gamma^v)} \right\} d\mathbf{U}d\mathbf{V}d\gamma^u d\gamma^v.$$

The optimal distribution  $q^*(\mathbf{U})$  satisfies

$$\ln q^*(\mathbf{U}) = E_{\mathbf{V}, \gamma^u, \gamma^v} [\ln \{ p(\mathbf{R} \mid \mathbf{U}, \mathbf{V}) p(\mathbf{U} \mid \gamma^u) p(\mathbf{V} \mid \gamma^v) p(\gamma^u) p(\gamma^v) \}] + \text{const.}$$

which is non-conjugate due to the form of  $p(\mathbf{R} \mid \mathbf{U}, \mathbf{V})$  and therefore the integral is intractable. However, by using Taylor approximation on the symmetrized logistic function (Jaakkola’s bound [104, 113])

$$\sigma(z) \geq \tilde{\sigma}(z, \xi) = \sigma(\xi) \exp \left\{ \frac{z - \xi}{2} - \frac{1}{2\xi} \left( \sigma(\xi) - \frac{1}{2} \right) (z^2 - \xi^2) \right\},$$

we can lower bound  $p(\mathbf{R} \mid \mathbf{U}, \mathbf{V})$  at the cost of introducing local variational parameters  $\xi_{ij}$ , yielding a new bound  $\tilde{\mathcal{L}}$  which contains at most quadratic terms. Collecting the terms containing  $\mathbf{U}$  gives (see the proof in Additional file 2):

$$\begin{aligned} \ln q^*(\mathbf{U}) = & -\frac{1}{2} \text{tr}(\mathbf{U}^T \mathbf{Q}^u \mathbf{U}) + \sum_i \mathbf{u}_i^T \left( \sum_j \hat{\mathbf{R}}_{ij} \hat{\xi}_{ij} E[\mathbf{v}_j \mathbf{v}_j^T] \right) \mathbf{u}_i \\ & + \sum_i \mathbf{u}_i^T \left( \sum_j \mathbf{R}'_{ij} E[\mathbf{v}_j] \right) \end{aligned}$$

where

$$\begin{aligned} \mathbf{Q}^u &= \frac{E[\gamma_u]}{2} (\mathbf{K}^{uT} \mathbf{1} - \mathbf{K}^u) + \frac{\alpha_u}{2} \mathbf{I}, \\ \hat{\xi}_{ij} &= -\frac{1}{2\xi_{ij}} \left( \sigma(\xi_{ij}) - \frac{1}{2} \right), \\ \hat{\mathbf{R}}_{ij} &= \mathbf{m}_i^u \mathbf{m}_j^v (c - 1) \mathbf{R}_{ij} + 1, \\ \mathbf{R}'_{ij} &= \mathbf{m}_i^u \mathbf{m}_j^v c \mathbf{R}_{ij} + \frac{1}{2} \hat{\mathbf{R}}_{ij}. \end{aligned}$$

Since this expression is quadratic in  $\text{vec}(\mathbf{U})$ , we conclude that  $q^*$  is Gaussian and the parameters can be found by completing the square. In particular,

$$q^*(\text{vec}(\mathbf{U})) = \mathcal{N}(\text{vec}(\mathbf{U}) \mid \phi, \Lambda^{-1})$$

$$\Lambda = \mathbf{Q}^u \otimes \mathbf{I} - 2 \cdot \text{blkdg}_i \left( \sum_j \hat{\mathbf{R}}_{ij} \hat{\xi}_{ij} E[\mathbf{v}_j \mathbf{v}_j^T] \right), \tag{4}$$

$$\phi = \Lambda^{-1} \text{vec}_i \left( \sum_j \mathbf{R}'_{ij} E[\mathbf{v}_j] \right), \tag{5}$$

where  $\text{blkdg}_i$  denotes the operator creating an  $L \cdot I \times L \cdot I$  block-diagonal matrix from  $I \cdot L \times L$ -sized blocks. The variational update for  $q(\mathbf{V})$  can be derived similarly. The most computationally intensive operation is computing

$$E[\mathbf{v}_j \mathbf{v}_j^T] = \text{Cov}(\mathbf{v}_j) + E[\mathbf{v}_j] E[\mathbf{v}_j]^T \tag{6}$$

which requires the inversion of the precision matrix, performed using blocked Cholesky decomposition.

The optimal value of the local variational parameters  $\xi_{ij}$  can be computed by writing the expectation of the joint distribution in terms of  $\xi$  and setting its derivative to zero. In particular,

$$\begin{aligned} \tilde{\mathcal{L}}(\xi) = & \sum_i \sum_j \hat{\mathbf{R}}_{ij} \left( \ln \sigma(\xi_{ij}) - \frac{\xi_{ij}}{2} - \frac{1}{2\xi_{ij}} \left( \sigma(\xi_{ij}) - \frac{1}{2} \right) \right. \\ & \left. \times \left( \xi_{ij}^2 - E \left[ (\mathbf{u}_i^T \mathbf{v}_j)^2 \right] \right) \right), \end{aligned}$$

from which [104, 112]

$$\begin{aligned} \xi_{ij}^2 = & E \left[ (\mathbf{u}_i^T \mathbf{v}_j)^2 \right] \\ = & \left( E[\mathbf{u}_i]^T E[\mathbf{v}_j] \right)^2 + \sum_l E[\mathbf{U}_{li}]^2 V[\mathbf{V}_{lj}] + V[\mathbf{U}_{li}] E[\mathbf{V}_{lj}]^2 \\ & + V[\mathbf{U}_{li}] V[\mathbf{V}_{lj}]. \end{aligned} \tag{7}$$

Since the model is conjugate with respect to the kernel weights, we can use the standard update formulas for the Gamma distribution

$$\begin{aligned} q^*(\gamma_n^u) &= \text{Gamma}(\gamma_n^u \mid a', b') \\ a' &= a + \frac{l^2}{2} \\ b' &= b + \frac{1}{2} E_{\mathbf{U}} \left[ \sum_i \sum_k \mathbf{K}_{n,ik}^u \|\mathbf{u}_i - \mathbf{u}_k\|^2 \right] \\ &= b + \frac{1}{2} \sum_i \sum_k \mathbf{K}_{n,ik}^u \left( E[\mathbf{u}_i^T \mathbf{u}_i] - 2E[\mathbf{u}_i^T \mathbf{u}_k] \right. \\ & \quad \left. + E[\mathbf{u}_k^T \mathbf{u}_k] \right), \end{aligned} \tag{8}$$

which also requires the explicit inversion of  $\Lambda$ . Figure 2 shows the pseudocode of the algorithm.

**Require:**

Data:  $\mathbf{R} \in \{0, 1\}^{I \times J}$ ,  $\mathbf{K}_u \in (\mathbb{R}^{I \times I})^{n_u}$ ,  $\mathbf{K}_v \in (\mathbb{R}^{J \times J})^{n_v}$ .  
 Parameters:  $L, D \in \mathbb{N}$ ,  $a_u, b_u, a_v, b_v, \alpha_u, \alpha_v, c \in \mathbb{R}_+$ .

Calculate  $\mathbf{m}^u \in \{0, 1\}^I$ ;  $\mathbf{m}^v \in \{0, 1\}^J$ .  
 Initialize  $q^*(\gamma^u)$  and  $q^*(\gamma^v)$  using Eq. 3.  
 Initialize  $q^*(\mathbf{U})$  and  $q^*(\mathbf{V})$  using Eq. 2.  
**for**  $d$  in  $1 \dots D$  **do**  
   Update Jaakkola's bound  $\xi$  using Eq. 7.  
   Update  $q^*(\mathbf{U})$  using Eq. 4 and 5.  
   Update Jaakkola's bound  $\xi$  using Eq. 7.  
   Update  $q^*(\mathbf{V})$  the same way as  $q^*(\mathbf{U})$ .  
   Update  $q^*(\gamma^u)$  and  $q^*(\gamma^v)$  using Eq. 8 and 9.  
**end for**  
**return**  $q^*(\mathbf{U})$ ,  $q^*(\mathbf{V})$ ,  $q^*(\gamma^u)$ ,  $q^*(\gamma^v)$ .

**Fig. 2** Pseudocode of the VB-MK-LMF algorithm**Results**

We present the results of a systematic comparison with KBMF-MKL [27], NRLMF [48] and KronRLS-MKL [51] using their provided implementations. Subsequently, our results show the effect of prior knowledge fading with increasing data size.

**Experimental settings**

Predictive performance was evaluated in a  $5 \times 10$ -fold cross-validation framework. To maintain consistency with the evaluations in earlier works, we utilized the CVS1-CVS2-CVS3 settings as presented in [48] and calculated the average AUROC and AUPRC values in each scenario. In particular, CVS1 corresponds to evaluating predictive performance after randomly blinding 10% of the interactions and using them as test entities. CVS2 corresponds to random drugs (entire rows blinded) and CVS3 corresponds to random targets. We used the same folds as the PyDTI tool to maximize comparability.

