

Sequential Feature Classification in the Context of Redundancies

Lukas Pfannschmidt*, Barbara Hammer†

Machine learning group

Bielefeld University

Bielefeld, Germany

**lpfannschmidt@techfak.uni-bielefeld.de*, †*bhammer@techfak.uni-bielefeld.de*

Abstract—The problem of all-relevant feature selection is concerned with finding a relevant feature set with preserved redundancies. There exist several approximations to solve this problem but only one could give distinction between strong and weak relevance. This approach was limited to the case of linear problems. In this work we present a new solution for this distinction in the non-linear case through the use of random forest models and statistical methods.

Index Terms—feature selection, interpretability, redundancies, strong and weak relevance

1. Background

To interpret the behaviour of machine learning models it is important to find the original input features which correspond to the output. Often the majority of input features is irrelevant, and we seek a feature subset which only contains relevant features in relation to an unknown underlying process. Furthermore, a feature subset without redundancies can be beneficial by increasing robustness. Such a subset of features with minimal redundant information is called the minimal feature subset and the majority of existing feature selection methods belong to the minimal feature selection problem.

In many cases, the function or relation between input features is unknown. This is often encountered in the life sciences where machine learning is used on data sets without identifying all variables beforehand. Here, the main goal of machine learning shifts from making perfect predictions to identifying the contributing elements. This is in part contrary to the minimal feature selection problem which removes redundant features in the process. These features could lead to additional insights in the biological mechanisms.

The *all-relevant feature selection problem* (ARFS) is the task of finding a relevant feature subset including redundant features. To find the all-relevant feature set, denoted as \mathcal{A} , is often much harder than finding the minimal relevant set \mathcal{M} . To identify \mathcal{M} many approaches exist, such as the Lasso [Tib96] and ElasticNet [ZH05], which exploit regularization to achieve sparseness in the model’s use of features and thus

produce a compressed subset. In the presence of redundancies these subset and model solutions only represent one of many possible solutions and are not unique. To find \mathcal{A} one would have to find all possible models, which is infeasible in practise for the exact solution.

Consequently, several approximations exist and were compared in [DSS]. One of the best performing approaches is the Boruta wrapper method [KR10]. It utilizes a random forest model together with random *shadow* features, which are used as comparison variables, to filter out all irrelevant features while keeping redundant features. Another new approach not compared in [DSS] utilized feature relevance bounds to get an approximation of all possible feature contributions and then utilized random contrast features for approximation of a relevance threshold [Göp+18]. Not only were these bounds able to produce an approximation of \mathcal{A} but could also classify for each relevant feature if other redundant features existed (weak relevance) or if they were unique in their information content (strongly relevant). This made it possible to get insights into possible relations between weakly relevant features [Pfa+19]. The feature relevance bounds could be phrased in terms of efficient convex linear problems and were released as a standalone analysis tool named FRI [PJ20]. While the linear problems are efficient the approach does not scale to bigger data sets and can not be adapted to non-linear problems without losing convexity of the optimization.

In this work we present a method which extends the Boruta method with a feature classification step which produces the same discrimination between strong and weak relevance as seen in [Pfa+19]. We therefore combine the advantages of Boruta and FRI while achieving more efficient run times. Relevance decisions can be made based on the accuracy of the model when excluding a single feature in question. In section 2.1 we describe our approach of efficiently decomposing the overall feature set into subsets with different characteristics to produce the distinction between strong and weak relevant features without testing all available features. We improve existing methods to find set \mathcal{M} by using statistical based thresholds and introduce a robust score testing scheme in section 2.2 to improve feature classification. In section 3 we analyse the general feature selection performance against other established approaches in context of redundant features.

This contribution has been made possible through the Bielefeld Young Researchers’ Fund and the BMBF (grant number 01S18041A)

2. Methods

2.1. Redundant Feature Selection

Let X be data set $X := \{x_i \in \mathbb{R}^d; i = 1 \dots n\}$ with n samples and with $\mathcal{G} := \{j \in \mathbb{Z}; j = 1 \dots d\}$ as the set of all features such that cardinality $|\mathcal{G}| = d$. Target variable $y \in \mathbb{R}^n$ is distributed according to some unknown function dependent on X such that $g(X) = y$. Without loss of generality we limit y to be continuous and g to be a regression function. Let $\bar{f} := \hat{g}$ be the optimal estimator of g with minimal estimation loss. Consider the estimation loss as

$$\mathcal{L}(X, f) := \sum_{\forall i} |f(x_i) - y_i|$$

where $x \in X$ which denotes the deviation from g if no random noise is involved.

We are interested in the input features from X which are relevant to f as defined by Nilsson, Bjorkegren, and Tegner [NBT]. One can define the relevance of a feature j for a single model such as \bar{f} or for all models in a given set e.g. $f \in \mathcal{F}$ as the set of all possible models. In the context of redundancies we are interested in all functions with similar \mathcal{L} and possibly different composition of input features:

$$\bar{\mathcal{F}} := \{f \in \mathcal{F} \mid \mathcal{L}(X, f) \approx \mathcal{L}(X, \bar{f})\}$$

Before we observed the set of data X with all features in \mathcal{G} , i.e. $X = X_{\mathcal{G}}$, now we also consider the data set with specific subsets of features, e.g. the dataset with feature j removed is denoted as $X_{\mathcal{G} \setminus j}$

Kohavi and John [KJ97] introduced several classes of feature relevance and the resulting subsets which the general feature set is composed of. These are the set of strongly relevant features (\mathcal{S}), the set of weakly relevant features (\mathcal{W}) and irrelevant features (\mathcal{I}). The union of \mathcal{S} and \mathcal{W} is the set of all-relevant features \mathcal{A} :

$$\mathcal{A} := \mathcal{S} \cup \mathcal{W}.$$

The goal of all relevant feature selection is finding all features belonging to \mathcal{A} , and not to identify the detailed composition of \mathcal{S} and \mathcal{W} . Most existing methods only produce information about membership of \mathcal{A} , except the feature relevance bounds method [Göp+18], which yields both \mathcal{S} and \mathcal{W} in the linear case.

