

RESEARCH

Open Access



# A top-down supervised learning approach to hierarchical multi-label classification in networks

Miguel Romero\* , Jorge Finke and Camilo Rocha

\*Correspondence:  
miguelangelromero@javerianacali.edu.co  
Department of Electronics  
and Computer Science,  
Pontificia Universidad Javeriana,  
Calle 18 N 118-250, Cali 760031,  
Colombia

## Abstract

Node classification is the task of inferring or predicting missing node attributes from information available for other nodes in a network. This paper presents a general prediction model to *hierarchical* multi-label classification, where the attributes to be inferred can be specified as a strict poset. It is based on a top-down classification approach that addresses hierarchical multi-label classification with supervised learning by building a local classifier per class. The proposed model is showcased with a case study on the prediction of gene functions for *Oryza sativa Japonica*, a variety of rice. It is compared to the Hierarchical Binomial-Neighborhood, a probabilistic model, by evaluating both approaches in terms of prediction performance and computational cost. The results in this work support the working hypothesis that the proposed model can achieve good levels of prediction efficiency, while scaling up in relation to the state of the art.

**Keywords:** Hierarchical classification, Supervised learning, XGBoost, Top-down approach, *Oryza sativa*, Gene function prediction

## Introduction

Network representations provide a formal framework to specify relationships between interconnected entities (nodes). In a number of scenarios, such frameworks annotate nodes with attributes to naturally identify groups whose members are related by, e.g., a particular similarity measure. In analyzing networks with node attributes, most studies assume that a node can take a finite number of possible values, each of which is called a *class*. A class may represent the gender of a user in a social network or the function associated to a protein in a protein-protein interaction network.

## The problem

The task of inferring or predicting missing node attributes from information available for other nodes in a network is called *node classification* (also referred to as *attribute prediction*) (Bhagat et al. 2011). Several formulations to the node classification problem have been proposed over the past decades. The problem of inferring an attribute from exactly one of two classes is referred to as *binary* classification (Khan and Madden 2010).

The extension of the binary classification problem to *any* finite (non-zero) number of classes is referred to as *multi-class* classification (Mills 2021). Furthermore, each type of problem is categorized as a *multi-label* classification, if a node is allowed to be simultaneously associated to more than one class (Prajapati et al. 2012).

For most techniques that address the above-mentioned problems, node classification is generally carried out independently for each class (Abu-El-Haija et al. 2019; Bhagat et al. 2011; Hamilton et al. 2017; Kipf and Welling 2017; Xiao et al. 2021). The main limitation of such compartmentalized approaches is that they ignore hidden relationships among classes, even when certain class relationships may serve as an input to improve the accuracy of attribute prediction.

In practice, classes may have explicit relations specifying their dependencies. For example, this is the case with the Gene Ontology hierarchy because genes and proteins associated to a function must also be associated to the ancestors of such function. The authors in Silla and Freitas (2011), amid this limitation, define dependencies between classes as ancestral relations by means of a hierarchy represented by a directed acyclic graph. A connection from class  $C_1$  to class  $C_2$  in the hierarchy means that every node with attribute  $C_1$  also has attribute  $C_2$  (i.e., the nodes having attribute  $C_1$  are a subclass of the nodes having attribute  $C_2$ ).

Formally, a classification problem is considered *hierarchical* if and only if its hierarchy of classes is a strict partial order (i.e., a *strict poset* or, equivalently, a directed acyclic graph). A strict poset  $(C, <)$  over a finite set  $C$  of classes defines a binary relation  $<$  on  $C$  that is asymmetric, anti-reflexive, and transitive. For instance, the hierarchy of biological processes can be defined over a strict poset according to which the functions *cell death*, *programmed cell death*, and *apoptotic process* are ordered by *cell death*  $<$  *programmed cell death*  $<$  *apoptotic process* (Gene Ontology Consortium 2019). Note that the transitive property of the order  $<$  guarantees that *cell death* is also the ancestor of *apoptotic process*. Note also that since the strict poset is anti-reflexive, no process can be ancestor of itself. Finally, the asymmetry property guarantees that *apoptotic process* cannot be ancestor of *programmed cell death*. Indeed, any strict poset is isomorphic to a directed acyclic graph (DAG), a concept more closely related to graph and network analysis.

Given a graph with some labeled nodes (i.e., nodes associated to classes) and a class hierarchy, the expected outcome of a hierarchical classification problem is a collection of predicted associations between nodes and classes. An *inconsistent prediction* for a hierarchical multi-label classification problem refers to the fact that a node is inferred to have a particular class  $C$ , but the outcome of the classifier fails to infer the node's association to all ancestor classes of  $C$ . In other words, an inconsistent prediction states that the prediction does not satisfy the ancestral relations for some class  $C$ . In many scenarios, it is desirable to rule out inconsistent prediction: that is, if a classifier predicts a particular class  $C$  for a node, then it should also predict all the ancestors of  $C$  for that node; conversely, if a classifier does not predict  $C$  for a node, then it should not predict any of  $C$ 's descendants for that node. This constraint is often referred to as the *true-path rule* in Gene Ontology (Ashburner et al. 2000; Valentini 2009).

The efforts to classify nodes generally aim to comply with the underlying ancestral relations between classes in scenarios where such a hierarchy is known (and thus

avoid inconsistent prediction). Consider a social network where nodes represent individuals and node attributes represent different levels of education. If a predictor outputs an individual without an undergraduate degree as candidate to have graduate degree, it fails to comply with the hierarchical organization of educational levels, in many reasonable scenarios, thereby violating the true-path rule.

Hierarchical classification problems may further be categorized depending on the approach used for training the underlying model. On the one hand, a *top-down* approach involves a binary classifier for each class in the hierarchy. In this case, the classifier associated to each class is trained iteratively from the roots (i.e., the classes without ancestors in the hierarchy) to the leaves (i.e., the classes without descendants in the hierarchy). In addition, local information about the ancestors and descendants of a class in the hierarchy is used to avoid independent predictions. On the other hand, a *big-bang* approach involves a multi-label classifier that considers the entire hierarchy of ancestral relations at once. The multi-label classifier is trained just once with the information of every class in the hierarchy and its dependencies.

#### Related work

Several studies have applied both top-down and big-bang approaches across different domains (Jiang et al. 2008; Dimitrovski et al. 2010; Bi and Kwok 2011; Ramírez-Corona et al. 2016). The authors in Jiang et al. (2008) propose a top-down approach, called Hierarchical Binomial-Neighborhood (HBN), to predict protein functions in yeast *Saccharomyces cerevisiae*. It is shown by the authors that the hierarchical structure of functions can be exploited to completely avoid inconsistent predictions and, at the same time, outperform approaches based on independent class prediction. However, they point out that the main limitation of their approach is the high computational effort required for assigning probability weights to every protein-function pair. The authors in Ramírez-Corona et al. (2016) introduce a top-down approach based on Chained Path Evaluation (CPE), which uses a classifier to train each non-leaf class (i.e., each class with at least one descendant) in the hierarchy. Information on ancestral relations is included in the classifier by adding an extra feature with the prediction of parents of each class. As in Jiang et al. (2008), the computational cost of the CPE model grows exponentially as a function of the number of paths in the hierarchy. The use of big-bang approaches is, in general, also limited by their high computational demands. In Dimitrovski et al. (2010), for example, the authors present a big-bang approach that addresses hierarchical multi-label classification based on Predictive Clustering Trees (PCTs). The computational cost of the PCT approach is directly proportional to the size of the hierarchy.

