Journal of Machine Learning Research 20 (2019) 1-34 Submitted 1/18; Revised 6/19; Published 6/19

Complete Search for Feature Selection in Decision Trees

Salvatore Ruggieri SALVATORE.RUGGIERI@UNIPI.IT
Department of Computer Science
University of Pisa
Largo B. Pontecorvo 3, 56127, Pisa, Italy

Editor: Inderjit Dhillon

Abstract

The search space for the feature selection problem in decision tree learning is the lattice of subsets of the available features. We design an exact enumeration procedure of the subsets of features that lead to *all and only the* distinct decision trees built by a greedy top-down decision tree induction algorithm. The procedure stores, in the worst case, a number of trees linear in the number of features. By exploiting a further pruning of the search space, we design a complete procedure for finding δ-acceptable feature subsets, which depart by at most δ from the best estimated error over any feature subset. Feature subsets with the best estimated error are called *best feature subsets*. Our results apply to any error estimator function, but experiments are mainly conducted under the wrapper model, in which the misclassification error over a search set is used as an estimator. The approach is also adapted to the design of a computational optimization of the sequential backward elimination heuristic, extending its applicability to large dimensional datasets. The procedures of this paper are implemented in a multi-core data parallel C++ system. We investigate experimentally the properties and limitations of the procedures on a collection of 20 benchmark datasets, showing that oversearching increases both overfitting and instability.

Keywords: Feature Selection, Decision Trees, Wrapper models, Complete Search

1. Introduction

Feature selection is essential for optimizing the accuracy of classifiers, for reducing the data collection effort, for enhancing model interpretability, and for speeding up prediction time (Guyon et al., 2006b). In this paper, we will consider decision tree classifiers $DT(S)$ built on a subset S of available features. Our results will hold for any top-down tree induction algorithm that greedily selects a split attribute at every node by maximizing a quality measure. The well-known C4.5 (Quinlan, 1993) and CART systems (Breiman et al., 1984) belong to this class of algorithms. Although more advanced learning models achieve better predictive performance (Caruana and Niculescu-Mizil, 2006; Delgado et al., 2014), decision trees are worth investigation, because they are building-blocks of the advanced models, e.g., random forests, or because they represent a good trade-off between accuracy and interpretability (Guidotti et al., 2018; Huysmans et al., 2011).

In this paper, we will consider the search for a *best feature subset* S, i.e., such that the estimated error $err(DT(S))$ on $DT(S)$ is minimum among all possible feature subsets. The size of the lattice of feature subsets is exponential in the number of available features. The complete search of the lattice is known to be an NP-hard problem (Amaldi and Kann,

1998). For this reason, heuristic searches are typically adopted in practice. For instance, the sequential backward elimination (**SBE**) heuristic starts with all features and repeatedly eliminates one feature at a time while error estimation does not increase. However, complete strategies do not have to be exhaustive. In particular, feature subsets that lead to duplicate decision trees can be pruned from the search space. A naïve approach that stores all distinct trees found during the search is, however, unfeasible, since there may be an exponential number of such trees. *Our first contribution* is a non-trivial enumeration algorithm **DTdistinct** of all distinct decision trees built using subsets of the available features. The procedure requires the storage of a linear number of decision trees in the worst case. The starting point is a recursive procedure for the visit of the lattice of all subsets of features. The key idea is that a subset of features is denoted by the union $R \cup S$ of two sets, where elements in R must *necessarily* be used as split attributes, and elements in S may be used or not. Pruning of the search space is driven by the observation that if a feature $a \in S$ is not used as split attribute by a decision tree built on $R \cup S$, then the feature subset $R \cup S \setminus \{a\}$ leads to the same decision tree. Duplicate decision trees that still pass such a (necessary but not sufficient) pruning condition can be identified through a test on whether or not they use all features in R. An intruiguing contribution of this paper consists in a specific order of visit of the search space, for which a negligible fraction of trees actually built are duplicates.

Enumeration of distinct decision trees can be used for finding the best feature subsets with reference to an error estimation function. Our results will hold for any error estimation function $err(DT(S))$. In experiments, we mainly adhere to the wrapper model (John et al., 1994; Kohavi and John, 1997), and consider the misclassification error on a search set that is not used for building the decision tree. The wrapper model for feature selection has shown superior performance in many contexts (Doak, 1992; Bolón-Canedo et al., 2013). We introduce the notion of a δ-acceptable feature subset, which leads to a decision tree with an estimated error that departs by at most δ from the minimum estimated error over any feature subset. *Our second contribution* is a complete search procedure $\mathbf{DTaccept}_\delta$ of δ-acceptable and best (for $\delta = 0$) feature subsets. The search builds on the enumeration of distinct decision trees. It relies on a key pruning condition that is a conservative extension of the condition above. If, for a feature $a \in S$, we have that $err(DT(R \cup S)) \leq \delta + err(DT(R \cup S \setminus \{a\}))$, then $R \cup S \setminus \{a\}$ can be pruned from the search with the guarantee of only missing decision trees whose error is at most δ from the best error of visited trees. Hence, visited feature subsets include acceptable ones.

Coupled with the tremendous computational optimization and multi-core parallelization of greedy decision tree induction algorithms, our approach makes it possible to increase the limit of practical applicability of theoretically hard complete searches. We show experimentally that the best feature subsets can be found in reasonable time for up to 60 features for small-sized datasets. Beyond such a limit, even heuristic approaches may require a large amount of time. We devise a white-box implementation **DTsbe** of **SBE**, specific for greedy decision tree algorithms, that exploits some of the pruning and computational optimization ideas. *Our third contribution* consists of a white-box optimization of **SBE** which extends its applicability to large dimensional datasets, and which exhibits, for medium and low dimensional datasets, a computational speedup of up to $100\times$.

Both **DTaccept**$_\delta$ and **DTsbe** are implemented in a multi-core data parallel C++ system, which is made publicly available. We report experiments on 20 benchmark datasets of small-to-large dimensionality. Results confirm previous studies that oversearching increases overfitting. In addition, they also highlight that oversearching increases instability, namely variability of the subset of selected features due to perturbation of the training set. Moreover, we show that sequential backward elimination can improve the generalization error of random forests for medium to large dimensional datasets. Such an experiment is made possible only thanks to the computational speedup of **DTsbe** over **SBE**.

This paper is organized as follows. First, we recall related work in Section 2. The visit of the lattice of feature subsets is based on a generalization of binary counting enumeration of subsets devised in Section 3. Next, Section 4 introduces a procedure for the enumeration of distinct decision trees as a pruning of the feature subset lattice. Complete search of best and acceptable feature subset is then presented in Section 5. Optimization of the sequential backward elimination heuristic is discussed in Section 6. Experimental results are presented in Section 7, with additional tables reported in Appendix A. Finally, we summarize the contribution of the paper in the conclusions.

2. Related Work

Blum and Langley (1997); Dash and Liu (1997); Guyon and Elisseeff (2003); Liu and Yu (2005); Bolón-Canedo et al. (2013) provide a categorization of approaches of feature subset selection along the orthogonal axes of the evaluation criteria, the search strategies, and the machine learning tasks. Common evaluation criteria include filter models, embedded approaches, and wrapper approaches. *Filters* are pre-processing algorithms that select a subset of features by looking at the data distribution, independently from the induction algorithm (Cover, 1977). *Embedded* approaches perform feature selection in the process of training and are specific to the learning algorithm (Lal et al., 2006). *Wrappers* approaches optimize induction algorithm performances as part of feature selection (Kohavi and John, 1997). In particular, training data is split into a building set and a search set, and the space of feature subsets is explored. For each feature subset considered, the building set is used to train a classifier, which is then evaluated on the search set. Search space exploration strategies include (Doak, 1992): *hill-climbing* search (forward selection, backward elimination, bidirectional selection, beam search, genetic search), *random* search (random start hill-climbing, simulated annealing, Las Vegas), and *complete* search. The aim of complete search is to find a feature subset that optimizes an evaluation metric. Typical objectives include minimizing the size of the feature subset provided that the classifier built from it has an accuracy greater or equal to a given threshold (*dimensionality reduction*), or minimizing the empirical misclassification error of the classifier on the search set (*performance maximization*). Finally, feature subset selection has been considered for classification, regression, and clustering tasks. Machine learning models and algorithms can be either treated as *black-boxes* or, instead, feature selection methods can be specific to the model and/or algorithm at hand (*white-box*). White-box approaches are less general, but can exploit assumptions on the model or algorithm to direct and speed up the search. For instance, the best k-subset problem for linear regression (Miller, 2002) smoothly generalizes the linear re-

gression problem to find out the subset of up to k features that best predict an independent variable.

Only complete space exploration can provide the guarantee of finding best feature subsets with respect to a given error estimation function. Several estimators have been proposed in the literature, including: the empirical misclassification error on the training set or in the search dataset; estimators adopted for tree simplification (Breslow and Aha, 1997; Esposito et al., 1997); bootstrap and cross-validation (Kohavi, 1995; Stone, 1997); and the recent *jeff* method (Fan, 2016), which is specific to decision tree models. Heuristic search approaches can lead to results arbitrarily worse than the best feature subset (Murthy, 1998). Complete search is known to be NP-hard (Amaldi and Kann, 1998). However, complete strategies do not need to be exhaustive in order to find a best feature subset. For instance, filter models can rely on monotonic evaluation metrics to support Branch & Bound search (Liu et al., 1998). Regarding wrapper approaches,the empirical misclassification error lacks the monotonicity property that would allow for pruning the search space in a complete search. Approximate Monotonicity with Branch & Bound (AMB&B) (Foroutan and Sklansky, 1987) tries and tackles this limitation, but it provides no formal guarantee that a best feature subset is found. Another form of search space pruning in wrapper approaches for decision trees has been pointed out by Caruana and Freitag (1994), who examine five hillclimbing procedures. They adopt a caching approach to prevent re-building duplicate decision trees. The basic property they observe is reported in a generalized form in this paper as Remark 6. While caching improves on the efficiency of a limited search, in the case of a complete search, it requires an exponential number of decision trees to be stored in cache, while our approach requires a linear number of them. We will also observe that Remark 6 may still leave duplicate trees in the search space, i.e., it is a necessary but not sufficient condition for enumerating distinct decision trees, while we will provide an exact enumeration and, in addition, a further pruning of trees that cannot lead to best/acceptable feature subsets.

A problem related to the focus of this paper regards the construction of optimal decision trees using non-greedy algorithms. In such a problem, the structure of a decision tree and the split attributes are determined at once as a global optimization problem. Bertsimas and Dunn (2017) and Menickelly et al. (2016); Verwer and Zhang (2017) formulate tree induction as a mixed-integer optimization problem and as an integer programming problem respectively. The optimization function is the misclassification error on the training set, possibly regularized with respect to decision tree size. Other approaches, e.g. Narodytska et al. (2018), encode tree construction as a constraint solving problem, with the aim of minimizing tree size. The search space of the optimal decision tree problem is larger than in the best feature subset problem. The former is exponential in the product of the number of features and the maximal tree depth, while the latter is exponential only in the number of features. An optimal decision tree may not be produceable by a fixed greedy algorithm, for any feature subset.

This paper significantly extends the preliminary results that appeared in Ruggieri (2017) in several directions. First, it improves on the enumeration procedure of distinct decision trees. The new ordering of visits in the search space has a clear theoretical justification, and an overhead (duplicated trees built) close to zero for all experimental datasets. Second, the paper introduces the notion of δ-acceptable feature subsets, which depart from best feature

subsets by at most δ in estimated error, and a novel algorithm that further prunes the enumeration of distinct decision trees to find out δ-acceptable feature subsets. Moreover, our approach applies to any error estimation function. Third, the experimental section (and an appendix with additional tables) now includes a comprehensive set of results on a larger collection of benchmark datasets. Fourth, the implementation of all proposed algorithms is now multi-core parallel, and it is publicly available. It reaches run-time efficiency improvements of up to 7× on an 8-core computer.

3. Enumerating Subsets

Let $S = \{a_1, \ldots, a_n\}$ be a set of n elements, with $n \geq 0$. The powerset of S is the set of its subsets: $Pow(S) = \{S' \mid S' \subseteq S\}$. There are 2^n subsets of S, and, for $0 \leq k \leq n$, there are $\binom{n}{k}$ subsets of size k. Figure 1 (left) shows the lattice (w.r.t. set inclusion) of subsets for $n = 3$. The order of visit of the lattice, or, equivalently, the order of enumeration of elements in $Pow(S)$, can be of primary importance for problems that explore the lattice as a search space. Well-known algorithms for subset generation produce lexicographic ordering, Grey code ordering, or binary counting ordering (Skiena, 2008). Binary counting maps each subset into a binary number with n bits by setting the i^{th} bit to 1 iff a_i belongs to the subset, and generating subsets by counting from 0 to $2^n - 1$. Subsets for $n = 3$ are generated as $\{\}$, $\{a_3\}$, $\{a_2\}$, $\{a_2, a_3\}$, $\{a_1\}$, $\{a_1, a_3\}$, $\{a_1, a_2\}$, $\{a_1, a_2, a_3\}$. In this section, we introduce a recursive algorithm for a generalization of reverse binary counting (namely, counting from $2^n - 1$ down to 0) that will be the building block for solving the problem of generating distinct decision trees. Let us start by introducing the notation $R \bowtie P = \cup_{S' \in P}\{R \cup S'\}$ to denote sets obtained by the union of R with elements of P. In particular:

$$R \bowtie Pow(S) = \cup_{S' \subseteq S}\{R \cup S'\}$$

consists of the subsets of $R \cup S$ that *necessarily* include R. This generalization of powersets will be crucial later on when we have to distinguish predictive attributes that *must* be used in a decision tree from those that *may* be used. A key observation of binary counting is that subsets can be partitioned between those including the value a_1 and those not including it. For example, $Pow(\{a_1, a_2, a_3\}) = (\{a_1\} \bowtie Pow(\{a_2, a_3\})) \cup (\emptyset \bowtie Pow(\{a_2, a_3\}))$. We can iterate the observation for the leftmost occurrence of a_2 and obtain:

$$Pow(\{a_1, a_2, a_3\}) = (\{a_1, a_2\} \bowtie Pow(\{a_3\})) \cup (\{a_1\} \bowtie Pow(\{a_3\})) \cup \ (\emptyset \bowtie Pow(\{a_2, a_3\})).$$

By iterating again for the leftmost occurrence of a_3, we conclude:

$$\begin{aligned} Pow(\{a_1, a_2, a_3\}) \ = \ & (\{a_1, a_2, a_3\} \bowtie Pow(\emptyset)) \cup (\{a_1, a_2\} \bowtie Pow(\emptyset)) \cup \\ & (\{a_1\} \bowtie Pow(\{a_3\})) \cup (\emptyset \bowtie Pow(\{a_2, a_3\})) \end{aligned}$$

Since $R \bowtie Pow(\emptyset) = \{R\}$, the leftmost set in the above union is $\{\{a_1, a_2, a_3\}\}$. In general, the following recurrence relation holds.

