

# Deep tree-ensembles for multi-output prediction

Felipe Kenji Nakano<sup>a,b,\*</sup>, Konstantinos Pliakos<sup>a,b</sup>, Celine Vens<sup>a,b</sup>

<sup>a</sup>*KU Leuven, Campus KULAK, Department of Public Health and Primary Care, Etienne Sabbelaan 53, 8500 Kortrijk, Belgium*

<sup>b</sup>*Itec, imec research group at KU Leuven, Etienne Sabbelaan 53, 8500 Kortrijk, Belgium*

---

## Abstract

Recently, deep neural networks have expanded the state-of-art in various scientific fields and provided solutions to long standing problems across multiple application domains. Nevertheless, they also suffer from weaknesses since their optimal performance depends on massive amounts of training data and the tuning of an extended number of parameters. As a countermeasure, some deep-forest methods have been recently proposed, as efficient and low-scale solutions. Despite that, these approaches simply employ label classification probabilities as induced features and primarily focus on traditional classification and regression tasks, leaving multi-output prediction under-explored. Moreover, recent work has demonstrated that tree-embeddings are highly representative, especially in structured output prediction. In this direction, we propose a novel deep tree-ensemble (DTE) model, where every layer enriches the original feature set with a representation learning component based on tree-embeddings. In this paper, we specifically focus on two structured output prediction tasks, namely multi-label classification and multi-target regression. We conducted experiments using multiple benchmark datasets and the obtained results confirm that our method provides superior results to state-of-the-art methods in both tasks.

*Keywords:* Ensemble Learning, Deep-Forest, Multi-output prediction, Multi-target regression, Multi-label classification

---

\*Corresponding author

*Email addresses:* `felipekenji.nakano@kuleuven.be` (Felipe Kenji Nakano), `konstantinos.pliakos@kuleuven.be` (Konstantinos Pliakos), `celine.vens@kuleuven.be` (Celine Vens)

---

## 1. Introduction

Modern technological advances have opened new horizons in many areas of science. Along with new opportunities and possibilities that were deemed unfeasible in the past, this ongoing progress has given birth to new challenges for the scientific community, stimulating new lines of research [1]. When it comes to the field of machine learning, traditional approaches need to adapt and tackle the new challenges that have emerged (e.g. increased data volume and complexity, need for scalability, etc.) [2].

In this direction, deep learning has arisen as a cutting edge methodology. In recent years, deep learning has expanded the state-of-the-art in diverse fields of knowledge and application domains, such as medicine [3], genomics [4], and self driving cars [5].

Apart from its blatant success, deep neural networks may also present some drawbacks. In most of the scenarios, the training consists of tuning a massive number of parameters which demands considerably large training datasets. This often leads to an extreme computational complexity. In addition, deep neural networks often fail to meet high prediction performance when it comes to small scale datasets, which are still mainstream in many application domains. For instance, studies that follow participants over time (e.g. patients in the context of clinical studies) typically have a limited number (often in the range of a few hundreds) of subjects.

In typical supervised learning, a prediction function is formed on a training set of instances in order to predict a target value [6]. Each instance is represented by a feature vector and associated with a target, which can be categorical (classification) or numerical (regression). Single target prediction tasks adopt the assumption that an instance is associated with a single output variable. However, this assumption does not hold in many real-world problems, as instances may belong to more than one class at the same time. For example, an article can be associated with multiple topics simultaneously. Multi-output prediction

models that learn to predict multiple output variables at the same time are of high importance in this setting. Multi-output prediction, which is also denoted as multi-target prediction [7], is a generalization of multi-target regression and multi-target classification. Multi-label classification can be considered as a special case of multi-target regression (or classification) with only two numerical values (classes) for each target (0 or 1) [8]. Nowadays, due to the increasing volume and complexity of the stored data, multi-output prediction methods draw growing attention. Such approaches boost prediction performance, exploiting the structure of the target space, while being computationally efficient.

In this direction, here we focus on multi-output prediction tasks. In particular, we tackle the bottlenecks from deep neural networks by proposing a novel deep tree-ensemble method, where tree-ensembles are built in multiple layers following a deep learning architecture. Our work is motivated by the deep forest approach proposed by Zhou & Feng [9], where the model also consists of several layers. In each layer of [9], a random forest [10] model outputs a class probability vector which is then used to extend the original input feature set. Different from [9], in our approach we include a tree-based representation learning step in every layer extending the original input space with low-dimensional tree-embeddings. More specifically, instead of simply using the predictions of the previous layer as extra attributes for the next one, we learn a high dimensional representation based on the decision paths of the trees in the tree-ensemble collection. Next, we project these high dimensional and sparse vectors to a low-dimensional space, yielding compact tree-embeddings which we use to extend the input feature set.

Our main hypothesis is that adding a representation learning component to the deep forest methodology will boost the predictive performance of the method while keeping the computational complexity low. The nodes of each tree of the ensemble are considered clusters that comprise the instances that are associated to them by the recursive partitioning procedure. The ensemble is converted into a binary vector where each item is associated with a cluster node. If a training instance has passed through this node the corresponding item receives the value of 1, otherwise 0. Tree-ensemble models are trained

over the original feature set in a supervised learning way and therefore the constructed features incorporate label information. It has been shown that such representations benefit from the learning mechanism of tree-ensembles and are apt to boost the performance of subsequently applied prediction models [11]. Furthermore, we filter and apply weights to this binary representation, based on the number of instances that are contained in a cluster (tree node). The decision paths of the individual trees of the ensemble are transformed to high dimensional and label-aware representations. Finally, we transfer these high dimensional feature vectors to a low dimensional space, removing noise and making the resulting compact tree-embeddings computationally efficient. Low-dimensional tree-embeddings were proved highly representative and effective in multi-output prediction tasks [12].

Driven by the great importance and potential of multi-output prediction as well as by the relative lack of deep-forest methods for multi-output tasks, we deploy our proposed deep tree-ensemble (DTE) method to specifically address the tasks of multi-target regression and multi-label classification. Experiments on benchmark datasets from both domains demonstrate the superiority of our method over state-of-the-art ones.

The main contributions of this work are summarized below:

- A new state-of-the-art deep tree-ensemble method for multi-target regression and multi-label classification;
- The first study that incorporates representation learning and specifically low-dimensional tree-embeddings into a deep-forest architecture;
- An effective stopping (pruning) criterion to determine the optimal number of layers as well as mechanisms to surpass overfitting;
- An extensive evaluation on 41 datasets, comparing our approach to state-of-the-art methods;

The remainder of this paper is organized as follows: Section 2 brings a brief overview on multi-target regression, multi-label classification and deep-forest

based methods; Section 3 introduces our proposed method; Section 4 presents our experimental setup, and Section 5 contains our results and discussion. Finally, Section 6 contains our final conclusions and future work directions;

## 2. Related work

In this section, we present a brief overview on relevant literature in multi-label classification and multi-target regression, covering the most popular methods in these fields. For a complete literature survey, we refer the reader to [8, 13, 14, 15].

Despite the great variety of methods, they can be roughly categorized into two approaches: global and local. The global approach corresponds to models that learn to predict all targets at once. Typically, a traditional algorithm, such as decision trees [16], is adapted to handle the constraints of the problem. Different from that, the local approach builds several models which learn to predict different subsets of outputs [17, 18]. In this case, the predictions of all models are combined to provide the final prediction.

### 2.1. Multi-target regression

Given a set of instances  $X$  and a continuous output set  $Y \in R^m$ , multi-target regression is defined as the task of learning a function  $f$  which, for every unseen instance  $x \in X$ , outputs  $y \in Y$ , with  $y$  being the target set associated with  $x$  [13].

Kocev *et al.* [19, 20] proposed global ensembles of predictive clustering trees for structured output tasks. Their experiments showed random forests [10] as the ensemble method that generally yields superior results. Using deep neural networks, Du and Xu [21] proposed a hierarchical neural networks method which employs a divide-and-conquer (local) approach. More specifically, it decomposes the original multi-target problem into several local sub-tasks, and combines their solutions. Also using deep neural networks, Zhen *et al.* [22] proposed a general model that employs a non-linear layer and a linear low-rank layer to perform

direct face alignment via multi-target regression. ERC [23] is an ensemble local approach where multiple regressors are chained in a random order. In this method, a separate regressor is built for each target and its predictions are used as extra attributes for the next regressor. Similarly to ERC, DSTARS [17] also chains multiple stacked regressors. Its main difference consists of using the feature importance to identify target correlations, in the sense that, targets with positive feature importance values are used as extra attributes for the next regressors.

## 2.2. Multi-label classification

Given a set of instances  $X$  and a discrete output set  $Y$ , multi-label classification is defined as the task of learning a function  $f$  which, for every unseen instance  $x \in X$ , outputs  $y \in Y$ , with  $y$  being the label set associated with  $x$  [13].

Based on lazy learning, ML-KNN [24] is an extended version of the traditional k-nearest neighbours, which in addition to finding the closest instances from the training dataset, also employs the concept of maximum a posteriori to predict the labels of new instances. Proposed by Tsoumakas *et al.* [18], Rakel is a local ensemble method which aims to reduce the size of the label set by separating it into smaller random subsets. Next, one model is built per set and the predictions are combined. As an extension of Rakel, Wang *et al.* [25] proposed to select the subsets of labels based on an active learning algorithm, rather than at random. Their results show that the performance can be improved by selecting the best group of labels per subset. Adopting the notion of missing labels, Ma and Chen [26] proposed a method that imposes low-rank matrix structures in both global and local manner. More precisely, it enforces local low-rank on predictions with the same labels, whereas the rank is expanded for predictions with different labels.

In the domain of image multi-label classification, several studies have successfully employed convolutional neural networks associated with a label dependence extraction method. For instance, as a recent influential approach, Chen

*et al.* [27] proposed a model that exploits graph convolutional networks to build stacked label dependence graphs. Furthermore, Chen *et al.* [28] proposed a semantic decoupling method which extracts label specific features and their dependence at the same time. Lastly, You *et al.* [29] proposed an adjacency-based method to extract dependence associated to a cross-modality attention mechanism. Despite their wide success, these methods are specifically designed for computer vision tasks, thus we do not consider them further in this work.

### 2.3. Deep-Forest based methods

Similarly to the single output deep forest approach [9], its multi-label version (MLDF) proposed in [30] employs multiple layers of random forests and extra trees [31]. However, prediction probabilities are evaluated using measured-based confidence values. More specifically, if the prediction probabilities from the current layer have a better confidence than the ones from the previous layer, they are updated. As an extension of the multi-label deep-forest, Wang *et al.* [32] addressed weak-label learning by using a label complement procedure. At each layer, the label set of the training dataset is complemented using an inner cross-validation scheme, that is, if the label is predicted as positive its label is changed in the training dataset.

Gao *et al.* [33] proposed a deep-forest variant that includes the oversampling method SMOTE [34] to address imbalanced tasks. All instances that belong to the minority class and were misclassified in the previous layer are used as input to SMOTE, the new generated instances are concatenated with the dataset and a new layer is built.

In the oncology field, Su *et al.* [35] proposed a deep forest model to predict anti-cancer drug response. The original deep forest was extended by allowing the usage of two feature sets and including a feature optimization method after each layer. Ma *et al.* [36] proposed a cost-sensitive version for price prediction. First, the prices are discretized using K-means, next the authors employed a deep-forest method which penalizes wrong predictions according to how inaccurate they are. In the context of hashing for image retrieval, Zhou *et al.* [37]

demonstrated that the original deep-forest is capable of learning a hashing function that outperforms deep learning hashing methods.

