

Received May 13, 2019, accepted July 5, 2019, date of publication July 22, 2019, date of current version August 8, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2930235

A Novel Hyperparameter-Free Approach to Decision Tree Construction That Avoids Overfitting by Design

RAFAEL GARCÍA LEIVA, ANTONIO FERNÁNDEZ ANTA¹, (Senior Member, IEEE),
VINCENZO MANCUSO², (Member, IEEE), AND PAOLO CASARI³, (Senior Member, IEEE)

IMDEA Networks Institute, 28918 Madrid, Spain

Corresponding author: Rafael García Leiva (rafael.garcia@imdea.org)

This work was supported in part by the European Union's Horizon 2020 Research and Innovation Programme (RECAP) under Grant 732667, in part by the Spanish Ministry of Science, Innovation and Universities under Grant TIN2017-88749-R (DisCoEdge), in part by the Comunidad de Madrid under Grant P2018/TCS-4499 (EdgeData-CM) and Grant S2018/TCS-4496 (TAPIR-CM), and in part by the NSF of China under Grant 61520106005. The work of V. Mancuso was supported the Spanish Ministry of Economy and Competitiveness through Ramon-y-Cajal under Grant RYC-2014-16285.

ABSTRACT Decision trees are an extremely popular machine learning technique. Unfortunately, overfitting in decision trees still remains an open issue that sometimes prevents achieving good performance. In this paper, we present a novel approach for the construction of decision trees that avoids the overfitting by design, without losing accuracy. A distinctive feature of our algorithm is that it requires neither the optimization of any hyperparameters, nor the use of regularization techniques, thus significantly reducing the decision tree training time. Moreover, our algorithm produces much smaller and shallower trees than traditional algorithms, facilitating the interpretability of the resulting models. For reproducibility, we provide an open source version of the algorithm.

INDEX TERMS Decision trees, regularization, interpretability, Kolmogorov complexity.

I. INTRODUCTION

Decision trees are a highly popular machine learning technique that has been successfully applied to solve a number of practical problems in areas such as natural language processing, information extraction, data mining, and pattern recognition [1]. Decision trees are also the foundation of more advanced machine learning methods, such as random forests or boosted trees [2].

One of the most challenging tasks for decision tree-building algorithms is to decide when to stop the tree growing process. Most of the tree building methods produce very complex models that *overfit* training data. Overfitted trees not only have poor predictive capabilities on new, previously unseen data, but can also be exceedingly difficult to interpret, which is a key barrier against the adoption of these models in practical applications [3]. A common approach to avoid overfitting in decision trees is *early stopping*, which forces the construction algorithm to stop before the tree becomes too complex. Popular stopping criteria include limiting the maximum depth of the tree, requiring a minimum number of sample points at leaf nodes, or computing the accuracy

The associate editor coordinating the review of this manuscript and approving it for publication was Bo Jin.

gain yielded by new nodes [1]. However, these heuristics require the optimization of (possibly multiple) hyperparameters, which makes the training process computationally expensive.

An alternative approach to avoid overfitting in decision trees is to apply a post-processing *pruning* algorithm. Pruning removes those branches that contribute the least to the accuracy of the final model [4]. Cost-complexity [2] is a common pruning technique, based on applying a cost function that takes into account how well the tree fits the training data and a regularization factor based on how complex the tree is. If we denote by $E(T)$ the fitness of a tree T , and by $L(T)$ its complexity, the cost function $C(T)$ could be given by

$$C(T) = E(T) + \lambda L(T), \quad (1)$$

where λ is a configurable parameter that controls the trade-off between fitness and complexity. In practice, the accuracy of the model is typically used as a fitness metric, and the number of leaf nodes in the tree as a complexity metric. Although, on average, cost complexity pruning produces better results than limiting the growth of the tree via early stopping [5], this technique requires the optimization of the hyperparameter λ .

In this paper we propose a decision tree construction algorithm with early stopping properties, that does not require hyperparameter tuning or pruning in order to get good accuracy and trees of small size. We provide an open source version of the algorithm in order to facilitate the reproducibility of our results.

A. RELATED WORK

Earlier algorithms for the construction of decision trees were based on an exhaustive search of the space of possible trees (see [6] for a historical review of the discipline). However, due to the computational complexity of those algorithms, alternative approaches were investigated. Most of the modern algorithms used in practice are based on greedy search techniques (see for example [7]), that significantly reduce the training time at the cost of producing sub-optimal trees.

Among the algorithms to build decision trees, the Iterative Dichotomiser 3 algorithm, ID3 [8], is based on an information gain splitting criteria, which stops when all the training samples belong to a single target value, or when no possible split has a positive information gain. Classification And Regression Trees, CART [9], builds binary trees based on a growing splitting criteria and applying a cost-complexity pruning. CHi-squared Automatic Interaction Detector, CHAID [10], relies on an adjusted significance tests to find the least significant values with respect to the target attribute, and the process is repeated until no significant pairs are found. C4.5 [11] is an evolution of ID3 that uses gain ratio as a splitting criteria, and ceases to grow the tree when the number of instances to split is below a given threshold. It includes an accuracy based post-pruning. The Quick Unbiased Efficient Statistical Tree algorithm, QUEST [12], uses quadratic discriminant analysis for splits, the stopping is based on statistical tests, and the resulting trees are pruned based on cross-validation. Multivariate Adaptive Regression Splines, MARS [13], is based on a two-phase (forward and backward pass) algorithm that uses regression functions approximated using linear splines and their tensor products. This two-phase approach is the same as that used by recursive partitioning trees. PUBLIC [14] is based on minimizing the sum of the length of the shortest computer program that can print the data given as input to the tree, and the length of the shortest program that can print the tree. Recent trends in tree building algorithms include deep trees [15] of hundreds of levels, and oblique trees [16] with more advanced if-else conditions. A detailed comparison of decision trees algorithms can be found in [17], and an up-to-date review including other classification methods in [18].

The above approaches rely on the right setting of one or more hyper-parameters, or require to run post-pruning processes in order to provide good performance while preventing the overfitting of the training data. Hyperparameter optimization can be achieved through an exhaustive search when a single tree is being trained, but would be prohibitively complex when training large ensembles of models, as in the case of random forests. Another critical element of decision

tree building algorithms is their ability to take those local decisions that yield the best global forecast performance. In this work, we apply string complexity concepts in order to achieve a global performance given local information, as explained in the following sections.