One existing all-relevant selection method is called Boruta [KR10]. It builds on the information metrics acquired by observing the single trees of a random forest model. Because the scores are not consistent with their real significance [Rud+06], Boruta employs an extended information system. Extended in the way, that additional to normal features Boruta adds shadow features. Shadow features are randomly shuffled clones of existing features to remove correlation with the target. Through the addition of those, one can estimate a contrast distribution of features, which are by design irrelevant. It then tests the real features against the shadow features iteratively, increasing the significance threshold, until all features are tested conclusively. Boruta's

comparison with a null distribution inspired one aspect of our proposed method, which we come back to later. More importantly, by comparing with a null distribution multiple times and introducing stochastic noise by repeatedly running a random forest model, Boruta can identify the \mathcal{A} set with high precision.

Having found \mathcal{A} , one can trivially classify all features in \mathcal{G} which are not in \mathcal{A} , as irrelevant:

$$\mathcal{I} := \mathcal{G} \setminus \mathcal{A}$$

To decompose \mathcal{A} into \mathcal{S} and \mathcal{W} we have to identify membership with at least one of them. We can check for strong relevance by repeatedly fitting models and checking the behaviour of the loss function similar to the approach in [NBT]. A feature j is strongly relevant if

$$\min_f \mathcal{L}(X_{\mathcal{G} \setminus j}, f) > \min_h \mathcal{L}(X, h).$$

This comparison would have to be performed for all $j \in \mathcal{G}$. While $|\mathcal{A}| \ll |\mathcal{G}|$ in most cases, i.e. the feature space is often sparse and most of the features can be ignored, we still can improve on this naive approach.

Earlier, we considered the specific problem of all-relevant feature selection. In general when feature selection is performed, one considers the *minimal-optimal* feature set (\mathcal{M}). The problem of finding \mathcal{M} can be stated as finding the minimal non-redundant feature subset with the maximum amount of information for some function class. This type of feature selection is easier than the ARFS [NBT] and many efficient methods exist such as greedy recursive feature elimination (RFE) [GMS17].

In the following we are using a new efficient threshold based approach which decides based on the models' importance values which features are important. Importance values are internal parameters which correspond to input features similar to the weights of linear models and schemes such as Lasso. Because we utilize Borute, we consider the importance values of a random forest model. In section 2.3 we describe in detail how we find \mathcal{M} with a sparse parameterization and statistical method to overcome inconsistent importance values. For now, we consider \mathcal{M} as given by this efficient approach, but alternatives such as RFE could also be applied.

By definition, $\mathcal{M} \subset \mathcal{A}$ and $\mathcal{S} \subset \mathcal{M}$. All features in the minimal optimal set \mathcal{M} are also included in \mathcal{A} and furthermore all strongly relevant features are included in \mathcal{M} . Instead of iterating and testing all features in \mathcal{A} we only have to consider the features in \mathcal{M} . In most cases when redundant relevant features are preset, i.e. $|\mathcal{M}| \ll |\mathcal{A}|$, it is much more efficient to only check features in the subset \mathcal{M} .

Therefore, it is sufficient to identify \mathcal{S} in \mathcal{M} through comparing the loss after the elimination of each feature and identify \mathcal{W} through

$$\mathcal{W} := \mathcal{A} \setminus \mathcal{S}. \quad (1)$$

Having said that, comparison of the loss is not straightforward as it requires a robust threshold to test against.

2.2. Robust Loss Comparison

In practise, we cannot perform a simple comparison between a reduced (set with feature j removed) and a normal feature set. Due to the stochastic nature of random forests, even fitting the same data set without feature removal can lead to variable models and thus to differing average losses. To counteract this, we estimate a distribution of \mathcal{L} . The distribution should represent a range of possible values which we regard as insignificant changes, similar to how our model would change, if an irrelevant feature would be removed.

Remember that we fit a RF model on a reduced dataset where one feature is eliminated and then observe the loss of the resulting model. The distribution should emulate this given setting and as such we would have to eliminate one feature in our sampling procedure as well. However, we cannot remove any feature from the initial feature set or we could possibly remove another relevant feature by chance. To emulate the changes in model size we therefore permute a randomly chosen feature and add it to the dataset. Permutation of feature j is denoted as $p_j(X)$ and an extended dataset with such a permuted feature as

$$X^* := \{X \cup p_j(X) \mid j \sim U(d)\}.$$

$p_j(X)$ has no dependence on the target variable and represents an irrelevant feature.

We then define the random population

$$\hat{\pi}(\bar{f}, X) := \{\mathcal{L}(X^*, \bar{f})\}^\alpha$$

with parameter $\alpha \in \mathbb{Z}$ as the number of samples used.

We then define an interval of plausible values based on a t-distribution

$$\Pi(\cdot, \alpha) := \overline{\hat{\pi}(\cdot)}_\alpha \pm T_{\alpha-1}(p) \cdot \sigma(\hat{\pi}(\cdot)) \sqrt{1 + (1/\alpha)}.$$

Here $\overline{\hat{\pi}}_\alpha$ denotes the sample mean and $\sigma(\cdot)$ the standard deviation, and T represents Student's t-distribution with $\alpha - 1$ degrees of freedom. The size of Π depends on parameter $p \ll 1$ and the number of samples. In our experiments $\alpha \geq 50$ already yields robust thresholds for common feature set sizes. The interval only has to be computed once and is valid for all feature comparisons as it represents the distribution of irrelevant features and not a feature specific one.