In the single-kernel setting, we compared the performance of the proposed method to KBMF, NRLMF and KronRLS. The optimal parameters for NRLMF were obtained from the original publication [48]. KBMF and KronRLS were parameterized using a grid search method. VB-MK-LMF was used with 3 neighbors in each kernel,  $\alpha_u = \alpha_v = 0.1$ ,  $a_u = a_v = 1$ ,  $b_u = b_v = 10^3$  and  $c = 10$ . The number of latent factors was set to  $L = 10$  in the Nuclear Receptor dataset and  $L = 15$  in the others, and a more detailed investigation of this parameter was also conducted. The number of iterations was chosen manually as 20 since the variational parameters usually converged between 20 – 50 iterations.

In the multiple-kernel setting, we compared the performance of the proposed method to KBMF-MKL and KronRLS-MKL using MACCS and Morgan fingerprints with RBF and Tanimoto similarities. Target kernels provided by KronRLS-MKL did not improve the results in

either case, thus only the ones computed by Yamanishi et al. were utilized. We also investigated the weights assigned to the kernels and tested robustness by introducing kernels with random values.

**Systematic evaluation**

Single-kernel results are shown in Table 1. In most cases, VB-MK-LMF significantly outperforms NRLMF and one-kernel KBMF in terms of AUROC and AUPRC according to a pairwise  $t$ -test. Overall, the improvement is more modest on the Enzyme dataset, although still significant in some cases. This can be attributed to the fact that this dataset is by far the largest, which can mitigate the benefits of Bayesian model averaging and side information. On average, VB-MK-LMF yields 4.7% higher AUPRC values in the pairwise cross-validation setting than the second best method. In the drug and target settings, this is 2% and 7.6%, respectively. The lower AUROC and AUPRC values in these scenarios are explained by the lack of observations for the test drugs or targets in the training set, resulting in a harder task than in the pairwise scenario.

Following earlier investigations, we examined the number of latent factors, which has a crucial role from computational, statistical and interpretational aspects. Contrary to earlier works [44], which recommend 50 – 100 as the number of latent factors, we found that these values do not yield better results; in fact, the AUPRC values quickly become saturated. Conceptually, it is unclear what is to be gained going beyond the rank of the original matrix, which corresponds to perfect factorization with respect to the Frobenius norm when using SVD, and is also known to lead to serious overfitting in unregularized cases [99, 101]. Although overfitting is usually less of an issue with variational Bayesian approximations, a large number of latent factors significantly increases computational time. Figure 3 depicts the AUPRC values on the smaller datasets with varying number of latent factors. The Enzyme and Kinase datasets were not included in this experiment due to the rapidly increasing runtime.

Multi-kernel AUPRC values are shown in Table 2. Compared to the previous Table, it is clear that both VB-MK-LMF and KBMF benefits from using multiple kernels. Moreover, there is also an improvement in predictive performance when one combines instances of the same kernel but with different neighbor truncation values. However, advantages of using both of these combination schemes simultaneously are unclear as the results usually do not improve or even get worse (except for the Kinase dataset). This is a known property of linear kernel combinations, i.e. using large linear kernel combinations may not improve predictive performance beyond that of the best individual kernels in the combination [114].

Table 3 shows the normalized kernel weights in each of the datasets. For illustration purposes, we also included a

**Table 1** Single-kernel results on gold standard data sets (maximum values are denoted by bold face)

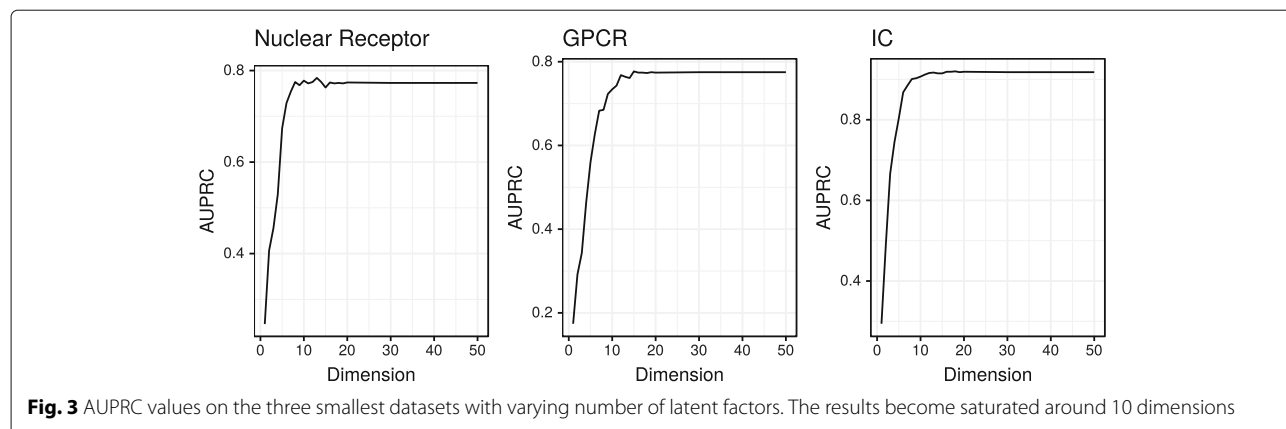
	VB-MK-LMF	NRLMF	KBMF
AUROC (CV1)			
Nuclear Receptor	<b>0.957</b> ± 0.010	0.949 ± 0.011	0.860 ± 0.024
GPCR	<b>0.976</b> ± 0.003	0.960 ± 0.004	0.911 ± 0.004
Ion Channel	<b>0.989</b> ± 0.001	0.984 ± 0.002	0.941 ± 0.003
Enzyme	<b>0.987</b> ± 0.001	0.976 ± 0.002	0.887 ± 0.003
Kinase	<b>0.921</b> ± 0.002	0.919 ± 0.001	0.916 ± 0.001
AUPRC (CV1)			
Nuclear Receptor	<b>0.773</b> ± 0.030	0.723 ± 0.042	0.533 ± 0.047
GPCR	<b>0.777</b> ± 0.016	0.703 ± 0.023	0.541 ± 0.012
Ion Channel	<b>0.916</b> ± 0.007	0.863 ± 0.012	0.763 ± 0.009
Enzyme	<b>0.890</b> ± 0.006	0.876 ± 0.007	0.656 ± 0.008
Kinase	<b>0.850</b> ± 0.003	0.845 ± 0.003	0.844 ± 0.003
AUROC (CV2)			
Nuclear Receptor	<b>0.939</b> ± 0.021	0.896 ± 0.023	0.845 ± 0.023
GPCR	0.878 ± 0.014	<b>0.883</b> ± 0.012	0.847 ± 0.018
Ion Channel	<b>0.812</b> ± 0.026	0.800 ± 0.026	0.785 ± 0.021
Enzyme	<b>0.851</b> ± 0.021	0.811 ± 0.024	0.718 ± 0.028
Kinase	<b>0.894</b> ± 0.004	0.891 ± 0.004	0.838 ± 0.004
AUPRC (CV2)			
Nuclear Receptor	<b>0.593</b> ± 0.058	0.547 ± 0.053	0.447 ± 0.048
GPCR	<b>0.368</b> ± 0.023	0.363 ± 0.023	0.365 ± 0.024
Ion Channel	<b>0.345</b> ± 0.035	0.343 ± 0.033	0.287 ± 0.035
Enzyme	0.349 ± 0.042	<b>0.360</b> ± 0.041	0.269 ± 0.037
Kinase	<b>0.803</b> ± 0.009	0.797 ± 0.010	0.735 ± 0.009
AUROC (CV3)			
Nuclear Receptor	<b>0.917</b> ± 0.026	0.847 ± 0.029	0.735 ± 0.050
GPCR	<b>0.941</b> ± 0.009	0.920 ± 0.014	0.839 ± 0.020
Ion Channel	<b>0.966</b> ± 0.007	0.958 ± 0.008	0.911 ± 0.012
Enzyme	<b>0.962</b> ± 0.005	0.947 ± 0.006	0.859 ± 0.012
Kinase	<b>0.767</b> ± 0.018	0.763 ± 0.018	0.740 ± 0.022
AUPRC (CV3)			
Nuclear Receptor	<b>0.601</b> ± 0.081	0.456 ± 0.079	0.352 ± 0.070
GPCR	<b>0.596</b> ± 0.040	0.553 ± 0.040	0.437 ± 0.047
Ion Channel	<b>0.826</b> ± 0.021	0.788 ± 0.028	0.695 ± 0.024
Enzyme	0.794 ± 0.017	<b>0.808</b> ± 0.018	0.573 ± 0.028
Kinase	<b>0.608</b> ± 0.039	0.597 ± 0.038	0.594 ± 0.039

CV indicates the cross-validation setting (pairwise, drug and target, respectively). AUROC and AUPRC values were averaged over 5 × 10 runs and 95% confidence intervals were computed. In most cases, VB-MK-LMF significantly outperforms the other methods using *t*-test

unit-diagonal positive definite kernel matrix with random values. In the first four datasets, the algorithm assigned more or less uniform weights to the real kernels and a lower one to the random kernel. In the Kinase dataset, the random kernel is almost zeroed out. This underlines the validity of VB-MK-LMF's kernel combination scheme.

