

rmcfs: An *R* Package for Monte Carlo Feature Selection and Interdependency Discovery

Michał Dramiński¹ and Jacek Koronacki¹

¹ Institute of Computer Science, Polish Academy of Sciences,
Jana Kazimierza 5, 01-248 Warsaw, Poland, E-mail:
`michal.draminski@ipipan.waw.pl`, URL:
`http://www.ipipan.eu/staff/m.draminski/`

February 7, 2018

Abstract

We describe an *R* package *rmcfs* that implements an algorithm for ranking features from high dimensional data according to their importance for a given supervised classification task. The ranking is performed prior to addressing the classification task per se. This *R* package is the new and extended version of the MCFS (Monte Carlo feature selection) algorithm whose early version was published in 2005. The package provides an easy *R* interface, a set of tools to review results and the new ID (Interdependency Discovery) component. The algorithm can be used on continuous and/or categorical features (e.g., gene expression and phenotypic data) to produce objective ranking of features with a statistically well-defined cutoff between informative and non-informative ones. Moreover, the directed ID-Graph that presents interdependencies between informative features is provided.

Keywords: MCFS-ID, feature selection, high-dimensional problems, *Java*, *R*, ID-Graph

1 Introduction

In the area of feature ranking and selection for high-dimensional supervised classification, a very significant progress has been achieved in the past two

decades. For a brief account, up to 2002, see [Dudoit and Fridlyand(2003)] and for an extensive survey and somewhat later developments see [Saeys *et al.*(2007)]. Without coming to details let us note that feature selection can be *wrapped* around the classifier construction or directly built (*embedded*) into the classifier construction, and not performed prior to addressing the classification task per se by *filtering* out noisy features first and keeping only informative ones for building a classifier. An early and successful method with embedded feature selection included, not mentioned by [Saeys *et al.*(2007)], was developed by Tibshirani *et al.* (see [Tibshirani *et al.*(2002)], [Tibshirani *et al.*(2003)]). More recently and within non-filter approaches, a Bayesian technique of automatic relevance determination, the use of support vector machines, and the use of ensembles of classifiers, all these either alone or in combination, have proved promising. For further details see [Li *et al.*(2002)], [Lu *et al.*(2007)], [Chrysostomou *et al.*(2008)] and the literature there.

Moreover, the last developments by the late Leo Breiman deserve special attention. In his Random Forests (RFs), he proposed to make use of the so-called variable (i.e., feature) importance for feature selection. Determination of the importance of the variable is not necessary for random forest construction, but it is a subroutine performed in parallel to building the forest; cf. [Breiman and Cutler(2008)]. Ranking features by variable importance can thus be considered to be a by-product of building the classifier. At the same time, nothing prevents one from using such variable importances within, say, the embedded approach; cf., e.g., [Díaz-Uriarte and De Andres(2006)]. In any case, feature selection by measuring variable importance in random forests should be seen as a very promising method, albeit under one proviso. Namely, the problem with variable importance as originally defined is that it is biased towards variables with many categories; cf. [Strobl *et al.*(2007)], [Archer and Kimes(2008)], [Nicodemus *et al.*(2010)]. Accordingly, proper debiasing is needed, in order to obtain true ranking of features; cf. [Strobl *et al.*(2008)]. And, however sound such debiasing may be, it incurs much additional computational cost. For an excellent and recent survey on RFs, their properties and capabilities, see [Ziegler and König(2014)].

Most recently, much work has been done to: (i) give embedded feature selection procedures, in particular those used within RFs (whether biased or unbiased), a clear statistical meaning; and (ii) better understand RFs' capability of discovering interdependencies between features; cf. [Paul and Dupont(2015)] (see also [Huynh-Thu *et al.*(2012)] and the literature there for (i)) and [Wright *et al.*(2016)] and the literature there for (ii).

In 2005, a novel, effective and reliable method for ranking features according to their importance for a given supervised classification task has been introduced by [Drański *et al.*(2005)]. The method, which relies on

Monte Carlo approach to select informative features and was fully developed in [Dramiński *et al.*(2008)] as Monte Carlo feature selection algorithm or MCFS, is capable of incorporating interdependencies between features. It bears some remote similarity to the RF methodology, but differs entirely in the way features ranking is performed. Specifically, our method does not require debiasing (cf. [Dramiński *et al.*(2010)]) and is conceptually simpler. A more important and newer result is that it provides explicit information about interdependencies among features (cf. [Dramiński *et al.*(2010)]), where an early version of the MCFS algorithm with an interdependency discovery, or ID, component has been introduced; see also [Kierczak *et al.*(2009)], [Kierczak *et al.*(2010)]).

In this paper, we present an *R* package that implements the most recent version of the MCFS-ID algorithm. In particular, the ideas from [Dramiński *et al.*(2010)] have been substantially expanded in [Dramiński *et al.*(2016)] by providing not only the ranking of features but also the directed ID-Graph that presents interdependencies between informative features. Within our approach, discovering interdependencies builds on identifying features which "cooperate" in determining that some samples belong to one class, another samples to another class, still another to still another class and so on. It is worthwhile to emphasize that this is completely different from the usual approach which aims at finding features that are similar in some sense. Instead, it can be said that our way to discover interdependencies between features amounts to determining multidimensional dependency between the classes and sequences of features. In this sense, we are in fact interested in *contextual (or predictive) interdependencies between features*, since the dependency in question requires the context of class information.

Let us emphasize that we do not aim at classification. While in our approach we heavily rely on using classifiers, we do not use them for the classification task per se. Indeed, we use classifiers only to: (i) rank features according to their importance with respect to their discriminative power to distinguish between classes; (ii) discover interdependencies between features. Given the top features found in step (i), one can later use them for classification by any classifier, but this is neither required nor of our interest. And clearly, step (ii) is aimed at something vastly different from sheer solving the classification task.

The procedure is particularly well suited to dealing with high-dimensional data including "small n large p problems", i.e., those with a small number of objects (records, samples) versus several orders of magnitude greater number of features for each object.

The procedure from [Dramiński *et al.*(2008)] and [Dramiński *et al.*(2016)] for Monte Carlo feature selection and interdependency discovery is briefly

recapitulated in Section 2. In Section 3, an overview of *rmcfs* package and its major *R* functions is provided. In Section 4 we demonstrate the use of the package *rmcfs* by performing mcfs algorithm and building ID-Graph for simulated data. We close with concluding remarks in Section 5.

2 MCFS-ID algorithm

2.1 Feature selection - the MCFS part

We begin with a brief recapitulation of our MCFS; see [Dramiński *et al.*(2008)], which can be consulted for details as well as rationale and statistical validation of our approach to feature selection.

We consider a particular feature to be important, or informative, if it is likely to take part in the process of classifying samples into classes "more often than not". This "readiness" of a feature to take part in the classification process, termed relative importance of a feature, is measured via intensive use of classification trees. In the main step of the procedure, we estimate relative importance of features by constructing thousands of trees for randomly selected subsets of features.

More precisely, out of all d features, s subsets of m features are randomly selected, m being fixed and $m \ll d$, and for each subset of features, t trees are constructed and their performance is assessed (one can easily see that the procedure is essentially the same as that of the Random Subspace Method, the fact the authors were not aware of at the time they wrote their proposal; cf [Ho(1998)]).

Each of the t trees in the inner loop is trained and evaluated on a different, randomly selected training and test sets that come from a split of the full set of training data into two subsets: each time, out of all n samples, $2/3$ of the samples are drawn at random for training in such a way as to preserve proportions of classes from the full set of training data, and the remaining samples are used for testing. See Figure 1 for a block diagram of the procedure.