Feature j is strongly relevant if

$$\mathcal{L}(X_{G \setminus j}, \bar{f}) < \Pi(\bar{f}, X)$$

which means we only have to check the lower bound of the prediction interval.

Using this procedure and checking all features in \mathcal{M} leads to the complete set of strongly relevant features \mathcal{S} and therefore also \mathcal{W} as seen in Equation 1.

2.3. Feature Importance Relevance Determination

As mentioned earlier in section 2.1 we do not use existing minimal feature selection methods to decide which

Algorithm 1: Estimating loss and irrelevant importance distribution

```

Data:  $X, y$ 
Input: Model,  $N_{\text{Samples}}$ 
 $\hat{\pi} \leftarrow \emptyset$  (empty set)
 $\hat{\gamma} \leftarrow \emptyset$ 
 $s \leftarrow 0$ 
while  $s < N_{\text{Samples}}$ ;  $s++$  do
     $s \leftarrow \text{generatePermFeature}(X)$ 
     $X^* \leftarrow X \cup s$ 

    // Fit model
    ext-model  $\leftarrow$  Model( $Xs, y$ )

    // loss samples section 2.2
    loss  $\leftarrow$  loss(ext-model,  $X^*, y$ )
    // importance samples
    section 2.3
    irrel_imp  $\leftarrow$  importance(ext-model,  $s$ )

     $\hat{\pi} \leftarrow \hat{\pi} \cup \text{loss}$ 
     $\hat{\gamma} \leftarrow \hat{\gamma} \cup \text{irrel\_imp}$ 
end
LossBounds  $\leftarrow$  t-statistic( $\hat{\pi}$ )
IrrelBounds  $\leftarrow$  t-statistic( $\hat{\gamma}$ )
Output: LossBounds, IrrelBounds

```

features are part of \mathcal{M} . Instead, we propose to use feature importance values from the Random forest model to efficiently decide which features are relevant. In contrast to Boruta, we are interested in a sparse solution of the feature set without redundant features, which represents \mathcal{M} .

Deep learning models with many hidden layers can be opaque in their attribution of the input features in relation to the output layer. In the case of random forest there exist several measures of feature importance. They commonly express a feature's importance by averaging an information measure over all splits in the decision forest, which were part of the ensemble. Examples are the average information gain of the objective function or the number of correct classifications with and without the feature. In the following let's consider the information gain measure as $\text{imp}(X, \bar{f}) \in \mathbb{R}^d$. $\text{imp}(X, \bar{f})_j$ is then the improvement of the splitting criterion averaged over all trees and all splits where feature j is used as the split feature.

If an importance measure would correlate to the relevance of the input feature, we could use these measures in deciding which features are relevant.

For example, we expect correlated or identical features to exhibit the same feature importance but in practice some features are implicitly preferred by small differences in information or pure stochastic reasons. We demonstrate this in Figure 2 (b) where some features have much bigger importance where others only have small or zero importance. This is an example of correlation bias [TL11] which is a common problem for many importance measures in general.

Algorithm 2: Iterative Decomposition Algorithm

```
Data:  $X, y$ 
Input: RF
 $S \leftarrow \emptyset$ 
 $\mathcal{W} \leftarrow \emptyset$ 
 $\mathcal{A} \leftarrow \text{Boruta}(\text{RF}, X, y)$ 
 $\mathcal{M} \leftarrow \text{ImpSelection}(\text{RF}, X, y)$  (Equation 2)

// Reduce dataset  $X$  to relevant
// features
 $\mathcal{V} \leftarrow \text{select}(X, \mathcal{A})$ 

// loss bounds using algorithm 1
 $\Pi \leftarrow \text{lossBounds}(\text{RF}, \mathcal{V}, y)$ 

// iterate over subset  $\mathcal{M}$ 
for feature  $j$  in  $\mathcal{M}$  do
  // Remove current feature
   $\mathcal{I} \leftarrow \mathcal{V}$  without feature  $j$ 
  // find best model without  $j$ 
   $\text{reduced\_model} \leftarrow \text{fit}(\text{RF}, \mathcal{I}, y)$ 
  // compute score
   $\text{loss}_j \leftarrow \text{loss}(\text{reduced\_model}, \mathcal{I}, y)$ 
  // add current feature to  $S$  if
  // significantly worse
  if  $\text{loss}_j$  not in  $\Pi$  then
    |  $S \leftarrow S \cup \{j\}$ 
  end
end

// decomposition (Equation 1)
 $\mathcal{W} \leftarrow \mathcal{A} \setminus S$ 
Output:  $S, \mathcal{W}$ 
```

One important parameter in fitting a random forest is the *feature fraction* which denotes how many features are included for each bootstrap tree. When this fraction is high ($\approx 100\%$) most features are included. To circumvent correlation bias, we use a feature fraction of only 10% in our experiments which yields a more even distribution of importances and is close to the recommended optimum of \sqrt{d} [DA06] as is demonstrated in Figure 2 (a).