Setting  $L$  to  $I$  (the rank of the kernels) yields an almost zero weight to the random kernel, i.e. allowing larger dimensions also allows sufficient separation of the latent representations, which makes spotting kernels with erroneous values easier for the algorithm. This property might also justify increasing the number of latent factors beyond





**Fig. 3** AUPRC values on the three smallest datasets with varying number of latent factors. The results become saturated around 10 dimensions

**Table 2** Multiple Kernel AUPRC values on gold standard data sets in the pairwise cross-validation setting (maximum values are denoted by bold face (maximum values are denoted by bold face))

Neighbors	MrgRbf	MrgTan	McsRbf	McsTan	Orig	All
Nuclear Receptor (KBMF-MKL: 0.566, KronRLS-MKL: 0.522)						
2	0.749	0.758	0.742	0.735	0.754	<b>0.779</b>
3	0.744	0.771	0.761	0.734	0.773	0.775
5	0.732	0.757	0.739	0.724	0.755	0.756
2+3	0.750	0.765	0.754	0.736	0.757	0.758
2+3+5	0.760	0.765	0.740	0.738	0.764	0.760
GPCR (KBMF-MKL: 0.622, KronRLS-MKL: 0.696)						
2	0.743	0.759	0.754	0.762	0.764	0.793
3	0.755	0.774	0.772	0.780	0.777	<b>0.802</b>
5	0.762	0.787	0.782	0.783	0.787	0.796
2+3	0.763	0.782	0.781	0.786	0.785	<b>0.802</b>
2+3+5	0.777	0.798	0.793	0.789	0.796	0.800
Ion Channel (KBMF-MKL: 0.826, KronRLS-MKL: 0.885)						
2	0.909	0.911	0.910	0.911	0.910	0.909
3	0.911	0.914	0.915	0.914	0.912	0.916
5	0.915	0.914	0.913	0.916	0.916	<b>0.917</b>
2+3	0.912	0.914	0.916	0.914	0.913	0.909
2+3+5	0.912	0.915	0.915	0.915	0.916	0.906
Enzyme (KBMF-MKL: 0.704, KronRLS-MKL: 0.893)						
2	0.885	0.887	0.879	0.883	0.888	0.884
3	0.885	0.890	0.885	0.882	0.890	<b>0.895</b>
5	0.883	0.886	0.880	0.881	0.884	0.883
2+3	0.888	0.889	0.880	0.881	0.888	0.881
2+3+5	0.887	0.889	0.881	0.878	0.888	0.875
Kinase (KBMF-MKL: 0.846, KronRLS-MKL: 0.561)						
Neighbors	-	-	2D	3D	ECFP	All
2	-	-	0.850	0.849	0.849	0.850
3	-	-	0.850	0.848	0.850	0.851
5	-	-	0.850	0.849	0.850	0.851
2+3	-	-	0.850	0.850	0.850	0.853
2+3+5	-	-	0.851	0.851	0.850	<b>0.854</b>

The table headers indicate the best AUPRC values obtained using the KBMF-MKL and KronRLS-MKL tools, utilizing all kernels and a grid search method for parameterization. The table bodies show AUPRC values from the VB-MK-LMF method in a cumulative manner. In particular, rows correspond to the cut-off value of the number of closest neighbors and the combinations of the resulting truncated kernels. Columns correspond to individual kernels. The last column was obtained by combining all kernels

**Table 3** Normalized kernel weights with an extra positive definite, unit-diagonal, random valued kernel matrix

	MrgRbf	MrgTan	McsRbf	McsTan	Orig	Random
Nuclear Receptor	0.175	0.176	0.175	0.175	0.175	0.123
GPCR	0.173	0.173	0.172	0.172	0.172	0.138
Ion Channel	0.176	0.176	0.176	0.176	0.176	0.120
Enzyme	0.176	0.176	0.176	0.176	0.176	0.119
		-	2D	3D	ECFP	Random
Kinase		-	0.300	0.283	0.398	0.019

The number of latent factors was not altered in this experiment. Setting the number of latent factors to  $l$  (the rank of the kernel matrix) zeroes out the weight of the random kernel

the rank of the interaction matrix in the multi-kernel setting.

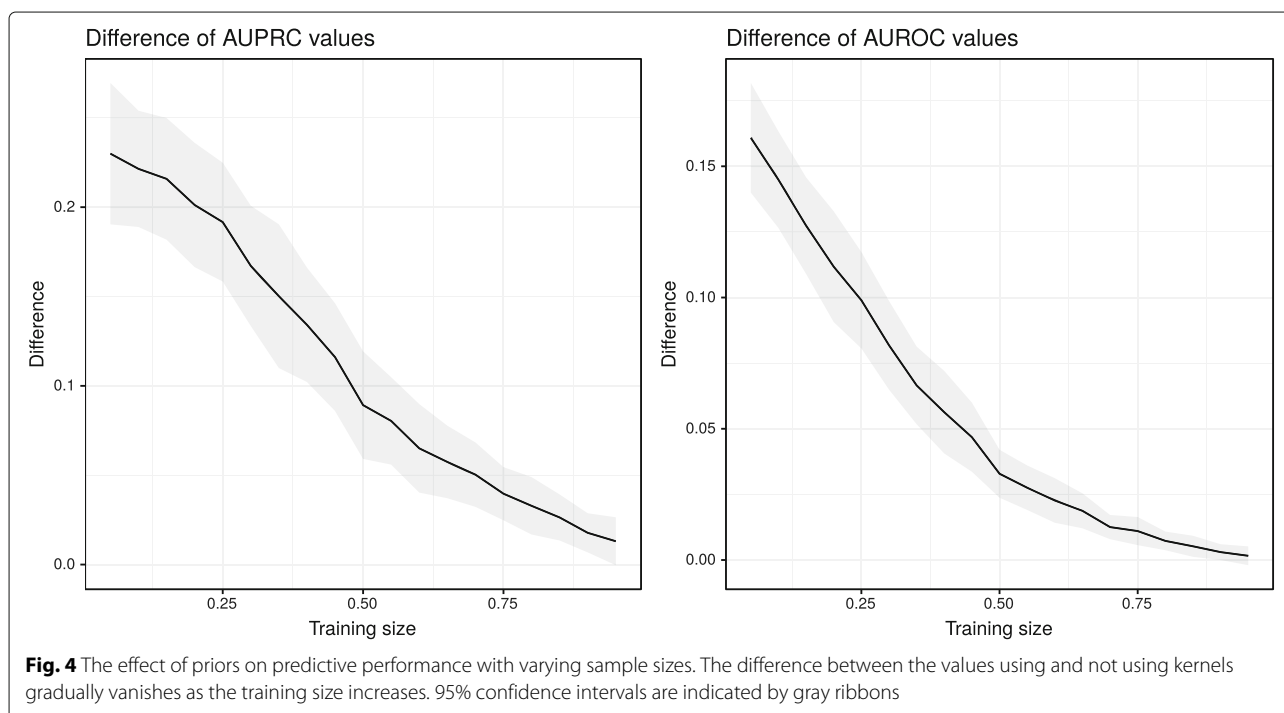
To understand the effect of priors behind the significantly improved performance, which is especially pronounced at smaller sample sizes, we investigated the difference in AUPRC and AUROC values while using and ignoring kernels, at varying training set sizes. The results suggest the existence of a “small sample size” region where using side information offer significant gains, and after which the effect of priors gradually vanishes. Figure 4 depicts the learning curves.


## Discussion

VB-MK-LMF introduces a matrix factorization model incorporating multiple kernel learning, Laplacian regularization and the explicit modeling of interaction probabilities, for which a variational Bayesian inference method

is proposed. The algorithm maps each drug and target into a joint vector space and interaction probabilities are derived from the inner products of the latent representations. Despite the suggested applicability of the unified “pharmacological space” [7], its semantics is still unexplored (for an early application in a ligand-receptor space, see [95], for a proof-of-concept illustration, see [22]). To facilitate a deeper understanding, we provide visual analytics tools alongside the factorization algorithm and allow arbitrary annotations to be mapped onto the latent representations.

