

Hierarchical ensemble methods for protein function prediction

Giorgio Valentini
AnacletoLab - Dipartimento di Informatica
Università degli Studi di Milano
e-mail: valentini@di.unimi.it

Abstract

Protein function prediction is a complex multi-class multi-label classification problem, characterized by multiple issues such as the incompleteness of the available annotations, the integration of multiple sources of high dimensional biomolecular data, the unbalance of several functional classes, and the difficulty of univocally determining negative examples. Moreover the hierarchical relationships between functional classes, that characterize both the Gene Ontology and FunCat taxonomies, motivate the development of hierarchy-aware prediction methods that showed significantly better performances than hierarchical-unaware “flat” prediction methods. In this paper we provide a comprehensive review of hierarchical methods for protein function prediction based on ensembles of learning machines. According to this general approach a separate learning machine is trained to learn a specific functional term and then the resulting predictions are assembled in a “consensus” ensemble decision, taking into account the hierarchical relationships between classes. The main hierarchical ensemble methods proposed in the literature are discussed in the context of existing computational methods for protein function prediction, highlighting their characteristics, advantages and limitations. Open problems of this exciting research area of computational biology are finally considered, outlining novel perspectives for future research.

1 Introduction

Exploiting the wealth of biomolecular data accumulated by novel high-throughput biotechnologies, “in silico” protein function prediction can generate hypotheses to drive the biological discovery and validation of protein functions [69]. Indeed, “in vitro” methods are costly in time and money, and automatic prediction methods can support the biologist in understanding the role of a protein or of a biological process, or in annotating a new genome at high level of accuracy, or more in general in solving problems of functional genomics [125].

The *Automated Function Prediction (AFP)* is a multiclass, multilabel classification problem characterized by hundreds or thousands of functional classes structured according to a predefined hierarchy. Even in principle also unsupervised methods can be

applied to *AFP*, due to the inherent difficulty of extracting functional classes without exploiting any available a priori information [177, 19], usually supervised or semi-supervised learning methods are applied in order to exploit the available a priori information about gene annotations.

From a computational standpoint *AFP* is a challenging problem for several reasons:

- *The number of functional classes is usually large*: hundreds for the Functional Catalogue (FunCat) [146] or thousands for the Gene Ontology (GO) [167].
- *Proteins may be annotated for multiple functional classes*: since each protein may belong to more than one class at the same time, the classification problem is multilabel.
- *Multiple sources of data are available for each protein*: high-throughput biotechnologies make available an increasing number of sources of genomic and proteomic data. Hence, in order to exploit all the information available for each protein, we need learning methods that are able to integrate different data sources [176].
- *Functional classes are hierarchically related*: annotations are not independent because functional classes are hierarchically organized; in general, known functional relationships (such as taxonomies) can be exploited to incorporate a priori knowledge in learning algorithms or to introduce explicit constraints between labels.
- *Small number of annotations for each class*: typically functional classes are severely unbalanced, with a small number of available “positive” annotations.
- *Multiple possible definitions of negative examples*: since we only have positive annotations¹, the notion of negative example is not uniquely determined, and different strategies of choosing negative examples can be in principle applied [195].
- *Different reliability of functional labels*: functional annotations have different degrees of evidence; that is, each label is assigned to a gene with a specific level of reliability.
- *Complex and noisy data*: data are usually complex (e.g., high-dimensional, large-scale, graph-structured) and noisy.

Most of the computational methods for *AFP* have been applied to unicellular organisms (e.g. *S. cerevisiae*) [50, 21, 100], but recently several approaches have been applied to multi-cellular organisms (such as *M. musculus* or the *A. thaliana* plant model organisms [125, 168, 85, 114, 102, 129]).

Several computational approaches, and in particular machine learning methods, have been proposed to deal with the above issues, ranging from sequence-based methods [83], to network-based methods [153], structured output algorithm based on kernels [160] and hierarchical ensemble methods [12].

¹The total number of GO negative annotations is about 2500, considering all species (August 2013).

Other approaches focused primarily on the integration of multiple sources of data, since each type of genomic data captures only some aspects of the genes to be classified, and a specific source can be useful to learn a specific functional class while being irrelevant to others. In the literature, many approaches have been proposed to deal with this topic: for example, functional linkage networks integration [49], kernel fusion [100], vector space integration [124], and ensemble systems [137].

Extensive experimental studies showed that flat prediction, i.e. predictions for each class made independently of the other classes, introduce significant inconsistencies in the classification, due to the violation of the *true path rule*, that governs the functional annotations of genes both in the GO and in FunCat taxonomies [118]. According to this rule, positive predictions for a given term must be transferred to its “ancestor” terms and negative predictions to its descendants (see Appendix A and Section 7 for more details about the GO and the true path rule). Moreover flat predictions are difficult to interpret because they may be inconsistent with one another. A method that claims, for example, that a protein has homodimerization activity but does not have dimerization activity is clearly incorrect, and a biologist attempting to interpret these results would likely not trust either prediction [118].

It is worth noting that the results of the Critical Assessment of Functional Annotation (CAFA) challenge, a recent comprehensive critical assessment and comparison of different computational methods for AFP [129], showed that AFP is characterized by multiple complex issues, and one of the best performing CAFA methods corrected flat predictions taking into account the hierarchical relationships between functional terms, with an approach similar to that adopted by hierarchical ensemble methods [52]. Indeed hierarchical ensemble methods embed in the learning process the relationships between functional classes. Usually this is performed in a second “reconciliation” step, where the predictions are modified to make them consistent with the ontology [82, 43, 36, 151]. More in general these methods exploit the relationships between ontology terms, structured according to a forest of trees [146] or a directed acyclic graph [167] to significantly improve prediction performances with respect to “flat” prediction methods [76, 179, 3].

Hierarchical classification and in particular ensemble methods for hierarchical classification have been applied in several domains different from protein function prediction, ranging from text categorization [92, 44, 200] to music genre classification [32, 54, 171], hierarchical image classification [22, 11] and video annotation [62], and automatic classification of World Wide Web documents [127, 33]. The present review focuses on hierarchical ensemble methods for AFP. For a more general review on hierarchical classification methods and their applications in different domains, see [155].

The paper is structured as follows. In Section 2 we provide a synthetic picture of the main categories of protein function methods, to properly position hierarchical ensemble methods in the context of computational methods for AFP. In section 3 the main common characteristics of hierarchical ensemble algorithms, as well as a general taxonomy of these methods are proposed. The following five sections focus on the main families of hierarchical methods for AFP, and discuss their main characteristics: Section 4 introduces Hierarchical Top-Down methods, Section 5 Bayesian ensemble approaches, Section 6 Reconciliation methods, Section 7 True Path Rule ensemble methods and the last one (Section 8) ensembles based on decision trees. Section 9 critically dis-

cuss the main issues and limitations of hierarchical ensemble methods and shows that this approach, such as the other current approaches for *AFP*, cannot be successfully applied without considering the large set of complex learning issues that characterize the *AFP* problem. The last two sections discuss the open problems and future possible research lines in the context of hierarchical ensemble methods and summarize the main findings in this exciting research area. In the Appendix some basic information about the FunCat and the GO, i.e. the two main hierarchical ontologies that are widely used to annotate proteins in all organisms, are provided, as well as the characteristics of the hierarchical-aware performance measures proposed in the literature to assess the accuracy and the reliability of the predictions made by hierarchical computational methods.

2 A taxonomy of protein function prediction methods

Several computational methods for the *AFP* problem have been proposed in the literature. Some methods provided predictions a relatively small set of functional classes [100, 172, 175], while others considered predictions extended to larger sets, using Support Vector Machines and semidefinite programming [100], artificial neural networks [193], functional linkage networks [84, 49], Bayesian Networks [172], or methods that combine functional linkage networks with learning machines using a logistic regression model [168] or simple algebraic operators [85].

Other research lines for *AFP* explicitly take into account the hierarchical nature of the multi-label classification problem. For instance, structured-output methods are based on the joint kernelization of both input variables and output labels, using e.g. perceptron-like learning algorithms [159] or maximum-margin algorithms [7]. Other approaches improve the prediction of GO annotations by extracting implicit semantic relationships between genes and functions [63]. Finally other methods adopted an ensemble approach [181] to take advantage of the intrinsic hierarchical nature of protein function prediction, explicitly considering the relationships between functional classes [152, 185, 118, 178].

Computational methods for *AFP*, mostly based on machine learning methods, can be schematically grouped in four families:

1. Sequence-based methods
2. Network-based methods
3. Kernel methods for structured output spaces
4. Hierarchical ensemble methods

This grouping is neither exhaustive nor strict, meaning that certain methods do not belong to any of these groups, and others belong to more than one.

2.1 Sequence-based methods

Algorithms based on alignment of sequences represent the first attempts to computationally predict the function of proteins [5, 6]: similar sequences are likely to share

common functions, even if it is well-known that secondary and tertiary structure conservation are usually more strictly related to protein functions. However algorithms able to infer similarities between sequences are today standard methods of assigning functions to proteins in newly sequenced organisms [51, 83]. Of course, global or local structure comparison algorithms between proteins can be applied to detect functional properties [105], and in this context, the integration of different sequence and structure-based prediction methods represents a major challenge [126].

Even if most of the research efforts for the design and development of AFP methods concentrated on machine learning methods, it is worth noting that in the AFP 2011 challenge [129] one of the best performing method is represented by a sequence-based algorithm [65]. Indeed, when the only information available is represented by a raw sequence of aminoacids or nucleotides, sequence-based methods can be competitive with state-of-the-art machine learning methods by exploiting homology-based inference [78].

2.2 Network-based methods.

These methods usually represent each dataset through an undirected graph $G = (V, E)$, where nodes $v \in V$ correspond to gene/gene products, and edges $e \in E$ are weighted according to the evidence of co-functionality implied by data source [106, 184]. These algorithms are able to transfer annotations from previously annotated (labeled) nodes to unannotated (unlabeled) ones by exploiting “proximity relationships” between connected nodes. Basically, these methods are based on transductive label propagation algorithms that predict the labels of unannotated examples without using a global predictive model [172, 49, 114]. Several method exploited the semantic similarity between GO terms [191, 194] to derive functional similarity measures between genes to construct functional terms, using then supervised or semi-supervised learning algorithm to infer GO annotations of genes [199, 166, 123, 201].

Different strategies to learn the unlabeled nodes have been explored by “label propagation” algorithms, that is methods able to “propagate” the labels of annotated proteins across the networks, by exploiting the topology of the underlying graph. For instance, methods based on the evaluation of the functional flow in graphs [184, 116], methods based on Hopfield networks [84, 17, 67], methods based on Markov [57, 93] and Gaussian Random Fields [175, 114], but also simple “guilt-by-association” methods [120, 107], based on the assumption that connected nodes/proteins in the functional networks are likely to share the same functions. Recently methods based on kernelized score functions, able to exploit both local and global semi-supervised learning strategies have been successfully applied to *AFP* [131] as well as to disease gene prioritization [138] and drug repositioning problems [140, 141].

[16] showed that different graph-based algorithms can be cast into a common framework where a quadratic cost objective function is minimized. In this framework closed form solutions can be derived by solving a linear system of size equal to the cardinality of nodes (proteins), or using fast iterative procedures such as the Jacobi method [148]. A network-based approach, alternative to label propagation and exhibiting strong theoretical predictive guarantees in the so-called mistake bound model, has been recently proposed by [39].

2.3 Kernel methods for structured output spaces.

By extending kernels to the output space, the multilabel hierarchical classification problem is solved globally: the multilabels are viewed as elements of a structured space modeled by suitable kernel functions [173, 144, 98], and structured predictions are viewed as a maximum a-posteriori prediction problem [8].

Given a feature space \mathcal{X} and a space of structured labels \mathcal{Y} , the task is to learn a mapping $f : \mathcal{X} \rightarrow \mathcal{Y}$ by an induced joint kernel function k that computes the “compatibility” of a given input-output pair (x, y) : for each test example $x \in \mathcal{X}$ we need to determine the label $\bar{y} \in \mathcal{Y}$ such that $\bar{y} = \arg \max_{y \in \mathcal{Y}} k(x, y)$, for any $x \in \mathcal{X}$. By modeling probabilities by a log-linear model, and using a suitable feature map $\phi(x, y)$, we can define an induced joint kernel function that uses both inputs and outputs to compute the “compatibility” of a given input-output pair [8]:

$$k : (\mathcal{X} \times \mathcal{Y}) \times (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R} \quad (1)$$

Structured output methods infer a label \hat{y} by finding the maximum of a function g that uses the previously defined joint kernel (1):

$$\hat{y} = \arg \max_{y \in \mathcal{Y}} g(x, y) \quad (2)$$

The *GOstruct* system implemented a structured Perceptron, and a variant of the structured support vector machine [173]. This approach has been successfully applied to the prediction of GO terms in mouse and other model organisms [160]. Structured output maximum-margin algorithms have been also applied to the tree-structured prediction of enzyme functions [7, 144].

2.4 Hierarchical ensemble methods

Other approaches take explicitly into account the hierarchical relationships between functional terms [64, 24, 152, 185, 82, 151]. Usually they modify the “flat” predictions (i.e. predictions made independently of the hierarchical structure of the classes) and correct them improving accuracy and consistency of the multilabel annotations of proteins [118].

The flat approach makes predictions for each term independently and, consequently, the predictor may assign to a single protein a set of terms that are inconsistent with one another. A possible solution for this problem is to train a classifier for each term of the reference ontology, to produce a set of prediction at each term and, finally, to reconcile the predictions by tacking into account the relationships between the classes of the ontology. Different ensemble based algorithms have been proposed ranging from methods restricted to multilabels with single and no partial paths [55] to methods extended to multiple and also partial paths [40]. Many recent published works clearly demonstrated that this approach ensures an increment in precision, but this comes at expenses of the overall recall [125, 76].

In the next section we discuss in detail hierarchical ensemble methods, since they constitute the main topic of this review.




Figure 1: Ensemble of classifiers

3 Hierarchical ensemble methods: exploiting the hierarchy to improve protein function prediction

Ensemble methods are one of the main research areas of machine learning [60, 181, 119, 139]. From a general standpoint ensembles of classifiers are sets of learning machines that work together to solve a classification problem (Fig. 1). Empirical studies showed that both in classification and regression problems ensembles improves on single learning machines, and moreover large experimental studies compared the effectiveness of different ensemble methods on benchmark data sets [14, 61, 9, 158], and they have been successfully applied to several computational biology problems [59, 143, 182, 1, 145]. Ensemble methods have been also successfully applied in an unsupervised setting [170, 18]. Several theories have been proposed to explain the characteristics and the successful application of ensembles to different application domains. For instance, Allwein, Schapire and Singer interpreted the improved generalization capabilities of ensembles of learning machines in the framework of large margin classifiers [150, 4], Kleinberg in the context of Stochastic Discrimination Theory [89], and Breiman and Friedman in the light of the bias–variance analysis borrowed from classical statistics [29, 70]. The interest in this research area is motivated also by the availability of very fast computers and networks of workstations at a relatively low cost that allow the implementation and the experimentation of complex ensemble methods using off-the-shelf computer platforms.

Constraints between labels and, more in general, the issue of label dependence have been recognized to play a central role in multilabel learning [56]. Protein function prediction can be regarded as a paradigmatic multilabel classification problem, where the exploitation of a priori knowledge about the hierarchical relationships between the labels can dramatically improve classification performance [118, 112, 43].

In the context of AFP problems, ensemble methods reflect the hierarchy of functional terms in the structure of the ensemble itself: each base learner is associated to a node of the graph representing the functional hierarchy and learns a specific GO term or FunCat category. The predictions provided by the trained classifiers are then combined by exploiting the hierarchical relationships of the taxonomy.

In their more general form hierarchical ensemble methods adopt a two-steps learn-

ing strategy:

1. In the first step each base learner separately or interacting with connected base learners learns the protein functional category on a per-term basis. In most cases this yields a set of independent classification problems, where each base learning machine is trained to learn a specific functional term, independently of the other base learners.
2. In the second step the predictions provided by the trained classifiers are combined by considering the hierarchical relationships between the base classifiers modeled according to the hierarchy of the functional classes.

Fig. 2 depicts the two learning steps of hierarchical ensemble methods. In the first step a learning algorithm (a square object in Fig. 2 a) is applied to train the base classifiers associated to each class (represented with numbers from 1 to 9). Then the resulting base classifiers (circles) in the prediction phase exploit the hierarchical relationships between classes to combine its predictions with those provided by the other base classifiers (Fig. 2 b). Note that the dummy 0 node is added to obtain a rooted hierarchy. Up and down arrows represent the possibility of combining predictions by exploiting those provided respectively by children and parents classifiers, according to a bottom-up or top-down learning strategy. Note that both “local” combinations are possible (e.g. the prediction of node 5 may depend only on the prediction of node 1), but also “global” combinations can be considered, by taking into account the predictions across the overall structure of the graph (e.g. predictions for node 9 can depend on all the predictions made by all the other base classifiers from 1 to 8). Moreover both top-down propagation of the predictions (down arrows, Fig. 2 b) and bottom-up propagation (up arrows) can be considered, depending on the specific design of the hierarchical ensemble algorithm.




Figure 2: Schematic representation of the two main learning steps of hierarchical ensemble methods. (a) Training of base classifiers (b) Top-down and/or bottom up propagation of the predictions.

This ensemble approach is highly modular: in principle any learning algorithm can be used to train the classifiers in the first step, and both annotation decisions, probabilities or whatever scores provided by each base learner can be combined, depending on the characteristics of the specific hierarchical ensemble method.

In this section we provide some basic notation and an ensemble taxonomy that will be used to introduce the different hierarchical ensemble methods for AFP.

3.1 Basic notation

A gene/gene product g can be represented through a vector $x \in \mathbb{R}^d$ having d different features (e.g., gene expression levels across d different conditions, sequence similarities with other genes/proteins, or presence or absence of a given domain in the corresponding protein or genetic or physical interaction with other proteins). Note that we, for a sake of simplicity and with a certain approximation, we refer in the same way to genes and proteins, even if it is well-known that a given gene may correspond to multiple proteins. A gene g is assigned to one or more functional classes in the set $C = \{c_1, c_2, \dots, c_m\}$ structured according to a FunCat forest of trees T or a directed acyclic graph G of the Gene Ontology². The assignments are coded through a vector of multilabels $y = (y_1, y_2, \dots, y_m) \in \{0, 1\}^m$, where g belongs to class c_i if and only if $y_i = 1$.

In both the *Gene Ontology* (*GO*) and *FunCat* taxonomies the functional classes are structured according to a hierarchy and can be represented by a directed graph, where nodes correspond to classes, and edges to relationships between classes. Hence the node corresponding to the class c_i can be simply denoted by i . We represent the set of children nodes of i by $\text{child}(i)$, and the set of its parents by $\text{par}(i)$. Moreover $y_{\text{child}(i)}$ denotes the labels of the children classes of node i and analogously $y_{\text{par}(i)}$ denotes the labels of the parent classes of i . Note that in FunCat only one parent is permitted, since the overall hierarchy is a tree forest, while in the GO, more parents are allowed, because the relationships are structured according to a directed acyclic graph.