B. CONTRIBUTIONS

The main contribution of this work is a novel algorithm to construct decision trees that, by design, do not overfit the training data. The algorithm is based on a novel cost function that is used to decide when to stop the building process, rather than to indicate how to prune the generated tree. Our algorithm does not depend on the optimization of hyperparameters, thus considerably reducing the training time. Moreover, the absence of hyperparameters means that we can process a very significant amount of data with reduced complexity, e.g., by using the full dataset for training, with no need to enact a training / validation split.

The models generated with our algorithm are significantly smaller in terms of number of nodes, and shallower than the models generated by other decision tree building algorithms; at the same time, the obtained models present similar, and sometimes better accuracy. Shallower models allow us to make faster predictions when they are used as part of large ensembles of trees, and largely improve interpretability, e.g., when examined by domain experts. Additional advantages of our new algorithm include lower sensitivity to the presence of errors in the training dataset and to the non-linearity of the hyper-boundaries between different classes.

II. NOTATION AND BACKGROUND

A. NOTATION

Let \mathcal{X} be a dataset composed of n training input vectors $x_i \in \mathbb{R}^m$, where $1 \leq i \leq n$, and let $y \in \mathcal{G}^n$ be a category vector, where \mathcal{G} is the set of classification labels ($\mathcal{G} = \{0, 1, \dots, G\}$). The problem at hand is to find the best model f , from a given family of models \mathcal{F} , such that $f(x_i) = y_i$ for many $x_i \in \mathcal{X}$. That is, we are interested in solving a *supervised classification problem*. We are also interested in the capability of the model f to generalize to previously unseen data, that is, to correctly classify input vectors extracted from a universe \mathcal{U} of input vectors, not necessarily included in the training dataset \mathcal{X} . For this reason, the best solution is usually not the overfitted model f that guarantees $f(x_i) = y_i$ for all $x_i \in \mathcal{X}$.

B. KOLMOGOROV COMPLEXITY

The novel cost function introduced in this paper to construct decision trees is based on the concept of Kolmogorov complexity [19]–[21], also known as Algorithmic Information Theory. The application of Kolmogorov complexity to the search of optimal statistical models is implemented by the Minimum Description Length (MDL) principle [22] and the Minimum Message Length (MML) [23]. In particular, minimum length techniques have been applied to the problem of inferring decision trees in [24], later on clarified and

extended in [25], as well as in [26] as a technique for pruning; the PUBLIC algorithm is based on the MML concept and is described in [14]. Although the underlining concepts behind our own cost function are the same (namely, that learning is equivalent to the capability to compress), our approach is very different from the ones described above.

According to the Kolmogorov complexity, the amount of information of an object, encoded as a finite string, is given by the length of the shortest computer program that is able to print (output) that string. Kolmogorov complexity does not require to know the set of valid strings in advance, or to make assumptions about their probability distribution. Therefore, it provides a universal definition for the amount of information contained in an object. However, the Kolmogorov complexity is a non-computable quantity [27], hence it must be approximated in practice.

The *Kolmogorov complexity* $K(s)$ of a string s composed of symbols from a fixed alphabet Σ is defined as¹

$$K(s) = \min_{p,v} \{|p| + |v| : U(p, v) = s\}, \quad (2)$$

where U is a universal computer, p is a program in a prefix-free language interpreted by U ,² and v is the input to the program. It can be shown that the Kolmogorov complexity $K(s)$ of a string s does not depend on the programming language used [27], i.e., any reasonable and sufficiently powerful computer language provides the same description length, up to a fixed additive constant that depends on the selected language, but not on the string itself.

Sometimes, the description of a string can be greatly reduced if we assume the knowledge of another string. The *conditional Kolmogorov complexity* of a string s given the string s' , denoted by $K(s|s')$, is defined as³

$$K(s|s') = \min_{p,v} \{|p| + |v| : U(p, \langle v, s' \rangle) = s\}. \quad (3)$$

III. DECISION TREE CONSTRUCTION ALGORITHM

A. KEY IDEA

In this paper, we provide an algorithm to construct decision trees that, by design, does not overfit training data, and that has no hyperparameters to be optimized. To achieve this, the algorithm must automatically understand when growing the decision tree adds needless complexity, and must measure such complexity in a way that is commensurate to some prediction quality aspect, e.g., inaccuracy. We argue that a natural way to achieve the above objectives is to define both the inaccuracy and the complexity (conveyed by the so-called *surfeit*, introduced later on) using the concept of Kolmogorov complexity [19].

The key insight behind the application of Kolmogorov complexity to machine learning is that the more patterns we can find in a dataset, the more we have learned about the

data. This means both identifying the dataset attributes that best explain the outcome, and devising a model that provides this explanation without exceeding complexity. Moreover, data compression is about finding and exploiting patterns and regularities in the data. A practical implementation that puts together the key features of Kolmogorov complexity and data compression follows the lines of the Minimum Description Length principle [22], that proposes to minimize the length of the model M plus the length of the data D given the model, namely $L(M) + L(D | M)$, where $L(\cdot)$ returns the length of its argument. Unfortunately, such an approach tends to favor simple models, and is not widely used in practice due to this limitation.

In this work, we propose to replace the minimization of $L(M) + L(D | M)$ by a multi-objective optimization problem [28] that reconciles two conflicting goals: to minimize the complexity of the model, and to minimize the inaccuracy of the model's predictions. Moreover, instead of working directly with the length of models $L(M)$ we propose to work with their *surfeit*, which evaluates the unnecessary complexity of a model used to represent a given dataset. The *surfeit* can be computed as the length of the model minus the length of the minimum computer program that can print the dataset, i.e., the Kolmogorov complexity of the data. Namely, this is $L(M) - K(D)$. Using the *surfeit* avoids the rejection of correct but complex models. Our approach is validated in practice by means of its application to a family of decision tree models.

B. OVERVIEW OF THE ALGORITHM

We describe now the algorithm (called Minimum Surfeit and Inaccuracy, or MSI for short) to build a decision tree given a training dataset \mathcal{X} with the help of the pseudocode reported in Algorithm 1 (an open source implementation of the algorithm is available at <https://gitlab.com/rleiva/mnpdt>). The algorithm is based on a breadth-first tree traversal, see for example [7]. The algorithm requires a function called `bestSplit()`, that returns the best way to split a given subset of the training data into two subsets, and a second function called `costFunc()`, that provides a quantitative evaluation of the cost of a tree in terms of complexity and accuracy. The details of these two functions are given in Sections III-C and III-D, respectively. A third function `forecast()` computes the most likely class of a given subset of the training data. Algorithm 1 is based on two nested loops: the external **while** loop keeps a set *Candidates* of the candidate tree nodes (leaves) to grow, whereas the internal **for each** loop finds the best such node from which the tree should be grown further. The latter operation requires to check all possible options and select the one that minimizes the cost of the resulting tree. The exit point in the algorithm is at the end of the **while** loop, where the current tree T is returned if there are no more candidate nodes to further grow the tree, or the nodes in the set *Candidates* generate trees that do not reduce the cost. The nodes of the tree are represented with the *TreeNode* data structure composed of the elements (1) *LChild*, that points to the left child node

¹ $|x|$ denotes the number of symbols, or length, of a string x .

