

Ranking to Learn:

Feature Ranking and Selection via Eigenvector Centrality

Giorgio Roffo^{1,2}(✉) and Simone Melzi²

¹ School of Computing Science, University of Glasgow, Glasgow, UK
Giorgio.Roffo@glasgow.ac.uk

² Department of Computer Science, University of Verona, Verona, Italy
Simone.Melzi@univr.it

Abstract. In an era where accumulating data is easy and storing it inexpensive, feature selection plays a central role in helping to reduce the high-dimensionality of huge amounts of otherwise meaningless data. In this paper, we propose a graph-based method for feature selection that ranks features by identifying the most important ones into arbitrary set of cues. Mapping the problem on an affinity graph - where features are the nodes - the solution is given by assessing the importance of nodes through some indicators of centrality, in particular, the *Eigenvector Centrality (EC)*. The gist of EC is to estimate the importance of a feature as a function of the importance of its neighbors. Ranking central nodes individuates candidate features, which turn out to be effective from a classification point of view, as proved by a thoroughly experimental section. Our approach has been tested on 7 diverse datasets from recent literature (e.g., biological data and object recognition, among others), and compared against filter, embedded and wrappers methods. The results are remarkable in terms of accuracy, stability and low execution time.

Keywords: Feature selection · Ranking · High dimensionality · Data mining

1 Introduction

As data collection technologies advance and computer power grows, a torrent of data is generated in almost every field computers are used [5]. Because the volume, velocity, variety and complexity of datasets is continuously increasing, pattern recognition methodologies have become indispensable in order to extract useful information from huge amounts of otherwise meaningless data.

Feature Selection (FS) is one of the long existing methods that deals with these problems [14]. Its objective is to select a minimal subset of those attributes that allows a problem to be clearly defined. By choosing a minimal subset of features, irrelevant and redundant features are removed according to some reasonable criteria so that the original task can be achieved equally well, if not better.

FS techniques can be partitioned into three classes [14]: *wrappers* (see Fig. 1), which use classifiers to score a given subset of features; *embedded* methods (in Fig. 3), which inject the selection process into the learning of the classifier; and *filter* methods (see Fig. 2), which analyze intrinsic properties of data, ignoring the classifier. Filters are also (relatively) robust against overfitting.

Most of these methods can perform two operations, *ranking* and *subset selection*: in the former, the importance of each individual feature is evaluated, usually by neglecting potential interactions among the elements of the joint set [8]; in the latter, the final subset of features to be selected is provided. In some cases, these two operations are performed sequentially (first the ranking, then the selection) [7, 12, 17, 24, 35]; in other cases, only the selection is carried out [13]. Usually, the subset selection is supervised, while in the ranking case, methods can be supervised or not. FS is *NP-hard* [14]; if there are n features in total, the goal is to select the optimal subset of $m \ll n$, by evaluating $\binom{n}{m}$ combinations; therefore, suboptimal search strategies are considered (see Sect. 2 for an overview). With the filters, features are first considered individually, ranked, and then a subset is extracted, some examples are Mutual Information [35], Relief-F [24], Inf-FS [30, 31] unsupervised and not [26], and mRMR [27]. Conversely, with wrapper and embedded methods, subsets of features are sampled, evaluated, and finally kept as the final output, for instance, FSV [7, 12] and SVM-RFE [17].

In this work, we propose a novel graph-based feature selection algorithm that ranks features according to a graph centrality measure (Eigenvector centrality [6]). The main idea behind the method is to map the problem on an affinity graph, and to model pairwise relationships among feature distributions by weighting the edges connecting them.

The novelty of the proposed method in terms of the state of the art is that it assigns a score of “importance” to each feature by taking into account all the other features mapped as nodes on the graph, bypassing the combinatorial problem in a methodologically sound fashion. Indeed, eigenvector centrality differs from other measurements (e.g., degree centrality) since a node - feature - receiving many links does not necessarily have a high eigenvector centrality. The reason is that not all nodes are equivalent, some are more relevant than others, and, reasonably, endorsements from important nodes count more (see Sect. 3.2). Noteworthy, another important contribution of this work is the scalability of the method. Indeed, centrality measurements can be implemented using the Map Reduce paradigm [20, 23, 34], which makes the algorithm prone to a possible distributed version [29].