The relative importance of feature g_k , RI_{g_k} , is defined as:

$$\text{RI}_{g_k} = \sum_{\tau=1}^{s \cdot t} \text{wAcc}_{\tau}^u \sum_{n_{g_k}(\tau)} \text{GR}(n_{g_k}(\tau)) \left(\frac{\text{no. in } n_{g_k}(\tau)}{\text{no. in } \tau} \right)^v, \quad (1)$$

where summation is over all $s \cdot t$ trees and, within each τ -th tree, over all nodes $n_{g_k}(\tau)$ of that tree on which the split is made on feature g_k , wAcc_{τ}^u stands for the weighted accuracy of the τ 's tree, $\text{GR}(n_{g_k}(\tau))$ stands for gain ratio for node $n_{g_k}(\tau)$, $(\text{no. in } n_{g_k}(\tau))$ denotes the number of samples

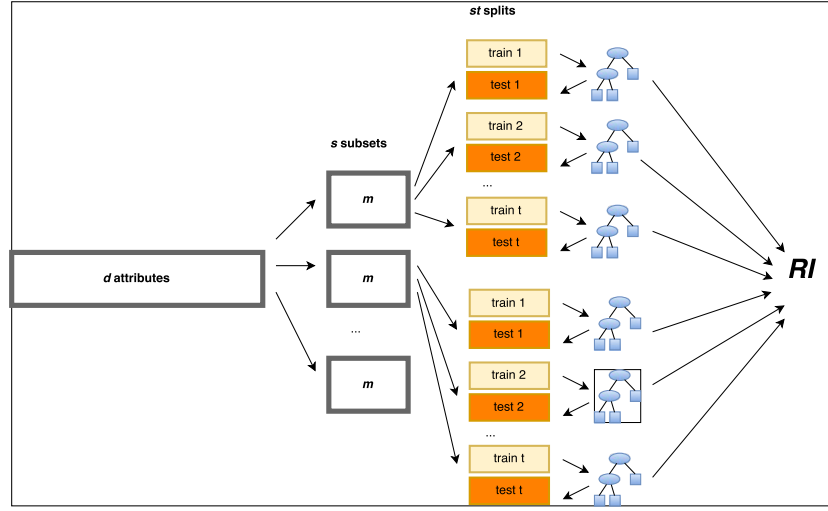


Figure 1: Block diagram of the main step of the MCFS procedure.

in node $n_{g_k}(\tau)$, (no. in τ) denotes the number of samples in the root of the τ -th tree, and u and v are fixed positive reals (now set to 1 by default; cf. [Dramiński *et al.*(2010)]). The normalizing factor (no. in τ), which has the same value for all τ , has been included mainly for computational reasons.

With u and v set to 1, there are three parameters, m , s and t to be set by an experimenter. Note that, overall, $s \cdot t$ trees are constructed and evaluated in the main step of the procedure. Both s and t should be sufficiently large, so that each feature has a chance to appear in many different subsets of features and randomness due to inherent variability in the data is properly accounted for. The choice of subset size m of features selected for each series of t experiments should take into account the trade-off between the need to prevent informative features from being masked too severely by the relatively most important ones and the natural requirement that s be not too large. Indeed, the smaller m , the smaller the chance of masking the occurrence of a feature. However, a larger s is then needed, since all features should have a high chance of being selected into many subsets of the features. For classification problems of dimension d ranging from several thousands to hundreds of thousands, we have found that taking m equal to a few hundreds (say, $m = 300 - 500$) and t equal to maximum 20 (even $t = 5$ usually suffices) is a good choice in terms of reliability and overall computational cost of the procedure. Finally, for a given m , s can be made a running parameter of the procedure, and the procedure executed for increasing s until the rankings of top scoring $p\%$ features prove (almost) the same for successive values of the s .

2.2 Determining the cutoff value

The above procedure provides one with a ranking of features. However, this ranking as such does not enable one to discern between informative and non-informative features. A cutoff between these two types of features is needed and we propose to determine it by one of 5 different methods (one of them being still under development). Note that finding the cutoff point is related to separating features with large enough RI values from the rest.

Available methods:

- **criticalAngle** - The critical angle method is based on the plot of the features' RIs in decreasing order of size, with the corresponding features equally spaced along the abscissa. The plot can be seen as piecewise linear function, where each linear segment joins two neighboring RIs. Roughly speaking, the cutoff (placed on the abscissa) corresponds to this point on the plot where the slope of consecutive segments changes significantly and lastingly from large to small.
- **Kmeans** - The method is based on clustering the RI values into two clusters by the k-means algorithm. It sets the cutoff where the two clusters are separated. This method is quite valuable when data contains a subset of very informative features.
- **Permutations (max RI)** - The method consists in permuting the decision attribute at least 20 times and running the mcfs algorithm for each permutation. The set of the maximal RIs from all these experiments is assumed approximately normally distributed and a critical value based on the one-sided (upper-tailed) Student's t test (at 95% significance level) is provided. A feature is declared informative if its RI in the original ranking (without any permutation) exceeds the obtained critical value. A more detailed description of this method is included in [Dramiński *et al.*(2010)].
- **Permutations (z-score)** - The method consists in permuting the decision attribute at least 30 times and running the mcfs algorithm for each permutation. For each feature, a separate distribution of the obtained RI values is constructed. For each feature, its RI in the original ranking (without any permutation) is compared against the corresponding distribution of RIs from the experiments with permuted decision attribute (the z test is used this time). This method has been described and studied in [Bornelöv and Komorowski(2016)]. It gives similar results to the previous one but needs more mcfs runs. Along with that, however, it gives a separate p value for each feature.

- `contrastAttributes` - The method consists in permuting all d features to obtain the so-called contrast features, including them into the dataset and running the `mcf`s. Given RIs of the contrast features, a statistical test can be used to find a cutoff between the informative and non-informative original features. This method is under development and is currently off.

Given our experience and taking into account its sound statistical foundation, we recommend most the Permutations (max RI) method. However, it may sometimes prove very conservative in the sense that it may give a highly limited set of informative features. On the other hand, if there are no meaningful features, the method will discover this fact and the set of informative features will prove empty. Finally, one has to admit that using the Permutations method requires much caution. Actually, proper significance level should be chosen adaptively, depending on the data under consideration. Instead of adding this kind of complexity to the way the cutoff is determined, we set (by default) the final cutoff value to be the mean of all the obtained cutoffs.

2.3 Interdependency discovery - the ID part

Once features are ranked by the MCFS procedure, a natural issue to be raised concerns possible interdependencies among the informative features. In this subsection, we describe shortly a way to present such interdependencies in the form of a directed graph.

The interdependencies among features are often modeled using interactions, similarly as in experimental design and analysis of variance. Perhaps the most widely used approach to recognizing interdependencies is finding correlations between features or finding groups of features that behave in some sense similarly across samples. A classical bioinformatics example of this problem is finding co-regulated features, most often genes or, rather more precisely, their expression profiles. Searching groups of similar features is usually done with the help of various clustering techniques, frequently specially tailored to a task at hand. See [Smyth *et al.*(2003)], [Hastie *et al.*(2001)], [Saeys *et al.*(2007)], [Gyenesi *et al.*(2007)] and the literature there. Our approach to interdependency discovery is significantly different in that we focus on identifying features that "cooperate" in determining that a sample belongs to a particular class. The initial idea was presented in [Dramiński *et al.*(2010)] but the current version provides directed ID-Graphs and is based on an improved interdependency measure (cf. [Dramiński *et al.*(2016)]).

Our ID-Graph is based on aggregating information provided by all the $s \cdot t$ trees (see Figure 1). However, let us begin by noting that, for a single classification tree, a set of decision rules defining each class is provided by this tree's paths. Each decision rule is produced as a conjunction of conditions imposed on the particular features which in fact point to conditional interdependencies between the features. Our trust in the decision rules that are learned by any single rule-based classifier, and thus in the discovered interdependencies, is naturally limited. Moreover, the classifier is trained on just one training set and therefore our conclusions are necessarily dependent on the classifier and are conditional upon the training set. In the case of classification trees, the problem is aggravated by their high variance, i.e., their tendency to provide varying results even for slightly different training sets. Accordingly, in order to overcome these problems and provide more objective results, the MCFS-ID procedure rests on building thousands of classification trees.

