

Noise tolerance of Multiple Classifier Systems in data integration-based gene function prediction

Matteo Rè and Giorgio Valentini

Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano
v. Comelico 39 Milano, Italy, <http://www.dsi.unimi.it>

Summary

The availability of various high-throughput experimental and computational methods developed in the last decade allowed molecular biologists to investigate the functions of genes at system level opening unprecedented research opportunities. Despite the automated prediction of genes functions could be included in the most difficult problems in bioinformatics, several recently published works showed that consistent improvements in prediction performances can be obtained by integrating heterogeneous data sources. Nevertheless, very few works have been dedicated to the investigation of the impact of noisy data on the prediction performances achievable by using data integration approaches.

In this contribution we investigated the tolerance of multiple classifier systems (MCS) to noisy data in gene function prediction experiments based on data integration methods. The experimental results show that performances of MCS do not undergo a significant decay when noisy data sets are added. In addition, we show that in this task MCS are competitive with kernel fusion, one of the most widely applied technique for data integration in gene function prediction problems.

1 Introduction

Recently, several works highlighted the central role played by data integration methods for the gene function prediction problem [14, 1, 5, 22]. Indeed it is well-known that each source of biomolecular data can reveal specific features of gene/gene products, and this information can be useful to characterize their functional role in living systems.

Gene function prediction in its general formulation is a complex classification problem characterized by the following items:

- each gene/gene product can be assigned to multiple terms/classes (a multiclass, multilabel classification problem)
- classes are structured according to a predefined hierarchy (a directed acyclic graph for the Gene Ontology (GO) [23] or a tree forest for FunCat [19]);
- classes are usually unbalanced (with positive examples usually less than negatives)
- gene labels are in several cases uncertain or largely incomplete
- multiple sources of data can be used to predict gene functions.

Even if the integration of multiple sources of data has enjoyed a certain attention in the computational biology community [9, 1, 14], at best of our knowledge, only few works tried to evaluate the impact of noisy data in data integration methods for gene function prediction problems [11].

According to [14], the main supervised gene function prediction methods based on heterogeneous data integration proposed in the literature can be schematically subdivided in three main categories: functional linkage networks, vector subspace integration and kernel fusion methods. Modeling 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 [6, 1]. In vector space integration (VSI) different vectorial data are concatenated [3], 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 [9].

Quite surprisingly, as observed in [14], only little attention has been devoted to multiple classifier systems (MCS) 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 [4], 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 [15]. We recently demonstrated 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 [17]. 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. 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.

In this context, the resistance of prediction methods to noise is a fundamental issue for at least two reasons. At first, it should be noticed that any single data set involved in a data integration experiment, aimed at predicting gene functions at a whole ontology level, is likely to represent noise in a consistent fraction of the nodes composing the functional hierarchy. Indeed, a data source that can be predictive for a functional class can be non-predictive for other classes, considering the extreme complexity of the most widely used functional ontologies (GO and FunCat), both composed by hundreds or thousands of nested terms. Moreover, most of the commonly used biomolecular data are noisy due to the machinery and the complex procedures involved in the generation of the data through high-throughput biotechnologies.

In this paper we study the noise tolerance of data integration methods in the context of supervised gene function prediction problems, with a particular focus on multiple classifier systems.

The paper is structured as follows. In Section 2 two MCS-based and a kernel fusion-based method for biomolecular data integration are presented. Section 3 and 4 show the experimental set-up and results relative to the analysis of noise tolerance of multiple classifier systems in two different gene function prediction problems involving the integration of multiple data sources. The results are discussed in Section 5 and the conclusions end the paper.

2 Methods

2.1 Multiple Classifier Systems

An ensemble system can integrate multiple sources of biomolecular data by training different learning machines on different "views" of the data, and then by combining the outputs of the component learners. In this work we programmatically considered simple methods: Weighted majority voting and Decision Templates.

