Simple ensemble methods are competitive with state-of-the-art data integration methods for gene function prediction

Matteo Ré RE@DSI.UNIMI.IT

Giorgio Valentini VALENTINI@DSI.UNIMI.IT

DSI, Dipartimento di Scienze dell' Informazione,
Università degli Studi di Milano,
Via Comelico 39, 20135 Milano, Italia.

Editor: Saso Dzeroski, Pierre Geurts, and Juho Rousu

Abstract

Several works showed that biomolecular data integration is a key issue to improve the prediction of gene functions. Quite surprisingly only little attention has been devoted to data integration for gene function prediction through ensemble methods. In this work we show that relatively simple ensemble methods are competitive and in some cases are also able to outperform state-of-the-art data integration techniques for gene function prediction.

Keywords: Gene function prediction, ensemble methods, biomolecular data integration

1. Introduction

The availability of an ever increasing amount of data sources due to recent advances in high throughput biotechnologies opens unprecedented opportunities for genome-wide gene function prediction. Indeed several works showed that biomolecular data integration plays an essential role in the prediction of gene functions.

Gene function prediction in its general formulation is a complex classification problem characterized by the following items: a) each gene/gene product can be assigned to multiple terms/classes (a multiclass, multilabel classification problem); b) classes are structured according to a predefined hierarchy: a directed acyclic graph for the Gene Ontology (GO) (The Gene Ontology Consortium, 2000) or a tree forest for FunCat (Ruepp et al., 2004); c) classes are usually unbalanced (with positive examples usually less than negatives); d) known gene labels are in several cases uncertain; e) multiple sources of data can be used to predict gene functions.

In this paper we focus on the last item, considering the problem of the prediction of a subset of FunCat and GO classes in the model organism *S. cerevisiae*.

The main approaches proposed in the literature can be schematically subdivided in three categories: functional linkage networks, vector subspace integration and kernel fusion methods (Noble and Ben-Hur, 2007). Modelling interactions between gene products using functional linkage networks is realized through graphs, where gene products are modeled as nodes and relationships between genes through edges (Karaoz et al., 2004; Chua et al., 2007).

In vector space integration (VSI) different vectorial data are concatenated (desJardins et al., 1997), while kernel methods, by exploiting the closure property with respect to the sum or other meaningful algebraic operators represent another valuable research direction for the integration of biomolecular data (Lanckriet et al., 2004).

All these methods suffer of limitations and drawbacks, due to their limited modularity when new data sources are added (e.g. vector-space integration methods), or when data are not available as relational data (e.g. functional linkage networks), or to their limited scalability to multiple data sources (e.g. Kernel integration methods based on semidefinite programming (Lanckriet et al., 2004)). About this last item, it is worth noting that quite recently more efficient methods for multiple kernel learning have been proposed (Sonnenburg et al., 2006; Rakotomamonjy et al., 2007), but in any case they are slower than solving single SVMs and are relatively complicated to program.

Quite surprisingly, as observed by Noble and Ben-Hur (2007), only little attention has been devoted to ensemble methods as a mean to integrate multiple biomolecular sources of data for gene function prediction. To our knowledge only few works very recently considered ensemble methods in this specific bioinformatics context: Naive-Bayes integration of the outputs of SVMs trained with multiple sources of data (Guan et al., 2008), and logistic regression for combining the output of several SVMs trained with different data and kernels in order to produce probabilistic outputs corresponding to specific GO terms (Obozinski et al., 2008).

The main aim of this work consists in showing that simple ensemble methods can obtain results comparable with state-of-the-art data integration methods, exploiting at the same time the modularity and scalability that characterize most of the ensemble algorithms. Indeed biomolecular data differing for their structural characteristics (e.g. sequences, vectors, graphs) can be easily integrated, because with ensemble methods the integration is performed at the decision level, combining the outputs produced by classifiers trained on different datasets. Moreover, as new types of biomolecular data, or updates of data contained in public databases, are made available to the research community, ensembles of learning machines are able to embed new data sources or to update existing ones by training only the base learners devoted to the newly added or updated data, without retraining the entire ensemble. In other words ensemble methods scale well with the number of the available data sources, and problems that characterize other data fusion approaches are thus avoided.

2. Methods

Data fusion can be realized by means of an ensemble system composed by learners trained on different "views" of the data and then combining the outputs of the component learners. Each type of data may capture different and complementary characteristics of the objects to be classified and the resulting ensemble may obtain better prediction capabilities through the diversity and the anti-correlation of the base learner responses.

We programmatically considered simple methods: Weighted majority voting, Naive-Bayes and Decision Templates.

2.1 Weighted majority voting

Majority voting is based on the *Condorcet Jury Theorem*, proposed in the context of social sciences since the end of XVIII century (de Condorcet, 1785). The theorem proves that the judgment of a committee is superior to those of individuals, provided the individuals have reasonable competence.