To see how an ID-Graph is built, let us recall that each node in each of the multitude of classification trees represents a feature on which a split is made. Now, for each node in each classification tree its all antecedent nodes can be taken into account along the path to which the node belongs; note that each node in a tree has only one parent and thus for the given node we simply consider its parent, then the parent of the parent and so on. In practice, the maximum possible depth of such analysis, i.e., the number of antecedents considered, if available before the tree's root is attained, is set to some predetermined value, which is the procedure's parameter (its default value being 5). For each pair [*antecedent node* \rightarrow *given node*] we add one directed edge to our ID-Graph from *antecedent node* to *given node*. Let us emphasize again that a node is equated with the feature it represents and thus any directed edge found is in fact an edge joining two uniquely determined features in a directed way. The edges are found along the paths in all the $s \cdot t$ MCFS-ID trees. Clearly, the same edge can appear more than once even in a single tree.

The strength of the interdependence between two nodes, actually two features, connected by a directed edge, termed ID weight of a given edge, or ID weight for short, is equal to the gain ratio (GR) in the given node multiplied by the fraction of objects in the given node and the antecedent node. Thus, for node $n_k(\tau)$ in τ -th tree, $\tau = 1, \dots, s \cdot t$, and its antecedent node $n_i(\tau)$, ID weight of the directed edge from $n_i(\tau)$ to $n_k(\tau)$, denoted $w[n_i(\tau) \rightarrow n_k(\tau)]$, is equal to

$$w[n_i(\tau) \rightarrow n_k(\tau)] = \text{GR}(n_k(\tau)) \left(\frac{\text{no. in } n_k(\tau)}{\text{no. in } n_i(\tau)} \right), \quad (2)$$

where $\text{GR}(n_k(\tau))$ stands for gain ratio for node $n_k(\tau)$, (no. in $n_k(\tau)$) denotes the number of samples in node $n_k(\tau)$ and (no. in $n_i(\tau)$) denotes the number of samples in node $n_i(\tau)$.

The final ID-Graph is based on the sums of all ID weights for each pair [*antecedent node* \rightarrow *given node*]; i.e., for each directed edge found, its ID weights are summed over all occurrences of this edge in all paths of all MCFS classification trees. For a given edge, it is this sum of ID weights which becomes the ID weight of this edge in the final ID-Graph. The following pseudo code (alg. 1) describes the calculation:

Algorithm 1 ID-Graph building procedure

```

 $w[n_\delta \rightarrow n] = 0$ 
for  $\tau_k \in T$  do
  for  $n \in \tau_k$  do
    for  $\delta \in D$  do
       $n_\delta = \delta$ -th antecedent of  $n$ 
       $w[n_\delta \rightarrow n] = w[n_\delta \rightarrow n] + \text{GR}(n) \left( \frac{\text{no. in } n}{\text{no. in } n_\delta} \right)$ 
    end for
  end for
end for

```

where T denotes the set of all $s \cdot t$ trees and $D = \{1, 2, \dots, \text{depth}\}$ with *depth* being the predetermined number of antecedents considered.

Note that an edge $n_i \rightarrow n_k$ from node n_i to node n_k is directed as is the edge (if found) from n_k to n_i , ($n_k \rightarrow n_i$). Interestingly, in most cases of ID-Graphs, we find that one of such two edges is dominating, i.e., has a much larger ID weight than the other. Whenever it happens, it means that not only n_i and n_k form a sound partial decision (a part of a conjunction rule) but also that their succession in the directed rule is not random.

The ID-Graph is a way to present interdependencies that follow from all of the MCFS classification trees. Each path in a tree represents a decision rule and by analyzing all tree paths we in fact analyze decision rules to find the most frequently observed features that along with other features form good decision rules. The ID-Graph thus presents some patterns that frequently occur in thousands of classification trees built by the MCFS procedure.

In sum, an ID-Graph provides a general roadmap that not only shows all the most variable attributes that allow for efficient classification of the objects but, moreover, it points to possible interdependencies between the attributes and, in particular, to a hierarchy between pairs of attributes. High differentiation of the values of ID weights in the ID-Graph gives strong evidence that some interdependencies between some features are much stronger than

others and that they create some patterns/paths calling for interpretation based on background knowledge.

Let us conclude this subsection with a short comparison of our approach to discovering interdependencies between features and that based on the RFs (cf. [Wright *et al.*(2016)]). Most importantly, we take full advantage of the Monte Carlo mechanism, what makes direct examination of possible interactions superfluous. Indeed, nothing like, e.g., pairwise permutation importance for pairs of features needs to be calculated. Moreover, due to the form of (2), in particular since the summands in the ID weights of pairs of features do not depend in any way on prediction performance of the trees involved in calculation of these summands, we can hope for no masking of interaction effects measured by ID weights by marginal effects (as yet, this last claim has not been confirmed by a thorough simulation study). And, last but not least, ours are directed interactions.

Clearly, only ID Graphs for strong interdependencies (interactions) are of interest. It is readily seen that such graphs can equally easily be constructed for discovering two-way interactions when each of the two features involved or only one of them, or even none of them is of large relative importance, i.e., gives a strong marginal effect (except for the XOR interaction of two features with no marginal effects).

It should be noticed, however, that the ID weights are calculated in a way which makes them incomparable with relative importances of individual features. To put it otherwise, both measures give us only rankings (from the greatest to the smallest) of, respectively: features w.r.t. their relative importance; and strengths of pairwise interactions, albeit without direct information on predictive ability of any interaction.

3 The *R* package *rmcfs*

The *R* package *rmcfs* contains 13 user functions and one example dataset borrowed from [Alizadeh *et al.*(2000)]. All of them are described in standard *R* package documentation (pdf file available as supplementary material) and their description is also available by *R* built-in help system. Here we will focus only on major *rmcfs* functionality covered by the following set of functions:

- ***mcfs()*** performs Monte Carlo feature selection and interdependence discovery (MCFS-ID) on a given data set. It uses C4.5 trees (*j48*) as implemented in WEKA 3-6-10, see [Hall *et al.*(2009)].
- ***plot.mcfs()*** S3 method for class *mcfs* equivalent to ***plot()***. It plots various aspects of the MCFS-ID result.

- ***print.mcfs()*** S3 method for class *mcfs* equivalent to ***print()***. It prints basic information of the MCFS-ID results: top features, cutoff values, confusion matrix obtained for *s·t* trees and classification rules obtained by RIPPER (*JRip*) algorithm.
- ***build.idgraph()*** constructs the ID-Graph based on the result returned by the ***mcfs*** function.
- ***plot.idgraph()*** visualizes the ID-Graph. It is a S3 method for object class *idgraph* equivalent to ***plot()***.

3.1 Function *mcfs*

The function ***mcfs*** can be used to build feature ranking, find the cutoff and evaluate classification performance of a set of top features. It is used as:

```

1 mcfs(formula , data ,
2     projections = "auto",
3     projectionSize = "auto",
4     featureFreq = 150,
5     splits = 5,
6     splitSetSize = 1000,
7     balance = "auto",
8     cutoffMethod = c("permutations", "criticalAngle", "kmeans", "
9         mean"),
10    cutoffPermutations = 20,
11    buildID = TRUE,
12    finalRuleset = TRUE,
13    finalCV = TRUE,
14    finalCVSetSize = 1000,
15    finalCVRepetitions = 3,
16    seed = NA,
17    threadsNumber = 2)

```

and takes the following arguments:

- ***formula*** specifies decision attribute and relation between class and other attributes (e.g., ***class .***).
- ***data*** defines input *data.frame* containing all features with decision attribute included.
- ***projections*** defines the number of subsets with randomly selected features. This parameter is usually set to a few thousands and is denoted in Equation 1 as *s*. By default it is set to ***"auto"*** and this value is based on size of input data set and ***featureFreq*** parameter.