2.1.1 Weighted majority voting

It is well-known, at least from the XVIII century, that the judgment of a committee is superior to those of individuals, provided the individuals have reasonable competence (*Condorcet Jury Theorem* [2]). *Weighted majority voting* ensembles are ensemble methods based on the cited theorem. They simply integrate multiple sources of data by aggregating classifiers trained on multiple "views" of the data, and by weighting the decision of each base classifier on the basis of its "competence" on the learned data.

More precisely, 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 [7], by which the posterior probability \hat{P} of the resulting ensemble is estimated as follows:

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

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

$$w_t = \frac{F_t}{\sum_{t=1}^L F_t} \quad (2)$$

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

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

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

2.1.2 Decision Templates

The main idea behind Decision Templates consists in the 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) [8]. 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 :

$$DT_j = \frac{1}{|\mathbf{X}_j|} \sum_{\mathbf{x} \in \mathbf{X}_j} DP(\mathbf{x}) \quad (4)$$

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 [8]:

$$D(\mathbf{x}) = \arg \max_j \mathcal{S}_j(\mathbf{x}) \quad (5)$$

In multi-class classification problems a decision template can be represented through a matrix with a number of columns equal to the number of classes, but it is easy to see that with dichotomic problems the decision templates are reduced to one-column matrices (since $DT_1(t, 2) = 1 - DT_1(t, 1)$), where $DT_1(t, y)$, $y \in \{1, 2\}$ represents the element (t, y) of the decision template for class 1.

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:

$$\mathcal{S}_1(\mathbf{x}) = 1 - \frac{1}{n} \sum_{t=1}^n [DT_1(t, 1) - d_{t,1}(\mathbf{x})]^2 \quad (6)$$

$$\mathcal{S}_2(\mathbf{x}) = 1 - \frac{1}{n} \sum_{t=1}^n [DT_2(t, 1) - d_{t,1}(\mathbf{x})]^2 \quad (7)$$

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

$$D(\mathbf{x}) = \arg \max_{\{1,2\}} (\mathcal{S}_1(\mathbf{x}), \mathcal{S}_2(\mathbf{x})) \quad (8)$$

2.2 Kernel fusion methods

Even if kernel functions are widely adopted in machine learning because they can be used with standard algorithms (i.e. SVM) to solve binary and multiclass classification problems, they can also be thought as a way to encode similarities among non-vectorial and heterogeneous data. Indeed as long as our datasets can be represented as a square kernel matrix, then any kernel method can be applied to the data. A key feature of kernels making them very attractive in problems requiring the integration of heterogeneous data is that their mathematics allow us to derive new kernels by combining two or more kernels each representing diverse types of data. Many algebraic operators are closed under positive semidefiniteness. Among them, the most important operation, w.r.t. data integration problems, is the sum: if K_1 and K_2 are valid kernel functions, then we can demonstrate that $K(x, y) = K_1(x, y) + K_2(x, y)$ is still a valid kernel.

Also, for positive coefficients, a weighted combination of kernels $(\mu_1 K_1(x, y) + \mu_2 K_2(x, y))$ preserves the Mercer's conditions [13].

Despite kernel combination methods have been successfully applied in a broad range of bioinformatics problems, the most effective way to compute the weights involved in the kernel fusion is still an open question. Quite interestingly, a recent work pointed out that a kernel fusion obtained simply by averaging the kernel matrices associated to each data source is able to perform equally well (if not better) than more sophisticated methods (such as multiple kernel methods based on semidefinite programming (SDP) [9]) that introduces weights on each kernel matrix [11]. In this work we compared the performances of less computationally intensive data integration methods, the ensemble systems, with performances achievable by a kernel fusion based on the average of kernel matrices.

3 Experimental setup

3.1 Data sets and functional labeling

Considering that a recent study showed that the fraction of genes annotated with experimental evidence in *S.cerevisiae* is one of the largest among many model organisms [18], we chose the yeast for our experiments. For our experiments we considered two groups of data sets: a) a small set composed by protein sequence and protein structural data, previously analyzed in [11]; b) a larger set that includes six data sets previously collected in [17] (see Tab. 1 in the supplementary materials for details about these data).