### 3. The proposed method

In this section, we thoroughly present our deep tree-ensemble (DTE) method. At first, we introduce the general framework, focusing on its main components. Next, we present a detailed description of the proposed representation learning components. Finally, we introduce the stopping (pruning) criterion that was employed as well as specific characteristics of our architecture.

#### 3.1. The general framework

Inspired by recent advances on deep-forest models, our proposed method also employs a cascade of layers of tree ensembles. The main difference relies on the introduced representation learning component at each layer. More specifically, our proposed model extends the original framework by including low-dimensional tree-embeddings, a promising representation learning method, which is further explained in the next subsection.

It is widely known that diversity among the underlying models is a crucial component in any ensemble method [38]. Hence, each layer of our model consists of a set of different tree-ensemble models. Let  $B = \{\beta_1, \dots, \beta_i, \dots, \beta_{|B|}\}$  be the different models. Each  $\beta_i$  corresponds to a tree-ensemble model such as random forest (RF) [10], extremely randomized trees also known as extra-trees (ET) [31], gradient boosting, etc. The number  $|B|$  also varies across different applications. Here, we build our deep tree-ensemble (DTE) method using one RF and one ET model ( $|B| = 2$ ) per layer. We employ these two models as they are well established and also powerful. Moreover, we follow the setting of predictive clustering trees [16] due to its natural clustering and variance reduction characteristics.

Our proposed framework is illustrated in Figure 1. Initially, the original feature representation ( $X$ ) is used as input. RF and ET are employed to create



their respective tree-embeddings  $F$ . Let  $F_{RF}$  and  $F_{ET}$  be the RF and ET based tree-embeddings, respectively.

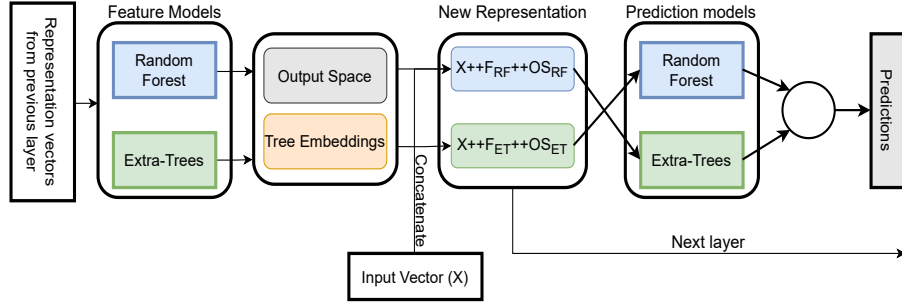


Figure 1: An example of a layer of our proposed model. The box *Output Space* (OS) refers to the prediction features as explained in the text. The arrow *Next layer* links the *New Representations* of this layer to the input box of the next layer, which has the same structure as the illustrated one.

Next, the original feature representation ( $X$ ) is concatenated with the generated representations ( $F$ ), and crossed over to be used as input for the models that perform predictions. That is, the  $F_{RF}$  is generated in the current layer and given as input to ET, and vice-versa. Our main motivation behind this is that each  $\beta_i$  benefits from learning from representations that were yielded by  $\beta_j$ , where  $i \neq j$ . This way, we increase the diversity between the predictors. Each model  $\beta_i$  is always built using representations that stem from different models. Besides boosting diversity, we assume that this step contributes to avoid overfitting.

Thus, for all layers  $L$  with  $L \geq 1$ , a RF model is trained using  $X \# F_{ET}$  features, and an ET model with  $X \# F_{RF}$ , where  $\#$  corresponds to vector concatenation. This can be generalized to  $|B|$  tree-ensemble models;  $\forall \beta_i \in B$  the generated feature set is  $X \# F_{\beta_1} \# F_{\beta_j} \# \dots \# F_{\beta_{|B|}}$ , where  $j \neq i$ . These models are chained to the DTE, building a new layer. This procedure is repeated until the stopping criterion is triggered.

When it comes to final predictions, our method employs all the models  $\beta_i \in B$  of its last layer (w.r.t the stopping criterion). More specifically, it averages

the output of all trees in all tree-ensemble models  $\beta_i$  to perform predictions.

It is possible to implement different feature representation combinations. For instance, output features as used in [9, 30] could be included in our architecture, allowing us to propose four variants of our method:

- X\_TE: Original feature representation and `tree-embeddings` ( $X \# F$ );
- X\_OS: Original features and prediction (i.e., `output space`) features. This variant is identical to the original deep-forest [9], except for the feature cross-over.
- X\_OS\_TE: Original features, tree-embeddings and prediction features;
- TE: Only tree-embeddings. A baseline variant which is used to evaluate the importance of concatenating the original feature representation with the generated one.

### 3.2. Low Dimensional Tree-embeddings

Originally proposed for biomedical data, tree-embeddings provide more informative and variant representations by exploiting the inner properties of decision trees, in particular, the clustering performed at each split is converted to a new feature representation [12].

Given a trained ensemble of decision trees, tree-embeddings are generated using the following procedure. Initially, all the trees are converted to a binary vector where each position is associated with a node. These vectors are concatenated, resulting in  $C = \{c_1, c_2, \dots, c_{|C|}\}$ ,  $|C|$  being the total number of nodes in the ensemble. By treating  $c \in C$  as a feature, we create a new representation  $F \in R^{N \times |C|}$  where  $N$  stands for the number of instances in the dataset. The values in  $F$  are instantiated according to how instances traverse the trees. Namely,  $F_{ij}$  equals to 1 if such instance  $i$  belongs to node (cluster)  $j$ , otherwise its value is set to 0, as exemplified in Figure 2. In more detail, in 2, we show a hypothetical instance being classified into leaf nodes 4 and 10. Next,

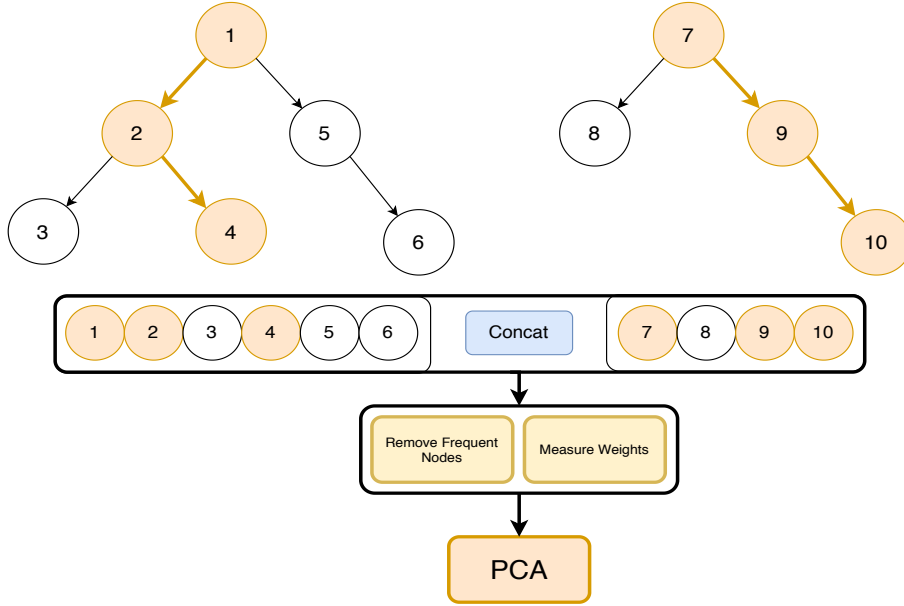


Figure 2: Hypothetical example being classified by leaf nodes 4 and 10.

we show the binary vector representation for the instance, where colored circles correspond to vector positions set to 1.

Secondly, we discard the nodes in  $F$  (columns  $F_{.j}$ ) which are present in more than  $p\%$  of the instances, removing non-meaningful nodes, such as the ones located close to the root and the root itself. Next, the remaining nodes are weighted according to Equation 1 where nodes that contain many instances are credited a lower weight, and the opposite is valid for more representative nodes. In Equation 1,  $w_j$  stands for the weight assigned to node  $j$ ,  $|c_j|$  is the number of instances that traverse  $j$  and  $e$  is a small positive value for numerical stability.

$$w_j = \frac{1}{\log(|c_j|) + e} \quad (1)$$

Lastly, the filtered and weighted vectors are transferred from their high dimensional space to a lower dimensional one. Pliakos and Vens [12] evaluated several standard methods in this context, such as: PCA, kernel PCA, independent component analysis and truncated singular value decomposition. The

results obtained from their experiments have showed that PCA is often associated with superior results. Hence, here we also employ PCA to obtain the final tree-embeddings. Recall that, the input to PCA should be centered before being processed.

This way, through this representation learning component, we yield a compact but yet informative and target-aware representation. By transferring the tree-path features to a low dimensional space we make our approach computationally more efficient removing as well any existing noise. Moreover, by tuning the number of kept components (final dimensionality of the  $F$  space) we implicitly assign a latent weight to the added feature set, in relation to the original space  $X$ . Thus, we can control the number of added features. Recall that, this procedure is performed separately for each ensemble model ( $\beta_i$ ) due to the embeddings cross-over mechanism, as shown in Figure 1.

Furthermore, considering that tree-embeddings are very representative and label-aware, overfitting can be an issue. Thus, we employ the following procedure:

*Overfitting prevention: The ensembles in the model itself are not used to generate tree-embedding features. A separate RF and an ET are built using a random sample of the training dataset, these models are used exclusively for the generation of tree-embeddings, and, thus, are not used to perform predictions.*

### 3.3. Output space features

Recently, features based on the output of models have been vastly exploited in structured output prediction [17, 30, 39]. More precisely, given a chain of models, the outputs (i.e., predictions) of the previous models are used as extra (induced) features for the next one.

These features are known to enrich the representational power by exploiting label correlations among the output set [17]. Furthermore, it also allows the models to learn from each other’s mistakes [40], enhancing the overall performance. Likewise, output features are commonly exploited in deep-forest models [9, 30].

Thus, we also incorporate them into our model. Moreover, analogously to tree-embeddings, we also employ an overfitting prevention measure to obtain output features.

*Overfitting prevention: The ensembles in the model itself are not used to generate output features. Similarly to [9, 30], an inner  $K$ -fold cross validation is employed, resulting in  $2K$  separate models (RF + ET). Their averaged outputs, either prediction probabilities or target values, are the output features.*

Using hard predictions (0 or 1) as extra attributes, especially for multi-label classification, could also be implemented, as proposed by Read [39]. Nonetheless, there has been an increasing tendency to use the raw output of the models due to superior results [41].

#### 3.4. Pruning criterion

As for finding the optimal number of layers, different options can be explored. At a first glance, the training could be stopped immediately after any deterioration in the performance. By being more flexible, it is also possible to establish a tolerance level where the training is interrupted if the performance does not improve after a certain number of layers, as proposed by Yang [30]. Both measures may lead to sub-optimal solutions due to premature stops. As a countermeasure, we propose a pruning criterion.

Given a pre-defined number of iterations, we select the layer with the best performance on the training dataset. Thus, all the ensembles trained after the layer with the best performance are pruned. The final model consists of all tree-ensembles up to the layer with the best performance, and the output of this layer yields the final predictions. This pruning criterion is exemplified in Figure 3.