Other studies address the node classification problem and obtain state-of-the-art performance for different case studies (see, e.g., Abu-El-Haija et al. 2019; Chen et al. 2021; Hamilton et al. 2017; Kipf and Welling 2017; Makrodimitris et al. 2020; Xiao et al. 2021). However, they do not take into account dependencies between classes (hierarchical or not), for they focus on multi-class instead of multi-label problems. For this reason, such developments can not be compared directly to assess hierarchical multi-label classification prediction.

### Main contribution

This work introduces a top-down classification approach that addresses hierarchical multi-label classification (HMC) using supervised learning. Given a network  $G = (V_G, E_G)$ , an assignment of classes to nodes in the network, and a class hierarchy specified as a directed acyclic graph  $H = (V_H, E_H)$ , the hierarchical multi-label classification problem is addressed by building a binary classifier for each class. Classifiers are built iteratively from the roots of the hierarchy to the leaves. The approach uses a correction mechanism to guarantee that the true-path rule is satisfied by the classifier's outcome; it is enforced by computing cumulative probabilities along the paths of classes in the input hierarchy.

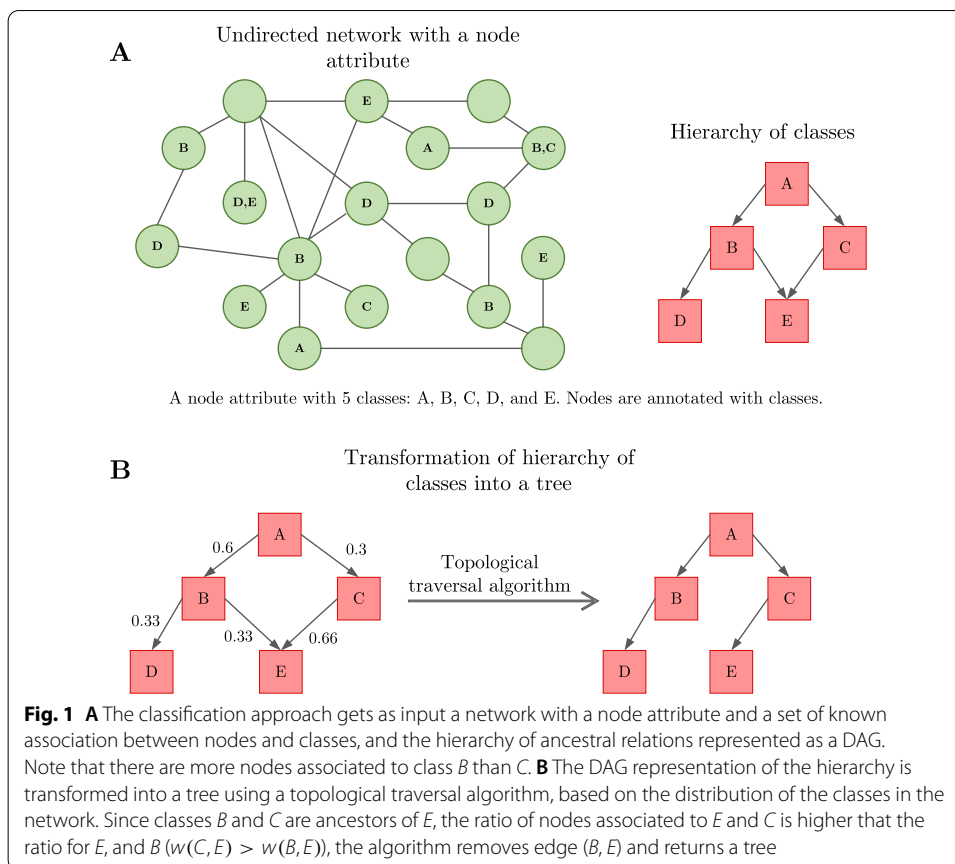
The results in this work support the working hypothesis that the proposed approach can achieve good levels of prediction efficiency, while scaling up in relation to the state of the art. This approach is showcased with a case study on the prediction of gene functions for *Oryza sativa Japonica*, a variety of rice. It is compared to the probabilistic HBN model (Jiang et al. 2008), by evaluating both approaches in terms of prediction performance (by means of the true positive and true negative rates) and in terms of their computational cost (by means of a comparison of the execution time). In the case study, the prediction task uses two inputs. Namely, (i) a gene co-expression network (GCN), in which a node represents a gene and a class of a node represents a gene function; and (ii) the hierarchical structure of biological processes defined in Gene Ontology Consortium (2019). The goal of the prediction task is to infer gene attributes from 15 sub-hierarchies grouping 1938 biological processes associated to 19663 genes.

**Outline.** The remainder of the paper is organized as follows. “Section [Hierarchical classification](#)” introduces the approach for node classification where classes have a hierarchical organization. “Section [Gene function prediction](#)” describes the problem of predicting gene functions. It also presents the results of applying the proposed model to *Oryza sativa Japonica*. Finally, “Section [Conclusion and future work](#)” draws some concluding remarks and future research directions.

### Hierarchical classification

This section presents a top-down classification approach in the form of a supervised learning model for hierarchical multi-label classification.

The input of the model are a graph  $G = (V_G, E_G)$  specifying an undirected network with nodes  $V_G$  and edges  $E_G$ , a directed acyclic graph  $H = (V_H, E_H)$ , with vertices  $V_H$  and edges  $E_H$  disjoint from  $V_G$  and  $E_G$ , respectively, representing the hierarchy of classes, and a function  $\phi : V_G \mapsto 2^{V_H}$  with a partial assignment of classes to nodes in the network. For  $v \in V_G$ , the set  $\phi(v) \subseteq V_H$  is the collection of classes initially associated to  $v$ . It is assumed that  $\phi$  satisfies the true-path rule for the hierarchy  $H$ , meaning that if a node  $v$  satisfies  $C \in \phi(v)$  for a class  $C \in V_H$ , then  $\phi(v)$  must contain all the ancestors of  $C$  in  $H$ . As mentioned in the introduction, the DAG  $H$  uniquely represents a strict poset. The goal of the model is then to build a function  $\phi' : V_G \mapsto 2^{V_H}$  extending  $\phi$  with new assignments of nodes in  $V_G$  to classes in  $V_H$ . Figure 1A depicts an example of the input of the model where the nodes of the network  $G$  are labeled with classes  $A-E$  and the hierarchy of classes  $H$  is a DAG. According to the true-path rule, nodes labeled with class  $E$  are also



related to classes *A*, *B*, and *C*. The objective is to predict new associations between nodes and classes for either nodes with or without labels.

The rest of this section is devoted to describe the main steps behind the construction of the supervised learning model.

### Hierarchy normalization

Hierarchies are represented as directed acyclic graphs where, in general, nodes can have any (finite) number of parents. Since the approach presented in this work assumes that every node has at most one parent, a topological traversal algorithm for directed graphs (see, e.g., Knuth 1997) is used to transform *H* into a tree, when required. In this way, the resulting model can take as input any hierarchy.

This algorithm uses the structure of *H* (not its tree version) and its distribution of classes. Given an ancestral relation  $A \rightarrow B$  (i.e., class *A* is a direct ancestor of class *B*), a weight  $w(A, B)$  for such an edge is defined as the ratio between the number of nodes associated to *B* (i.e., the size of the set  $\phi^{-1}(B)$ ) and the number of nodes associated to *A* (i.e., the size of the set  $\phi^{-1}(A)$ ). Since all nodes associated to *B* must be associated to *A*, then by definition each weight  $w(A, B)$  is in the range  $[0, 1]$ . For any node *B* with  $n \geq 1$  parents  $A_1, \dots, A_n$  in *H* (i.e.,  $A_i \rightarrow B$  in  $E_H$ , for  $1 \leq i \leq n$ ), the parent of *B* in the resulting tree is the node  $A_j$  maximizing  $w(A_j, B)$  among all the  $A_i$ 's. Ties are broken arbitrarily. This process can be effectively computed in time and space  $O(|V_H| + |E_H|)$ , namely, in

resources linear in the size of  $H$ . Such a process, based on a topological-sorting traversal, is described in Algorithm 1. Finally, note that the topological sorting of the vertices of  $H$  in Algorithm 1, can be exploited to compute the value of function  $w(\_, \_)$  by dynamic programming in space  $\Theta(|V_H|)$ . More precisely, a function  $\rho : V_H \rightarrow \mathbb{N}$  assigning to each class  $B_i$  its number of descendants  $\rho(B_i)$  in  $H$  can be computed from the direct descendants of  $B_i$ , which are processed before  $B_i$  in the topological sorting of  $V_H$ .