Sequence and structural data have been labeled according to the GO term prediction benchmark, originally proposed in [11]. It involves 5323 yeast genes labeled according to 56 GO classes: 27, 22 and 7 of them belong to the MF, BP and CC ontologies respectively. The number of positives genes in the considered functional class ranges from 101 to 516. The negatives sets in the training sets were constructed in order to equal the number of positives examples, using the same strategy proposed in [11]. For the second group of data we labeled the genes with the 15 top-level FunCat functional classes (see Tab. 2 in the supplementary materials for more details). For this second task, considering that the classes are less unbalanced, we use all the available negative examples.

3.2 Experiments with noisy data

We performed two sets of experiments, respectively using the first and the second group of biomolecular data sets (Section 3.1). In both cases we analyzed the tolerance of ensemble systems and kernel fusion methods to noisy data, by adding to the original data noisy kernel matrices.

3.2.1 Experiments with sequence and structural data

The performances of the evaluated methods for the prediction of any of the 56 considered GO terms have been evaluated using a 5-fold cross validation scheme repeated three times

as in [11] in order to allow a direct comparison. We then produced two noisy kernels by a random permutation of the rows and the corresponding columns of the structure kernel matrix. These two noisy kernels have been used to train two additional component classifiers and to evaluate the performances of data integration methods. At first, we trained two SVMs using the original no-noisy data: protein sequence similarities encoded through a mismatch kernel [10] and protein structure similarities encoded through an empirical kernel map [20]. In particular the structural alignments used in the construction of the structure kernel have been obtained using the MAMMOTH structural aligner [16]. The E-values produced by MAMMOTH have been converted in a valid kernel matrix using an empirical kernel map. Prior to integration, each kernel is centered around the origin in the feature space and each example is projected onto the unit sphere using a cosine normalization. We integrated the two kernel matrices by using the simple average kernel fusion method, an ensemble system based on the weighted average (using linear weights based on the F-scores obtained by the SVM associated to the structure and mismatch kernels), and a second ensemble that does not require the definition of explicit weights (the Decision Templates combiner, Sect. 2).

Then we analyzed the resistance of ensemble systems to noisy data. We tested the average kernel fusion (KF), the weighted average ensemble and the Decision Templates ensemble for the integration of the structure and mismatch kernels adding one or two noisy kernels. Following the experimental setup of [11] we used, as component classifiers for the ensemble systems, linear SVMs with a fixed C regularization parameter set to 10. To compare our results with those published in [11], we computed the AUC scores for each prediction task, averaged across the 15 splits resulting from the 5-fold cross-validation repeated 3 times.

3.2.2 Experiments with six different types of biomolecular data

In these experiments we tested the tolerance to noise of MCS using a larger set of data. For each available data set, we generated a noisy copy by applying the same procedures described in Section 3.2.1, and we compared the performances of ensemble methods and KF using: a) the six original no-noisy data sets described in Section 3.1; b) the same six no-noisy data plus six noisy copies of all the available data. We evaluated the performances on 15 top-level FunCat classes using 5-fold cross-validation techniques. We used as base learners linear probabilistic SVMs [12] with a fixed C regularization parameter set to 10.

4 Results

4.1 Results with sequence and structural data

Fig. 1 shows the comparative results, in absence of noisy data, obtained by the average kernel fusion, the weighted average ensemble and the decision template combiner for the integration of the structure and mismatch kernels. Each point in the graph represents the AUC results of one of the 56 classification tasks: Fig. 1 (a) represents KF vs. weighted average ensembles, and Fig. 1 (b) KF vs. decision templates.