#### 3.5. Computational complexity

First, we evaluate the complexity of obtaining tree-embeddings and output space features. Next, we use them to define the complexity of our proposed

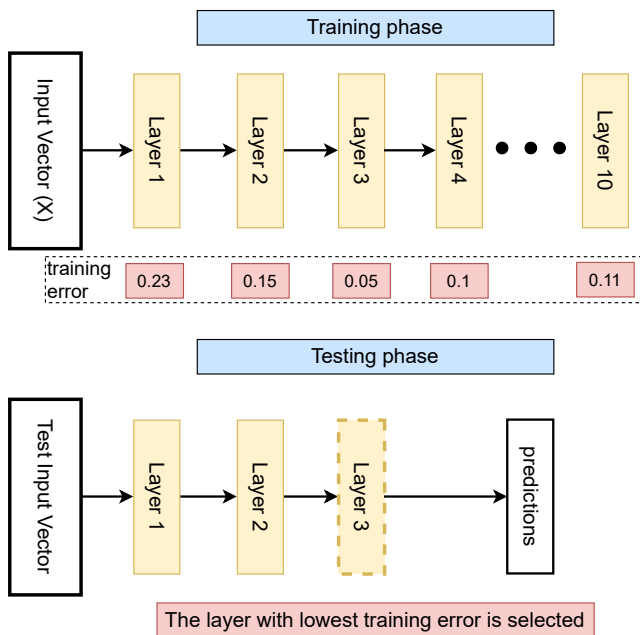


Figure 3: Proposed stopping (pruning) criterion. Select the layer with the lowest training error given a pre-defined number of layers

method.

As mentioned in the previous section, obtaining tree-embedding features relies mainly on the following steps:

1. Building the tree ensembles at each layer;
2. Traverse the trees to obtain decision paths;
3. Apply PCA to obtain the tree-embedding features;

The complexity of measuring tree-embedding features is defined as follows. 1) The cost for building a tree ensemble may vary according to the ensemble itself, hence we generalize this as  $Comp_{Model}$ ; 2) Obtaining decision paths has the same complexity as performing predictions which is  $\mathcal{O}(n_{trees} \times N \log(N))$ , where  $N$  is the number of instances; Finally, 3) PCA is applied to the matrix from the previous step, resulting in  $\mathcal{O}(C^2N + C^3)$ ,  $C$  being the number of nodes in the ensemble. Such complexity is associated to PCA because it requires measuring

the covariance matrix ( $C^2N$ ) and its eigenvalue decomposition ( $C^3$ ) [42]. Thus, the total complexity,  $Comp_{TE}$ , is  $\mathcal{O}(Comp_{Model} + n_{trees} \times N \log(N) + C^2N + C^3)$ .

When it comes to output space features, its complexity depends on:

1. Building the tree ensembles at each layer;
2. Traverse the trees to obtain predictions;

As previously discussed, 1) the complexity of the first step is  $Comp_{Model}$ . 2) Likewise, the complexity of performing predictions is defined as  $\mathcal{O}(n_{trees} \times N \log(N))$ . Hence, its total complexity,  $Comp_{OS}$ , is  $\mathcal{O}(Comp_{Model} + (n_{trees} \times N \log(N)))$ .

The complexity of each variant of our proposed method is defined by combining  $Comp_{OS}$  and  $Comp_{TE}$ . At each layer, it is necessary to compute the features and build the ensembles. Furthermore, a different set of extra features is built for each ensemble. Thus, the complexity of obtaining the features multiplied by the number of models. Formally, we state the complexity of each of the variants as follows. Here,  $N_{models}$  stands for the number of ensembles in each layer.

- $Comp_{X_{TE}} = \mathcal{O}(N_{models} \times Comp_{Model} + N_{models} \times Comp_{TE})$ ;
- $Comp_{X_{OS}} = \mathcal{O}(N_{models} \times Comp_{Model} + N_{models} \times Comp_{OS})$ ;
- $Comp_{X_{TE_{OS}}} = \mathcal{O}(N_{models} \times Comp_{Model} + N_{models} \times (Comp_{TE} + Comp_{OS}))$ ;
- $Comp_{TE} = \mathcal{O}(N_{models} \times Comp_{Model} + N_{models} \times Comp_{TE})$  ;

## 4. Experimental Setup

### 4.1. Multi-target regression

In Section 4.1.1, more specifically, in Table 1, we present the multi-target regression datasets used in this work.

Datasets	#Instances	#Features	#Targets	#Domain
atp1d	337	370	6	Air ticket price over a day
atp7d	296	370	6	Air ticket price over 7 days
oes97	334	263	16	Number of employees per job measured in 1997
oes10	403	298	16	Number of employees per job measured in 2010
rf1	9125	64	8	Water flow in the Mississippi River
rf2	9125	576	8	Water flow in the Mississippi River including forecasting features
scm1d	9803	280	16	Product prices after 1 day
scm20d	8966	61	16	Product prices after 20 days
edm	154	16	2	Electrical discharge of machinery
sf1	323	10	3	Occurrence of solar flares over a period of 24h
sf2	1066	9	3	Occurrence of solar flares over a period of 24h(v2)
jura	359	15	3	Heavy metal concentration
wq	1060	16	14	Fauna and flora in Slovenian rivers
enb	768	8	2	Heating and cooling of efficient buildings
slump	103	7	3	Concrete quality
andro	49	30	6	Water quality
osales	639	401	12	Online sales of a product over its first month
scpf	1137	23	3	Number of views, clicks and comments
musicOrigin1	1059	68	2	Geographical origin of music
musicOrigin2	1059	116	2	Geographical origin of music
friedman	500	25	6	Data sampled from an uniform distribution
mp5spec	80	700	3	Measurements of corn using a mp5 spectrometer
mp6spec	80	700	3	Measurements of corn using a mp6 spectrometer
polymer	41	10	4	Polymer processing plants

Table 1: Multi-target regression datasets characteristics

#### 4.1.1. Datasets

We have collected a total of 24 benchmark datasets from public repositories<sup>1</sup> and split them using 10 fold cross-validation. These datasets are available at<sup>3</sup>. In terms of characteristics, they range from a very small number of features and targets (enb) to 700 features (mp5spec and mp6spec) and 16 targets (oes97, oes10, scm1d and scm20d).

#### 4.1.2. Comparison methods

- DSTARS: A recently proposed approach which combines multiple stacked models into a deep structure [17], available at <sup>4</sup>.
- ERC: An ensemble of regressor chains where one regressor is trained for

<sup>1</sup><http://mulan.sourceforge.net/datasets-mtr.html>

<sup>2</sup><http://people.vcu.edu/~acano/MTR-SVRCC/>

<sup>3</sup>[https://itec.kuleuven-kulak.be/?page\\_id=8493](https://itec.kuleuven-kulak.be/?page_id=8493)

<sup>4</sup><https://github.com/smastelini/mtr-toolkit>



each target, and outputs are used as extra attributes for the next regressors [23];

- MLP: A comparative multi-layer perceptron neural network regressor <sup>5</sup>;
- RF + ET: A shallow version of our proposed method which consists of a single layer with one RF and one ET; <sup>6</sup>;
- Our proposed method: We compare all the four variants of our method: TE, X\_TE, X\_OS, X\_TE\_OS <sup>7</sup>;

For RF + ET, we have fixed the number of trees to 150. For the MLP, we have also employed an inner 5-fold cross validation to optimize the number of layers and neurons. More specifically, we have considered the following values: number of neurons {50,100} and number of layers {1,2}. The activation function and solver were fixed to ReLu and ADAM, respectively.

In its publicly available implementation, DSTARS normalizes all features and targets in the dataset. Consequently, its predictions are also normalized, making them incomparable with the other methods. To avoid this situation, we have removed any normalization procedure involved in DSTARS. We have also set the number of trees to 150.

As for the ERCs [23], we have employed regression trees as base models, and 10 randomly ordered chains. In this case, predictions are performed by averaging the output of each model.

For the number of PCA components in our method, we have optimized the number of components considering a percentage of the total number of decision path features {1 component, 1% , 5%, 20%, 40%, 60%, 80%, 95%} multiplied by  $\min(N, |c|)$ ,  $N$  being the number of instances in the dataset and  $|c|$  the number of nodes in the ensemble. The number of trees was fixed to 150 and output features are generated using 5-fold cross validation. As for preventing the

---

<sup>5</sup><https://scikit-learn.org/stable/>

<sup>6</sup><https://scikit-learn.org/stable/>

<sup>7</sup>[https://itec.kuleuven-kulak.be/?page\\_id=8493](https://itec.kuleuven-kulak.be/?page_id=8493)

overfitting which may be associated with tree-embeddings, we have employed random samples which contain 50% of the training dataset. Finally, we run our model with a maximum of 10 layers.

#### 4.1.3. Evaluation Measures

Similar to related literature, we are employing the aRRMSE (Average Relative Root Mean Squared Error). RRMSE is the square root of the mean squared error (RMSE) of the model divided by the RMSE of a default model that predicts the average target value in the training set. Likewise, aRRMSE is computed by taking the average RRMSE of all targets. Formally, Equation 2 presents aRRMSE where  $\hat{y}$  stands for predicted values,  $\bar{y}$  represents the average value,  $M$  is total the number of targets and  $N$  is the number of testing instances.

$$aRRMSE = \frac{1}{M} \sum_{t=1}^M \sqrt{\frac{\sum_{i=1}^N (y_t^i - \hat{y}_t^i)^2}{\sum_{i=1}^N (y_t^i - \bar{y}_t)^2}} \quad (2)$$

## 4.2. Multi-label classification

In this section, we provide details on the multi-label experimental setup.

### 4.2.1. Datasets

In Table 2, we present statistical information about the dataset used.

As can be seen, the datasets belong to different domains of applications, including sound recognition (birds, emotions and CAL500), image annotation (scene), text classification (enron, LLOG) and medical and biomedical applications such as CHD\_49, gram\_negative, gram\_positive, gram\_positivePseAAC, plantGO, virusGO, virusPseAAC, medical, genbase and yeast.

There is also a wide variation in the number of features and labels. Certain datasets present fewer features such as flags and CHD\_49 with 19 and 49, respectively, whereas plantGO presents 1763 features. The number of labels also varies significantly, ranging from 4 up to 174.

Datasets	#Instances	#Features	#Labels	# Cardinality	#Application
birds	645	260	19	1.01	Species recognition based on chirping
CAL500	502	68	174	26.04	Instrument, genre and voice identification
CHD_49	555	49	6	2.58	Coronary heart disease prediction
emotions	593	72	6	1.87	Emotion recognition from songs
enron	1702	1001	53	3.38	Email classification
flags	194	19	7	3.392	Country flag prediction
genbase	662	1186	27	1.25	Protein function prediction
Gram_negative	1392	1717	8	1.04	Gram negative bacterial species sub-cellular location
Gram_positive	519	912	4	1.04	Gram positive bacterial species sub-cellular location
Gram_positivePseAAC	519	912	4	1.04	Gram positive bacterial species sub-cellular location
LLOG	1460	1004	75	1.18	Text classification
medical	978	1449	45	1.25	Diagnosis based on clinical text
PlantGO	978	3091	12	1.07	Plant species sub-cellular location
scene	2407	294	6	1.07	Image annotation
VirusGO	207	749	6	1.21	Virus species sub-cellular location
VirusPseAAC	207	440	6	1.21	Virus species sub-cellular location
yeast	2417	103	14	4.24	Protein function prediction

Table 2: Multi-label classification datasets characteristics

We have collected them from public repositories<sup>8 9</sup>, concatenated the train and test subsets, and performed 10 fold cross-validation. These datasets are available at<sup>10</sup>.