Algorithm 2.1: Topological-sorting based traversal for hierarchy normalization

---

```

1  input :  $H$ 
2  output :  $T = (V_T, E_T)$ 
3  set  $E_T = \{\}$ 
4  compute a topological sorting  $B_1, \dots, B_n$  of  $V_H$  (leaves first)
5  foreach class  $B_i$  with  $1 \leq i \leq n$ 
6    identify all parents  $A_1, \dots, A_m$  of  $B_i$  in  $H$ 
7    foreach parent  $A_j$  of  $B_i$  with  $1 \leq j \leq m$ 
8      compute the weight  $w(A_j, B_i)$ 
9      identify  $A_k$  with  $w(A_k, B_i) \leq w(A_j, B_i)$  for  $1 \leq j \leq m$ 
10     extend  $E_T$  with  $(A_k \rightarrow B_i)$ 
11 return the tree  $T$ 

```

---

As an example, consider the hierarchy depicted in Fig. 1B. Note that class  $E$  has more than one parent, there exists an ancestral relation from  $B$  to  $E$  and from  $C$  to  $E$  (i.e.,  $B \rightarrow E$  and  $C \rightarrow E$ , respectively). By the true-path rule, nodes associated to class  $D$  are also associated to class  $B$ , and the ones associated to class  $E$  are associated to both  $B$  and  $C$ . Since there are 4 nodes associated to  $E$ , 4 to  $D$ , 4 to  $B$ , and 2 to  $C$ , the weight  $w(B, E)$  is 0.33 and the weight  $w(C, E)$  is 0.66. Therefore, the topological-sorting traversal will remove edge  $B \rightarrow E$ .

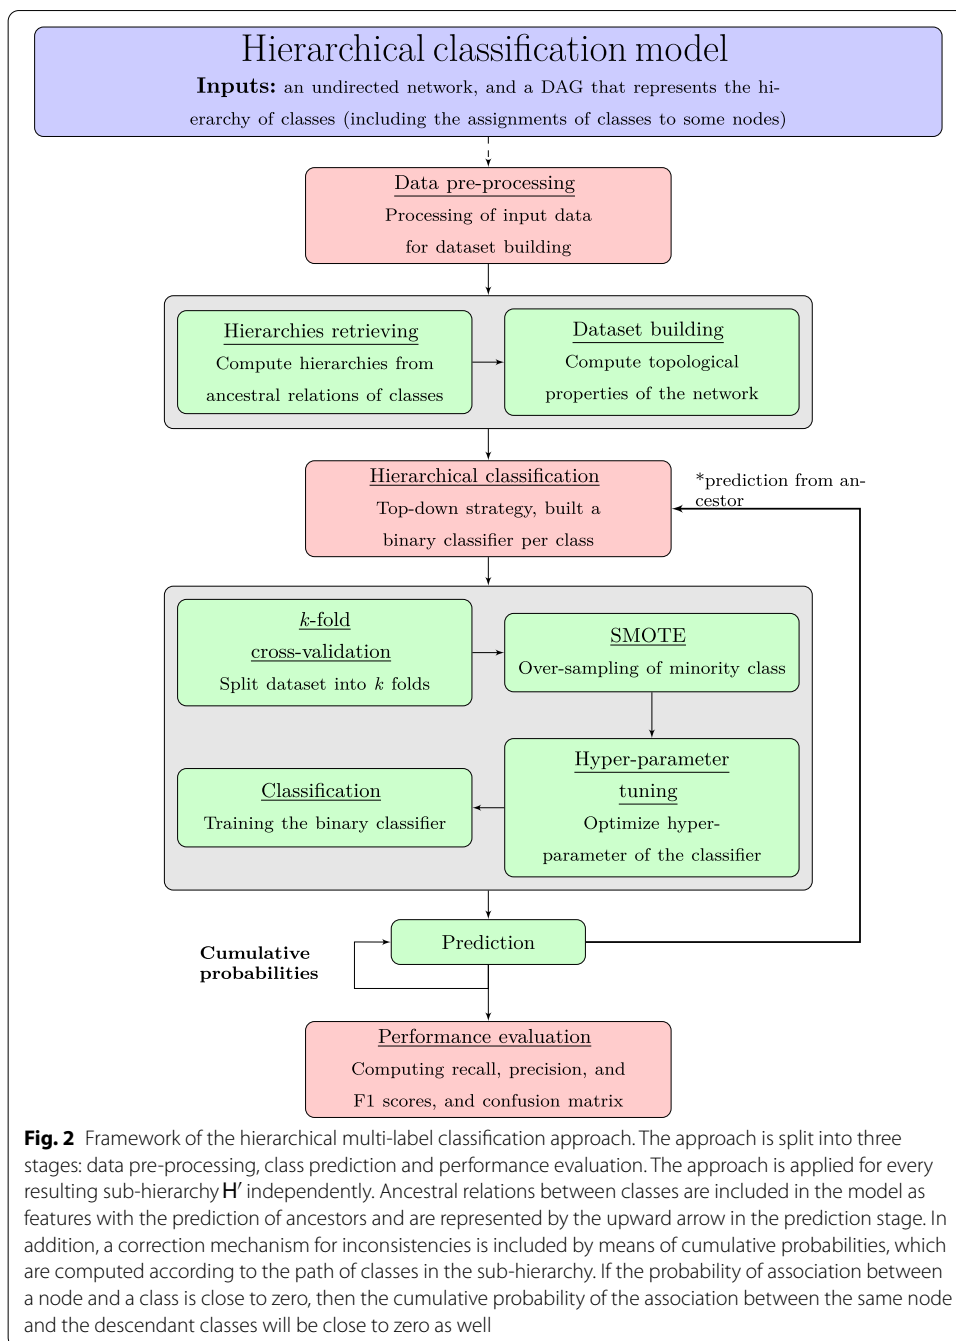
In the rest of this paper, it is assumed that hierarchy  $H$  is indeed a tree  $T$ .

### The model

Given the network  $G$  and the hierarchy tree  $T$ , the model is built in a process comprising three stages. Figure 2 depicts the general approach.

**Stage 0: data pre-processing.** In this stage, topological features of  $G$  and  $T$ , and hierarchical information in  $T$  are readied and combined for supervised learning.

Classes that are too specific or too general are ignored in the prediction to avoid overfitting and learning bias. In the case study presented in “Section [Gene function prediction](#)”, a class is defined as too specific or too general if it is associated to less than 5 or more than 300 genes, respectively (Jiang et al. 2008). As a result, the input hierarchy  $T$  can be split into several sub-trees, each one representing a sub-hierarchy  $T' = (V_{T'}, E_{T'})$  with  $V_{T'} \subseteq V_T$  and  $E_{T'} \subseteq E_T$ , over which the model is applied independently. That is, a sub-hierarchy  $T'$  is a subset of the classes and ancestral relations in  $T$ . As a matter of fact, this situation arises in the case study presented in “Section [Gene function prediction](#)”. Furthermore, each sub-hierarchy  $T'$  is associated to the subgraph of  $G$  consisting of all nodes labeled with the root class of  $T'$ , that is, each sub-hierarchy  $T'$  is related to a different subgraph  $G' = (V'_G, E'_G)$  with  $V'_G \subseteq V_G$  and  $E'_G \subseteq E_G$ . In this way, sub-hierarchies are considered independent problems with smaller inputs (*à la* divide and conquer).