- ***projectionSize*** defines the number of features in one subset. It can be defined by an absolute value (e.g., 100 denotes 100 randomly selected features) or by a fraction of input attributes (e.g., 0.05 denotes 5% of input features). This parameter is denoted in Equation 1 as m . If it is set to "***auto***" then ***projectionSize*** equals to \sqrt{d} , where d is the number of input features. Minimum number of input features in one subset is 1.
- ***featureFreq*** determines how many times each input feature should be randomly selected when ***projections*** = "***auto***". By default each feature should be drawn 150 times.
- ***splits*** defines the number of splits of each subset. This parameter is denoted in Equation 1 as t and by default set on 5.
- ***splitSetSize*** determines whether to limit input dataset size. It helps to speedup computation for data sets with a large number of objects. If the parameter is larger than 1, it determines the number of objects that are drawn at random for each of the $s \cdot t$ decision trees. If ***splitSetSize*** = 0 then the mcfs uses all objects in each iteration.
- ***balance*** determines the way to balance classes. It should be set to 2 or higher if input dataset contains heavily unbalanced classes. Each subset s will contain all the objects from the least frequent class and randomly selected set of objects from each of the remaining classes. This option helps to select features that are important for discovering a relatively rare class. The parameter defines the maximal imbalance ratio. If the ratio is set to 2, then subset s will contain the number of objects from each class (but the least frequent one) proportional to the square root of the class size $\sqrt{size(c)}$. If ***balance*** = 0 then balancing is turned off. If ***balance*** = 1 it is on but does not change the size of classes. Default value is "***auto***".
- ***cutoffMethod*** determines the final cutoff method. Default value is "***permutations***".
- ***cutoffPermutations*** determines the number of permutation runs. It needs at least ***cutoffPermutations*** = 20 for a statistically significant result. Minimum is 3 however 0 turns off permutation method.
- ***buildID*** if = ***TRUE***, Interdependencies Discovery is on and all ID-Graph edges are collected.

- ***finalRuleset*** if $= \text{TRUE}$, classification rules (by RIPPER algorithm) are created on the basis of the final set of features.
- ***finalCV*** if $= \text{TRUE}$, it runs cross validation (cv) experiments on the final set of features. The following set of classifiers is used: C4.5, NB, SVM, kNN, logistic regression and RIPPER.
- ***finalCVSetSize*** limits the number of objects used in the final cv experiment. For each cv repetition, the objects are selected randomly from the uniform distribution.
- ***finalCVRepetitions*** defines the number of repetitions of the cv experiment. The more repetitions, the more stable result.
- ***seed*** seed for random number generator in *Java*. By default seed is random. Replication of the result is possible only if ***threadsNumber*** $= 1$.
- ***threadsNumber*** number of threads to use in computation. More threads needs more CPU cores as well as memory usage is a bit higher. It is recommended to set this value equal to or less than CPU available cores.

Function ***mcfs*** produces S3 object that is a list of the following components:

- ***datainput*** *data.frame* limited to the top important features set.
- ***target*** decision attribute name.
- ***RI*** *data.frame* that contains all features with relevance score sorted from the most relevant to the least relevant. This is the ranking of features.
- ***ID*** *data.frame* that contains features interdependencies as graph edges. It can be converted into a graph object by ***build.idgraph*** function.
- ***distances*** *data.frame* that contains convergence statistics of subsequent projections.
- ***cmatrix*** confusion matrix obtained from all $s \cdot t$ decision trees.
- ***cutoff*** *data.frame* that contains cutoff values obtained by the following methods: criticalAngle, kmeans, permutations (max RI) and the last one *mean* value based of all cutoff values.

- ***cutoff_value*** the number of features chosen as informative by the method defined by parameter ***cutoffMethod***.
- ***cv_accuracy*** *data.frame* that contains classification results obtained by cross validation performed on ***cutoff_value*** features. This *data.frame* exists if ***finalCV = TRUE***.
- ***permutation*** this *data.frame* contains results of permutation experiments: all RI values obtained from all permutation experiments; RI obtained for reference MCFS experiment (i.e, the experiment on the original data); *p* values from Anderson-Darling normality test applied separately for each feature to the ***cutoffPermutations*** RI set; *p* values from Student-t test applied separately for each feature to the ***cutoffPermutations*** RI vs. reference RI. All these *p* values are related to Permutation (z-score) method (see Section 2.2). This *data.frame* exists if ***cutoffPermutations > 0***.
- ***jrip*** classification rules (produced by RIPPER algorithm) and related cv statistics obtained for ***cutoff_value*** features for this algorithm.
- ***params*** all settings used by MCFS-ID.
- ***exec_time*** execution time of MCFS-ID.

Admittedly, one can become overwhelmed by the sheer number of input parameters but in most applications the ***mcfs*** function can be used with the set of default/recommended parameters (as in Section 4).

3.2 Function *plot.mcfs*

Function to plot ***mcfs*** result is defined as follows:

```

1 plot.mcfs(x,
2   type = c("ri", "id", "distances", "features", "cv", "cmatrix",
3     "heatmap"),
4   size = NA,
5   ri_permutations = c("max", "all", "sorted", "none"),
6   diffBars = TRUE,
7   features_margin = 10,
8   cv_measure = c("wacc", "acc", "pearson", "MAE", "RMSE", "SMAPE",
9     " "),
10  heatmap_norm = c("none", "norm", "scale"),
11  heatmap_fun = c("median", "mean"),
12  heatmap_colors = c("white", "red"),
13  cex = 1, ...)
```

and takes the following arguments:

- ***x*** *mcfs* S3 object - the result of the MCFS-ID experiment returned by *mcfs* function.
- ***type***
 - ***ri*** plots top features set with their RIs as well as max RI obtained from permutation experiments. Red color denotes important features.
 - ***id*** plots top ID values obtained from the MCFS-ID.
 - ***distances*** plots distances (convergence diagnostics of the algorithm) between subsequent feature rankings obtained during the MCFS-ID experiment.
 - ***features*** plots top features set along with their RI. It is a horizontal barplot that shows important features in red color and non important in grey.
 - ***cv*** plots cross validation results based on the top features.
 - ***cmatrix*** plots the confusion matrix obtained on all $s \cdot t$ trees.
 - ***heatmap*** plots heatmap results based on top features. Only numeric features can be presented on the heatmap.
- ***size*** number of features to plot.
- ***ri_permutations*** if ***type*** = "***ri***" and ***ri_permutations*** = "***max***", then it additionally shows horizontal lines that correspond to max RI values obtained from each single permutation experiment.
- ***diffBars*** if ***type*** = "***ri***" or ***type*** = "***id***" and ***diffBars*** = ***TRUE***, then it shows difference values for RI or ID values.
- ***features_margin*** if ***type*** = "***features***", then it determines the size of the left margin of the plot.
- ***cv_measure*** if ***type*** = "***cv***", then it determines the type of accuracy shown in the plot: weighted or unweighted/balanced accuracy ("***wacc***" or "***acc***"). If target attribute is numeric it is possible to review one of the following prediction quality measures: ("***pearson***", "***MAE***", "***RMSE***", "***SMAPE***").

- ***heatmap_norm*** if ***type = "heatmap"***, then it defines type of input data normalization ***"none"*** - without any normalization, ***"norm"*** - normalization within range [-1,1], ***"scale"*** - standardization/centering by mean and stdev.
- ***heatmap_fun*** if ***type = "heatmap"***, then it determines calculation ***"mean"*** or ***"median"*** within the class to be shown as heatmap color intensity.
- ***heatmap_colors*** if ***type = "heatmap"***, then it defines low and hi colors on the heatmap.
- ***cex*** size of fonts.
- ... additional plotting parameters.