Given a set of k classes whose labels $\omega_j \in \Omega$, $1 \leq j \leq k$, we denote by $d_{t,j}(x) \in [0,1]$ the support (e.g. the probability) estimated by the base t^{th} classifier of an ensemble composed by L base learners, that a given example x belongs to the class ω_j. For brevity we denote $d_{t,j}(x)$ as $d_{t,j}$. A simple way to integrate different data sources is represented by the weighed linear combination rule (Kittler et al., 1998), by which the posterior probability $\hat{P}$ of the resulting ensemble is estimated as follows:

$$\hat{P}(\omega_j|x) = \sum_{t=1}^{L} w_t d_{t,j}(x) \tag{1}$$

Considering that gene classes are largely unbalanced (positive examples are largely less than negative ones), we chose the F-measure to compute the weights:

$$w_t^f = \frac{F_t}{\sum_{t=1}^{L} F_t} \qquad\qquad w_t^{log} \propto log\frac{F_t}{1-F_t} \tag{2}$$

where F_t is the F-measure assessed on the training data for the t^{th} base learner. The w_t^f weights are obtained by a linear combination of the F-measures, and w_t^{log} by a logarithmic transformation. The decision $D_j(x)$ of the ensemble about the class ω_j is taken using the estimated probability $\hat{P}$ (eq. 1):

$$D_j(x) = \begin{cases} 1, & \text{if } \hat{P}(\omega_j|x) > 0.5 \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

where output 1 correspond to positive predictions for ω_j and 0 to negatives.

2.2 Naive-Bayes combination

The Naive-Bayes combination assumes independence between classifiers, and estimates the class-conditional support given the observed vector of categorized component classifiers outputs (Titterington et al., 1981). We denote by $s_t \in \Omega$ the class predicted by the t^{th} classifier, that is, in our setting, $s_t = \arg\max_j d_{t,j}$, and let be $\mathbf{s} =< s_1, s_2, \ldots, s_L >$ the vector of the classes predicted by the L base learners. By assuming conditional independence between classifiers, the class conditional probability for the class $\omega_j \in \Omega$ is:

$$P(\mathbf{s}|\omega_j) = P(s_1, s_2, \ldots, s_L|\omega_j) = \prod_{i=1}^{L} P(s_i|\omega_j) \tag{4}$$

and by applying the Bayes theorem we obtain the posterior probability for class ω_j:

$$P(\omega_j|\mathbf{s}) = \frac{P(\omega_j)P(\mathbf{s}|\omega_j)}{P(\mathbf{s})} = \frac{P(\omega_j)\prod_{i=1}^{L} P(s_i|\omega_j)}{P(\mathbf{s})} \tag{5}$$

By estimating the class conditional probability $P(s_i|\omega_j)$ for each base learner through the confusion matrix M^i computed on the training set, we can obtain an estimate of the posterior probability of the ensemble for the class ω_j. Finally, for each class ω_j, the Bayes rule is applied to choose the class predicted by the Naive-Bayes ensemble:

$$D_j(x) = \begin{cases} 1, & \text{if } \hat{P}(\omega_j|\mathbf{s}) > 0.5 \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

In our implementation we applied the Titterington's modification of the class conditional probability estimate to regularize and avoid zero probabilities in the computation of eq. 5:

$$\hat{P}(s_i|\omega_j) = \left(\frac{m^i_{\omega_j,s_i} + 1/k}{N_j + 1} \right) \tag{7}$$

where $m^i_{\omega_j,s_i}$ is the entry for the true class ω_j and predicted class s_i of the confusion matrix M^i for the i^{th} base learner, k is the number of classes and N_j the number of examples belonging to class ω_j.

2.3 Decision Templates

Decision Templates are a combination method based on the comparison of a "prototypical answer" of the ensemble for the examples belonging to a given class (the template) with the current answer of the ensemble to a specific example whose class needs to be predicted (the decision profile) (Kuncheva et al., 2001). The decision profile DP($\mathbf{x}$) for an instance $\mathbf{x}$ is a matrix composed by $d_{t,j} \in$[0,1] elements representing the support (e.g. the probability) given by the t^{th} classifier to class ω_j. Decision templates DT_j are the averaged decision profiles obtained from $\mathbf{X}_j$, the set of training instances belonging to the class ω_j:

$$DT_j = \frac{1}{|\mathbf{X}_j|} \sum_{\mathbf{x} \in \mathbf{X}_j} DP(\mathbf{x}) \tag{8}$$