Our approach is extensively tested on 7 benchmarks of cancer classification and prediction on genetic data (*Colon* [2], *Prostate* [11], *Leukemia* [11], *Lymphoma* [11]), handwritten recognition (GINA [1]), generic feature selection datasets (MADELON [15]), and object recognition (PASCAL VOC 2007 [9]). We compare the proposed method on these data, while comparing it against seven efficient approaches under different conditions (number of features selected and number of training samples considered), overcoming all of them in terms of ranking stability or classification accuracy.

Finally, we provide an open and portable library of feature selection algorithms, integrating the methods with uniform input and output formats to facilitate large scale performance evaluation. The *Feature Selection Library* (FSLib Matlab Toolbox¹) and interfaces are fully documented. The library integrates directly with MATLAB, a popular language for machine learning and pattern recognition research.

The rest of the paper is organized as follows. A brief overview of the related literature is given in Sect. 2, mostly focusing on the comparative approaches we consider in this work. Our feature selection algorithm is described in Sect. 3. Graph construction and weighting are presented in Sects. 3.1 and 3.2 respectively, while the employed Eigenvector centrality is discussed in Sect. 3.3. Section 4 contains the experimental evaluations and results. Finally, conclusions are provided in Sect. 6.

2 Related Work

Since the mid-1990s, few domains explored used more than 50 features. The situation has changed considerably in the past few years and most papers explore domains with hundreds to tens of thousands of features. New approaches are proposed to address these challenging tasks involving many irrelevant and redundant cues and often comparably few training examples. Among the most used FS strategies, *Relief-F* [24] is an iterative, randomized, and supervised approach that estimates the quality of the features according to how well their values differentiate data samples that are near to each other; it does not discriminate among redundant features (i.e., may fail to select the most useful features), and performance decreases with few data. Similar problems affect SVM-RFE (RFE) [17], which is a wrapper method (see Fig. 1) that selects features in a sequential, backward elimination manner, ranking high a feature if it strongly separates the samples by means of a linear SVM.




Fig. 1. Wrapper models involve optimizing a predictor as part of the selection process. They tend to give better results but filter methods are usually computationally less expensive than wrappers.

¹ The FSLib is publicly available on File Exchange - MATLAB Central at: <https://it.mathworks.com/matlabcentral/fileexchange/56937-feature-selection-library>.




Fig. 2. Filter methods: the selection of features is independent of the classifier used. They rely on the general characteristics of the training data to select features with independence of any predictor.

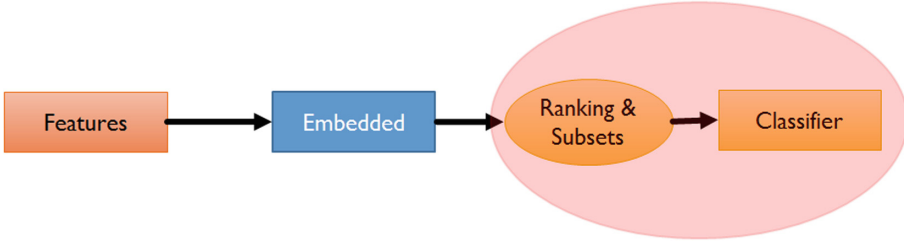


Fig. 3. In embedded methods the learning part and the feature selection part can not be separated.

Batti [4] has developed the Mutual Information-Based Feature Selection (MIFS) criterion, where the features are selected in a greedy manner. Given a set of existing selected features, at each step it locates the feature x_i that maximizes the relevance to the class. The selection is regulated by a proportional term β that measures the overlap information between the candidate feature and existing features. In [36] the authors proposed a *graph-based* filter approach to feature selection, that constructs a graph in which each node corresponds to each feature, and each edge has a weight corresponding to mutual information (MI) between features connected by that edge. This method performs dominant set clustering to select a highly coherent set of features and then it selects features based on the multidimensional interaction information (MII). Another effective yet fast filter method is the *Fisher* method [13], it computes a score for a feature as the ratio of interclass separation and intraclass variance, where features are evaluated independently, and the final feature selection occurs by aggregating the m top ranked ones. Other widely used filters are based on mutual information, dubbed *MI* here [35], which considers as a selection criterion the mutual information between the distribution of the values of a given feature and the membership to a particular class. Mutual information provides a principled way of measuring the mutual dependence of two variables, and has been used by a number of researchers to develop information theoretic feature selection criteria. Even in the last case, features are evaluated independently, and the final feature selection occurs by aggregating the m top ranked ones. Maximum-Relevance Minimum-Redundancy criterion (MRMR) [27] is an efficient incremental search algorithm. Relevance scores are assigned by maximizing the joint mutual information between the class variables and the subset of selected features. The computation of the information between high-dimensional vectors is impractical, as