We demonstrate this on the Ion Channel dataset. Using  $L = 2$ , the resulting latent representations can be visualized in a 2D Cartesian coordinate system as shown in Fig. 5. Drugs are colored on the basis of their respective ATC classes, where only the classes with more than 5 members were used. Targets are colored according to



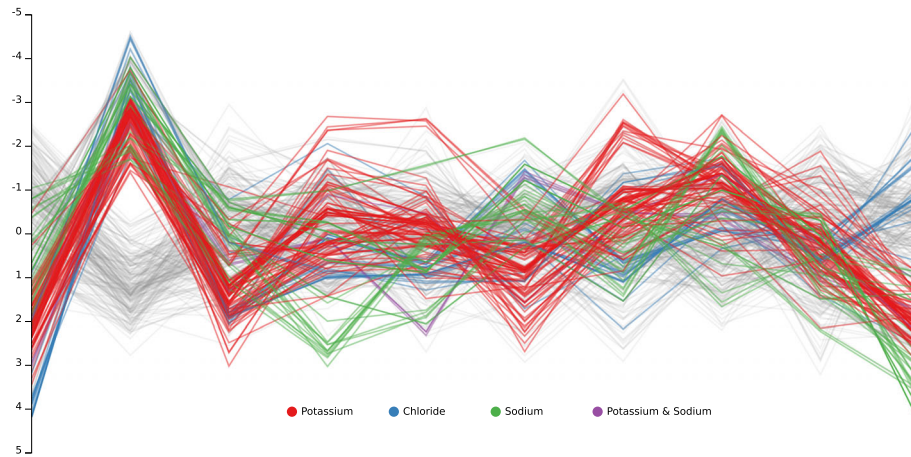


their ion transporter activity as obtained from the Gene Ontology. Known interactions are represented as edges. Even in this low-dimensional case, drugs in the same class tend to cluster together. The only exception is the “Other antiepileptics” class, which is easily explained by its heterogeneity, also indicated by the name. Targets also cluster fairly nicely, albeit with somewhat more outliers. It can be also observed that the targets exhibiting potassium and sodium transporter activity are placed halfway between the sodium and potassium groups.

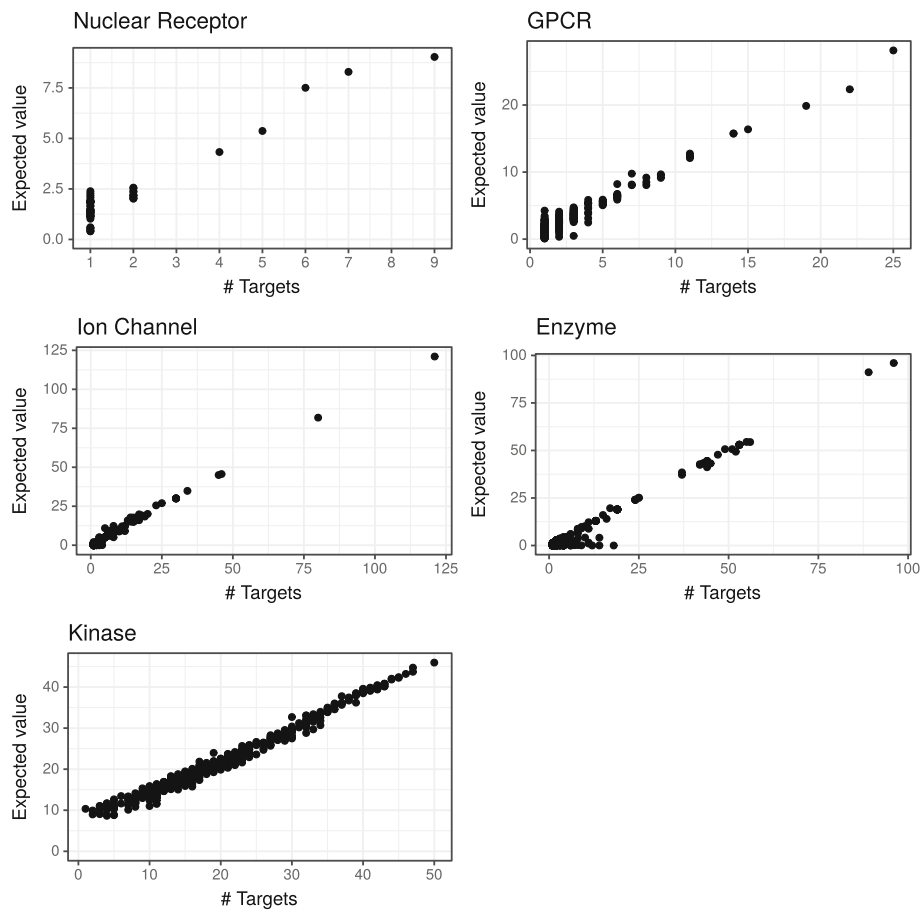
Similarly, Fig. 6 depicts the joint space using a parallel coordinates visualization with  $L = 10$ , where ion transporter activity is denoted by different colors. Most of the dimensions tend to separate at least one class from the others and many of them seem to distinguish between more than two classes. This indicates that the algorithm manages to find biologically meaningful latent dimensions, possibly encoding pharmacophore properties and

the properties of binding sites, but we leave it for further exploration.

From a more practical viewpoint, it is important to touch on the issue of drug promiscuity and polypharmacology. This refers to the observation that some drugs tend to act on multiple targets leading to distinct pharmacological effects, which is often considered an undesirable property [86], although partly unavoidable and potentially utilizable [115]. In either case, predicting the expected number of interactions in a restricted set of targets is a unique property of probabilistic DTI predictors, e.g. compared to ranking approaches. To illustrate this ability of VB-MK-LMF, we computed the expected value of the total number of interactions for every drug in all datasets, treating them independently, shown in Fig. 7 together with the number known targets. Overall, the expected value of further hits approximates the number of interactions already discovered rather closely, although



**Fig. 6** Parallel coordinates visualization of 10 latent dimensions in the Ion Channel dataset. Each curve corresponds to a latent representation of a drug or a target. Targets are colored on the basis of their ion transporter activity




**Fig. 7** Drug promiscuity vs. the expected number of interactions. The number of targets of each drug in the datasets are depicted on the horizontal axis. The expected number of interactions as predicted by VB-MK-LMF are depicted on the vertical axis

it tends to over-estimate, especially when only one or two interactions are known. We also conducted a  $10\times$  cross-validation experiment for each drug in the GPCR dataset and performed the same comparison with similar results (Fig. 8). It is worth to mention that the number of currently unobserved positive interactions in large-scale settings and in comprehensive DTI repositories is vital for the pharmaceutical industry and an open scientific question, as indicated by research on drug-likeness and druggability. Assuming total independence, the expected value provides a raw estimate for this. However, as the relative frequency of positive interactions among the unobserved cases should influence the selection of weight for the observed cases ( $c$ ), and the value of  $c$  influences the expected value, resolving this circular situation and tuning  $c$  requires further investigations.

We also performed a case-based evaluation by obtaining the top 5 novel predictions in the incomplete datasets and examining whether they are present in the current version of the DrugBank database. Most interactions were confirmed and some of the unconfirmed hits are known

to bind to other members of that particular protein family. This shows the ability of VB-MK-LMF to predict novel interactions. The predicted lists are similar to those of the NRLMF method. Table 4 illustrates these results and also contains the rank of the predicted interactions among the NRLMF predictions.

Finally, we discuss computational issues. Due to the explicit computation of inverse matrices, the variational approximation is highly compute-intensive, however, it is straightforward to parallelize and many steps can be written as BLAS operations. GPUs are particularly well-suited for this task. All computations presented in this work can be performed on a mid-range graphics card. Figure 9 shows the runtime of GPU and CPU implementations in terms of latent factors  $200 \times 200$  matrix factorization task, which showed a  $30\times$  speedup using an NVIDIA Titan X graphics card. However, in larger dimensions or with many latent factors, one can quickly run out of GPU memory, i.e. scaling remains an open question. Although GPUs provide excellent performance with single precision, double precision performance typically lags



**Fig. 8** Expected number of interactions as predicted by VB-MK-LMF for each drug in the GPCR dataset. The number of targets are depicted on the horizontal axis. A  $10\times$  cross-validation setting was used