² A programming language is prefix-free if no program can be a prefix of another program.

³ $\langle v, s' \rangle$ denotes the concatenation of the strings v and s' .

Algorithm 1 Minimum Surfeit and Inaccuracy (MSI)

```

1: procedure buildTree(data)
2:   Create TreeNode r; r.LChild, r.RChild, r.Split  $\leftarrow$  none; r.Data  $\leftarrow$  data; r.Class  $\leftarrow$  forecast(data)
3:   Create Tree T; T.Root  $\leftarrow$  r;
4:   bestCost  $\leftarrow$  costFunc(T)
5:   Candidates  $\leftarrow$  {r}
6:   while Candidates  $\neq$   $\emptyset$  do
7:     bestSol  $\leftarrow$  none
8:     for each  $\ell \in$  Candidates do
9:        $\Theta \leftarrow$  bestSplit( $\ell$ .Data)
10:      if  $\Theta = \emptyset$  then Candidates  $\leftarrow$  Candidates  $\setminus$  { $\ell$ } ▷ Discard  $\ell$  as candidate to grow the tree.
11:      else
12:        Create TreeNode  $\ell L$ ,  $\ell R$  ▷ Tree nodes  $\ell L$ ,  $\ell R$  will be the children of  $\ell$ .
13:         $\ell L$ .LChild,  $\ell L$ .RChild,  $\ell L$ .Split,  $\ell R$ .LChild,  $\ell R$ .RChild,  $\ell R$ .Split  $\leftarrow$  none
14:        (dataL, dataR)  $\leftarrow$  split( $\Theta$ ,  $\ell$ .Data)
15:         $\ell L$ .Data  $\leftarrow$  dataL
16:         $\ell R$ .Data  $\leftarrow$  dataR
17:         $\ell L$ .Class  $\leftarrow$  forecast(dataL)
18:         $\ell R$ .Class  $\leftarrow$  forecast(dataR)
19:        ( $\ell$ .LChild,  $\ell$ .RChild,  $\ell$ .Split)  $\leftarrow$  ( $\ell L$ ,  $\ell R$ ,  $\Theta$ ) ▷ Grow the tree at node  $\ell$ .
20:        C  $\leftarrow$  costFunc(T); ▷ C is the cost of the tree if grown at  $\ell$ .
21:        ( $\ell$ .LChild,  $\ell$ .RChild,  $\ell$ .Split)  $\leftarrow$  (none, none, none) ▷ Leave the tree as it was.
22:        if C < bestCost then
23:          bestCost  $\leftarrow$  C
24:          bestSol  $\leftarrow$  ( $\ell$ ,  $\Theta$ ,  $\ell L$ ,  $\ell R$ )
25:        if bestSol  $\neq$  none then
26:          ( $\ell$ ,  $\Theta$ ,  $\ell L$ ,  $\ell R$ )  $\leftarrow$  bestSol
27:          ( $\ell$ .LChild,  $\ell$ .RChild,  $\ell$ .Split)  $\leftarrow$  ( $\ell L$ ,  $\ell R$ ,  $\Theta$ )
28:          Candidates  $\leftarrow$  Candidates  $\setminus$  { $\ell$ }  $\cup$  { $\ell L$ ,  $\ell R$ }
29:        else return T ▷ No candidate reduces the cost of the tree.
30:      return T ▷ No more candidates.

```

in the tree, (2) *RChild*, that points to the right child node, (3) *Data*, that contains the subset of training data of the node, (4) *Split*, which is a pair composed by the attribute and a threshold used for the data splitting, and (5) *Class*, which is the most likely class for this node (based on its *Data*). A tree node could be either an internal node having a split criterion and two children, or a leaf node with a predicted class and no child.

The main difference between our algorithm and other decision tree algorithms is in the **for each** loop. In traditional algorithms, the order in which the branches are evaluated is irrelevant. However, since our algorithm could stop the growing process at any point, at each iteration we explicitly select the best candidate node to grow the tree.

C. SPLITTING CRITERION

A decision tree is an algorithm that recursively partitions the training vectors of \mathcal{X} in such a way that the same values y_i are grouped together. Given a subset $Q \subseteq \mathcal{X}$ we have to find the optimal split for Q . A split is a pair $\theta = (j, w)$ where $1 \leq j \leq m$ is an index, m is the number of attributes, and $w \in \mathbb{R}$ a threshold. A split partitions the set Q into two disjoint subsets $Q_\ell = \{x_i \in Q : x_i \leq w\}$, and $Q_r = Q \setminus Q_\ell$. We use

the *minimal weighted entropy* as the splitting criterion.⁴ The weighted entropy of a split, denoted by \tilde{H} , is defined as:

$$\tilde{H}(Q, \theta) = \frac{d(Q_\ell)}{d(Q)} H(Q_\ell) + \frac{d(Q_r)}{d(Q)} H(Q_r) \quad (4)$$

where $d(S)$ is the number of elements (or diameter) of set S , and $H(S)$ is the entropy of S , i.e.,

$$H(S) = - \sum_{x \in S} p(x) \log p(x), \quad (5)$$

being $p(x)$ the probability of getting x if we select a random element from S . The function $\text{bestSplit}(Q)$ returns the best split θ^* of a given subset of data Q , defined as $\theta^* = \arg \min_{\theta} \tilde{H}(Q, \theta)$. If θ^* does not split Q , $\text{bestSplit}(Q)$ returns \emptyset .

D. COST FUNCTION

For every possible branch to grow, we have to compute how good the resulting tree would be if we add the new nodes, compared to the same tree without them. In this section we introduce a novel cost function to evaluate and compare

⁴Other metrics, like Gini impurity or information gain could be used for the same purpose.

candidate trees. A problem with traditional cost functions for tree evaluation is that they are based on incommensurable quantities, such as the accuracy of the model and the number of nodes. Our goal here is to introduce two new metrics that are conceptually equivalent to the traditional ones while still being commensurable, i.e., based on the same units. We have also designed them to have the same scale, in order to enable a direct comparison. These two metrics will convey how badly the tree fits the data (*inaccuracy*), and how unnecessarily complex the model is (what we call *surfeit*).