Hierarchical ensemble methods train a set of calibrated classifiers, one for each node of the taxonomy T . These classifiers are used to derive estimates $\hat{p}_i(g)$ of the probabilities $p_i(g) = \mathbb{P}(V_i = 1 \mid V_{\text{par}(i)} = 1, g)$ for all g and i , where $(V_1, \dots, V_m) \in \{0, 1\}^m$ is the vector random variable modeling the unknown multilabel of a gene g , and $V_{\text{par}(i)}$ denotes the random variables associated to the parents of node i . Note that $p_i(g)$ are probabilities conditioned to $V_{\text{par}(i)} = 1$, that is the probability that a gene is annotated to a given term i , given that the gene is just annotated to its parent terms, thus respecting the true path rule. Ensemble methods infer a multilabel assignment $\hat{y} = (\hat{y}_1, \dots, \hat{y}_m) \in \{0, 1\}^m$ based on estimates $\hat{p}_1(g), \dots, \hat{p}_m(g)$.

3.2 A taxonomy of hierarchical ensemble methods

Hierarchical ensemble methods for AFP share several characteristics, from the two-steps learning approach to the exploitation of the hierarchical relationships between classes. For these reasons is quite difficult to clearly and univocally individuate a taxonomy of hierarchical ensemble methods. Here we show a taxonomy useful mainly to describe and discuss existing methods for AFP. For a recent review and taxonomy of hierarchical ensemble methods, not specific for AFP problems, we refer the reader to the comprehensive Silla's and others review [155].

²Usually a dummy root class c_0 , which every gene belongs to, is added to T or G to facilitate the processing

In the following sections we discuss the following groups of hierarchical ensemble methods:

- Top-down ensemble methods. These methods are characterized by a simple top-down approach in the second step: only the output of the parent node/base classifier influences the output of the children, thus resulting in a top-down propagation of the decisions.
- Bayesian ensemble methods. These are a class of methods theoretically well-founded and in some cases they are optimal from a Bayesian standpoint.
- Reconciliation methods. This is a heterogeneous class of heuristics by which we can combine the predictions of the base learners, by adopting different “local” or “global” combination strategies.
- True Path Rule ensembles. These methods adopt an heuristic approach based on the “true path rule” that governs both the GO and FunCat ontologies.
- Decision tree-based ensembles. These methods are characterized by the application of decision trees as base learners or by adopting decision tree-like learning strategies to combine predictions of the base learners.

Despite this general characterization, several methods could be assigned to different groups, and for several hierarchical ensemble methods it is difficult to assign them to any the introduced classes of methods.

For instance, in [81, 97, 20] the authors used the hierarchy only to construct training sets different for each term of the Gene Ontology, by determining positive and negative examples on the basis of the relationships between functional terms. In [64] for each classifier associated with a node, a gene is labeled as positive (i.e. belonging to the term associated with that node) if it actually belongs to that node, or as negative if it does not belong to that node or to the ancestors or descendants of the node.

Other approaches exploited the correlation between nearby classes [152, 26, 3]. Shahbaba et al. [152] take into account the hierarchy to introduce correlation between functional classes, using a multinomial logit model with Bayesian priors in the context of *E. coli* functional classification with Riley’s hierarchies [142]. Bogdanov et al. incorporated functional interrelationships between terms during the extraction of features based on annotations of neighbouring genes and then applied a nearest-neighbour classifier to predict protein functions [26]. The HiBLADE method (Hierarchical multi-label Boosting with LAbel DEpendency) [3], takes advantage of not only the pre-established hierarchical taxonomy of the classes, but also effectively exploits the hidden correlation among the classes that is not shown through the class hierarchy, thereby improving the quality of the predictions. In particular the dependencies of the children for each label in the hierarchy are captured and analyzed using Bayes method and instance-based similarity. Experiments using the FunCat taxonomy and the yeast model organism show that the proposed method is competitive with TPR-w (Section 7.2) and HABYES-CS (Section 5.3) hierarchical ensemble methods.

An adaptation of a classical multiclass boosting algorithm [149] has been adapted to fit the hierarchical structure of the FunCat taxonomy [2]: the method is relatively

simple and straightforward to be implemented and achieves competitive results for the AFP in the yeast model organism.

Finally other hierarchical approaches have been proposed in the context of competitive networks learning framework. Competitive networks are well-known unsupervised and supervised methods able to map the input space into a structured output space where clusters or classes are usually arranged according to a grid topology, and where learning adopts at the same way a competition, cooperation and adaptation strategy [91]. Interestingly enough, in [27] the authors adopted this approach to predict the hierarchy of gene annotations in the yeast model organism, by using a tree-topology according to the FunCat taxonomy: each neuron is connected with its parent or with its children. Moreover each neuron in tree-structured output layer is connected to all neurons of the input layer, representing the instances, i.e. the set of genomic features associated to each gene to be classified. Results obtained with the hierarchy of Enzyme Commission Codes showed that this approach is competitive with those obtained with hierarchical decision trees ensembles [151] (Section. 8).

To provide a general picture of the methods discussed in the following sections, Table 1 summarizes their main characteristics. The first two columns report the name and a reference to the method, the third whether multiple or single paths across the taxonomy are allowed, the next whether partial paths are considered (i.e. paths that do not end with a leaf). The successive columns refer to the class structure (a tree or a DAG), to the adoption or not of cost-sensitive (i.e. unbalance-aware) classification approaches, to the adoption of strategies to properly select negative examples in the training phase. Finally the last three columns summarize the type of the base learner used (“spec” means that only a specific type of base learner is allowed, “any” that any type of learner can be used within the method), whether the method improves or not with respect to the flat approach and the mode of processing of the nodes (“TD” – top-down approach, and “TD&BUP” – adopting both top-down and bottom-up strategies). Of course methods having more checkmarks are more flexible and in general methods that can process a DAG can also process tree-structured ontologies, but the viceversa is not guaranteed, while the type of node processing relies on the way the information is propagated across the ontology. It is worth noting that all the considered methods improves on baseline “flat” classification methods.

4 Hierarchical Top-Down (HTD) ensembles

These ensemble methods exploit the hierarchical relationships between functional terms in a top-to-bottom fashion, i.e. considering only the relationships denoted by the down arrows in Fig. 2 b. The basic Hierarchical Top-Down ensemble method (HTD) algorithm is straightforward: for each gene g , starting from the set of nodes at the first level of the graph G (denoted by $\text{root}(G)$), the classifier associated to the node $i \in G$ computes whether the gene belongs to the class c_i . If yes, the classification process continues recursively on the nodes $j \in \text{child}(i)$; otherwise, it stops at node i , and the nodes belonging to the descendants rooted at i are all set to 0. To introduce the method we use probabilistic classifiers as base learners trained to predict class c_i associated to the node i of the hierarchical taxonomy. Their estimates $\hat{p}_i(g)$ of $\mathbb{P}(V_i = 1 \mid V_{\text{par}(i)} = 1, g)$

Table 1: Characteristics of some of the main hierarchical ensemble methods for AFP

Methods	Ref.	multi path	partial path	class struct	cost sens	sel neg	base learn	improves on flat	node process
HMC-LMLP	[35, 37]	✓	✓	TREE			ANY	✓	TD
HTD-CS	[43]	✓	✓	TREE	✓		ANY	✓	TD
HTD-MULTI	[79]		✓	TREE			ANY	✓	TD
HTD-PERLEV	[122]			TREE			SPEC	✓	TD
HTD-NET	[82]	✓	✓	DAG			ANY	✓	TD
BAYES NET-ENS	[12]	✓	✓	DAG		✓	SPEC	✓	TD&BUP
HIER-MB AND BFS	[76]	✓	✓	DAG			ANY	✓	TD&BUP
HBAYES	[38, 40]	✓	✓	TREE		✓	ANY	✓	TD&BUP
HBAYES-CS	[42]	✓	✓	TREE	✓	✓	ANY	✓	TD&BUP
RECONC-HEURISTIC	[118]	✓	✓	DAG			ANY	✓	TD
CASCADED LOG	[118]	✓	✓	DAG			ANY	✓	TD
PROJECTION-BASED	[118]	✓	✓	DAG			ANY	✓	TD&BUP
TPR	[183, 179]	✓	✓	TREE		✓	ANY	✓	TD&BUP
TPR-w	[179]	✓	✓	TREE	✓	✓	ANY	✓	TD&BUP
TPR-w WEIGHTED	[47]	✓	✓	TREE	✓		ANY	✓	TD&BUP
DECISION-TREE-ENS	[151]	✓	✓	DAG			SPEC	✓	TD&BUP

are used by the HTD ensemble to classify a gene g as follows

$$\hat{y}_i = \begin{cases} \{\hat{p}_i(g) > \frac{1}{2}\} & \text{if } i \in \text{root}(G) \\ \{\hat{p}_i(g) > \frac{1}{2}\} & \text{if } i \notin \text{root}(G) \wedge \{\hat{p}_{\text{par}(i)}(g) > \frac{1}{2}\} \\ 0 & \text{if } i \notin \text{root}(G) \wedge \{\hat{p}_{\text{par}(i)}(g) \leq \frac{1}{2}\} \end{cases} \quad (3)$$

where $\{x\} = 1$ if $x > 0$ otherwise $\{x\} = 0$, and $\hat{p}_{\text{par}(i)}$ is the probability predicted for the parent of the term i . It is easy to see that this procedure ensures that the predicted multilabels $\hat{y} = (\hat{y}_1, \dots, \hat{y}_m)$ are consistent with the hierarchy. We can apply the same top-down procedure also using non probabilistic classifiers, i.e. base learners generating continuous scores, or also discrete decisions, by slightly modifying (3).

In [42] a cost-sensitive version of the basic top-down hierarchical ensemble method HTD has been proposed: by assigning \hat{y}_i before the label of any j in the subtree rooted at i , the following rule is used:

$$\hat{y}_i = \left\{ \hat{p}_i \geq \frac{1}{2} \right\} \times \{ \hat{y}_{\text{par}(i)} = 1 \} \quad (4)$$

for $i = 1, \dots, m$ (note that the guessed label \hat{y}_0 of the root of G is always 1). Then the cost-sensitive variant HTD-CS introduces a single cost sensitive parameter $\tau > 0$ which replaces the threshold $\frac{1}{2}$. The resulting rule for HTD-CS is then

$$\hat{y}_i = \{ \hat{p}_i \geq \tau \} \times \{ \hat{y}_{\text{par}(i)} = 1 \} \quad (5)$$

By tuning τ we may obtain ensembles with different precision/recall characteristics. Despite the simplicity of the hierarchical top-down methods, several works showed their effectiveness for AFP problems [179, 36].

For instance, Cerri and de Carvalho experimented different variants of Top-down hierarchical ensemble methods for AFP [35, 36, 37]. The HMC-LMLP (Hierarchical

Multilabel Classification with Local Multilayer Perceptron) successively trains a local MLP network for each hierarchical level, using the classical backpropagation algorithm [147]. Then the output of the MLP for the first layer is used as input to train the MLP that learns the classes of the second level and so on (Fig. 3). A gene is annotated to a class if its corresponding output in the MLP is larger than a predefined threshold, then in a post-processing phase (second-step of the hierarchical classification) inconsistent predictions are removed (i.e. classes predicted without the prediction of their superclasses) [37]. In practice, instead of using a dichotomic classifier for each node,

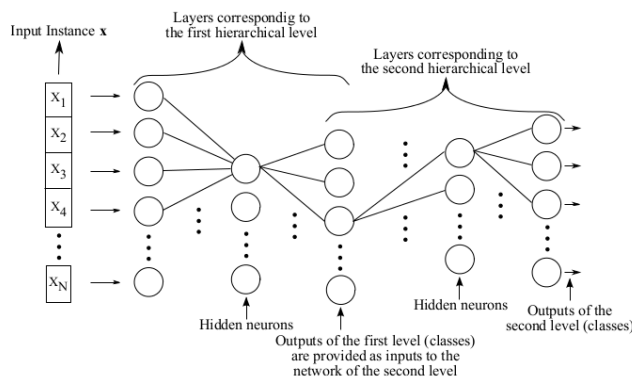


Figure 3: HMC-LMLP: Outputs of the MLP responsible for the predictions in the first level are used as input to another MLP for the predictions in the second level (adapted from [37]).

the HMC-LMLP algorithm applies a single multi-class multi-layer perceptron for each level of the hierarchy.

A related approach adopts multi-class classifiers (HTD-MULTI) for each node, instead of a simple binary classifier, and tries to find the most likely path from the root to the leaves of the hierarchy, considering simple techniques, such as the multiplication or the sum of the probabilities estimated at each node along the path [79]. The method has been applied to the Cell Cycle branch of the FunCat hierarchy with the yeast model organism, showing improvements with respect to classical Hierarchical Top-Down methods, even if the proposed approach can only predict classes along a single “most likely path”, thus not considering that in *AFP* we may have annotations involving multiple and partial paths.

Another method that introduces multi-class classifiers instead of simple dichotomic classifiers has been proposed by Paes et al. [122]: local per level multiclass classifiers (HTD-PERLEV) are trained to distinguish between the classes of a specific level of the hierarchy, and two different strategies to remove inconsistencies are introduced. The method has been applied to the hierarchical classification of enzymes using the EC taxonomy for the hierarchical classification of enzymes, but unfortunately this algorithm is not well-suited to *AFP*, since leaf nodes are mandatory (that is partial path annotations are not allowed) and multi-label annotations along multiple paths are not

allowed.

Another interesting Top-down hierarchical approach proposed by the same authors, is HMC-LP (Hierarchical Multilabel Classification Label-Powerset), a hierarchical variation of the *Label-Powerset* non-hierarchical multilabel method [174], that has been applied to the prediction of gene function of the yeast model organism using 10 different data sets and the FunCat taxonomy [35]. According to the *Label-Powerset* approach, the method is based on a first label-combination step by which, for each example (gene), all classes assigned to the example are combined into a new and unique class, and this process is repeated for each level of the hierarchy. In this way the original problem is transformed into a hierarchical single-label problem. In both the training and test phases the Top-Down approach is applied, and at the end of the classification phase the original classes can be easily reconstructed [35]. In an experimental comparison using the FunCat taxonomy for *S. cerevisiae* results showed that hierarchical Top-down ensemble methods significantly outperform decision trees based hierarchical methods, but no significant difference between different flavours of Top-down hierarchical ensembles has been detected [36].

Top-down algorithms can be conceived also in the context of network-based methods (HTD-NET). For instance, in [82] a probabilistic model, that combines relational protein-protein interaction data and the hierarchical structure of GO to predict true-path consistent function labels, obeys the true path rule by setting the descendants of a node as negative whenever that node is set to negative. More precisely the authors at first compute a local hierarchical conditional probability, in the sense that for any non-root GO term, only the parents affect its labelling. This probability is computed within a network-based framework assuming that the labeling of a gene is independent of any other genes given that of its neighbours (a sort of Markov property with respect to gene functional interaction networks), and assuming also a binomial distribution for the number of neighbours labeled with child terms with respect to those labeled with the parent term. These assumptions are quite stringent, but are necessary to make the model tractable. Then a global hierarchical conditional probability is computed by recursively applying the previously computed local hierarchical conditional probability by considering all the ancestors. More precisely, by assuming that $\mathbb{P}(\hat{y}_i = 1 | g, N(g))$, i.e. the probability that a gene g is annotated for a node i , given the status of the annotations of its neighborhood $N(g)$ in the functional networks, the global hierarchical conditional probability factorizes according to the GO graph:

$$\mathbb{P}(\hat{y}_i = 1 | g, N(g)) = \prod_{j \in \text{anc}(i)} \mathbb{P}(\hat{y}_j = 1 | \hat{y}_{\text{par}(j)} = 1, N_{loc}(g)) \quad (6)$$

where $N_{loc}(g)$ represents the local hierarchical neighborhood information on the parent-child GO term pair $\text{par}(j)$ and j [82]. This approach guarantees to produce GO term label assignments consistent with the hierarchy, without the need of a post-processing step.

Finally in [190] the author applied a hierarchical method to the classification of yeast FunCat categories. Despite its well-founded theoretical properties based on large margin methods, this approach is conceived for one path hierarchical classification, and hence it results unsuited for hierarchical AFP, where usually multiple paths in the

hierarchy should be considered, since in most cases genes can play different functional roles in the cell.

5 Ensemble based Bayesian approaches for hierarchical classification

These methods introduce a Bayesian approach to the hierarchical classification of proteins, by using the classical Bayes theorem, or Bayesian networks to obtain tractable factorizations of the joint conditional probabilities from the original “full Bayesian” setting of the hierarchical AFP problem [12, 76], or to achieve “Bayes-optimal” solutions with respect to loss functions well-suited to hierarchical problems [40, 43].

5.1 The solution based on Bayesian networks

One of the first approaches addressing the issue of inconsistent predictions in the Gene Ontology is represented by the Bayesian approach proposed in [12] (BAYES NET-ENS). According to the general scheme of hierarchical ensemble methods, two main steps characterize the algorithm:

1. Flat prediction of each term/class (possibly inconsistent)
2. Bayesian hierarchical combination scheme to allow collaborative error-correction over all nodes

After training a set of base classifiers on each of the considered GO terms (in their work the authors applied the method to 105 selected GO terms), we may have a set of (possibly inconsistent) \hat{y} predictions. The goal consists in finding a set of consistent y predictions, by maximizing the following equation derived from the Bayes theorem:

$$\mathbb{P}(y_1, \dots, y_n | \hat{y}_1, \dots, \hat{y}_n) = \frac{\mathbb{P}(\hat{y}_1, \dots, \hat{y}_n | y_1, \dots, y_n) \mathbb{P}(y_1, \dots, y_n)}{Z} \quad (7)$$

where n is the number of GO nodes/terms and Z is a constant normalization factor.

Since the direct solution of (7) is too hard, i.e. exponential in time w.r.t to the number of nodes, the authors proposed a Bayesian network structure to solve this difficult problem, in order to exploit the relationships between the GO terms. More precisely, to reduce the complexity of the problem, the authors imposed the following constraints:

1. y_i nodes conditioned to their children (GO structure constraints)
2. \hat{y}_i nodes conditioned on their label y_i (Bayes rule)
3. \hat{y}_i are independent from both $\hat{y}_j, i \neq j$ and $y_j, i \neq j$, given y_i

In other words we can ensure that a label is 1 (positive) when any one of its children is 1 and the edges from y_i to \hat{y}_i assure that a classifier output \hat{y}_i is a random variable independent of all other classifier outputs \hat{y}_j and labels y_j , given its true label y_i (Fig. 4).




Figure 4: Bayesian network involved in the hierarchical classification (adapted from [12]).