**Table 4** Top 5 predicted interactions which are not present in the datasets

Probability	Drug	Target	Drug name	Target name	DrugBank	NRLMF
Nuclear Receptor						
0.943	D00316	hsa6096	Etretinate	RARB	Yes	1
0.671	D01132	hsa6097	Tazarotene	RARC	<sup>a</sup> RARB	6
0.662	D01132	hsa190	Tazarotene	NROB1		18
0.529	D00898	hsa2100	Dienestrol	ESR2	Yes	7
0.445	D00094	hsa6095	Tretinoin	RARA	Yes	26
GPCR						
0.966	D00283	hsa1814	Clozapine	DRD3	Yes	1
0.956	D00110	hsa1813	Cocaine	DRD2		188
0.938	D02358	hsa154	Metoprolol	ADRB2	Yes	2
0.937	D02614	hsa154	Denopamine	ADRB2	Yes	4
0.937	D04625	hsa154	Isoetharine	ADRB2	Yes	3
Ion Channel						
0.990	D00538	hsa6331	Zonisamide	SCN5A	Yes	9
0.986	D00294	hsa3767	Diazoxide	KCNJ11	Yes	244
0.985	D00552	hsa6331	Tetracaine	SCN5A	Yes	5
0.983	D00438	hsa779	Nimodipine	CACNA1S	Yes	2
0.983	D00649	hsa8911	Amiloride	CACNA1I		83
Enzyme						
0.999	D00542	hsa1571	Halothane	CYP2E1	Yes	1
0.995	D00097	hsa5743	Salicylic acid	PTGS2	Yes	4
0.995	D00437	hsa1559	Nifedipine	CYP2C9	Yes	5
0.987	D00501	hsa50940	Pentoxifylline	PDE11A	<sup>a</sup> PDE5A	2
0.986	D00501	hsa5150	Pentoxifylline	PDE7A	<sup>a</sup> PDE5A	3

Many of the hits were confirmed by the current version of DrugBank. The <sup>a</sup>symbol indicates a known interaction with another member of the protein family. The last column denotes the rank of the interaction among the NRLMF predictions

far behind, especially with modern consumer-level graphics cards. This raises the issue of numerical stability. To cope with the memory footprint of the algorithm, we provide a sparse implementation beside the standard dense solver. To address the issue of numerical stability, we also provide a QR factorization-based implementation which is more stable but significantly slower than the default Cholesky-based method. The computation in VB-MK-LMF is dominated by the inversion in Eq. 6, which gives  $\mathcal{O}(DL^3 \max(I^3, J^3))$  for the total time complexity ( $D$  is the number of iterations). Comparison with the time complexity of NRLMF,  $\mathcal{O}(DLIJ)$ , clearly shows the burden of Bayesian computation in the current implementation and calls for the usage of approximative inversion techniques, which we consider as a future work.


## Conclusion

We presented Variational Bayesian Multiple Kernel Logistic Matrix Factorization (VB-MK-LMF), integrating multiple kernel learning, weighted observations, graph

Laplacian regularization, and explicit modeling of probabilities of binary drug-target interactions. Compared to other state-of-the-art methods, VB-MK-LMF achieved significantly better predictive performance in standard benchmarks.

Admittedly, benchmarking the pure predictive performance on a given dataset gives a very focused view about the real-world applicability of the methods, but helps comparability. On the other hand, the release of new and updated datasets as shown in Additional file 1 in fact quickly create an impractical fragmentary situation. In general, the definition of a standard background knowledge pool for a benchmarking is even more complicated, as earlier attempts show in computational fusion methods for gene prioritization [116, 117].

Additionally, currently the possible utilizations of a DTI prediction method in real-world applications are at least as diverse as the methodological repertoire. For example, DTI prediction methods could be applied in data



**Fig. 9** Runtime of the GPU and CPU implementations in terms of the number of latent factors. This benchmark was conducted on a  $200 \times 200$  matrix factorization. The GPU implementation brings a 30 $\times$  speedup on an NVIDIA GTX Titan X graphics card

quality control phase for anomaly detection, especially in the case of merging different bioactivity values from public and private sources. Screening design, hit triage and prioritization for further validation [118], possibly in an active learning framework [16, 119], are standard usages. Finally, DTI prediction methods may also provide essential data to support visualization and visual data analytics, as we demonstrated in a new range of dimensionality (10 – 20), which proved to be sufficient with VB-MK-LMF.

Another key property of VB-MK-LMF is the explicit modeling of probabilities, which allows the prediction of interaction probabilities and their credibility. We demonstrated the use of probabilistic predictions by proposing DTI dataset specific versions of promiscuity and druggability, through the expected number of hits in a dataset for a drug or a target respectively. In general, the predicted posteriors for the interactions can be seen as a probabilistic “data-analytic” knowledge base, which allows new functionalities in post-processing, beyond enrichment methods available for ranking methods [33, 37]. To utilize the Bayesian predictions of VB-MK-LMF, we also plan to investigate their decision theoretic usage, when certainty for expected gains and losses of prioritization of interactions is expected, e.g. in functional validations.

Further interesting research directions are the regression version of VB-MK-LMF directly approximating the continuous activity data [8, 52] and the use of multiple instances of VB-MK-LMF for overlapping DTI matrices, which are linked to each other by weighted common

observations. The latter could improve the scalability of the method using parallel implementations for mid-sized DTI tasks with  $10^5$  drugs and  $10^4$  targets, going beyond the current benchmarks.

## Additional files

**Additional file 1:** The properties of DTI methods related to the development or evaluation of VB-MK-LMF. (PDF 124 kb)

**Additional file 2:** Derivation of the lower bound using Jaakkola’s bound on the logistic sigmoid. (PDF 107 kb)

## Abbreviations

ADME: Absorption, distribution, metabolism, and excretion; AUPRC: Area under the precision-recall curve; AUROC: Area under the receiver operating characteristic curve; CPI: Compound-protein interaction; CVS: Cross-validation setting; DTI: Drug-target interaction; FDA: Food and drug administration; GPCR: G-protein coupled receptor; GP-GPU: General purpose computing on graphics processing unit; KBMF: Kernelized Bayesian matrix factorization; KRM: Kernel regression-based method; MACCS: Molecular access system; MF: Matrix factorization; MKL: Multiple Kernel learning; MLP: Multi-layer perceptron; NRLMF: Neighborhood regularized logistic matrix factorization; RLS-KF: Regularized least squares kernel fusion; SVD: Singular value decomposition; SVM: Support vector machine; VB-MK-LMF: Variational Bayesian multiple kernel logistic matrix factorization

## Acknowledgements

Not applicable.

## Funding

This work was supported by the ÚNKP-16-3-III. New National Excellence Program of the Ministry of Human Capacities (BB), OTKA 119866 (PA) and the János Bolyai Research Scholarship (PA).

## Availability of data and materials

The code and data used in the current study are available at <http://bioinformatics.mit.bme.hu/VB-MK-LMF/>.

**Authors' contributions**

BB and AP designed the experiments. BB developed the software and performed the experiments. BB and AP analyzed the data and wrote the paper. Both authors read and approved the final manuscript.

**Ethics approval and consent to participate**

Not applicable.

**Consent for publication**

Not applicable.

**Competing interests**

The authors declare that they have no competing interests.

**Publisher's Note**

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 30 June 2017 Accepted: 21 September 2017