Similar to linear models, using a sparse regularization we could force redundant features to exhibit low importance in the model. Lasso uses L1 regularization which leads to many zero entries in the model weights and features with non-zero weights are considered part of the feature set. By parameterizing the RF to have a high *feature fraction* we force a similar sparsity as can be seen in Figure 2. But apparent in the figure is that even irrelevant features do not have an importance value equal to zero which makes a simple threshold at zero is impossible. Thus, important for this approach is a well-defined threshold to decide which value is considered irrelevant. A simple measure like the mean of all importance values (blue horizontal line) does not work in general. We therefore propose to replicate the statistical comparison with a null distribution from section 2.2 in the

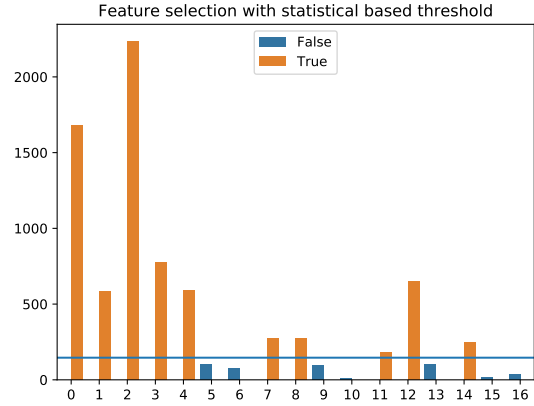


Figure 1: Importance values per feature (bars) of random forest model and upper feature selection threshold of Γ_s . The lower threshold is < 0 and excluded in this figure. Colors denote the membership of features to \mathcal{M} .

context of importance values.

2.3.1. Minimal Feature Set. In [Pfa+19] the authors proposed a statistics based approach to estimate a dynamic threshold for linear models. They use randomly permuted real features to simulate irrelevant variables and observe their weights in the model over many samples. We replicate this and extend the statistics from section 2.2 to include the importance values of randomly permuted features.

Because we create samples X^* identical to the samples in $\hat{\pi}()$ for the loss distribution already, we can efficiently combine both samples and model fitting at the same time. For the distribution in section 2.2 we focus on the score, whereby here we focus on the feature importance values, i.e. we use the same models to create the distributions.

The samples are then defined as

$$\hat{\gamma}_s(\bar{f}, X) := \{\text{imp}(p_j(X), \bar{f}) \mid j \sim U(d)\}^\alpha$$

where p is the permutation function and $p_j(X)$ is the shadow feature. We therefore sample the importance of the shadow feature in the model \bar{f} . The bounds are then defined as

$$\Gamma_s(\cdot, \alpha) := \overline{\hat{\gamma}_s(\cdot)}_\alpha \pm T_{\alpha-1}(p) \cdot \sigma(\hat{\gamma}_s(\cdot)) \sqrt{1 + (1/\alpha)}.$$

To produce the minimal-optimal feature set \mathcal{M} we fit a simple random forest with a high allowed feature fraction. This leads to the behaviour seen in Figure 2 (b) where only a subset of correlated features shows significant importance. We then compare each feature's importance value with Γ_s which represents the distribution of importance values we consider as irrelevant. The minimal-optimal set is then given as

$$\mathcal{M} := \{\text{imp}(X, \bar{f})_j > \Gamma_s \mid \forall j \in \mathcal{G}\}. \quad (2)$$

Figure 1 shows the upper bound of the interval in use with random forest model on toy data.

3. Results

3.1. Implementation

For the experiments we implemented the algorithms in section 2 in Python utilizing several existing libraries. The methods such as Boruta or our importance selection method are wrappers which require the definition of an inner model. Because we are fitting this model many times, we opted for a very efficient random forest implementation, which is provided by the *LightGBM* library [Ke+17].

The Boruta [KR10] method is implemented in the *boruta_py* library¹. Other utility functions are used from *scikit-learn* [Ped+11]. The complete implementation (nick-named "*Squamish*") and source code is publicly available².

3.2. Evaluation

In this section we try to characterize the behaviour of our proposed method. All scripts and data generation methods are publicly available and results can be reproduced³.

3.2.1. Selection Methods. To show the characteristics of our methods we ran several benchmarks against established feature selection methods. Here we specifically focus on the performance of the all-relevant feature selection in a setting where redundant features are present.

We compare against the following approaches

- *ElasticNet*: A linear model using the *ElasticNet* method which combines both L_2 and L_1 regularization. The combination can be weighted linearly and allows for more sensitivity of redundant features. In our experiments we fit this parameter through grid search in combination with cross validation. Features are selected according to recursive feature elimination which is guided by a cross-validated model performance [ZH05].
- *RF*: A random forest model (*LightGBM*) with recursive feature elimination as the selection method where number of features is also decided by cross-validation.
- *FRI*: The feature relevance interval method [PJ20; Pfa+19] as the only representative with distinction between strong and weak relevance.
- *SQ*: The sequential feature class decomposition method presented in algorithm 2 and implemented as described in section 3.1. Parameters for the statistical test where $\alpha = 50$ and $p = 10^{-6}$

Hyper-parameters for all methods were decided by cross-validation. The tree models used the default parameters from *LightGBM* except:

```
num_leaves: 32
```

1. https://github.com/scikit-learn-contrib/boruta_py

2. <https://github.com/lpfann/squamish>

3. instructions available at https://github.com/lpfann/squamish_experiments

```
max_depth: 5
boosting_type: rf
bagging_fraction: 0.632
bagging_freq: 1
```

The `feature_fraction` parameter was left to default (1) for the *RF* method. In *SQ* for the loss comparison it was set to 0.8. *SQ* set Boruta's tree model to `feature_fraction=0.1`.

3.2.2. Stability of Feature Importance Values. The proposed method is utilizing random forest importance values to decide about the relevance of features. In the following we perform a short analysis of the variance of these importance values over multiple model fits on the same dataset. We employ a random forest model as described in the section before with two different parameter choices. The first choice is a low `feature_fraction = 0.1` and the second the maximum `feature_fraction = 1` such that all features are allowed in each tree generation.