For each sub-hierarchy  $T'$  in  $T$ , datasets are built based on two types of topological properties, namely, hand-crafted features and node embeddings. For the first type, properties of nodes  $V_G$  such as degree, average neighbor degree, centrality, and eccentricity are computed. Additionally, for each class  $C$  in  $T'$ , two features are computed to represent the probability of a node being associated to  $C$  and its parent in  $H$  based on the information of the neighborhood. For  $C$  and its parent, a node and its neighbors, and the associations between the neighbors and both classes, these new features represent the ratio between the number of neighbors associated to each class and the total number of

neighbors. For the second type of properties, continuous representations capturing the characteristics of the nodes in  $G'$  (i.e. node embeddings) are computed using `node2vec` (Grover and Leskovec 2016).

**Stage 1: hierarchical classification.** This stage comprises a top-down approach combining different supervised machine learning techniques/tools. It builds prediction classifiers for each sub-hierarchy  $T'$  independently. The approach uses stratified  $k$ -fold cross-validation, the Synthetic Minority Over-sampling Technique (SMOTE) (Chawla et al. 2002), hyper-parameter tuning (Bergstra and Bengio 2012), and a binary classifier, [e.g. XGBoost (Chen and Guestrin 2016) or graph convolutional networks (Kipf and Welling 2017)]. These techniques are combined sequentially in a pipeline, which is used iteratively from the root to the leaves of each sub-hierarchy  $T'$ . Note that, since the top-down approach builds a different classifier for each class in the sub-hierarchy, the proposed model can be used for multi-class and multi-label problems. As a result, nodes can be independently associated to multiple classes.

The combination of the above-mentioned techniques/tools makes up the core of the approach; and each technique has a different objective. Stratified  $k$ -fold aims to overcome overfitting by randomly selecting independent  $k$  subsets of the dataset where the distribution of the labels is similar for all folds. In this approach, 5 folds are used for cross-validation, that is the train-test ratio is 80/20. Over-sampling aims to overcome learning bias handling imbalanced datasets for underrepresented classes. SMOTE synthesizes new examples of the minority class from the existing ones. Hyper-parameter tuning aims to improve the performance of the prediction by optimizing parameters of the classifier such as, e.g., learning rate, number of estimators, and maximum depth of trees.

Two types of classifiers were used; namely, the XGBoost (Chen and Guestrin 2016) gradient boosting decision trees and graph convolutional networks (Kipf and Welling 2017). XGBoost was chosen for interpretability (Elshawi et al. 2019; Rudin 2019) and graph convolutional networks for state-of-the-art performance. In general, any other binary classifier can be used in this stage. The typical parameter values used for XGBoost classifiers are: *gbtree* booster, area under Precision-Recall (*aucpr*) evaluation metric, learning rate (*eta*) of 0.05, maximum tree depth (*max\_depth*) of 6, subsample ratio (*subsample*) of 0.9, and minimum sum of instance weight in a child (*min\_child\_weight*) of 3. For the graph convolutional networks, the implementation by Data61 (2018) was used with the following parameters: 16 layers of 16 units each, ReLU activation function, dropout rate of 50%, learning rate of 0.01, and binary cross-entropy loss function. Further details of the implementation can be founded in the repository [https://github.com/migueleci/node\\_classification](https://github.com/migueleci/node_classification).

Classifiers for each class in  $T'$  are built independently, so that there is no relation between their predictions. Including information from the ancestors of a class  $C$  into its classifier is not enough to avoid inconsistent predictions. For this reason, a correction mechanism is included in this stage. Since ensuring the true-path rule is key in the proposed approach, this stage computes cumulative probabilities along the paths of classes in  $T'$ . Namely, the probability of association between a node  $v$  and  $C$  is directly related to the predicted probabilities of the node being associated to all ancestors of  $C$ . Intuitively, the principle is as follows: if the probability of association of  $C$  to  $v$  is close to zero, then



the probability of association for all descendant classes of  $C$  to  $v$  will be close to zero as well. The main consequence of enforcing the principle is that the classification computed from the cumulative probability satisfies the true-path rule and removes the inconsistencies in the prediction.

**Stage 2: performance evaluation.** This stage comprises the evaluation of the metrics used for measuring the prediction performance of the classifiers. Performance evaluation focuses on recall (true positive rate) and precision scores. It also evaluates the precision-recall curve instead of the accuracy, loss, or ROC curves. This is mainly because datasets are often imbalanced (w.r.t. the positive class in a binary classification), thus both positive and negative classes of the binary classifier need to be analyzed separately. Recall and precision scores are computed from the predicted cumulative probabilities as a function of the *optimum threshold*, which is defined as the threshold that maximizes the F1 score from the precision-recall curve for the cumulative probabilities.

### Gene function prediction

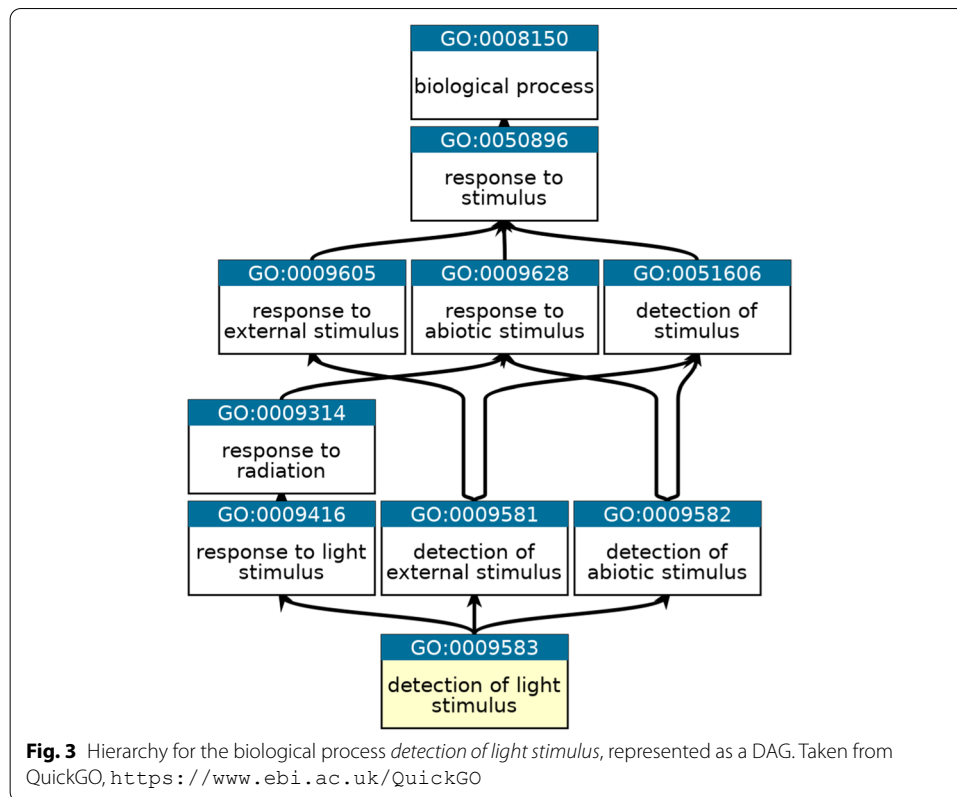
This section presents a case study on the prediction of gene functions (i.e., biological processes in which genes are involved) for the *Oryza sativa Japonica* rice variety. First, the problem of predicting gene functions is introduced. Then, the results after applying the approach proposed in Section 2 to this problem are described. The probabilistic approach, proposed in Jiang et al. (2008), is used to compare the novel results.