Our measure of the inaccuracy of a model M for a dataset \mathcal{X} is the length of the shortest computer program that, given as input the model M , is able to print the dataset \mathcal{X} . We seek to measure how difficult (in terms of the size of a program, not of its running time) it is to fix the errors introduced by the model. Intuitively, in this sense, it is more difficult to fix a model that makes several different mistakes than a model that makes the same mistake several times. Formally, the *inaccuracy* of a dataset \mathcal{X} and a candidate model M is

$$\mathcal{I}(\mathcal{X}, M) = \frac{K(\mathcal{X}|M)}{K(\mathcal{X})}, \quad (6)$$

where $K(\mathcal{X}|M)$ is the conditional Kolmogorov complexity of the dataset \mathcal{X} given the model M . The normalization factor $K(\mathcal{X})$, the Kolmogorov complexity of the dataset \mathcal{X} , is introduced to guarantee that both the inaccuracy and the surfeit have the same scale. We recall that Kolmogorov complexity is a non-computable quantity. Therefore, we approximate $\mathcal{I}(\mathcal{X}, M)$ by the ratio $|Comp(E)|/|Comp(\mathcal{X})|$, where $|Comp(E)|$ is the length of the compressed version of the subset $E \subset \mathcal{X}$ composed by those points that have been misclassified by the tree, and $|Comp(\mathcal{X})|$ is the length of the compressed version of the full dataset.⁵

To assess the model's complexity, we seek to measure the amount of surfeit introduced by our current model with respect to the shortest possible model. Formally, the *surfeit* of a dataset \mathcal{X} and a candidate model M is

$$\mathcal{S}(\mathcal{X}, M) = 1 - \frac{K(\mathcal{X})}{|M|}. \quad (7)$$

The length of the shortest possible model for the dataset \mathcal{X} is given by its Kolmogorov complexity $K(\mathcal{X})$. The normalization factor is given by the length of the current model being evaluated, $|M|$. Formally, exceedingly short models M such that $|M| < K(\mathcal{X})$, are not considered candidate models. Again, because $K(\mathcal{X})$ is a non-computable quantity, in practice we approximate the surfeit of a model by its redundancy, that is, by computing $1 - |Comp(M)|/|M|$, where $|Comp(M)|$ is the length of a compressed version of a string describing the tree M . We remark that the shortest possible model M for a dataset \mathcal{X} must be incompressible, but an incompressible model for \mathcal{X} is not necessarily the shortest possible one. That is, it might happen that a model is not redundant but still presents some surfeit.

⁵In our implementation the dataset \mathcal{X} is stored in a numpy array. The subset E is the numpy array composed of those rows of \mathcal{X} for which the model provides a wrong prediction.

Both quantities, inaccuracy and surfeit, have to be combined into a single value by applying a function

$$\mathcal{N}(\mathcal{X}, M) = g(\mathcal{I}(\mathcal{X}, M), \mathcal{S}(\mathcal{X}, M)). \quad (8)$$

Since both the inaccuracy and the surfeit are relative quantities, a natural way to combine them is via their harmonic mean. Other candidate functions (arithmetic mean, geometric mean, Euclidean distance, product, and addition) will be evaluated in the following.

IV. PRACTICAL IMPLEMENTATION

The abstract concept of Kolmogorov complexity is usually approximated in practice with compression algorithms [27]. In this work we have tested three different algorithms: the Lempel-Ziv-Markov chain algorithm, LZMA [29] that uses dictionaries for compression, zlib [30] that combines dictionaries with Huffman encodings, and bzip2 a compressor based on the Burrows-Wheeler transform [31]. All compressors have been configured to use the maximum compression level allowed, in order to avoid problems due to small window size buffers [32].

For the representation of a tree as a string we use the following template:

```
def tree{[attrs]}:
    if [attr] <= [thresh]:
        return [label] || [subtree]
    else:
        return [label] || [subtree]
```

where $[attrs]$ is the list of the only attributes used in the model,⁶ $[attr]$ is a single attribute represented by the letter X followed by a number (e.g. X1), $[thresh]$ is the threshold used for the split, $[label]$ is one of the valid labels from the set \mathcal{G} , and $|| [subtree]$ means that the return statement can be replaced by another level of $[if - else]$ conditions. We could have used a much shorter description of trees by replacing word tokens with symbols, e.g., via the ternary conditional operators $?$ and $:$ used in modern programming languages, or by dropping the return statement. This would produce shorter trees, but the complexity of the models would remain the same, up to an additive constant that does not depend on the models themselves. Since, e.g., the harmonic mean compares relative values instead of absolute ones, this additive constant can be safely ignored.

V. RESULTS

In this section we evaluate our new algorithm, and compare its performance against the well-known algorithms CART, C4.5, CHAID, and PUBLIC. In order to measure the prediction performance of the tree classifiers, we used the mean accuracy, defined as one minus the mean error. Other performance parameters taken into account include the total size of the tree, and the maximum depth.

⁶If the dataset contains many attributes, listing all of them when dealing with very short models would make the length of the model's header greater than the length of the body.




FIGURE 1. Synthetic dataset with splits computed by CART.




FIGURE 2. Decision tree obtained by our algorithm.

Figure 1 shows a synthetic dataset consisting of 100 random points lying on a two-dimensional plane, where all the points with an X_1 attribute smaller than 50 are colored blue, and the rest as red. We artificially introduce a red point, simulating a measurement error, in the blue area. The gray lines correspond to the classification performed by the CART algorithm, as implemented by the scikit-learn toolkit [33]. CART will not stop until all the points have been correctly classified, so we require to have at least five points in order to make a node split, a de facto standard suggested, for example, in Section 9.2.2 of [2] (for a more theoretical justification of this minimum value see the expected count condition in [34]).

The tree obtained by applying our algorithm to Figure 1 can be seen in Figure 2. Our algorithm does not try to model the error point, since the gain due to an increment in the accuracy does not compensate the redundancy introduced in the model. Recall that our algorithm stops growing the tree when the cost function, based on the measures of inaccuracy and surfeit, does not decrease when adding new nodes to the tree. Our algorithm presents a lower sensitivity to the errors found in datasets, at least if the number of errors is small compared to the number of valid points.