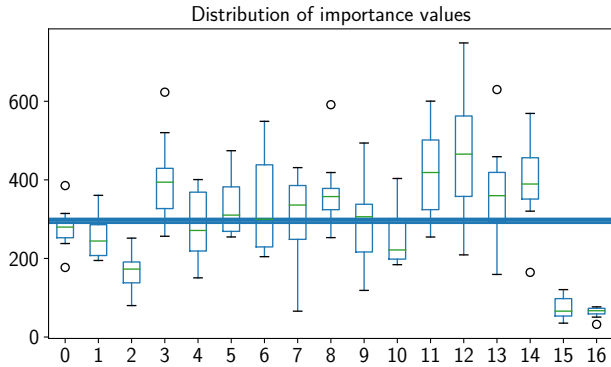
We generate a simple linear classification dataset with 17 features and 300 samples. Features 0-4 are considered strongly relevant, features 5-14 are weakly relevant with in-between correlations, and features 15-16 are irrelevant.

The test consisted of fitting the model on the dataset 10 times in a row and record the feature importance gain measure as defined in section 2.3. The resulting distributions are visualized in Figure 2 for each parameter choice. One can see, that even without any variation of the data, the values show high variance in both settings and do not correlate to the real mutual information with the target variable. Furthermore, the choice of a high feature fraction leads to some variables overshadowing the importance of others such as seen in feature 12 which contains the same information as all other weakly relevant in this case. A lower fraction leads to a more evenly distributed importance signature.

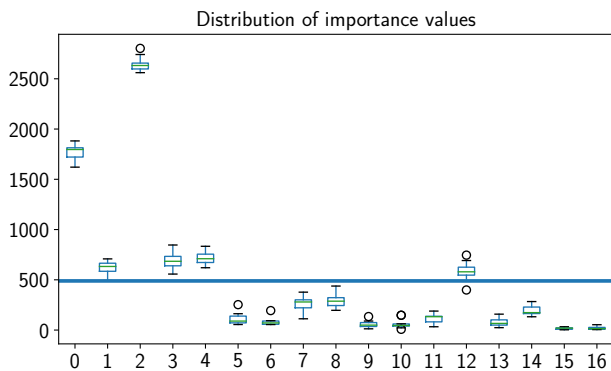
3.2.3. Parameterization for Feature Selection. Based on the evaluation in section 3.2.2 we also perform further analysis on the consequences of the parameter for feature selection. We extend the experiment with a feature selection step. Compared to recursive feature elimination guided by cross validation with Boruta.

We record the number of times each feature was selected in the feature set. This results in the frequency of selection or the probability that a feature is selected. The frequencies are given in Figure 3 for both models and both parameter settings for `feature_fraction`. We can see that a lower feature fraction is beneficial in the case of all-relevant feature selection where the Boruta model recognizes all relevant features 0-14 without selecting random features. The RFE procedure on the other hand suffers in this case and loses precision by selecting irrelevant features 50% of the time. It performs better with a high feature fraction parameter and selects strongly relevant features (0-4) consistently but shows higher variance in the case of weakly relevant features.

3.2.4. General Feature Selection Accuracy. The most relevant metric for feature selection methods is the accuracy



(a) 0.1 feature fraction



(b) 1.0 feature fraction

Figure 2: Distribution of feature gain importance values of random forest classifier over multiple bootstrap iterations on toy example where features 0-14 are correlated and 15-16 are irrelevant. Mean of all importance values is given as blue horizontal line. Subplots (a) and (b) represent different fraction of features allowed in tree construction.

of selected features. When the ground truth is known, we can explicitly evaluate the validity of the selected features. We focus on the all-relevant feature selection problem where we use the following measures to evaluate the match of the detected feature set and the known ground truth of all relevant features: precision and recall. Recall is defined by $TP / (TP+FN)$ with TP = number of true positives and FN = number of false negatives. It denotes how many of the relevant features were selected which is crucial when looking for the all relevant feature set. Precision is defined by $TP / (TP+FP)$ with FP = number of false positives and describes the frequency of false positives part of the feature set. One can use the F_1 measure as the harmonic mean of precision and recall:

$$F_1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

In this evaluation we compare the methods from section 3.2.1 by highlighting the F_1 measure.

TABLE 1: Parameters of synthetically generated datasets for a *linear separable* classification problem as described in section 3.2.4. Columns denote number of features with corresponding characteristics: n (number of samples), *strong* (number of strongly relevant features), *weak* (number of weakly relevant features), *irr* (number of irrelevant features).

Set	n	strong	weak	irr
Set 1	150	6	0	6
Set 2	150	0	6	6
Set 3	150	3	4	3
Set 4	256	6	6	6
Set 5	512	1	2	11
Set 6	200	1	20	0
Set 7	200	1	20	20
Set 8	2000	10	10	50

First we compare all methods on a linear classification dataset to allow a fair comparison with the linear models. To generate a multidimensional classification problem, we use a randomly generated prototype vector which defines a hyperplane. The defining features of this plane are strongly relevant. Sample points are generated in this feature space and the class is determined by the side of the hyperplane the points lie on. Weakly relevant features are constructed by replacing a feature of the original feature space with its linear combination. The elements of this combination are highly correlated and produce a set of redundant features. By removing the original feature and replacing it with those elements we achieve weak relevance by definition. Irrelevant features are sampled from a standard normal distribution.

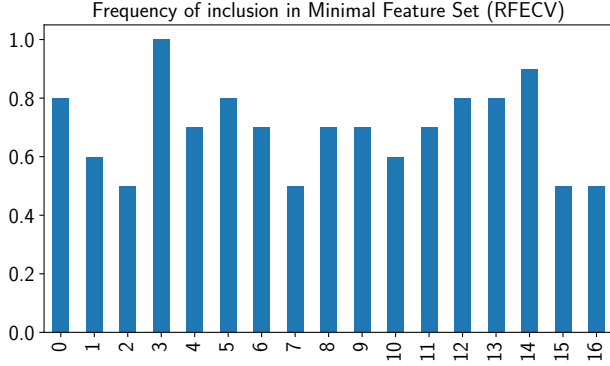
We generate 8 datasets with different feature set compositions as given by Table 1. For example, Set 1 consists of 150 samples (n), 6 strongly relevant features (*strong*), no weak relevant features (*weak*) and 6 irrelevant random features (*irr*).