#### 4.2.2. Comparison methods

We have compared the following multi-label methods:

- Rakel: An ensemble approach which builds a base classifier per subset of labels [18]<sup>11</sup>.
- MLKNN: A multi-label version of the baseline K-Nearest-Neighbours classifier [24]<sup>12</sup>;
- MLDF: A recently proposed multi-label deep-forest version proposed in [30], available at<sup>13</sup>;

<sup>8</sup><http://mulan.sourceforge.net/datasets-mlc.html>

<sup>9</sup><http://www.uco.es/kdis/mlresources/>

<sup>10</sup>[https://itec.kuleuven-kulak.be/?page\\_id=8493](https://itec.kuleuven-kulak.be/?page_id=8493)

<sup>11</sup><http://scikit.ml/>

<sup>12</sup><http://scikit.ml/>

<sup>13</sup>[http://www.lamda.nju.edu.cn/code\\_MLDF.ashx](http://www.lamda.nju.edu.cn/code_MLDF.ashx)

- MLP: A comparative multi-layer perceptron neural network multi-label classifier <sup>14</sup>;
- ECC: An ensemble of classifier chains where one classifier is trained for each class, and outputs are used as extra attributes for the next classifiers [39];
- RF + ET: A shallow version of our proposed method which consists of a single layer with one RF and one ET; <sup>15</sup>;
- Our proposed method: We compare all the four variants of our method: TE, X\_TE, X\_OS, X\_OS\_TE <sup>16</sup>;

For Raket, we have employed SVMs with linear kernel as base classifiers, random overlapping subsets of size 3 and  $2 \times M$  classifiers, M being the number of possible labels per dataset. For MLKNN, we have optimized K considering the values {3,5,7} and S {0.5, 0.7, 1.0} using a 5-fold inner cross validation. Similarly to our MTR experiments, we have employed 10 chains following a random order of targets with decision trees as regressors in the ECCs [43, 44].

For our main competitor, MLDF, we have employed the same parameters from their original manuscript. That is, one random forest and one extra tree model, 40 trees at the first layer, increasing 20 trees per layer up to 100, 5-fold inner cross-validation to generate the output features. As the stopping criteria, the training is interrupted after the addition of 3 layers without improvement. Since this method is measure dependent, we have compared it using the measures with which it is defined.

The parameters related to the MLP, RF + ET and the variants of our proposed method were optimized in the same manner as in the multi-target regression experiments. In these experiments, however, we employ multiple evaluation measures. We have picked one measure, namely the microAUC,

---

<sup>14</sup><https://scikit-learn.org/stable/>

<sup>15</sup><https://scikit-learn.org/stable/>

<sup>16</sup>[https://itec.kuleuven-kulak.be/?page\\_id=8493](https://itec.kuleuven-kulak.be/?page_id=8493)

for optimization, meaning that the same model is compared for all evaluation measures.

As demonstrated in [27, 28, 29], deep convolutional neural networks have been recently applied for image multi-label classification. Despite being very powerful, these approaches are designed for computer vision applications, meaning that raw images are necessary as input. Consequently, they are regarded as out-of-scope.

#### 4.2.3. Evaluation Measures

In Table 3, we present the symbols used to describe our multi-label evaluation measures.

$Y$	True label set
$X$	Instance set
$Z$	Predicted labels
TP	True positive
FP	False positive
FN	False negative
$ S $	Size of set S
$N$	Number of instances ( $ X $ )
$M$	Number of labels in Y ( $ Y $ )
$\theta$	Ranking function
$I$	Indicator function, returns 1 if its argument evaluates to True
$\hat{Y}_i$	Complement of Y for an instance $i$
$\oplus$	Boolean XOR
$T$	Threshold values
$t$	Iterator for threshold values

Table 3: Table with symbols used to describe the multi-label evaluation measures.

The hamming loss (Equation 3) measures the average number of labels

wrongly predicted. In this case, we have applied a threshold of 0.5 to obtain  $Z$ .

$$HammingLoss = \sum_i^N Z_i \oplus Y_i \quad (3)$$

Based on the ranking of the predicted labels, one error, Equation 4, measures the average number of times that the label with highest prediction probability (lowest rank) is not a true label.

$$OneError = \frac{1}{N} \sum_i^N I(\operatorname{argmin}_{y \in Y} \theta(y) \notin Y_i) \quad (4)$$

Also based on ranking, the ranking error (Equation 5) measures the average proportion of labels that were ranked incorrectly.

$$RankingError = \frac{1}{N} \frac{1}{|Y_i| |\hat{Y}_i|} |(y_1, y_2) : \theta_i(y_1) > \theta_i(y_2), (y_1, y_2) \in Y_i \times \hat{Y}_i| \quad (5)$$

We are also employing the average precision score [45]. Due to possible imbalance in the label set, we are employing its micro variant described at Equations 6, 7 and 8.

$$MicroAverageRecall = \sum_j^M \frac{TP_j}{TP_j + FN_j} \quad (6)$$

$$MicroAveragePrecision = \sum_j^M \frac{TP_j}{TP_j + FP_j} \quad (7)$$

$$MicroAveragePrecisionScore = \sum_t^T \frac{(MicroAverageRecall_t - MicroAverageRecall_{t-1})}{MicroAveragePrecision_t} \quad (8)$$

The microAUC is obtained by measuring the area under the curve generated by using the recall (Equation 9) in the x-axis and the false positive rate (FPR, Equation 10) in the y-axis.

$$Recall = \frac{TP}{TP + FN} \quad (9)$$

$$FPR = \frac{FP}{FP + TP} \quad (10)$$

## 5. Results and Discussion

Initially, we present the results on multi-target regression, followed by multi-label classification. More results concerning other stopping criteria are available in the Appendix.

### 5.1. Multi-target regression

Datasets ↓	RF+ET	MLP	ERC	X_TE	TE	X_OS	X_OS_TE	DSTARS
atp1d	0.488	<b>0.463</b>	0.526	0.486 <sup>5.7</sup>	0.509 <sup>2.0</sup>	0.481 <sup>5.6</sup>	0.479 <sup>7.2</sup>	0.479
atp7d	0.814	0.957	1.09	0.814 <sup>5.1</sup>	<b>0.785</b> <sup>1.5</sup>	0.815 <sup>5.8</sup>	0.804 <sup>4.7</sup>	1.143
oes97	0.563	0.611	0.648	0.566 <sup>5.6</sup>	0.625 <sup>2.8</sup>	0.575 <sup>5.3</sup>	0.575 <sup>5.4</sup>	<b>0.508</b>
oes10	0.441	0.512	0.494	0.44 <sup>4.7</sup>	0.461 <sup>2.7</sup>	0.447 <sup>6.2</sup>	0.444 <sup>4.8</sup>	<b>0.399</b>
edm	0.867	1.021	1.096	0.838 <sup>7.6</sup>	0.817 <sup>6.9</sup>	0.827 <sup>8.8</sup>	0.819 <sup>8.9</sup>	<b>0.799</b>
sf1	<b>0.981</b>	1.5	1.299	0.997 <sup>7.2</sup>	1.006 <sup>7.3</sup>	1.071 <sup>8.6</sup>	1.055 <sup>9.0</sup>	1.111
sf2	1.402	2.766	2.824	1.423 <sup>6.0</sup>	1.385 <sup>5.2</sup>	1.49 <sup>7.8</sup>	<b>1.356</b> <sup>8.0</sup>	1.615
jura	0.664	0.733	0.685	0.633 <sup>4.8</sup>	0.689 <sup>2.0</sup>	0.615 <sup>7.0</sup>	0.637 <sup>7.7</sup>	<b>0.56</b>
wq	<b>0.94</b>	1.165	0.969	0.942 <sup>8.3</sup>	0.954 <sup>8.2</sup>	1.001 <sup>9.0</sup>	0.996 <sup>8.6</sup>	0.96
enb	0.178	0.153	0.371	0.165 <sup>6.4</sup>	0.191 <sup>4.3</sup>	<b>0.152</b> <sup>8.8</sup>	0.181 <sup>8.7</sup>	2.529
slump	0.851	1	<b>0.792</b>	0.857 <sup>5.8</sup>	0.881 <sup>3.2</sup>	0.895 <sup>9.0</sup>	0.89 <sup>8.7</sup>	0.814
andro	0.939	0.909	1.829	0.931 <sup>3.8</sup>	1.035 <sup>3.5</sup>	0.939 <sup>5.9</sup>	0.933 <sup>7.3</sup>	<b>0.79</b>
osales	0.776	0.739	1.158	0.759 <sup>6.1</sup>	0.801 <sup>4.2</sup>	0.716 <sup>8.2</sup>	0.712 <sup>8.2</sup>	<b>0.653</b>
scpf	0.819	1.17	1.038	0.807 <sup>4.3</sup>	<b>0.804</b> <sup>2.9</sup>	0.912 <sup>8.8</sup>	0.873 <sup>7.6</sup>	0.856
musicOrigin1	0.877	1.058	0.875	0.873 <sup>8.2</sup>	0.912 <sup>8.4</sup>	<b>0.839</b> <sup>7.9</sup>	0.863 <sup>9.0</sup>	1.493
musicOrigin2	0.873	1.045	0.854	0.873 <sup>7.9</sup>	0.916 <sup>8.1</sup>	<b>0.851</b> <sup>8.8</sup>	0.873 <sup>8.9</sup>	1.493
friedman	0.946	1.11	1.02	<b>0.936</b> <sup>8.0</sup>	0.95 <sup>7.8</sup>	1.051 <sup>7.9</sup>	1.093 <sup>7.3</sup>	0.936
mp5spec	1.142	1.218	<b>1.08</b>	1.139 <sup>7.8</sup>	1.124 <sup>6.5</sup>	1.173 <sup>9.0</sup>	1.216 <sup>9.0</sup>	2.038
mp6spec	1.135	1.185	<b>1.054</b>	1.133 <sup>7.3</sup>	1.118 <sup>4.8</sup>	1.153 <sup>8.9</sup>	1.183 <sup>9.0</sup>	2.036
polymer	0.722	0.679	0.683	0.678 <sup>5.4</sup>	0.672 <sup>5.3</sup>	0.67 <sup>4.4</sup>	<b>0.644</b> <sup>3.0</sup>	1.563
rf1	0.643	0.656	0.743	<b>0.624</b> <sup>2.4</sup>	0.691 <sup>1.0</sup>	0.656 <sup>3.0</sup>	0.644 <sup>3.1</sup>	1.664
rf2	0.606	0.7	0.635	0.601 <sup>3.4</sup>	0.696 <sup>1.0</sup>	<b>0.596</b> <sup>2.7</sup>	0.619 <sup>3.7</sup>	1.537
scm1d	0.536	0.498	<b>0.429</b>	0.554 <sup>3.4</sup>	0.641 <sup>1.0</sup>	0.536 <sup>4.0</sup>	0.548 <sup>4.0</sup>	0.471
scm20d	<b>0.748</b>	0.786	0.777	0.766 <sup>3.6</sup>	0.81 <sup>1.0</sup>	0.76 <sup>4.0</sup>	0.76 <sup>4.0</sup>	0.749
mean	0.789	0.943	0.958	<b>0.785</b> <sup>5.78</sup>	0.811 <sup>4.23</sup>	0.801 <sup>6.89</sup>	0.8 <sup>6.91</sup>	1.133

Table 4: Performance in terms of aRRMSE for multi-target regression. The numbers above the measures correspond to the mean number of layers in the final model.

As can be seen in Table 4, our proposed X\_TE provides the best results on average. When compared to the main competitor (DSTARS), X\_TE was ranked in a higher position according to the Friedman-Nemenyi test presented in Figure 4.