More precisely, from the previous constraints we can derive the following equations:

$$\text{from the first constraint: } \mathbb{P}(y_1, \dots, y_n) = \prod_{i=1}^n \mathbb{P}(y_i | \text{child}(y_i)) \quad (8)$$

$$\text{from the last two constraints: } \mathbb{P}(\hat{y}_1, \dots, \hat{y}_n | y_1, \dots, y_n) = \prod_{i=1}^n \mathbb{P}(\hat{y}_i | y_i) \quad (9)$$

Note that (8) can be inferred from training labels simply by counting, while (9) can be inferred by validation during training, by modeling the distribution of \hat{y}_i outputs over positive and negative examples, by assuming a parametric model (e.g. Gaussian distribution –see Fig. 5).

For the implementation of their method the authors adopted bagged ensemble of SVMs [28] to make their predictions more robust and reliable at each node of the GO hierarchy, and median values of their outputs on out-of-bag examples have been used to estimate means and variances for each class. Finally means and variances have been used as parameters of the gaussian models used to estimate the conditional probabilities of (9).

Results with the 105 terms/nodes of the GO BP (model organism *S.cerevisiae*) showed substantial improvements with respect to non hierarchical “flat” predictions: the hierarchical approach improves AUC results on 93 of the 105 GO terms (Fig. 6).

5.2 The Markov blanket and approximated breadth first solution

In [76] the authors proposed an alternative approximated solution to the complex equation (7) by introducing two variants of the Bayesian integration:

1. HIER-MB: Hierarchical Bayesian combination involving nodes in the Markov Blanket.

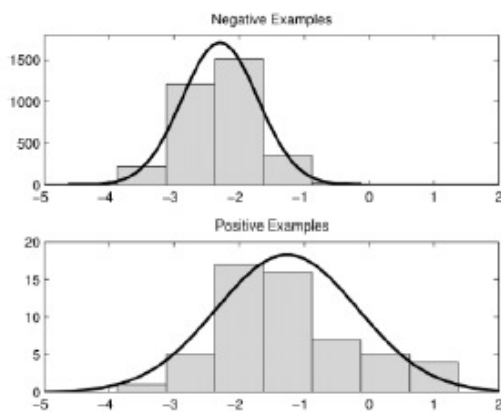


Figure 5: Distribution of positive and negative validation examples (a gaussian distribution is assumed – adapted from [12]).




Figure 6: Improvements induced by the hierarchical prediction of the GO terms. Darker shades of blue indicate largest improvements, and darker shades of red indicate largest deterioration; white means no change (adapted from [12]).

2. HIER-BFS: Hierarchical Bayesian combination involving the 30 first nodes visited through a Breadth-First-Search (BFS) in the GO graph.

The method has been applied to the prediction of more than 2000 GO terms for the mouse model organism, and performed among the top methods in the *MouseFunc* challenge [125].

The first approach (HIER-MB) modifies the output of the base learners (SVMs in the Guan et al. paper) taking into account the Bayesian network constructed using the Markov blanket surrounding the GO term of interest (Fig. 7). In a Bayesian network the Markov blanket of a node i is represented by its parents ($\text{par}(i)$), its children ($\text{child}(i)$) and its children's other parents. The Bayesian network involving the Markov blanket of node i is used to provide the prediction \hat{y}_i of the ensemble, thus leveraging the local relationships of node i and the predictions for the nodes included in its Markov blanket.




Figure 7: Markov blanket surrounding the GO term Y_1 . Each GO term is represented as a blank node, while the SVM classifier output is represented as a gray node (adapted from [76]).




Figure 8: The breadth-first subnetwork stemming from Y_1 . Each GO term is represented through a blank node and the SVM outputs as gray nodes (adapted from [76]).

To enlarge the size of the Bayesian subnetwork involved in the prediction of the node of interest, a variant based on the Bayesian networks constructed by applying a classical breadth-first search is the basis of the HIER-BFS algorithm. To reduce the complexity at most 30 terms are included (i.e. the first 30 nodes reached by the breadth-first algorithm). In the implementation ensembles of 25 SVMs have been trained for each node, using vector space integration techniques [58] to integrate multiple sources of data.

Note that with both HIER-MB and HIER-BFS methods we do not take into account the overall topology of the GO network, but only the terms related to the node for which we perform the prediction. Even if this general approach is reasonable and achieve good results, its main drawback is represented by the locality of the hierarchical integration (limited to the Markov blanket and the first 30 BFS nodes). Moreover in previous works it has been shown that the adopted integration strategy (vector space in-

tegration) is in most cases worse than Kernel fusion [100] and than ensemble methods for data integration [137].

In the same work [76] the authors propose also a sort of “test and select” method [154], by which three different classification approaches: a) single flat SVMs; b) Bayesian hierarchical correction; c) Naive Bayes combination are applied, and for each GO term the best one is selected by internal cross-validation (Fig. 9).

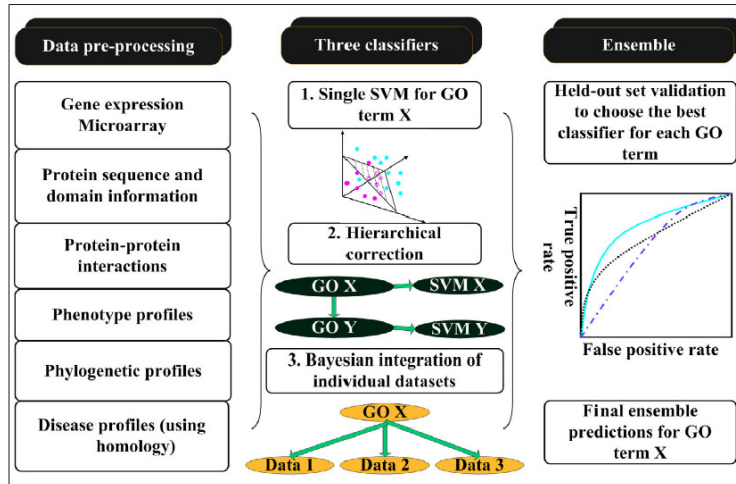


Figure 9: Integration of diverse methods and diverse sources of data in an ensemble framework for AFP prediction. The best classifier for each GO term is selected through held-out set validation (adapted from [76]).

It is worth noting that other approaches adopted Bayesian networks to resolve the hierarchical constraints underlying the GO taxonomy. For instance in the *FALCON* algorithm the GO is modeled as a Bayesian network and for any given input the algorithm returns the most probable GO term assignment in accordance with the GO structure, by using an evolutionary-based optimization algorithm [94].

5.3 HBAYES: an “optimal” Hierarchical Bayesian ensemble approach

The HBAYES ensemble method [38, 40] is a general technique for solving hierarchical classification problems on generic taxonomies G structured according to forest of trees. The method consists in training a calibrated classifier at each node of the taxonomy. In principle any algorithm (e.g. Support Vector Machines or Artificial Neural Networks) whose classifications are obtained by thresholding a real prediction \hat{p} , e.g., $\hat{y} = \text{SGN}(\hat{p})$, can be used as base learner. The real-valued outputs $\hat{p}_i(g)$ of the calibrated classifier for node i on the gene g are viewed as estimates of the probabilities $p_i(g) = \mathbb{P}(y_i = 1 |$

$y_{\text{par}(i)} = 1, g$), The distribution of the random boolean vector Y is assumed to be

$$\mathbb{P}(Y = y) = \prod_{i=1}^m \mathbb{P}(Y_i = y_i \mid Y_{\text{par}(i)} = 1, g) \quad \text{for all } y \in \{0, 1\}^m$$

where, in order to enforce that only multilabels Y that respect the hierarchy have nonzero probability, it is imposed that $\mathbb{P}(Y_i = 1 \mid Y_{\text{par}(i)} = 0, g) = 0$ for all nodes $i = 1, \dots, m$ and all g . This implies that the base learner at node i is only trained on the subset of the training set including all examples (g, y) such that $y_{\text{par}(i)} = 1$.

5.3.1 HBAYES ensembles for protein function prediction

H-loss is a measure of discrepancy between multilabels based on a simple intuition: *if a parent class has been predicted wrongly, then errors in its descendants should not be taken into account*. Given fixed cost coefficients $\theta_1, \dots, \theta_m > 0$, the H-loss $\ell_H(\hat{y}, v)$ between multilabels \hat{y} and v is computed as follows: all paths in the taxonomy T from the root down to each leaf are examined and, whenever a node $i \in \{1, \dots, m\}$ is encountered such that $\hat{y}_i \neq v_i$, then θ_i is added to the loss, while all the other loss contributions from the subtree rooted at i are discarded. This method assumes that, given a gene g , the distribution of the labels $V = (V_1, \dots, V_m)$ is $\mathbb{P}(V = v) = \prod_{i=1}^m p_i(g)$ for all $v \in \{0, 1\}^m$, where $p_i(g) = \mathbb{P}(V_i = v_i \mid V_{\text{par}(i)} = 1, g)$. According to the true path rule, it is imposed that $\mathbb{P}(V_i = 1 \mid V_{\text{par}(i)} = 0, g) = 0$ for all nodes i and all genes g .

In the evaluation phase, HBAYES predicts the Bayes-optimal multilabel $\hat{y} \in \{0, 1\}^m$ for a gene g based on the estimates $\hat{p}_i(g)$ for $i = 1, \dots, m$. By definition of Bayes-optimality, the optimal multilabel for g is the one that minimizes the loss when the true multilabel V is drawn from the joint distribution computed from the estimated conditionals $\hat{p}_i(g)$. That is,

$$\hat{y} = \underset{y \in \{0, 1\}^m}{\operatorname{argmin}} \mathbb{E}[\ell_H(y, V) \mid g] \quad (10)$$

In other words the ensemble method HBAYES provides an approximation of the optimal Bayesian classifier w.r.t. the H-loss [38]. More precisely, as shown in [43] the following theorem holds:

Theorem 1 *For any tree T and gene g the multilabel generated according to the HBAYES prediction rule is the Bayes-optimal classification of g for the H-loss.*

In the evaluation phase the uniform cost coefficients $\theta_i = 1$, for $i = 1, \dots, m$, are used. However, since with uniform coefficients the H-loss can be made small simply by predicting sparse multilabels (i.e., multilabels \hat{y} such that $\sum_i \hat{y}_i$ is small), in the training phase the cost coefficients are set to $\theta_i = 1/|\text{root}(G)|$, if $i \in \text{root}(G)$, and to $\theta_i = \theta_j/|\text{child}(j)|$ with $j = \text{par}(i)$ otherwise. This normalizes the H-loss, in the sense that the maximal H-loss contribution of all nodes in a subtree excluding its root equals that of its root.

Let $\{E\}$ be the indicator function of event E . Given g and the estimates $\hat{p}_i = \hat{p}_i(g)$ for $i = 1, \dots, m$, the HBAYES prediction rule can be formulated as follows:

HBAYES prediction rule

Initially, set the labels of each node i to

$$\hat{y}_i = \operatorname{argmin}_{y \in \{0,1\}} \left(\theta_i \hat{p}_i (1-y) + \theta_i (1-\hat{p}_i) y + \hat{p}_i \{y=1\} \sum_{j \in \text{child}(i)} H_j(\hat{y}) \right) \quad (11)$$

where

$$H_j(\hat{y}) = \theta_j \hat{p}_j (1-\hat{y}_j) + \theta_j (1-\hat{p}_j) \hat{y}_j + \hat{p}_j \{\hat{y}_j=1\} \sum_{k \in \text{child}(j)} H_k(\hat{y})$$

is recursively defined over the nodes j in the subtree rooted at i with each \hat{y}_j set according to (11).

Then, if \hat{y}_i is set to zero, set all nodes in the subtree rooted at i to zero as well.

It is worth noting that \hat{y} can be computed for a given g via a simple bottom-up message-passing procedure. It can be shown that if all child nodes k of i have \hat{p}_k close to a half, then the Bayes-optimal label of i tends to be 0 irrespective of the value of \hat{p}_i . Vice versa, if i 's children all have \hat{p}_k close to either 0 or 1, then the Bayes-optimal label of i is based on \hat{p}_i only, ignoring the children. This behaviour can be intuitively explained in the following way: the estimate \hat{p}_k is built based only on the examples on which the parent i of k is positive; hence, a “neutral” estimate $\hat{p}_k = \frac{1}{2}$ signals that the current instance is a negative example for the parent i . Experimental results show that this approach achieves comparable results with the TPR method (Section 7), an ensemble approach based on the “true path rule” [134].

5.3.2 HBAYES-CS: the cost-sensitive version

The HBAYES-CS is the cost-sensitive version of HBAYES proposed in [43]. By this approach the misclassification cost coefficient θ_i for node i is split in two terms θ_i^+ and θ_i^- for taking into account misclassifications respectively for positive and negative examples. By considering separately these two terms, (11) can be rewritten as:

$$\hat{y}_i = \operatorname{argmin}_{y \in \{0,1\}} \left(\theta_i^- \hat{p}_i (1-y) + \theta_i^+ (1-\hat{p}_i) y + \hat{p}_i \{y=1\} \sum_{j \in \text{child}(i)} H_j(\hat{y}) \right) \quad (12)$$

where the expression for $H_j(\hat{y})$ gets changed correspondingly. By introducing a factor $\alpha \geq 0$ such that $\theta_i^- = \alpha \theta_i^+$ while keeping $\theta_i^+ + \theta_i^- = 2\theta_i$, the relative costs of false positives and false negatives can be parametrized, thus allowing us to further rewrite the hierarchical Bayesian rule (Section 5.3.1):

$$\hat{y}_i = 1 \iff \hat{p}_i \left(2\theta_i - \sum_{j \in \text{child}(i)} H_j \right) \geq \frac{2\theta_i}{1+\alpha}. \quad (13)$$

By setting $\alpha = 1$ we obtain the original version of the hierarchical Bayesian ensemble and by incrementing α we introduce progressively lower costs for positive predictions. In this way we can obtain that the recall of the ensemble tends to increase,

eventually at the expenses of the precision, and by tuning the α parameter we can obtain different combinations of precision/recall values.

In principle, a cost factor α_i can be set for each node i to explicitly take into account the unbalance between the number of positive n_i^+ and negative n_i^- examples, estimated from the training data:

$$\alpha_i = \frac{n_i^-}{n_i^+} \Rightarrow \theta_i^+ = \frac{2}{\frac{n_i^-}{n_i^+} + 1} \theta_i = \frac{2n_i^+}{n_i^- + n_i^+} \theta_i. \quad (14)$$

The decision rule (13) at each node then becomes:

$$\hat{y}_i = 1 \iff p_i \left(2\theta_i - \sum_{j \in \text{child}(i)} H_j \right) \geq \frac{2\theta_i}{1 + \alpha_i} = \frac{2\theta_i n_i^+}{n_i^- + n_i^+}. \quad (15)$$

Results obtained with the yeast model organism showed that HBAYES-CS significantly outperform HTD methods [43, 134].

6 Reconciliation methods

Hierarchical ensemble methods are basically two-steps methods, since at first provide predictions for the single classes and then arrange these predictions to take into account the functional relationships between GO terms. Noble and colleagues name this general approach *reconciliation methods* [118]: they proposed methods for calibrating and combining independent predictions to obtain a set of probabilistic predictions that are consistent with the topology of the ontology. They applied their ensemble methods to the genome-wide and ontology-wide function prediction with *M. musculus*, involving about 3000 GO terms.

Their goal consists in providing consistent predictions, that is predictions whose confidence (for example, posterior probability) increases as we ascend from more specific to more general terms in the GO. Moreover another important issue of these methods is the availability of confidence values associated to the predictions, that can be interpreted as probabilities that a protein has a certain function given the information provided by the data.

The overall reconciliation approach can be summarized in four basic steps (Fig. 10):

1. *Kernel computation.* At first a set of kernels is computed from the available data. We may choose kernel specific for each source of data (e.g. diffusion kernels for protein-protein interaction data [157], linear or gaussian kernel for expression data, string kernel for sequence data [103]). Multiple kernels for the same type of data can also be constructed [118].
2. *SVM learning.* SVMs are used as base learners using the kernels selected at the previous step: the training is performed by internal cross-validation to avoid overfitting, and a local cost-sensitive strategy is applied, by tuning separately the C regularization factor for positive and negative examples. Note that the authors in their experiments used SVMs as base learners but any meaningful classifier could be used at this step.

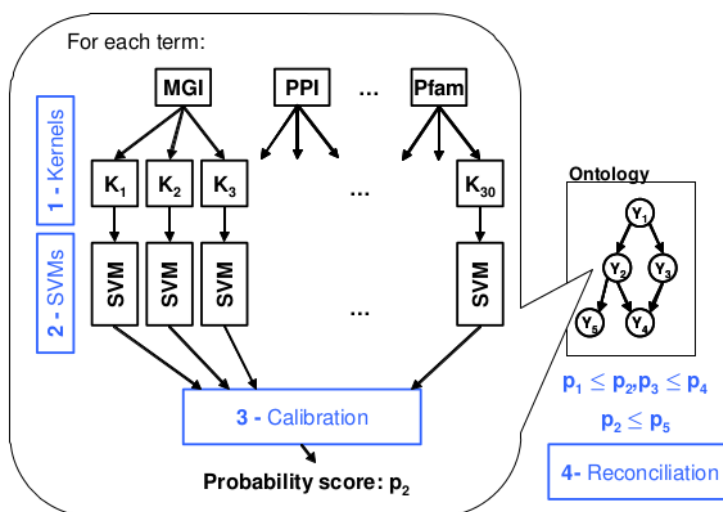


Figure 10: The overall scheme of reconciliation methods (adapted from [118]).

- 3. Calibration.** To produce individual probabilistic outputs from the set of SVM outputs corresponding to one GO term, a logistic regression approach is applied. In this way a calibration of the individual SVM outputs is obtained, resulting in a probabilistic prediction of the random variable Y_i , for each node/term i of the hierarchy, given the outputs \hat{y}_i of the SVM classifiers.
- 4. Reconciliation.** the first three steps generate unreconciled outputs, that is in practice a “flat” ensemble is applied that may generate inconsistent predictions with respect to the given taxonomy. In this step, the outputs of step three are processed by a ‘reconciliation method’. The goal of this stage is to combine predictions for each term to produce predictions that are consistent with the ontology, meaning that all the probabilities assigned to the ancestors of a GO term are larger than the probability assigned to that term.

The first three steps are basically the same (or very similar) for each reconciliation ensemble method. The crucial step is represented by the fourth, i.e. the reconciliation step, and different ensemble algorithms can be designed to implement it. The authors proposed 11 different ensemble methods for the reconciliation of the base classifier outputs. Schematically, they can be subdivided in four main classes of ensembles:

1. Heuristic methods
2. Bayesian network-based methods
3. Cascaded logistic regression
4. Projection-based methods

6.1 Heuristic methods

These approach preserve the "reconciliation property":

$$\forall i, j \in G, (i, j) \in G \Rightarrow \hat{p}_i \geq \hat{p}_j \quad (16)$$

through simple heuristic modifications of the probabilities computed at step 3 of the overall reconciliation scheme.