All models given in section 3.2.1 are repeatedly fit on bootstraps of these datasets, resulting in 10 results per dataset per model which are averaged in the following. The prediction accuracy on the datasets is listed in Table 2 with sufficient accuracy for all models.

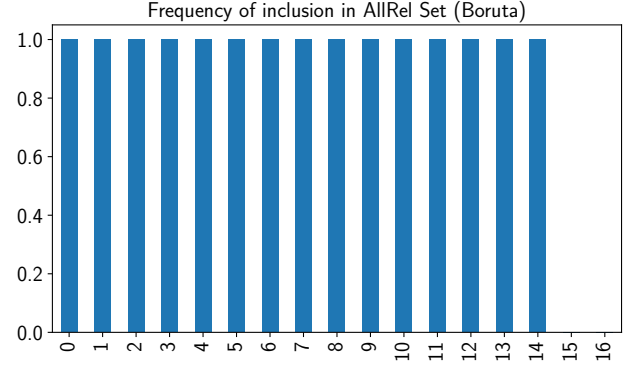
The results of the F_1 measure evaluation are given in Table 3. Additionally, we recorded the runtime of all methods while performing feature selection which is given in Table 4.

Most evident is the perfect score of FRI in this setting while being the slowest method. SQ follows second and performs very good feature selection in all cases while being the second fastest. The RFE scheme using random forests (RF) performed worst, not selecting weakly relevant features such as in Set 6 and 7. The ElasticNet also not as sensitive in this experiment, but showing the fastest runtime given its simplicity.

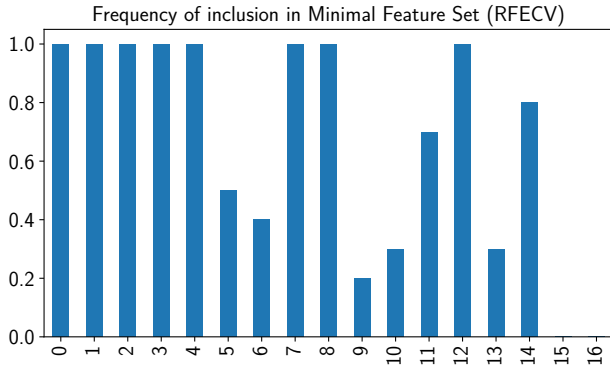
3.2.5. Non-linear problems. While many problems can be tackled using linear models, many relations are non-linear in



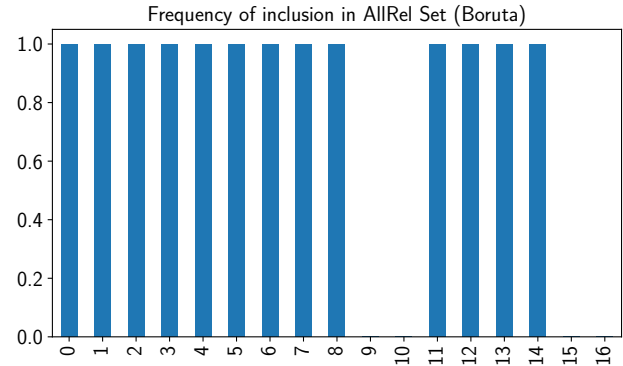
(a) RFECV, feature_fraction = 0.1



(b) Boruta, feature_fraction = 0.1



(c) RFECV, feature_fraction = 1



(d) Boruta, feature_fraction = 1

Figure 3: Frequency of feature selection for dataset with 5 strongly relevant features (0-4), 10 weakly relevant features (5-14) and 2 irrelevant features (15-16) as described in section 3.2.2. Vertical bars represent probability that each feature was included in the selected feature set for RFECV and Boruta. The random forest model used different setting for feature_fraction.

TABLE 2: Training accuracy of models on *linear separable* classification data generated according to Table 1.

	ElasticNet	FRI	RF	SQ
Set 1	0.97	0.99	0.83	1.00
Set 2	0.99	0.99	1.00	1.00
Set 3	0.98	0.99	0.86	1.00
Set 4	0.98	0.99	0.87	1.00
Set 5	0.98	1.00	0.95	1.00
Set 6	0.97	1.00	0.89	0.99
Set 7	0.98	0.99	0.90	1.00
Set 8	0.98	1.00	0.91	1.00

TABLE 3: Average F_1 measure on *linear separable* data sets regarding feature classification.

	model	ElasticNet	FRI	RF	SQ
score	type				
	data				
	Set 1	0.91	1.00	0.86	0.98
	Set 2	0.75	1.00	0.29	0.92
	Set 3	0.83	1.00	0.67	0.97
	Set 4	0.86	1.00	0.68	0.93
	Set 5	0.85	1.00	0.77	0.99
	Set 6	0.52	1.00	0.17	0.99
	Set 7	0.38	1.00	0.17	0.95
Set 8	0.83	1.00	0.65	0.99	

nature. In this experiment we generate data which can not be separated with a linear hyperplane. Our assumption is, that the linear models ElasticNet and FRI should not perform good in this case.

We utilize the classification data generation function from scikit-learn⁴ to create binary classification data

4. https://scikit-learn.org/stable/modules/generated/sklearn.datasets.make_classification.html

with multiple opposing cluster of samples (parameter `n_clusters_per_class = 2`). We then process the informative features to produce weakly relevant (redundant) features and additional irrelevant features⁵.