The second best method was the baseline ensemble RF + ET. In few cases, it managed to be the best method, nonetheless, on average, it is slightly worse than our X\_TE variant. These results are very surprising, since RF + ET is

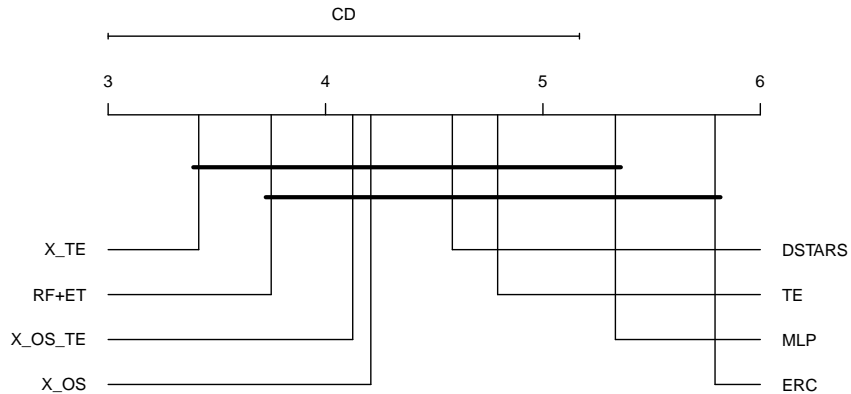


Figure 4: Friedman-Nemenyi comparing the multi-target regression methods.

regarded as a relatively simple approach. Despite being very similar between themselves, when compared to each other, X\_TE provided the best performance in more datasets than RF + ET. In 15 out of 24 datasets, X\_TE was the superior method, whereas RF + ET yielded the best performance in 7 datasets and tied in 2 cases. These results are reinforced by the Friedman-Nemenyi diagram shown in Figure 4, where X\_TE is ranked higher than RF + ET. This finding also validates our choice of using RF + ET as building blocks for more complex models. Despite its relative simplicity, it managed to be superior than complex models such as the state-of-art comparison methods.

The performance of the comparison methods was unexpected. Even though ERC managed to achieve the highest performance in some cases, on average it is considerably worse than X\_TE. Likewise, DSTARS also had the upper-hand in several cases, nonetheless when all datasets are considered, its performance is underwhelming. This is due to their oscillating performance where they perform better in some cases, but very poorly in others. A similar behaviour was also noticed for the MLP.

When comparing our variants with each other, the TE variant provided inferior results in many cases. It also limited itself to shallower models due to its performance sequentially decreasing on the training dataset. This is an indication that tree-embeddings should be used to augment the original features



and not entirely replace them. Hence, it highlights the necessity of using the original features (X) with the extra features.

Moreover, the variants that employ output space features, X\_OS and X\_OS.TE, were also associated to the best performance in a few datasets. Despite being almost equivalent among themselves, they are not competitive when compared to X\_TE. Such evidence attests our hypothesis that learning a representation component at each layer of our proposed method (DTE) leads to superior predictive performance.

These results also validate our stopping (pruning) criterion. In comparison to the results obtained with different stopping criteria, our performance is superior to immediately stopping the training after deterioration (Appendix Table 1), and it approaches the best performance in the test (Appendix Table 2).

### 5.2. Multi-label classification

Datasets ↓	RF+ET	ML-KNN	Rakel	MLP	ECC	MLDF	X_TE	TE	X_OS	X_OS.TE
birds	0.047	0.061	0.053	0.067	0.055	0.044 <sup>6.9</sup>	0.047 <sup>4.78</sup>	0.05 <sup>2.89</sup>	0.042 <sup>4.56</sup>	<b>0.041</b> <sup>5.33</sup>
CAL500	0.138	0.161	<b>0.138</b>	0.139	0.16	0.227 <sup>1.1</sup>	<b>0.138</b> <sup>1.44</sup>	<b>0.138</b> <sup>1.56</sup>	0.144 <sup>2.33</sup>	0.142 <sup>1.67</sup>
CHD_49	<b>0.286</b>	0.327	0.287	0.306	0.338	0.314 <sup>3.0</sup>	0.29 <sup>6.0</sup>	0.29 <sup>5.67</sup>	0.312 <sup>9.0</sup>	0.311 <sup>9.0</sup>
emotions	<b>0.182</b>	0.27	0.25	0.216	0.243	0.198 <sup>4.4</sup>	0.189 <sup>6.22</sup>	0.192 <sup>5.67</sup>	0.192 <sup>8.78</sup>	0.191 <sup>8.89</sup>
enron	0.047	0.06	0.051	0.055	0.056	0.063 <sup>2.1</sup>	0.047 <sup>4.78</sup>	0.048 <sup>7.22</sup>	<b>0.046</b> <sup>9.0</sup>	<b>0.046</b> <sup>8.44</sup>
flags	<b>0.248</b>	0.315	0.321	0.297	0.301	0.278 <sup>3.9</sup>	0.24 <sup>6.67</sup>	0.251 <sup>6.67</sup>	0.285 <sup>9.0</sup>	0.27 <sup>8.78</sup>
genbase	0.047	<b>0.001</b>	0.047	<b>0.001</b>	<b>0.001</b>	0.025 <sup>5.8</sup>	0.017 <sup>4.67</sup>	0.005 <sup>4.44</sup>	0.042 <sup>4.89</sup>	0.029 <sup>6.0</sup>
Gram_negative	0.046	0.021	0.076	0.014	<b>0.012</b>	0.017 <sup>7.9</sup>	0.018 <sup>5.56</sup>	0.016 <sup>3.89</sup>	0.033 <sup>5.89</sup>	0.017 <sup>7.67</sup>
Gram_positive	0.067	0.04	0.149	0.038	<b>0.033</b>	0.048 <sup>4.7</sup>	0.035 <sup>5.0</sup>	<b>0.033</b> <sup>7.0</sup>	0.036 <sup>5.89</sup>	0.034 <sup>6.33</sup>
GrampositivePseAAC	0.239	0.27	0.269	0.274	0.303	0.239 <sup>4.5</sup>	<b>0.233</b> <sup>5.78</sup>	0.249 <sup>4.11</sup>	0.251 <sup>7.11</sup>	0.248 <sup>7.78</sup>
LLOG	0.153	<b>0.108</b>	0.16	0.11	0.124	0.154 <sup>3.4</sup>	0.123 <sup>3.22</sup>	0.124 <sup>1.44</sup>	0.146 <sup>6.22</sup>	0.155 <sup>6.33</sup>
medical	0.022	0.017	0.028	<b>0.012</b>	0.01	0.014 <sup>10.7</sup>	0.016 <sup>6.33</sup>	0.017 <sup>1.0</sup>	0.017 <sup>6.56</sup>	0.015 <sup>7.11</sup>
PlantGO	0.062	0.054	0.09	0.045	0.046	0.045 <sup>10.4</sup>	0.051 <sup>5.22</sup>	<b>0.043</b> <sup>2.67</sup>	0.056 <sup>4.89</sup>	0.046 <sup>7.78</sup>
scene	0.123	0.115	0.155	0.105	0.151	0.114 <sup>7.3</sup>	0.119 <sup>2.22</sup>	0.123 <sup>2.11</sup>	0.095 <sup>8.78</sup>	<b>0.091</b> <sup>8.89</sup>
VirusGO	0.119	0.066	0.246	0.054	<b>0.035</b>	0.054 <sup>5.3</sup>	0.077 <sup>3.44</sup>	0.069 <sup>2.78</sup>	0.095 <sup>5.67</sup>	0.071 <sup>4.0</sup>
VirusPseAAC	0.21	0.231	0.24	0.242	0.288	0.226 <sup>2.5</sup>	<b>0.206</b> <sup>5.56</sup>	0.22 <sup>7.0</sup>	0.222 <sup>4.22</sup>	0.215 <sup>5.67</sup>
yeast	0.195	0.205	0.221	0.199	0.219	0.204 <sup>7.0</sup>	<b>0.194</b> <sup>8.44</sup>	0.202 <sup>6.22</sup>	0.204 <sup>8.89</sup>	0.2 <sup>9.0</sup>
mean	0.131	0.137	0.164	0.128	0.14	0.133 <sup>5.35</sup>	<b>0.12</b> <sup>5.02</sup>	0.122 <sup>4.26</sup>	0.13 <sup>6.57</sup>	0.125 <sup>6.98</sup>

Table 5: Hamming Distance using MicroAUC to prune

As detailed in Tables 5, 6, 7, 8 and 9, our variants are often associated to the best performances in most of the measures. Similar to the MTR experiments, the X\_TE performs better in the majority of the experiments as stated by the hamming loss, label ranking, micro average precision score and microAUC.

Datasets ↓	RF+ET	ML-KNN	Rakel	MLP	ECC	MLDF	X_TE	TE	X_OS	X_OS_TE
birds	0.068	0.473	0.537	0.427	0.274	0.329 <sup>4.6</sup>	0.073 <sup>4.78</sup>	0.094 <sup>2.89</sup>	<b>0.062</b> <sup>4.56</sup>	0.07 <sup>5.33</sup>
CAL500	<b>0.175</b>	0.707	0.782	0.762	0.33	0.643 <sup>1.5</sup>	<b>0.175</b> <sup>1.44</sup>	0.177 <sup>1.56</sup>	0.196 <sup>2.33</sup>	0.191 <sup>1.67</sup>
CHD_49	0.203	0.534	0.449	0.485	0.316	0.508 <sup>4.5</sup>	0.203 <sup>6.0</sup>	<b>0.201</b> <sup>5.67</sup>	0.224 <sup>9.0</sup>	0.249 <sup>0</sup>
emotions	<b>0.145</b>	0.576	0.52	0.538	0.268	0.405 <sup>4.6</sup>	0.149 <sup>6.22</sup>	0.154 <sup>5.67</sup>	0.159 <sup>8.78</sup>	0.171 <sup>8.89</sup>
enron	<b>0.076</b>	0.553	0.734	0.488	0.258	0.311 <sup>5.9</sup>	<b>0.076</b> <sup>4.78</sup>	0.088 <sup>7.22</sup>	0.081 <sup>9.0</sup>	0.085 <sup>8.44</sup>
flags	<b>0.185</b>	0.51	0.484	0.499	0.324	0.606 <sup>2.5</sup>	<b>0.185</b> <sup>6.67</sup>	0.201 <sup>6.67</sup>	0.233 <sup>9.0</sup>	0.241 <sup>8.78</sup>
genbase	0.034	0.017	1.0	0.011	0.004	0.273 <sup>5.5</sup>	<b>0.007</b> <sup>4.67</sup>	<b>0.007</b> <sup>4.44</sup>	0.016 <sup>4.89</sup>	0.009 <sup>6.0</sup>
Gram_negative	0.017	0.081	0.482	0.057	0.033	0.222 <sup>9.1</sup>	<b>0.01</b> <sup>5.56</sup>	0.013 <sup>3.89</sup>	0.011 <sup>5.89</sup>	0.009 <sup>7.67</sup>
Gram_positive	0.051	0.088	0.297	0.086	0.049	0.531 <sup>6.9</sup>	0.036 <sup>5.0</sup>	<b>0.035</b> <sup>7.0</sup>	0.042 <sup>5.89</sup>	0.032 <sup>6.33</sup>
GrampositivePseAAC	0.26	0.696	0.537	0.611	0.448	0.805 <sup>2.4</sup>	0.25 <sup>5.78</sup>	<b>0.249</b> <sup>4.11</sup>	0.259 <sup>7.11</sup>	0.263 <sup>7.78</sup>
LLOG	0.157	0.527	0.701	0.521	0.314	0.69 <sup>1.0</sup>	<b>0.131</b> <sup>3.22</sup>	0.134 <sup>1.44</sup>	0.152 <sup>6.22</sup>	0.191 <sup>6.33</sup>
medical	0.035	0.344	1.0	0.312	0.108	0.185 <sup>11.1</sup>	0.03 <sup>6.33</sup>	0.036 <sup>1.0</sup>	0.033 <sup>6.56</sup>	<b>0.028</b> <sup>7.11</sup>
PlantGO	0.041	0.347	1.0	0.316	0.149	0.132 <sup>10.4</sup>	0.036 <sup>5.22</sup>	0.045 <sup>2.67</sup>	0.034 <sup>4.89</sup>	<b>0.037</b> <sup>7.78</sup>
scene	0.113	0.382	0.664	0.314	0.298	0.475 <sup>12.2</sup>	0.109 <sup>2.22</sup>	0.13 <sup>2.11</sup>	<b>0.085</b> <sup>8.78</sup>	0.098 <sup>8.89</sup>
VirusGO	0.088	0.168	0.796	0.152	0.071	0.193 <sup>5.6</sup>	0.039 <sup>3.44</sup>	0.052 <sup>2.78</sup>	0.066 <sup>5.67</sup>	<b>0.034</b> <sup>4.0</sup>
VirusPseAAC	<b>0.222</b>	0.7	0.806	0.732	0.496	0.511 <sup>3.9</sup>	<b>0.222</b> <sup>5.56</sup>	0.241 <sup>7.0</sup>	0.233 <sup>4.22</sup>	0.261 <sup>5.67</sup>
yeast	0.167	0.451	0.577	0.44	0.238	0.462 <sup>7.5</sup>	<b>0.166</b> <sup>8.44</sup>	0.176 <sup>6.22</sup>	0.187 <sup>8.89</sup>	0.195 <sup>9.0</sup>
mean	0.12	0.421	0.669	0.397	0.234	0.428 <sup>5.84</sup>	<b>0.112</b> <sup>5.02</sup>	0.12 <sup>4.26</sup>	0.122 <sup>6.57</sup>	0.126 <sup>6.98</sup>