- The MAX method simply chooses the the largest logistic regression value for the node i and all its descendants desc:

$$p_i = \max_{j \in \text{desc}(i)} \hat{p}_j \quad (17)$$

- The AND method implements the notion that the probability of all ancestral GO terms $\text{anc}(i)$ of a given term/node i are large, assuming that, conditional on the data, all predictions are independent:

$$p_i = \prod_{j \in \text{anc}(i)} \hat{p}_j \quad (18)$$

- OR estimates the probability that the node i is annotated at least for one of the descendant GO terms, assuming again that, conditional on the data, all predictions are independent:

$$1 - p_i = \prod_{j \in \text{desc}(i)} (1 - \hat{p}_j) \quad (19)$$

6.2 Cascaded logistic regression

Instead of modeling class-conditional probabilities, as required by the Bayesian approach, logistic regression can be used instead to directly model posterior probabilities. Considering that modelling conditional densities is in most cases difficult (also using strong independence assumptions as shown in Section 5.1), the choice of logistic regression could be a reasonable one. In [118] the authors embedded in the logistic regression setting the hierarchical dependencies between terms. By assuming that a random variable X whose values represent the features of the gene g of interest is associated to a given gene g , and assuming that $\mathbb{P}(Y = y|X = x)$ factorizes according to the GO graph, then it follows:

$$\mathbb{P}(Y = y|X = x) = \prod_i \mathbb{P}(Y_i = y_i | \forall j \in \text{par}(i) Y_j = y_j, X_i = x_i) \quad (20)$$

with $\mathbb{P}(Y_i = 1 | \forall j \in \text{par}(i) Y_j = 0, X_i = x_i) = 0$. The authors estimated $\mathbb{P}(Y_i = 1 | \forall j \in \text{par}(i) Y_j = 1, X_i = x_i)$ with logistic regression. This approach is quite similar to fitting independent logistic regressions, but note that in this case only examples of proteins having all parents GO terms are used to fit the model, thus implicitly taking into account the hierarchical relationships between GO terms.

6.3 Bayesian network-based methods

These methods are variants of the Bayesian network approach proposed in [12] (Section 5.1): the GO is viewed as a graphical model where a joint Bayesian prior is put on the binary GO term variables y_i . The authors proposed four variants that can be summarized as follows:

- The BPAL is a belief propagation approach with asymmetric Laplace likelihoods. The graphical model has edges directed from more general terms to more specific terms. Differently from [12] the distribution of each SVM output is modeled as an asymmetric Laplace distribution, and a variational inference algorithm, that solves an optimization problem whose minimizer is the set of marginal probabilities of the distribution, is used to estimate the posterior probabilities of the ensemble [189].
- The BPALF approach is similar to BPAL, but with edges inverted and directed from more specific terms to more general terms.
- The BPLR is a heuristic variant of BPAL, where in the inference algorithm, the Bayesian log posterior ratio for Y_i is replaced by the marginal log posterior ratio obtained from the logistic regression (LR).
- The BPLRF is equal to BPLR but with reversed edges.

6.4 Projection-based methods

A different approach is represented by methods that directly use the calibrated values obtained from logistic regression (step 3 of the overall scheme of the reconciliation methods) to find the closest set of values that are consistent with the ontology. This approach leads to a constrained optimization problem. The main contribution of the Obozinski et al. work [118] is represented by the introduction of projection reconciliation techniques based on isotonic regression [10] and the Kullback-Leibler divergence.

The *Isotonic regression* method tries to find a set of marginal probabilities p_i that are close to the set of calibrated values \hat{p}_i obtained from the logistic regression. The euclidean distance is used as a measure of closeness. Hence, considering that the “reconciliation property” requires that $p_i \geq p_j$ when $(i, j) \in E$, this approach yields the following quadratic program:

$$\begin{aligned} \min_{p_i, i \in I} \quad & \sum_{i \in I} (p_i - \hat{p}_i)^2 \\ \text{s.t.} \quad & p_j \leq p_i, \quad (i, j) \in E \end{aligned} \quad (21)$$

This problem is the classical isotonic regression problems that can be solved using an interior point solver or also approximated algorithm when the number of edges of the graph is too large [31].

Considering that we deal with probabilities, a natural measure of distance between probability density functions $f(x)$ and $g(x)$ defined with respect to a random variable x

is represented by the Kullback-Leibler divergence $D_{f_x||g_x}$:

$$D_{f_x||g_x} = \int_{-\infty}^{\infty} f(x) \log \left(\frac{f(x)}{g(x)} \right) dx \quad (22)$$

In the context of reconciliation methods we need to consider a discrete version of the Kullback-Leibler divergence, yielding the following optimization problem:

$$\begin{aligned} \min_p D_{\hat{p}||p} &= \min_{p_i, i \in I} \sum_{i \in I} \hat{p}_i \log \left(\frac{\hat{p}_i}{p_i} \right) \\ \text{s.t.} \quad &p_j \leq p_i, \quad (i, j) \in E \end{aligned} \quad (23)$$

The algorithm finds the probabilities closest to the probabilities \hat{p} obtained from logistic regression according to the Kullback-Leibler divergence and obeying the constraints that probabilities cannot increase descending the hierarchy underlying the ontology.

The extensive experiments exploited in [118] show that, among the reconciliation methods, isotonic regression is the most generally useful. Across a range of evaluation modes, term sizes, ontologies and recall levels, isotonic regression yields consistently high precision. On the other hand, isotonic regression is not always the “best method”, and a biologist with a particular goal in mind may apply other reconciliation methods. For instance, with small terms usually Kullback-Leibler projections achieve the best results, but considering average ”per term” results heuristic methods yields precision at a given recall comparable with projection methods and better than that achieved with Bayes-net methods.

This ensemble approach achieved excellent results in the prediction of protein function in the mouse model organism, demonstrating that hierarchical multilabel methods can play a crucial role for the improvement of protein function prediction performances [118]. Nevertheless the approach suffers from some drawbacks. Indeed the paper focuses on the comparison of hierarchical multilabel methods, but it does not analyze impact of the concurrent use of data integration and hierarchical multilabel methods on the overall classification performances. Moreover, potential improvements could be introduced by applying cost sensitive variants of hierarchical multilabel predictors, able to effectively calibrate the precision/recall trade-off at different levels of the functional ontology.

7 True Path Rule hierarchical ensembles

These ensemble methods exploit at the same time the downward and upward relationships between classes, thus considering both the parent-to-child and child-to-parent functional links (Fig. 2 b).

The True Path Rule (TPR) ensemble method [183, 179] is directly inspired by the *true path rule* that governs both GO and FunCat taxonomies. Citing the curators of the Gene Ontology [73]: “An annotation for a class in the hierarchy is automatically transferred to its ancestors, while genes unannotated for a class cannot be annotated for its descendants”. Considering the parents of a given node i , a classifier that respects

the true path rule needs to obey the following rules:

$$\begin{cases} y_i = 1 & \Rightarrow y_{\text{par}(i)} = 1 \\ y_i = 0 & \not\Rightarrow y_{\text{par}(i)} = 0. \end{cases} \quad (24)$$

On the other hand, considering the children of a given node i , a classifier that respects the true path rule needs to obey the following rules:

$$\begin{cases} y_i = 1 & \not\Rightarrow y_{\text{child}(i)} = 1 \\ y_i = 0 & \Rightarrow y_{\text{child}(i)} = 0. \end{cases} \quad (25)$$

From (24) and (25) we observe an asymmetry in the rules that govern the assignments of positive and negative labels. Indeed, we have a propagation of positive predictions from bottom to top of the hierarchy in (24), and a propagation of negative labels from top to bottom in (25). Conversely, negative labels cannot propagate from bottom to top, and positive predictions cannot propagate from top to bottom.

The “true path rule” suggests algorithms able to propagate “positive” decisions from bottom to top of the hierarchy, and negative decisions from top to bottom (Fig. 11)




Figure 11: The asymmetric flow of information suggested by the true path rule.

7.1 The True Path Rule ensemble algorithm

The TPR algorithm puts together the predictions made at each node by local “base” classifiers to realize an ensemble that obeys the “true path rule”.

The basic ideas behind the method can be summarized as follows:

1. Training of the base learners: for each node of the hierarchy a suitable learning algorithm (e.g. a multi-layer perceptron or a support vector machine) provides a classifier for the associated functional class.

2. In the evaluation phase the trained classifiers associated to each class/node of the graph provide a local decision about the assignment of a given example to a given node.
3. Positive decisions, that is annotations to a specific functional class, may propagate from bottom to top across the graph: they influence the decisions of the parent nodes and of their ancestors in a recursive way, by traversing the graph towards higher level nodes/classes. Conversely, negative decisions do not affect decisions of the parent node (that is they do not propagate from bottom to top, eq. 24).
4. Negative predictions for a given node (taking into account the local decision of its descendants) are propagated to the descendants, to preserve the consistency of the hierarchy according to the true path rule, while positive decisions do not influence decisions of child nodes (eq. 25).

The ensemble combines the local predictions of the base learners associated to each node with the positive decisions that come from the bottom of the hierarchy, and with the negative decisions that spring from the higher level nodes. More precisely, base classifiers estimate local probabilities $\hat{p}_i(g)$ that a given example g belongs to class θ_i , but the core of the algorithm is represented by the evaluation phase, where the ensemble provides an estimate of the “consensus” global probability $\bar{p}_i(g)$.

It is worth noting that instead of a probability, $\hat{p}_i(g)$ may represent a score associated to the likelihood that a given gene/gene product belongs to the functional class i .

Let us consider the set $\phi_i(g)$ of the children of node i for which we have a positive prediction for a given gene g :

$$\phi_i(g) = \{j : j \in \text{child}(i), \hat{y}_j = 1\} \quad (26)$$

The global consensus probability $\bar{p}_i(g)$ of the ensemble depends both on the local prediction $\hat{p}_i(g)$ and on the prediction of the nodes belonging to $\phi_i(g)$:

$$\bar{p}_i(g) = \frac{1}{1 + |\phi_i(g)|} \left(\hat{p}_i(g) + \sum_{j \in \phi_i(g)} \bar{p}_j(g) \right). \quad (27)$$

The decision $\hat{y}_i(g)$ at node/class i is set to 1 if $\bar{p}_i(g) > t$, and to 0 otherwise (a natural choice for t is 0.5), and only children nodes for which we have a positive prediction can influence their parent. In the leaf nodes the sum of eq. 27 disappears and eq. 27 becomes $\bar{p}_i(g) = \hat{p}_i(g)$. In this way positive predictions propagate from bottom to top, and negative decisions are propagated to its descendants when for a given node $\hat{y}_i(g) = 0$.

The bottom-up per level traversal of the tree assures that all the offsprings of a given node i are taken into account for the ensemble prediction. For the same reason we can safely set the classes belonging to the subtree rooted at i to negative, when \hat{y}_i is set to 0. It is worth noting that we have a two-way asymmetric flow of information across the tree: positive predictions for a node influence its ancestors, while negative predictions influence its offsprings.

The algorithm provides both the multilabels \hat{y}_i and an estimate of the probabilities \bar{p}_i that a given example g belongs to the class $i = 1, \dots, m$.

7.2 The cost-sensitive variant

Note that in the *TPR* algorithm there is no way to explicitly balance the local prediction $\hat{p}_i(g)$ at node i with the positive predictions coming from its offsprings (eq. 27). By balancing the local predictions with the positive predictions coming from the ensemble, we can explicitly modulate the interplay between local and descendant predictors. To this end a *weight* w , $0 \leq w \leq 1$ is introduced, such that if $w = 1$ the decision at node i depends only by the local predictor, otherwise the prediction is shared proportionally to w and $1 - w$ between respectively the local parent predictor and the set of its children:

$$\bar{p}_i = w \hat{p}_i + \frac{1 - w}{|\Phi_i|} \sum_{j \in \Phi_i} \bar{p}_j \quad (28)$$

This variant of the *TPR* algorithm is the *weighted True Path Rule (TPR-w)* hierarchical ensemble algorithm: by tuning the w parameter we can modulate the precision/recall characteristics of the resulting ensemble. More precisely, for $w \rightarrow 0$ the weight of the parent local predictor is small, and the ensemble decision mainly depends on the positive predictions of the offsprings nodes (classifiers). Conversely, $w \rightarrow 1$ corresponds to a higher weight of the parent predictor; then less weight is given to possible positive predictions of the children, and the decision depends mainly on the local/parent base classifier. In case of a negative decision all the subtree is set to zero, causing the precision to increase. Note that for $w \rightarrow 1$ the behaviour of *TPR-w* becomes similar to that of *HTD* (Section 4).

A specific advantage of the *TPR-w* ensembles is the capability of tuning precision and recall rates, through the parameter w (28). For small values of w the weight of the decision of the parent local predictor is small, and the ensemble decision depends mainly by the positive predictions of the offsprings nodes (classifiers), and higher values of w correspond to a higher weight of the “parent” local predictor, with a resulting higher precision. In [179] the author shows that the w parameter highly influences the precision/recall characteristics of the ensemble: low w values yield a higher recall, while high values improve the precision of the *TPR-w* ensemble.

Recently Chen and Hu proposed a method that applies the *TPR-w* hierarchical strategy, but using composite kernel SVMs as base classifiers, and a supervised clustering with over-sampling strategy to solve the imbalance data set learning problem, showed that the proper selection of base learners, as well as unbalance-aware learning strategies can further improve the results in terms of hierarchical precision and recall [46].

The same authors proposed also an enhanced version of the *TPR-w* strategy to overcome a limitation of this bottom-up hierarchical method for *AFP*. Indeed for some classes at the lower levels of the hierarchy, the classifier performances are sometimes quite poor, due to both noisy data and the relatively low number of available annotations. More precisely, in the basic *TPR* ensemble, the probabilities \bar{p}_j computed by the children of the node i (28) contributes in equal way to the probability \bar{p}_i computed by the ensemble at node i , independently of the accuracy of the predictions made

by its children classifiers. This “unweighted” mechanism may generate error propagation of the errors across the hierarchy: a poor performance child classifier may for instance with high probability predict as positive a negative example and this error may propagate to its parent node and recursively to its ancestor nodes. To try to alleviate this possible bottom-up error propagation in [47] Chen and Hu proposed an improved *TPR* ensemble (*TPR-w weighted*), based on classifier performance. To this end they weighted the contribution of each child classifier on the basis of their performance evaluated on a validation data set, by adding to (28) another weight v_j

$$\bar{p}_i = w \hat{p}_i + \frac{1-w}{|\phi_i|} \sum_{j \in \phi_i} v_j \cdot \bar{p}_j \quad (29)$$

where v_j is computed on the basis of some accuracy metric A_j (e.g. the F-score) estimated for the child classifiers associated with node j :

$$v_j = \frac{A_j}{\sum_{k \in \phi_i} A_k} \quad (30)$$

In this way the contribution of “poor” classifier is reduced, while “good” classifiers weight more in the final computation of \bar{p}_i (29). Experiments with the “Protein Fate” subtree of the FunCat taxonomy with the yeast model organism show that this approach improves prediction with respect to the “vanilla” *TPR-w* hierarchical strategy [47].

7.3 Advantages and drawbacks of *TPR* methods

While the propagation of negative decisions from top to bottom nodes is quite straightforward and common to the hierarchical *Top-Down* algorithm, the propagation of positive decisions from bottom to top nodes of the hierarchy is specific to the *TPR* algorithm. For a discussion of this item, see Appendix C.

Experimental results show that *TPR-w* achieves equal or better results than the *TPR* and *Top-down* hierarchical strategy, and both hierarchical strategies achieve significantly better results than *Flat* classification methods [178, 42]. The analysis of the per-level classification performances shows that *TPR-w*, by exploiting a global strategy of classification, is able to achieve a good compromise between precision and recall, enhancing the F-measure at each level of the taxonomy [179].

Another advantage of *TPR-w* consists in the possibility of tuning precision and recall by using a global strategy: large values of the w parameter improve the precision, and small values the recall.

Moreover *TPR* and *TPR-w* ensembles provide also a probabilistic estimate of the prediction reliability for each functional class of the overall taxonomy.

The decisions performed at each node of the hierarchical ensemble are influenced by the positive decisions of its descendants. More precisely, the analyses performed in [179] showed that:

- Weights of descendants decrease exponentially w.r.t their depth. As a consequence the influence of descendant nodes decays quickly with their depth.

- The parameter w plays a central role in balancing the weight of the parent classifier associated to a given node with the weights of its positive offsprings: small values of w increase the weight of descendant nodes, large values the weight of the local parent predictor associated to that node.
- The effect on the overall probability predicted by the ensemble is the result of the choice of the w parameter, the strength of the prediction of the local learners and of its descendants.

These characteristics of $TPR-w$ ensembles are well-suited for the hierarchical classification of protein functions, considering that annotations of deeper nodes are likely to have less experimental evidence than higher nodes. Moreover, by enforcing the strength of the descendant nodes through low w values, we can improve the recall characteristics of the overall system (at the expense of a possible reduction in precision).

Unfortunately the method has been conceived and applied only to the FunCat taxonomy, structured according to a tree forest (Section 11), while no applications have been performed using the GO, structured according to a directed acyclic graph (Section 11).

8 Ensembles based on decision trees

Another interesting research line is represented by hierarchical methods based on inductive decision trees [128]. The first attempts to exploit the hierarchical structure of functional ontologies for *AFP* simply used different decision tree models for each level of the hierarchy [86], or investigated a modified decision tree model, in which the assignment to a node is propagated toward the parent nodes [50], by extending the classical C4.5 decision tree algorithm for multiclass classification.

In the context of the predictive clustering tree framework [23], Blockeel et al. proposed an improved version which they applied to the prediction of gene function in the yeast [25].

More recent approaches, always based on modified decision trees, used distance measure derived from the hierarchy, and significantly improved previous methods [185]. The authors showed that separate decision tree models are less accurate than a single decision tree trained to predict all classes at once, even when they are built taking into account the hierarchy.

Nevertheless the previously proposed decision tree-based methods achieve often results not comparable with state-of-the-art hierarchical ensemble methods. To overcome this limitation Schietgat et al. showed that ensembles of hierarchical multilabel decision trees are competitive with state-of-the-art statistical learning methods for DAG-structured prediction of protein function in *S. cerevisiae*, *A. thaliana* and *M. musculus* model organisms [151]. A further work explored the suitability of different ensemble methods based on predictive clustering trees, ranging from global ensembles that learn ensembles of predictive models, each able to predict the entire structure of the hierarchy (i.e. all the GO terms for a given gene) and local ensembles that train an

entire ensemble as a classifier for each branch of the taxonomy. Recently a novel approach used PPI network autocorrelation in hierarchical multi-label classification trees to improve gene function prediction [164].

In [121], methods related to decision trees, in the sense that interpretable classification rules to predict all functions at all levels of the GO hierarchy, have been proposed, but using an Ant Colony optimization classification algorithm to discover classification rules.

Finally, bagging and random forest ensembles [30] have been applied to the *AFP* in yeast, showing that both local and global hierarchical ensemble approaches perform better than the single model counterparts in terms of predictive power [90].