Again, we generate sets with different feature configurations given in Table 5. We fit all models on 20 newly

5. Generation function available at https://github.com/lpfnan/arfs_gen

TABLE 4: Runtime in seconds (rounded) for experiment described in section 3.2.4.

	model data	ElasticNet	FRI	RF	SQ
runtime	Set 1	0	2	0	1
	Set 2	0	2	0	1
	Set 3	0	2	0	1
	Set 4	0	3	1	3
	Set 5	0	5	2	6
	Set 6	0	4	1	2
	Set 7	0	6	3	3
	Set 8	2	201	138	80

TABLE 5: Parameters of generated datasets for *non-linearly* separable classification data. Numeric difference between $n_features$ and strong (n_strel) and weak ($n_redundant$) relevant features is filled with irrelevant features, i.e. NL 1 contains 10 irrelevant features.

Set	NL 1	NL 2	NL 3	NL 4
n_features	20	20	50	80
n_strel	10	4	10	10
n_redundant	0	10	10	10

generated sets and compute the average metric values.

The combined metrics are given in Table 6 per dataset and more concise in Table 7 averaged over all sets. Both linear models show low training accuracy at 70% which hints that the linear models can not replicate the non-linear relation. The random forest based models (SQ, RF) fare better with accuracies $\geq 83\%$. While not exceptional, it highlights the difficulty of this toy classification problem.

First we analyse the general feature selection accuracy without discriminating between the relevance class subsets. Overall the selection accuracy of all methods get worse which is apparent in Table 7. The ElasticNet scores last, with an average recall of 0.65 followed by the recursive feature elimination method (RF) with 0.71. It is beat by FRI with recall of 0.8 even though it can not handle non-linear separable data. SQ has the highest recall with 91% of relevant features recognized. When also considering the precision, we see that FRI scores perfect precision with no false positives such that the overall result actually is better than SQ here with an F_1 of 0.88 which is slightly better than SQ.

3.2.6. Accuracy of Relevance Classification. The general feature selection accuracy evaluation in the sections before does not consider the difference between strong and weak relevance. From all models considered before, only FRI and SQ provide can provide this distinction. We now present the precision and recall on the subsets \mathcal{S} and \mathcal{W} recorded in the experiments in the previous evaluations on linearly and non-linearly separable datasets. For that we reuse the same metrics as before and independently evaluate each subset.

TABLE 6: Statistics of benchmark with *non-linearly* separable classification data as generated according to Table 5. *Precision*, *recall* and *f1* quantify the feature selection performance, whereby *accuracy* denotes the training accuracy (quality of model fit).

	model dataset	ElasticNet	FRI	RF	SQ
precision	NL 1	0.88	1.00	0.71	0.97
	NL 2	0.86	1.00	0.67	1.00
	NL 3	0.41	1.00	1.00	0.80
	NL 4	0.21	1.00	1.00	0.62
recall	NL 1	0.70	0.53	1.00	1.00
	NL 2	0.86	1.00	0.86	1.00
	NL 3	0.55	0.77	0.55	0.63
	NL 4	0.50	0.89	0.45	1.00
f1	NL 1	0.78	0.69	0.83	0.98
	NL 2	0.86	1.00	0.75	1.00
	NL 3	0.47	0.87	0.71	0.70
	NL 4	0.30	0.94	0.62	0.77
accuracy	NL 1	0.66	0.67	0.79	0.82
	NL 2	0.70	0.75	0.81	0.83
	NL 3	0.60	0.66	0.87	0.89
	NL 4	0.73	0.74	0.86	0.90

TABLE 7: Mean over all datasets of benchmark with non-linearly separable classification data as in Table 6.

model	precision	recall	f1	accuracy
ElasticNet	0.59	0.65	0.60	0.67
FRI	1.00	0.80	0.88	0.71
RF	0.85	0.71	0.73	0.83
SQ	0.85	0.91	0.86	0.86

TABLE 8: Analysis of feature selection accuracy grouped by relevance class subsets on *linearly separable data*.

		Weakly		Strongly	
		FRI	SQ	FRI	SQ
precision	Set 1	1.00	1.00	1.00	1.00
	Set 2	1.00	0.87	1.00	1.00
	Set 3	1.00	0.99	1.00	0.74
	Set 4	1.00	0.99	1.00	0.86
	Set 5	0.99	0.99	1.00	0.49
	Set 6	1.00	1.00	1.00	0.50
	Set 7	1.00	0.99	1.00	0.46
	Set 8	0.99	1.00	1.00	0.90
recall	Set 1	1.00	1.00	1.00	0.95
	Set 2	1.00	1.00	1.00	1.00
	Set 3	1.00	0.72	1.00	0.99
	Set 4	1.00	0.68	1.00	0.99
	Set 5	1.00	0.50	1.00	1.00
	Set 6	1.00	0.93	1.00	1.00
	Set 7	1.00	0.88	1.00	1.00
	Set 8	1.00	0.85	1.00	1.00

TABLE 9: Mean feature selection accuracy metrics grouped by relevance class subsets and averaged over all *linearly separable* datasets. (Detailed: Table 8)

	Weakly		Strongly	
	FRI	SQ	FRI	SQ
precision	1.00	0.98	1.00	0.74
recall	1.00	0.82	1.00	0.99

TABLE 10: Analysis of feature selection accuracy grouped by relevance class subsets on *non-linearly separable* data.

dataset	model	Weakly		Strongly	
		precision	recall	precision	recall
NL 1	FRI	-	-	1.00	0.11
	SQ	-	-	1.00	0.59
NL 2	FRI	0.83	1.00	1.00	0.50
	SQ	0.99	1.00	1.00	0.97
NL 3	FRI	0.67	1.00	1.00	0.04
	SQ	0.20	0.14	0.86	0.87
NL 4	FRI	0.57	1.00	1.00	0.01
	SQ	0.18	0.39	0.65	0.87

In Table 8 we see the metrics for all linear datasets and in Table 9 their mean. The recall for strongly relevant features is near perfect for our proposed method (SQ). The precision is not perfect though and sometimes FP selections occur, as can be seen in Set 2 where the recall is 100% (all relevant features were selected) but the precision is at 87% which hints at irrelevant features being selected as well. Additionally, in some cases such as in Set 5 SQ tends to select weakly relevant features as strongly relevant.