By computing the similarity $\mathcal{S}$ between $DP(\mathbf{x})$ and the decision template DT_j for each class ω_j, from a set of c classes, the final decision of the ensemble is taken by assigning a test instance $\mathbf{x}$ to a class with the largest similarity (Kuncheva et al., 2001):

$$D(\mathbf{x}) = \arg\max_j \mathcal{S}_j(\mathbf{x}) \tag{9}$$

It is easy to see that with dichotomic problems the decision templates are reduced to two-columns matrices, and the similarity ($\mathcal{S}_1$) for the positive class and the similarity ($\mathcal{S}_2$) for the negative class can be computed as 1 minus the normalized squared euclidean distance:

$$\mathcal{S}_1(\mathbf{x}) = 1 - \frac{1}{n} \sum_{t=1}^{n} [DT_1(t,1) - d_{t,1}(\mathbf{x})]^2 \tag{10}$$

$$\mathcal{S}_2(\mathbf{x}) = 1 - \frac{1}{n} \sum_{t=1}^{n} [DT_2(t,1) - d_{t,1}(\mathbf{x})]^2 \tag{11}$$

where DT_1 is the decision template for the positive and DT_2 for the negative class. The final decision of the ensemble is:

$$D(\mathbf{x}) = \arg\max_{\{1,2\}}(\mathcal{S}_1(\mathbf{x}), \mathcal{S}_2(\mathbf{x})) \tag{12}$$

2.4 Kernel fusion and vector space integration

Kernel fusion (KF) for data integration is based on the closure property of kernels with respect to the sum or other algebraic operators. For instance, if K_a and K_b are both kernel functions, $K(x, y) = K_a(x, y) + K_b(x, y)$ is another valid kernel, as well as a weighted combination of kernels $w_a K_a(x, y) + w_b K_b(x, y), \quad w_a \geq 0, w_b \geq 0$.

In our experiments we integrated the different data sets by simply summing their Gram matrices, and then we trained the SVMs directly with the resulting matrix. Moreover in Sect. 3.2 we considered also semi-definite programming methods that allow the joint optimization of both the SVM margin and of the weights on each individual kernel (Lanckriet et al., 2004).

Vector space integration (VSI) consists in concatenating the vectors of the different data sets (desJardins et al., 1997). The resulting concatenated vectors are used to train a SVM. Note that training a linear SVM with concatenated vectors (VSI) is equivalent to kernel fusion with linear kernels. In our experiments we used gaussian kernels.

3. Experimental results

Even if the growing rate of the amount of biomolecular data available for many species was constantly increasing in the last years, the model organisms with a consistent amount of literature inherent to data fusion based gene function prediction are actually reduced to S.*cerevisiae* and M.*musculus*. Despite the availability of a well established public benchmark dataset, such as the one provided during the MouseFunc contest (Pena-Castillo et al., 2008), a recent comparison between many model organisms showed that the fraction of genes annotated with experimental evidence is about 30% larger in S.*cerevisiae* than in M.*musculus* (85.4% and 57.8% respectively for the yeast and mouse model organisms, Rhee et al. (2008)). We thus decided to use yeast data for our experiments. In order to maximize the effective use of the larger experimental coverage of gene functional annotations available for the yeast, we also adopted as a reference functional ontology, the MIPS Functional Catalogue (FunCAT), which is composed by annotations mainly based on experimental evidences (Ruepp et al., 2004), allowing us to minimize the impact of non experimental functional annotations.

3.1 Prediction of top-level FunCat classes in yeast

We predicted the top-level 15 functional classes of the FunCat taxonomy of the model organism *S. cerevisiae*, using 6 different sources of data (Table 1). Two of the considered datasets are devoted to the characterization of each gene in terms of the protein-domain architecture of its protein product (D_{pfam1} and D_{pfam2}). One dataset describes the expression pattern of the genes in several experimental conditions (D_{expr}). The last three data sources involved in this experiment represent different kinds of relationships between pro-

Table 1: Datasets

Code	Dataset	examples	features	description
D_{ppi1}	PPI - STRING	2338	2559	protein-protein interaction data from (vonMering et al., 2003)
D_{ppi2}	PPI - BioGRID	4531	5367	protein-protein interaction data from the *BioGRID* database (Stark et al., 2006)
D_{pfam1}	Protein domain log-E	3529	5724	Pfam protein domains with log E-values computed by the *HMMER* software toolkit
D_{pfam2}	Protein domain binary	3529	4950	protein domains obtained from *Pfam* database (Finn et al., 2008)
D_{expr}	Gene expression	4532	250	merged data of Spellman and Gasch experiments
D_{seq}	Pairwise similarity	3527	6349	Smith and Waterman log-E values between all pairs of yeast sequences

teins: experimentally supported (D_{ppi1}) and predicted (D_{ppi2}) protein-protein interactions, and evolutionary relationships expressed in terms of protein sequence conservation (D_{seq}). Each dataset was split into a training set and a test set (composed, respectively, by the 70% and 30% of the available samples), considering yeast genes common to all data sets (about 1900) and with at least 1 FunCat annotation. A 3-fold stratified cross-validation has been performed on the training data for model selection, using gaussian SVMs with probabilistic output (Lin et al., 2007) as base learners for both ensemble methods, and for VSI and KF data integration. More precisely, we applied a grid search by varying both the C regularization parameter and the σ parameter of the gaussian kernel between 10^{-3} and 10^{3}, in order to select the best model for each classification task. We compared the performances of single gaussian SVMs trained on each data set with those obtained with vector-space-integration (VSI) techniques, kernel fusion through the sum of gaussian kernels, and with the ensembles described in Sect. 2.