A second experiment with synthetic datasets is depicted in Figure 3. There, we create two isotropic Gaussian blobs that partially overlap. We start with a standard deviation of 2.5 for each cluster, so they are easy to separate, and increase the standard deviation in increments of 0.01, until we reach 4.5, which leads to significant overlap. For each value of the standard deviation, we run the experiment 100 times and we compute the average accuracy for the two algorithms using different datasets for training and testing. For this experiment, the hyperparameter “minimum number of samples per leaf




FIGURE 3. Example of isotropic Gaussian blobs (std. dev. = 3).




FIGURE 4. Accuracy for isotropic Gaussian blobs.

node” of the CART algorithm has been optimized in order to achieve the maximum accuracy (in this case, the best value was obtained with a minimum size of 26 samples). The results of this experiment are shown in Figure 4. On average, our algorithm provides the same accuracy (0.872) as the optimized version of the CART algorithm.

For each iteration of the experiment, we have also computed the average number of nodes, including internal and leaf nodes, required by the models to properly classify the clouds in the dataset. The results of this measurement are show in Figure 5. Our algorithm requires an average of 5.7 nodes compared to 23 nodes for the CART algorithm. Moreover, our algorithm is more stable than CART, in the sense that it produces models of similar complexity when it gets similar input datasets (a standard deviation of 0.3 compared to 3.9 for CART).

In Figure 6 we show the maximum depth of the tree, defined as the longest path from the root of the tree to any of its leaves. The maximum depth of the tree is a good measure of the maximum time it will require for the model to provide a classification. Our algorithm has an average depth of 2.2 nodes, whereas the average depth yielded by the CART algorithm is 4.8 nodes. We emphasize that the CART algorithm requires to optimize a hyperparameter in order to obtain these optimal results, whereas




FIGURE 5. Number of tree nodes for isotropic Gaussian blobs.

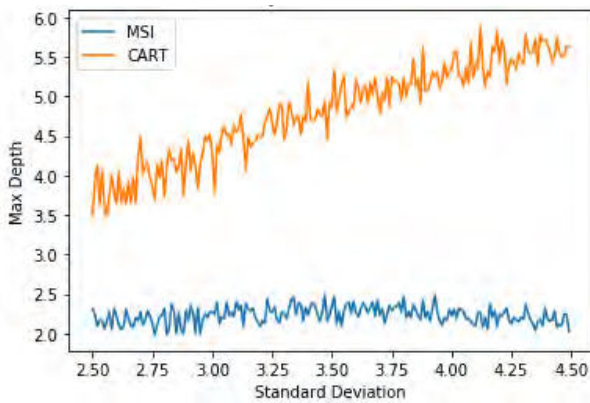


FIGURE 6. Maximum tree depth for isotropic Gaussian blobs.

the algorithm we propose, by design, does not require this optimization.

Before proceeding, we would like to stress that neither the cost function nor the compressor are hyperparameters. We demonstrate this by showing that the choice of different cost functions and compressors leads to substantially similar results. In Figure 7, we apply our algorithm to the discussed partially overlapping isotropic Gaussian blobs described above (and illustrated with Figure 3), and evaluate different alternatives for the definition of the cost function $\mathcal{N}(\mathcal{X}, M)$ in (8): arithmetic mean $(\mathcal{I}(\mathcal{X}, M) + \mathcal{S}(\mathcal{X}, M))/2$, geometric mean $(\mathcal{I}(\mathcal{X}, M) \times \mathcal{S}(\mathcal{X}, M))^{1/2}$, harmonic mean $2/(\mathcal{I}(\mathcal{X}, M)^{-1} + \mathcal{S}(\mathcal{X}, M)^{-1})$, Euclidean distance $(\mathcal{I}(\mathcal{X}, M)^2 + \mathcal{S}(\mathcal{X}, M)^2)^{1/2}$, sum $\mathcal{I}(\mathcal{X}, M) + \mathcal{S}(\mathcal{X}, M)$, and product $\mathcal{I}(\mathcal{X}, M) \times \mathcal{S}(\mathcal{X}, M)$. The figure shows the extremely limited difference between the different functions (for an easier interpretability of the results, the mean has been subtracted from the values). Similarly, Figure 8 shows the performance of our algorithm when using the LZMA, zlib, and bz2 compressors. We observe that all of them yield similar performance. The above results suggest that the performance our algorithm is independent of the specific choice made for either implementation aspect.

In the following, we employ the bz2 as the compressor to approximate the computation of the Kolmogorov complex-

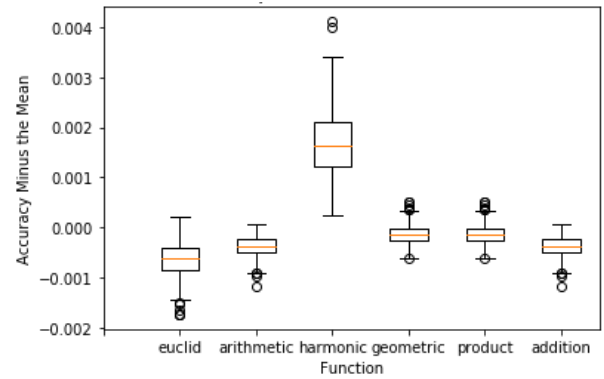


FIGURE 7. Evaluation of cost functions.

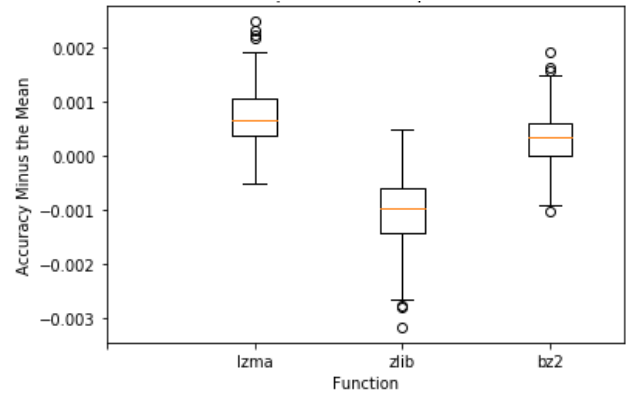


FIGURE 8. Evaluation of different compressors.

ity, and the harmonic mean as the cost function. The latter has the additional advantage of comparing relative values instead of absolute ones, and therefore makes it possible to get rid of additive constants related to the choice of a specific template for the representation of the tree, as discussed in Section IV.