Lemma 1 *Let $S = \{a_1, \ldots, a_n\}$. We have:*

$$R \bowtie Pow(S) = \{R \cup S\} \cup \bigcup_{i=n,\ldots,1} (R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\})$$

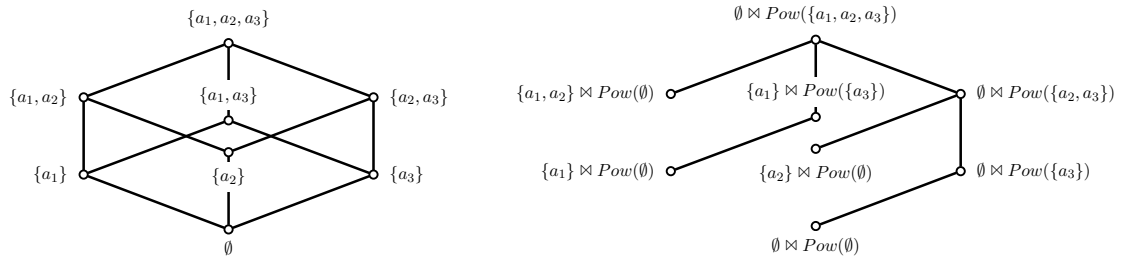

Figure 1: Lattice of subsets and reverse binary counting.

Proof The proof is by induction on n. The base case $n = 0$ is trivial: $R \bowtie Pow(\emptyset) = \{R\}$ by definition. Consider now $n > 0$. Since $Pow(S) = (\{a_1\} \bowtie Pow(S \setminus \{a_1\})) \cup Pow(S \setminus \{a_1\})$, we have: $R \bowtie Pow(S) = R \bowtie ((\{a_1\} \bowtie Pow(\{a_2, \ldots, a_n\})) \cup (\emptyset \bowtie Pow(\{a_2, \ldots, a_n\})))$. Since the $\bowtie$ operator satisfies:

$$R \bowtie (P_1 \cup P_2) = (R \bowtie P_1) \cup (R \bowtie P_2) \quad \text{and} \quad R_1 \bowtie (R_2 \bowtie P) = (R_1 \cup R_2) \bowtie P$$

we have: $R \bowtie Pow(S) = ((R \cup \{a_1\}) \bowtie Pow(\{a_2, \ldots, a_n\})) \cup (R \bowtie Pow(\{a_2, \ldots, a_n\}))$. By induction hypothesis on the leftmost occurrence of $\bowtie$:

$$\begin{aligned} R \bowtie Pow(S) &= \{R \cup \{a_1\} \cup \{a_2, \ldots, a_n\}\} \cup \\ &\quad \bigcup_{i=n,\ldots,2} (R \cup \{a_1\} \cup \{a_2, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}) \cup \\ &\quad R \bowtie Pow(\{a_2, \ldots, a_n\}) \\ &= \{R \cup S\} \cup \bigcup_{i=n,\ldots,1} (R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}) \end{aligned}$$

■

This result can be readily translated into a procedure **subset(**R**,** S**)** for the enumeration of elements in $R \bowtie Pow(S)$. In particular, since $\emptyset \bowtie Pow(S) = Pow(S)$, **subset(**$\emptyset$**,** S**)** generates all subsets of S. The procedure is shown as Algorithm 1. The search space of the procedure is the tree of the recursive calls of the procedure. The search space for $n = 3$ is reported in Figure 1 (right). According to line 1 of Algorithm 1, the subset outputted at a node labelled as $R \bowtie Pow(S)$ is $R \cup S$. Hence, the output for $n = 3$ is the reverse counting ordering: $\{a_1, a_2, a_3\}$, $\{a_1, a_2\}$, $\{a_1, a_3\}$, $\{a_1\}$, $\{a_2, a_3\}$, $\{a_2\}$, $\{a_3\}$, $\{\}$. Two key properties of Algorithm 1 will be relevant for the rest of the paper.

Remark 2 *A set $R' \cup S'$ generated at a non-root node of the search tree of Algorithm 1 is obtained by removing an element from the set $R \cup S$ generated at its father node. In particular, $R' \cup S' = R \cup S \setminus \{a\}$ for some $a \in S$.*

The invariant $|R' \cup S'| = |R \cup S|$ readily holds for the loop at lines 4–8 of Algorithm 1. Before the recursive call at line 6, an element of S is removed from R', hence the set $R' \cup S'$ outputted at a child node has one element less than the set $R \cup S$ outputted at its father node.

Algorithm 1 subset(R, S) enumerates $R \bowtie Pow(S)$

1: **output** $R \cup S$
2: $R' \leftarrow R \cup S$
3: $S' \leftarrow \emptyset$
4: **for** $a_i \in S$ **do**
5: $\quad R' \leftarrow R' \setminus \{a_i\}$
6: $\quad$ subset(R', S')
7: $\quad S' \leftarrow S' \cup \{a_i\}$
8: **end for**

Remark 3 *The selection order of $a_i \in S$ at line 4 of Algorithm 1 is irrelevant.*

The procedure does not rely on any specific order of selecting members of S, which is a form of *don't care non-determinism* in the visit of the lattice. Any choice generates all elements in $R \bowtie Pow(S)$. In case of an *a priori* positional order of attributes, namely line 4 is "**for** $a_i \in S$ **order by** i **desc do**", Algorithm 1 produces precisely the reversed binary counting order. However, if the selection order varies from one recursive call to another, then the output is still an enumeration of subsets.

4. Generating All Distinct Decision Trees

We build on the subset generation procedure to devise an algorithm for the enumeration of all distinct decision trees built on subsets of the predictive features.

4.1. On Top-Down Greedy Decision Tree Induction

Let us first introduce some notation and assumptions. Let $F = \{a_1, \ldots, a_N\}$ be the set of predictive features, and $S \subseteq F$ a subset of them. We write $T = DT(S)$ to denote the decision tree built from features in S on a fixed training set. Throughout the paper, we make the following assumption on the node split criterion in top-down greedy decision tree induction.

Assumption 4 *Let $T = DT(S)$. A split attribute at a decision node of T is chosen as $argmax_{a \in S} f(a, C)$, where $f()$ is a quality measure and C are the cases of the training set reaching the node.*

While the assumption regards univariate splits, it can be restated for bi-variate or multivariate split conditions, and the theoretical results in this paper can be adapted to such general cases. However, since we build on a software that deals with univariate splits only (see Section 7), experiments are restricted to such a case. Moreover, the results will hold for any quality measure $f()$ as far as the split attributes are chosen as the ones that maximize $f()$. Examples of quality measures used in this way include Information Gain (IG), Gain Ratio[1] (GR), and the Gini index, which are adopted in the C4.5 (Quinlan, 1993) and in the CART systems (Breiman et al., 1984). A second assumption regards the stopping criterion

1. Gain Ratio normalizes Information Gain over the Split Information (SI) of an attribute, i.e., GR = IG/SI. This definition does not work well for attributes which are (almost) constants over the cases C,

in top-down decision tree construction. Let $stop(S, C)$ be the boolean result of the stopping criterion at a node with cases C and predictive features S.

Assumption 5 *If $stop(S, C) = true$ then $stop(S', C) = true$ for every $S' \subseteq S$.*

The assumption states that either: (1) the stopping criterion does not depend on S; or, if it does, then (2) stopping is monotonic with regard to the set of predictive features. (1) is a fairly general assumption, since typical stopping criteria are based on the size of cases C at a node and/or on the purity of the class attribute in C. We will later on consider the stopping criterion of C4.5 which halts tree construction if the number of cases of the training set reaching the current node is lower than a minimum threshold m (formally, $stop(S, C)$ is true iff $|C| < m$). Another widely used stopping criterion satisfying (1) consists of setting a maximum depth of the decision tree. (2) applies to criteria which require minimum quality of features for splitting a node. E.g., the C4.5 additional criterion of stopping if IG of all features is below a minimum threshold satisfies the assumption. The following remark, which is part of the decision tree folklore (see e.g., Caruana and Freitag (1994)), states a useful consequence of Assumptions 4 and 5. Removing any feature not used in a decision tree from the initial set of features does not affect the result of tree building.

Lemma 6 *Let $features(T)$ denote the set of split attributes in a decision tree $T = DT(S)$. For every S' such that $S \supseteq S' \supseteq features(T)$, $DT(S') = T$.*

Proof If a decision tree T built from S uses only features from $U = features(T) \subseteq S$, then at any decision node of T it must be true that $argmax_{a \in S} f(a, C) = argmax_{a \in U} f(a, C)$. Hence, removing any unused attribute in $S \setminus U$ will not change the result of maximizing the quality measure and then, by Assumption 4, the split attribute at a decision node. Moreover, by Assumption 5, a leaf node in T will remain a leaf node for any subset of S. ■

4.2. Enumerating Distinct Decision Trees

Consider a subset of features $R \cup S$, where R must be necessarily used by a decision tree and $S = \{a_1, \ldots, a_n\}$ may be used or not. Let $T = DT(R \cup S)$, and $U = features(T) \supseteq R$ be the features used in split nodes of T. Therefore, $S \cap U = \{a_1, \ldots, a_k\}$ is the set of features in S actually selected as split features, and $S \setminus U = \{a_{k+1}, \ldots, a_n\}$ is the set of features never selected as split features. By Lemma 6, the decision tree T is equal to the one built starting from features $R \cup \{a_1, \ldots, a_k\}$ plus any subset of $\{a_{k+1}, \ldots, a_n\}$. In symbols, all the decision trees for feature subsets in $(R \cup \{a_1, \ldots, a_k\}) \bowtie Pow(\{a_{k+1}, \ldots, a_n\})$ do coincide with T. We will use this observation to remove from the recurrence relation of Lemma 1 some sets in $R \bowtie Pow(S)$ which lead to duplicate decision trees. Formally, when searching for feature subsets that lead to distinct decision trees, the recurrence relation can be modified as:

$$R \bowtie Pow(S) = \{R \cup S\} \cup \bigcup_{i=k,\ldots,1} (R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\})$$

i.e., when SI ≈ 0. Quinlan (1986) proposed the heuristic of restricting the evaluation of GR only to attributes with above *average* IG. The heuristic is implemented in the C4.5 system (Quinlan, 1993). It clearly breaks Assumption 4, making the selection of the split attribute dependent on the set S. An heuristic that satisfies Assumption 4 consists of restricting the evaluation of GR only for attributes with IG higher than a *minimum* threshold.

Algorithm 2 DTdistinct(R, S) enumerates distinct decision trees using feature subsets in $R \bowtie Pow(S)$.

```
1: build tree T = DT(R ∪ S)
2: U ← features(T)
3: if R ⊆ U then
4:    output T
5: end if
6: R' ← R ∪ (S ∩ U)
7: S' ← S \ U
8: for a_i ∈ S ∩ U order by rk_frontier(T) do
9:    R' ← R' \ {a_i}
10:   DTdistinct(R', S')
11:   S' ← S' ∪ {a_i}
12: end for
```

since the missing union:

$$\bigcup_{i=n,\ldots,k+1} (R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}) \tag{1}$$

is included in $(R \cup \{a_1, \ldots, a_k\}) \bowtie Pow(\{a_{k+1}, \ldots, a_n\})$, and then it contains sets of features V such that $DT(V) = DT(R \cup S)$. In particular, this implies the following property, for any error estimation function $err()$, which will be useful later on:

$$err(DT(R \cup \{a_1, \ldots, a_k\})) = err(DT(V)) \tag{2}$$
$$\text{for all } V \in \textstyle\bigcup_{i=n,\ldots,k+1}(R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}).$$

The simplified recurrence relation prunes from the the search space feature subsets that lead to duplicated decision trees. However, we will show in Example 1 that such a pruning alone is not sufficient to generate distinct decision trees only, i.e., duplicates may still exist.

Algorithm 2 provides an enumeration of *all and only the distinct decision trees*. It builds on the subset generation procedure. Line 1 constructs a tree T from features $R \cup S$. Features in the set $S \setminus U$ of unused features in T are not iterated over in the loop at lines 8–12, since those iterations would yield the same tree as T. This is formally justified by the modified recurrence relation above. The tree T is outputted at line 4 only if $R \subseteq U$, namely features *required* to be used (i.e., R) are *actually* used in decision node splits. We will shows that such a test characterizes a uniqueness condition for all feature subsets that lead to a same decision tree. Hence, it prevents outputting more than once a decision tree that can be obtained from multiple paths of the search tree.