the time required becomes prohibitive. To face this problem the mRMR propose to estimate the mutual information for continuous variables using Parzen Gaussian windows. This estimate is based on a heuristic framework to minimize redundancy and uses a series of intuitive measures of relevance and redundancy to select features. Note, it is equivalent to MIFS with $\beta = \frac{1}{n-1}$, where n is the number of features. Selecting features in unsupervised learning scenarios is a much harder problem, due to the absence of class labels that would guide the search for relevant information. In this scenario, we compare our approach against the recent unsupervised graph-based filter dubbed Inf-FS [31]. In the Inf-FS formulation, each feature is a node in the graph, a path is a selection of features, and the higher the centrality score, the most important (or most different) the feature. It assigns a score of “importance” to each feature by taking into account all the possible feature subsets as paths on a graph. Another unsupervised method is the Laplacian Score (LS) [19], where the importance of a feature is evaluated by its power of locality preserving. In order to model the local geometric structure, this method constructs a nearest neighbor graph. LS algorithm seeks those features that respect this graph structure. Finally, for the embedded method (see Fig. 3), we include the *feature selection via concave minimization* (FSV) [7], where the selection process is injected into the training of an SVM by a linear programming technique.

3 Proposed Method

3.1 Building the Graph

Given a set of features $X = \{x^{(1)}, \dots, x^{(n)}\}$ we build an undirected graph $G = (V, E)$; where V is the set of vertices corresponding, one by one, to each variable x . E codifies (weighted) edges among features. Let the adjacency matrix A associated with G define the nature of the weighted edges: each element a_{ij} of A , $1 \leq i, j \leq n$, represents a pairwise potential term. Potentials can be represented as a binary function $\varphi(x^{(i)}, x^{(j)})$ of the nodes $x^{(k)}$ such as:

$$a_{ij} = \varphi(x^{(i)}, x^{(j)}). \quad (1)$$

The graph can be weighted according to different heuristics, therefore the function φ can be handcrafted or automatically learned from data.

3.2 φ -Design

The design of the φ function is a crucial operation. In this work, we weight the graph according to good reasonable criteria, related to class separation, so as to address the classification problem. In other words, we want to rank features according to how well they discriminate between two classes. Hence, we draw upon best-practice in FS and propose an ensemble of two different measures capturing both relevance (supervised) and redundancy (unsupervised) proposing

a kernelized-based adjacency matrix. Before continuing with the discussion, note that each feature distribution $x^{(i)}$ is normalized so as to sum to 1.

Firstly, we apply the Fisher criterion:

$$f_i = \frac{|\mu_{i,1} - \mu_{i,2}|^2}{\sigma_{i,1}^2 + \sigma_{i,2}^2},$$

where $\mu_{i,C}$ and $\sigma_{i,C}$ are the mean and standard deviation, respectively, assumed by the i -th feature when considering the samples of the C -th class. The higher f_i , the more discriminative the i -th feature. However, a natural generalization of this score into a multi-class framework is given by

$$f_i = \frac{\sum_{c=1}^C (\mu_{i,c} - \mu_i)^2}{\sigma_i^2}$$

where μ_i and σ_i denote the mean and standard deviation of the whole data set corresponding to the i -th feature (i.e., $\sigma_i^2 = \sum_{c=1}^C (\sigma_{i,c})^2$).

Because we are given class labels, it is natural that we want to keep only the features that are related to or lead to these classes. Therefore, we use mutual information to obtain a good feature ranking that score high features highly predictive of the class.

$$m_i = \sum_{y \in Y} \sum_{z \in x^{(i)}} p(z, y) \log \left(\frac{p(z, y)}{p(z)p(y)} \right),$$

where Y is the set of class labels, and $p(\cdot, \cdot)$ the joint probability distribution.

A kernel k is then obtained by the matrix product

$$k = (f \cdot m^\top),$$

where f and m are $n \times 1$ column vectors normalized in the range 0 to 1, and k results in a $n \times n$ matrix.

To boost the performance, we introduce a second feature-evaluation metric based on standard deviation [17] – capturing the amount of variation or dispersion of features from average – as follows:

$$\Sigma(i, j) = \max \left(\sigma^{(i)}, \sigma^{(j)} \right),$$

where σ being the standard deviation over the samples of x , and Σ turns out to be a $n \times n$ matrix with values $\in [0, 1]$.