3.3 Function ***build.idgraph***

Function to build ID-Graph is defined as follows:

```
1 build.idgraph(mcfs_result ,  
2   size = NA,  
3   size_ID = NA,  
4   self_ID = FALSE,  
5   plot_all_nodes = FALSE,  
6   size_ID_mult = 3,  
7   size_ID_max = 100)
```

and takes the following arguments:

- ***mcfs_result*** results returned by ***mcfs*** function.
- ***size*** number of top features to select. If ***size = NA***, then ***size*** is defined by ***mcfs_result\$cutoff_value*** parameter.
- ***size_ID*** number of interdependencies (edges in ID-Graph) to be included. If ***size_ID=NA***, then parameter ***size_ID*** is defined by multiplication ***size_ID_mult*size***.
- ***self_ID*** if = ***TRUE***, then include self-loops from ID-Graph.
- ***plot_all_nodes*** if = ***TRUE***, then include all nodes, even if they are not connected to any other node (isolated nodes).

- ***size_ID_mult*** If ***size_ID_mult*** = 3 there will be 3 times more edges than features (nodes) presented in the ID-Graph. It works only if ***size*** = ***NA*** and ***size_ID*** = ***NA***.
- ***size_ID_max*** maximum number of interactions to be included in the ID-Graph (the upper limit).

It produces S3 *idgraph/igraph* object that can be: plotted in R, exported to graphML (in XML format) or saved as csv or rds files.

3.4 Function *plot.idgraph*

Function to plot ID-Graph is defined as follows:

```
1 plot.idgraph(x,  
2   label.dist = 0.5,  
3   cex = 1)
```

and takes the following arguments:

- ***graph*** *idgraph/igraph* S3 object representing feature interdependencies. This object is produced by ***build.idgraph*** function.
- ***label.dist*** space between the node's label and the corresponding node in the plot.
- ***cex*** size of fonts.

4 Example

First of all make sure you have *Java* installed on your computer and then install *rmcfs* from CRAN repository:

```
1 R> install.packages("rmcfs")
```

Before loading the package, set *Java* parameters (we set the maximum size of memory allocated by *Java* at 2 GB) and then load the library.

```
1 R> options(java.parameters = "-Xmx2g")  
2 R> library("rmcfs")
```

When this is done, then the entire *R* environment is ready to run the example below.

4.1 The artificial data

To review and understand the MCFS-ID algorithm we suggest to create and use an extraordinarily simple example data set. It consists of objects from 3 classes, A , B and C , that contain 40, 20 and 10 objects, respectively (70 objects altogether). For each object, we create 6 binary features ($A1$, $A2$, $B1$, $B2$, $C1$ and $C2$) that are 'ideally' or 'almost ideally' correlated with *class* feature. If an object's '*class*' equals ' A ', then its features $A1$ and $A2$ are set to class value ' A '; otherwise $A1 = A2 = 0$. If an object's '*class*' is ' B ' or ' C ', we proceed analogously, but we introduce some random corruption to 2 observations from class ' B ' and to 4 observations from class ' C ': in the former case, for each of the two observations and both attributes $B1/B2$, we randomly replace their value ' B ' by '0' and in the latter case, again for each of the four observations and both attributes $C1/C2$, we randomly replace their value ' C ' by '0'. The data also contains additional 500 random numerical features with uniformly $[0,1]$ distributed values. Thus we end up with 6 nominal important features (3 pairs with different levels of importance for classification) and 500 randomly distributed.

```

1 R> set.seed(0)
2 R> class <- c(rep("A",40), rep("B",20), rep("C",10))
3 R> A <- B <- C <- rep("0", length(class))
4 R> A[class == "A"] <- "A"
5 R> B[class == "B"] <- "B"
6 R> C[class == "C"] <- "C"
7 R> rnd <- runif(length(class))
8 R> B[class == "B"][rnd[class == "B"] <= (sort(rnd[class == "B"])
9   ) [2])] <- "0"
9 R> C[class == "C"][rnd[class == "C"] <= (sort(rnd[class == "C"])
10  ) [4])] <- "0"
10 R> d <- data.frame(matrix(runif(500 * length(class)), ncol =
11   500))
11 R> d <- cbind(d, data.frame(A1 = A, A2 = A, B1 = B, B2 = B,
12 + C1 = C, C2 = C, class))

```

Once input data is created, we can review, e.g., the last few columns and rows of it:

```

1 R> d[50:70, 497:507]
2
3           X497      X498      X499      X500 A1 A2 B1 B2 C1 C2
4 50 0.20466809 0.56092946 0.95223478 0.70965238  0  0  B  B  0  0
      B

```

5	51	0.11404115	0.29816202	0.04477183	0.03345120	0	0	B	B	0	0
		B									
6	52	0.02938875	0.21069288	0.20364813	0.94183044	0	0	B	B	0	0
		B									
7	53	0.49005094	0.12659981	0.22824847	0.81885672	0	0	B	B	0	0
		B									
8	54	0.12656223	0.71003559	0.12826985	0.57745653	0	0	B	B	0	0
		B									
9	55	0.84182080	0.02133940	0.98461836	0.18093354	0	0	B	B	0	0
		B									
10	56	0.24221154	0.32765974	0.36310019	0.78545327	0	0	0	0	0	0
		B									
11	57	0.49214327	0.12442556	0.61966539	0.58656922	0	0	B	B	0	0
		B									
12	58	0.74438287	0.81923355	0.87711323	0.45680401	0	0	B	B	0	0
		B									
13	59	0.06014164	0.95796935	0.81141427	0.65610882	0	0	B	B	0	0
		B									
14	60	0.34402144	0.46230511	0.98609784	0.65171840	0	0	B	B	0	0
		B									
15	61	0.01665507	0.35895054	0.92959504	0.66303803	0	0	0	0	C	C
		C									
16	62	0.53326256	0.54672994	0.94477699	0.53642020	0	0	0	0	C	C
		C									
17	63	0.92763065	0.15789376	0.33911702	0.37232281	0	0	0	0	0	0
		C									
18	64	0.42980634	0.15969091	0.11927503	0.61437815	0	0	0	0	C	C
		C									
19	65	0.20276997	0.29115349	0.23746162	0.79282321	0	0	0	0	0	0
		C									
20	66	0.76193072	0.18045207	0.13270288	0.26172409	0	0	0	0	C	C
		C									
21	67	0.12037553	0.13130389	0.68910685	0.99061888	0	0	0	0	0	0
		C									
22	68	0.88857332	0.95222348	0.07483475	0.94459348	0	0	0	0	C	C
		C									
23	69	0.20723421	0.42698189	0.77679508	0.49228194	0	0	0	0	C	C
		C									
24	70	0.15157902	0.08401528	0.94564002	0.03452377	0	0	0	0	0	0
		C									

The number of objects in each class is equal to:

```

1 R> table(d$class)
2
3 A B C
4 40 20 10

```

Package *rmcfs* includes function ***artificial.data*** that creates above example dataset.

```
1 R> d <- artificial.data(rnd_features = 500, corruption = c  
  (0,2,4), seed = 0)
```

4.2 MCFS-ID on artificial data

Let us run *mcfs* on the created data. If the given CPU is HT (Hyper-Threading) quad core, we recommend setting up to 8 threads for processing because the efficiency gain of multithreaded processing is limited by the number of physical CPU cores. However, in the example below, we use 1 thread to get perfect reproducibility, since the implemented multithreaded mechanism is asynchronous. For a standard PC and 1 thread, it may take around 2 minutes, but we observe nearly linear dependence between number of threads running on physical cores and speed of calculation (see Section 4.3). Parameter ***cutoffPermutations*** is set to 5 for this simple example dataset, but for real data we recommend default 20.

```
1 R> result <- mcfs(class~., d, cutoffPermutations = 5, seed = 0,  
2 + threadsNumber = 1)  
3 R> result$exec_time  
4  
5 Time difference of 2.622107 mins
```

After successfully running MCFS-ID algorithm we can check convergence of the algorithm (see Figure 2). The distance function shows the difference between two consecutive rankings - zero means no changes between two rankings (see the left Y axis). Common part gives the fraction of features that overlap for two different rankings (see the right Y axis). Ranking stabilizes over a number of iterations: distance tends to zero and common part tends to 1. Beta1 shows the slope of the tangent of a smoothed distance function. If Beta1 tends to 0 (the right Y axis) then the distance is given by a flat line.

```
1 R> plot(result, type = "distances")
```

Now we can check the cutoff value for various methods and review the result for the default one.