9 The hierarchical classification alone is not enough

Several works showed the in protein function prediction problems we need to consider several learning issues [69, 153, 129]. In particular in [140] the authors showed that even if hierarchical ensemble methods are fundamental to improve the accuracy of the predictions, their mere application is not enough to assure state-of-the-art results if we at the same time do not consider other important learning issues related to *AFP*. Indeed in [42] it has been shown a significant synergy between hierarchical classification, data integration methods and cost-sensitive techniques, highlighting that hierarchical ensemble methods should be designed taking into account different learning issues essential for the *AFP* problem.

9.1 Hierarchical methods and data integration

Several works, and the recently published results of the CAFA 2011 (Critical Assessment of Functional Annotation) challenge showed that data integration plays a central role to improve the predictions of protein functions [117, 129, 52, 99, 161].

Indeed, high-throughput biotechnologies make available increasing quantities of biomolecular data of different types, and several works pointed out that data integration is fundamental to improve the accuracy in *AFP* [69].

According to [117], we may subdivide the main approaches to data integration for *AFP* in four groups:

1. Vector subspace integration
2. Functional association networks integration
3. Kernel fusion
4. Ensemble methods

Vector Space Integration. This approach consists in concatenating vectorial data to combine different sources of biomolecular data [58]. For instance, [124] concatenate different vectors, each one corresponding to a different source of genomic data, in order to obtain a larger vector that is used to train a standard SVM. A similar approach has

been proposed by [76], but each data source is separately normalized in order to take into account the data distribution in each individual vector space.

Functional association networks integration. In functional association networks, different graphs are combined to obtain the composite resulting network [84, 49]. The simplest approaches adopt conjunctive/disjunctive techniques [106], that is respectively adding an edge when in all the networks two genes are linked together or when a link between the two genes is present in at least one functional network, or probabilistic evidence integration schemes [172].

Other methods differentially weight each data source using techniques ranging from Gaussian Random fields [175], to Naive-Bayes integration [115], constrained linear regression [114], or by merging data taking into account the GO hierarchy [113], or by applying XML-based techniques [108].

Kernel fusion. These techniques at first construct a separated Gram matrix for each available data source using appropriate kernels representing similarities between genes/gene products. Then, by exploiting the closure property with respect to the sum and other algebraic operators, the Gram matrices are combined to obtain a “consensus” global integrated matrix.

Besides combining kernels linearly with fixed coefficients [124], one may also use semidefinite programming to learn the coefficients [100]. As methods based on semidefinite programming do not scale well to multiple data sources, more efficient methods for multiple kernel learning have been recently proposed [162, 130]. Kernel fusion methods, both with and without weighting the data sources, have been successfully applied to the classification of protein functions [101, 104, 41, 198]. Recently a novel method proposed an enhanced kernel integration approach by which the weights are iteratively optimized by reducing the empirical loss of a multi-label classifier for each of the labels simultaneously, using a combined objective function [198].

Ensemble methods. Genomic 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.

Some examples of ensemble methods for data combination include “late integration” of kernels trained on different sources [124], Naive Bayes integration [169] of the outputs of SVMs trained with multiple sources [76], and logistic regression for combining the output of several SVMs trained with different biomolecular data and kernels [118].

Recently, in [137] the authors showed that simple ensemble methods, such as weighted voting [53, 88] or Decision Templates [96] give results comparable to state-of-the-art data integration methods, exploiting at the same time the modularity and scalability that characterize most ensemble algorithms. Another work showed that ensemble methods are also resistant to noise [136].

Using an ensemble approach, 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 [132, 133, 135].

As an example of the effectiveness of the integration of hierarchical ensemble methods with data fusion techniques, in [42] six different sources of yeast biomolecular data have been integrated, ranging from protein domain data (PFAM BINARY and PFAM LOGE) [66], gene expression measures (EXPR) [72], predicted and experimentally supported protein-protein interaction data (STRING and BioGRID) [163, 188] to pairwise sequence similarity data (SEQ. SIM.). Kernel fusion integration (sum of Gram matrices) has been applied, and pre-processing has been performed using the the *HCgene* R package [180].

Table 2 summarizes the results of the comparison across about two hundreds of FunCat classes, including single-source and data integration approaches together with both flat and hierarchical ensembles.


Table 2: Comparison of the results (average per-class F-scores) achieved with single sources and multi-source (data fusion) techniques. FLAT, HTD, HTD-CS, HB (HBAYES), HB-CS (HBAYES-CS), TPR and TPR-W ensemble methods are compared with and without data integration. In the last row the number in parentheses refers to the percent relative increment in F-score performance achieved with data fusion techniques with respect to the best single source of evidence (Biogrid).

METHODS	FLAT	HTD	HTD-CS	HB	HB-CS	TPR	TPR-W
SINGLE-SOURCE							
BIOGRID	0.2643	0.3759	0.4160	0.3385	0.4183	0.3902	0.4367
STRING	0.2203	0.2677	0.3135	0.2138	0.3007	0.2801	0.3048
PFAM BINARY	0.1756	0.2003	0.2482	0.1468	0.2407	0.2532	0.2738
PFAM LOGE	0.2044	0.1567	0.2541	0.0997	0.2847	0.3005	0.3160
EXPR.	0.1884	0.2506	0.2889	0.2006	0.2781	0.2723	0.3053
SEQ. SIM.	0.1870	0.2532	0.2899	0.2017	0.2825	0.2742	0.3088
MULTI-SOURCE (DATA FUSION)							
KERNEL FUSION	0.3220(22)	0.5401(44)	0.5492(32)	0.5181(53)	0.5505(32)	0.5034(29)	0.5592(28)


Data fusion techniques improve average per class F-score across classes in FLAT ensembles (first column of Table 2), and significantly boost multilabel hierarchical methods (columns HTD, HTD-CS, HB, HB-CS, TPR, TPR-W of Table 2).

Fig. 12 depicts the classes (black nodes) where Kernel Fusion achieves better results than the best single-source data set (BIOGRID). It is worth noting that the number of black nodes is significantly larger in TPR-W (Fig. 12 b) w.r.t. FLAT methods (Fig. 12 a).

Hierarchical multilabel ensembles largely outperform FLAT approaches [76, 118], but Table 2 and Fig.12 also reveal a synergy between hierarchical ensemble methods and data fusion techniques.



(a)



(b)

Figure 12: FunCat trees to compare F-scores achieved with data integration (KF) to the best single-source classifiers trained on BIOGRID data. Black nodes depict functional classes for which KF achieves better F-scores. (a) FLAT, (b) TPR-W ensembles.

9.2 Hierarchical methods and cost-sensitive techniques

According to [183, 43], cost-sensitive approaches boost predictions of hierarchical methods when single-sources of data are used to train the base learners. These re-

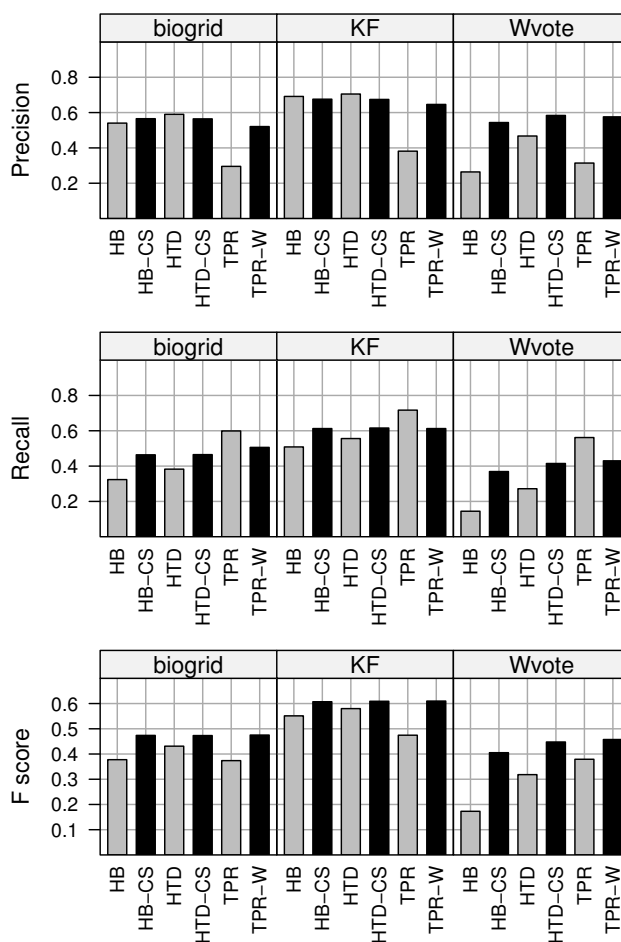


Figure 13: Comparison of hierarchical precision, recall and F-score among different hierarchical ensemble methods using the best source of biomolecular data (BIOGRID), Kernel Fusion (KF), and Weighted Voting (WVOTE) data integration techniques. HB stands for HBAYES.

sults are confirmed when cost-sensitive methods (HBAYES-CS – Section 5.3.2, HTD-CS – Section 4 and TPR-W – Section 7.2) are integrated with data fusion techniques, showing a synergy between multilabel hierarchical, data fusion (in particular kernel fusion), and cost-sensitive approaches (Fig. 13) [42].

Per-level analysis of the F-score in HBAYES-CS, HTD-CS, and TPR-W ensembles shows a certain degradation of performance w.r.t. the depth of nodes, but this degradation is significantly lower when data fusion is applied. Indeed, the per-level F-score achieved by HBAYES-CS and HTD-CS when a single source is used consistently decreases from the top to the bottom level, and it is halved at level 5 w.r.t. to the first

level. On the other hand, in the experiments with Kernel Fusion the average F-score at level 2, 3 and 4 is comparable, and the decrement at level 5 w.r.t. level 1 is only about 15% (Fig. 14). Similar results are reported also with TPR-w ensembles.




Figure 14: Comparison of per level average precision, recall and F-score across the five levels of the FunCat taxonomy in HBAYES-CS using single data sets (single) and kernel fusion techniques (KF). Performance of “single” are computed by averaging across all the single data sources.

In conclusion, the synergic effects of hierarchical multilabel ensembles, cost-sensitive, and data fusion techniques significantly improve the performance of *AFP*. Moreover, these enhancements allow to obtain better and more homogeneous results at each level of the hierarchy. This is of paramount importance, because more specific annotations are more informative, and can get more biological insights about the functions of genes.

9.3 Different strategies to select “negative” genes

In both GO and FunCat only positive annotations are usually available, while negative annotations are very reduced. More precisely, in the GO only about 2500 negative annotations are available, and surely this amount does not allow a sufficient coverage of negative examples

Moreover, some seminal works in functional genomics pointed out that the strategy of choosing negative training examples does affect the classifier performance [15, 104, 156, 195].

In [42] two strategies for choosing negative examples have been compared: the *Basic (B)* and the *Parent Only (PO)* strategy.

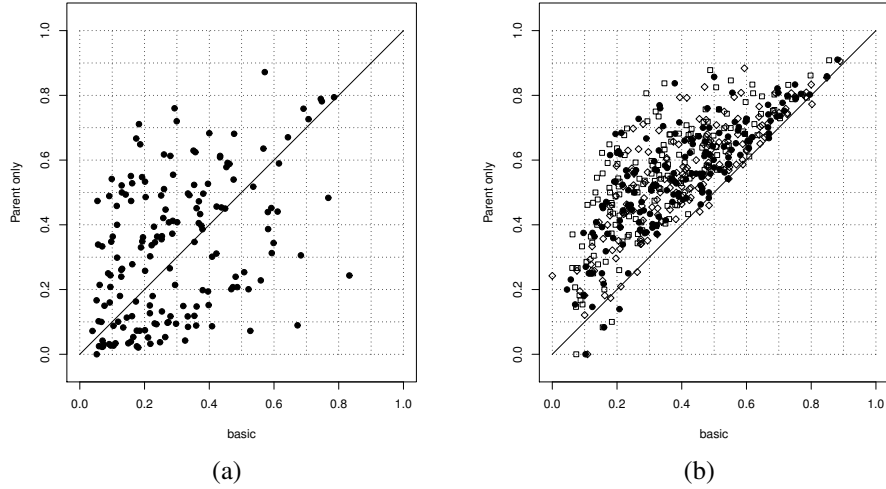


Figure 15: Comparison of average per-class F-score between *Basic* and *PO* strategies. (a) FLAT ensembles; (b) Hierarchical cost-sensitive strategies: HTD-CS (squares), TPR-W (triangles), HBAYES-CS (filled circles). Abscissa: per-class F-score with base learners trained according to the *Basic* strategy; ordinate: per-class F-score with base learners trained according to the *PO* strategy.

According to the *B* strategy, the set of negative examples are simply those genes g that are not annotated for class c_i , that is:

$$N_B = \{g : g \notin c_i\} \quad (31)$$

The *PO* selection strategy chooses as negatives for the class c_i only those examples that are non annotated to c_i but are annotated for a parent class. More precisely, for a given class c_i corresponding to node i in the taxonomy, the set of negative examples is:

$$N_{PO} = \{g : g \notin c_i, g \in \text{par}(i)\} \quad (32)$$

Hence, this strategy selects negative examples for training that are in a certain sense “close” to positives. It is easy to see that $N_{PO} \subseteq N_B$, hence this strategy selects for training a large set of generic negative examples, possibly annotated with classes that are associated with faraway nodes in the taxonomy. Of course, the set of positive examples is the same for both strategies.

The *B* strategy worsens the performance of hierarchical multilabel methods, while for FLAT ensembles there is no clear trend. Indeed in Fig. 15, comparing the F-scores obtained with *B* to those obtained with *PO*, using both hierarchical cost-sensitive (Fig. 15 (a)) and FLAT (Fig. 15 (b)) methods. Each point represents the F-score for a

specific FunCat class achieved by a specific method with B (abscissa) and PO (ordinate) strategy for the selection of negative examples. In Figure 15 (a) most points lie above the bisector independently of the hierarchical cost-sensitive method being used. This shows that hierarchical methods gain in performance when using the PO strategy as opposed to the B strategy (p-value = 2.2×10^{-16} according to the Wilcoxon signed-ranks test). This is not the case for FLAT methods (Fig. 15 (b)).

These results can be explained by considering that the PO strategy takes into account the hierarchy to select negatives, while the B strategy does not. More precisely, FLAT methods having no information about the hierarchical structure of classes, may fail to distinguish negative examples belonging to very distant classes, thus resulting in a high false positive rate, while hierarchical methods, which know the taxonomy, can use the information coming from other base classifiers to prevent a local base learner from incorrectly classifying “distant” negative examples.

In conclusion these seminal works show that the strategy to choose negative examples exerts a significant impact on the accuracy of the predictions of hierarchical ensemble methods, and more research work is needed to explore this topic.

10 Open problems and future trends

In the previous section we showed that different learning issues should be considered to improve the effectiveness and the reliability of hierarchical ensemble methods. Most of these issues and others related to hierarchical ensemble methods and to AFP represent challenging problems that have been only partially considered by previous work. For these reasons we try to delineate some of the open problem and research trends in the context of this research area.

For instance the selection strategies for negative examples have been only partially explored, even if some seminal works show that this item exerts a significant impact on the accuracy of the predictions [15, 104, 156, 42]. Theoretical and experimental comparison of different strategies should be performed in a systematic way, to assess the impact of the different strategies on different hierarchical methods, considering also the characteristics of the learning machines used as base learners.

Some works showed also the cost-sensitive strategies are needed to significantly improve predictions, especially in a hierarchical context [42], but new research could be considered for both applying or designing cost-sensitive base learners and to develop novel hierarchical ensemble unbalance-aware. Cost-sensitive methods have been applied to both the single base learners and also to the overall hierarchical ensemble strategy [179, 42], and recently a hierarchical variant of $SMOTE$ (Synthetic Minority Over-sampling Technique) [45] has been applied to hierarchical protein function prediction, showing very promising results [47]. In principle classical “balancing” strategies should be explored to improve the accuracy and the reliability of the base learners and hence of the overall hierarchical classification process. For instance, random undersampling or oversampling techniques could be applied: the former augments the annotations by exactly duplicating the annotated proteins, whereas the latter randomly takes away some unannotated examples [13]. Other approaches could be considered such as heuristic resampling methods [95], or embedding resampling methods into

data mining algorithms [77], or ensemble methods tailored to imbalanced classification problems [165, 202, 71].

Since functional classes are unbalanced, precision/recall analysis plays a central role in *AFP* problems, and often drives “in vitro” experiments that provide biological insights about specific functional genomics problems [69]. Only a few hierarchical ensemble methods, such as HBAYES-CS [43] and TPR-W [179] can tune their precision/recall characteristics through a single global parameter. In HBAYES-CS, by incrementing the cost factor $\alpha = \theta_i^- / \theta_i^+$ we introduce progressively lower costs for positive predictions, thus resulting in an increment of the recall (at the expenses of a possibly lower precision). In TPR-W, by incrementing w we can reduce the recall and enhance the precision. Parametric versions of other hierarchical ensemble methods could be developed, in order to design ensemble methods with “tunable” precision/recall characteristics.

Another important issue that should be considered in the design of novel hierarchical ensemble methods is the incompleteness of the available annotations and its impact on the performance of computational methods for *AFP*. Indeed the successful application of supervised and semi-supervised machine learning methods to these tasks requires a gold standard for protein function, i.e. a trusted set of correct examples, but unfortunately the annotations is incomplete and undergoes frequent updates, and also the GO is frequently updated. Some seminal works showed that on the one hand current machine learning approaches are able to generalize and predict novel biology from an incomplete gold standard and on the other hand incomplete functional annotations adversely affect the evaluation of machine learning performance [80]. A very recent work addressed these items by proposing methods based on *Weak-label learning* specifically designed to replenish the functions of proteins under the assumption that proteins are partially annotated. More precisely two new algorithms have been proposed: *ProWL*, Protein Function prediction with Weak-label learning, which can recover missing annotations by using the available *relevant* annotations, that is a set of trusted annotations for a given protein, and *ProWL-IF*, Protein Function prediction with Weak-label learning and Knowledge of Irrelevant Function, by which also *irrelevant* functions, i.e. functions that cannot be associated with the protein of interest are exploited to replenish the missing functions [196, 197]. The results show that these items should be considered in future works for hierarchical multi-label predictions of protein functions in model organisms.

Another issue is represented by the reliability of the annotations. Usually only experimental evidence is used to annotate the proteins for training *AFP* methods, but most of the available annotations are computationally predicted annotations without any experimental validation [74]. To at least partially exploit this huge amount of information, computational methods able to take into account the different reliability of the available annotations should be developed and integrated into hierarchical ensemble algorithms.

A quite neglected item is the interpretability of the hierarchical models. Nevertheless the generation of comprehensible classification models is of paramount importance for biologists in order to provide new insights about the correlation of protein features and their functions [68]. A first step in this direction is represented by the work of Cerri et al. that exploits the advantages of grammar-based evolutionary algorithms to

incorporate prior knowledge with the simplicity of genetic algorithms for optimization problems in order to produce interpretable rules for hierarchical multi-label classification [34].