Finally, the adjacency matrix A of the graph G is given by

$$A = \alpha k + (1 - \alpha) \Sigma, \tag{2}$$

where α is a loading coefficient $\in [0, 1]$. The generic entry a_{ij} accounts for how much discriminative are the feature i and j when they are jointly considered; at the same time, a_{ij} can be considered as a weight of the edge connecting the nodes i and j of a graph, where the i -th node models the i -th feature distribution (we report the sketch of our method in Algorithm 1).

Algorithm 1. Eigenvector Centrality Feature Selection (EC-FS)**Input:** $X = \{x^{(1)}, \dots, x^{(n)}\}$, $Y = \{y^{(1)}, \dots, y^{(n)}\}$ **Output:** v_0 ranking scores for each feature**- Building the graph** C_1 positive class, C_2 negative class**for** $i = 1 : n$ **do** Compute $\mu_{i,1}$, $\mu_{i,2}$, $\sigma_{i,1}$, and $\sigma_{i,2}$ Fisher score: $f(i) = \frac{(\mu_{i,1} - \mu_{i,2})^2}{\sigma_{i,1}^2 + \sigma_{i,2}^2}$ Mutual Information: $m(i) = \sum_{y \in Y} \sum_{z \in x^{(i)}} p(z, y) \log \left(\frac{p(z, y)}{p(z)p(y)} \right)$ **end for****for** $i = 1 : n$ **do** **for** $j = 1 : n$ **do** $k(i, j) = f(i)m(j)$, $\Sigma(i, j) = \max(\sigma^{(i)}, \sigma^{(j)})$, $A(i, j) = \alpha k(i, j) + (1 - \alpha)\Sigma(i, j)$ **end for****end for****- Ranking**Compute eigenvalues $\{A\}$ and eigenvectors $\{V\}$ of A $\lambda_0 = \max_{\lambda \in A} (\text{abs}(\lambda))$ **return** v_0 the eigenvector associated to λ_0

3.3 Eigenvector Centrality

From a graph theory perspective identifying the most important nodes corresponds to individuate some indicators of centrality within a graph (e.g., the relative importance of nodes). A first way used in graph theory is to study accessibility of nodes, see [10, 28] for example. The idea is to compute A^l for some suitably large l (often the diameter of the graph), and then use the row sums of its entries as a measure of accessibility (i.e. $\text{scores}(i) = [A^l \mathbf{e}]_i$, where \mathbf{e} is a vector with all entries equal to 1). The accessibility index of node i would thus be the sum of the entries in the i -th row of A^l , and this is the total number of paths of length l (allowing stopovers) from node i to all nodes in the graph. One problem with this method is that the integer l seems arbitrary. However, as we count longer and longer paths, this measure of accessibility converges to a index known as eigenvector centrality measure (EC) [6].

The basic idea behind the EC is to calculate v_0 the eigenvector of A associated to the largest eigenvalue. Its values are representative of how strongly each node is connected to the other nodes. Since the limit of A^l as l approaches a large positive number L converges to v_0 ,

$$\lim_{l \rightarrow L} [A^l \mathbf{e}] = v_0, \quad (3)$$

the EC index makes the estimation of indicators of centrality free of manual tuning over l , and computationally efficient.

Let us consider a vector, for example \mathbf{e} , that is *not* orthogonal to the principal vector v_0 of A . It is always possible to decompose \mathbf{e} using the eigenvectors as basis with a coefficient $\beta_0 \neq 0$ for v_0 . Hence:

$$\mathbf{e} = \beta_0 v_0 + \beta_1 v_1 + \dots + \beta_n v_n, \quad (\beta_0 \neq 0). \quad (4)$$

Then

$$\begin{aligned} A\mathbf{e} &= A(\beta_0 v_0 + \beta_1 v_1 + \dots + \beta_n v_n) = \beta_0 A v_0 + \beta_1 A v_1 + \dots + \beta_n A v_n \\ &= \beta_0 \lambda_0 v_0 + \beta_1 \lambda_1 v_1 + \dots + \beta_n \lambda_n v_n. \end{aligned} \quad (5)$$