Fig. 2 and 3 compare the AUC results achieved by adding respectively one and two noisy data sets. In both cases KF average methods are compared with weighted linear and decision tem-

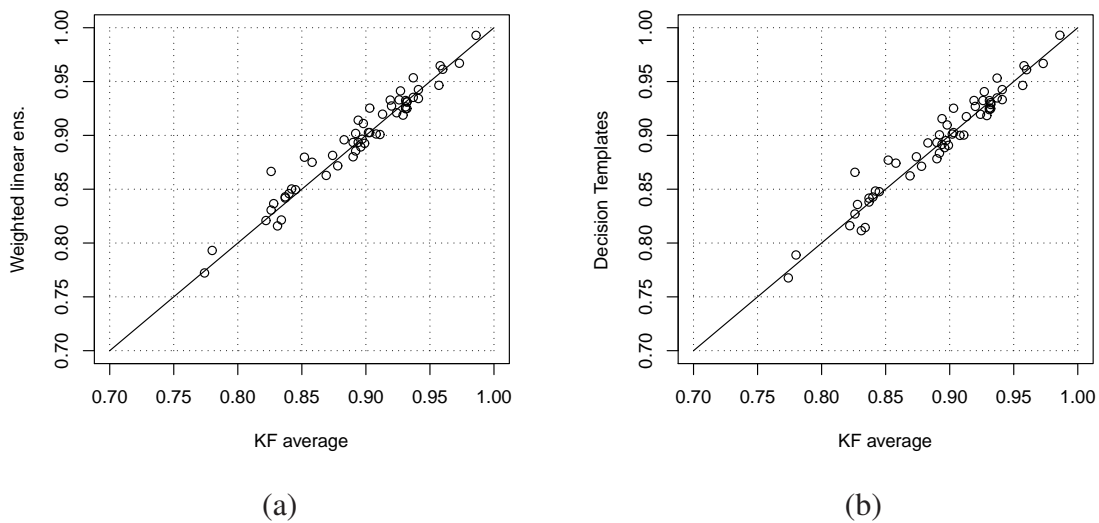


Figure 1: Comparison of AUC results between average kernel fusion and ensemble data integration methods without noisy data. Points represent the AUC score of kernel fusion (abscissa) and ensemble (ordinate) methods for 56 GO terms. (a) Average KF vs Weighted ensembles (b) Average KF vs Decision Templates.

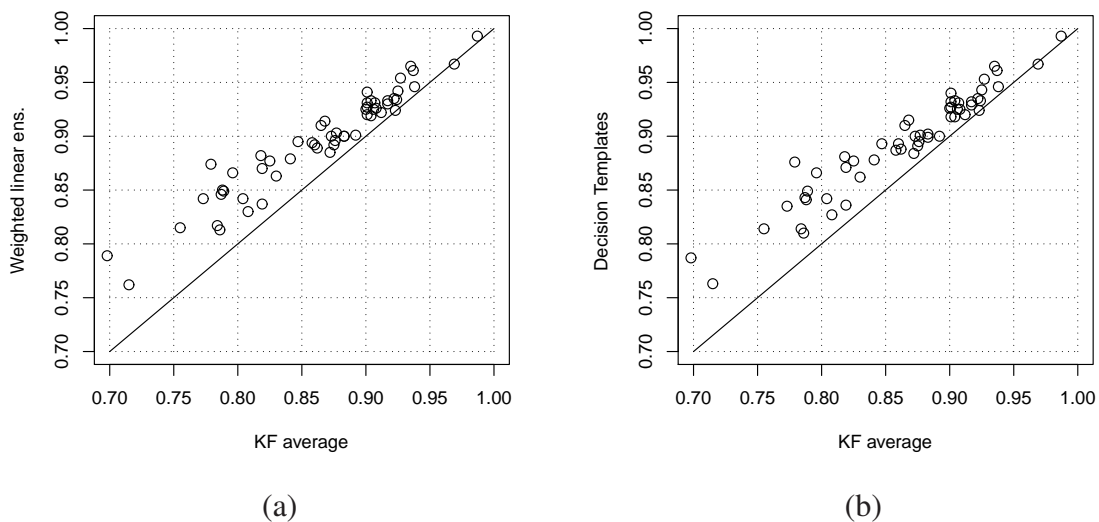


Figure 2: Comparison of AUC results between average kernel fusion and ensemble data integration methods with 1 noisy data set. Points represent the AUC score of kernel fusion (abscissa) and ensemble (ordinate) methods for 56 GO terms. (a) Average KF vs Weighted ensembles (b) Average KF vs Decision Templates.

plate ensembles. Detailed results for each GO term are shown in Table 3 in the Supplementary Information.