Other issues depend on “strength” or the general rule that relates the predictions made by the base learner at a given term/node of the hierarchy with the predictions made by the other base learners of the hierarchical ensemble. For instance the TPR algorithm (Section 7) weights the positive predictions of deeper nodes with an exponentially decrement with respect to their depth (Sect. 11), but other rules (e.g. linear or polynomial) could be considered as the basis for the development of new algorithms that put more weight on the decisions of deep nodes of the hierarchy. Other enhancements could be introduced with the TPR-w algorithm (Section 7.2): indeed we can note that positive children of a node at level i of the hierarchy have the same weight, independently of the size of their hanging subtree. In some cases this could be useful, but in other cases it could be desirable to directly take into account the fact that a positive prediction is maintained along a path of the tree: indeed this witnesses for a positive annotation of the node at level i .

More in general, in the spirit of a work recently proposed [42], the analysis of the synergy between the issues introduced above could be of great interest to better understand the behaviour of hierarchical ensemble methods.

Finally we introduce some problems that could open new and interesting research lines in the context of hierarchical ensemble methods.

At first an important issue could be represented by the design and development of multi-task learning strategies [192] able to exploit the relationships between functional classes already during the learning phase, in order to establish a functional connection between learning processes associated to hierarchically related classes of the functional taxonomy. In this way just during the training of the base learners, the learning processes will be dependent on each other (at least for nearby nodes/classes), enabling “mutual learning” of related classes in the taxonomy.

A second, to my knowledge not explored learning issue is represented by the meta-integration of hierarchical predictions. Considering that there is no “killer” hierarchical ensemble method, a meta-combination of the hierarchical predictions could be explored to enhance the overall performances.

A last issue is represented by multi-species predictions in a hierarchical context. By exploiting homology relationships between proteins of different species we could enhance the prediction for a particular species by using predictions or data available for other species. This is a common practice with e.g. sequence-based methods, but novel research is needed to extend this homology-based approach in the context of hierarchical ensemble methods for multi-species prediction. It is worth noting that this multi-species approach yields to big-data analysis with the associated problems of scalability of existing algorithms. A possible solution to this last problem could be represented by distributed parallel computation [75], or by the adoption of secondary memory-based computational techniques [109]

11 Conclusions

Hierarchical ensemble methods represent one of the main research lines for *AFP*. Their two-steps learning strategy introduces a high modularity in the prediction system: in the first step different base learners can be trained to individually learn the functional classes, and in the second step different algorithms can be chosen to hierarchically combine the predictions provided by the base classifiers. The best results can be obtained when the global topology of the ontology is exploited, and when both top-down and bottom-up learning strategies are applied [76, 118, 179, 43].

Nevertheless a hierarchical learning strategy alone is not enough to achieve state-of-the-art results for *AFP*. Indeed we need to design hierarchical ensemble methods in the context of the learning issues strictly related to the *AFP* problem.

The first one is represented by data fusion, since each source of biomolecular data may provide different and often complementary information about a protein and an integration of data fusion methods with hierarchical ensembles is mandatory to improve *AFP* results.

The second one is represented by the cost-sensitive techniques needed to take into account the usually small number of positive annotations: data unbalance-aware methods should be embedded in hierarchical methods to avoid solutions biased toward low sensitivity predictions.

Other issues, ranging from the proper choice of negative examples to the reliability and the incompleteness of the available annotation, the balance between local and global learning strategies and the meta-integration of hierarchical predictions have been only partially addressed in previous work. More in general, the synergy between hierarchical ensemble methods, data integration algorithms, cost-sensitive techniques and other related issues is the key to improve *AFP* methods and to drive experiments aimed at discovering previously unannotated or partially annotated protein functions [42].

Indeed, despite their successful application to protein function prediction in different model organisms, as outlined in Section 9, there is large room for future research in this challenging area of computational biology.

In particular the development of multi-task learning methods to jointly learn related GO terms in a hierarchical context, and the design of multi-species hierarchical algorithms, able to scale with millions of proteins, represent a compelling challenge for the computational biology and bioinformatics community.

Appendix

A. Gene Ontology and FunCat

The two main taxonomies of gene functional classes are represented by the Gene Ontology (GO) [167], and The Functional Catalogue (FunCat) [146]. In the former the functional classes are structured according to a directed acyclic graph, i.e. a DAG (Fig. 17), while the latter through a forest of trees (Fig. 18). The GO is composed by thousands of functional classes, and it is set out in three separated ontologies: "Biological Processes", "Molecular Function" and "Cellular Component" (Fig. 16). Indeed a




Figure 16: The Gene Ontology

gene can participate to specific biological processes (e.g. cell cycle, metabolism, nucleotide biosynthesis) and at the same time can perform specific molecular functions (e.g. catalytic or binding activities, that occur at the molecular level) in specific cellular components (e.g. mitochondrion or rough endoplasmic reticulum).

GO: the Gene Ontology

The Gene Ontology (GO) project began as a collaboration between three model organism databases, FlyBase (*Drosophila*), the Saccharomyces Genome Database (SGD) and the Mouse Genome Database (MGD), in 1998. Now it includes several of the world's major repositories for plant, animal and microbial genomes. The GO project has developed three structured controlled vocabularies (ontologies) that describe gene products in terms of their associated biological processes, cellular components and molecular functions in a species-independent manner (Fig. 16). Biological Process (BP) represents series of events accomplished by one or more ordered assemblies of molecular functions that exploit a specific biological function. For instance lipid metabolic process or tricarboxylic acid cycle. Molecular Function (MF) describes activities that occur at the molecular level, such as catalytic or binding activities. An example of MF is glycine dehydrogenase activity or glucose transporter activity. Cellular Component (CC) represents just parts or components of a cell, such as organelles or physical places or compartments in which a specific gene product is located. An example is the endoplasmic reticulum or the ribosome.

The ontologies of the GO are structured as a directed acyclic graph (DAG) $G = \langle V, E \rangle$, where $V = \{t | \text{terms of the } GO\}$ and $E = \{(t, u) | t, u \in V\}$. Relations between GO terms are also categorized in three main groups:

- *is-a* (subtype relations): if we say the term/node A is a B , we mean that node A is a subtype of node B . For example, mitotic cell cycle is a cell cycle, or lyase

activity is a catalytic activity.

- *part-of* (part-whole relations): *A* is part of *B* means that whatever *A* exists, it is part of *B*, and the presence of *A* implies the presence of *B*. For instance, mitochondrion is part of cytoplasm.
- *regulates* (control relations): if we say that *A* regulates *B* we mean that *A* directly affects the manifestation of *B*, i.e. the former regulates the latter.

While *is-a* and *part-of* are transitive, *regulates* is not. Moreover in some cases regulatory proteins are not expected to have the same properties as the proteins they regulate and hence predicting regulation may require other data and assumptions than predicting function similarity. For these reasons usually in *AFP* *regulates* relations (that however are a minority of the existing relations) are usually not used.

Each annotation is labeled with an *evidence code*, that indicates how the annotation to a particular term is supported. They are subdivided in several categories ranging from experimental evidence codes, used when experimental assays have been applied for the annotation, e.g. Inferred from Physical Interaction (IPI) or Inferred from Mutant Phenotype (IMP) or Inferred from Genetic Interaction (IGI), to author statement codes, such as Traceable Author Statement (TAS), that indicate that the annotation was made on the basis of a statement made by the author(s) in the cited reference, to computational analysis evidence codes, based on an in silico analyses manually reviewed (e.g. Inferred from Sequence or Structural Similarity (ISS)). For the full set of available evidence codes, please see the GO web site (<http://www.geneontology.org>).

A GO graph for the yeast model organism is represented in Fig. 17. It is worth noting that, despite the complexity of the represented graph, Fig. 17 does not show all the available terms and the relationships involved in the GO BP ontology with the yeast model organism.

FunCat: the Functional Catalogue

The FunCat taxonomy started with *Saccharomyces cerevisiae* genome project at MIPS (<http://mips.gsf.de/>): at the beginning, FunCat contained only those categories required to describe yeast biology [110, 111], but successively its content has been extended to plants to annotate genes from the *Arabidopsis thaliana* genome project and furthermore to cover prokaryotic organisms and finally animals too [146].

The *FunCat* represents a relatively simple and concise set of gene functional classes: it consists of 28 main functional categories (or branches) that cover general fields like cellular transport, metabolism and cellular communication/signal transduction. These main functional classes are divided into a set of subclasses with up to six levels of increasing specificity, according to a tree-like structure that accounts for different functional characteristics of genes and gene products. Genes may belong at the same time to multiple functional classes, since several classes are subclasses of more general ones, and because a gene may participate in different biological processes and may perform different biological functions.

Taking into account the broad and highly diverse spectrum of known protein functions, the FunCat annotation scheme covers general features, like cellular transport,




Figure 17: GO BP DAG for the yeast model organism (realized through the *HCGene* software [180], involving more than 1000 terms and more than 2000 edges

metabolism and protein activity regulation. Each of its main 28 functional branches is organized as a hierarchical, tree like structure, thus leading to a tree forest with hundreds of functional categories.

Differently from the GO, the FunCat is more compact and does not intend to classify protein functions down to the most specific level. From a general standpoint it can be compared to parts of the Molecular Function and Biological Process terms of the GO system.

One of the main advantages of FunCat is its intuitive category structure. For instance, the annotation of yeast uses only 18 of the main categories and less than 300 distinct categories (Fig. 18), while the *Saccharomyces Genome Database* (SGD) [48] uses more than 1500 GO terms in its yeast annotation. Indeed FunCat focuses on the functional process and in part to the molecular function, while GO aims at representing a fine granular description of proteins that provides annotations with a wealth of detailed information. However, to achieve this goal, the detailed description offered by GO leads to a large number of terms (e.g. the ontology for biological processes alone contains more than 10000 terms), and such a huge amount of terms is very difficult to be handled for annotators. Moreover, we may have a very large number of possible assignments, that may lead to erroneous or inconsistent annotations. FunCat is simpler: its tree structure, compared with the DAG structure of the GO, leads to both simple procedures for annotation, and less difficult computational-based classification tasks.




Figure 18: FunCat tree for the yeast model organism (realized through the *HCGene* software [180]). A “dummy” root node has been added to obtain a single tree from the tree forest.

In other words, it represents a well-balanced compromise between extensive depth, breadth and resolution but without being too granular and specific.

It is worth noting that both FunCat and GO ontologies undergo modifications between different releases, and at the same time the annotations are also subjected to changes, since they represent the results of the knowledge of the scientific community at a given time. As a consequence, predictions resulting e.g. in false positives for a given release of the GO may become true positive in future releases, and more in general we should keep in mind that the available annotations are always partial and incomplete and depend on the knowledge available for the species under study. Nevertheless, even if some works pointed out the inconsistency of current GO taxonomies through the analysis of violations of terms univocality [187], GO and FunCat are considered the ground truth to evaluate *AFP* methods, since they represent the main effort of the scientific community to organize a commonly accepted taxonomy of protein functions [74]

B. *AFP* performance assessment in a hierarchical context

In the context of ontology-wide protein function prediction problems, where negative examples are usually a lot more than positives, accuracy is not a reliable measure to as-

sess the classification performance. For this reason the classical F-score is used instead, to take into account the unbalance of functional classes. If TP represents the positive examples correctly predicted as positive, FN , the positive examples incorrectly predicted as negative and FP the negatives incorrectly predicted as positives, then the precision P , and the recall R are:

$$P = \frac{TP}{TP+FP} \quad R = \frac{TP}{TP+FN} \quad (33)$$

The F-score F is the harmonic mean between precision and recall:

$$F = \frac{2 \cdot P \cdot R}{P+R} \quad (34)$$

If we need to evaluate the correct ranking of annotated proteins with respect to a specific functional class, a valuable measure is represented by the area under the receiving operating characteristic curve (AUC). A random ranking corresponds to $AUC \simeq 0.5$, while values close to 1 correspond to near optimal ranking, i.e. $AUC = 1$ if all the annotated genes are ranked before the unannotated ones.

In order to better capture the hierarchical and sparse nature of the protein function prediction problem, we also need specific measures that estimate how far a predicted structured annotation is from the correct one. Indeed, functional classes are structured according to a direct acyclic graph (Gene Ontology) or to a tree (FunCat), and we need measures to accommodate not just “exact matches” but also “near misses” of different sorts.

For instance, correctly predicting a parent or ancestor annotation, while failing to predict the most specific available annotation should be “partially correct”, in the sense that we can gain information about the more general functional characteristics of a gene, missing only its most specific functions.

More precisely, given a general taxonomy G representing the graph of the functional classes, for a given gene/gene product x consider the graph $P(x) \subset G$ of the predicted classes and the graph $C(x)$ of the correct classes associated to x , and let be $l(P)$ the set of the leaves (nodes without children) of the graph P . Given a leaf $p \in P(x)$, let be $\uparrow p$ the set of ancestors of the node p that belong to $P(x)$, and given a leaf $c \in C(x)$, let be $\uparrow c$ the set of ancestors of the node c that belong to $C(x)$, the Hierarchical Precision (HP), Hierarchical Recall (HR) and Hierarchical F-score (HF) are defined as follows [186]:

$$\begin{aligned} HP &= \frac{1}{|l(P(x))|} \sum_{p \in l(P(x))} \max_{c \in l(C(x))} \frac{|\uparrow c \cap \uparrow p|}{|\uparrow p|} \\ HR &= \frac{1}{|l(C(x))|} \sum_{c \in l(C(x))} \max_{p \in l(P(x))} \frac{|\uparrow c \cap \uparrow p|}{|\uparrow c|} \\ HF &= \frac{2 \cdot HP \cdot HR}{HP + HR} \end{aligned} \quad (35)$$

In the case of the FunCat taxonomy, since it is structured as a tree, we can simplify

HP , HR and HF as follows:

$$\begin{aligned} HP &= \frac{1}{|I(P(x))|} \sum_{p \in I(P(x))} \frac{|C(x) \cap \uparrow p|}{|\uparrow p|} \\ HR &= \frac{1}{|I(C(x))|} \sum_{c \in I(C(x))} \frac{|\uparrow c \cap P(x)|}{|\uparrow c|} \\ HF &= \frac{2 \cdot HP \cdot HR}{HP + HR} \end{aligned} \quad (36)$$

An overall high hierarchical precision is indicative that the predictor is able to detect the most general functions of genes/gene products. On the other hand a high average hierarchical recall indicates that the predictors are able to detect the most specific functions of the genes. The hierarchical F-measure expresses the correctness of the structured prediction of the functional classes, taking into account also partially correct paths in the overall hierarchical taxonomy, thus providing in a synthetic way the effectiveness of the structured hierarchical prediction.

Another variant of hierarchical classification measure is represented by the Hierarchical F-measure proposed by Kiritchenko et al. [87]. Let $P(x)$ the set of classes predicted in the overall hierarchy for a given gene/gene product x , and $C(x)$ the corresponding set of “true” classes. Then the Hierarchical Precision HP_K and the Hierarchical recall HR_K according to Kiritchenko are defined as:

$$HP_K = \sum_x \frac{|P(x) \cap C(x)|}{|P(x)|} \quad HR_K = \sum_x \frac{|P(x) \cap C(x)|}{|C(x)|} \quad (37)$$

Note the these definitions do not explicitly consider the paths included in the predicted subgraphs, but simply the ratio between the number of common classes and respectively the predicted (HP_K) and the true classes (HR_K). The Hierarchical F-measure HF_K is the harmonic mean between HP_K and HR_K :

$$HF_K = \frac{2 \cdot HP_K \cdot HR_K}{HP_K + HR_K} \quad (38)$$

C. Effect of the propagation of the positive decisions in TPR ensembles

In TPR ensembles a *generic* node at level k is any node whose distance from the root is equal to k . The posterior probability computed by the ensemble for a generic node at level k is denoted by q_k . More precisely q_k denotes the probability computed by the ensemble and \hat{q}_k the probability computed by the base learner local to a node at level k . Moreover we define q_{k+1}^j as the probability of a child j of a node at level k , where the index $j \geq 1$ refers to different children of a node at level k . From (28) we can derive the following expression for the probability q_k computed for a generic node at level k of the hierarchy [179]:

$$q_k(g) = w \cdot \hat{q}_k(g) + \frac{1-w}{|\Phi_k(g)|} \sum_{j \in \Phi_k(g)} q_{k+1}^j(g) \quad (39)$$

To simplify the notation we can introduce the following expression to indicate the average of the probabilities computed by the positive children nodes of a generic node at level k :

$$a_{k+1} = \frac{1}{|\phi_k|} \sum_{j \in \phi_k} \hat{q}_{k+1}^j \quad (40)$$

and we can introduce similar notations for a_{k+2} (average of the probabilities of the grandchildren), and more in general for a_{k+j} (descendants at level j of a generic node). By extending these definition across levels, we can obtain the following theorem:

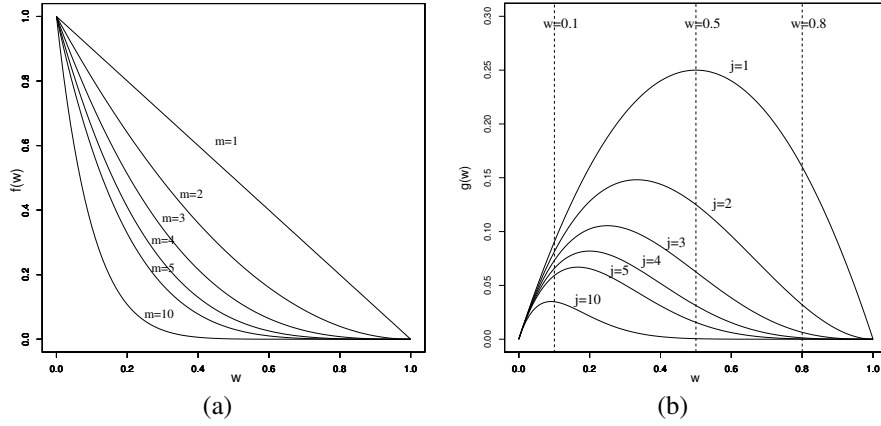


Figure 19: (a) Plot of $f(w) = (1-w)^m$, while varying m from 1 to 10. (b) Plot of $g(w) = w(1-w)^j$, while varying j from 1 to 10. The integers j refer to internal nodes at distance j from the reference node at level k

Theorem 2 (Influence of positive descendant nodes) *In a TPR- w ensemble, for a generic node at level k , with a given parameter $w, 0 \leq w \leq 1$, balancing the weight between parent and children predictors, and having a variable number larger or equal than 1 of positive descendants for each of the m lower levels below, the following equality holds for each $m \geq 1$:*

$$q_k = w\hat{q}_k + \sum_{j=1}^{m-1} w(1-w)^j a_{k+j} + (1-w)^m a_{k+m} \quad (41)$$

For the full proof see [179].

Theorem 2 shows that the contribution of the descendant nodes decays exponentially with their depth and depends critically on the choice of the w parameter. To get more insights into the relationships between w and its effect on the influence of positive decisions on a generic node at level k (Theorem 1), Fig. 19 (a) shows the function that governs the decay of the influence of leaf nodes at different depths m , and Fig. 19 (b) the function responsible of the influence of nodes above the leaves.