So in the same way:

$$\begin{aligned} A^l \mathbf{e} &= A^l(\beta_0 v_0 + \beta_1 v_1 + \dots + \beta_n v_n) = \beta_0 A^l v_0 + \beta_1 A^l v_1 + \dots + \beta_n A^l v_n \\ &= \beta_0 \lambda_0^l v_0 + \beta_1 \lambda_1^l v_1 + \dots + \beta_n \lambda_n^l v_n, \quad (\beta_0 \neq 0). \end{aligned} \quad (6)$$

Finally we divide by the constant $\lambda_0^l \neq 0$ (see Perron-Frobenius theorem [25]),

$$\frac{A^l \mathbf{e}}{\lambda_0^l} = \beta_0 v_0 + \frac{\lambda_1^l \beta_1 v_1}{\lambda_0^l} + \dots + \frac{\lambda_n^l \beta_n v_n}{\lambda_0^l}, \quad (\beta_0 \neq 0). \quad (7)$$

The limit of $\frac{A^l \mathbf{e}}{\lambda_0^l}$ as l approaches infinity equals $\beta_0 v_0$ since $\lim_{l \rightarrow \infty} \frac{\lambda_i^l}{\lambda_0^l} = 0, \forall i > 0$. What we see here is that as we let l increase, the ratio of the components of $A^l \mathbf{e}$ converges to v_0 . Therefore, marginalizing over the columns of A^l , with a sufficiently large l , corresponds to calculate the principal eigenvector of matrix A [6]. Figure 4 illustrates a toy example of three random planar graphs. Graphs are made of 700 nodes and they are weighted by the Euclidean distance between each pair of points. In the example, high scoring nodes are those ones farther from the mean (i.e., the distance is conceived as quantity to maximize), the peculiarity of the eigenvector centrality is that a node is important if it is linked to other important nodes (higher scores).




Fig. 4. Eigenvector centrality plots for three random planar graphs. On the left, a simple Gaussian distribution where central nodes are at the peripheral part of the distribution as expected. The central and right plots, some more complicated distributions, a node receiving many links does not necessarily have a high eigenvector centrality. Best viewed in color.

To the aim of this work, the use of eigenvector centrality allows to individuate candidate features, which turn out to be effective from a classification point of view, since indicators of centrality characterize the *global* (as opposed to local) prominence of a feature in the graph. Summarizing, the gist of eigenvector centrality is to compute the centrality of a node as a function of the centralities of its neighbors.

4 Experiments and Results

4.1 Datasets and Comparative Approaches

The datasets are chosen for letting the proposed method deal with diverse FS scenarios, as shown on Table 1. In the details, we consider the problems of dealing with few training samples and many features (*few train* in the table), unbalanced classes (*unbalanced*), or classes that severely overlap (*overlap*), or whose samples are noisy (*noise*) due to: (a) complex scenes where the object to be classified is located (as in the VOC series) or (b) many outliers (as in the genetic datasets, where samples are often *contaminated*, that is, artefacts are injected into the data during the creation of the samples). Lastly we consider the *shift* problem, where the samples used for the test are not congruent (coming from the same experimental conditions) with the training data.

Table 1. This table reports several attributes of the datasets used. The abbreviation *n.s.* stands for *not specified* (for example, in the object recognition datasets, the features are not given in advance).

Name	# samples	# classes	# feat.	<i>few train</i>	<i>unbal.</i> (+/-)	<i>overlap</i>	<i>noise</i>	<i>shift</i>
GINA [1]	3153	2	970			X		
MADOLON [16]	4.4K	2	500			X		
<i>Colon</i> [2]	62	2	2K	X	(40/22)		X	
<i>Lymphoma</i> [11]	45	2	4026	X	(23/22)			
<i>Prostate</i> [33]	102	2	6034	X	(50/52)			
<i>Leukemia</i> [11]	72	2	7129	X	(47/25)		X	X
VOC 2007 [9]	10K	20	n.s.		X		X	

Table 2 lists the methods in comparison, whose details can be found in Sect. 2. Here we just note their *type*, that is, f = filters, w = wrappers, e = embedded methods, and their *class*, that is, s = supervised or u = unsupervised (using or not using the labels associated with the training samples in the ranking operation). Additionally, we report their computational complexity (if it is documented in the literature). The computational complexity of our approach is $O(Tn + n^2)$.

The term Tn is due to the computation of the mean values among the T samples of every feature (n). The n^2 concerns the construction of the matrix A .