Example 1 *Let $F = \{a_1, a_2, a_3\}$. Assume that a_1 has no discriminatory power unless data has been split by a_3. More formally, $DT(S) = DT(S \setminus \{a_1\})$ if $a_3 \notin S$. The visit of feature subsets of Figure 1 (right) gives rise to the trees built by **DTdistinct($\emptyset$, F)** as shown in Figure 2 (left). For instance, the subset $\{a_1, a_2\}$ visited at the node labelled $\{a_1, a_2\} \bowtie \emptyset$ in Figure 1 (right), produces the decision tree $DT(\{a_1, a_2\})$. By assumption, such a tree*

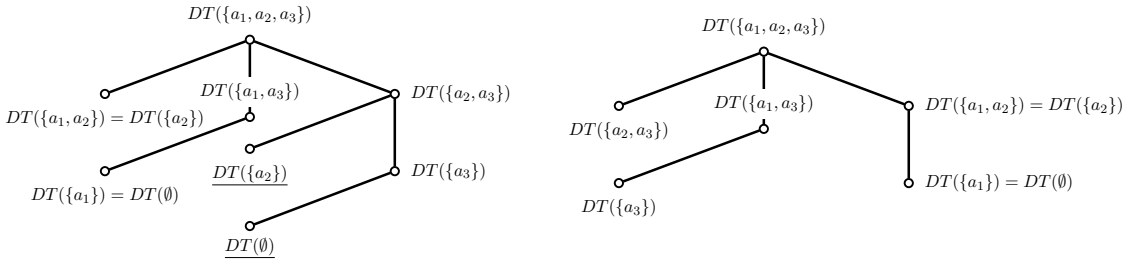

Figure 2: Search spaces of Algorithm 2 for different selection orders.

is equal to $DT(\{a_2\})$, which is a duplicate tree produced in another node – underlined in Figure 2 (left) – corresponding to the feature set visited at the node labelled $\{a_2\} \bowtie \emptyset$. Another example regarding $DT(\{a_1\}) = DT(\emptyset)$ is shown in Figure 2 (left), together with its underlined duplicated tree. Unique trees for two or more duplicates are characterized by the fact that features appearing to the left of $\bowtie$ must necessarily be used as split features by the constructed decision tree. In the two previous example cases, the nodes underlined output their decision trees, while the other duplicates do not pass the test at line 3 of Algorithm 2.

The following non-trivial result holds.

Theorem 7 ***DTdistinct(**R**,** S**)*** *outputs the distinct decision trees built on sets of features in $R \bowtie Pow(S)$.*

Proof The search space of **DTdistinct** is a pruning of the search space of **subset**. Every tree built at a node and outputted is then constructed from a subset in $R \bowtie Pow(S)$. By Remark 3, the order of selection of $a_i \in S \cap U$ at line 8 is irrelevant, since any order will lead to the same space $R \bowtie Pow(S)$.

Let us first show that decision trees in output are all distinct. The key observation here is that, by line 4, all features in R are used as split features in the outputted decision tree. The proof proceeds by induction on the size of S. If $|S| = 0$, then there is at most one decision tree in output, hence the conclusion. Assume now $|S| > 0$, and let $S = \{a_1, \ldots, a_n\}$. By Lemma 1, any two recursive calls at line 10 have parameters $(R \cup \{a_1, \ldots, a_{i-1}\}, \{a_{i+1}, \ldots, a_n\})$ and $(R \cup \{a_1, \ldots, a_{j-1}\}, \{a_{j+1}, \ldots, a_n\})$, for some $i < j$. Observe that a_i is missing as a predictive attribute in the trees in output from the first call, while by inductive hypothesis it must be a split attribute in the trees in output by the second call. Hence, the trees in output from recursive calls are all distinct among them. Moreover, they are all different from $T = DT(R \cup S)$, because recursive calls do not include some feature $a_i \in S \cap U$ that is by definition used in T.

Let us now show that trees pruned at line 8 or at line 4 are already outputted elsewhere, which implies that every distinct decision tree is outputted at least once. First, by Lemma 6, the trees that would have been outputted in the pruned iterations at line 8 (i.e., for $a_i \in S \setminus U$) are equal to the tree of $T = DT(R \cup S)$. Second, if the tree T is not outputted at line 4, because $R \not\subseteq U$, we have that it is outputted at another node of the search tree. The proof is by induction on $|R|$. For $|R| = 0$ it is trivial, because the the premise $R \not\subseteq U$ does not hold. Let $R = \{a_1, \ldots, a_n\}$, with $n > 0$, and let $a_1, \ldots, a_n$ be

in the order they have been added by recursive calls. Fix $R' = \{a_1, \ldots, a_{i-1}\}$ such that $a_i \notin U$ and $R' \subseteq U$. There is a sibling node or a sibling of an ancestor node in the search tree corresponding to a call with parameters R' and $S' \supseteq \{a_{i+1}, \ldots, a_n\} \cup S$. By inductive hypothesis on $|R'| < |R|$, the distinct decision trees with features in $R' \bowtie Pow(S')$ are all outputted, including T because T has split features in $R \cup S \setminus \{a_i\}$ which belongs to $R' \bowtie Pow(S')$. ■

The proof of Theorem 7 does not assume any specific order at line 8 of Algorithm 2. Any order would produce the enumeration of distinct decision trees – this is a consequence of Remark 3. However, the order may impact on the size of the search space.

Example 2 *Reconsider Example 1. The order of selection of a_i's in the visit of Figure 2 (left) is by descending i's. This ordering does not take into account the fact that a_3 has more discriminatory power than a_1, i.e., its presence gives rise to more distinct decision trees. Consider instead having a_3 removed in the rightmost child of the root, e.g., the selection order a_1, a_2, and a_3. The search space of* ***DTdistinct(∅, F)*** *is reported in Figure 2 (right). Notice that no duplicate decision tree is built here, and that the size of the search space is smaller than in the previous example. In fact, the node labelled as $DT(\{a_1, a_2\}) = DT(\{a_2\})$ corresponds to the exploration of $\emptyset \bowtie \{a_1, a_2\}$. The a_1 attribute is unused and hence, it is pruned at line 8 of Algorithm 2. The sub-space to be searched consists then of only the subsets of $\{a_1\}$, not all subsets of $\{a_1, a_2\}$. The child node finds $DT(\{a_1\}) = DT(\emptyset)$ and then it stops recursion since there are no used attributes to iterate over at line 8.*

Algorithm 2 adopts a specific order intended to be effective in pruning the search space. In particular, our objective is to minimize the number of duplicated trees built. In fact, even though duplicates are not outputted, building them has a computational cost that should be minimized. Duplicated decision trees are detected through the test $R \subseteq U$ at line 4 of Algorithm 2. Thus, we want to minimize the number of recursive calls **DTdistinct(R', S')** where attributes in R' have lower chances of being used. Since required attributes are removed from R' one at a time at line 9 (and added to the set of possibly used attributes S' at line 11), this means ranking attributes in $S \cap U$ by increasing chances of being used in decision trees built in recursive calls. How do we estimate such chances?

Example 3 *Consider the sample decision tree at Figure 3(a). It is built on the set of features $F = \{a_1, a_2, a_3, a_4\}$. Which attributes have the highest chance of being used if included in a subset of F? Whenever included in the subset, a_1 will be certainly used at the root node. In fact, it already maximizes the quality measure on F, hence by Assumption 4 it will also maximizes the quality measure over any subset of F. Assume now a_1 is selected for inclusion. Both a_2 and a_3 will be certainly selected as split attributes, for the same reason as above. Which one should be preferred first? Let us look at the sub-trees rooted at a_2 and a_3. If a_2 is selected, then the sub-tree rooted at a_2 uses no further attribute. In fact, a_1 cannot be counted as a further attribute because it is known to be already used. Conversely, if a_3 is selected, the sub-tree rooted at a_3 ensures that attribute a_4 can be used. Therefore, having a_3 gives more chances of using further attributes in decision tree building. Therefore, a_3 should be selected for inclusion before a_2. Finally, sub-trees rooted at a_2 and a_4 use no*

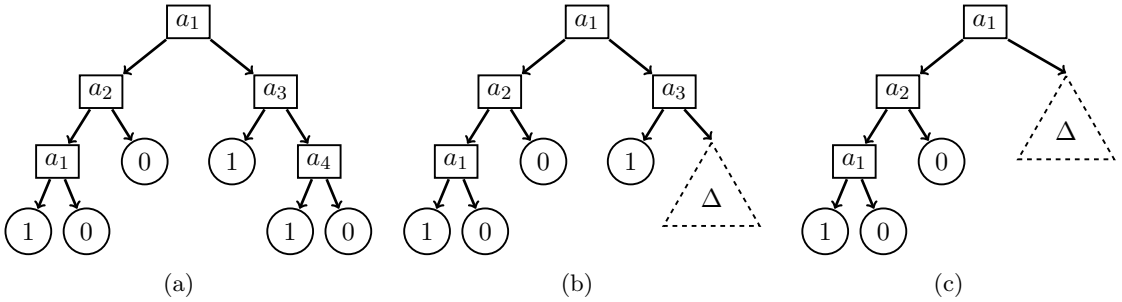

Figure 3: A sample decision tree, and two subtrees replaced with an oracle Δ. Internal nodes are labelled with split attributes, and leaves are labelled with class value.

further attributes, and we break the tie by selecting a_2 before a_4. In summary, the rank of attributes in F with increasing chances of being used is a_4, a_2, a_3, a_1.

Let us formalize the intuition of this example. We define an *a-frontier node* of a decision tree T w.r.t. a set R' of features, as a decision node that uses $a \notin R'$ for the first time in a path from the root to a node. A frontier node is any a-frontier node, for any feature a. Based on the previous example, we should count the number of attributes that are used in sub-trees rooted at frontier nodes.

Definition 8 *We define $frontier(T, R', a)$ as the number of distinct features not in R' that are used in sub-trees of T rooted at a-frontier nodes of T w.r.t. R'.*

Notice that attributes in R' are excluded from the counting in $frontier()$. The idea is that R' will include attributes that must already appear somewhere in a decision tree (in between the root and frontier nodes), and thus their presence in the sub-tree rooted at a does not imply further usability of such attributes. We are now in the position to introduce the ranking $rk_{frontier}$ based on ascending $frontier()$.

Definition 9 *Let $T = DT(R \cup S)$ and $U = features(T)$. $rk_{frontier}(T)$ is the order $r_1, \ldots, r_k$ of elements in $S \cap U$ such that, for $i = k, \ldots, 1$:*

$$r_i = argmax_{a \in (S \cap U) \setminus \{r_{i+1}, \ldots, r_k\}} \; frontier(T, R \cup \{r_{i+1}, \ldots, r_k\}, a).$$

The definition of the ranking iterates from the last to the first position, and at each step selects the feature which maximizes the $frontier()$ measure. Iteration is necessary due to the fact that the frontier nodes depend on the features selected at the previous step. Intuitively, the ordering tries to keep as much as possible sub-trees with large sets of not yet ranked attributes. This is in line with the objective of ranking features based on increasing chances of being used in decision trees of recursive calls.

Example 4 *Reconsider Example 3 and the decision tree $T = DT(F)$ in Figure 3(a). Assume that $R = \emptyset$ and $S = F = \{a_1, a_2, a_3, a_4\}$. The set of used features is $U = F$.*

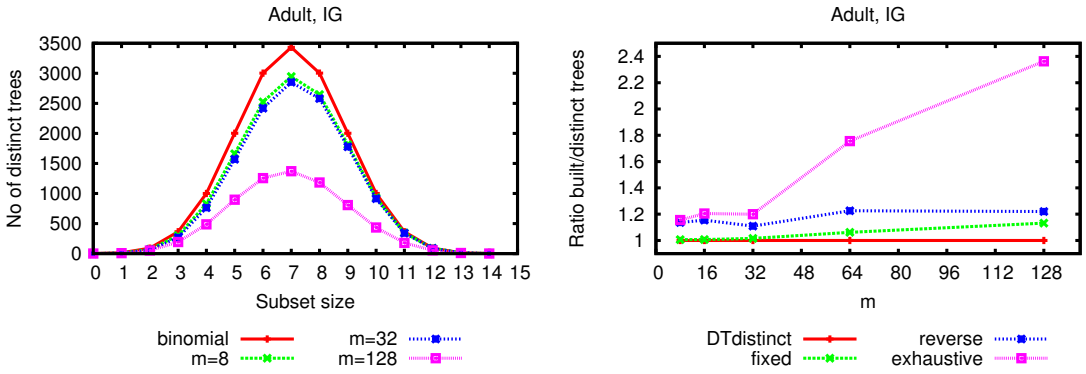

(a) Distribution of distinct decision trees.

(b) Ratio duplicate/distinct decision trees.

Figure 4: Distinct decision trees and overhead of **DTdistinct** on the Adult dataset.

The rank of attributes starts by defining $r_4 = argmax_{a\in\{a_1,a_2,a_3,a_4\}}\ frontier(T,\emptyset,a)$ *which is trivially the split attribute* a_1 *at root node of* T *– the only frontier node of* T *w.r.t.* $\emptyset$*. Next,* $r_3 = argmax_{a\in\{a_2,a_3,a_4\}}\ frontier(T,\{a_1\},a)$ *is defined by looking at the frontier nodes, which are those using* a_2 *and* a_3*. As discussed in Example 3,* $frontier(T,\{a_1\},a_2) = 1$ *and* $frontier(T,\{a_1\},a_3) = 2$*. Thus, we have* $r_3 = a_3$*. At the third step,* $r_2 = argmax_{a\in\{a_2,a_4\}}$ $frontier(T,\{a_1,a_2\},a)$ *is defined by looking at the frontier nodes using* a_2 *and* a_4*. Both have a frontier of* 1*, so we fix* $r_2 = a_2$*. Finally,* r_1 *must be necessarily be* a_4*. Summarizing, the ordering provided by* $rk_{frontier}(T)$ *is* a_4, a_2, a_3, a_1 *as stated in Example 3.*

Let us now point out some properties of **DTdistinct**.

Property 1: linear space complexity. Consider the set F of all features, with N elements. **DTdistinct($\emptyset$, F)** is computationally linear in space (per number of trees built) in N. In fact, there are at most N nested calls, since the size of the second parameter decreases at each call. Summarizing, at most N decision trees are built and stored in the nested calls, i.e., space complexity is linear per number of trees built. An exhaustive search would instead keep in memory the distinct decision trees built in order to check whether a new decision tree is a duplicate. Similarly, so will do applications based on complete search that exploit duplicate pruning through caching of duplicates (Caruana and Freitag, 1994). Those approaches would require exponential space, since the number of distinct trees can be exponential as shown in the next example.