### Gene co-expression networks

High-throughput sequencing technologies have enabled the identification of numerous genes and gene products. However, biological processes in which many such genes are involved remain largely unknown (i.e., relations between genes and biological processes have not been comprehensively validated through in vivo experimentation) (Ranganathan et al. 2019). Identifying the functions of genes is key to enhance the understanding on how to characterize the genome of a particular organism. In general, traditional in silico approaches to predict gene functions consider each function as an independent class. The task is generally defined as a binary classification problem based on gene expression.

Genes (or gene products) can be associated to more than one biological process and such processes may be related (e.g., by ancestral relationships). The assignment of functions to genes obeys the true-path rule (Valentini 2009). Consequently, efforts to predict whether a gene is associated to a particular function should consider the ancestral relations of that function. Ignoring such a hierarchical structure leads to biological inconsistencies in the outcome of the prediction. On the contrary, when a gene is associated to multiple biological processes without ancestral relations, the prediction is done for each one of the functions independently.

As an example, consider two biological processes in which a gene may be involved: *response to external stimulus* and *detection of light stimulus*. According to the hierarchy of biological processes in Fig. 3, the former function is an ancestor of the latter (Ashburner et al. 2000). By the true-path rule, if a gene is associated to *detection of light stimulus*, it must be associated to *response to external stimulus*.



In addition, a common approach to integrate large volumes of transcriptional data and synthesize the hierarchical structure of gene functions is to characterize gene co-expression networks (GCNs). GCNs have been used to infer biological processes and pathways based on highly correlated expression patterns between genes (Oti et al. 2008; Dam et al. 2017; Vandepoele et al. 2009). It is well-known that co-expressed genes, i.e., genes with similar expression profiles, tend to share the same function or be related to the same regulatory pathway (Emamjomeh et al. 2017; Serin et al. 2016; Zhou et al. 2002).

#### Predicting gene functions in *Oryza sativa Japonica*

The goal of this case study is to predict gene functions, that is, the biological processes in which some genes are involved. The problem is tackled by using the model proposed in Section 2 on the GNC of *Oryza sativa Japonica* (Obayashi et al. 2018) and a hierarchy of biological processes for this organism (Sakai et al. 2013). The computational experiments supporting the results in this section have been executed in a cluster with 5 nodes, each one with 64GB of memory and a AMD Opteron™ Processor 6376 with 64 CPU cores.

Formally, a gene co-expression network is represented as a undirected, weighted graph  $G = (V_G, E_G, f)$ , built from empirical data, where genes are represented by nodes  $V_G$ , edges  $E_G$  denote co-expression relationships, and the weight  $f : E_G \rightarrow \mathbb{R}_{\geq 0}$  measures the level of co-expression between genes. Additionally, the graph  $H = (V_H, E_H)$  is a directed acyclic graph (DAG) which represents the hierarchical organization of biological processes, where  $E_H$  represents the ancestral relations between functions. Genes are associated to one or more biological processes through a function  $\phi : V_G \rightarrow 2^{V_H}$ ,

where  $V_H$  denotes the set of all biological processes. The predictive model combines the existing set of labels in  $\phi$ , topological properties of  $G$  and the hierarchical information of  $H$  to obtain a new labeling function  $\phi'$  using the hierarchical multi-label classification approach. As a result, the function  $\phi'$  contains suggestions of previously unidentified associations between genes and functions satisfying the true-path rule.

The set of known associations between genes and functions used in this work contains 19663 rice genes, 550813 co-expression relations, 3743 biological processes, 220598 assignments of functions to genes, and 7185 ancestral relations of functions (all biological processes belong to the same hierarchy). To avoid overfitting and learning bias in the proposed model, only those functions associated to more than 4 and at most 300 genes are considered (Jiang et al. 2008). Under this criterion, 1938 functions (52%) are used for prediction. As a result, the function hierarchy breaks down into 27 sub-hierarchies, from which 12 correspond to isolated functions or small sub-hierarchies (fewer than 7 functions). The 15 remaining sub-hierarchies are described in Table 1, sorted from the smallest to the largest in terms of the number of functions  $V_{\Gamma}$  and number genes associated with each of them.

The prediction performance of the proposed approach is compared with the HBN model presented in Jiang et al. (2008). The HBN model uses a *top-down* approach that integrates relational data of protein-protein interaction network (PPI) with the hierarchical data of biological processes with the objective of predicting protein functions. For this case study, the HBN model is adapted to the problem of predicting gene functions based on GCNs. To predict the probability of a gene  $g$  being associated to function  $A$ , the local neighborhood information of  $g$  in the GCN and the ancestors of  $A$  in the hierarchy are considered. The HBN model computes the probability of gene  $g$  being associated to function  $A$  obeying the true-path rule.

The figures in this section show the mean performance for the proposed approach and the HBN model between multiple experiments, in which each experiment represents the mean performance between the  $k$  folds used for cross-validation. In all of them, the

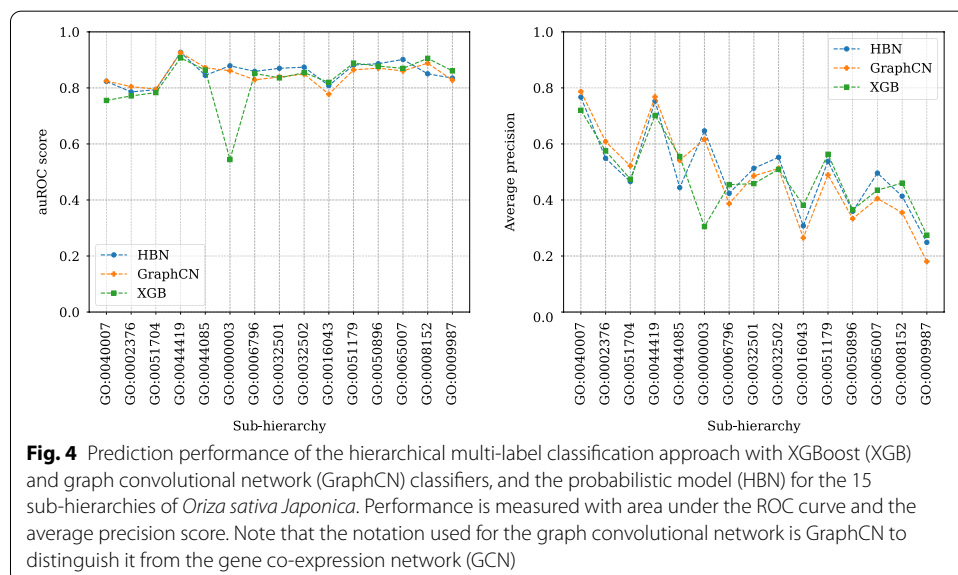
**Table 1** Sub-hierarchies generated for the gene co-expression network of *Oryza sativa Japonica*