Table 2 shows the average F-measure, recall, precision and AUC across the 15 selected FunCat classes, obtained through the evaluation of the test sets (each constituted by 570 genes). The four first columns refer respectively to the weighted linear, weighted logarithmic, Decision Template and Naive-Bayes ensembles; VSI and KF stands respectively for vector space integration and kernel fusion, D_{avg} represents the average results of the single SVMs across the six datasets, and D_{ppi2} represents the single SVM that achieved the best performance, i.e. the one trained using protein-protein interactions data collected from BioGrid. Table 3 shows the same results obtained by each single SVM trained on a specific biomolecular data set.

Looking at the values presented in Table 2, on the average, data integration through simple ensemble methods provide better results than single SVMs, VSI and Kernel fusion, independently of the applied combination rule. In particular, Decision Templates achieved the best average F-measure, and ensemble methods as a whole the best AUC. Among the ensemble of classifiers, with respect to the AUC, the worst performing method is the

Naive-Bayes combiner albeit its performances are still, on the average, higher than the ones reported for VSI, Kernel fusion and the single classifiers. Precision of the ensemble methods is relatively high: this is of paramount importance to drive the biological validation of "in silico" predicted functional classes: considering the high costs of biological experiments, we need to obtain a high precision (and possibly recall) to be sure that positive predictions are actually true with the largest confidence.

To understand whether the differences between AUC scores in the 15 dichotomic tasks are significant, we applied a non parametric test based on the Mann-Whitney statistic (Delong et al., 1988), using a recently proposed software implementation (Vergara et al., 2008). Table 4 shows that at 0.01 significance level in most cases there is no significant difference between AUC scores of the weighted linear and logarithmic ensembles (E_{lin} and E_{log}) and the Decision Template (E_{dt}) combiner. A different behavior is observed for the Naive-Bayes combiner: its performances are comparable to the ones obtained by the other ensemble methods only in 2 over 15 classification tasks and worse in the remaining 13.

Most interestingly, ensemble methods significantly outperform the other data integration methods. For instance, wins-ties-losses of E_{lin} vs VSI are $13-2-0$, and $9-6-0$ vs KF; Naive-Bayes, the worst performing ensemble method, achieves $9-6-0$ wins-ties-losses with VSI and $5-10-0$ with KF. It is worth noting that, among the tested ensemble methods, E_{lin}, E_{log} and E_{dt} undergo no losses when compared with single SVMs (Table 4, bottom): we can safely choose any ensemble method (but not the Naive-Bayes combiner) to obtain equal or better results than any of the single SVMs. On the contrary in many cases VSI, E_{NB} and the kernel fusion methods obtained worse results than single SVMs, although performances achieved by the Naive-Bayes combiner and the kernel fusion methods are, in general, better than those obtained by VSI. Nevertheless, we can observe that a single SVM trained with Ppi-2 data achieves good results (11 ties with ensembles and an average AUC $\simeq 0.81$ w.r.t. 0.86 of the ensembles, Table 2 and 4), showing that large protein-protein interactions data sets alone provide information sufficient to correctly predict several FunCat classes.

Figure 1 compares the ROC curves of the different data integration methods used in our experiments. ROC curves of weighted majority voting (E_{lin}) are consistently above the corresponding ROC curves of kernel fusion and vector space integration for all the considered FunCat classes. ROC curves of Naive-Bayes combiner are below those of kernel fusion only

Table 2: Ensemble methods, kernel fusion and vector space integration: average F-score, recall, precision and AUC (Area Under the Curve) across the data sets.

Metric	E_{lin}	E_{log}	E_{dt}	E_{NB}	VSI	KF	D_{avg}	D_{ppi2}
F	0.4347	0.4111	0.5302	0.5174	0.3213	0.3782	0.3544	0.4818
rec	0.3304	0.2974	0.4446	0.6467	0.2260	0.3039	0.2859	0.3970
prec	0.8179	0.8443	0.7034	0.5328	0.6530	0.6293	0.5823	0.6157
AUC	0.8642	0.8653	0.8613	0.7933	0.7238	0.7775	0.7265	0.8170