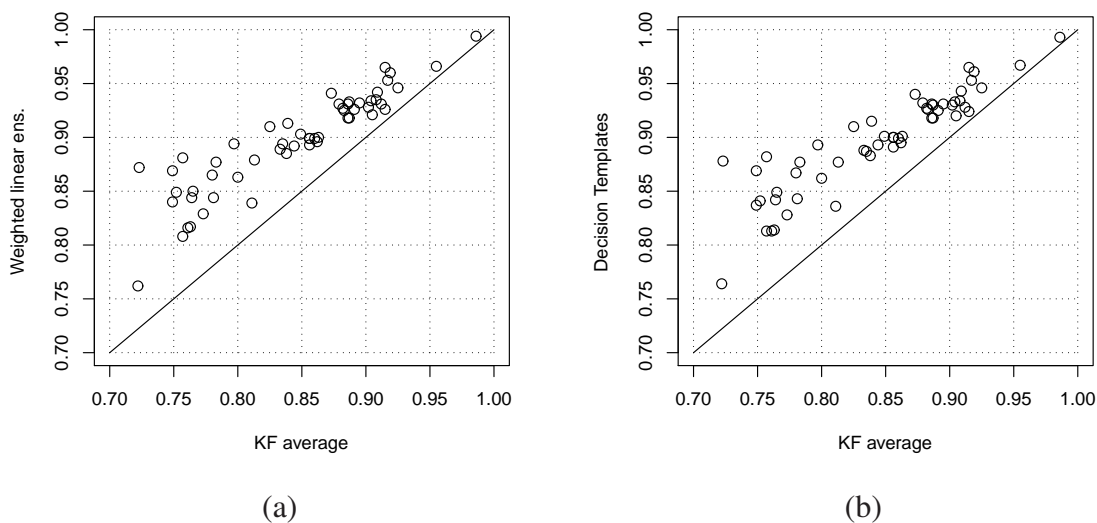


Figure 3: Comparison of AUC results between average kernel fusion and ensemble data integration methods with 2 noisy data sets. Points represent the AUC score of kernel fusion (abscissa) and ensemble (ordinate) methods for 56 GO terms. (a) Average KF vs Weighted ensembles (b) Average KF vs Decision Templates.

4.2 Results with six different types of biomolecular data

Figure 4 presents in a synthetic way the results obtained with no-noisy and noisy data by average KF and the ensemble methods described in Section 2. Each point represents the AUC score for a specific FunCat class, achieved with no-noisy (circular points) and noisy data (triangular points) using KF and respectively weighted linear (Figure 4 (a)) and decision templates (Figure 4 (b)) ensembles. Full results for each ensemble, kernel fusion and single SVMs trained on a single source of data for each FunCat class are available in Table 4 in the Supplementary Information.

5 Discussion

Considering no noisy data integration, the performances obtained by the weighted average ensembles appear comparable with those achieved by the average kernel fusion (Fig. 1(a) and Fig. 4 (a)), while decision templates achieve comparable (Fig. 1 (b)) or slightly worse (Fig. 4 (b)) performances w.r.t. KF. These visual clues are confirmed by the Wilcoxon signed-ranks test [24]: at 0.01 significance level no difference can be registered between weighted linear and KF methods for both the considered tasks, while a significant difference can be reported in favour of KF with respect to decision templates only for the second task (integration of the six original no-noisy data sets).

The AUC scores obtained by the ensemble and kernel fusion methods when one noisy kernel is added to the structure and mismatch kernels are compared in Fig. 2. Looking at the data plotted in Fig. 2, it is clear that, in presence of one noisy kernel, the performances achieved by the ensemble systems are higher than the ones obtained by using the average KF approach.