Figure 2: Distance function and common part.

```
1 R> print(result$cutoff)
```

```
2
3      method      minRI  size      minID
4 1 criticalAngle 0.02667534   22        NA
5 2      kmeans 0.33329248    6        NA
6 3 permutations 0.07494005    6 5.029552
7 4         mean 0.03540977   11        NA
```

```
1 R> result$cutoff_value
```

```
2
3 [1] 6
```

By default, the final cutoff value is equal to the value obtained from the method based on permutations of the decision attribute. The mean obtained from all methods is in our example equal to 11. The mean cutoff value is larger than actual number of informative features since the critical angle method takes into consideration only the shape of the RI distribution. Input parameter *cutoffMethod* for the *mcfs* function determines the method which we want to use as the oracle. Since cutoff value is determined by *result\$cutoff_value*, the user may change it according to his/her will.




Figure 3: RIs and maximal RIs from permutation experiments for top 50 features.

This value is used by *plot* function.

We may plot RI values in decreasing order for the top, e.g., 50 features. See Figure 3. The line with red/grey dots gives RI values, the blue vertical barplot gives difference δ between consecutive RI values. Informative features are separated from non-informative ones by the cutoff value and are presented in the plot as red and grey dots, respectively. We can also view all maximal RIs obtained from all permutation experiments (parameter *plot_permutations* = *TRUE*) - Figure 3 shows these 5 maximal RIs as horizontal red lines. The distribution of the max RI values determines the cutoff value.

```
1 R> plot(result, type = "ri", size = 50, plot_permutations = TRUE)
```

Similarly, we can plot ID weights in decreasing order for the top, e.g., 50 ID-Graph edges.

```
1 R> plot(result, type = "id", size = 50)
```

Now, we can review labels and RIs of the top features. The resulting plot is presented in Figure 4. One can see that all six features are highly




Figure 4: Top features selected by MCFS-ID.

important and their RIs are much higher than those of other features. The set of informative features is flagged in red in the plot. Features *A1* and *A2* have substantially higher RI values because they 'ideally' separate the largest class and in most cases they appear in the root node. The second level of a tree may be determined by either of *B1*, *B2*, *C1* and *C2*, but in our case features *B1/B2* are less corrupted by "noise" and hence they are much more informative than features *C1/C2*. Notice that all pairs (*A1/A2*, *B1/B2*, *C1/C2*) have similar RIs and all the resulting importance levels are consistent with the corruption level introduced to our artificial data.

```
1 R> plot(result, type = "features", size = 10)
```

Finally, we can build and visualize ID-Graph. By default, we plot all the edges that connect 6 informative features as defined by ***result\$cutoff_value***.

```
1 R> gid <- build.idgraph(result)
2 R> plot(gid, label_dist = 1)
```

In the ID-Graphs, as seen in Figure 5, some additional information is conveyed with the help of suitable graphical means. The color intensity of a node is proportional to the corresponding feature's RI. The size of a node is




Figure 5: ID-Graph for artificial data.

proportional to the number of edges related to this node. The width and level of darkness of an edge is proportional to the ID weight of this edge. Since we would like to review only the strongest ID weights let us plot ID-Graph with only 12 top edges (***size_ID = 12***).

```

1 R> gid <- build.idgraph(result, size_ID = 12)
2 R> plot(gid, label_dist = 1)

```

In Figure 6, top 6 features along with top 12 ID weights are presented. Notice that the two most important features, *A1* and *A2*, point to all another ones, while *B1* and *B2* point to *C1* and *C2*. The directions of edges reproduce paths in decision trees. If features *A1*, *A2* determine the root nodes, then the leafs are *B1*, *B2*, *C1* or *C2*. If features *B1*, *B2* determine the root nodes then the leafs are *C1*, *C2*. Interestingly and correctly, the ID-Graph (in both Figures 5 and 6) does not show any connection between identical features (e.g., between *A1* and *A2*). Indeed, such features do not cooperate in distinguishing between classes, since they are the same!

For top features set, when the execution of MCFS-ID has been finished, the procedure runs 10 fold cross validation (cv) on 6 different classifiers (see Figure 7). Each cv is repeated 3 (defined by ***finalCVRepetitions*** parameter) times and the mean value of accuracy and weighted accuracy are gathered. Since the weighted accuracy is equal to the mean over all true positive rates (TPR), it is more meaningful for datasets with unbalanced classes. In




Figure 6: ID-Graph for artificial data, limited to top 6 features and top 12 ID weights.

our example, given the first two features from the ranking, $A1$ and $A2$, only 66% weighted accuracy can be achieved, since only class A can then be separated and the remaining objects can be labeled as B - we perfectly classify 2 out of 3 classes (and hence $wacc=66\%$). For 6 and more top features the accuracy depends on an algorithm and a given cv experiment. CV plot presents the result for *result\$cutoff_value* features (red label on X axis) and its multiple (0.25, 0.5, 0.75, 1, 1.25, 1.5, 2).

```
1 R> plot(result, type = "cv", measure = "wacc")
```

The RIPPER algorithm is a rule-based classifier and thus it provides the user with classification rules. Function *print()* provides in particular the rules and their classification ability.

```
1 R> print(result)
2
3 JRIP classification rules created on top 6 features:
4 JRIP rules:
5 =====
6
7 (A1 = 0) and (B1 = 0) => class=C (12.0/2.0)
8 (A1 = 0) => class=B (18.0/0.0)
9 => class=A (40.0/0.0)
10
11 Number of Rules : 3
```




Figure 7: Cross Validation results for top features. Selected cutoff value is flagged by red color on the X axis.

```

12 RIPPER CV Result (10 folds repeated 3 times)
13 Confusion Matrix
14   predicted
15   class B C A
16 B 54.0  6.0  0.0
17 C  2.0 26.0  2.0
18 A  0.0  0.0 120.0
19
20
21 Accuracy = 0.9523
22 WeightedAccuracy = 0.9222
23
24 True Positive Rate
25   A: 1.0
26   B: 0.9
27   C: 0.8666
28 False Positive Rate
29   A: 0.0222
30   B: 0.0133
31   C: 0.0333

```




Figure 8: Prediction result on validation set vs the number of top features used to train SVM classifier.

4.3 MCFS-ID on real data

In order to present practical usability of the algorithm we used Arcene dataset downloaded from UCI repository (<https://archive.ics.uci.edu/ml/datasets/Arcene>). Arcene's task is to distinguish cancer versus normal patterns from mass-spectrometric data. This is a two-class classification problem with 10000 continuous input variables and 100 objects used for training. This dataset was one of 5 datasets of the NIPS 2003 feature selection challenge; see [Guyon *et al.*(2004)]. After running MCFS-ID on its default parameters we obtained the ranking of features that we used for prediction on the validation set. Starting from 1 up to 200 top features we repeatedly trained SVM and applied it on 100 new objects (unseen during the feature selections and training the model) from the validation dataset. For over 100 top features classification accuracy is almost 99% and for over 175 top features it achieves 100% (see Figure 8; this almost or strictly perfect accuracy follows from the task's low complexity and small size of the validation set - unfortunately, a larger test set with class labels is unavailable).