Table 2. List of the FS approaches considered in the experiments, specified according to their *Type*, class (*Cl.*), and complexity (*Compl.*). As for the complexity, T is the number of samples, n is the number of initial features, K is a multiplicative constant, i is the number of iterations in the case of iterative algorithms, and C is the number of classes. N/A indicates that the computational complexity is not specified in the reference paper.

Acronym	Type	Cl	Compl.
Fisher [13]	f	s	$\mathcal{O}(Tn)$
FSV [7, 12]	e	s	N/A
Inf-FS [31]	f	u	$\mathcal{O}(n^{2.37}(1+T))$
MI [35]	f	s	$\sim \mathcal{O}(n^2 T^2)$
LS [19]	f	u	N/A
Relief-F [24]	f	s	$\mathcal{O}(iTnC)$
RFE [17]	w/e	s	$\mathcal{O}(T^2 n \log_2 n)$
Ours	f	s	$\mathcal{O}(Tn + n^2)$





















As for the computation of the leading eigenvector, it costs $\mathcal{O}(m^2 n)$, where m is a number much smaller than n that is selected within the algorithm [22]. In the case that the algorithm can not be executed on a single computer, we refer the reader to [20, 23, 29, 34] for distributed algorithms.

4.2 Exp. 1: Deep Representation (CNN) with Pre-training

This section proposes a set of tests on the PASCAL VOC-2007 [9] dataset. In object recognition VOC-2007 is a suitable tool for testing models, therefore, we use it as reference benchmark to assess the strengths and weaknesses of using our approach regarding the classification task. For this reason, we compare our approach against 8 state-of-the-art FS methods reported in Table 2. This experiment considers as features the cues extracted with a deep convolutional neural network architecture (CNN). We selected the pre-trained model called very deep ConvNets [32], which performs favorably to the state of the art for classification and detection in the ImageNet Large-Scale Visual Recognition Challenge 2014 (ILSVRC). We use the 4,096-dimension activations of the last layer as image descriptors (i.e., 4,096 features in total). The VOC-2007 edition contains about 10,000 images split into train, validation, and test sets, and labeled with twenty object classes. A one-vs-rest SVM classifier for each class is learnt (where cross-validation is used to find the best parameter C and α mixing coefficient in Eq. 2 on the training data) and evaluated independently and the performance is measured as mean Average Precision (mAP) across all classes.

Table 3 serves to analyze and empirically clarify how well important features are ranked high by several FS algorithms. The amount of features used for the two experiments is very low: $\approx 3\%$ and $\approx 6\%$ of the total. The results are significant: our method achieved the best performance in terms of mean average

Table 3. Varying the cardinality of the selected features. The image classification results achieved in terms of average precision (AP) scores while selecting the first 128 (3%) and 256 (6%) features from the total 4,096.