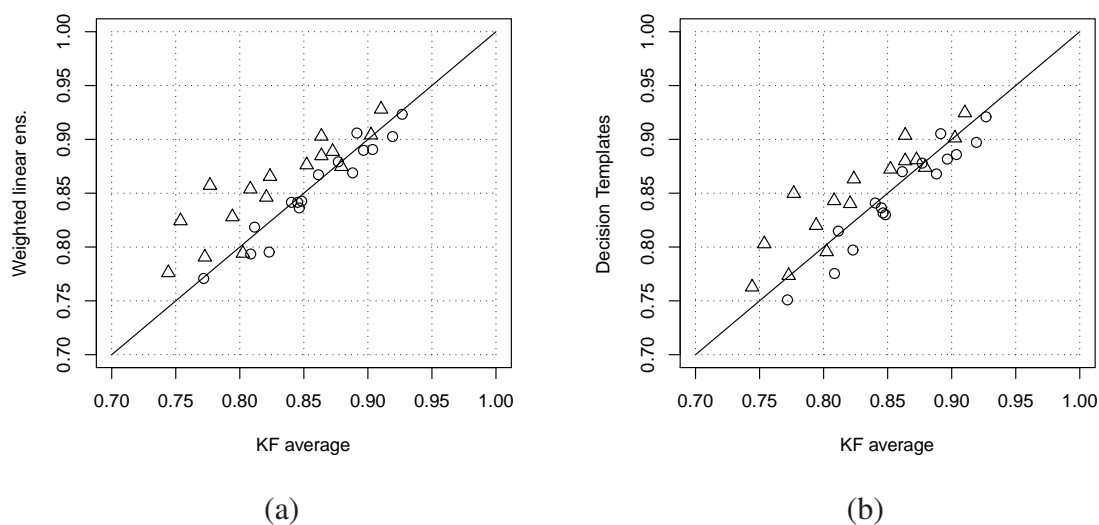


Figure 4: Experiments with 6 no-noisy and 6 noisy data sets. Points represent the AUC score of kernel fusion (abscissa) and ensemble (ordinate) methods for 15 FunCat functional classes. Circular points: results with no-noisy data. Triangular points: results with noisy data. (a) Average KF vs Weighted ensembles (b) Average KF vs Decision Templates.

Most of the points lie above the bisector in Fig. 2 (a) and (b), showing that both Weighted linear and Decision Templates ensembles tend to outperform the average based kernel fusion. Using the Wilcoxon signed-ranks test we register a significant difference in favour of Weighted linear and Decision Template w.r.t. average KF ($p\text{-value} \simeq 4.30 \cdot 10^{-11}$). The resistance to noise of the tested ensemble methods is also confirmed by the results collected after the inclusion in the data sources collection of the second noisy kernel (Fig. 3). In this extreme test in which half of the considered datasets are merely representing noise, the AUC averaged across both the cross validation splits and the GO functional classes is 0.895, 0.894 and 0.841 for the weighted linear ensemble, the decision templates ensemble and the average kernel fusion respectively. Again the Wilcoxon signed-ranks test shows a significant difference in favour of Weighted linear and Decision Template w.r.t. average KF ($p\text{-value} \simeq 3.87 \cdot 10^{-11}$). These results indicate that ensemble systems are more robust to noisy data than the KF method.

The experiments with the larger group of biomolecular data sets enforce these findings. Indeed, while the results between the different integration methods are quite comparable with the no-noisy data (circular points, Fig. 4), when the six noisy data sets are added to the original data, ensemble methods clearly outperform KF (triangular points, Fig. 4). Interestingly enough, it seems that adding noisy data sets does not significantly worsen the performances of weighted linear ensembles ($p\text{-value} \simeq 0.803$) and decision templates ($p\text{-value} \simeq 0.561$), while a significant decrement can be detected in KF ($p\text{-value} \simeq 1.52 \cdot 10^{-5}$).

These results, at least for weighted linear ensembles, can be explained by considering that these methods are able to learn which are the noisy data and penalize them by assigning low weights to base classifiers trained with noisy data. From this standpoint, kernel fusion methods based on semi-definite or semi-infinite programming techniques [21] could be considered for future experiments, because they also are able to learn weights from the data.