Example 5 *Let us consider the well-known Adult dataset*[2] *(Lichman, 2013), consisting of 48,842 cases, and with* $N = 14$ *predictive features and a binary class. Figure 4(a) shows, for the IG split criterion, the distributions of the number of distinct decision trees w.r.t. the size of feature subset. The distributions are plotted for various values of the stopping parameter* m *(formally,* $stop(S,C)$ *is true iff* $|C| < m$*). For low* m *values, the distribution approaches the binomial; hence, the number of distinct decision trees approaches* 2^N*.*

2. See Section 7 for the experimental settings.

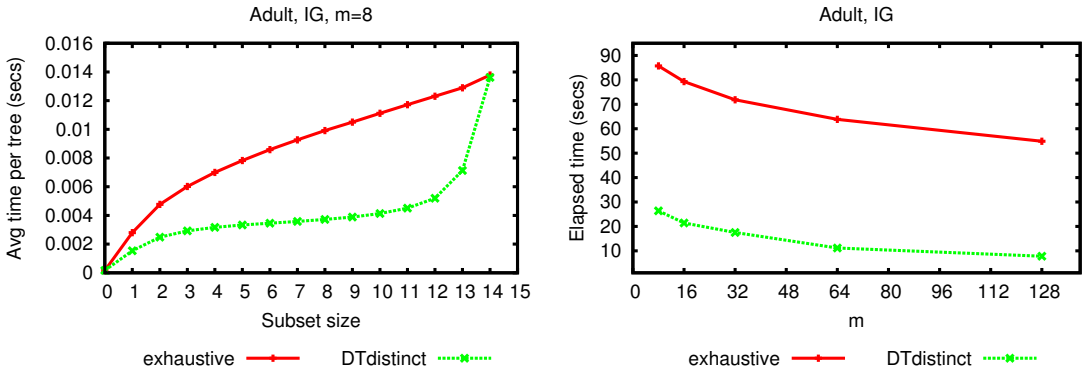

(a) Average elapsed time for tree building.

(b) Total elapsed time.

Figure 5: **DTdistinct** elapsed times on the Adult dataset.

Property 2: reduced overhead. Algorithm 2 may construct duplicated decision trees at line 1, which, however, are not outputted due to the test at line 3. It is legitimate to ask ourselves how many duplicates are constructed. Or, in other words, how effective is the selection order at line 8 based on $rk_{frontier}()$. Formally, we measure such an overhead as the ratio of all decision trees constructed at line 1 over the number of distinct decision trees. An ideal ratio of 1 means that no duplicate decision tree is constructed at all.

Example 6 (Ctd.) *Figure 4(b) shows the overhead at the variation of m for three possible orderings of selection at line 8 of Algorithm 2. One is the the ordering stated by* ***DTdistinct****, based on $rk_{frontier}()$. The second one is the reversed order, namely $a_n, \ldots, a_1$ for $rk_{frontier}()$ being $a_1, \ldots, a_n$. The third one is based on assigning a static index $i \in [1, N]$ to features a_i's, and then ordering over i. The $rk_{frontier}()$ ordering used by* ***DTdistinct*** *is impressively effective, with a ratio of almost* 1 *everywhere.*

The effectiveness of the $rk_{frontier}()$ ordering will be confirmed in the experimental section. Figure 4(b) also reports the ratio of the number of trees in an exhaustive search (which are 2^N for N features) over the number of distinct trees. Smaller m's lead to a smaller ratio, because built trees are larger in size and hence there are more distinct ones. Thus, for small m values, pruning duplicate trees does not guarantee alone a considerably more efficient enumeration than exhaustive search. The next property will help in such cases.

Property 3: feature-incremental tree building. The construction of each single decision tree at line 1 of Algorithm 2 can be sped up by Remark 2. The decision tree T' at a child node of the search tree differs from the decision tree T built at the father node by one missing attribute a_i. The construction of T' can then benefit from this observation. In the implementation of Algorithm 2, we first recursively clone T and then re-build only sub-trees rooted at nodes whose split attribute is a_i. This requires maintaining in memory the trees built along recursive calls, which gives rise to the linear space complexity (in the number of trees) of the algorithm. However, it allows for incrementally building T' from T.

Example 7 (Ctd.) *Figure 5(a) shows the average elapsed time required to built a decision tree w.r.t. the size of the feature subset, for the fixed parameter $m = 8$. The exhaustive search requires an average time linear in the size of the feature subset. Due to incremental tree building,* ***DTdistinct*** *requires instead a sub-linear time.*

Property 4: multi-core parallelization. The parallelization of greedy decision tree algorithms is non-trivial, due to their recursive nature, with some upper bounds on the maximum speedup achievable (Aldinucci et al., 2014). Exhaustive exploration of feature subsets is, instead, parallelizable with perfect scalability, due the exponential number of independent tasks of tree construction. Regarding our pruned search, the loop at lines 8–12 of Algorithm 2 has no strong dependency between iterations, and it can also be easily parallelized on multi-core platforms. Our implementation of Algorithm 2 runs the loop at lines 8–12 in task parallel threads. The construction of the tree at line 1 is also executed using nested task parallelism. For fair comparison, the implementation of exhaustive search exploits nested parallelism as well in the enumeration and in the construction of trees.

Example 8 (Ctd.) *Figure 5(b) contrasts the total elapsed times of exhaustive search and* ***DTdistinct****. For small values of m, the number of trees built by exhaustive search approaches the number of distinct decision trees (see Figure 4(b)). Nevertheless, the running time of* ***DTdistinct*** *is constantly better than the exhaustive search. This improvement is due to the incremental building of decision trees. The computational efficiency in terms of absolute elapsed times is, in addition, due the effectiveness of parallel implementation, which in the example at hand runs on an low-cost 8-core machine reaching a 7× speedup.*

5. Best and Acceptable Feature Subsets

Consider an error estimation function $err(T)$ for a decision tree T. A best feature subset is such that the estimated error of the tree built on such features is minimum among decision trees built on any feature subset. It is δ-acceptable, if the estimated error is lower or equal than the minimum plus δ.

Definition 10 *Let F be a set of features. We define $err_{bst} = min_{S \subseteq F} err(DT(S))$, and call it the best estimated error w.r.t. feature subsets. For $\delta \geq 0$, a feature subset $S \subseteq F$ is δ-acceptable if:*

$$err(DT(S)) \leq \delta + err_{bst}.$$

When $\delta = 0$, we call S a best feature subset. Finally, $DT(S)$ is called a δ-acceptable decision tree.

The *δ-acceptable feature subset problem* consists of finding a δ-acceptable feature subset. In particular, for $\delta = 0$, it consists of finding a best feature subset.

We make no assumption on the error estimation function $err(T)$. In experiments, unless otherwise stated, we adhere to the wrapper model (John et al., 1994; Kohavi and John, 1997), by assuming that the available training set is split into a *building* set, used to build the decision trees $T = DT(S)$ on a subset S of features F, and a *search* set. Error is

Algorithm 3 DTaccept$_\delta(R, S)$ finds a δ-acceptable feature subset S_{acc} in $R \bowtie Pow(S)$.

```
1: // e_acc is initialized outside to ∞
2: build tree T = DT(R ∪ S)
3: U ← features(T)
4: if err(T) ≤ e_acc then
5:    e_acc ← err(T)
6:    S_acc ← U
7: end if
8: U ← greedy_δ(U, R, S)
9: R' ← R ∪ (S ∩ U)
10: S' ← S \ U
11: for a_i ∈ S ∩ U order by rk_frontier(T) do
12:    R' ← R' \ {a_i}
13:    DTaccept(R', S')
14:    S' ← S' ∪ {a_i}
15: end for
```

Algorithm 4 greedy$_\delta(U, R, S)$

```
1: W ← U
2: for a_i ∈ S ∩ U order by rk_frontier(T) do
3:    Ŵ ← W \ {a_i}
4:    if e_acc ≤ lberr(R, S ∩ Ŵ) + δ then
5:       W ← Ŵ
6:    end if
7: end for
8: return W
```

estimated as the empirical misclassification error on the search set, and it is computed using the C4.5's distribution imputation method[3].

Algorithm 3 builds on the procedure for the enumeration of distinct decision trees by implementing a further pruning of the search space. In particular, a call **DTaccept$_\delta(R, S)$** searches for a δ-acceptable feature subset S_{acc} among all subsets in $R \bowtie Pow(S)$. The global variable e_{acc} stores the best error estimate found so far, and it is initialized outside the call to ∞. The structure of **DTaccept$_\delta$** follows the one of Algorithm 2, from which it differs in two main points.

The first difference regards lines 4-8, which instead of just outputting the feature subset, they update the best error estimation found so far in case the estimated error of T is lower or equal[4] than it. The set of features S_{acc} is also updated.

3. Predictions of instances with no missing value follows a path from the decision tree root to a leaf node. For instances with missing value of the attribute tested at a decision node, several options are available (Saar-Tsechansky and Provost, 2007; Twala, 2009). In C4.5, all branches of the decision node are followed, and the prediction of a leaf in a branch contributes in proportion to the weight of the branch's child node (fraction of cases reaching the decision node that satisfy the test outcome of the child). The class value predicted for the instance is the one with the largest total contribution.
4. We break ties in favor of smaller feature subsets.

The second difference regards the set U of used features which, at line 8, is possibly pruned by the call $\mathbf{greedy}_\delta(U, R, S)$. Such a function tries and relax the pruning of features subsets in formula (2) in order to include additional attributes. Let $S = \{a_1, \ldots, a_n\}$. In particular, we aim at finding a minimal set $W = \{a_1, \ldots, a_k\} \subseteq S \cap U$ such that:

$$e_{acc} \leq err(DT(V)) + \delta \quad (3)$$
$$\text{for all } V \in \bigcup_{i=n,\ldots,k+1}(R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}).$$

If such a condition holds[5], we can prune from the search space the feature subsets in the quantifier of the condition, because even in the case that a feature subset with the global best error is pruned this way, the best error found so far e_{acc} is within the δ bound from the error of the pruned decision tree. Practically, this means that we can continue using W in the place of U in the rest of the search – and, in fact, lines 9–15 of $\mathbf{DTaccept}_\delta$ coincide with lines 6–12 of **DTdistinct**. The search space is pruned if $W \subset U$, because the loop at lines 11-15 iterates over a smaller set of features. The approach that $\mathbf{greedy}_\delta$ adopts for determining $\{a_1, \ldots, a_k\}$ is a greedy one, which tries and removes one candidate feature from $S \cap U$ at a time. The order of removal is by[6] $rk_{frontier}(T)$ as in the main loop of **DTdistinct**. The function $\mathbf{greedy}_\delta$ relies on a lower-bound function $lberr()$ for which the test condition[7] $e_{acc} \leq lberr(R, \{a_1, \ldots, a_k\}) + \delta$ is required to imply that (3) holds. This is obviously true when:

$$lberr(R, \{a_1, \ldots, a_k\}) \leq err(DT(V)) \quad (4)$$
$$\text{for all } V \in \bigcup_{i=n,\ldots,k+1}(R \cup \{a_1, \ldots, a_{i-1}\}) \bowtie Pow(\{a_{i+1}, \ldots, a_n\}).$$

hence, the adjective "lower-bound" function for $lberr()$. Our lower-bound function is defined as follows for a candidate $\{a_1, \ldots, a_k\}$. We start from the decision tree $T = DT(R \cup S) = DT(R \cup (S \cap U))$ and remove sub-trees in T rooted at frontier nodes with split features in $(S \cap U) \setminus \{a_1, \ldots, a_k\}$. Let T' be the partial tree obtained and call the removed frontier nodes the "to-be-expanded" nodes. The features used in T' belong to $R \cup \{a_1, \ldots, a_k\}$, which is included in any V quantified over in (4). Due to the Assumptions 4-5, any tree $DT(V)$ may differ from T only at the nodes to-be-expanded and their sub-trees, i.e., T' is a sub-tree of $DT(V)$. At the best, the error of the sub-trees in V that expand T' will be zero[8]. Thus, we have $err(T') \leq err(DT(V))$, where $err(T')$ adds 0 as estimated error at nodes to-be-expanded. Since T' is defined only starting from T, and not from any specific

5. A direct extension of (2) is to require $err(DT(R \cup S)) \leq err(DT(V)) + \delta$, i.e., that the error of the last built tree is within the δ bound from the error of any tree over quantified attributes V. The relaxed condition (3) allows for a better pruning, namely that the best error found so far is within the δ bound.
6. The rationale is to try and remove first features that lead to minimal changes in the decision tree, and, consequently, to minimal differences in its misclassification error. This is a direct generalization of the approach of **DTdistinct** of removing unused features, which lead to no change in misclassification.
7. Cf. line 4 of Algorithm 4, where $\{a_1, \ldots, a_k\}$ is $S \cap \hat{W}$.
8. Consider the wrapper model, namely error is estimated as the empirical misclassification error over a search set. Since decision tree error is always not lower than the Bayes error (Devroye et al., 1996), a better lower bound can be obtained at each node to-be-expanded as: $1 - \frac{1}{|I|}\sum_{\mathbf{x} \in I} g(\mathbf{x})$. Here, I includes the instances in the search set reaching the node to-be-expanded, and $g(\mathbf{x}) = max_c\ k^c_{\mathbf{x}}/n_{\mathbf{x}}$, where $n_{\mathbf{x}}$ is the number of instances in I whose predictive attribute values are the same as for $\mathbf{x}$, and $k^c_{\mathbf{x}}$ is the number of such instances that also have class value equal to c. Unfortunately, the computational cost of calculating this lower bound function makes it impractical.

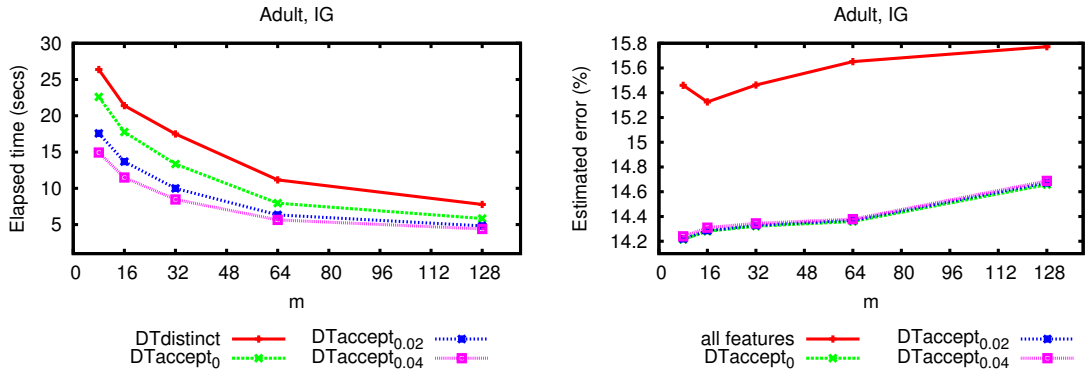

(a) Elapsed running time.