Table 6: Label ranking using MicroAUC to prune

Datasets ↓	RF+ET	ML-KNN	Rakel	MLP	ECC	MLDF	X_TE	TE	X_OS	X_OS_TE
birds	0.509	0.533	0.537	0.545	0.549	0.549 <sup>7.0</sup>	0.509 <sup>4.78</sup>	0.532 <sup>2.89</sup>	0.468 <sup>4.56</sup>	<b>0.463</b> <sup>5.33</sup>
CAL500	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b>	<b>1.0</b> <sup>1.0</sup>	<b>1.0</b> <sup>1.44</sup>	<b>1.0</b> <sup>1.56</sup>	1.0 <sup>2.33</sup>	<b>1.0</b> <sup>1.67</sup>
CHD_49	0.836	0.886	<b>0.804</b>	0.848	0.89	0.885 <sup>3.2</sup>	0.822 <sup>6.0</sup>	0.834 <sup>5.67</sup>	0.812 <sup>9.0</sup>	0.814 <sup>9.0</sup>
emotions	0.689	0.82	0.77	0.77	0.792	0.727 <sup>3.7</sup>	0.682 <sup>6.22</sup>	0.695 <sup>5.67</sup>	0.661 <sup>8.78</sup>	<b>0.654</b> <sup>8.89</sup>
enron	0.898	0.917	0.989	0.879	0.884	0.916 <sup>11.3</sup>	0.899 <sup>4.78</sup>	0.905 <sup>7.22</sup>	0.879 <sup>9.0</sup>	<b>0.866</b> <sup>8.44</sup>
flags	0.84	0.932	0.861	0.915	0.828	0.86 <sup>2.7</sup>	<b>0.754</b> <sup>6.67</sup>	0.823 <sup>6.67</sup>	0.766 <sup>9.0</sup>	<b>0.754</b> <sup>8.78</sup>
genbase	1.0	0.029	1.0	0.024	<b>0.022</b>	0.404 <sup>4.8</sup>	0.31 <sup>4.67</sup>	0.099 <sup>4.44</sup>	0.931 <sup>4.89</sup>	0.713 <sup>6.0</sup>
Gram_negative	0.326	0.113	0.489	0.081	<b>0.07</b>	0.198 <sup>8.4</sup>	0.121 <sup>5.56</sup>	0.096 <sup>3.89</sup>	0.23 <sup>5.89</sup>	0.113 <sup>7.67</sup>
Gram_positive	0.256	0.096	0.301	0.1	<b>0.09</b>	0.102 <sup>5.7</sup>	0.113 <sup>5.0</sup>	0.096 <sup>7.0</sup>	0.111 <sup>5.89</sup>	0.096 <sup>6.33</sup>
GrampositivePseAAC	0.87	0.699	<b>0.541</b>	0.652	0.761	0.919 <sup>2.7</sup>	0.806 <sup>5.78</sup>	0.763 <sup>4.11</sup>	0.598 <sup>7.11</sup>	0.592 <sup>7.78</sup>
LLOG	1.0	<b>0.859</b>	1.0	<b>0.859</b>	<b>0.859</b>	1.0 <sup>2.8</sup>	<b>0.859</b> <sup>3.22</sup>	<b>0.859</b> <sup>1.44</sup>	0.948 <sup>6.22</sup>	0.979 <sup>6.33</sup>
medical	0.804	0.513	1.0	0.411	<b>0.322</b>	0.45 <sup>15.0</sup>	0.569 <sup>6.33</sup>	0.614 <sup>1.0</sup>	0.627 <sup>6.56</sup>	0.512 <sup>7.11</sup>
PlantGO	0.658	0.416	1.0	0.367	<b>0.323</b>	0.421 <sup>16.5</sup>	0.526 <sup>5.22</sup>	0.393 <sup>2.67</sup>	0.556 <sup>4.89</sup>	0.443 <sup>7.78</sup>
scene	0.619	0.456	0.69	0.413	0.605	0.491 <sup>9.0</sup>	0.58 <sup>2.22</sup>	0.594 <sup>2.11</sup>	0.368 <sup>8.78</sup>	<b>0.333</b> <sup>8.89</sup>
VirusGO	0.567	0.251	0.862	0.23	<b>0.134</b>	0.258 <sup>5.3</sup>	0.337 <sup>3.44</sup>	0.278 <sup>2.78</sup>	0.443 <sup>5.67</sup>	0.299 <sup>4.0</sup>
VirusPseAAC	0.952	<b>0.791</b>	0.856	0.871	0.909	0.976 <sup>3.1</sup>	0.92 <sup>5.56</sup>	0.867 <sup>7.0</sup>	0.883 <sup>4.22</sup>	0.856 <sup>5.67</sup>
yeast	0.855	0.799	0.97	0.814	0.858	0.831 <sup>7.1</sup>	0.828 <sup>8.44</sup>	0.84 <sup>6.22</sup>	0.769 <sup>8.89</sup>	<b>0.761</b> <sup>9.0</sup>
mean	0.746	0.595	0.804	<b>0.575</b>	0.582	0.646 <sup>6.43</sup>	0.626 <sup>5.02</sup>	0.605 <sup>4.26</sup>	0.65 <sup>6.57</sup>	0.603 <sup>6.98</sup>

Table 7: One error using MicroAUC to prune

Datasets $\uparrow$	RF+ET	ML-KNN	Rakel	MLP	ECC	X_TE	TE	X_OS	X_OS_TE
birds	0.599	0.09	0.053	0.12	0.315	0.595 <sup>4.78</sup>	0.502 <sup>2.89</sup>	<b>0.628</b> <sup>4.56</sup>	0.618 <sup>5.33</sup>
CAL500	<b>0.485</b>	0.244	0.253	0.259	0.334	0.484 <sup>1.44</sup>	0.477 <sup>1.56</sup>	0.453 <sup>2.33</sup>	0.461 <sup>1.67</sup>
CHD_49	<b>0.707</b>	0.549	0.589	0.57	0.621	0.706 <sup>6.0</sup>	0.697 <sup>5.67</sup>	0.669 <sup>9.0</sup>	0.649 <sup>9.0</sup>
emotions	0.751	0.433	0.466	0.514	0.601	0.752 <sup>6.22</sup>	0.739 <sup>5.67</sup>	<b>0.754</b> <sup>8.78</sup>	0.718 <sup>8.89</sup>
enron	0.615	0.271	0.269	0.32	0.434	<b>0.618</b> <sup>4.78</sup>	0.588 <sup>7.22</sup>	0.598 <sup>9.0</sup>	0.591 <sup>8.44</sup>
flags	<b>0.83</b>	0.624	0.621	0.638	0.737	0.826 <sup>6.67</sup>	0.791 <sup>6.67</sup>	0.782 <sup>9.0</sup>	0.769 <sup>8.78</sup>
genbase	0.686	0.977	0.047	0.977	0.988	0.979 <sup>4.67</sup>	<b>0.982</b> <sup>4.44</sup>	0.886 <sup>4.89</sup>	0.974 <sup>6.0</sup>
Gram_negative	0.954	0.86	0.501	0.901	0.937	<b>0.981</b> <sup>5.56</sup>	0.976 <sup>3.89</sup>	0.966 <sup>5.89</sup>	0.987 <sup>6.7</sup>
Gram_positive	0.958	0.87	0.574	0.875	0.94	0.979 <sup>5.0</sup>	0.976 <sup>7.0</sup>	0.975 <sup>8.9</sup>	<b>0.981</b> <sup>6.33</sup>
GrampositivePseAAC	0.513	0.317	0.354	0.335	0.338	<b>0.516</b> <sup>5.78</sup>	0.495 <sup>4.11</sup>	0.499 <sup>7.11</sup>	0.492 <sup>7.78</sup>
LLOG	0.656	0.528	0.308	0.521	0.625	<b>0.7</b> <sup>3.22</sup>	0.686 <sup>1.44</sup>	0.605 <sup>6.22</sup>	0.542 <sup>6.33</sup>
medical	0.718	0.468	0.028	0.588	0.77	<b>0.796</b> <sup>6.33</sup>	0.784 <sup>1.0</sup>	0.745 <sup>6.56</sup>	0.776 <sup>7.11</sup>
PlantGO	0.795	0.503	0.09	0.562	0.68	0.834 <sup>5.22</sup>	<b>0.846</b> <sup>2.67</sup>	0.813 <sup>4.89</sup>	0.844 <sup>7.78</sup>
scene	0.711	0.514	0.37	0.561	0.523	0.719 <sup>2.22</sup>	0.689 <sup>2.11</sup>	<b>0.746</b> <sup>8.78</sup>	0.714 <sup>8.89</sup>
VirusGO	0.804	0.734	0.231	0.779	0.89	<b>0.929</b> <sup>3.44</sup>	0.918 <sup>2.78</sup>	0.853 <sup>5.67</sup>	0.927 <sup>4.0</sup>
VirusPseAAC	0.435	0.268	0.236	0.261	0.279	<b>0.439</b> <sup>5.56</sup>	0.431 <sup>7.0</sup>	0.425 <sup>4.22</sup>	0.417 <sup>5.67</sup>
yeast	0.724	0.532	0.49	0.543	0.638	<b>0.727</b> <sup>8.44</sup>	0.705 <sup>6.22</sup>	0.702 <sup>8.89</sup>	0.684 <sup>9.0</sup>
mean	0.702	0.517	0.322	0.548	0.626	<b>0.74</b> <sup>5.02</sup>	0.722 <sup>4.26</sup>	0.711 <sup>6.57</sup>	0.714 <sup>6.98</sup>