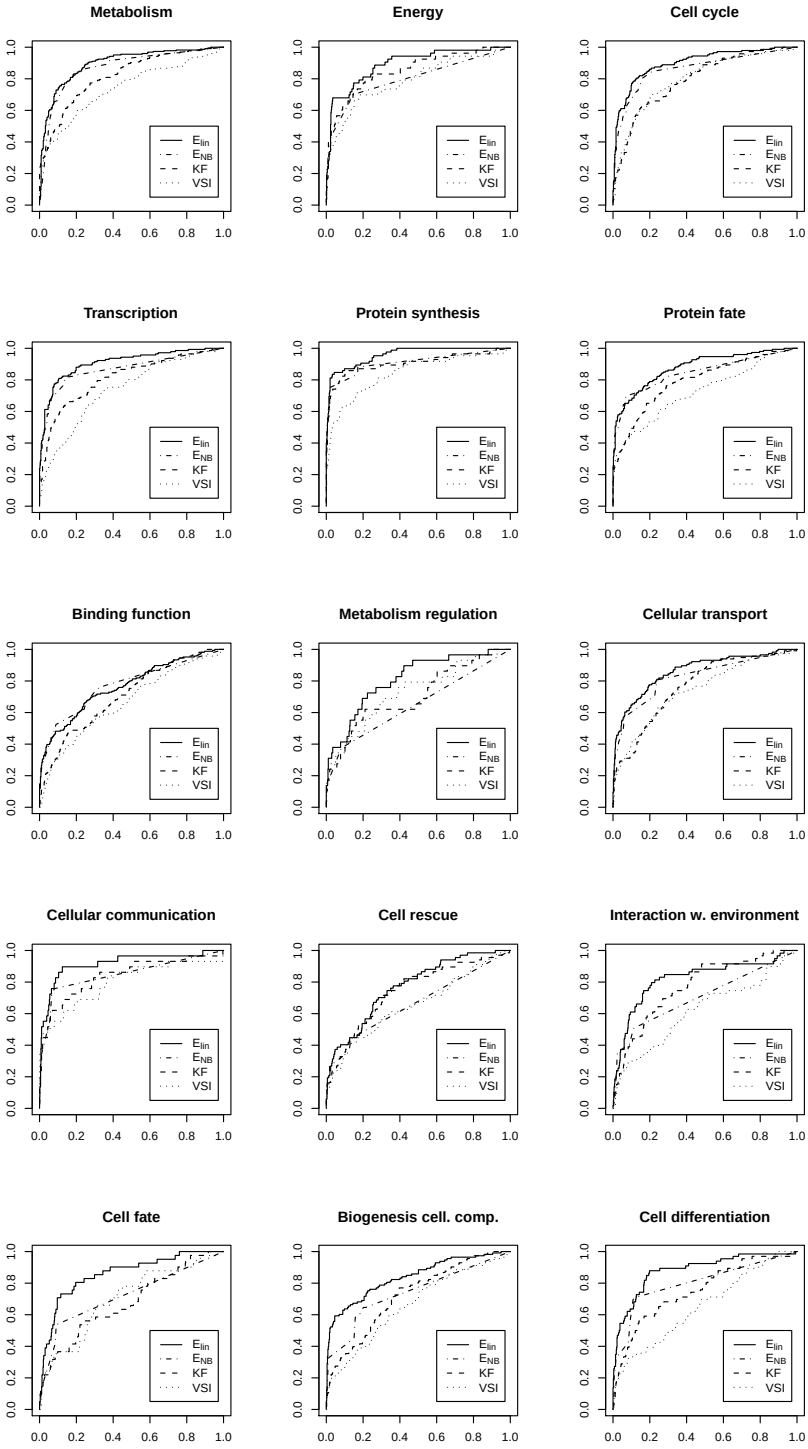

Figure 1: Comparison of ROC curves between different data integration methods. E_{lin}: ensemble weighted majority voting; E_{NB}: Naive-Bayes ensemble integration; KF: kernel fusion; VSI: vector space integration.

Table 3: Single SVMs: average F-score, recall, precision and AUC. Each SVM is identified by the same name of the data set used for its training (Table 1).

Metric	D_{ppi1}	D_{ppi2}	D_{pfam1}	D_{pfam2}	D_{expr}	D_{seq}
F	0.3655	0.4818	0.2363	0.3391	0.2098	0.4493
rec	0.2716	0.3970	0.1457	0.2417	0.1571	0.5019
prec	0.6157	0.6785	0.7154	0.6752	0.3922	0.4162
AUC	0.7501	0.8170	0.6952	0.6995	0.6507	0.7469

Table 4: Results of the non-parametric test based on Mann-Whitney statistics to compare AUCs between ensembles, VSI, Kernel fusion and single SVMs. Each entry represents wins-ties-losses between the corresponding row and column at 0.01 significance level. Top: Comparison between ensemble methods, VSI and kernel fusion; Bottom: Comparison between data integration methods and single SVMs.