It should be expected that some combination of top features from the ranking and features that reveal strongest two-way interdependencies (interactions) with those top ones may comprise a very good set of features to classify on. Note that such strong interdependencies, measured by the ID

weights, may occur between pairs of features each of which is of high relative importance and, perhaps equally or almost equally likely, between features only one of which has high value of RI. Indeed, a feature with low value of RI may prove a good helper of the other feature to increase predictive power of the latter through a directed interaction with the former. In order to confirm this claim we first selected 200 features with the highest RI and then built ID-Graph for 50 edges with the highest ID weights and such that each edge connects a feature from the aforementioned set of 200 features with any other feature (see Figure 9). Concatenation of features present in the ID Graph with top 10 features gave us 58 features with 9 of them present both in the top 10 and in the graph:

1	[1]	"V1184"	"V5473"	"V698"	"V8502"	"V8806"	"V7197"	"V4290"	"
		"V1476"	"V7899"						
2	[10]	"V9743"	"V7542"	"V4183"	"V9050"	"V3206"	"V6292"	"V4352"	"
		"V2448"	"V4557"						
3	[19]	"V893"	"V504"	"V4542"	"V2227"	"V2278"	"V7696"	"V4970"	"
		"V8623"	"V9082"						
4	[28]	"V130"	"V5112"	"V2080"	"V3170"	"V9213"	"V2057"	"V4895"	"
		"V7530"	"V7014"						
5	[37]	"V8976"	"V8904"	"V4322"	"V1046"	"V4513"	"V1322"	"V9104"	"
		"V5417"	"V9358"						
6	[46]	"V4738"	"V4301"	"V8132"	"V9319"	"V787"	"V9477"	"V3956"	"
		"V3014"	"V9884"						
7	[55]	"V8506"	"V7319"	"V2247"	"V7748"				

Thus obtained set of features was used to train SVM model and determine its classification quality on the validation set. For the 58 features we obtained 98% accuracy.

To verify convergence of our MCFS heuristic procedure we have run the algorithm three times on the same training data - each time with a different seed. All three runs were made using the same default parameters as before. Having obtained all three rankings, we compared the three sets of top 100 features. As many as 92 features were present in all three rankings, no matter what was the starting seed. Below we present top twenty features from the first run and their positions in the rankings for the remaining two runs.

	attribute	position_1	position_2	position_3	RI_1	RI_2	RI_3
1	2						
2	1	V1184	1	1	1	0.1627335	
		0.1667576	0.1608779				
3	2	V8502	2	2	2	0.1393340	
		0.1379264	0.1378373				
4	3	V5473	3	3	3	0.1260721	
		0.1306319	0.1259599				

17	16	V7696	16	15	19	0.0744866
		0.0752617	0.0705253			
18	17	V436	17	13	13	0.0733441
		0.0780559	0.0812494			
19	18	V6986	18	24	27	0.0726183
		0.0676880	0.0655208			
20	19	V3339	19	19	18	0.0718252
		0.0691084	0.0720879			
21	20	V9082	20	28	24	0.0715442
		0.0659364	0.0686405			




Figure 10: Execution time vs number of threads (Arcene data).

To analyze scalability of our current implementation we run *mcfs* procedure for a multiple number of threads for three different Intel CPUs machines. The result presented in Figure 10 shows an initially nearly linear relation between time and threads number: $time(th) = time(th = 1)/th$, where th is the number of threads, for 14 cores CPU. Experiments show that the gain in speed highly depends on the architecture of the given CPU. Since Hyper-Threading (HT) technology is based on utilizing the same physical core by two threads simultaneously, it can only give a boost to the speed when the CPU is not well utilized by a single thread. For a multicore system (2x14), we noticed that there is no gain in speed from 16 threads on. We profiled the

application to find possible concurrency between threads and locks for critical section execution but it has proved not to be the case. There are many other issues that can affect *Java* multithreading program such as: memory bandwidth, L3 cache size limitation, or frequent Garbage Collector execution (see [Qian *et al.*(2015)]). However multithreading allows to run calculation 7 times faster by using 8 threads vs 1 on 14 cores CPU (49 sec vs 340 sec respectively).

5 Summary

In this paper, we have presented *rmcfs* package that can be successfully used for feature selection and interdependencies discovery. Here, we have discussed one extraordinarily simple usage example and only one example using a real-life dataset, but the MCFS-ID algorithm has been proved in the past studies as very useful technique to limit the number of features and select the informative ones in various real high-dimensional problems; see [Dramiński *et al.*(2008)], [Kierczak *et al.*(2009)], [Kierczak *et al.*(2010)], [Khaliq *et al.*(2015)], [Enroth *et al.*(2012)], [Dramiński *et al.*(2016)]. The ranking of features can be used to review features one by one (starting from the top) even in the case of BIG DATA. Using the MCFS, one can build rule sets on top features and review their meaning; see [Bornelov *et al.*(2014)]. The ID part presents interdependencies between features in a consistent and comprehensive way. It shows not simple correlations but nonlinear relations between features that may reveal causality in the data (to be inferred from or confirmed by suitable background knowledge); see [Dramiński *et al.*(2016)]. The novelty of the paper lies in presenting: (i) a concise and comprehensive description of an *R* package that automatically implements proper (default) settings of all crucial MCFS-ID parameters to obtain reliable results for data of any given size; (ii) ways of determining the cut-off value between informative and non-informative features; (iii) scalability analysis of current implementation of the *rmfcs*; (iv) an illustration of using discovered interdependencies (interactions) to advantage in classification tasks.

Acknowledgements

We thank our close collaborators, Jan Komorowski, Michal J. Dabrowski, Klev Diamanti, Marcin Kierczak and Marcin Kruczyk, who have built on the MCFS-ID, most notably by providing a host of new insights and results within the area of bioinformatics. Our thanks also to Julian Zubek who wrote

some pieces of the *R* code. Last but not least, we thank the anonymous reviewers for most valuable comments and insightful suggestions.

References

- [Alizadeh *et al.*(2000)] Alizadeh AA, Eisen MB, Davis RE, Ma C, Lossos IS, Rosenwald A, Boldrick JC, Sabet H, Tran T, Yu X, Powell JJ, Yang L, Marti GE, Moore T, Hudson JJ, Lu L, Lewis DB, Tibshirani R, Sherlock G, Chan WC, Greiner TC, Weisenburger DD, Armitage JO, Warnke R, Levy R, Wilson W, Grever MR, Byrd JC, Botstein D, Brown PO, Staudt LM (2000). “Distinct Types of Diffuse Large B-cell Lymphoma Identified by Gene Expression Profiling.” *Nature*, **403**(6769), 503–511.
- [Archer and Kimes(2008)] Archer KJ, Kimes RV (2008). “Empirical Characterization of Random Forest Variable Importance Measures.” *Comput. Stat. Data Anal.*, **52**(4), 2249–2260. ISSN 0167-9473. doi:10.1016/j.csda.2007.08.015. <http://dx.doi.org/10.1016/j.csda.2007.08.015>.
- [Bornelöv and Komorowski(2016)] Bornelöv S, Komorowski J (2016). “Selection of Significant Features Using Monte Carlo Feature Selection.” *Challenges in Computational Statistics and Data Mining*, **605**, 25–38. ISSN 1860-949X. doi:10.1007/978-3-319-18781-5_2.
- [Bornelov *et al.*(2014)] Bornelov S, Marillet S, Komorowski J (2014). “Ciruviz: A Web-based Tool for Rule Networks and Interaction Detection Using Rule-based Classifiers.” *BMC Bioinformatics*, **15**(1), 139+. ISSN 1471-2105. doi:10.1186/1471-2105-15-139. <http://dx.doi.org/10.1186/1471-2105-15-139>.
- [Breiman and Cutler(2008)] Breiman L, Cutler A (2008). *Random Forests - Classification/Clustering Manual*. http://www.math.usu.edu/~adele/forests/cc_home.htm.
- [Chrysostomou *et al.*(2008)] Chrysostomou K, Chen SY, Liu X (2008). “Combining Multiple Classifiers for Wrapper Feature Selection.” *International Journal of Data Mining, Modelling and Management*, **1**(1), 91–102.
- [Díaz-Uriarte and De Andres(2006)] Díaz-Uriarte R, De Andres SA (2006). “Gene Selection and Classification of Microarray Data Using Random Forest.” *BMC Bioinformatics*, **7**(1), 1. doi:10.1186/1471-2105-7-3.