Table 8: Micro average precision score using MicroAUC to prune

Datasets $\uparrow$	RF+ET	Rakel	MLP	ECC	X_TE	TE	X_OS	X_OS_TE
birds	0.925	0.5	0.601	0.744	0.921 <sup>4.78</sup>	0.879 <sup>2.89</sup>	<b>0.931</b> <sup>4.56</sup>	0.923 <sup>5.33</sup>
CAL500	<b>0.823</b>	0.597	0.605	0.741	<b>0.823</b> <sup>1.44</sup>	0.82 <sup>1.56</sup>	0.8 <sup>2.33</sup>	0.805 <sup>1.67</sup>
CHD_49	<b>0.794</b>	0.711	0.688	0.722	0.793 <sup>6.0</sup>	0.788 <sup>5.67</sup>	0.759 <sup>9.0</sup>	0.74 <sup>9.0</sup>
emotions	<b>0.876</b>	0.688	0.71	0.789	0.875 <sup>6.22</sup>	0.864 <sup>5.67</sup>	0.872 <sup>8.78</sup>	0.858 <sup>8.89</sup>
enron	<b>0.923</b>	0.644	0.734	0.839	<b>0.923</b> <sup>4.78</sup>	0.909 <sup>7.22</sup>	0.919 <sup>9.0</sup>	0.914 <sup>8.44</sup>
flags	<b>0.84</b>	0.678	0.702	0.774	<b>0.84</b> <sup>6.67</sup>	0.812 <sup>6.67</sup>	0.799 <sup>9.0</sup>	0.791 <sup>8.78</sup>
genbase	0.962	0.5	0.991	0.997	0.989 <sup>4.67</sup>	<b>0.992</b> <sup>4.44</sup>	0.984 <sup>4.89</sup>	0.986 <sup>6.0</sup>
Gram_negative	0.99	0.745	0.964	0.979	<b>0.995</b> <sup>5.56</sup>	0.993 <sup>3.89</sup>	0.993 <sup>5.89</sup>	<b>0.995</b> <sup>7.67</sup>
Gram_positive	0.973	0.802	0.946	0.975	0.989 <sup>5.0</sup>	0.989 <sup>7.0</sup>	0.983 <sup>5.89</sup>	<b>0.992</b> <sup>6.33</sup>
GrampositivePseAAC	0.762	0.642	0.62	0.631	<b>0.77</b> <sup>5.78</sup>	0.764 <sup>1.11</sup>	0.755 <sup>7.11</sup>	0.754 <sup>7.78</sup>
LLOG	0.855	0.6	0.758	0.815	<b>0.881</b> <sup>3.22</sup>	0.877 <sup>1.44</sup>	0.836 <sup>6.22</sup>	0.77 <sup>6.33</sup>
medical	0.965	0.5	0.834	0.941	0.971 <sup>6.33</sup>	0.966 <sup>1.0</sup>	0.97 <sup>6.56</sup>	<b>0.972</b> <sup>7.11</sup>
PlantGO	0.968	0.5	0.827	0.914	0.974 <sup>5.22</sup>	0.972 <sup>2.67</sup>	0.974 <sup>4.89</sup>	<b>0.978</b> <sup>7.78</sup>
scene	0.907	0.646	0.817	0.801	<b>0.909</b> <sup>2.22</sup>	0.89 <sup>2.11</sup>	0.92 <sup>8.78</sup>	0.894 <sup>8.89</sup>
VirusGO	0.926	0.555	0.904	0.963	<b>0.977</b> <sup>3.44</sup>	0.973 <sup>2.78</sup>	0.946 <sup>5.67</sup>	<b>0.977</b> <sup>4.0</sup>
VirusPseAAC	<b>0.774</b>	0.552	0.584	0.628	<b>0.774</b> <sup>5.56</sup>	0.762 <sup>7.0</sup>	0.757 <sup>4.22</sup>	0.743 <sup>5.67</sup>
yeast	0.847	0.686	0.745	0.807	<b>0.848</b> <sup>8.44</sup>	0.835 <sup>6.22</sup>	0.827 <sup>8.89</sup>	0.815 <sup>9.0</sup>
mean	0.889	0.62	0.766	0.827	<b>0.897</b> <sup>5.02</sup>	0.887 <sup>4.26</sup>	0.884 <sup>6.57</sup>	0.877 <sup>6.98</sup>

Table 9: Micro AUC

According to both hamming and ranking losses, X\_TE surpasses the competitors in a large number of datasets, leading to a better performance overall. The superiority of our method is more noticeable according to the micro average precision score, as shown in Table 8. In this case, X\_TE yields a considerably higher mean micro average precision score.

A similar behaviour was perceived for microAUC where X\_TE normally provides the best results. Moreover, we could notice that the inclusion of tree-embeddings leads to an improvement in almost all cases when compared to the original representation (RF + ET).

When analyzing one error, the neural network, MLP, was superior on average, followed by another competitor method, ECC. These results are unexpected since these methods failed to achieve a competitive performance according to the other measures. In particular, these methods were exceptionally better than X\_TE in some datasets, such as genbase, medical and plantGO, nonetheless our variant X\_OS\_TE provides better results in more datasets to the point that it is ranked higher according to its Frieman-Nemenyi test (Appendix Figure 1).

Our main competitor, MLDF, struggled to achieve a competitive performance. With the exception of a few cases, its performance was underwhelming even when compared to the baselines for all three available evaluation measures <sup>17</sup>. This finding is rather surprising at first since it is fairly similar to our variant X\_OS. However, there are some major differences between the two methods. First, there is a different optimization mechanism for the number of trees. Second, the extra features used at each layer are different, since X\_OS directly concatenates the output features, and MLDF employs a measure-aware feature reuse mechanism. Third, our method also employs cross-over between the extra feature sets to prevent overfitting. The difference is highlighted in Figure 5 where MLDF is associated with a low rank, making it statistically significantly different from all variants of our method.

---

<sup>17</sup>The extra features employed by MLDF are measure specific, and they are not defined for micro average precision score and microAUC.

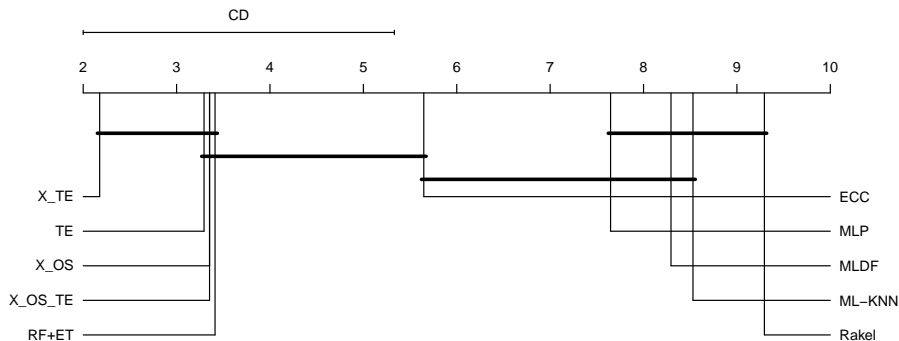


Figure 5: Friedman-Nemenyi test considering the Label Ranking error.

Moreover, the TE variant struggles again to perform competitively against X\_TE, instigating the necessity of incorporating the original features. Likewise, the stopping criterion adopted by our method approximates the best possible performance (Appendix Table 4).

When analyzing the CAL500 dataset, we have noticed that the X\_TE variant is constantly superior than the variants X\_OS and X\_OS\_TE in all evaluation measures. We believe that this is related to the number of labels in its label set. Since CAL500 has a relatively large label set, its output features are often similar to sparse vectors, meaning that, in this case, they may hinder the models performance. This is further exemplified by the performance provided by the baseline (RF + ET) which is always superior to all variants with output features (X\_OS and X\_OS\_TE) in this dataset. When compared to each other, X\_TE is superior to RF + ET. This motivates the usage of tree-embeddings in other applications with larger label sets. Furthermore, at this point, it can be deduced that methods which use directly the output variables to augment the input feature space, such as MLDF, could not be applied to extreme multi-label data. They would lead to extremely high dimensional feature sets, dramatically increasing their computational demands.

In some datasets, such as gram\_negative, gram\_positive and grampositivePseAAC, the performance of most of evaluated methods was close to flawless. Differently

from that, unexpected results were obtained in the CAL500 dataset considering the one error measure. This could be explained by its substantial number of possible labels (174) and its relatively low label cardinality (26.04) which can often lead to wrong predictions being ranked as first.

### 5.3. Effect of the number of layers

In Figure 6, we present the progression of microAUC, our optimized measure, throughout the training of our proposed model for some datasets. The values correspond to the mean value of 10 folds measured in the test partitions.

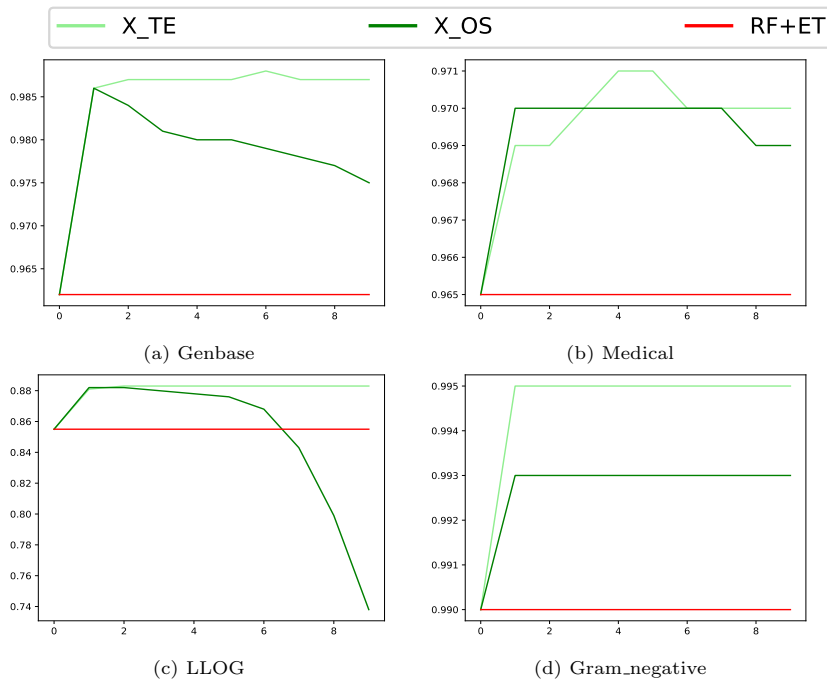


Figure 6: MicroAUC progression throughout the building of our model. The x-axis corresponds to the number of layers in the model, whereas the y-axis corresponds to the MicroAUC measured in the test partitions.

As can be seen, the addition of the tree-embeddings is associated to an immediate improvement in the performance, attesting their representation power. Such increase in the performance can be followed by a plateau, where both

training and testing performances present a static behaviour, as seen in Figure 6c. An oscillating trend was also noticed in some datasets, nonetheless such fluctuation tends to be minor.

The inclusion of output space (OS) features also contributes to the performance. However, its improvement is less pronounced. This difference is visible in Figure 6d where X\_TE approximates a flawless performance. In other situations, such as the ones depicted in Figures 6a and 6c, X\_OS fails to hold a competitive performance, leading often to sub-optimal models.

## 6. Conclusion

In this paper, we have proposed a novel deep tree-ensemble model for multi-output prediction tasks which can integrate tree-embeddings as well as output features in the learning process. Our results have shown that our model provides superior results in both multi-target regression and multi-label classification.

More precisely, they reveal that integrating a representation learning component in each layer can significantly boost prediction performance, affirming our original hypothesis. Tree-embeddings can be more meaningful than just output features in both tasks. Additionally, in some cases, the combination of both can boost the performance even further. We also highlight that preserving the original feature set is essential to achieve high performance, since the tree-embeddings variant rarely yields superior results by itself.