Conflict of interests

The author has no direct financial relations that might lead to a conflict of interest associated with this paper.

Acknowledgements

The author thanks the reviewers for their comments and suggestions and acknowledges partial support from the *PRIN* project “Automi e linguaggi formali: aspetti matematici e applicativi”, funded by the Italian Ministry of University.

References

- [1] T Abeel, T Helleputte, Y Van de Peer, P Dupont, and Y Saeys. Robust biomarker identification for cancer diagnosis with ensemble feature selection methods. *Bioinformatics*, 26(3):392–398, 2010.
- [2] N. Alaydie, C.K. Reddy, and F. Fotouhi. Hierarchical multi-label boosting for gene function prediction. In *Proceedings of the International Conference on Computational Systems Bioinformatics (CSB)*, Stanford, CA, 2010.
- [3] N. Alaydie, C.K. Reddy, and F. Fotouhi. Exploiting Label Dependency for Hierarchical Multi-label Classification. In *Advances in Knowledge Discovery and Data Mining*, number 7301 in Lecture Notes in Computer Science, pages 294–305. Springer-Verlag, 2012.
- [4] E.L. Allwein, R.E. Schapire, and Y. Singer. Reducing multiclass to binary: a unifying approach for margin classifiers. *Journal of Machine Learning Research*, 1:113–141, 2000.
- [5] S.F. Altschul, W. Gish, W. Miller, E.W. Myers, and D.J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215, 1990.
- [6] SF Altschul, TL Madden, AA Schaffer, J Zhang, Z Zhang, W Miller, and DJ. Lipman. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic Acids Research*, 25(17):3389–3402, 1997.
- [7] K. Astikainen, L. Holm, E. Pitkanen, S. Szedmak, and J. Rousu. Towards structured output prediction of enzyme function. *BMC Proceedings*, 2(Suppl 4:S2), 2008.
- [8] G. Bakir, T. Hoffman, B. Scholkopf, B. Smola, A.J. Taskar, and S.V.N. Vishwanathan. *Predicting structured data*. MIT Press, Cambridge, MA, 2007.
- [9] R.E. Banfield, L.O. Hall, O. Lawrence, K.W. Bowyer, W. Kevin, and P. Kegelmeyer. A comparison of decision tree ensemble creation techniques. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 29(1):173–180, 2007.

- [10] R.E. Barlow and H.D. Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337):140–147, 1972.
- [11] Z. Barutcuoglu and C. DeCoro. Hierarchical shape classification using bayesian aggregation. In *Proc. of the IEEE Conf. on Shape Modeling and Applications*, 2006.
- [12] Z. Barutcuoglu, R.E. Schapire, and O.G. Troyanskaya. Hierarchical multi-label prediction of gene function. *Bioinformatics*, 22(7):830–836, 2006.
- [13] G. Batista, R. Prati, and M. Monard. A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1):20–29, 2004.
- [14] E. Bauer and R. Kohavi. An empirical comparison of voting classification algorithms: Bagging, Boosting and variants. *Machine Learning*, 36(1/2):105–139, 1999.
- [15] A. Ben-Hur and W.S. Noble. Choosing negative examples for the prediction of protein-protein interactions. *BMC Bioinformatics*, 7(Suppl 1/S2), 2006.
- [16] Y. Bengio, O. Delalleau, and N. Le Roux. Label Propagation and Quadratic Criterion. In O. Chapelle, B. Scholkopf, and A. Zien, editors, *Semi-Supervised Learning*, pages 193–216. MIT Press, 2006.
- [17] A. Bertoni, M. Frasca, and G. Valentini. COSNet: a cost sensitive neural network for semi-supervised learning in graphs. In *European Conference on Machine Learning, ECML PKDD 2011*, volume 6911 of *Lecture Notes on Artificial Intelligence*, pages 219–234. Springer, 2011.
- [18] A. Bertoni and G. Valentini. Ensembles based on random projections to improve the accuracy of clustering algorithms. In *Neural Nets, WIRN 2005*, volume 3931 of *Lecture Notes in Computer Science*, pages 31–37. Springer, 2006.
- [19] A. Bertoni and G. Valentini. Discovering multi-level structures in bio-molecular data through the Bernstein inequality. *BMC Bioinformatics*, 9(S2), 2008.
- [20] R. Bi, Y. Zhou, F. Lu, and W. Wang. Predicting Gene Ontology functions based on Support Vector Machines and statistical significance estimation. *Neurocomputing*, 70:718–725, 2007.
- [21] Y. Bilu and M. Linial. The advantage of functional prediction based on clustering of yeast genes and its correlation with non-sequence based classification. *Journal of Computational Biology*, 9:193–210, 2002.
- [22] A. Binder, M. Kawanabe, and U. Brefeld. Bayesian aggregation for hierarchical genre classification. In *Proc. of the 9th Asian Conf. on Computer Vision*, 2009.
- [23] H. Blockeel, M. Bruynooghe, S. Dzeroski, J. Ramon, and J. Struyf. Top-down induction of clustering trees. In *Proc. of the 15th Int. Conf. on Machine Learning*, pages 55–63, 1998.

- [24] H. Blockeel, L. Schietgat, and A. Clare. Hierarchical multilabel classification trees for gene function prediction. In J. Rousu, S. Kaski, and E. Ukkonen, editors, *Probabilistic Modeling and Machine Learning in Structural and Systems Biology*, Tuusula, Finland, 2006. Helsinki University Printing House.
- [25] H. Blockeel, L. Schietgat, S. Dzeroski, and A. Clare. Decision trees for hierarchical multilabel classification: a case study in functional genomics. In *Proc. of the 10th European Conference on Principles and Practice of Knowledge Discovery in Databases*, volume 4213 of *Lecture Notes in Artificial Intelligence*, pages 18–29. Springer-Verlag, 2006.
- [26] P. Bogdanov and A.K. Singh. Molecular function prediction using neighborhood features. *IEEE ACM Transactions on Computational Biology and Bioinformatics*, 7(2):208–217, 2011.
- [27] H. Borges and J. Nievola. Hierarchical classification using a competitive neural network. In *International Conference on Computing, Networking and Communications - ICNC 2012*, pages 172–177, 2012.
- [28] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [29] L. Breiman. Bias, variance and arcing classifiers. Technical Report TR 460, Statistics Department, University of California, Berkeley, CA, 1996.
- [30] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [31] O. Burdakov, O. Sysoev, A. Grimvall, and M. Hussian. An $O(n^2)$ algorithm for isotonic regression. In *Large-Scale Nonlinear Optimization*, number 83 in *Nonconvex Optimization and Its Applications*, pages 25–33. Springer-Verlag, 2006.
- [32] A. Burred, J. and Lerch. A hierarchical approach to automatic musical genre classification. In *Proc. of the 6th Int. Conf. on Digital Audio Effects*, pages 8–11, 2003.
- [33] M. Ceci and D. Malerba. Classifying web documents in a hierarchy of categories: A comprehensive study. *Journal of Intelligent Information Systems*, 28(1):1–41, 2007.
- [34] R. Cerri, R. Barros, A. de Carvalho, and A. Freitas. A grammatical evolution algorithm for generation of hierarchical multi-label classification rules. In *2013 IEEE Congress on Evolutionary Computation*, 2013.
- [35] R. Cerri and A. de Carvalho. Hierarchical multilabel classification using top-down label combination and artificial neural networks. In *SBRN '10 Proceedings of the 2010 Eleventh Brazilian Symposium on Neural Networks*, pages 253–258. IEEE Computer Society, 2010.
- [36] R. Cerri and A. de Carvalho. New top-down methods using SVMs for hierarchical multilabel classification problems. In *IJCNN 2010*, pages 1–8. IEEE Computer Society, 2010.

- [37] R. Cerri and A. de Carvalho. Hierarchical multilabel protein function prediction using local neural networks. In *Advances in Bioinformatics and Computational Biology*, number 6832 in Lecture Notes in Computer Science, pages 10–17. Springer-Verlag, 2011.
- [38] N. Cesa-Bianchi, C. Gentile, A. Tironi, and L. Zaniboni. Incremental algorithms for hierarchical classification. In *Advances in Neural Information Processing Systems*, volume 17, pages 233–240. MIT Press, 2005.
- [39] N. Cesa-Bianchi, C. Gentile, F. Vitale, and G. Zappella. Random spanning trees and the prediction of weighted graphs. In *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel, 2010.
- [40] N. Cesa-Bianchi, C. Gentile, and L. Zaniboni. Hierarchical classification: Combining Bayes with SVM. In *Proc. of the 23rd Int. Conf. on Machine Learning*, pages 177–184. ACM Press, 2006.
- [41] N. Cesa-Bianchi, M. Re, and G. Valentini. Functional inference in FunCat through the combination of hierarchical ensembles with data fusion methods. In *ICML-MLD 2nd International Workshop on learning from Multi-Label Data*, pages 13–20, Haifa, Israel, 2010.
- [42] N. Cesa-Bianchi, M. Re, and G. Valentini. Synergy of multi-label hierarchical ensembles, data fusion, and cost-sensitive methods for gene functional inference. *Machine Learning*, 88(1):209–241, 2012. 10.1007/s10994-011-5271-6.
- [43] N. Cesa-Bianchi and G. Valentini. Hierarchical cost-sensitive algorithms for genome-wide gene function prediction. *Journal of Machine Learning Research, W&C Proceedings, Machine Learning in Systems Biology*, 8:14–29, 2010.
- [44] S. Chakrabarti, B. Dom, R. Agrawal, and P. Raghavan. Scalable feature selection, classification and signature generation for organizing large text databases into hierarchical topic taxonomies. *The VLDB Journal*, 7:163–178, 1998.
- [45] NV Chawla, KW Bowyer, LO Hall, and Kegelmeyer WP. SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
- [46] B. Chen, L. Duan, and J. Hu. Composite kernel based svm for hierarchical multi-label gene function classification. In *Proc. of the Twenty-Sixth International Florida Artificial Intelligence Research Society Conference*, pages 1380–1385, 2012.
- [47] B. Chen and J. Hu. Hierarchical multi-label classification based on over-sampling and hierarchy constraint for gene function prediction. *IEEE Transactions on Electrical and Electronic Engineering*, 7:183–189, 2012.
- [48] K.R. Christie et al. Saccharomyces genome database (SGD) provides tools to identify and analyze sequences from Saccharomyces cerevisiae and related sequences from other organisms. *Nucleic Acids Res.*, 32:D311–D314, 2004.

- [49] 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.
- [50] A. Clare and R.D. King. Predicting gene function in *saccharomices cerevisiae*. *Bioinformatics*, 19(Supp.2):II42–II49, 2003.
- [51] A. Conesa, S. Gotz, JM Garcia-Gomez, J. Terol, Talon M., and M. Robles. Blast2GO: a universal tool for annotation, visualization and analysis in functional genomics research. *Bioinformatics*, 21(18):3674–3676, 2005.
- [52] D. Cozzetto, D. Buchan, K. Bryson, and D. Jones. Protein function prediction by massive integration of evolutionary analyses and multiple data sources. *BMC Bioinformatics*, 14(Suppl 3:S1), 2013.
- [53] 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.
- [54] C. DeCoro, Z. Barutcuoglu, and R. Fiebrink. Bayesian aggregation for hierarchical genre classification. In *Proc. of the 8th Int. Conf. on Music Information Retrieval*, pages 77–80, 2007.
- [55] O. Dekel, J. Keshet, and Y. Singer. Large margin hierarchical classification. In *Proc. of the 21st Int. Conf. on Machine Learning*, pages 209–216. Omnipress, 2004.
- [56] K. Dembczynski, W. Waegeman, W. Cheng, and E. Hullermeier. On label dependence in multi-label classification. In *ICML-MLD: 2nd International Workshop on learning from Multi-Label Data*, pages 5–12, Haifa, Israel, 2010.
- [57] M. Deng, T. Chen, and F. Sun. An integrated probabilistic model for functional prediction of proteins. *J. Comput. Biol.*, 11:463–475, 2004.
- [58] 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.
- [59] M. Dettling and P. Buhlmann. Boosting for tumor classification with gene expression data. *Bioinformatics*, 19(9):1061–1069, 2003.
- [60] T.G. Dietterich. Ensemble methods in machine learning. In J. Kittler and F. Roli, editors, *Multiple Classifier Systems. First International Workshop, MCS 2000, Cagliari, Italy*, volume 1857 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag, 2000.
- [61] T.G. Dietterich. An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting and randomization. *Machine Learning*, 40(2):139–158, 2000.

- [62] A. Dimou, G. Tsoumakas, V. Mezaris, I. Kompatsiaris, and I. Vlahavas. An empirical study of multi-label methods for video annotation. In *Proc. 7th International Workshop on Content-Based Multimedia Indexing, CBMI 09*, Chania, Greece, 2009.
- [63] B. Done, P. Khatri, A. Done, and S. Draghici. Predicting novel human gene ontology annotations using semantic analysis. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 7(1):91-99, 2010.
- [64] R. Eisner, B. Poulin, D. Szafron, and P. Lu. Improving protein prediction using the hierarchical structure of the Gene Ontology. In *IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology*, 2005.
- [65] M Falda, S Toppo, A Pescarolo, E Lavezzo, Di Camillo B., A Facchinetti, E Cilia, R Velasco, and P. Fontana. Argot2: a large scale function prediction tool relying on semantic similarity of weighted Gene Ontology terms. *BMC Bioinformatics*, 13(Suppl 4:S14), 2012.
- [66] R.D. Finn, J. Tate, J. Mistry, P.C. Coghill, 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.
- [67] M. Frasca, A. Bertoni, M. Re, and G. Valentini. A neural network algorithm for semi-supervised node label learning from unbalanced data. *Neural Networks*, 43:84–98, 2013.
- [68] A. Freitas, D.C. Wieser, and R. Apweiler. On the importance of comprehensible classification models for protein function prediction. *IEEE ACM Trans. Comput. Biol. Bioinformatics*, 7(1):172–182, 2010.
- [69] I. Friedberg. Automated protein function prediction-the genomic challenge. *Brief. Bioinformatics*, 7:225–242, 2006.
- [70] J. Friedman and P. Hall. On Bagging and Nonlinear Estimation. Technical Report Tech. Report, Statistics Department, University of Stanford, CA, 2000.
- [71] M. Galar, A. Fernandez, E. Barrenechea, and H. Bustince. A review on ensembles for the class imbalance problem: Bagging-, boosting-, and hybrid-based approaches. *IEEE Transacriion on Systems, Man and Cybernetics, Part C*, 2013.
- [72] P. Gasch et al. Genomic expression programs in the response of yeast cells to environmental changes. *Mol.Biol.Cell*, 11:4241–4257, 2000.
- [73] Gene Ontology Consortium. True path rule, 2010. <http://www.geneontology.org/GO.usage.shtml#truePathRule>.
- [74] Gene Ontology Consortium. Gene Ontology annotations and resources. *Nucleic Acids Research*, 41:D530–535, 2013.

- [75] J. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. PowerGraph: Distributed graph-parallel computation on natural graphs. In *OSDI'12 Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation*, pages 17–30, Hollywood, CA, 2012.
- [76] 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.
- [77] H. Guo and H. Viktor. Learning from imbalanced data sets with boosting and data generation: the Databoost-IM approach. *ACM SIGKDD Explorations Newsletter*, 6(1):30–39, 2004.
- [78] T. Hamp et al. Homology-based inference sets the bar high for protein function prediction. *BMC Bioinformatics*, 14(Suppl 3:S7), 2013.
- [79] J. Hernandez, L.E. Sucar, and E. Morales. A hybrid global-local approach for hierarchical classification. In *Proc. of the Twenty-Sixth International Florida Artificial Intelligence Research Society Conference*, pages 432–437, 2013.
- [80] C. Huttenhower, M.A. Hibbs, C.L. Myers, A.A. Caudy, D.C. Hess, and O.G. Troyanskaya. The impact of incomplete knowledge on evaluation: an experimental benchmark for protein function prediction. *Bioinformatics*, 25(18):2404–2410, 2009.
- [81] L. Jensen, R. Gupta, H. Staerfeldt, and S. Brunak. Prediction of human protein function according to the Gene Ontology categories. *Bioinformatics*, 19(5):635–642, 2003.
- [82] 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.
- [83] A. Juncker, L. Jensen, A. Perleoni, A. Bernsel, M. Tress, P. Bork, G. von Heijne, A. Valencia, A. Ouzounis, R. Casadio, and S. Brunak. Sequence-based feature prediction and annotation of proteins. *Genome Biology*, 10:206, 2009.
- [84] U. Karaoz et al. Whole-genome annotation by using evidence integration in functional-linkage networks. *Proc. Natl Acad. Sci. USA*, 101:2888–2893, 2004.
- [85] W.K. Kim, C. Krumpelman, and E. Marcotte. Inferring mouse gene functions from genomic-scale data using a combined functional network/classification strategy. *Genome Biology*, 9:S5, 2008.
- [86] R.D. King, A. Karwath, A. Clare, and L. Dehaspe. The utility of different representations of protein sequence for predicting functional class. *Bioinformatics*, 17(5):445–454, 2001.

- [87] S. Kiritchenko, S. Matwin, and A. Famili. Hierarchical text categorization as a tool of associating genes with Gene Ontology codes. In *Proceedings of the 2nd European Workshop on Data Mining and Text Mining for Bioinformatics*, pages 26–30, 2004.
- [88] 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.
- [89] E.M. Kleinberg. On the Algorithmic Implementation of Stochastic Discrimination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(5):473–490, 2000.
- [90] D. Kocev, C. Vens, J. Struyf, and S. Dzeroski. Tree ensembles for predicting structured outputs. *Pattern Recognition*, 46(3):817–833, 2013.
- [91] T. Kohonen. The self-organizing map. *Neurocomputing*, 21:1–6, 1998.
- [92] D. Koller and M. Sahami. Hierarchically classifying documents using very few words. In *Proc. of the 14th Int. Conf. on Machine Learning*, pages 170–178, 1997.
- [93] Y. Kourmpetis, A. Van Dijk, M. Bink, R. Van Ham, and C. Ter Braak. Markov random field analysis for protein function prediction based on network data. *PLoS ONE*, 5(2:e9293), 2010.
- [94] Y. Kourmpetis, A. van Dijk, and C. ter Braak. Gene Ontology consistent protein function prediction: the FALCON algorithm applied to six eukaryotic genomes. *Algorithms for Molecular Biology*, 8, 2013.
- [95] M. Kubat and S. Matwin. Addressing the curse of imbalanced training sets: one-sided selection. In *Proceedings of the 14th International Conference on Machine Learning (ICML)*, pages 179–187, Nashville, USA, 1997.
- [96] 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.
- [97] A. Laegreid, H. Hvidsten, T. Midelfart, J. Komorowski, and A. Sandvik. Predicting Gene Ontology biological process from temporal gene expression patterns. *Genome Research*, 13(5):965–979, 2003.
- [98] C.H. Lampert and M.B. Blaschko. Structured prediction by joint kernel support estimation. *Machine Learning*, 77:249–269, 2009.
- [99] L. Lan, N. Djuric, Y. Guo, and Vucetic S. MS-kNN: protein function prediction by integrating multiple data sources. *BMC Bioinformatics*, 14(Suppl 3:S8), 2013.
- [100] 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.