Finally, we have compared the performance of our algorithm with the performance of four popular decision trees algorithms, CART, C4.5, CHAID, and PUBLIC, with a collection of real datasets. More specifically, we have selected 12 well known datasets from the UCI Machine Learning Repository [35], among those with the largest amount of data. The selected ones are: diagnosis of breast cancer (*cancer*), optical recognition of handwritten digits (*digits*), predicting protein localization sites in gram-negative bacteria (*yeast*), classification of NASA space shuttle data (*shuttle*), classification of blocks in web pages (*page*), segmentation of outdoor images (*image*), predicting the age of abalones from physical measurements (*abalone*), predicting the quality of red and white variants of Portuguese wine (*wine*) [36], filter spam emails (*spam*), wall-following robot navigation (*wall*), classification of land use based on Landsat satellite images (*landsat*), and distinguishing signals from background noise in the MAGIC gamma telescope images (*magic*). For each dataset, we have repeated

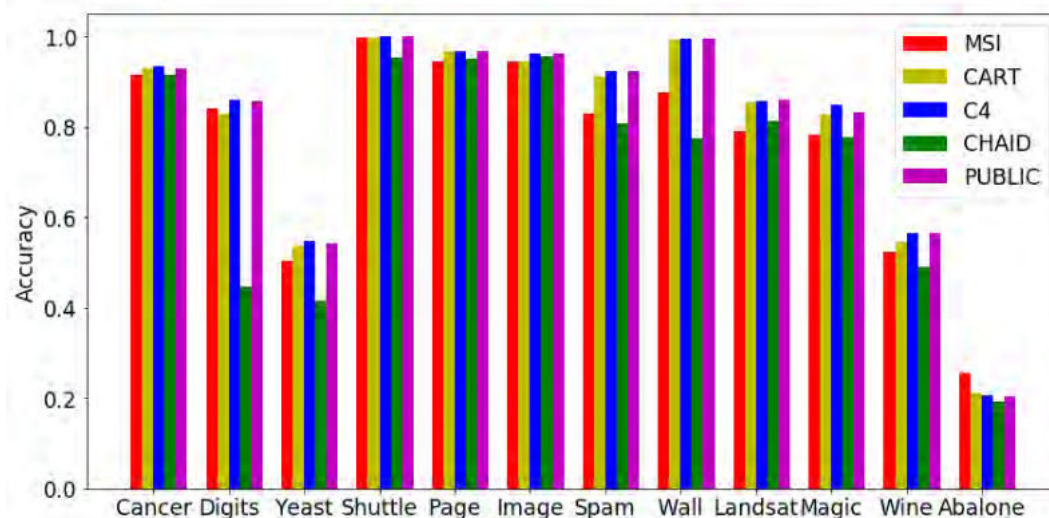


FIGURE 9. Accuracy over example datasets.

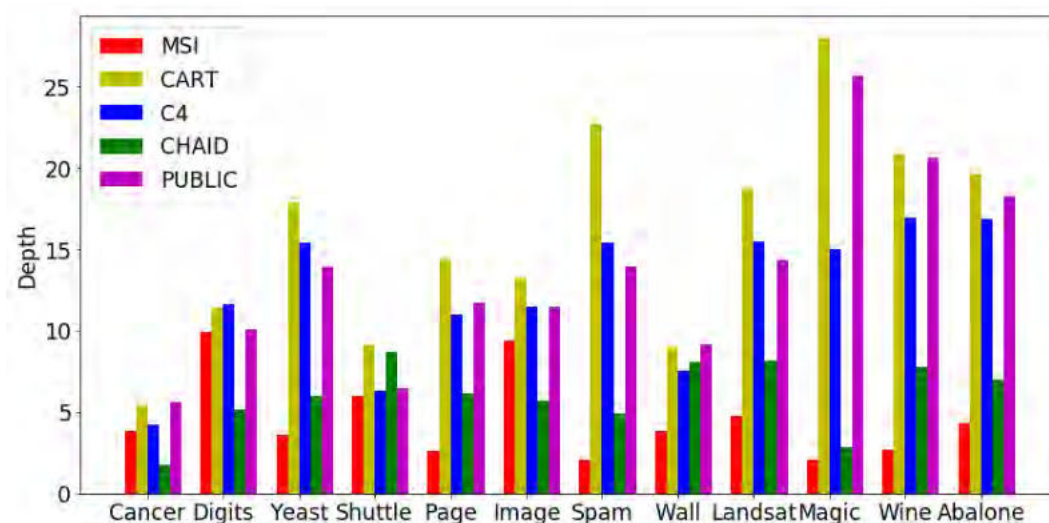


FIGURE 10. Depth of models over example datasets.

the experiment 100 times,⁷ by randomly selecting the training (70%) and testing (30%) subsets at each iteration.

In Figure 9 we compare the accuracy of the resulting models obtained by applying CART, C4.5, CHAID, as well as our own new algorithm to the above real datasets. In 3 of the 12 datasets, our algorithm provides better accuracy than CART. In the remaining 9 cases, the accuracy is fully comparable, and less than 2% smaller than the best of the other algorithms, on average. In Figure 11 we provide a comparison of the depth of the resulting models. Our algorithm yields a shallower tree than all other algorithms, except for datasets *Cancer*, *Digits*, and *Image*, confirming the non-overfitting properties of our tree construction process. In Figure 10, we show a comparison of the number of nodes of the resulting models. Only for one of the datasets (*Digits*),

⁷In order to limit the computational complexity, the number of repeated experiments was limited to less than 100 for the CHAID algorithm when applied to some particularly heavy datasets.

does our model produce a slightly larger tree than those generated by CART, C4.5, CHAID and PUBLIC. In most of the cases, the trees generated by our algorithm have one half (worst case) up to about two orders of magnitude fewer nodes (best case) than the other algorithms considered in our evaluation.

Finally, we have compared our algorithm with a post-processing pruning algorithm applied to the tree obtained by CART. In particular, we have applied a cost-complexity pruning guided by a cross-validation to the *Shuttle* dataset, as implemented by the *rpart* library in R. The optimal value of the cost-complexity metric is achieved for a tree of 59 nodes, whereas our algorithm, without requiring the optimization of any hyperparameters, obtains a tree of 28 nodes. Both trees achieve the same accuracy.