As for future work, we could investigate the incorporation of an extra steps in the prediction phase, such as the outlier removal performed by Moyano [43]. In this case, models whose predictions deviate from the ensemble are not incorporated in the final prediction.

Moreover, this paper focuses on fundamental multi-output prediction. In the future, it would be very interesting to investigate the use of the proposed method in specialized application domains, such as the field of large scale image MLC, where our approach could potentially get combined with state of the art computer vision techniques.

In addition, the investigation of the performance of our method in the field of large-scale data or big data would be a very interesting topic for future research. For example, we could investigate replacing PCA with a computationally more efficient feature transformation method.

Furthermore, we would like to exploit different ensembles of trees. That is, explore if the inclusion of other tree-based ensembles can improve the results, such as predictive bi-clustering trees [46] [47], extreme gradient boosting [48] and oblique forests [49]. Alternative representation components in the construction of the layers of our approach as well as deeper architectures for specific cases could be interesting research topics in the future. Finally, other multi-output tasks such as hierarchical multi-label classification [50] and extreme multi-label classification [51] may also be addressed with our method, specially because the MLDF would result in a very high number of features whereas the tree-embeddings would be unaffected.

### **Acknowledgements**

The authors acknowledge the support from the Research Fund Flanders (through research project G080118N) and from the Flemish Government (AI Research Program).

### **References**

- [1] S. Yin, X. Li, H. Gao, O. Kaynak, Data-based techniques focused on modern industry: An overview, *IEEE Transactions on Industrial Electronics* 62 (1) (2015) 657–667.
- [2] M. I. Jordan, T. M. Mitchell, Machine learning: Trends, perspectives, and prospects., *Science* 349 (6245) (2015) 255–260.
- [3] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, S. Thrun, Dermatologist-level classification of skin cancer with deep neural networks, *nature* 542 (7639) (2017) 115–118.



- [4] G. Eraslan, Ž. Avsec, J. Gagneur, F. J. Theis, Deep learning: new computational modelling techniques for genomics, *Nature Reviews Genetics* 20 (7) (2019) 389–403.
- [5] C. Badue, R. Guidolini, R. V. Carneiro, P. Azevedo, V. B. Cardoso, A. Forechi, L. Jesus, R. Berriel, T. M. Paixão, F. Mutz, L. de Paula Veronese, T. Oliveira-Santos, A. F. De Souza, Self-driving cars: A survey, *Expert Systems with Applications* 165 (2021) 113816. doi:<https://doi.org/10.1016/j.eswa.2020.113816>.  
URL <http://www.sciencedirect.com/science/article/pii/S095741742030628X>
- [6] I. H. Witten, E. Frank, M. A. Hall, *Data Mining: Practical Machine Learning Tools and Techniques*, 4th Edition, Morgan Kaufmann, 2016.
- [7] W. Waegeman, K. Dembczyński, E. Hüllermeier, Multi-target prediction: a unifying view on problems and methods, *Data Mining and Knowledge Discovery* 33 (2) (2019) 293–324.
- [8] G. Tsoumakas, I. Katakis, I. Vlahavas, *Mining Multi-label Data*, Springer US, Boston, MA, 2010, pp. 667–685. doi:10.1007/978-0-387-09823-4\_34.  
URL [https://doi.org/10.1007/978-0-387-09823-4\\_34](https://doi.org/10.1007/978-0-387-09823-4_34)
- [9] Z.-H. Zhou, J. Feng, Deep forest, *National Science Review* 6 (1) (2019) 74–86.
- [10] L. Breiman, Random forests, *Machine learning* 45 (1) (2001) 5–32.
- [11] C. Vens, F. Costa, Random forest based feature induction, in: *2011 IEEE 11th International Conference on Data Mining*, IEEE, 2011, pp. 744–753.
- [12] K. Pliakos, C. Vens, Mining features for biomedical data using clustering tree ensembles, *Journal of biomedical informatics* 85 (2018) 40–48.

- [13] D. Xu, Y. Shi, I. W. Tsang, Y.-S. Ong, C. Gong, X. Shen, Survey on multi-output learning, *IEEE Transactions on Neural Networks and Learning Systems* 31 (7) (2020) 2409–2429. doi:10.1109/TNNLS.2019.2945133.
- [14] H. Borchani, G. Varando, C. Bielza, P. Larrañaga, A survey on multi-output regression, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 5 (5) (2015) 216–233.
- [15] M.-L. Zhang, Z.-H. Zhou, A review on multi-label learning algorithms, *IEEE transactions on knowledge and data engineering* 26 (8) (2013) 1819–1837.
- [16] H. Blockeel, L. De Raedt, Top-down induction of first-order logical decision trees, *Artificial intelligence* 101 (1-2) (1998) 285–297.
- [17] S. M. Mastelini, E. J. Santana, R. Cerri, S. Barbon, Dstars: A multi-target deep structure for tracking asynchronous regressor stacking, *Applied Soft Computing* 91 (2020) 106215. doi:<https://doi.org/10.1016/j.asoc.2020.106215>.  
URL <http://www.sciencedirect.com/science/article/pii/S1568494620301551>
- [18] G. Tsoumakas, I. Katakis, I. Vlahavas, Random k-labelsets for multilabel classification, *IEEE transactions on knowledge and data engineering* 23 (7) (2010) 1079–1089.
- [19] D. Kocev, M. Ceci, T. Stepišnik, Ensembles of extremely randomized predictive clustering trees for predicting structured outputs, *Machine Learning* (2020) 1–29.
- [20] D. Kocev, C. Vens, J. Struyf, S. Džeroski, Tree ensembles for predicting structured outputs, *Pattern Recognition* 46 (3) (2013) 817–833.
- [21] J. Du, Y. Xu, Hierarchical deep neural network for multivariate regression, *Pattern Recognition* 63 (2017) 149–157.

- [22] X. Zhen, M. Yu, Z. Xiao, L. Zhang, L. Shao, Heterogenous output regression network for direct face alignment, *Pattern Recognition* (2020) 107311.
- [23] E. Spyromitros-Xioufis, G. Tsoumakas, W. Groves, I. Vlahavas, Multi-target regression via input space expansion: treating targets as inputs, *Machine Learning* 104 (1) (2016) 55–98.
- [24] M.-L. Zhang, Z.-H. Zhou, MI-knn: A lazy learning approach to multi-label learning, *Pattern Recognition* 40 (7) (2007) 2038 – 2048.  
doi:<https://doi.org/10.1016/j.patcog.2006.12.019>.  
URL <http://www.sciencedirect.com/science/article/pii/S0031320307000027>
- [25] R. Wang, S. Kwong, X. Wang, Y. Jia, Active k-labelsets ensemble for multi-label classification, *Pattern Recognition* 109 (2020) 107583.
- [26] Z. Ma, S. Chen, Expand globally, shrink locally: Discriminant multi-label learning with missing labels, *Pattern Recognition* (2020) 107675.
- [27] Z. Chen, X. Wei, P. Wang, Y. Guo, Learning graph convolutional networks for multi-label recognition and applications, *IEEE Transactions on Pattern Analysis & Machine Intelligence* doi:10.1109/TPAMI.2021.3063496.
- [28] T. Chen, M. Xu, X. Hui, H. Wu, L. Lin, Learning semantic-specific graph representation for multi-label image recognition, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 522–531.
- [29] R. You, Z. Guo, L. Cui, X. Long, Y. Bao, S. Wen, Cross-modality attention with semantic graph embedding for multi-label classification, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, 2020, pp. 12709–12716.
- [30] L. Yang, X.-Z. Wu, Y. Jiang, Z.-H. Zhou, Multi-label learning with deep forest, in: *24th European Conference on Artificial Intelligence (ECAI' 20)*, Santiago de Compostela, Spain, 2020, pp. 1634–1641.

- [31] P. Geurts, D. Ernst, L. Wehenkel, Extremely randomized trees, *Machine learning* 63 (1) (2006) 3–42.
- [32] Q.-W. Wang, L. Yang, Y.-F. Li, Learning from weak-label data: A deep forest expedition., in: *AAAI*, 2020, pp. 6251–6258.
- [33] J. Gao, K. Liu, B. Wang, D. Wang, Q. Hong, An improved deep forest for alleviating the data imbalance problem, *Soft Computing* (2020) 1–17.
- [34] N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, Smote: synthetic minority over-sampling technique, *Journal of artificial intelligence research* 16 (2002) 321–357.
- [35] R. Su, X. Liu, L. Wei, Q. Zou, Deep-resp-forest: A deep forest model to predict anti-cancer drug response, *Methods* 166 (2019) 91–102.
- [36] C. Ma, Z. Liu, Z. Cao, W. Song, J. Zhang, W. Zeng, Cost-sensitive deep forest for price prediction, *Pattern Recognition* 107 (2020) 107499.
- [37] M. Zhou, X. Zeng, A. Chen, Deep forest hashing for image retrieval, *Pattern Recognition* 95 (2019) 114–127.
- [38] O. Sagi, L. Rokach, Ensemble learning: A survey, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 8 (4) (2018) e1249.
- [39] J. Read, B. Pfahringer, G. Holmes, E. Frank, Classifier chains for multi-label classification, *Machine Learning* 85 (3) (2011) 333–359. doi:10.1007/s10994-011-5256-5.  
URL <https://doi.org/10.1007/s10994-011-5256-5>
- [40] D. H. Wolpert, Stacked generalization, *Neural networks* 5 (2) (1992) 241–259.
- [41] J. Read, B. Pfahringer, G. Holmes, E. Frank, Classifier chains: A review and perspectives, *J. Artif. Int. Res.* 70 (2021) 683–718. doi:10.1613/jair.1.12376.  
URL <https://doi.org/10.1613/jair.1.12376>

- [42] G. H. Golub, C. F. Van Loan, *Matrix computations*, Vol. 3, JHU press, 2013.
- [43] J. M. Moyano, E. L. Gibaja, S. Ventura, An evolutionary algorithm for optimizing the target ordering in ensemble of regressor chains, in: 2017 IEEE congress on evolutionary computation (CEC), IEEE, 2017, pp. 2015–2021.
- [44] J. M. Moyano, E. L. Gibaja, K. J. Cios, S. Ventura, Review of ensembles of multi-label classifiers: models, experimental study and prospects, *Information Fusion* 44 (2018) 33–45.
- [45] R. Padilla, S. L. Netto, E. A. da Silva, A survey on performance metrics for object-detection algorithms, in: 2020 International Conference on Systems, Signals and Image Processing (IWSSIP), IEEE, 2020, pp. 237–242.
- [46] K. Pliakos, C. Vens, Network inference with ensembles of bi-clustering trees, *BMC bioinformatics* 20 (1) (2019) 525.
- [47] B. Zamith, F. K. Nakano, R. Cerri, C. Vens, Predictive bi-clustering trees for hierarchical multi-label classification, *ECML PKDD 2020*.
- [48] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, 2016, pp. 785–794.
- [49] R. Katuwal, P. N. Suganthan, L. Zhang, Heterogeneous oblique random forest, *Pattern Recognition* 99 (2020) 107078.
- [50] F. K. Nakano, M. Lietaert, C. Vens, Machine learning for discovering missing or wrong protein function annotations, *BMC bioinformatics* 20 (1) (2019) 485.
- [51] K. Bhatia, H. Jain, P. Kar, M. Varma, P. Jain, Sparse local embeddings for extreme multi-label classification, in: *Advances in neural information processing systems*, 2015, pp. 730–738.