- [101] G.R. Lanckriet, R. G. Gert, M. Deng, N. Cristianini, M. Jordan, and W.S. Noble. Kernel-based data fusion and its application to protein function prediction in yeast. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 300–311, 2004.
- [102] I. Lee, B. Ambaru, P. Pranjali-Thakkar, E.M. Marcotte, and S.Y. Rhee. Rational association of genes with traits using a genome-scale gene network for arabidopsis thaliana. *Nat. Biotechnol.*, 28:149–156, 2010.
- [103] 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.
- [104] 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.
- [105] Y. Loewenstein, D. Raimondo, O. Redfern, J. Watson, D. Frishman, M. Linial, C. Orengo, J. Thornton, and A. Tramontano. Protein function annotation by homology-based inference. *Genome Biology*, 10:207, 2009.
- [106] E.M. Marcotte, M. Pellegrini, M.J. Thompson, T.O. Yeates, and D. Eisenberg. A combined algorithm for genome-wide prediction of protein function. *Nature*, 402:83–86, 1999.
- [107] R. McDermott, J. Bumgarner and R. Samudrala. Functional annotation from predicted protein interaction networks. *Bioinformatics*, 21(15):3217–3226, 2005.
- [108] M. Mesiti, E. Jimenez-Ruiz, I. Sanz, R. Berlanga-Llavori, P. Perlasca, G. Valentini, and D. Manset. XML-Based Approaches for the Integration of Heterogeneous Bio-Molecular Data. *BMC Bioinformatics*, 10(S7):S12, 2009.
- [109] M. Mesiti, M. Re, and G. Valentini. Scalable network-based learning methods for automated function prediction based on the neo4j graph-database. In *Automated Function Prediction SIG 2013 - ISMB Special Interest Group Meeting*, pages 6–7, Berlin, 2013.
- [110] H.W Mewes et al. Overview of the yeast genome. *Nature*, 387:7–8, 1997.
- [111] H.W Mewes et al. MIPS: a database for genomes and protein sequences. *Nucleic Acids Res.*, 28:37–40, 2000.
- [112] S. Mostafavi and Q. Morris. Using the gene ontology hierarchy when predicting gene function. In *Proceedings of the Twenty-Fifth Annual Conference on Uncertainty in Artificial Intelligence (UAI-09)*, pages 419–427, Corvallis, Oregon, 2009. AUAI Press.
- [113] S. Mostafavi and Q. Morris. Fast integration of heterogeneous data sources for predicting gene function with limited annotation. *Bioinformatics*, 26(14):1759–1765, 2010.

- [114] S. Mostafavi, D. Ray, D. Warde-Farley, C. Grouios, and Q. Morris. GeneMANIA: a real-time multiple association network integration algorithm for predicting gene function. *Genome Biology*, 9(S4), 2008.
- [115] C. Myers and O. Troyanskaya. Context-sensitive data integration and prediction of biological networks. *Bioinformatics*, 23:2322–2330, 2007.
- [116] E. Nabieva, K. Jim, A. Agarwal, B. Chazelle, and M. Singh. Whole-proteome prediction of protein function via graph-theoretic analysis of interaction maps. *Bioinformatics*, 21(S1):302–310, 2005.
- [117] 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.
- [118] G. Obozinski, G. Lanckriet, C. Grant, Jordan. M., and W.S. Noble. Consistent probabilistic output for protein function prediction. *Genome Biology*, 9(S6), 2008.
- [119] O. Okun, G. Valentini, and M. Re. *Ensembles in Machine Learning Applications*, volume 373 of *Studies in Computational Intelligence*. Springer, Berlin, 2011.
- [120] S. Oliver. Guilt-by-association goes global. *Nature*, 403:601–603, 2000.
- [121] F. Otero, A. Freitas, and C. Johnson. A hierarchical classification ant colony algorithm for predicting gene ontology terms. In *Evolutionary Computation, Machine Learning and Data Mining in Bioinformatics*, number 5483 in Lecture Notes in Computer Science, pages 68–79. Springer-Verlag, 2009.
- [122] B. Paes, A. Plastion, and A. Freitas. Improving local per level hierarchical classification. *Journal of Information and Data Management*, 3(3):394–409, 2012.
- [123] G. Pandev, C.L. Myers, and V. Kumar. Incorporating functional inter-relationships into protein function prediction algorithms. *BMC Bioinformatics*, 10(1-142), 2009.
- [124] P. Pavlidis, J. Weston, J. Cai, and W.S. Noble. Learning gene functional classification from multiple data. *J. Comput. Biol.*, 9:401–411, 2002.
- [125] L. Pena-Castillo et al. A critical assessment of *Mus musculus* gene function prediction using integrated genomic evidence. *Genome Biology*, 9:S1, 2008.
- [126] A. Prlic, T.A. Down, E. Kulesha, R.D. Finn, A. Kahari, and T. Hubbard. Integrating sequence and structural biology with DAS. *BMC Bioinformatics*, 8(233), 2007.
- [127] K. Punera and J. Ghosh. Enhanced hierarchical classification via isotonic smoothing. In *WWW 2008*, pages 151–160, Beijing, China, 2008. ACM.

- [128] J.R. Quinlan. Induction of decision trees. *Machine Learning*, (1):81–106, 1986.
- [129] P. Radivojac et al. A large-scale evaluation of computational protein function prediction. *Nature Methods*, 10(3):221–227, 2013.
- [130] 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.
- [131] M. Re, M. Mesiti, and G. Valentini. A Fast Ranking Algorithm for Predicting Gene Functions in Biomolecular Networks. *IEEE ACM Transactions on Computational Biology and Bioinformatics*, 9(6):1812–1818, 2012.
- [132] M. Re and G. Valentini. Ensemble based data fusion for gene function prediction. In J. Kittler, J. Benediktsson, and F. Roli, editors, *Multiple Classifier Systems. Eighth International Workshop, MCS 2009, Reykjavik, Iceland*, volume 5519 of *Lecture Notes in Computer Science*, pages 448–457. Springer, 2009.
- [133] M. Re and G. Valentini. Prediction of gene function using ensembles of SVMs and heterogeneous data sources. In *Applications of supervised and unsupervised ensemble methods*, volume 245 of *Computational Intelligence Series*, pages 79–91. Springer, 2009.
- [134] M. Re and G. Valentini. An experimental comparison of Hierarchical Bayes and True Path Rule ensembles for protein function prediction. In N. El-Gayar, Kittler J., and F. Roli, editors, *Multiple Classifier Systems. Ninth International Workshop, MCS 2010, Cairo, Egypt*, volume 5997 of *Lecture Notes in Computer Science*, pages 294–303. Springer-Verlag, 2010.
- [135] M. Re and G. Valentini. Integration of heterogeneous data sources for gene function prediction using Decision Templates and ensembles of learning machines. *Neurocomputing*, 73(7-9):1533–37, 2010.
- [136] M. Re and G. Valentini. Noise tolerance of Multiple Classifier Systems in data integration-based gene function prediction. *Journal of Integrative Bioinformatics*, 7(3:139), 2010.
- [137] 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, Machine Learning in Systems Biology*, 8:98–111, 2010.
- [138] M. Re and G. Valentini. Cancer module genes ranking using kernelized score functions. *BMC Bioinformatics*, 13(Suppl 14/S3), 2012.
- [139] M. Re and G. Valentini. Ensemble methods: a review. In *Advances in Machine Learning and Data Mining for Astronomy*, Data Mining and Knowledge Discovery, pages 563–594. Chapman & Hall, 2012.

- [140] M. Re and G. Valentini. Large scale ranking and repositioning of drugs with respect to drugbank therapeutic categories. In *International Symposium on Bioinformatics Research and Applications (ISBRA 2012)*, volume 7292 of *Lecture Notes in Computer Science*, pages 225–236. Springer, 2012.
- [141] M. Re and G. Valentini. Network-based drug ranking and repositioning with respect to drugbank therapeutic categories. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2014, <http://doi.ieeecomputersociety.org/10.1109/TCBB.2013.62>
- [142] M. Riley. Functions of the gene products of Escherichia coli. *Microbiology Review*, 57:862–952, 1993.
- [143] V. Robles, P Larranaga, JM Pena, et al. Bayesian network multi-classifiers for protein secondary structure prediction. *Artificial Intelligence in Medicine*, 31(2):117–136, 2004.
- [144] J. Rousu, C. Saunders, S. Szedmak, and J. Shawe-Taylor. Kernel-based learning of hierarchical multilabel classification models. *Journal of Machine Learning Research*, 7:1601–1626, 2006.
- [145] A. Rozza, G. Lombardi, M. Re, E. Casiraghi, G. Valentini, and Campadelli P. A novel ensemble technique for protein subcellular location prediction. In *Ensembles in Machine Learning Applications*, volume 373 of *Studies in Computational Intelligence*, pages 151–177. Springer, Berlin, 2011.
- [146] 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.
- [147] D.E. Rumelhart, R. Durbin, and Y. Chauvin. Backpropagation: the basic theory. In Rumelhart D.E. Chauvin Y., editor, *Backpropagation: Theory, Architectures and Applications*, pages 1–34. Lawrence Erlbaum, Hillsdale, NJ, 1995.
- [148] Y. Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing Company, Boston, MA, 1996.
- [149] R. Schapire and Y. Singer. Boostexter: A boosting-based system for text categorization. *Machine Learning*, 39(2/3):135–168, 2000.
- [150] R.E. Schapire, Y. Freund, P. Bartlett, and W. Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651–1686, 1998.
- [151] L. Schietgat, C. Vens, J. Struyf, H. Blockeel, and S. Dzeroski. Predicting gene function using hierarchical multi-label decision tree ensembles. *BMC Bioinformatics*, 11(2), 2010.

- [152] B. Shahbaba and M. Neal. Gene function classification using Bayesian models with hierarchy-based priors. *BMC Bioinformatics*, 7(448), 2006.
- [153] R. Sharan, I. Ulitsky, and R. Shamir. Network-based prediction of protein function. *Mol. Sys. Biol.*, 8(88), 2007.
- [154] A. Sharkey, N. Sharkey, U. Gerecke, and G. Chandroth. The test and select approach to ensemble combination. In J. Kittler and F. Roli, editors, *Multiple Classifier Systems. First International Workshop, MCS 2000, Cagliari, Italy*, volume 1857 of *Lecture Notes in Computer Science*, pages 30–44. Springer-Verlag, 2000.
- [155] C. Silla and A. Freitas. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery*, 22(1-2):31–72, 2011.
- [156] N. Skunca et al. Quality of computationally inferred gene ontology annotations. *PLoS Computational Biology*, 8, 2012.
- [157] A.J. Smola and I.R. Kondor. Kernel and regularization on graphs. In B. Scholkopf and M.K. Warmuth, editors, *Proc. of the Annual Conf. on Computational Learning Theory*, Lecture Notes in Computer Science, pages 144–158. Springer, 2003.
- [158] S.Y. Sohna and H.W. Shinb. Experimental study for the comparison of classifier combination methods. *Pattern Recognition*, 40:33–40, 2007.
- [159] A. Sokolov and A. Ben-Hur. A structured-outputs method for prediction of protein function. In *MLSB08, the Second International Workshop on Machine Learning in Systems Biology*, 2008.
- [160] A. Sokolov and A. Ben-Hur. Hierarchical classification of Gene Ontology terms using the GOstruct method. *Journal of Bioinformatics and Computational Biology*, 8(2):357–376, 2010.
- [161] A. Sokolov, C. Funk, K. Graim, K. Verspoor, and A. Ben-Hur. Combining heterogeneous data sources for accurate functional annotation of proteins. *BMC Bioinformatics*, 14(Suppl 3:S10), 2013.
- [162] S. Sonnenburg, G. Ratsch, C. Schafer, and B. Scholkopf. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, 2006.
- [163] 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.
- [164] D. Stojanova, M. Ceci, D. Malerba and S. Deroski. Using PPI network autocorrelation in hierarchical multi-label classification trees for gene function prediction. *BMC Bioinformatics*, 14(285), 2013.

- [165] Y. Sun, M. Kamel, A. Wong, and Y. Wang. Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition*, 40:3358–3378, 2007.
- [166] Y. Tao, L. Sam, J. Li, C. Friedman, and Y.A. Lussier. Information theory applied to the sparse Gene Ontology annotation network to predict novel gene function. *Bioinformatics*, 23(13):i529–i538, 2007.
- [167] The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. *Nature Genet.*, 25:25–29, 2000.
- [168] W. Tian, L. Zhang, M. Tasan, F. Gibbons, O. King, J. Park, Z. Wunderlich, J.M. Cherry, and F.P. Roth. Combining guilt-by-association and guilt-by-profiling to predict *Saccharomyces cerevisiae* gene function. *Genome Biology*, 9:S7, 2008.
- [169] D. Titterton, 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.
- [170] A. Topchy, A.K. Jain, and W. Puch. Clustering Ensembles: Models of Consensus and Weak Partitions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1866–1881, 2005.
- [171] K. Trohidis, G. Tsoumahas, G. Kalliris, and I.P. Vlahavas. Multilabel classification of music into emotions. In *Proc. of the 9th International Conference on Music Information Retrieval*, pages 325–330, 2008.
- [172] O.G. Troyanskaya et al. A Bayesian framework for combining heterogeneous data sources for gene function prediction (in *saccharomyces cerevisiae*). *Proc. Natl Acad. Sci. USA*, 100:8348–8353, 2003.
- [173] I. Tsochantaridis, T. Joachims, T. Hoffman, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6:1453–1484, 2005.
- [174] G. Tsoumakas, I. Katakis, and I. Vlahavas. Random k-Labelsets for Multi-Label Classification. *IEEE Transactions on Knowledge and Data Engineering*, 23(7):1079–1089, 2011.
- [175] K. Tsuda, H. Shin, and B. Scholkopf. Fast protein classification with multiple networks. *Bioinformatics*, 21(Suppl 2):ii59–ii65, 2005.
- [176] A. Valencia. Automatic annotation of protein function. *Curr. Opin. Struct. Biol.*, 15:267–274, 2005.
- [177] G. Valentini. Mosclust: a software library for discovering significant structures in bio-molecular data. *Bioinformatics*, 23(3):387–389, 2007.
- [178] G. Valentini. True path rule hierarchical ensembles. In J. Kittler, J. Benediktsson, and F. Roli, editors, *Multiple Classifier Systems. Eighth International Workshop, MCS 2009, Reykjavik, Iceland*, volume 5519 of *Lecture Notes in Computer Science*, pages 232–241. Springer, 2009.

- [179] G. Valentini. True Path Rule hierarchical ensembles for genome-wide gene function prediction. *IEEE ACM Transactions on Computational Biology and Bioinformatics*, 8(3):832–847, 2011.
- [180] G. Valentini and N. Cesa-Bianchi. Hcgene: a software tool to support the hierarchical classification of genes. *Bioinformatics*, 24(5):729–731, 2008.
- [181] G. Valentini and F. Masulli. Ensembles of learning machines. In *Neural Nets WIRN-02*, volume 2486 of *Lecture Notes in Computer Science*, pages 3–19. Springer-Verlag, 2002.
- [182] G. Valentini, M. Muselli, and F. Ruffino. Cancer recognition with bagged ensembles of Support Vector Machines. *Neurocomputing*, 56:461–466, 2004.
- [183] 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.
- [184] A. Vazquez, A. Flammini, A. Maritan, and A. Vespignani. Global protein function prediction from protein-protein interaction networks. *Nature Biotechnology*, 21:697–700, 2003.
- [185] C. Vens, J. Struyf, L. Schietgat, S. Dzeroski, and H. Blockeel. Decision trees for hierarchical multi-label classification. *Machine Learning*, 73(2):185–214, 2008.
- [186] K. Verspoor, J. Cohn, S. Mnizewski, and C. Joslyn. A categorization approach to automated ontological function annotation. *Protein Science*, 15:1544–1549, 2006.
- [187] K. Verspoor, D. Dvorkin, K. Cohen, and L. Hunter. Ontology quality assurance through analysis of term transformations. *Bioinformatics*, 25(12):i77–i84, 2009.
- [188] C. von Mering, R. Krause, B. Snel, M. Cornell, S. Oliver, S. Fields, and P. Bork. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 417:399–403, 2002.
- [189] M.J. Wainwright and M.I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- [190] H. Wang, X. She, and W. Pan. Large margin hierarchical classification with mutually exclusive class membership. *Journal of Machine Learning Research*, 12:2649–2676, 2011.
- [191] J.Z. Wang, R. Du, R.S. Payattakil, P.S. Yu, and C.F. Chen. Improving GO semantic similarity measures by exploring the ontology beneath the terms and modelling uncertainty. *Bioinformatics*, 23(10):1274–1281, 2007.
- [192] C. Widmer and G. Ratsch. Multitask Learning in Computational Biology. *Journal of Machine Learning Research*, W&P, 27:207–216, 2012.

- [193] J. Xiong et al. Genome wide prediction of gene function via a generic knowledge discovery approach based on evidence integration. *BMC Bioinformatics*, 7(268), 2006.
- [194] H. Yang, T. Nepusz, and A. Paccanaro. Improving GO semantic similarity measures by exploring the ontology beneath the terms and modelling uncertainty. *Bioinformatics*, 28(10):1383–1389, 2012.
- [195] N. Youngs, D. Penfold-Brown, K. Drew, D. Shasha, and R. Bonneau. Parametric bayesian priors and better choice of negative examples improve protein function prediction. *Bioinformatics*, 29(9):1190–1198, 2013.
- [196] G. Yu, C. Domeniconi, H. Rangwala, G. Zhang, and Z. Yu. Transductive multi-label ensemble classification for protein function prediction. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1077–1085, 2012.
- [197] G. Yu, H. Rangwala, C. Domeniconi, G. Zhang, and Z. Yu. Protein function prediction with incomplete annotations. *IEEE Transactions on Computational Biology and Bioinformatics*, 2013. (in press).
- [198] G. Yu, H. Rangwala, C. Domeniconi, G. Zhang, and Z. Zhang. Protein function prediction by integrating multiple kernels. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, China, 2013.
- [199] H. Yu, L. Gao, K. Tu, and Z. Guo. Broadly predicting specific gene functions with expression similarity and taxonomy similarity. *Gene*, 352:75–81, 2005.
- [200] M.L. Zhang and Z.H. Zhou. Multi-label neural network with applications to functional genomics and text categorization. *IEEE Trans. on Knowledge and Data Engineering*, 18(10):1338–1351, 2006.
- [201] X. Zhang and D. Dai. A framework for incorporating functional interrelationships into protein function prediction algorithms. *IEEE ACM Transactions on Computational Biology and Bioinformatics*, 9(3):740–753, 2012.
- [202] Y. Zhang and D. Wang. Cost-sensitive ensemble method for class-imbalanced datasets. *Abstract and Applied Analysis*, 2013:6 pages, 2013.