Published online: 04 October 2017

**References**

- Williams AJ, Ekins S, Tkachenko V. Towards a gold standard: Regarding quality in public domain chemistry databases and approaches to improving the situation. *Drug Discov Today*. 2012;17(13-14):685–701. doi:10.1016/j.drudis.2012.02.013.
- Goldmann D, Montanari F, Richter L, Zdražil B, Ecker GF. Exploiting open data: a new era in pharmacoinformatics. *Future Med Chem*. 2014;6(5):503–14. doi:10.4155/fmc.14.13.
- Chen X, Yan CC, Zhang X, Zhang X, Dai F, Yin J, Zhang Y. Drug-target interaction prediction: Databases, web servers and computational models. *Brief Bioinform*. 2016;17(4):696–712. doi:10.1093/bib/bbv066.
- Zheng W, Thorne N, McKew JC. Phenotypic screens as a renewed approach for drug discovery. *Drug Discov Today*. 2013;18(21):1067–73.
- Orchard S, Al-Lazikani B, Bryant S, Clark D, Calder E, Dix I, Engkvist O, Forster M, Gaulton A, Gilson M, Glen R, Grigorov M, Hammond-Kosack K, Harland L, Hopkins A, Larmine C, Lynch N, Mann RK, Murray-Rust P, Lo Piparo E, Southan C, Steinbeck C, Wishart D, Hermjakob H, Overington J, Thornton J. Minimum information about a bioactive entity (MIABE). *Nat Rev Drug Discov*. 2011;10(9):661–9. doi:10.1038/nrd3503.
- Samwald M, Jentzsch A, Bouton C, Kallesøe CS, Willighagen E, Hajagos J, Scott Marshall M, Prud'hommeaux E, Hassanzadeh O, Pichler E, Stephens S. Linked Open drug data for pharmaceutical research and development. *J Cheminformatics*. 2011;3(5):19. doi:10.1186/1758-2946-3-19.
- Yamanishi Y, Araki M, Gutteridge A, Honda W, Kanehisa M. Prediction of drug-target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*. 2008;24(13):232–40. doi:10.1093/bioinformatics/btn162.
- Pahikkala T, Airola A, Pietilä S, Shakyawar S, Szwajda A, Tang J, Aittokallio T. Toward more realistic drug-target interaction predictions. *Brief Bioinform*. 2015;16(2):325–37. doi:10.1093/bib/bbu010.
- Davis MI, Hunt JP, Herrgard S, Ciceri P, Wodicka LM, Pallares G, Hocker M, Treiber DK, Zarrinkar PP. Comprehensive analysis of kinase inhibitor selectivity. *Nat Biotechnol*. 2011;29(11):1046–51. doi:10.1038/nbt.1990.0402594v3.
- Schomburg I, Chang A, Placzek S, Söhngen C, Rother M, Lang M, Munaretto C, Ulas S, Stelzer M, Grote A, Scheer M, Schomburg D. BRENDA in 2013: Integrated reactions, kinetic data, enzyme function data, improved disease classification: new options and contents in BRENDA. *Nucleic Acids Res*. 2013;41(D1):1–9. doi:10.1093/nar/gks1049.
- Lindh M, Svensson F, Schaal W, Zhang J, Sköld C, Brandt P, Karlén A. Toward a benchmarking data set able to evaluate ligand- and structure-based virtual screening using public HTS data. *J Chem Inf Model*. 2015;55(2):343–53. doi:10.1021/ci5005465.
- Mervin LH, Afzal AM, Drakakis G, Lewis R, Engkvist O, Bender A. Target prediction utilising negative bioactivity data covering large chemical space. *J Cheminformatics*. 2015;7(1):1–16. doi:10.1186/s13321-015-0098-y.
- Liu C, Su J, Yang F, Wei K, Ma J, Zhou X. Compound signature detection on LINCS L1000 big data. *Mol BioSyst*. 2015;11(3):714–22. doi:10.1039/C4MB00677A.
- Kövesdi I, Dominguez-Rodriguez MF, Örfi L, Náray-Szabó G, Varró A, Papp JG, Matyus P. Application of neural networks in structure–activity relationships. *Med Res Rev*. 1999;19(3):249–69.
- Burbidge R, Trotter M, Buxton B, Holden S. Drug design by machine learning: support vector machines for pharmaceutical data analysis. *Comput Chem*. 2001;26(1):5–14.
- Warmuth MK, Liao J, Rätsch G, Mathieson M, Putta S, Lemmen C. Active learning with support vector machines in the drug discovery process. *J Chem Inf Comput Sci*. 2003;43(2):667–73.
- Willett P, Barnard JM, Downs GM. Chemical similarity searching. *J Chem Inf Comput Sci*. 1998;38(6):983–96.
- Ginn CM, Willett P, Bradshaw J. Combination of molecular similarity measures using data fusion. In: *Virtual Screening: An Alternative or Complement to High Throughput Screening?* Netherlands: Springer; 2000. p. 1–16.
- Ding H, Takigawa I, Mamitsuka H, Zhu S. Similarity-based machine learning methods for predicting drug-target interactions: a brief review. *Brief Bioinform*. 2013056. doi:10.1093/bib/bbt056.
- Kitchen DB, Decornez H, Furr JR, Bajorath J. Docking and scoring in virtual screening for drug discovery: methods and applications. *Nat Rev Drug Discov*. 2004;3(11):935–49.
- Sousa SF, Fernandes PA, Ramos MJ. Protein–ligand docking: current status and future challenges. *Proteins Struct Funct Bioinform*. 2006;65(1):15–26.
- Gönen M. Predicting drug–target interactions from chemical and genomic kernels using bayesian matrix factorization. *Bioinformatics*. 2012;28(18):2304–310.
- Zheng X, Ding H, Mamitsuka H, Zhu S. Collaborative matrix factorization with multiple similarities for predicting drug-target interactions. In: *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '13*. Chicago; 2013. p. 1025. doi:10.1145/2487575.2487670.
- Waller CL, Shah A, Nolte M. Strategies to support drug discovery through integration of systems and data. *Drug Discov Today*. 2007;12(15):634–9.
- Muresan S, Petrov P, Southan C, Kjellberg MJ, Kogej T, Tyrchan C, Varkonyi P, Xie PH. Making every SAR point count: The development of Chemistry Connect for the large-scale integration of structure and bioactivity data. *Drug Discov Today*. 2011;16(23-24):1019–1030. doi:10.1016/j.drudis.2011.10.005.
- Agrafiotis DK, Alex S, Dai H, Derkinderen A, Farnum M, Gates P, Izrailev S, Jaeger EP, Konstant P, Leung A, Lobanov VS, Marichal P, Martin D, Rassokhin DN, Shemanarev M, Skalkin A, Stong J, Tabruyn T, Vermeiren M, Wan J, Xu XY, Yao X. Advanced Biological and Chemical Discovery (ABCD): Centralizing discovery knowledge in an inherently decentralized world. *J Chem Inf Model*. 2007;47(6):1999–2014. doi:10.1021/ci700267w.
- Gönen M, Khan S, Kaski S. Kernelized bayesian matrix factorization. In: *International Conference on Machine Learning*. Atlanta; 2013. p. 864–72.
- Cheng F, Liu C, Jiang J, Lu W, Li W, Liu G, Zhou W, Huang J, Tang Y. Prediction of drug-target interactions and drug repositioning via network-based inference. *PLoS Comput Biol*. 2012;8(5):doi:10.1371/journal.pcbi.1002503.
- Fu G, Ding Y, Seal A, Chen B, Sun Y, Bolton E. Predicting drug target interactions using meta-path-based semantic network analysis. *BMC Bioinformatics*. 2016;17(1):160.
- Ashburn TT, Thor KB. Drug repositioning: identifying and developing new uses for existing drugs. *Nat Rev Drug Discov*. 2004;3(8):673–83.
- Li J, Zheng S, Chen B, Butte AJ, Swamidass SJ, Lu Z. A survey of current trends in computational drug repositioning. *Brief Bioinform*. 2016;17(1):2–12.
- Arany A, Bolgár B, Balogh B, Antal P, Matyus P. Multi-aspect candidates for repositioning: data fusion methods using heterogeneous information sources. *Curr Med Chem*. 2013;20(1):95–107.
- Temesi G, Bolgár B, Arany Á, Szalai C, Antal P, Matyus P. Early repositioning through compound set enrichment analysis: a knowledge-recycling strategy. *Future Med Chem*. 2014;6(5):563–75.
- Liu Z, Guo F, Gu J, Wang Y, Li Y, Wang D, Lu L, Li D, He F. Similarity-based prediction for anatomical therapeutic chemical classification of drugs by integrating multiple data sources. *Bioinformatics*. 2015;31(11):1788–95.