(b) Misclassification error on search set.

Figure 6: **DTaccept**$_\delta$ elapsed times and estimated errors on the Adult dataset.

V, we then define $lberr(R, \{a_1, \ldots, a_k\}) = err(T')$ and have that (4) holds. In summary, we have the following result.

Theorem 11 *Let F be a set of features.* ***DTaccept***$_\delta$***(R, S)*** *finds a δ-acceptable feature subset among all subsets in $R \bowtie Pow(S)$, if it exists. In particular,* ***DTaccept***$_\delta$***(∅, F)*** *finds a δ-acceptable feature subset.*

Example 9 *Consider again the sample decision tree T at Figure 3(a). Assume $R = \emptyset$ and $S = S \cap U = \{a_1, a_2, a_3, a_4\}$. Also, let e_{acc} the best error estimate found so far.* ***greedy***$_\delta$ *will consider removing attributes in the order provided by $rk_{frontier}(T)$, which is a_4, a_2, a_3, a_1 (see Example 4). At the first step, the sub-tree rooted at a_4 is tentatively removed, and replaced with an oracle with zero error estimate, as shown in Figure 3(b). If the error $lb = lberr(\emptyset, \{a_4\})$ of such a tree is such that $e_{acc} \leq lb + \delta$, we can commit the removal of a_4. Assume this is the case. In the next step, the sub-tree rooted at a_2 is also tentatively removed and replaced with an oracle. Again, if the estimated error $lb = lberr(\emptyset, \{a_4, a_2\})$ of such a tree is such that $e_{acc} \leq lb + \delta$, we can commit the removal of a_2. Assume this is not the case, and a_2 is not removed. In the third step, the sub-tree rooted at a_3 is tentatively removed and replaced with an oracle, as shown in Figure 3(c). If the estimated error $lb = lberr(\emptyset, \{a_4, a_3\})$ of such a tree is such that $e_{acc} \leq lb + \delta$, we can commit the removal of a_3. Assume this is the case. In the last step, we try and remove the sub-tree rooted at a_1, and replace it with an oracle. The estimated error of such a tree is $lb = lberr(\emptyset, \{a_4, a_3, a_1\}) = 0$. Condition $e_{acc} \leq lb + \delta$ does not hold (otherwise, it would have been satisfied at the second step as well). In summary,* ***greedy***$_\delta$ *returns $\{a_2, a_1\}$, whilst $\{a_4, a_3\}$ can be safely not iterated over at step 11 of Algorithm 3.*

Example 10 *Reconsider Example 5. Figure 6(a) shows the elapsed running times of* ***DTaccept***$_\delta$***(∅, F)*** *where F is the set of all features of the Adult dataset. Results are shown for several values of the parameter δ. It is worth noting that, for $\delta = 0$, the elapsed time is smaller than the enumeration of distinct trees by* ***DTdistinct****. In fact, when we*

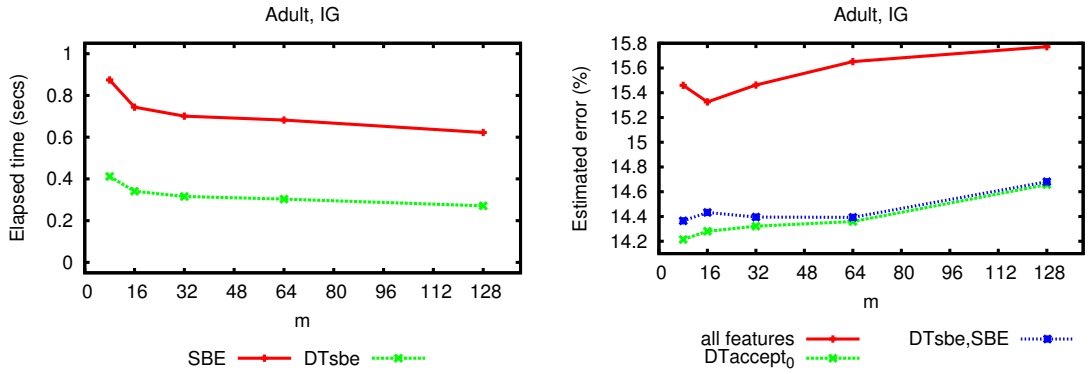

(a) Elapsed running time.

(b) Misclassification error on search set.

Figure 7: **SBE** and **DTsbe** elapsed times and estimated errors on the Adult dataset.

search for a best feature subset, $\boldsymbol{DTaccept_0}$ *prunes from the search space those (distinct) decision trees for which the lower bound on the estimated error is higher than the best error estimation found during the search. Figure 6(b) shows the estimated error (misclassification error on the search set) of the decision tree built on features returned by* $\boldsymbol{DTaccept_\delta}$ *and on all features* F*. The difference between the estimated error of* $DT(F)$ *and the estimated error of a best feature subset (*$\delta = 0$*) provides the range of error estimations that may be returned by heuristic feature selection approaches adhering to the wrapper model. Notice that the estimated error of* δ*-acceptable feature subsets for* $\delta = 0.02$ *and* $\delta = 0.04$ *is very close to the estimated error of a best feature subset (only* $+2\%$ *and* $+4\%$ *respectively).*

The example shows that the pruning strategy of $\mathbf{DTaccept}_\delta$ is effective in the specific case of the Adult dataset. In the worst case, however, the search space remains the one of distinct decision trees, which, for low m values, is exponential in the number of features. In Section 7, we will test performances on datasets of larger dimensionalities.

As a final note, we observe that our approach can be easily adapted to other variants of the feature selection problem. One variant consists of regularizing the error estimation with a penalty ϵ for every feature in a subset. In such a case, the only changes in Algorithm 3 would be: the test at line 4 becomes $err(T)+|U|\cdot\epsilon \leq e_{acc}$; the assignment at line 5 becomes $e_{acc} \leftarrow err(T) + |U| \cdot \epsilon$. Regarding the lower bound function $lberr(R, S \cap \hat{W})$ called at line 4 of $\mathbf{greedy}_\delta$, by adding the penalty $|S \cap \hat{W}| \cdot \epsilon$ we obtain a lower bound on the regularized error estimation of trees built from subsets V pruned in (4).

6. A White-Box Optimization of SBE

On the practical side, $\mathbf{DTaccept}_\delta$ does not scale to large dimensional datasets. Moreover, acceptability/best error on the search set may be obtained at the cost of overfitting (Doak, 1992; Reunanen, 2003) and instability (Nogueira and Brown, 2016). We will discuss these issues in the experimental section. The ideas underlying our approach, however, can impact also on the efficiency of well-performing heuristic searches. In particular, we consider here the cornerstone sequential backward elimination (**SBE**) heuristic. It starts building a classi-

fier T using the set S of all features, i.e., $S = F$. For every $a \in S$, a classifier T_a is built using features in $S \setminus \{a\}$. If no T_a's has a lower or equal estimated error than T, the algorithm stops returning S as the subset of selected features. Otherwise, the procedure is repeated removing $\hat{a}$ from S, where $T_{\hat{a}}$ is the classifier with the smallest estimated error. In summary, features are eliminated one at a time in a greedy way. **SBE** is a black-box approach. The procedure applies to any type of classifier. A white-box optimization can be devised for decision tree classifiers that satisfy the assumptions of Section 4.1. Let $U = \mathit{features}(T)$ be the set of features used in the current decision tree $T = DT(S)$. By Lemma 6, for a non-used feature $a \in S \setminus U$, it turns out that $T_a = DT(S \setminus \{a\}) = DT(S) = T$. Thus, only trees T_a for $a \in U$ need to be built and evaluated at each step, saving the construction of $|S \setminus U|$ decision trees. The following other pruning can be devised. Let $\hat{b}$ be the feature removed at the current step, and $a \in U = \mathit{features}(T)$ such that $\hat{b} \notin \mathit{features}(T_a)$, where $\mathit{features}(T_a)$ is known from the previous step. By Lemma 6, it turns out that $(T_{\hat{b}})_a = DT(S \setminus \{\hat{b}, a\}) = DT(S \setminus \{a\}) = T_a$. In summary, $(T_{\hat{b}})_a$ can be pruned if $a \notin \mathit{features}(T_{\hat{b}})$ or if $\hat{b} \notin \mathit{features}(T_a)$. The optimizations of incremental tree building and parallelization discussed for **DTdistinct** readily apply for such heuristic as well. We call the resulting white-box optimized algorithm **DTsbe**.

Example 11 *Figure 7(a) contrasts the elapsed running times of **SBE** and **DTsbe** on the Adult dataset. The efficiency improvement of **DTsbe** over **SBE** is consistently in the order of 2×. Figure 7(b) shows instead their estimated error (misclassification error on the search set), which are obviously the same. Estimates are very close to the estimated error of a best feature subset provided by $\boldsymbol{DTaccept}_0$.*

7. Experiments

7.1. Datasets and Experimental Settings

We perform experiments on 20 small and large dimensional benchmark datasets publicly available from the UCI ML repository (Lichman, 2013). Some of the datasets have been used in the NIPS 2003 challenge on feature selection, and are described in detail by Guyon et al. (2006a). Table 1 reports the number of instances and features of the datasets.

Following (Kohavi, 1995; Reunanen, 2003), the generalization error of a classifier is estimated by repeated stratified 10-fold cross validation[9]. Cross-validation is repeated 10 times. At each repetition, the available dataset is split into 10 folds, using stratified random sampling. Each fold is used to compute the misclassification error of the classifier built on the remaining 9 folds used as training set for building classification models. The generalization error is then the average misclassification error over the 100 classification models (10 models times 10 repetitions). The following classification models will be considered:

baseline: the decision tree $DT(F)$ built on all available features F;

DTsbe and SBE: the decision tree built on the feature subset selected by **DTsbe** (or equivalently by **SBE**);

9. Cross-validation is a nearly unbiased estimator (Kohavi, 1995), yet highly variable for small datasets. Kohavi's recommendation is to adopt a stratified version of it. Variability of the estimator is accounted for by adopting repetitions (Kim, 2009).

DTaccept$_0$: the decision tree built on the feature subset selected by **DTaccept$_0$** (i.e, a best feature subset);

RF: a random forest[10] of 100 decision trees;

DTsbe+RF: a random forest of 100 decision trees, where the available set of features is restricted to those selected by **DTsbe**.

Random forests are included in experiments in order to relate results to state-of-the-art classification models. With reference to the feature selection methods (**DTsbe**, **SBE** and **DTaccept$_0$**), the error estimator used adheres to the wrapper model. The training set is split into 70% building set and 30% search set using stratified random sampling. The building set is used for constructing decision trees over subsets of features, and the search set for estimating their error using empirical misclassification. The estimated (search) error of a feature selection strategy is the average misclassification error on the 100 search sets of the decision tree built on the feature subset selected by the strategy. Thus, best feature subsets are those with minimum estimated (search) error, but not necessarily with minimum generalization (cross-validation) error. Different feature selection strategies are applied to the same triples of building, search, and test sets. Hence, paired t-tests can be adopted to evaluate statistically significant differences of their means.

Information Gain (IG) is used as quality measure in node splitting during tree construction. No form of tree simplification (Breslow and Aha, 1997) is considered in this section. Appendix A reports additional results about the impact of C4.5 error-based pruning and the Gain Ratio quality measure.

All procedures described in this paper are implemented by extending the YaDT system (Ruggieri, 2002, 2004; Aldinucci et al., 2014). It is a state-of-the-art main-memory C++ implementation of C4.5 with many algorithmic and data structure optimizations as well as with multi-core data parallelism in tree building. The extended YaDT version is publicly available from the author's home page[11]. Tests were performed on a PC with Intel 8 cores i7-6900K at 3.7 GHz, without hyperthreading, 16 Gb RAM, and Windows Server 2016 OS.

7.2. On the Search Space Visit of DTdistinct

For the Adult dataset, the order $rk_{frontier}(T)$ used in **DTdistinct** was impressively effective in producing a visit of the feature subset space with a negligible fraction of duplicates (see Figure 4(b)). Figure 8(a) shows the ratio of the number of built trees over the number of distinct trees for other small to medium dimensionality datasets – namely, those for which **DTdistinct** terminates within a time-out of 1h. The ratios are below 1.00035, i.e., there are at most 0.035% duplicate trees built by **DTdistinct**.

7.3. Elapsed Running Times

Table 1 compares the elapsed running times of **SBE** and its computational optimization **DTsbe**. For all datasets, the tree building stopping parameter m is set to the small value 2, which is the default for C4.5. The ratio of elapsed times shows a speedup in the range

10. At each decision node, only a random subset of attributes are considered for splitting the node. The number of attributes in the subset is logaritmic in the number of available features.
11. http://pages.di.unipi.it/ruggieri

Table 1: Experimental datasets, and average elapsed running times (seconds). IG and m=2.

dataset	inst.	feat.	DTsbe	SBE	ratio	DTaccept$_0$	RF
Adult	48,842	14	0.63	1.24	0.510	32.42	1.94
Letter	20,000	16	0.39	0.69	0.567	231.64	0.87
Hypo	3,772	29	0.06	0.21	0.283	56.12	0.05
Ionosphere	351	34	0.01	0.09	0.109	0.74	0.01
Soybean	683	35	0.06	0.20	0.281	1843.12	0.03
Kr-vs-kp	3196	36	0.06	0.23	0.272	>1h	0.07
Anneal	898	38	0.01	0.08	0.152	18.27	0.01
Census	299,285	40	35.02	117.85	0.297	>1h	18.64
Spambase	4601	57	0.69	4.84	0.144	>1h	0.13
Sonar	208	60	0.01	0.40	0.031	18.77	0.01
Optdigits	3823	64	0.47	4.61	0.103	>1h	0.14
Coil2000	9,822	85	2.13	23.74	0.090	>1h	0.28
Clean1	476	166	0.13	20.43	0.006	>1h	0.03
Clean2	6,598	166	3.91	151.94	0.026	>1h	0.22
Madelon	2,600	500	9.79	>1h	-	>1h	0.27
Isolet	7,797	617	226.75	>1h	-	>1h	2.09
Gisette	7,000	5,000	87.40	>1h	-	>1h	11.69
P53mutants	31,420	5,408	616.67	>1h	-	>1h	9.58
Arcene	900	10,000	25.59	>1h	-	>1h	4.33
Dexter	2,600	20,000	1060.91	>1h	-	>1h	0.18

2× - 100×. For large dimensional datasets, the black-box **SBE** does not even terminate within a time-out of 1h on the first fold of the first iteration of repeated cross-validation. This is a relevant result for machine learning practitioners, extending the applicability of the **SBE** heuristic, a key reference in feature selection strategies.