VI. LESSONS LEARNED

The algorithm proposed in this paper has been designed to find a compact model that describes well a dataset

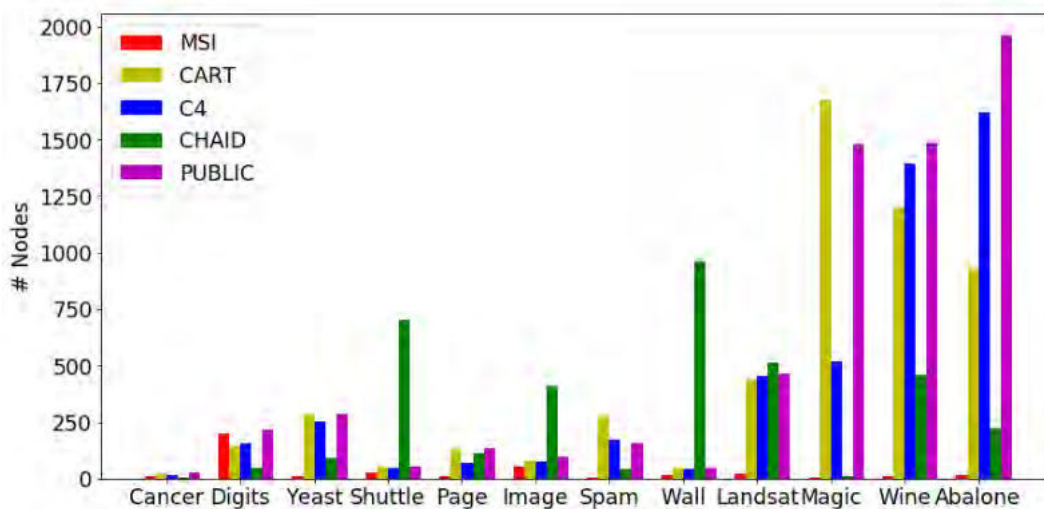


FIGURE 11. Number of nodes over example datasets.

(high accuracy) without over-fitting the data. The experiment described in Figure 1 suggests that our algorithm has lower sensitivity than CART to the errors found in datasets, at least if the number of errors is small compared with the number of valid points. The experiment of Figure 3 shows that the algorithm tends to produce simpler models when the classes that compose the dataset are not linearly separable. These two situations, errors and non-linearity, are common causes of model overfitting when using decision trees. In general, the CART algorithm produces much more complex models than our algorithm in those situations, even when configured to avoid overfitting as much as possible.

As we can see in Figures 9, 10 and 11, our new algorithm produces trees with a significantly smaller number of nodes (and depth) than those produced with standard algorithms used in practice, namely CART, C4, CHAID and PUBLIC, without significantly decreasing the accuracy. The experiments have been performed with a collection of datasets resulting from real experiments in order to test the applicability of the new algorithm to practical problems. Given the size of the trees produced, our algorithm is the ideal method to apply in those cases where the interpretability of the results is critical, or where there is a large risk of model overfitting.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a new algorithm to build decision trees based on the compressibility of candidate models and the compressibility of the errors generated by those models. The main advantage of the new algorithm is that it does not overfit the training data by design, and additionally it does not require external or ad-hoc procedures to control the overoptimization of the produced model. Moreover, our algorithm does not require the optimization of any hyperparameters. Experimental validation with synthetic and real datasets demonstrates that the accuracy of the new algorithm is similar to the accuracy of well-known decision tree algorithms,

like CART, C4.5, CHAID and PUBLIC. Our results show that the proposed algorithm produces models with a considerably smaller number of nodes, without any substantial accuracy decrease.

Future work includes extending the cost function to include a goodness measure for the attributes used at the nodes of the trees, e.g., based on how correlated the values of these attributes are with the target values. The early stopping properties of our algorithm can also be applied to other machine learning techniques, e.g., to find optimal deep neural network architectures.

REFERENCES

- [1] L. Rokach and O. Z. Maimon, *Data Mining with Decision Trees: Theory and Applications*, vol. 81. Singapore: World Scientific, 2014.
- [2] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning* (Springer Series in Statistics), vol. 1. New York, NY, USA: Springer, 2001.
- [3] M. Wu, M. C. Hughes, S. Parbhoo, M. Zazzi, V. Roth, and F. Doshi-Velez, "Beyond sparsity: Tree regularization of deep models for interpretability," 2017, *arXiv:1711.06178*. [Online]. Available: <https://arxiv.org/abs/1711.06178>
- [4] M. Bohanec and I. Bratko, "Trading accuracy for simplicity in decision trees," *Mach. Learn.*, vol. 15, no. 3, pp. 223–250, Jun. 1994.
- [5] B. D. Ripley, *Pattern Recognition and Neural Networks*. Cambridge, U.K.: Cambridge Univ. Press, 1996.
- [6] W.-Y. Loh, "Fifty years of classification and regression trees," *Int. Stat. Rev.*, vol. 82, no. 3, pp. 329–348, Dec. 2014.
- [7] T. H. Cormen, *Introduction to Algorithms*. Cambridge, MA, USA: MIT Press, 2009.
- [8] J. R. Quinlan, "Induction of decision trees," *Mach. Learn.*, vol. 1, no. 1, pp. 81–106, 1986.
- [9] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. New York, NY, USA: Kluwer, 1984.
- [10] G. V. Kass, "An exploratory technique for investigating large quantities of categorical data," *Appl. Statist.*, vol. 29, no. 2, pp. 119–127, 1980.
- [11] J. R. Quinlan, *C4.5: Programs for Machine Learning*, vol. 38. San Mateo, CA, USA: Morgan Kaufmann, 1993, p. 48.
- [12] W.-Y. Loh and Y.-S. Shih, "Split selection methods for classification trees," *Statistica Sinica*, vol. 7, no. 4, pp. 815–840, Oct. 1997.
- [13] J. H. Friedman, "Multivariate adaptive regression splines," *Ann. Statist.*, vol. 19, no. 1, pp. 1–67, 1991.
- [14] R. Rastogi and K. Shim, "Public: A decision tree classifier that integrates building and pruning," in *Proc. VLDB*, vol. 98, Aug. 1998, pp. 24–27.
- [15] D. Ignatov and A. Ignatov, "Decision stream: Cultivating deep decision trees," in *Proc. IEEE 29th Int. Conf. Tools Artif. Intell. (ICTAI)*, Nov. 2017, pp. 905–912.