35. Bleakley K, Yamanishi Y. Supervised prediction of drug-target interactions using bipartite local models. *Bioinformatics*. 2009;25(18):2397–403. doi:10.1093/bioinformatics/btp433.
36. Xia Z, Wu LY, Zhou X, Wong STC. Semi-supervised drug-protein interaction prediction from heterogeneous biological spaces. *BMC Syst Biol*. 2010;4(56):6. doi:10.1186/1752-0509-4-S2-S6.
37. Agarwal S, Dugar D, Sengupta S. Ranking chemical structures for drug discovery: A new machine learning approach. *J Chem Inf Model*. 2010;50(5):716–31. doi:10.1021/ci9003865.
38. van Laarhoven T, Nabuurs SB, Marchiori E. Gaussian interaction profile kernels for predicting drug-target interaction. *Bioinformatics*. 2011;27(21):3036–43. doi:10.1093/bioinformatics/btr500.
39. Perlman L, Gottlieb A, Atias N, Ruppín E, Sharan R. Combining Drug and Gene Similarity Measures for Drug-Target Elucidation. *Comput Biol*. 2011;18(2):133–45. doi:10.1089/cmb.2010.0213.
40. Chen B, Ding Y, Wild DJ. Improving integrative searching of systems chemical biology data using semantic annotation. *J Cheminformatics*. 2012;4(1):6. doi:10.1186/1758-2946-4-6.
41. Yu H, Chen J, Xu X, Li Y, Zhao H, Fang Y, Li X, Zhou W, Wang W, Wang Y. A systematic prediction of multiple drug-target interactions from chemical, genomic, and pharmacological data. *PLoS ONE*. 2012;7(5):. doi:10.1371/journal.pone.0037608.
42. Mei JP, Kwok CK, Yang P, Li XL, Zheng J. Drug-target interaction prediction by learning from local information and neighbors. *Bioinformatics*. 2013;29(2):238–45. doi:10.1093/bioinformatics/bts670.
43. van Laarhoven T, Marchiori E. Predicting drug-target interactions for new drug compounds using a weighted nearest neighbor profile. *PLoS ONE*. 2013;8(6):1–6. doi:10.1371/journal.pone.0066952.
44. Zheng W, Thorne N, McKew JC. Phenotypic screens as a renewed approach for drug discovery. *Drug Discov Today*. 2013;18(21–22):1067–73. doi:10.1016/j.drudis.2013.07.001.
45. Wang Y, Zeng J. Predicting drug-target interactions using restricted Boltzmann machines. *Bioinformatics*. 2013;29(13):126–34. doi:10.1093/bioinformatics/btt234.
46. Simm J, Arany A, Zakeri P, Haber T, Wegner JK, Chupakhin V, Ceulemans H, Moreau Y. Macau: Scalable Bayesian Multi-relational Factorization with Side Information using MCMC. In: Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing. Roppongi: IEEE; 2017.
47. Yuan Q, Gao J, Wu D, Zhang S, Mamitsuka H, Zhu S. DrugE-Rank: Improving drug-target interaction prediction of new candidate drugs or targets by ensemble learning to rank. *Bioinformatics*. 2016;32(12):18–27. doi:10.1093/bioinformatics/btw244.
48. Liu Y, Wu M, Miao C, Zhao P, Li XL. Neighborhood Regularized Logistic Matrix Factorization for Drug-Target Interaction Prediction. *PLoS Comput Biol*. 2016;12(2):1–26. doi:10.1371/journal.pcbi.1004760.
49. Hao M, Bryant SH, Wang Y, Iorio F, Rittman T, Ge H, Menden M, Saez-Rodriguez J, Bartlett JB, Dredge K, Dalglish AG, Steinbach G, Koehl GE, Schlittl HJ, Geissler EK, Cappelli C, Gu S, Keiser MJ, Wang L, Haupt VJ, Schroeder M, Ma DL, Chan DS, Leung CH, Yamanishi Y, Araki M, Gutteridge A, Honda W, Kanehisa M, Bleakley K, Yamanishi Y, van Laarhoven T, Nabuurs SB, Marchiori E, Mei JP, Kwok CK, Yang P, Li XL, Zheng J, Hao M, Wang Y, Bryant SH, Wang B, Liu Y, Wu M, Miao C, Zhao P, Li XL, Kanehisa M, Schomburg I, Günther S, Wishart DS, Kuang Q, Smith TF, Waterman MS, Hattori M, Okuno Y, Goto S, Kanehisa M, Ma H, King J, Lyu MR, Duchí J, Hazan E, Singer Y, Gonen M, Kaski S, Cao Y, Charisi A, Cheng LC, Jiang T, Girke T, Guha R, Sievers F, Leslie C, Eskin E, Noble WS, Langham JJ, Cleves AE, Spitzer R, Kirshner D, Jain AN, Collins I, von Coburg Y, Kottke T, Weizel L, Ligneau X, Stark H, Wishart D, Alaimo S, Sui J. Predicting drug-target interactions by dual-network integrated logistic matrix factorization. *Sci Rep*. 2017;7:40376. doi:10.1038/srep40376.
50. Hao M, Wang Y, Bryant SH. Improved prediction of drug-target interactions using regularized least squares integrating with kernel fusion technique. *Analytica Chimica Acta*. 2016;909:41–50. doi:10.1016/j.aca.2016.01.014.
51. Nascimento ACA, Prudêncio RBC, Costa IG. A multiple kernel learning algorithm for drug-target interaction prediction. *BMC Bioinformatics*. 2016;17(1):46. doi:10.1186/s12859-016-0890-3.
52. Bolgár B, Antal P. Bayesian matrix factorization with non-random missing data using informative Gaussian process priors and soft evidences In: Antonucci A, Corani G, Campos CP, editors. Proceedings of the Eighth International Conference on Probabilistic Graphical Models. Lugano: PMLR; 2016. p. 25–36.
53. Wu Z, Cheng F, Li J, Li W, Liu G, Tang Y. SDTNBI: an integrated network and cheminformatics tool for systematic prediction of drug–target interactions and drug repositioning. *Brief Bioinform*. 2016;17(1):1–12. doi:10.1093/bib/bbw012.
54. Keum J, Nam H. Self-blml: Prediction of drug-target interactions via self-training svm. *PLoS ONE*. 2017;12(2):0171839.
55. Visser U, Abeyruwan S, Vempati U, Smith AP, Lemmon V, Schürer SC. BioAssay Ontology (BAO): a semantic description of bioassays and high-throughput screening results. *BMC Bioinformatics*. 2011;12(1):257. doi:10.1186/1471-2105-12-257.
56. Chen B, Dong X, Jiao D, Wang H, Zhu Q, Ding Y, Wild DJ. Chem2Bio2RDF: a semantic framework for linking and data mining chemogenomic and systems chemical biology data. *BMC Bioinformatics*. 2010;11:255. doi:10.1186/1471-2105-11-255.
57. Gaulton A, Hersey A, Nowotka M, Bento AP, Chambers J, Mendez D, Motow P, Atkinson F, Bellis LJ, Cibrián-Uhalte E, et al. The ChEMBL database in 2017. *Nucleic Acids Res*. 2016;45(D1):945–54.
58. Mathias SL, Hines-Kay J, Yang JJ, Zahoransky-Kohalmi G, Bologa CG, Ursu O, Oprea TI. The CARLSBAD database: A federated database of chemical bioactivities. *Database*. 2013;2013:1–8. doi:10.1093/database/bat044.
59. Said A, Bellogín A. Comparative recommender system evaluation: benchmarking recommendation frameworks. In: Proceedings of the 8th ACM Conference on Recommender Systems. Foster City: ACM; 2014. p. 129–36.
60. Tiikkainen P, Bellis L, Light Y, Franke L. Estimating error rates in bioactivity databases. *J Chem Inf Model*. 2013;53(10):2499–505. doi:10.1021/ci400099q.
61. Hersey A, Chambers J, Bellis L, Patrícia Bento A, Gaulton A, Overington JP. Chemical databases: curation or integration by user-defined equivalence?. *Drug Discov Today Technol*. 2015;14:17–24. doi:10.1016/j.ddtec.2015.01.005.
62. Lipinski CA, Litterman NK, Southan C, Williams AJ, Clark AM, Ekins S. Parallel worlds of public and commercial bioactive chemistry data: Miniperspective. *J Med Chem*. 2015;58(5):2068.
63. Southan C, Vrkonyi P, Muresan S. Quantitative assessment of the expanding complementarity between public and commercial databases of bioactive compounds. *J Cheminformatics*. 2009;1(1):1–17. doi:10.1186/1758-2946-1-10.
64. Tiikkainen P, Franke L. Analysis of commercial and public bioactivity databases. *J Chem Inf Model*. 2012;52(2):319–26. doi:10.1021/ci2003126.
65. Hu Y, Bajorath J. Growth of ligand-target interaction data in ChEMBL is associated with increasing and activity measurement-dependent compound promiscuity. *J Chem Inf Model*. 2012;52(10):2550–558. doi:10.1021/ci3003304.
66. Johnson MA, Maggiora GM. Concepts and Applications of Molecular Similarity. New York: Wiley; 1990.
67. Maggiora G, Vogt M, Stumpfe D, Bajorath J. Molecular similarity in medicinal chemistry: miniperspective. *J Med Chem*. 2013;57(8):3186–204.
68. Lipinski CA. Lead-and drug-like compounds: the rule-of-five revolution. *Drug Discov Today Technol*. 2004;1(4):337–41.
69. Tian S, Wang J, Li Y, Li D, Xu L, Hou T. The application of in silico drug-likeness predictions in pharmaceutical research. *Adv Drug Deliv Rev*. 2015;86:2–10.
70. Rask-Andersen M, Masuram S, Schiöth HB. The druggable genome: evaluation of drug targets in clinical trials suggests major shifts in molecular class and indication. *Annu Rev Pharmacol Toxicol*. 2014;54:9–26.
71. Gao M, Skolnick J. A comprehensive survey of small-molecule binding pockets in proteins. *PLoS Comput Biol*. 2013;9(10):1003302.
72. Hopkins AL. Network pharmacology: the next paradigm in drug discovery. *Nat Chem Biol*. 2008;4(11):682–90.
73. Kubinyi H. Similarity and dissimilarity: a medicinal chemist's view. *Perspectives Drug Discov Des*. 1998;9:225–52.
74. Eckert H, Bajorath J. Molecular similarity analysis in virtual screening: foundations, limitations and novel approaches. *Drug Discov Today*. 2007;12(5):225–33.
75. Ding H, Takigawa I, Mamitsuka H, Zhu S. Similarity-based machine learning methods for predicting drug–target interactions: a brief review. *Brief Bioinform*. 2013;15(5):734–47.