Table 1 also reports the elapsed times for **DTaccept**$_0$. It makes it possibile to search for best feature subsets in a reasonable amount of time for datasets with up to 60 features. In general, the elapsed time depends on: (1) the size of the search space, which is bounded by the number of distinct decision trees; and on: (2) the size of the dataset and the stopping parameter m, which affect the time to build single decision trees. In turn, (1) depends on the number of available features and on (2) itself – in fact, decision trees for small datasets or for large m's values are necessarily small in size. In summary, efficiency of **DTaccept**$_0$ can be achieved for datasets with limited number of features *and* (limited number of instances *or* large values of the stopping parameter m). For example, time-out bounds in Table 1 are reached for medium-to-large number of features (>60) or for medium-to-large number of instances (e.g., the *Kr-vs-kp*, *Census*, and *Spambase* datasets). Notice that while large m's values can speed up the search, they negatively affect the predictive accuracy of decision trees (see next subsection).

Finally, Table 1 also includes the elapsed times of building random forests. They are very low compared to all the other times. In fact, the number of trees in a forest (set to 100) is much smaller than the number of trees built by sequential backward elimination (quadratic in the number of features, in the worst case) and by complete search (exponential, in the worst case).

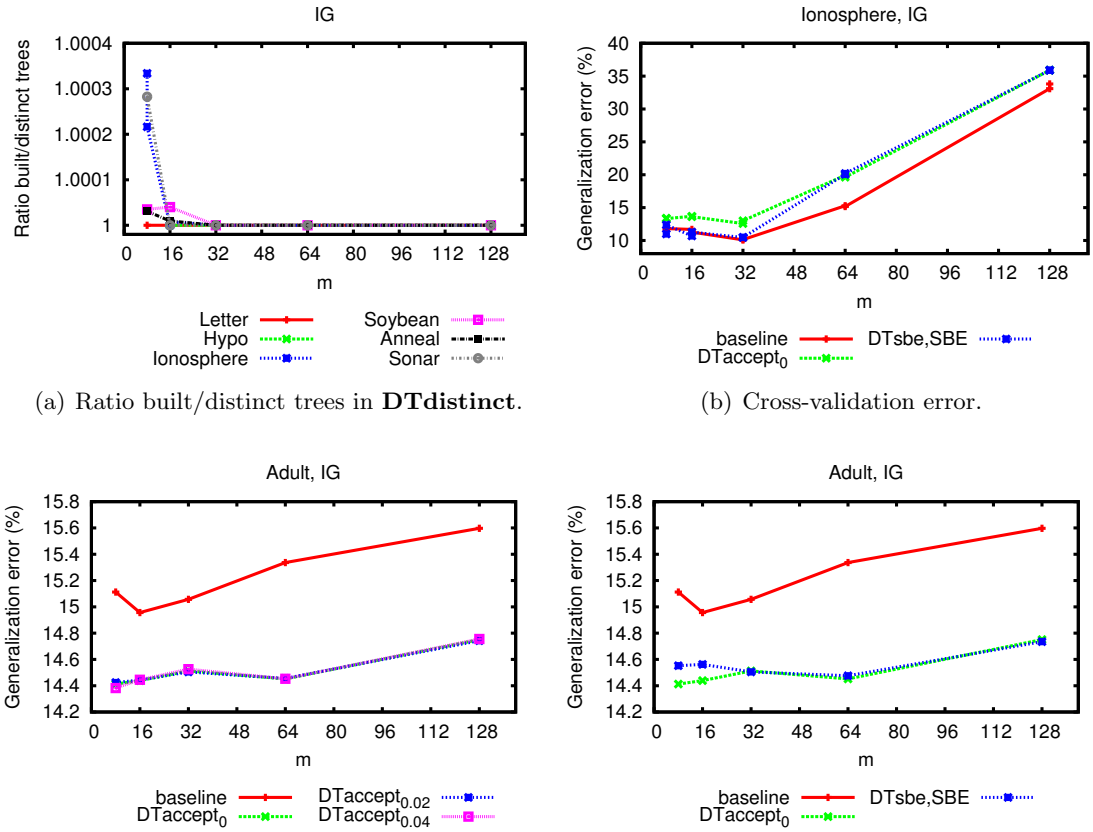

(a) Ratio built/distinct trees in **DTdistinct**.

(b) Cross-validation error.

(c) Cross-validation errors.

(d) Cross-validation error.

Figure 8: Procedure performances.

7.4. Estimated Errors and Generalization Errors

Best feature subsets found by complete search minimize the estimated error (misclassification on the search set), but there is no guarantee that this extends to unseen cases, i.e., to the generalization error evaluated with cross-validation.

Figures 8(c) and 8(d) show generalization errors on the Adult dataset for decision trees constructed on all available features (baseline), and on features selected by **DTaccept**$_0$ (best feature subset), **DTaccept**$_{0.02}$ and **DTaccept**$_{0.04}$ (0.02 and 0.04-acceptable feature subsets), and **SBE** (same as **DTsbe**). For such a dataset, the best feature subsets have the best generalization error, the 0.02- and 0.04-acceptable feature subsets are very close to it, and **SBE** has a similar performance except for low m values.

Figures 8(b) highlights a different result. For the Ionosphere dataset and large m values, the best generalization error is for the baseline, then for **SBE**, and finally for the best feature subsets. Ionosphere is a small dataset, hence using all instances for training (particularly for large m values) results in a better strategy than splitting training into building and search sets for feature selection.

Table 2: Estimated (search) errors and generalization (cross-validation) errors (mean ± stdev). IG and m=2. Best method in **bold**. "⋆" labels methods not different from the best one at 95% significance level.

	estimated (search) error			**generalization (cross-val.) error**		
dataset	**baseline**	**DTsbe**	**DTaccept$_0$**	**baseline**	**DTsbe**	**DTaccept$_0$**
Adult	16.13 ± 0.29	14.35 ± 0.29	**14.19 ± 0.21**	15.76 ± 0.43	14.45 ± 0.45	**14.30 ± 0.42**
Letter	14.82 ± 0.50	13.98 ± 0.47	**13.87 ± 0.44**	12.46 ± 0.76⋆	**12.38 ± 0.71**	12.38 ± 0.76⋆
Hypo	0.55 ± 0.26	0.32 ± 0.17	**0.29 ± 0.17**	**0.42 ± 0.32**	0.52 ± 0.38	0.57 ± 0.43
Ionosphere	11.47 ± 2.82	5.78 ± 1.86	**2.12 ± 1.14**	11.40 ± 5.69	**9.77 ± 4.67**	11.29 ± 5.61
Soybean	16.38 ± 2.68	7.78 ± 2.11	**5.75 ± 1.39**	13.12 ± 4.30	13.15 ± 4.00	**11.53 ± 4.15**
Kr-vs-kp	1.10 ± 0.40	**0.64 ± 0.28**	-	**0.55 ± 0.44**	1.17 ± 0.61	-
Anneal	1.39 ± 0.86	0.84 ± 0.53	**0.60 ± 0.45**	**0.99 ± 0.92**	1.61 ± 1.34	2.14 ± 1.86
Census	6.16 ± 0.07	**4.97 ± 0.13**	-	6.03 ± 0.12	**4.98 ± 0.16**	-
Spambase	8.14 ± 0.83	**5.58 ± 0.61**	-	7.46 ± 1.20⋆	**7.29 ± 1.21**	-
Sonar	28.02 ± 6.72	14.28 ± 3.77	**2.12 ± 1.30**	26.40 ± 9.51⋆	**25.84 ± 9.05**	26.40 ± 9.40⋆
Optdigits	11.11 ± 0.95	**7.79 ± 0.67**	-	**9.88 ± 1.22**	10.27 ± 1.66	-
Coil2000	9.32 ± 0.49	**7.16 ± 0.33**	-	9.03 ± 0.68	**8.76 ± 0.72**	-
Clean1	22.22 ± 3.90	**10.53 ± 2.33**	-	**17.38 ± 6.76**	19.45 ± 6.96	-
Clean2	4.21 ± 0.49	**1.54 ± 0.37**	-	3.31 ± 0.67	**2.80 ± 0.71**	-
Madelon	29.76 ± 3.74	**15.67 ± 2.17**	-	25.91 ± 3.53	**22.51 ± 3.20**	-
Isolet	18.65 ± 0.86	**11.95 ± 0.66**	-	17.10 ± 1.23	**16.80 ± 1.39**	-
Gisette	6.65 ± 0.62	**3.12 ± 0.43**	-	6.14 ± 0.85	**5.91 ± 0.94**	-
P53mutants	0.62 ± 0.06	**0.31 ± 0.04**	-	0.58 ± 0.10	**0.51 ± 0.11**	-
Arcene	48.75 ± 3.17	**30.83 ± 2.58**	-	**48.04 ± 5.46**	48.21 ± 5.08⋆	-
Dexter	48.37 ± 1.69	**31.82 ± 1.38**	-	48.88 ± 2.94⋆	**48.20 ± 3.47**	-
wins/⋆	0/0	13/0	7/0	6/4	12/1	2/2

Table 2 reports the estimated errors and the generalization errors for all datasets. Here and in the following tables, the best method is shown in bold. Other methods are labelled with "⋆" if the null hypothesis in a paired t-test with the best method at 95% significance level cannot be rejected. Three conclusions can be made from the table.

First, heuristic search performs well in comparison to complete search as far as estimated error is concerned with. In fact, the estimated error of **DTsbe** (or equivalently, of **SBE**) is close to the estimated error of best feature subsets in 3 cases out of 7, and is always better than the baseline. Only for Ionosphere and Sonar, it departs from the estimated error of the best feature subsets.

Second, heuristic search generalizes to unseen cases better than baseline and complete search. Globally, **DTsbe** wins in 12 cases and it is not statistically different from the winner in another case. The larger the dimensionality of datasets, the more clear is the advantage of **DTsbe** over the baseline. Comparing **DTsbe** and **DTaccept$_0$** between them, **DTsbe** wins in 2 cases (Ionosphere and Anneal), loses in 2 cases, and it is not statistically different in 3 cases (we count here Hypo, where the baseline wins). Thus we cannot conclude that complete search leads to better generalization errors than heuristic search.

Third, the difference between estimated and generalization error is greater for heuristic search than for the baseline, and for complete search compared to heuristic search. In the

Table 3: Generalization (cross-validation) errors (mean ± stdev). IG and m=2.

dataset	generalization (cross-val.) error RF	DTsbe+RF
Adult	**14.21 ± 0.43**	14.29 ± 0.43
Letter	**4.04 ± 0.46**	4.66 ± 0.63
Hypo	0.70 ± 0.39	**0.48 ± 0.36**
Ionosphere	**6.50 ± 4.01**	8.38 ± 4.35
Soybean	**7.37 ± 2.76**	10.95 ± 3.58
Kr-vs-kp	1.35 ± 0.67*	**1.26 ± 0.66**
Anneal	**0.94 ± 0.88**	1.69 ± 1.38
Census	**4.85 ± 0.07**	4.92 ± 0.12
Spambase	**5.63 ± 1.12**	5.79 ± 1.05
Sonar	**18.42 ± 8.62**	23.26 ± 8.90
Optdigits	**2.03 ± 0.64**	2.76 ± 0.88
Coil2000	**6.01 ± 0.12**	6.11 ± 0.20
Clean1	**10.84 ± 4.62**	15.24 ± 5.68
Clean2	2.46 ± 0.56*	**2.41 ± 0.54**
Madelon	36.25 ± 3.06	**20.93 ± 2.87**
Isolet	**6.34 ± 0.86**	6.46 ± 0.84*
Gisette	4.17 ± 0.78	**3.80 ± 0.71**
P53mutants	**0.45 ± 0.04**	0.45 ± 0.06*
Arcene	45.47 ± 2.61	**44.59 ± 2.59**
Dexter	48.95 ± 2.07	**43.59 ± 2.88**
wins/$\star$	13/2	7/2

case of Sonar, for instance, the differences are considerably large. This implies that oversearching may incur in feature subsets that overfit the data, thus reinforcing the conclusions of Doak (1992); Quinlan and Cameron-Jones (1995); Reunanen (2003) that oversearching increases overfitting.

7.5. Comparison with Random Forests

Random forests are a state-of-the-art classification model that is important to relate with. Single decision trees may be preferrable when interpretability of the model is a requirement that can be traded-off with a lower generalization error. How much lower? A first natural question is then how complete and heuristic searches perform compared to random forests. Table 3 reports the generalization errors of random forest models (**RF**). Contrasting them with the generalization errors in Table 2, random forests perform worse for Hypo and Madelon only. In all other cases, they perform better, often considerably better.