- [16] M. A. Carreira-Perpinán and P. Tavallari, "Alternating optimization of decision trees, with application to learning sparse oblique trees," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 1219–1229.
- [17] T.-S. Lim, W.-Y. Loh, and Y.-S. Shih, "A comparison of prediction accuracy, complexity, and training time of thirty-three old and new classification algorithms," *Mach. Learn.*, vol. 40, no. 3, pp. 203–228, 2000.
- [18] C. Zhang, C. Liu, X. Zhang, and G. Alpanidis, "An up-to-date comparison of state-of-the-art classification algorithms," *Expert Syst. Appl.*, vol. 82, pp. 128–150, Oct. 2017.
- [19] R. J. Solomonoff, "A formal theory of inductive inference. Part I," *Inf. Control*, vol. 7, no. 1, pp. 1–22, Mar. 1964.
- [20] A. N. Kolmogorov, "Three approaches to the quantitative definition of information," *Problems Inf. Transmiss.*, vol. 1, no. 1, pp. 1–7, 1965.
- [21] G. J. Chaitin, "On the simplicity and speed of programs for computing infinite sets of natural numbers," *J. ACM*, vol. 16, no. 3, pp. 407–422, Jul. 1969.
- [22] P. D. Grünwald, *The Minimum Description Length Principle*. Cambridge, MA, USA: MIT Press, 2007.
- [23] C. S. Wallace, *Statistical and Inductive Inference by Minimum Message Length*. New York, NY, USA: Springer, 2005.
- [24] J. R. Quinlan and R. L. Rivest, "Inferring decision trees using the minimum description length principle," *Inf. Comput.*, vol. 80, no. 3, pp. 227–248, Mar. 1989.
- [25] C. S. Wallace and J. Patrick, "Coding decision trees," *Mach. Learn.*, vol. 11, no. 1, pp. 7–22, Apr. 1993.
- [26] M. Mehta, J. Rissanen, and R. Agrawal, "MDL-based decision tree pruning," in *Proc. KDD*, Aug. 1995, vol. 21, no. 2, pp. 216–221.
- [27] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications* (Texts in Computer Science), vol. 9. New York, NY, USA: Springer, 2008.
- [28] K. Miettinen, *Nonlinear Multiobjective Optimization*, vol. 12. New York, NY, USA: Springer, 2012.
- [29] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans. Inf. Theory*, vol. 23, no. 3, pp. 337–343, May 1977.
- [30] P. Deutsch, "Deflate compressed data format specification version 1.3," IETF, Fremont, CA, USA, Tech. Rep. RFC 1951, May 1996.
- [31] M. Burrows and D. J. Wheeler, "A block-sorting lossless data compression algorithm," *Digit. Syst. Res. Center, Palo Alto, CA, USA, Tech. Rep. 124*, 1994.
- [32] M. Cebrián, M. Alfonso, and A. Ortega, "Common pitfalls using the normalized compression distance: What to watch out for in a compressor," *Commun. Inf. Syst.*, vol. 5, no. 4, pp. 367–384, 2005.
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, and J. Vanderplas, "Scikit-learn: Machine learning in python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [34] D. S. Starnes, D. Yates, and D. S. Moore, *The Practice of Statistics*. New York, NY, USA: Macmillan, 2010.
- [35] D. Dheeru and E. K. Taniskidou. (2017). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [36] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis, "Modeling wine preferences by data mining from physicochemical properties," *Decis. Support Syst.*, vol. 47, no. 4, pp. 547–553, 2009.



RAFAEL GARCÍA LEIVA received the bachelor's degree in computer science from the University of Córdoba, Spain, the master's degree in computational sciences from the University of Amsterdam, and the Diploma degree in advanced studies in telematics from the Universidad Autónoma de Madrid. He has worked for four years at the University of Córdoba as a Research Assistant and a Scientific Programmer in geographical information systems and remote sensing.

He further served for three years in the Universidad Autónoma de Madrid as a Research Engineer in high energy physics. He has worked during three years as a Research and Development Manager with Andago Ingeniería, coordinating the Research and Development activities of the company. In 2008, he funded Entropy Computational Services, where he worked in social networks, mobile applications, and quantitative trading over a period of five years. In 2014, he joined the IMDEA Networks Institute as a Research Engineer, working in big data and machine learning.



ANTONIO FERNÁNDEZ ANTA received the bachelor's degree from Universidad Politécnica de Madrid (UPM), having received awards at the university and national levels for his academic performance, and the M.Sc. and Ph.D. degrees from the University of Southwestern Louisiana in 1992 and 1994, respectively. He was a Full Professor with the Universidad Rey Juan Carlos (URJC) and was on the Faculty of UPM, where he received an award for his research productivity. He held a postdoctoral position with MIT from 1995 to 1997, and spent sabbatical years at Bell Labs, Murray Hill, and MIT Media Lab. He is currently a Research Professor with IMDEA Networks Institute. He has more than 25 years of research experience, and more than 200 scientific publications. He is a Senior Member of ACM. He was the Chair of the Steering Committee of DISC and has served in the TPC for numerous conferences and workshops.



VINCENZO MANCUSO received the M.Sc. and Ph.D. degrees in electronics, computer science, and telecommunications from the University of Palermo, Italy. He was with INRIA, Sophia Antipolis, France; Rice University, Houston, TX, USA; and the University of Palermo. He is currently a Research Associate Professor with the IMDEA Networks Institute, Madrid, Spain. He is currently working on analysis and optimization of wireless access networks and on the measurements and assessment of mobile broadband networks. He has authored more than 110 peer-reviewed publications focusing on Internet QoS and on the analysis, design, and experimental evaluation of opportunistic and adaptive protocols and architectures for wireless networks. He was a recipient of the Ramon y Cajal Research Grant from the Spanish Ministry of Science and Innovation.



PAOLO CASARI received the Ph.D. degree in information engineering from the University of Padova, Italy, in 2008. He was on leave at the Massachusetts Institute of Technology, in 2007, working on underwater communications and networks. He collaborated to several funded projects including CLAM (FP7), RACUN (EDA), and several US ARO, ONR, and NSF initiatives. He is the Principal Investigator of the NATO SPS project Threat-Detect and the Scientific Coordinator of the EU H2020 RECAP Project. In 2015, he joined the IMDEA Networks Institute, Madrid, Spain, where he leads the Ubiquitous Wireless Networks Group. His research interests include many aspects of networked communications, such as channel modeling, network protocol design, localization, simulation, and experimental evaluations. He received two best paper awards. He regularly serves in the organizing committee for several international conferences. He serves on the Editorial Boards of the IEEE TRANSACTIONS ON MOBILE COMPUTING and of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, and has been a Guest Editor of a special issue of IEEE ACCESS on Underwater Acoustic Communications and Networking, and of an earlier special issue of the *Hindawi Journal of Electrical and Computer Engineering* on the same topic.

...