We also compare both methods in the much harder task from section 3.2.5. The detailed results are given in Table 10 and their average in Table 11.

Here the results are mixed. FRI achieves perfect recall for weakly relevant features but misses a lot of strongly relevant ones. This is extreme in sets NL 3 and NL 4 where it has an average recall of 0.04 and 0.01. Apparently, it is more inclined to classify strongly relevant features as weakly relevant because the precision of the latter is decreased. SQ is also showing many false negative weakly relevant features while also selecting false positives which hurts its score

TABLE 11: Mean feature selection accuracy metrics grouped by relevance class subsets and averaged over all *non-linearly separable* datasets. (Detailed: Table 10)

model	Weakly		Strongly	
	precision	recall	precision	recall
FRI	0.69	1.00	1.00	0.16
SQ	0.46	0.51	0.88	0.82

it that setting. On the other hand it is quite balanced in the classification of strongly relevant features and correctly selects and classifies over 80% of them.

4. Conclusion

In this paper we presented a new all-relevant feature selection approach which builds upon existing methods. We combine several ideas to decompose the all-relevant feature set into strongly and weakly relevant features. We highlight the general selection accuracy in the linear and non-linear case which outperforms many existing approaches. While it is not as accurate as FRI in the linear case, our method is nearly equal to FRI in the non-linear case. In general, it performs faster and scales better to bigger datasets than FRI. Our method also enables the classification of feature relevance classes, which is new in the case of non-linear models. Here the results are more mixed and further analysis in correct parameterization to control false positives is needed. Though the ratio of false positives and false negatives is always a trade off which should be considered in application.

References

- [DA06] Ramón Díaz-Uriarte and Sara Alvarez de Andrés. “Gene Selection and Classification of Microarray Data Using Random Forest”. en: *BMC Bioinformatics* 7.1 (Jan. 2006), p. 3. ISSN: 1471-2105. DOI: 10.1186/1471-2105-7-3.
- [DSS] Frauke Degenhardt, Stephan Seifert, and Silke Szymczak. “Evaluation of Variable Selection Methods for Random Forests and Omics Data Sets”. en. In: *Briefings in Bioinformatics* (). DOI: 10.1093/bib/bbx124.
- [GMS17] Baptiste Gregorutti, Bertrand Michel, and Philippe Saint-Pierre. “Correlation and Variable Importance in Random Forests”. en. In: *Statistics and Computing* 27.3 (May 2017), pp. 659–678. ISSN: 1573-1375. DOI: 10.1007/s11222-016-9646-1.
- [Göp+18] Christina Göpfert et al. “Interpretation of Linear Classifiers by Means of Feature Relevance Bounds”. In: *Neurocomputing* 298 (July 2018), pp. 69–79. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2017.11.074.
- [Ke+17] Guolin Ke et al. “LightGBM: A Highly Efficient Gradient Boosting Decision Tree”. In: *Advances in Neural Information Processing Systems* 30. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 3146–3154.
- [KJ97] Ron Kohavi and George H. John. “Wrappers for Feature Subset Selection”. In: *Artif. Intell.* 97.1-2 (Dec. 1997), pp. 273–324. ISSN: 0004-3702. DOI: 10.1016/S0004-3702(97)00043-X.

- [KR10] Miron B. Kursa and Witold R. Rudnicki. “Feature Selection with the Boruta Package”. en. In: *Journal of Statistical Software* 36.11 (2010). ISSN: 1548-7660. DOI: 10.18637/jss.v036.i11.
- [NBT] Roland Nilsson, Johan Björkegren, and Jesper Tegner. “Consistent Feature Selection for Pattern Recognition in Polynomial Time”. en. In: (), p. 24.
- [Ped+11] F. Pedregosa et al. “Scikit-Learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [Pfa+19] Lukas Pfannschmidt et al. “FRI-Feature Relevance Intervals for Interpretable and Interactive Data Exploration”. In: *2019 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*. July 2019, pp. 1–10. DOI: 10.1109/CIBCB.2019.8791489.
- [PJ20] Lukas Pfannschmidt and Jonathan Jakob. *Lpfann/Fri: Feature Relevance Intervals*. Zenodo. Mar. 2020. DOI: 10.5281/zenodo.3734217.
- [Rud+06] Witold R. Rudnicki et al. “A Statistical Method for Determining Importance of Variables in an Information System”. en. In: *Rough Sets and Current Trends in Computing*. Ed. by Salvatore Greco et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2006, pp. 557–566. ISBN: 978-3-540-49842-1. DOI: 10.1007/11908029_58.
- [Tib96] Robert Tibshirani. “Regression Shrinkage and Selection via the Lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 58.1 (1996), pp. 267–288. ISSN: 00359246.
- [TL11] Laura Toloşi and Thomas Lengauer. “Classification with Correlated Features: Unreliability of Feature Ranking and Solutions”. In: *Bioinformatics* 27.14 (July 2011), pp. 1986–1994. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/btr300.
- [ZH05] Hui Zou and Trevor Hastie. “Regularization and Variable Selection via the Elastic Net”. In: *Journal of the Royal Statistical Society, Series B* 67 (2005), pp. 301–320.