A second question is whether the coupling of heuristic search with random forests may enhance the performance of the random forests alone. Table 3 shows also the generalization errors of **DTsbe+RF**. Recalling that datasets are listed in order of increasing dimensionality (see Table 1), it is immediate to note that for large dimensionality datasets, there is an advantage in doing feature selection before random forests. In particular, we correlate the dataset dimensionality with the difference of mean generalization errors normalized by standard deviation of **RF**. The rank correlation coefficient is $\rho = 0.49$, and the Spearman's test rejects the null hypotesis of zero correlation at 95% confidence level. Intuitively, ran-

Table 4: Number of features and sample Pearson's correlation coefficient in cross-validation (mean ± stdev). IG and m=2.

	number of features			$\phi_{Pearson}$		
dataset	**baseline**	**DTsbe**	**DTaccept$_0$**	**baseline**	**DTsbe**	**DTaccept$_0$**
Adult	13.92 ± 0.27	6.37 ± 1.81	**5.66 ± 1.18**	**0.85 ± 0.36**	0.45 ± 0.27	0.55 ± 0.22
Letter	16.00 ± 0.00	11.35 ± 1.47	**10.21 ± 1.02**	**1.00 ± 0.00**	0.53 ± 0.25	0.70 ± 0.17
Hypo	15.57 ± 2.02	**6.64 ± 1.23**	8.85 ± 3.05	**0.78 ± 0.16**	0.76 ± 0.13	0.51 ± 0.22
Ionosphere	11.44 ± 1.54	**5.35 ± 1.64**	7.20 ± 1.63	**0.43 ± 0.16**	0.33 ± 0.21	0.05 ± 0.19
Soybean	27.52 ± 0.95	**15.28 ± 2.02**	17.12 ± 2.67	**0.87 ± 0.10**	0.29 ± 0.19	0.27 ± 0.18
Kr-vs-kp	28.88 ± 1.07	**19.56 ± 2.10**	-	**0.84 ± 0.11**	0.63 ± 0.14	-
Anneal	10.13 ± 0.68	**6.65 ± 0.95**	10.26 ± 2.38	**0.95 ± 0.06**	0.67 ± 0.20	0.36 ± 0.20
Census	37.94 ± 0.28	**8.97 ± 3.62**	-	**0.97 ± 0.08**	0.50 ± 0.17	-
Spambase	48.89 ± 2.45	**30.98 ± 4.05**	-	**0.53 ± 0.15**	0.35 ± 0.12	-
Sonar	13.32 ± 1.79	**6.78 ± 1.23**	9.11 ± 1.38	**0.25 ± 0.16**	0.15 ± 0.19	0.03 ± 0.15
Optdigits	46.40 ± 1.68	**28.13 ± 2.68**	-	**0.75 ± 0.08**	0.38 ± 0.11	-
Coil2000	62.63 ± 2.14	**39.92 ± 4.65**	-	**0.77 ± 0.07**	0.43 ± 0.10	-
Clean1	28.53 ± 2.46	**14.78 ± 2.29**	-	**0.31 ± 0.13**	0.17 ± 0.11	-
Clean2	79.03 ± 4.67	**37.72 ± 4.67**	-	**0.31 ± 0.10**	0.20 ± 0.09	-
Madelon	113.99 ± 9.87	**49.25 ± 5.91**	-	**0.23 ± 0.05**	0.20 ± 0.06	-
Isolet	240.40 ± 7.61	**105.26 ± 7.71**	-	**0.36 ± 0.04**	0.27 ± 0.05	-
Gisette	129.40 ± 3.95	**56.62 ± 5.01**	-	**0.24 ± 0.04**	0.19 ± 0.06	-
P53mutants	73.73 ± 4.27	**22.05 ± 4.23**	-	**0.20 ± 0.05**	0.09 ± 0.06	-
Arcene	69.19 ± 2.21	**35.02 ± 4.57**	-	**0.05 ± 0.03**	0.02 ± 0.03	-
Dexter	284.85 ± 8.40	**126.10 ± 9.84**	-	**0.38 ± 0.03**	0.28 ± 0.04	-
wins/$\star$	0/0	18/0	2/0	20/0	0/0	0/0

dom forests are more robust to large dimensionality effects than single decision trees due to the (random) selection of a logarithmic number of features to be considered in splits at decision nodes. However, when the number of available features increases, the logarithmic reduction is not sufficient anymore. It is worth noting that such an experimental analysis can only be made due to the efficiency improvements of **DTsbe** over **SBE**. In fact, from Table 1, we know that **SBE** does not terminate within a reasonable time-out.

7.6. On Feature Reduction and Stability

Feature selection is typically a multi-objective problem. In addition to the selection of a feature subset with minimum error, one may be interested in other performance measures, e.g., in minimizing the size of the subset and in minimizing the variability due to perturbations in the training set (stability). The two objectives are contrasting between them, since one can reach low variability by always using all available features. Table 4 reports on the left hand side the mean and standard deviation of the number of features actually used by the decision trees built during cross-validation. The baseline method shows that the embedded feature selection used in tree construction uses a restricted number of features w.r.t. the available ones. **DTsbe** leads to using half or even fewer features. It is the best performing selection strategy with regard to such a measure. **DTaccept$_0$** wins in 2 datasets

Table 5: Elapsed running times (seconds) and generalization (cross-validation) errors (mean ± stdev). IG, m=2, max depth = 5, and misclassification on training set as error estimator in **DTsbe**. $^{(1)}$ one-hot encoding of discrete features required by OCT.

	elapsed time		generalization error		
dataset	**DTsbe**	**OCT**	**baseline**	**DTsbe**	**OCT**
Adult$^{(1)}$	0.34	5.53	14.91 ± 0.45*	**14.87 ± 0.45**	15.2 ± 0.24
Letter	0.03	0.68	49.85 ± 0.92	**49.44 ± 0.92**	55.3 ± 1.61
Hypo$^{(1)}$	0.01	0.20	**0.81 ± 0.41**	0.83 ± 0.44*	1.72 ± 0.72
Ionosphere	0.01	0.09	12.32 ± 4.64	**9.58 ± 4.60**	19.25 ± 5.93
Soybean$^{(1)}$	0.03	0.19	17.88 ± 3.88	**10.83 ± 3.00**	35.82 ± 5.26
Kr-vs-kp$^{(1)}$	0.02	0.17	5.91 ± 1.18	6.10 ± 1.22	**5.8 ± 0.71**
Anneal$^{(1)}$	0.01	0.33	**1.50 ± 1.51**	1.74 ± 1.42*	4.9 ± 2.0
Census$^{(1)}$	10.27	88.15	5.41 ± 0.06	**5.28 ± 0.08**	5.36 ± 0.05
Spambase	0.05	0.38	9.35 ± 1.30*	**9.32 ± 1.29**	11.58 ± 1.15
Sonar	0.02	0.09	26.29 ± 9.52*	**25.91 ± 9.86**	37.36 ± 6.95
Optdigits	0.07	0.07	20.03 ± 1.97	**18.64 ± 2.07**	66.32 ± 1.47
Coil2000	0.15	0.60	**6.01 ± 0.13**	6.11 ± 0.19	7.64 ± 0.5
Clean1	0.04	0.17	24.85 ± 6.17	**21.51 ± 6.99**	36.97 ± 3.61
Clean2	0.30	1.16	6.01 ± 1.16	**5.19 ± 0.72**	8.81 ± 1.04
Madelon	0.44	1.84	23.92 ± 3.83	**21.23 ± 2.84**	44.24 ± 3.71
Isolet	11.33	21.79	35.40 ± 1.50	**34.01 ± 1.37**	38.78 ± 1.56
Gisette	5.54	32.26	6.61 ± 0.91	**6.24 ± 0.96**	9.06 ± 0.74
P53mutants	150.99	142.69	0.47 ± 0.09*	**0.47 ± 0.09**	0.75 ± 0.11
Arcene	3.36	7.47	47.36 ± 5.55*	**47.22 ± 4.89**	48.89 ± 1.83
Dexter	3.02	25.53	46.02 ± 2.51*	**45.86 ± 2.59**	48.58 ± 1.41
wins/$\star$			3/6	16/2	1/0

out of 7. While the differences with **DTsbe** are statistically significant, they are not large – about 1-2 features.

Table 4 reports on the right hand side the mean and standard deviation of the sample Pearson's correlation coefficient over feature subsets used by decision trees[12] built during cross-validation. The mean value corresponds to the $\hat{\boldsymbol{\Phi}}_{\boldsymbol{Pearson}}$ measure of stability introduced by Nogueira and Brown (2016). Differently from other measures of stability, it is unbiased w.r.t. dimensionality of the dataset. From the table, we can conclude that the baseline method has the highest stability of selected features (also with the smallest variance), where the value 1 means that any two distinct folds in cross-validation produce decision trees that use the same subset of features. However, such a stability is obtained at the expenses of a greater number of used features. Feature selection strategies have a considerably lower stability, in some cases half of the values of the baseline. **DTsbe** wins over **DTaccept**$_0$ in 4 cases out of 7 and is equivalent in another. **DTaccept**$_0$ has a better stability only for the 2 lower dimensional datasets. In summary, we can conclude that, in the case of decision tree classifiers, oversearching increases instability as well.

12. Stability is calculated on the subset of features *used* by decision trees, not on the feature subsets *selected* by a strategy. This allows for measuring variability of the baseline approach, which otherwise would result to have zero variability, and to contrast it to the feature selection strategies.

7.7. Comparison with Optimal Decision Trees

The non-greedy OCT approach by Bertsimas and Dunn (2017) transforms decision tree learning into a mixed-integer optimization problem. The objective function to minimize is the empirical misclassification error on the training set regularized by a complexity penalty *cp* on tree size. Such a decision tree is optimal over all possible ways of setting the split attribute at all nodes. An optimal decision tree may not be produceable by greedy top-down algorithms for any feature subset. In this section, we compare OCT[13] with **DTsbe**. For uniformity of comparison, we set the error estimator in **DTsbe** to the misclassification error on the training set (instead of adhering to the wrapper model). In both cases, we set as stopping criterion a maximum tree depth of 5. This is useful for two reasons. First, decision trees are small and easily interpretable, which is the main advantage of using single decision trees over the more powerful model of random forests. Second, exploring the search space of optimal tree search becomes feasible.

Table 5 shows the elapsed times and generalization errors of the two approaches. It also reports the baseline generalization error of depth-bounded C4.5 decision tree. The elapsed times of **DTsbe** are less than the half of the ones of **OCT**, for all datasets except *P53mutants*. The generalization error of **DTsbe** is the best one in 16 out of 20 datasets. Contrasting with Table 2, the restriction on maximum tree depth leads to higher errors – considerably higher for *Letter*, *Kr-vs-kp*, *Optdigits*, and *Isolet*. The generalization errors of **OCT** are almost always higher than the ones of **DTsbe**, and even than the ones of the baseline. **OCT** wins for only one dataset. Such a low performance of **OCT** can be attributed to oversearching, as already observed for $\mathbf{DTaccept}_0$ in Table 2. The much better generalization error of **OCT** reported by Bertsimas and Dunn (2017) can instead be attributed to the parameter focusing approach used there to set the complexity penaly *cp* in the regularization of the empirical misclassification error.

8. Conclusions

We have introduced an original pruning algorithm of the search space of feature subsets which allows for enumerating all and only the distinct decision trees. The lattice of feature subsets is explored by distinguishing features that must be necessarily used from those that may be possibly used. The order of the visit is impressively effective, and it relies on an estimation of the chances of generating distinct trees. Based on the enumeration of distinct decision trees, we introduced an algorithm for finding δ-acceptable feature subsets, which depart by at most δ from the best estimated error of decision trees built from any feature subset. The framework is stated in general terms for any top-down greedy decision tree induction algorithm, any quality measure used to select split attributes, and any error estimation function. Coupled with a few computational optimizations and a multi-core parallel implementation, this makes it possible to investigate properties of complete search for datasets of up to 60 features. Beyond such a limit, we contributed by exploiting ideas and optimizations proposed in the paper to a white-box computational improvement of the sequential backward elimination heuristic, which extends its practical applicability to large

13. OCT parameters: 8 core processes in Julia, max depth = 5, other parameters set to default (min bucket = 1, local search = true, *cp* = 0.01). No focusing of parameters performed. Non-binary discrete features of experimental datasets are processed with one-hot encoding as required by OCT.

dimensional datasets. Experimental results reinforce, in the case of decision trees, previous findings that oversearching increases overfitting, and, in addition, they highlight that oversearching also increases instability. Sequential backward elimination performs better than complete search over the feature subsets, and also better than the OCT optimal (non-greedy) decision tree algorithm. It also improves the generalization error of random forest models for medium-to-large dimensional datasets.

Acknowledgments. I am grateful to the anonymous reviewers for several helpful comments. I also thank Jack Dunn for providing the OCT system.

References

M. Aldinucci, S. Ruggieri, and M. Torquati. Decision tree building on multi-core using FastFlow. *Concurrency and Computation: Practice and Experience*, 26(3):800–820, 2014.

E. Amaldi and V. Kann. On the approximation of minimizing non zero variables or unsatisfied relations in linear systems. *Theoretical Computer Science*, 209:237–260, 1998.

D. Bertsimas and J. Dunn. Optimal classification trees. *Machine Learning*, 106(7):1039–1082, 2017.

A. Blum and P. Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 97(1-2):245–271, 1997.

V. Bolón-Canedo, N. Sánchez-Maroño, and A. Alonso-Betanzos. A review of feature selection methods on synthetic data. *Knowledge and Information Systems*, 34(3):483–519, 2013.

L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth Publishing Company, 1984.

L. A. Breslow and D. W. Aha. Simplifying decision trees: A survey. *The Knowledge Engineering Review*, 12:1–40, 1997.

R. Caruana and D. Freitag. Greedy attribute selection. In *Proc. of the Int. Conf. on Machine Learning (ICML 1994)*, pages 28–36. Morgan Kaufmann, 1994.

R. Caruana and A. Niculescu-Mizil. An empirical comparison of supervised learning algorithms. In *Proc. of the Int. Conf. on Machine Learning (ICML 2006)*, volume 148, pages 161–168. ACM, 2006.

T. M. Cover. On the possible ordering on the measurement selection problem. *Trans. Systems, Man, and Cybernetics*, 9:657–661, 1977.

M. Dash and H. Liu. Feature selection for classification. *Intelligent Data Analysis*, 1(1-4): 131–156, 1997.

M. Fernández Delgado, E. Cernadas, S. Barro, and D. Gomes Amorim. Do we need hundreds of classifiers to solve real world classification problems? *Journal of Machine Learning Research*, 15(1):3133–3181, 2014.

L. Devroye, L. Györfi, and G. A. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.

J. Doak. An evaluation of feature selection methods and their application to computer security. Technical Report CSE-92-18, University of California Davis, 1992.

F. Esposito, D. Malerba, and G. Semeraro. A comparative analysis of methods for pruning decision trees. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(5):476–491, 1997.

L. Fan. Accurate robust and efficient error estimation for decision trees. In *Proc. of the Int. Conf. on Machine Learning (ICML 2016)*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 239–247, 2016.

I. Foroutan and J. Sklansky. Feature selection for automatic classification of non-gaussian data. *Trans. Systems, Man, and Cybernetics*, 17(2):187–198, 1987.

R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, D. Pedreschi, and F. Giannotti. A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5):93:1–93:42, December 2018.

I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.