	VSI	E_{log}	E_{lin}	E_{dt}	E_{NB}
E_{log}	13-2-0	-	-	-	-
E_{lin}	13-2-0	0-14-1	-	-	-
E_{dt}	13-2-0	1-13-1	1-11-3	-	-
E_{NB}	9-6-0	0-2-13	0-2-13	0-2-13	-
KF	3-12-0	0-6-9	0-6-9	0-6-9	0-10-5

	D_{ppi1}	D_{ppi2}	D_{pfam1}	D_{pfam2}	D_{expr}	D_{seq}
E_{lin}	11-4-0	4-11-0	15-0-0	14-1-0	15-0-0	13-2-0
E_{log}	11-4-0	4-11-0	15-0-0	14-1-0	15-0-0	13-2-0
E_{dt}	11-4-0	4-11-0	15-0-0	14-1-0	15-0-0	13-2-0
E_{NB}	5-10-0	2-11-2	9-6-0	8-7-0	12-3-0	7-8-0
VSI	1-11-3	0-8-7	2-11-2	1-14-0	4-11-0	0-12-3
KF	1-14-0	0-9-6	5-10-0	5-10-0	11-4-0	3-12-0

for four classes: "Energy", "Metabolism", "Regulation", "Cell rescue" and "Interaction with the environment".

3.2 Predicting GO terms using protein sequence and structural information

To compare our proposed ensemble methods against published results relative to published benchmark data sets, we chose the sequence and structural protein yeast data available from Lewis et al. (2006).

Our aim is to compare the results obtained with average kernel fusion (KF) and weighted average KF through semi-definite programming techniques (KF SDP) published in (Lewis

et al., 2006) with the results obtained with Weighted linear ensembles and Decision Templates using the same kernel machines as base learners.

Sequence data are represented through mismatch kernels (Leslie et al., 2003), that constitute a generalization upon the simpler spectrum kernel, which represents a string as a vector of counts of all possible substrings of a fixed length (Leslie et al., 2002). For structural data, empirical kernel maps (Scholkopf et al., 2004) have been applied to convert MAMMOTH scores to kernels. MAMMOTH (Ortiz et al., 2002) is an algorithm for structural alignment of proteins; it computes a score that reflects the alignment's quality. Unfortunately these scores cannot be used directly for kernels because the corresponding Gram matrix is not positive definite and the empirical kernel map needs to be obtained from the scores instead.

In the experiments we tried to recognize the same 56 GO terms considered in Lewis et al. (2006): they come from both the Biological Process, Molecular Function and Cellular Compartment ontologies. Moreover we adopted the same experimental set-up: 5-fold cross validation repeated three times, a fixed SVM regularization parameter $C = 10$ for all the classification tasks, and the AUC to measure classification performance. Note that for the average and SDP kernel fusion methods we used the AUC results published in Lewis et al. (2006).

Results are summarized in Figure 2. Each plot represents the comparison of the AUC scores between an ensemble and a kernel fusion method across all the 56 protein classification problems. Most of the points lie above the bisector in Figure 2 b) and d), showing that both Weighted linear and Decision Templates ensembles tend to outperform the SDP kernel fusion approach. This is less clear when we compare ensemble methods with the average kernel fusion technique: indeed while a certain prevalence can be observed for the Weighted linear ensemble, with Decision Templates the points seem to be quite equally distributed along the bisector (Figure 2 a and c). These visual clues are confirmed by the Wilcoxon signed-ranks test (Wilcoxon, 1945): we register a significant difference in favour of Weighted linear ($p\text{-}value \simeq 10^{-6}$) and Decision Template ($p\text{-}value \simeq 10^{-5}$) w.r.t. SDP KF, a significant difference between Weighted linear and average KF ($p\text{-}value \simeq 0.05$), but no significant difference between Decision Templates and average KF ($p\text{-}value \simeq 0.29$).

4. Conclusions

The main objective of this contribution is to demonstrate that simple ensemble methods are competitive with state-of-the-art methods for gene function prediction based on heterogeneous biomolecular data integration.

It is well-known that gene function prediction methods need to take into account the hierarchical relationships between classes to improve their predictions (Guan et al., 2008; Obozinski et al., 2008; Valentini and Re, 2009; Valentini, 2010). Nevertheless, in this investigation we focused on data integration, in order to study the improvement due to the usage of multiple sources of data, without exploiting any knowledge about the hierarchical relationships between classes. In this way we can separate the contribution due to data fusion techniques from the improvement due to hierarchical methods.

Considering the increasing growing rate of available biomolecular data, the modularity and scalability that characterize ensemble methods can favour an easy update of existing