- [Dramiński *et al.*(2016)] Dramiński M, Dąbrowski MJ, Diamanti K, Koronacki J, Komorowski J (2016). “Discovering Networks of Interdependent Features in High-Dimensional Problems.” *Big Data Analysis: New Algorithms for a New Society*, pp. 285–304. ISSN 2197-6503. doi:10.1007/978-3-319-26989-4_12.
- [Dramiński *et al.*(2010)] Dramiński M, Kierczak M, Koronacki J, Komorowski J (2010). “Monte Carlo Feature Selection and Interdependency Discovery in Supervised Classification.” *Advances in Machine Learning*, **2**, 371–385. ISSN 1860-949X. doi:10.1007/978-3-642-05179-1_17.
- [Dramiński *et al.*(2005)] Dramiński M, Koronacki J, Komorowski J (2005). “A study on Monte Carlo Gene Screening.” *Intelligent Information Processing and Web Mining*, **31**, 349–356. doi:10.1007/3-540-32392-9_36.
- [Dramiński *et al.*(2008)] Dramiński M, Rada-Iglesias A, Enroth S, Wadelius C, Koronacki J, Komorowski HJ (2008). “Monte Carlo Feature Selection for Supervised Classification.” *Bioinformatics*, **24**(1), 110–117. <http://dblp.uni-trier.de/db/journals/bioinformatics/bioinformatics24.html#DraminskiREWKK08>.
- [Dudoit and Fridlyand(2003)] Dudoit S, Fridlyand J (2003). “Classification in Microarray Experiments.” *Statistical analysis of gene expression microarray data*, **1**, 93–158.
- [Enroth *et al.*(2012)] Enroth S, Bornelöv S, Wadelius C, Komorowski J (2012). “Combinations of Histone Modifications Mark Exon Inclusion Levels.” *PloS one*, **7**(1), e29911. doi:10.1371/journal.pone.0029911.
- [Guyon *et al.*(2004)] Guyon I, Gunn SR, Ben-Hur A, Dror G (2004). “Result Analysis of the NIPS 2003 Feature Selection Challenge.” **4**, 545–552.
- [Gyenesei *et al.*(2007)] Gyenesei A, Wagner U, Barkow-Oesterreicher S, Stolte E, Schlapbach R (2007). “Mining Co-regulated Gene Profiles for the Detection of Functional Associations in Gene Expression data.” *Bioinformatics*, **23**(15), 1927–1935.
- [Hall *et al.*(2009)] Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten IH (2009). “The WEKA Data Mining Software: An Update.” *ACM SIGKDD explorations newsletter*, **11**(1), 10–18. <http://www.cs.waikato.ac.nz/ml/index.html>.

- [Hastie *et al.*(2001)] Hastie T, Tibshirani R, Botstein D, Brown P (2001). “Supervised Harvesting of Expression Trees.” *Genome biology*, **2**(1), 1–0003.
- [Ho(1998)] Ho TK (1998). “The Random Subspace Method for Constructing Decision Forests.” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, **20**(8), 832–844.
- [Huynh-Thu *et al.*(2012)] Huynh-Thu VA, Saeys Y, Wehenkel L, Geurts P (2012). “Statistical Interpretation of Machine Learning-based Feature Importance Scores for Biomarker Discovery.” *Bioinformatics*, **28**(13), 1766–1774.
- [Khaliq *et al.*(2015)] Khaliq Z, Leijon M, Belák S, Komorowski J (2015). “A Complete Map of Potential Pathogenicity Markers of Avian Influenza Virus Subtype H5 Predicted from 11 Expressed Proteins.” *BMC microbiology*, **15**(1), 1. doi:10.1186/s12866-015-0465-x.
- [Kierczak *et al.*(2010)] Kierczak M, Dramiński M, Koronacki J, Komorowski J (2010). “Computational Analysis of Molecular Interaction Networks Underlying Change of HIV-1 Resistance to Selected Reverse Transcriptase Inhibitors.” *Bioinformatics and biology insights*, **4**, 137–146.
- [Kierczak *et al.*(2009)] Kierczak M, Ginalski K, Dramiński M, Koronacki J, Rudnicki W, Komorowski J (2009). “A Rough Set-based Model of HIV-1 Reverse Transcriptase Resistome.” *Bioinformatics and biology insights*, **3**, 109–127.
- [Li *et al.*(2002)] Li Y, Campbell C, Tipping M (2002). “Bayesian Automatic Relevance Determination Algorithms for Classifying Gene Expression data.” *Bioinformatics*, **18**(10), 1332–1339.
- [Lu *et al.*(2007)] Lu C, Devos A, Suykens JA, Arús C, Huffel SV (2007). “Bagging Linear Sparse Bayesian Learning Models for Variable Selection in Cancer Diagnosis.” *Information Technology in Biomedicine, IEEE Transactions on*, **11**(3), 338–347.
- [Nicodemus *et al.*(2010)] Nicodemus KK, Malley JD, Strobl C, Ziegler A (2010). “The Behaviour of Random Forest Permutation-based Variable Importance Measures under Predictor Correlation.” *BMC Bioinformatics*, **11**(1), 1.

- [Paul and Dupont(2015)] Paul J, Dupont P (2015). “Inferring Statistically Significant Features from Random Forests.” *Neurocomputing*, **150**, 471–480.
- [Qian *et al.*(2015)] Qian J, Li D, Srisa-an W, Jiang H, Seth S (2015). “Factors affecting scalability of multithreaded Java applications on manycore systems.” pp. 167–168.
- [Saeys *et al.*(2007)] Saeys Y, Inza I, Larrañaga P (2007). “A Review of Feature Selection Techniques in Bioinformatics.” *Bioinformatics*, **23**(19), 2507–2517.
- [Smyth *et al.*(2003)] Smyth GK, Yang YH, Speed T (2003). “Statistical Issues in cDNA Microarray Data Analysis.” *Functional Genomics*, **224**, 111–136. ISSN 1064-3745. doi:10.1385/1-59259-364-X:111.
- [Strobl *et al.*(2008)] Strobl C, Boulesteix AL, Kneib T, Augustin T, Zeileis A (2008). “Conditional Variable Importance for Random Forests.” *BMC Bioinformatics*, **9**(1), 1. doi:10.1186/1471-2105-9-307.
- [Strobl *et al.*(2007)] Strobl C, Boulesteix AL, Zeileis A, Hothorn T (2007). “Bias in Random Forest Variable Importance Measures: Illustrations, Sources, and a Solution.” *BMC Bioinformatics*, **8**(1), 1. doi:10.1186/1471-2105-8-25.
- [Tibshirani *et al.*(2002)] Tibshirani R, Hastie T, Narasimhan B, Chu G (2002). “Diagnosis of Multiple Cancer Types by Nearest Shrunken Centroids of Gene Expressions.” *Proceedings of the National Academy of Sciences*, **99**(10), 6567–6572.
- [Tibshirani *et al.*(2003)] Tibshirani R, Hastie T, Narasimhan B, Chu G (2003). “Class Prediction by Nearest Shrunken Centroids, with Applications to DNA Microarrays.” *Statistical Science*, **18**(1), 104–117.
- [Wright *et al.*(2016)] Wright MN, Ziegler A, König IR (2016). “Do little interactions get lost in dark random forests?” *BMC bioinformatics*, **17**(1), 145.
- [Ziegler and König(2014)] Ziegler A, König IR (2014). “Mining data with random forests: current options for real-world applications.” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, **4**(1), 55–63.