I. Guyon, S. Gunn, A. Ben-Hur, and G. Dror. Design and analysis of the NIPS2003 challenge. In I. Guyon, M. Nikravesh, S. Gunn, and L. A. Zadeh, editors, *Feature Extraction: Foundations and Applications*, pages 237–263. Springer, 2006a.

I. Guyon, M. Nikravesh, S. Gunn, and L. A. Zadeh, editors. *Feature Extraction: Foundations and Applications*, volume 207 of *Studies in Fuzziness and Soft Computing*. Springer, 2006b.

J. Huysmans, K. Dejaeger, C. Mues, J. Vanthienen, and B. Baesens. An empirical evaluation of the comprehensibility of decision table, tree and rule based predictive models. *Decision Support Systems*, 51(1):141–154, 2011.

G. H. John, R. Kohavi, and K. Pfleger. Irrelevant features and the subset selection problem. In *Proc. of the Int. Conf. on Machine Learning (ICML 1994)*, pages 121–129. Morgan Kaufmann, 1994.

J.-H. Kim. Estimating classification error rate: Repeated cross-validation, repeated hold-out and bootstrap. *Computational Statistics & Data Analysis*, 53(11):3735–3745, 2009.

R. Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI 1995)*, pages 1137–1145. Morgan Kaufmann, 1995.

R. Kohavi and G. H. John. Wrappers for feature subset selection. *Artificial Intelligence*, 97 (1-2):273–324, 1997.

T. N. Lal, O. Chapelle, J. Weston, and A. Elisseeff. Embedded methods. In I. Guyon, M. Nikravesh, S. Gunn, and L. A. Zadeh, editors, *Feature Extraction: Foundations and Applications*, pages 137–165. Springer, 2006.

M. Lichman. UCI machine learning repository, 2013. http://archive.ics.uci.edu/ml.

H. Liu and L. Yu. Toward integrating feature selection algorithms for classification and clustering. *IEEE Transactions on Knowledge and Data Engineering*, 17(4):491–502, 2005.

H. Liu, H. Motoda, and M. Dash. A monotonic measure for optimal feature selection. In *Proc. of the European Conf. on Machine Learning (ECML 1998)*, volume 1398 of *Lecture Notes in Computer Science*, pages 101–106. Springer, 1998.

M. Menickelly, O. Günlük, J. Kalagnanam, and K. Scheinberg. Optimal generalized decision trees via integer programming. *CoRR*, abs/1612.03225, 2016.

A. Miller. *Subset Selection in Regression*. Chapman and Hall, 2 edition, 2002.

S. K. Murthy. Automatic construction of decision trees from data: A multi-disciplinary survey. *Data Mining and Knowledge Discovery*, 2:345–389, 1998.

N. Narodytska, A. Ignatiev, F. Pereira, and J. Marques-Silva. Learning optimal decision trees with SAT. In *Proc. of Int. Joint Conf. on Artificial Intelligence (IJCAI 2018)*, pages 1362–1368. ijcai.org, 2018.

S. Nogueira and G. Brown. Measuring the stability of feature selection. In *Proc. of Machine Learning and Knowledge Discovery in Databases (ECML-PKDD 2016) Part II*, volume 9852 of *LNCS*, pages 442–457, 2016.

J. R. Quinlan. Induction of decision trees. *Machine Learning*, 1:81–106, 1986.

J. R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, CA, 1993.

J. R. Quinlan and M. Cameron-Jones. Oversearching and layered search in empirical learning. In *Proc. of Int. Joint Conf. on Artificial Intelligence (IJCAI 1995)*, pages 1019–1024. Morgan Kaufmann, 1995.

J. Reunanen. Overfitting in making comparisons between variable selection methods. *Journal of Machine Learning Research*, 3:1371–1382, 2003.

S. Ruggieri. Efficient C4.5. *IEEE Transactions on Knowledge and Data Engineering*, 14: 438–444, 2002.

S. Ruggieri. YaDT: Yet another Decision tree Builder. In *Proc. of Int. Conf. on Tools with Artificial Intelligence (ICTAI 2004)*, pages 260–265. IEEE, 2004.

S. Ruggieri. Subtree replacement in decision tree simplification. In *Proc. of the SIAM Conf. on Data Mining (SDM 2012)*, pages 379–390. SIAM, 2012.

S. Ruggieri. Enumerating distinct decision trees. In *Proc. of the Int. Conf. on Machine Learning (ICML 2017)*, number 70 in JMLR Workshop and Conference Proceedings, pages 2960–2968, 2017.

M. Saar-Tsechansky and F. Provost. Handling missing values when applying classification models. *Journal of Machine Learning Research*, 8:1625–1657, 2007.

S. S. Skiena. *The Algorithm Design Manual*. Springer, 2 edition, 2008.

M. Stone. Asymptotics for and against cross-validation. *Biometrika*, 64:29–?35, 1997.

B. Twala. An empirical comparison of techniques for handling incomplete data using decision trees. *Applied Artificial Intelligence*, 23(5):373–405, 2009.

S. Verwer and Y. Zhang. Learning decision trees with flexible constraints and objectives using integer optimization. In *Proc. of Int. Conf. on Integration of AI and OR Techniques in Constraint Programming (CPAIOR 2017)*, volume 10335 of *Lecture Notes in Computer Science*, pages 94–103. Springer, 2017.

Table 6: Cross-validation errors (mean ± stdev). IG+EBP and m=2.

dataset	baseline	DTsbe	DTaccept$_0$	RF	DTsbe+RF
Adult	14.62 ± 0.38	14.18 ± 0.44	**14.09 ± 0.40**	**14.04 ± 0.38**	14.21 ± 0.45
Letter	11.91 ± 0.71*	**11.90 ± 0.70**	11.95 ± 0.75*	**4.13 ± 0.50**	4.77 ± 0.65
Hypo	**0.46 ± 0.31**	0.53 ± 0.38	0.57 ± 0.40	0.83 ± 0.44	**0.50 ± 0.38**
Ionosphere	11.48 ± 5.87	**9.58 ± 4.51**	11.06 ± 5.58	**6.53 ± 4.08**	8.52 ± 4.46
Soybean	**9.87 ± 3.70**	12.31 ± 3.80	11.51 ± 4.20	**8.64 ± 2.96**	11.83 ± 3.78
Kr-vs-kp	**0.64 ± 0.47**	1.06 ± 0.57	-	1.94 ± 0.81*	**1.82 ± 0.84**
Anneal	**0.93 ± 0.86**	1.41 ± 1.26	2.22 ± 1.99	**1.18 ± 1.02**	1.65 ± 1.50
Census	5.14 ± 0.10	**4.93 ± 0.10**	-	**4.87 ± 0.07**	4.91 ± 0.12
Spambase	7.06 ± 1.23*	**6.95 ± 1.21**	-	**5.64 ± 1.19**	5.81 ± 1.06
Sonar	**25.97 ± 9.34**	25.99 ± 9.08*	26.69 ± 9.40*	**17.50 ± 8.48**	22.88 ± 8.67
Optdigits	**9.63 ± 1.20**	9.90 ± 1.67*	-	**2.00 ± 0.71**	2.79 ± 0.84
Coil2000	5.98 ± 0.11*	**5.97 ± 0.05**	-	**5.97 ± 0.05**	5.97 ± 0.05*
Clean1	**17.53 ± 6.87**	19.41 ± 7.15	-	**10.73 ± 4.46**	15.35 ± 5.55
Clean2	3.29 ± 0.67	**2.76 ± 0.71**	-	2.41 ± 0.54*	**2.38 ± 0.55**
Madelon	25.73 ± 3.45	**21.77 ± 3.39**	-	35.72 ± 3.06	**20.03 ± 3.06**
Isolet	16.68 ± 1.15	**16.34 ± 1.36**	-	**6.11 ± 0.78**	6.35 ± 0.80
Gisette	5.94 ± 0.83	**5.57 ± 0.87**	-	4.01 ± 0.72	**3.76 ± 0.76**
P53mutants	0.52 ± 0.09	**0.49 ± 0.10**	-	0.46 ± 0.04*	**0.45 ± 0.07**
Arcene	**48.09 ± 5.39**	48.57 ± 4.96*	-	46.34 ± 3.25	**44.51 ± 2.99**
Dexter	48.92 ± 2.94*	**48.13 ± 3.51**	-	48.96 ± 1.89	**43.74 ± 3.00**
wins/$\star$	8/4	11/3	1/2	12/3	8/1

Appendix A. Additional Experimental Results

First, we consider tree simplification, a well-known strategy to address the common problems of overfitting the training data and of trading accuracy for simplicity. Decision tree pruning does not satisfy Assumption 5, hence it cannot be embedded within the heuristic or complete search procedures. We apply then the C4.5 Error Based Pruning (EBP) (Breslow and Aha, 1997) approach as a post-processing on the decision trees built over the feature subset selected by the strategies considered. The running time added by such a post-processing is negligible. Table 6 reports the generalization errors, including, for comparison purposes, also the random forest models. The winners/starred strategies are almost identical to Table 2. Generalization errors are now almost always smaller for all strategies (for *Coil2000* the improvement is vast), confirming that tree simplification is an effective strategy (Ruggieri, 2012). Finally, contrasting Table 6 with Table 3, there is no evidence that adding pruning to random forests improves their accuracy.

Next, we consider whether the results extend to quality measures other than IG. Table 7 and Table 8 report the elapsed times and the generalization errors for the Gain Ratio (GR) (see footnote 1). Running times are generally greater than the ones of IG shown in Table 1 – up to 2.5× for **DTsbe** and up to 10× for **DTaccept**$_0$. Regarding generalization errors, the baseline wins or is not stastitically different from the winner in 3 additional cases with respect to IG (cfr. Table 2), thus lowering the gap with the search heuristics. Also, the benefits of adding **DTsbe** to random forests for large dimensional datasets is not as marked as in Table 3 and in Table 6. The improved generalization performances can be attributed to the better discriminative power of Gain Ratio over Information Gain.

Table 7: Average elapsed running times (seconds). GR and m=2.

dataset	DTsbe	SBE	ratio	DTaccept$_0$	RF
Adult	0.78	1.40	0.559	48.68	1.97
Letter	0.47	0.77	0.617	302.62	0.90
Hypo	0.06	0.22	0.282	309.80	0.05
Ionosphere	0.01	0.10	0.106	1.01	0.01
Soybean	0.05	0.15	0.312	>1h	0.03
Kr-vs-kp	0.07	0.23	0.285	>1h	0.07
Anneal	0.03	0.14	0.196	185.51	0.02
Census	51.23	135.70	0.378	>1h	21.03
Spambase	0.85	5.11	0.167	>1h	0.13
Sonar	0.01	0.42	0.030	22.42	0.01
Optdigits	0.61	4.96	0.123	>1h	0.14
Coil2000	2.77	24.09	0.115	>1h	0.32
Clean1	0.19	25.79	0.007	>1h	0.03
Clean2	7.50	195.43	0.038	>1h	0.24
Madelon	21.81	>1h	-	>1h	0.28
Isolet	492.55	>1h	-	>1h	2.29
Gisette	211.04	>1h	-	>1h	11.08
P53mutants	1012.83	>1h	-	>1h	9.70
Arcene	66.52	>1h	-	>1h	4.58
Dexter	1656.07	>1h	-	>1h	0.18

Table 8: Cross-validation errors (mean ± stdev). GR and m=2.

dataset	baseline	DTsbe	DTaccept$_0$	RF	DTsbe+RF
Adult	14.62 ± 0.43	13.91 ± 0.50	**13.77 ± 0.44**	**13.50 ± 0.40**	13.78 ± 0.48
Letter	12.76 ± 0.78	12.25 ± 0.77	**12.14 ± 0.80**	**3.78 ± 0.41**	4.42 ± 0.66
Hypo	**0.47 ± 0.35**	0.52 ± 0.42*	0.54 ± 0.41*	0.54 ± 0.38*	**0.48 ± 0.39**
Ionosphere	10.57 ± 5.01*	**9.94 ± 4.93**	11.46 ± 5.05	**6.21 ± 3.83**	8.60 ± 4.70
Soybean	**8.45 ± 3.02**	9.18 ± 3.43	-	**5.69 ± 2.31**	7.92 ± 2.71
Kr-vs-kp	**0.63 ± 0.52**	1.05 ± 0.57	-	**1.03 ± 0.59**	1.12 ± 0.66*
Anneal	**1.34 ± 1.16**	1.85 ± 1.53	2.31 ± 1.85	**0.62 ± 0.88**	1.82 ± 1.65
Census	5.50 ± 0.10	**5.03 ± 0.29**	-	**4.51 ± 0.08**	4.67 ± 0.15
Spambase	7.57 ± 1.28*	**7.37 ± 1.25**	-	**5.44 ± 1.04**	5.63 ± 1.03
Sonar	25.63 ± 9.08*	**25.14 ± 8.83**	26.17 ± 8.97*	**19.70 ± 8.01**	23.61 ± 9.27
Optdigits	**10.23 ± 1.65**	10.60 ± 1.91	-	**2.01 ± 0.63**	2.59 ± 0.76
Coil2000	9.00 ± 0.70*	**8.88 ± 0.73**	-	**6.01 ± 0.10**	6.05 ± 0.14
Clean1	**17.68 ± 5.93**	18.59 ± 6.02*	-	**9.93 ± 4.20**	13.15 ± 4.91
Clean2	3.18 ± 0.72	**2.91 ± 0.71**	-	**2.36 ± 0.56**	2.38 ± 0.55*
Madelon	30.40 ± 4.05	**24.96 ± 3.73**	-	35.35 ± 2.53	**23.00 ± 3.73**
Isolet	**16.51 ± 1.27**	16.58 ± 1.16*	-	**6.38 ± 0.74**	6.44 ± 0.82*
Gisette	6.62 ± 0.99	**5.97 ± 1.02**	-	**4.10 ± 0.74**	4.20 ± 0.81*
P53mutants	0.59 ± 0.11	**0.52 ± 0.09**	-	**0.45 ± 0.04**	0.45 ± 0.06*
Arcene	48.67 ± 5.64*	**47.50 ± 4.85**	-	45.28 ± 2.52	**44.37 ± 1.91**
Dexter	48.55 ± 2.56*	**48.25 ± 3.29**	-	49.36 ± 1.56	**45.08 ± 2.80**
wins/$\star$	7/6	11/3	2/2	16/1	4/5