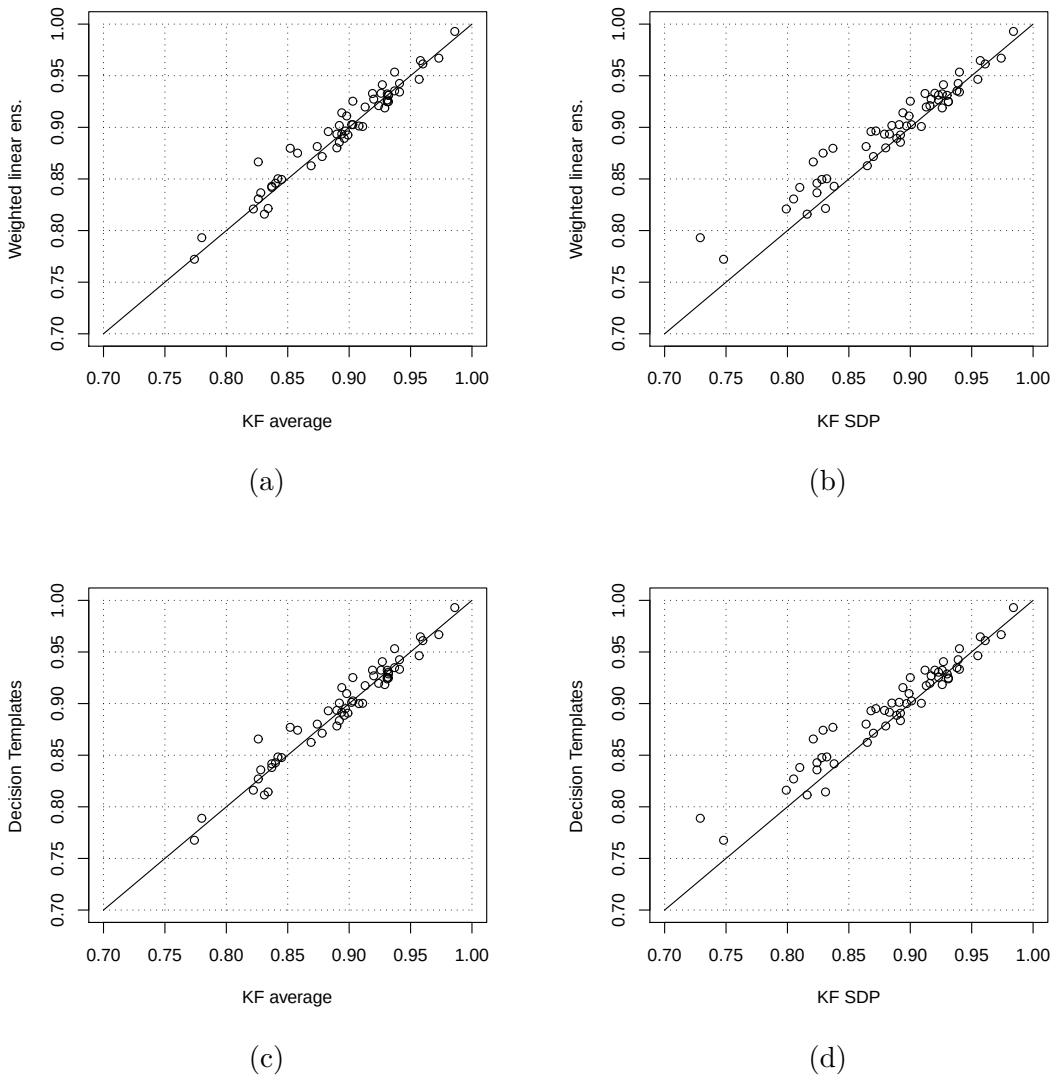

Figure 2: Comparison of AUC results between kernel fusion and ensemble data integration methods. Points represent the AUC score of kernel fusion (abscissa) and ensemble (ordinate) methods for 56 GO terms. (a) Average KF vs Weighted ensembles (b) SDP KF vs Weighted ensembles (c) Average KF vs Decision Templates (d) SDP KF vs Decision Templates.

sources of data and an easy integration of new ones. Our experiments show that relatively simple ensemble methods are competitive with kernel fusion and vector space integration, two of the most largely applied machine learning data integration techniques for gene function prediction. This could seem quite surprising, but considering the uncertainty that

characterize both annotations and measurements of data values, we can expect that relatively simple methods are able to nicely work in a similar context. Moreover it is worth noting that each type of data can only capture a particular characteristic of a protein, and for different functional classes the same type of data can be highly informative or completely unuseful to discriminate positive and negative examples. For these reasons the inherent modularity and adaptivity of ensemble systems can explain their effectiveness for the integration of multiple biomolecular data sources. In particular we think that ensemble methods devoted to biomolecular data integration can be a valuable research line to improve the accuracy of gene function prediction problems.

Acknowledgments

We would like to thank the reviewers for their comments. The authors gratefully acknowledge partial support by the PASCAL2 Network of Excellence under EC grant no. 216886. This publication only reflects the authors' views.

References

H.N. Chua, W. Sung, and L. Wong. An efficient strategy for extensive integration of diverse biological data for protein function prediction. *Bioinformatics*, 23(24):3364–3373, 2007.

N.C. de Condorcet. *Essai sur l' application de l' analyse à la probabilité des decisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, 1785.

ER Delong, DM Delong, and DL Clarke-Pearson. Comparing the areas under two or more or more correlated Receiver Operating Characteristics Curves: a non parametric approach. *Biometrics*, 44(3):837–845, 1988.

M. desJardins, P. Karp, M. Krummenacker, T.J. Lee, and C.A. Ouzounis. Prediction of enzyme classification from protein sequence without the use of sequence similarity. In *Proc. of the 5th ISMB*, pages 92–99. AAAI Press, 1997.