6 Conclusions and future work

The impact of noisy data in the performances of data integration methods for gene function prediction is of paramount importance for at least two reasons. At first, most of the available data from high throughput biotechnologies are inherently noisy. At second, the same source of data can be predictive for certain functional classes, but can be non-predictive for others: from this standpoint a certain type of data can be considered "noisy" for the prediction of certain subsets of terms in the GO or FunCat hierarchy.

In this work we evaluated the performances of simple ensemble systems in data integration based gene function prediction problems with and without noisy data. We compared their results with those obtained using state-of-the-art methods and previously published datasets. We found that, in these empirical benchmarks based on both real data and artificially generated noisy datasets, the ensemble systems are able to perform better than average kernel fusion systems, one of the most widely applied data integration methods. It is worth noting that the tolerance of ensemble methods to noisy data sets is confirmed for the integration of both a small and a relatively large number of different sources of biomolecular data. Future research could try to evaluate the impact of an even large number of noisy data taking into account also the hierarchical structure of gene functional classes.

Considering that ensemble systems are often less computationally intensive than other data integration approaches and that they perform data integration at decision level, avoiding all the potential problems affecting experiments aimed at integrating structurally different types of data, we believe that this class of data integration approaches constitutes a valuable research line in data integration based gene function prediction.

Motivated by these promising results, we plan to experiment with other more refined ensemble systems based on meta-learning or boosting techniques to integrate multiple sources of noisy data and to compare ensemble methods with other Multi-Kernel learning algorithms that showed tolerance to noisy data [11].

Another interesting research item could be the evaluation of the resistance to noisy datasets by gradually varying the amount of artificial noise. In this way we could characterize at a more refined level the response of data integration methods to noise, and we could better understand the behaviour of these methods in real-world whole-genome and whole-ontology gene function prediction experiments.

7 Acknowledgments

We would like to thank the reviewers for their comments. The authors gratefully acknowledge partial support by Università degli Studi di Milano (PUR project ATE-2009-0356, "Computational methods for biomedical pattern analysis") and by PASCAL2 Network of Excellence under EC grant no. 216886. This publication only reflects the authors' views.

References

- [1] 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.
- [2] N.C. de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, 1785.
- [3] 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.
- [4] 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.
- [5] X. Jiang, N. Nariai, M. Steffen, S. Kasif, and E. Kolaczyk. Integration of relational and hierarchical network information for protein function prediction. *BMC Bioinformatics*, 9(350), 2008.
- [6] U. Karaoz et al. Whole-genome annotation by using evidence integration in functional-linkage networks. *Proc. Natl Acad. Sci. USA*, 101:2888–2893, 2004.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] J. Mercer. Function of positive and negative type and their connection with the theory of integral equations. *Philos. Trans. Roy. Soc. London*, 209:415–446, 1909.
- [14] 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.

- [15] G. Obozinski, G. Lanckriet, C. Grant, Jordan. M., and W.S. Noble. Consistent probabilistic output for protein function prediction. *Genome Biology*, 9(S6), 2008.
- [16] A.R. Ortiz et al. MAMMOTH (matching molecular models obtained from theory): an automated method for model comparison. *Protein Sci.*, 11:2606–2621, 2002.
- [17] M. Re and G. Valentini. Simple ensemble methods are competitive with state-of-the-art data integration methods for gene function prediction. *Journal of Machine Learning Research, W&C Proceedings*, 2010. (in press).
- [18] S.Y Rhee et al. Use and misuse of the gene ontology annotations. *Nature Rev. Genetics*, 9(17-18):509–515, 2008.
- [19] 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.
- [20] B. Scholkopf, K. Tsuda, and J.P. Vert. *Kernel Methods in Computational Biology*. MIT Press, Cambridge, MA, 2004.
- [21] S. Sonnenburg, G. Ratsch, C. Schafer, and B. Scholkopf. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, 2006.
- [22] I. Tetko, I. Rodchenkov, M. Walter, T. Rattei, and H. Mewes. Beyond the ‘best’ match: machine learning annotation of protein sequences by integration of different sources of information. *Bioinformatics*, 24(5):621–628, 2008.
- [23] The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nature Genet.*, 25:25–29, 2000.
- [24] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics*, 1:80–83, 1945.