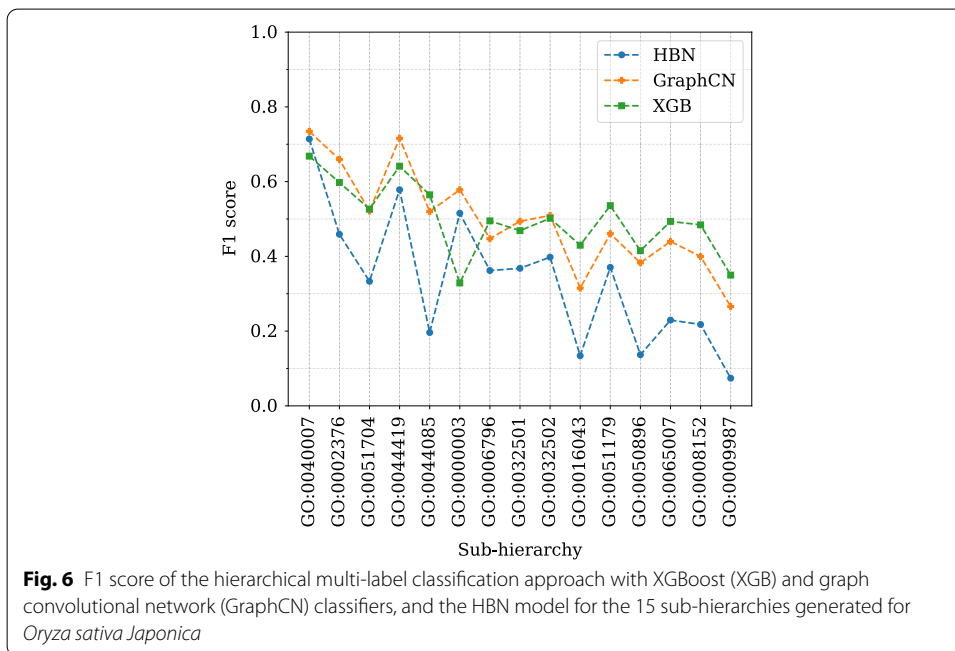
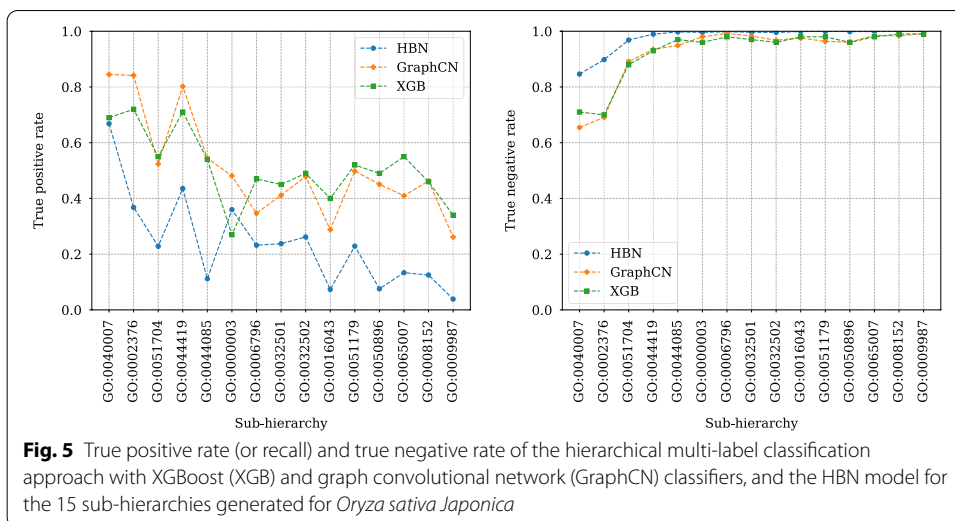
	Root	Func	Genes	Desc
1	GO:0040007	10	108	Growth
2	GO:0002376	11	131	Immune system process
3	GO:0051704	22	144	Multi-organism process
4	GO:0044419	37	777	Interspecies interaction between organisms
5	GO:0044085	50	377	Cellular component biogenesis
6	GO:0000003	72	648	Reproduction
7	GO:0006796	118	1270	Phosphate-containing compound metabolic process
8	GO:0032501	118	1043	Multicellular organismal process
9	GO:0032502	149	1063	Developmental process
10	GO:0016043	140	661	Cellular component organization
11	GO:0051179	164	1350	Localization
12	GO:0050896	261	3319	Response to stimulus
13	GO:0065007	485	2224	Biological regulation
14	GO:0008152	775	5862	Metabolic process
15	GO:0009987	925	5900	Cellular process

variation (error bar or standard deviation) is not included because it is negligible (and can add visual noise to the plots). Figure 4 illustrates the performance of the proposed approach using XGBoost (XGB) and graph convolutional network (GraphCN) classifiers and the HBN model measured with the area under the ROC curve and the average precision score. Note that their performance seem to be similar in most sub-hierarchies and it is not possible to conclude which one performs better from Fig. 4. However, since only biological processes associated to more than 4 and less than 300 genes are considered (less than 2% of the genes in the GCN), datasets generated for the filtered biological processes are highly imbalanced. For this reason, the area under the ROC curve is not suitable for the case study (this measured is biased for the over-represented class in the classification task), and the analysis should focus on other metrics such as recall and F1 score instead.

An outstanding difference between the performance of the proposed approach and the HBN model is observed when the confusion matrices are analyzed. Figure 5 shows the true positive rate (or the measure of recall) and the true negative rate for the 15 sub-hierarchies. Note that the true positive rate of the proposed approach is higher than the HBN model for most of the sub-hierarchies, whereas the true negative rate of the HBN model is higher for all sub-hierarchies. However, the HBN model is biased for the negative class because the probability predicted by the HBN model for most of the associations between genes and functions is close to zero. As the datasets are highly imbalanced, the performance in terms of the positive class are key to determine which approach is adequate. Recall that a dataset is said to be *imbalanced* for binary classification if one of the classes is under-represented in relation to the other one, i.e., the number of instances related to one class is much higher than the number of instances related to the other. For example, a dataset with 1000 instances that has 900 negative and 100 positive samples is imbalanced.

The true positive rate illustrated in Fig. 5 shows that the proposed approach outperforms the HBN model in the identification of the (positive) associations between genes





and functions. The performance varies between XGBoost and graph convolutional networks, but both classifiers have better overall performance than the HBN model. The results suggest that graph convolutional networks are better for small sub-hierarchies, while XGBoost is better for larger ones. Even though the true negative rate of the HBN model is close to 1 for all sub-hierarchies, as illustrated on Fig. 5, the performance of the proposed approach in terms of the average of both recall and precision (i.e., F1 score) is better than the HBN model. Figure 6 presents the F1 score of the proposed approach and the HBN model for the 15 sub-hierarchies. In this case study there is no observable correlation between the size/depth/span of a hierarchy and the prediction performance, according to the experiments. This is coherent with the overall computational

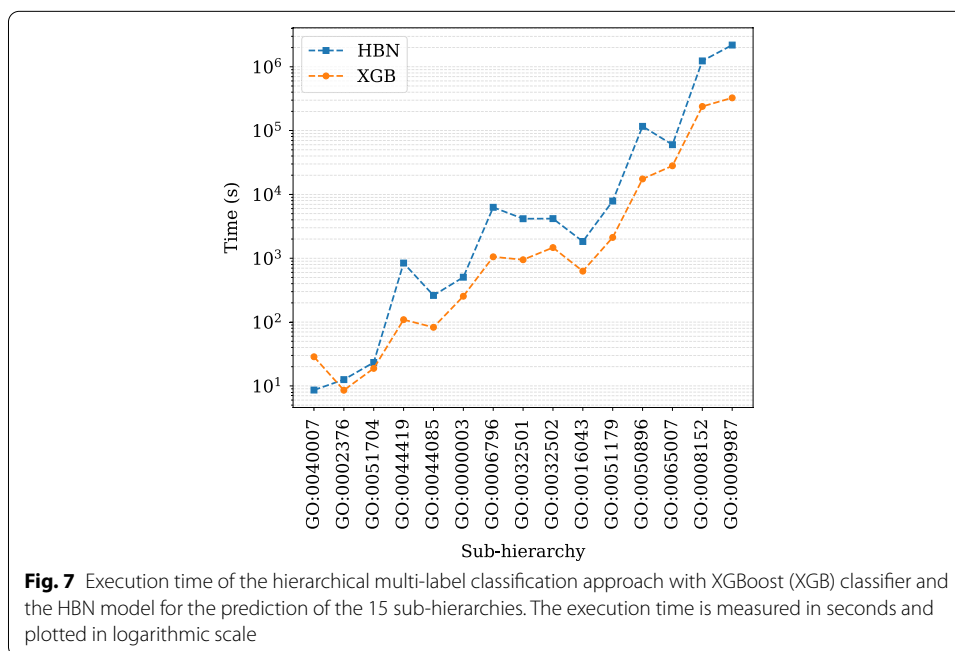
complexity of the algorithms. On the other hand, there is no experimental evidence to suggest that some degree of correlation exists between the number of label nodes and the prediction performance. However, these observations need to be further investigated with other case studies.