R.D. Finn, J. Tate, J. Mistry, P.C. Coggill, J.S. Sammut, H.R. Hotz, G. Ceric, K. Forslund, S.R. Eddy, E.L. Sonnhammer, and A. Bateman. The Pfam protein families database. *Nucleic Acids Research*, 36:D281–D288, 2008.

Y Guan, C.L. Myers, D.C. Hess, Z. Barutcuoglu, A. Caudy, and O.G. Troyanskaya. Predicting gene function in a hierarchical context with an ensemble of classifiers. *Genome Biology*, 9(S2), 2008.

U. Karaoz et al. Whole-genome annotation by using evidence integration in functional-linkage networks. *Proc. Natl Acad. Sci. USA*, 101:2888–2893, 2004.

J. Kittler, M. Hatef, R.P.W. Duin, and J. Matas. On combining classifiers. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 20(3):226–239, 1998.

L.I. Kuncheva, J.C. Bezdek, and R.P.W. Duin. Decision templates for multiple classifier fusion: an experimental comparison. *Pattern Recognition*, 34(2):299–314, 2001.

G.R. Lanckriet, T. De Bie, N. Cristianini, M. Jordan, and W.S. Noble. A statistical framework for genomic data fusion. *Bioinformatics*, 20:2626–2635, 2004.

C. Leslie, E. Eskin, and W.S. Noble. The spectrum kernel: a string kernel for SVM protein classification. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 564–575, New Jersey, 2002. World Scientific.

C. Leslie et al. Mismatch string kernels for SVM protein classification. In *Advances in Neural Information Processing Systems*, pages 1441–1448, Cambridge, MA, 2003. MIT Press.

D.P. Lewis, T. Jebara, and W.S. Noble. Support vector machine learning from heterogeneous data: an empirical analysis using protein sequence and structure. *Bioinformatics*, 22(22): 2753–2760, 2006.

H.T. Lin, C.J. Lin, and R.C. Weng. A note on Platt's probabilistic outputs for support vector machines. *Machine Learning*, 68:267–276, 2007.

W.S. Noble and A. Ben-Hur. Integrating information for protein function prediction. In T. Lengauer, editor, *Bioinformatics - From Genomes to Therapies*, volume 3, pages 1297–1314. Wiley-VCH, 2007.

G. Obozinski, G. Lanckriet, C. Grant, Jordan. M., and W.S. Noble. Consistent probabilistic output for protein function prediction. *Genome Biology*, 9(S6), 2008.

A.R. Ortiz et al. MAMMOTH (matching molecular models obtained from theory): an automated method for model comparison. *Protein Sci.*, 11:2606–2621, 2002.

L. Pena-Castillo et al. A critical assessment of Mus musculus gene function prediction using integrated genomic evidence. *Genome Biology*, 9:S1, 2008.

A. Rakotomamonjy, F. Bach, S. Canu, and Y. Grandvalet. More efficiency in multiple kernel learning. In *ICML '07: Proceedings of the 24th international conference on Machine learning*, pages 775–782, New York, NY, USA, 2007. ACM.

S.Y Rhee et al. Use and misuse of the gene ontology annotations. *Nature Rev. Genetics*, 9 (17-18):509–515, 2008.

A. Ruepp, A. Zollner, D. Maier, K. Albermann, J. Hani, M. Mokrejs, I. Tetko, U. Guldener, G. Mannhaupt, M. Munsterkotter, and H.W. Mewes. The FunCat, a functional annotation scheme for systematic classification of proteins from whole genomes. *Nucleic Acids Research*, 32(18):5539–5545, 2004.

B. Scholkopf, K. Tsuda, and J.P. Vert. *Kernel Methods in Computational Biology.* MIT Press, Cambridge, MA, 2004.

S. Sonnenburg, G. Ratsch, C. Schafer, and B. Scholkopf. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, 2006.

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

The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nature Genet.*, 25:25–29, 2000.

D. Titterington, G. Murray, D. Spiegelhalter, A. Skene, J. Habbema, and G. Gelpke. Comparison of discriminant techniques applied to a complex data set of head injured patients. *Journal of the Royal Statistical Society*, 144(2), 1981.

G. Valentini and M. Re. Weighted True Path Rule: a multilabel hierarchical algorithm for gene function prediction. In *MLD-ECML 2009, 1st International Workshop on learning from Multi-Label Data*, pages 133–146, Bled, Slovenia, 2009.

G. Valentini. True Path Rule hierarchical ensembles for genome-wide gene function prediction. *IEEE ACM Transactions on Computational Biology and Bioinformatics* (in press).

I Vergara, T Norambuena, E Ferrada, A Slater, and F. Melo. StAR: a simple tool for the statistical comparison of ROC curves. *BMC Bioinformatics*, 9(265), 2008.

C. vonMering et al. STRING: a database of predicted functional associations between proteins. *Nucleic Acids Research*, 31:258–261, 2003.

F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics*, 1:80–83, 1945.