76. Gönen M. Predicting drug-target interactions from chemical and genomic kernels using Bayesian matrix factorization. *Bioinformatics*. 2012;28(18):2304–10. doi:10.1093/bioinformatics/bts360.
77. Daina A, Michielin O, Zoete V. Swissadme: a free web tool to evaluate pharmacokinetics, drug-likeness and medicinal chemistry friendliness of small molecules. *Sci Rep*. 2017;7:42717.
78. Hopkins AL. Drug discovery: predicting promiscuity. *Nature*. 2009;462(7270):167–8.
79. Cereto-Massagué A, Guasch L, Valls C, Mulero M, Pujadas G, Garcia-Vallvé S. Decoyfinder: an easy-to-use python gui application for building target-specific decoy sets. *Bioinformatics*. 2012;28(12):1661–2.
80. Hussein HA, Geneix C, Petitjean M, Borrel A, Flatters D, Camproux AC. Global vision of druggability issues: applications and perspectives. *Drug Discov Today*. 2017;22(2):404–415. Elsevier.
81. Jamali AA, Ferdousi R, Razzaghi S, Li J, Safdari R, Ebrahimie E. Drugminer: comparative analysis of machine learning algorithms for prediction of potential druggable proteins. *Drug Discov Today*. 2016;21(5):718–24.
82. Hussein HA, Borrel A, Geneix C, Petitjean M, Regad L, Camproux AC. Pockdrug-server: a new web server for predicting pocket druggability on holo and apo proteins. *Nucleic Acids Res*. 2015;43(W1):W436–W442. Oxford University Press.
83. Chen X, Yan CC, Zhang X, Zhang X, Dai F, Yin J, Zhang Y. Drug–target interaction prediction: databases, web servers and computational models. *Brief Bioinform*. 2015;17(4):696–712.
84. Cheng T, Hao M, Takeda T, Bryant SH, Wang Y. Large-Scale Prediction of Drug-Target Interaction: a Data-Centric Review. *The AAPS Journal*. 2017;1–12. Springer.
85. Lavecchia A. Machine-learning approaches in drug discovery: methods and applications. *Drug Discov Today*. 2014;20(3):318–31. doi:10.1016/j.drudis.2014.10.012.
86. Lounkine E, Keiser MJ, Whitebread S, Mikhailov D, Hamon J, Jenkins JL, Lavan P, Weber E, Doak AK, Côté S, et al. Large-scale prediction and testing of drug activity on side-effect targets. *Nature*. 2012;486(7403):361–7.
87. Jacob L, Vert JP. Protein-ligand interaction prediction: an improved chemogenomics approach. *Bioinformatics*. 2008;24(19):2149–56.
88. Xu Q, Yang Q. A survey of transfer and multitask learning in bioinformatics. *J Comput Sci Eng*. 2011;5(3):257–68.
89. Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis vol. 2*. Boca Raton: Chapman & Hall/CRC; 2014.
90. Nagamine N, Sakakibara Y. Statistical prediction of protein–chemical interactions based on chemical structure and mass spectrometry data. *Bioinformatics*. 2007;23(15):2004–12.
91. van Laarhoven T, Nabuurs SB, Marchiori E. Gaussian interaction profile kernels for predicting drug–target interaction. *Bioinformatics*. 2011;27(21):3036–43. doi:10.1093/bioinformatics/btr500.
92. Wen M, Zhang Z, Niu S, Sha H, Yang R, Yun Y, Lu H. Deep-learning-based drug–target interaction prediction. *J Proteome Res*. 2017;16(4):1401–9.
93. Srebro N, Jaakkola T. Sparse matrix factorization of gene expression data; 2001. Internal report, MIT Artificial Intelligence Laboratory. Available at [www.Ai.Mit.Edu/~research/abstracts/abstracts2001/genomics/01srebro.Pdf](http://www.Ai.Mit.Edu/~research/abstracts/abstracts2001/genomics/01srebro.Pdf).
94. Dueck D, Morris QD, Frey BJ. Multi-way clustering of microarray data using probabilistic sparse matrix factorization. *Bioinformatics*. 2005;21(suppl 1):144–51.
95. Bock JR, Gough DA. A new method to estimate ligand-receptor energetics. *Mol Cell Proteomics*. 2002;1(11):904–10.
96. Agarwal P, Searls DB. Literature mining in support of drug discovery. *Brief Bioinform*. 2008;9(6):479–92.
97. Parsons AB, Lopez A, Givoni IE, Williams DE, Gray CA, Porter J, Chua G, Sopko R, Brost RL, Ho CH, et al. Exploring the mode-of-action of bioactive compounds by chemical-genetic profiling in yeast. *Cell*. 2006;126(3):611–25.
98. Takács G, Pilászy I, Németh B, Tikk D. Matrix factorization and neighbor based algorithms for the netflix prize problem. In: *Proceedings of the 2008 ACM Conference on Recommender Systems*. Lausanne: ACM; 2008. p. 267–74.
99. Srebro N, Jaakkola T, et al. Weighted low-rank approximations. In: *icml*. Washington; 2003. p. 720–7.
100. Pan R, Zhou Y, Cao B, Liu NN, Lukose R, Scholz M, Yang Q. One-class collaborative filtering. In: *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference On*. Pisa: IEEE; 2008. p. 502–11.
101. Salakhutdinov R, Mnih A. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. 2008880–7. doi:10.1145/1390156.1390267.
102. Severinski C, Salakhutdinov R. Bayesian probabilistic matrix factorization: a user frequency analysis. 2014. <http://adsabs.harvard.edu/abs/2014arXiv1407.78405>.
103. Zhou T, Shan H, Banerjee A, Sapiro G. Kernelized probabilistic matrix factorization: Exploiting graphs and side information. In: *SDM*. Anaheim: SIAM / Omnipress; 2012. p. 403–14.
104. Hernandez-Lobato JM, Houlby N, Ghahramani Z. Stochastic Inference for Scalable Probabilistic Modeling of Binary Matrices. In: *Proceedings of the 31st International Conference on Machine Learning (ICML)*; 2014. p. 379–387.
105. Gönen M, Kaski S. Kernelized bayesian matrix factorization. *IEEE Trans Pattern Anal Mach Intell*. 2014;36(10):2047–60.
106. Koutsoukas A, Lowe R, KalantarMotamedi Y, Mussa HY, Klaffke W, Mitchell JB, Glen RC, Bender A. In silico target predictions: defining a benchmarking data set and comparison of performance of the multiclass naïve bayes and parzen-rosenblatt window. *J Chem Inf Model*. 2013;53(8):1957–66.
107. Schomburg KT, Rarey M. Benchmark data sets for structure-based computational target prediction. *J Chem Inf Model*. 2014;54(8):2261–74. doi:10.1021/ci500131x.
108. Wale N, Karypis G. Target fishing for chemical compounds using target-ligand activity data and ranking based methods. *J Chem Inf Model*. 2009;49(10):2190–201. doi:10.1021/ci9000376. NIHMS150003.
109. Peón A, Dang CC, Ballester PJ. How reliable are ligand-centric methods for target fishing?., *Front Chem*. 2016;4(April):15. doi:10.3389/fchem.2016.00015.
110. Landrum G. Rdkit: Open-source cheminformatics. 2006;3(04):2012. Online. <http://www.rdkit.org>. Accessed.
111. Jordan MI, Ghahramani Z, Jaakkola TS, Saul LK. An introduction to variational methods for graphical models. *Machine learning*. 1999;37(2):183–233. Springer.
112. Bishop CM. Pattern recognition. *Mach Learn*. 2006;128:1–58.
113. Jaakkola TS, Jordan MI. Bayesian parameter estimation via variational methods. *Stat Comput*. 2000;10(1):25–37. doi:10.1023/A:1008932416310.
114. Cortes C, Mohri M, Rostamizadeh A. Learning non-linear combinations of kernels. In: *Proceedings of the 22Nd International Conference on Neural Information Processing Systems*. NIPS'09. USA: Curran Associates Inc.; 2009. p. 396–404. <http://dl.acm.org/citation.cfm?id=2984093>. 2984138.
115. Maggiora G, Gokhale V. Non-specificity of drug-target interactions–consequences for drug discovery. In: *Frontiers in Molecular Design and Chemical Information Science–Herman Skolnik Award Symposium 2015*. Jürgen Bajorath. Boston: ACS Publications; 2016. p. 91–142.
116. Börnigen D, Tranchevent LC, Bonachela-Capdevila F, Devriendt K, De Moor B, De Causmaecker P, Moreau Y. An unbiased evaluation of gene prioritization tools. *Bioinformatics*. 2012;28(23):3081–088.
117. Moreau Y, Tranchevent LC. Computational tools for prioritizing candidate genes: boosting disease gene discovery. *Nat Rev Genet*. 2012;13(8):523–36.
118. Paricharak S, Méndez-Lucio O, Chavan Ravindranath A, Bender A, IJzerman AP, van Westen GJP. Data-driven approaches used for compound library design, hit triage and bioactivity modeling in high-throughput screening. *Brief Bioinform*. 2016. In preparation doi:10.1093/bib/bbw105.
119. Cobanoglu MC, Liu C, Hu F, Oltvai ZN, Bahar I. Predicting drug–target interactions using probabilistic matrix factorization. *J Chem Inf Model*. 2013;53(12):3399–409.