Finally, the execution time of the proposed approach and the HBN model is illustrated on Fig. 7. The execution time for the graph convolutional network classifier is not included because the experiments were executed on CPUs rather than GPU. It is known that neural networks run much faster on GPUs; thus, it would not be fair to make a comparison with the available data. Note that the execution time is measured in seconds and plotted on a logarithmic scale. Except for the smallest sub-hierarchy (GO:0040007), the execution time of the proposed approach, using XGBoost classifier, is better than the HBN model. On average, the execution time of the HBN model is approximately 4 times as much of the proposed approach.

### Conclusion and future work

By combining different techniques from machine learning, the hierarchical multi-label classification model presented in this paper introduces an approach to address the node classification problem for scenarios in which nodes can have attributes obeying a hierarchical organization. Taken into account hierarchical dependencies is shown to be a key aspect for obtaining more consistent predictions that satisfy the true-path rule.

A baseline comparison between the proposed approach using two different classification methods, namely, gradient boosting decision trees and graph convolutional networks, and the HBN model introduced by Jiang et al. (2008) is presented. Both approaches are applied to the problem of predicting gene function on the variety of rice *Oryza sativa Japonica*. The proposed hierarchical multi-label classification approach outperforms the HBN model in two aspects. First, using topological information of



the network is a key feature to obtain the overall best performance of the prediction. In such setting, the true positive rates of the proposed approach are significantly higher than the HBN model, whereas the true negative rates yield similar values (close to one). This result suggests that the proposed approach can lead to good prediction of associations between genes and functions in *Oryza sativa Japonica* and, potentially, in other organisms.

For scenarios in which the classes of the hierarchy are under-represented, i.e., datasets are imbalanced, it is important to center the performance analysis on metrics that are not biased by the imbalanced dataset. Such metrics include the true positive rate (or the measure of recall), the true negative rate, and the F1-score. Other widely-used metrics, like the area under ROC curve and the measure of average precision, are misleading for evaluating the performance of a classifier under such conditions.

Second, the execution time of the proposed approach for the XGBoost classifier is, on average, 4 times better than that of the HBN model. The reduction in computational cost of the proposed top-down approach can be attributed to the fact that it predicts the probability of associations between a class and every node of the network at the same time. Also, the efficient computation of the DAG into a tree helps in making the proposed approach relevant to analyze larger networks and hierarchies.

Finally, although the performance of the proposed approach is promising, it requires to gather sufficient information from node classes, which in some cases is incomplete or unavailable. For example, information about gene functions is limited for many genes and gene products. For some organisms there is no such information available at all. The shortage of information may lead to over-fitting or learning bias in the approach, and consequently to misleading conclusions. Including other networks as additional sources of information for the classification problem seems to be interesting for future work. Other networks can be added with the help of transfer learning techniques. For example, by creating new features that aggregate the information extracted from other networks that can be integrated in the proposed approach as additional input to improve the prediction performance. Furthermore, other approaches such as semi-supervised and transductive learning can also be considered for future work to handle the amount of data required for training.

#### Abbreviations

DAG: Directed acyclic graph; GO: Gene ontology; GCN: Gene co-expression network; HBN: Hierarchical binomial-neighborhood; HMC: Hierarchical multi-label classification; ROC: Receiver operating characteristic; SMOTE: Synthetic minority over-sampling technique.

#### Acknowledgements

Not applicable.

#### Authors' contributions

M.R. proposed the original idea. J.F. and C.R. provide advice on algorithms concepts and implementation. M.R. structured the methodology and performed the analysis. M.R., J.F., and C.R. wrote the manuscript. All authors read and approved the final manuscript.

#### Authors' information

**Miguel Romero** Ph.D. Student in Engineering and Applied Sciences at the Pontificia Universidad Javeriana, in Cali (Colombia). He earned a B.S. degree in Economics and Systems Engineering from the Escuela Colombiana de Ingeniería, Bogotá (Colombia). He has experience in rewrite logic, network analysis, algorithms, and competitive programming. He works on the development of mathematical models and algorithms that allow identifying, from in-silico omic characterization, the expression of phenotypic traits in different varieties of crops.

**Jorge Finke** Professor in the Department of Electronics and Computer Science at the Pontificia Universidad Javeriana, in Cali (Colombia). He earned a B.S., a M.Sc. and a Ph.D. degree in Systems theory from The Ohio State University, Columbus, OH. More than 10 years of experience in developing representations of complex, large-scale data, and innovative approaches to predictive modeling and analysis, including cutting-edge research in academia as well as real-world consulting for companies across multiple industries. My motivation is to discover new and disruptive opportunities that enable organizations to solve problems and achieve strategic goals.

**Camilo Rocha** Associate Professor in the Department of Electronics and Computer Science at the Pontificia Universidad Javeriana, in Cali (Colombia). Starting February 2020, he have been appointed Dean of Engineering and Sciences. He earned a B.S. and a M.Sc. degree in Informatics from the Universidad de los Andes (Bogotá), and a M.Sc. degree in Mathematics and a Ph.D. degree in Computer Science from the University of Illinois at Urbana-Champaign. His main research interests are in formal methods, algorithms, and software engineering, more specifically on techniques for building reliable software systems.

#### Funding

This work was funded by the OMICAS program: Optimización Multiescala In-silico de Cultivos Agrícolas Sostenibles (Infraestructura y Validación en Arroz y Caña de Azúcar), anchored at the Pontificia Universidad Javeriana in Cali and funded within the Colombian Scientific Ecosystem by The World Bank, the Colombian Ministry of Science, Technology and Innovation, the Colombian Ministry of Education and the Colombian Ministry of Industry and Tourism, and ICETEX, under GRANT ID: FP44842-217-2018.

#### Availability of data and materials

The datasets analyzed for the current study are publicly available from different sources. They can be found in the following locations: (i) Gene co-expression data of *Oryza sativa Japonica* is available on ATTED-II (Obayashi et al. 2018). (ii) Functional data of rice genes is available on Sakai et al. (2013) and Kurata and Yamazaki (2006). The data collected, cleaned, and processed from the above sources as used in the case study can be requested to the authors. A workflow implementation is publicly available: (i) Project name: Node Classification (ii) Project home page: [https://github.com/miguelceci/node\\_classification](https://github.com/miguelceci/node_classification) (iii) Operating system(s): platform independent. (iv) Programming language: Python 3. (v) Other requirements: None. (vi) License: GNU GPL v3.

#### Declarations

##### Competing interests

The authors declare that they have no competing interests.

Received: 18 September 2021 Accepted: 28 January 2022

Published online: 10 February 2022

#### References

- Abu-El-Haija S, Perozzi B, Kapoor A, Lee J (2019) N-gcn: Multi-scale graph convolution for semi-supervised node classification. In: Conference on uncertainty in artificial intelligence (UAI)
- Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, IsseI-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G (2000) Gene ontology: tool for the unification of biology. *Nat Genet* 25(1):25–29. <https://doi.org/10.1038/75556>
- Bergstra J, Bengio Y (2012) Random search for hyper-parameter optimization. *J Mach Learn Res* 13(10):281–305
- Bhagat S, Cormode G, Muthukrishnan S (2011) Node classification in social networks. In: Aggarwal CC (ed) *Social network data analytics*. Springer US, Boston, pp 115–148. [https://doi.org/10.1007/978-1-4419-8462-3\\_5](https://doi.org/10.1007/978-1-4419-8462-3_5)
- Bi W, Kwok JT (2011) Multi-label classification on tree- and dag-structured hierarchies. In: Proceedings of the 28th international conference on international conference on machine learning. ICML'11. Omnipress, Madison, pp 17–24
- Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP (2002) SMOTE: synthetic minority over-sampling technique. *J Arti Intell Res* 16:321–357. <https://doi.org/10.1613/jair.953>. arXiv:1106.1813
- Chen T, Guestrin C (2016) XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp 785–794. <https://doi.org/10.1145/2939672.2939785>. arXiv:1603.02754
- Chen Q, Li Y, Tan K, Qiao Y, Pan S, Jiang T, Chen Y-PP (2021) Network-based methods for gene function prediction. *Brief Funct Genomics* 20(4):249–257. <https://doi.org/10.1093/bfgp/elab006>
- Data61 C (2018) StellarGraph Machine Learning Library. GitHub
- Dimitrovski I, Kocev D, Loskovska S, Džeroski S (2010) Detection of visual concepts and annotation of images using ensembles of trees for hierarchical multi-label classification. In: Ünay D, Çataltepe Z, Aksoy S (eds) *Recognizing patterns in signals, speech, images and videos*, vol 6388. Springer, Berlin, pp 152–161. [https://doi.org/10.1007/978-3-642-17711-8\\_16](https://doi.org/10.1007/978-3-642-17711-8_16)
- Elshawi R, Al-Mallah MH, Sakr S (2019) On the interpretability of machine learning-based model for predicting hypertension. *BMC Med Inform Decis Mak* 19(1):146. <https://doi.org/10.1186/s12911-019-0874-0>
- Emamjomeh A, Saboori Robat E, Zahiri J, Solouki M, Khosravi P (2017) Gene co-expression network reconstruction: a review on computational methods for inferring functional information from plant-based expression data. *Plant Biotechnol Rep* 11(2):71–86. <https://doi.org/10.1007/s11816-017-0433-z>
- Gene Ontology Consortium (2019) The Gene Ontology Resource: 20 years and still GOing strong. *Nucleic Acids Research* 47(D1):330–338. <https://doi.org/10.1093/nar/gky1055>



- Grover A, Leskovec J (2016) node2vec: Scalable feature learning for networks. [arXiv:1607.00653](https://arxiv.org/abs/1607.00653)
- Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. In: Proceedings of the 31st international conference on neural information processing systems, NIPS'17, pp. 1025–1035. Curran Associates Inc., Red Hook
- Jiang X, Nariai N, Steffen M, Kasif S, Kolaczyk ED (2008) Integration of relational and hierarchical network information for protein function prediction. *BMC Bioinform* 9(1):350. <https://doi.org/10.1186/1471-2105-9-350>
- Khan SS, Madden MG (2010) A survey of recent trends in one class classification. In: Coyle L, Freyne J (eds) *Artificial intelligence and cognitive science*, vol 6206. Springer, Berlin, pp 188–197. [https://doi.org/10.1007/978-3-642-17080-5\\_21](https://doi.org/10.1007/978-3-642-17080-5_21)
- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: International conference on learning representations (ICLR)
- Knuth DE (1997) *The art of computer programming*, 3rd edn. Addison-Wesley, Reading
- Kurata N, Yamazaki Y (2006) Oryzabase. An integrated biological and genome information database for rice. *Plant Physiol* 140(1):12–17. <https://doi.org/10.1104/pp.105.063008>
- Makrodimitis S, van Ham RCHJ, Reinders MJT (2020) Automatic gene function prediction in the 2020's. *Genes* 11(11):1264. <https://doi.org/10.3390/genes11111264>
- Mills P (2021) Solving for multi-class: a survey and synthesis. [arXiv:1809.05929](https://arxiv.org/abs/1809.05929) [cs, stat]. [arXiv:1809.05929](https://arxiv.org/abs/1809.05929)
- Obayashi T, Aoki Y, Tadaka S, Kagaya Y, Kinoshita K (2018) ATTED-II in 2018: a plant coexpression database based on investigation of the statistical property of the mutual rank index. *Plant Cell Physiol* 59(1):3–3. <https://doi.org/10.1093/pcp/pcx191>
- Oti M, van Reeuwijk J, Huynen MA, Brunner HG (2008) Conserved co-expression for candidate disease gene prioritization. *BMC Bioinform* 9(1):208. <https://doi.org/10.1186/1471-2105-9-208>
- Prajapati P, Thakkar A, Ganatra A (2012) A survey and current research challenges in multi-label classification methods. *Int J Soft Comput Eng (IJSCE)* 2(1):248–252
- Ramírez-Corona M, Sucar LE, Morales EF (2016) Hierarchical multilabel classification based on path evaluation. *Int J Approx Reason* 68:179–193. <https://doi.org/10.1016/j.ijar.2015.07.008>
- Ranganathan S, Gribskov MR, Nakai K, Schönbach C (2019) *Encyclopedia of bioinformatics and computational biology*. Elsevier, OCLC: 1052465484
- Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell* 1(5):206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Sakai H, Lee SS, Tanaka T, Numa H, Kim J, Kawahara Y, Wakimoto H, Yang C-C, Iwamoto M, Abe T, Yamada Y, Muto A, Inokuchi H, Ikemura T, Matsumoto T, Sasaki T, Itoh T (2013) Rice annotation project database (RAP-DB): an integrative and interactive database for rice genomics. *Plant Cell Physiol* 54(2):6–6. <https://doi.org/10.1093/pcp/pcs183>
- Serin EAR, Nijveen H, Hilhorst HWM, Ligterink W (2016) Learning from co-expression networks: possibilities and challenges. *Front Plant Sci*. <https://doi.org/10.3389/fpls.2016.00444>
- Silla CN, Freitas AA (2011) A survey of hierarchical classification across different application domains. *Data Min Knowl Disc* 22(1–2):31–72. <https://doi.org/10.1007/s10618-010-0175-9>
- Valentini G (2009) True path rule hierarchical ensembles. In: Benediktsson JA, Kittler J, Roli F (eds) *Multiple classifier systems*, vol 5519. Springer, Berlin, pp 232–241. [https://doi.org/10.1007/978-3-642-02326-2\\_24](https://doi.org/10.1007/978-3-642-02326-2_24)
- van Dam S, Vösa U, van der Graaf A, Franke L, de Magalhães JP (2017) Gene co-expression analysis for functional classification and gene-disease predictions. *Brief Bioinform*. <https://doi.org/10.1093/bib/bbw139>
- Vandepoele K, Quimbaya M, Casneuf T, De Veylder L, Van de Peer Y (2009) Unraveling transcriptional control in arabidopsis using cis-regulatory elements and coexpression networks. *Plant Physiol* 150(2):535–546. <https://doi.org/10.1104/pp.109.136028>
- Xiao S, Wang S, Dai Y, Guo W (2021) Graph neural networks in node classification: survey and evaluation. *Mach Vis Appl* 33(1):4. <https://doi.org/10.1007/s00138-021-01251-0>
- Zhou X, Kao M-CJ, Wong WH (2002) Transitive functional annotation by shortest-path analysis of gene expression data. *Proc Natl Acad Sci* 99(20):12783–12788. <https://doi.org/10.1073/pnas.192159399>

## Publisher's Note

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

Submit your manuscript to a SpringerOpen® journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)