PASCAL VOC 2007																
	First 128/4096 Features Selected								First 256/4096 Features Selected							
	Fisher	FSV	Inf-FS	LS	MI	ReliefF	RFE	Ours	Fisher	FSV	Inf-FS	LS	MI	ReliefF	RFE	Ours
	52.43	87.90	88.96	89.37	12.84	57.20	86.42	88.09	82.65	90.22	91.16	90.94	73.51	81.67	88.17	90.79
	13.49	80.74	80.43	80.56	13.49	49.10	82.14	80.94	83.21	80.07	83.36	84.21	75.04	71.27	83.30	84.72
	85.46	86.77	87.04	86.96	80.91	75.42	83.16	88.74	89.14	86.15	88.88	89.31	85.48	83.54	86.12	89.15
	79.04	83.58	85.31	83.51	61.50	63.75	78.55	86.90	87.05	80.68	87.24	87.84	75.25	73.30	86.13	87.42
	46.61	39.80	44.83	49.36	35.39	18.33	46.24	47.37	52.54	49.00	52.65	49.44	48.94	35.67	47.28	53.20
	12.29	72.89	76.69	76.98	12.29	31.54	74.68	76.27	77.32	78.69	79.23	79.97	59.23	63.83	79.38	80.57
	82.09	78.61	85.78	85.82	63.58	74.95	83.94	85.92	85.86	84.01	86.74	87.06	85.27	82.76	85.61	86.56
	75.29	82.25	83.34	81.81	40.96	66.95	81.02	83.29	83.46	83.49	85.61	84.98	79.16	76.78	84.50	85.57
	54.81	52.37	58.62	60.07	16.95	29.07	59.84	60.57	63.14	62.54	63.93	64.23	63.20	48.19	62.16	64.53
	47.98	61.68	59.23	65.50	11.42	11.42	62.96	60.55	66.51	70.18	67.96	71.54	22.96	51.28	64.20	69.71
	49.68	63.50	67.69	63.86	12.62	12.62	67.05	67.70	68.42	69.27	71.78	71.01	65.77	52.24	71.43	70.95
	81.06	80.57	83.16	83.21	70.70	68.12	80.07	83.00	84.24	84.15	85.08	85.20	82.03	74.85	83.52	85.20
	74.91	83.33	81.23	81.75	14.13	63.06	81.55	82.79	85.68	83.13	85.28	85.41	71.36	75.53	83.47	85.28
	13.18	71.42	81.32	80.24	13.18	34.43	76.57	82.20	84.29	81.16	84.20	83.81	81.01	70.68	82.97	84.12
	91.33	90.03	89.10	89.33	91.08	88.85	89.03	91.27	91.95	89.99	90.65	90.64	91.77	90.38	90.64	91.99
	47.89	39.40	45.38	47.94	13.23	13.30	48.61	49.05	54.94	47.95	53.86	54.31	48.98	34.74	50.18	55.88
	10.87	68.82	73.35	74.05	10.87	10.87	66.86	73.80	73.43	75.84	79.01	81.57	10.87	11.73	75.47	78.85
	45.87	56.08	58.94	58.92	13.30	13.31	62.06	61.32	66.46	59.77	63.07	63.92	58.78	44.74	66.68	64.86
	63.51	88.52	91.42	91.48	58.62	73.32	88.46	91.30	84.05	90.61	93.21	93.16	81.33	82.93	90.24	92.31
	64.29	65.61	66.79	62.99	47.25	24.96	67.10	67.30	71.44	69.19	70.56	70.75	71.39	55.59	73.17	72.49
mAP	54.60	71.69	74.43	74.69	34.72	44.03	73.32	75.42	76.79	75.80	78.17	78.47	66.57	63.09	76.73	78.71

precision (mAP) followed by the unsupervised filter methods LS and Inf-FS. As for the methods in comparison, one can observe the high variability in classification accuracy; indeed, results show that our method is robust to classes (i.e., by changing the testing class its performance is always comparable with the top scoring method).

4.3 Exp. 2: Testing on Microarray Databases

In application fields like biology is inconceivable to devise an analysis procedure which does not comprise a FS step. A clear example can be found in the analysis of expression microarray data, where the expression level of thousands of genes is simultaneously measured. Within this scenario, we tested the proposed approach on four well-known microarray benchmark datasets for two-class problems. Results are reported in Table 4. The testing protocol adopted in this experiment consists in splitting the dataset up to 2/3 for training and 1/3 for testing. In order to have a fair evaluation, the feature ranking has been calculated using only the training samples, and then applied to the testing samples. The classification

is performed using a linear SVM. For setting the best parameters (C of the linear SVM, and α mixing coefficient) we used a 5-fold cross validation on the training data. This procedure is repeated several times and results are averaged over the trials. Results are reported in terms of the Receiver Operating Characteristic or ROC curves. A widely used measurement that summarizes the ROC curve is the Area Under the ROC Curve (AUC) [3] which is useful for comparing algorithms independently of application. Hence, classification results for the datasets used show that the proposed approach produces superior results in all the cases. The overall performance indicates that our approach is more robust than the others, by changing the data it still produces high quality rankings.

Table 4. The tables show results obtained on the expression microarray scenario. Tests have been repeated 100 times, and the means of the computed AUCs are reported for each dataset.

Microarray databases												
Method	Colon						Leukemia					
	# Features					Time	# Features					Time
	50	100	150	200	Average		50	100	150	200	Average	
Fisher-S	91.25	88.44	89.38	87.81	89.22	0.02	99.33	99.78	99.78	99.78	99.66	0.01
FSV	85.00	88.12	89.38	89.69	88.04	0.18	98.22	98.44	99.11	99.33	98.77	0.37
Inf-FS	88.99	89.41	89.32	89.01	89.18	0.91	99.91	99.92	99.97	99.98	99.95	5.49
LS	90.31	89.06	89.38	90.00	89.68	0.03	98.67	99.33	99.56	99.56	99.28	0.07
MI	89.38	90.31	90.63	90.94	90.31	0.31	99.33	99.33	99.56	99.33	98.38	0.21
ReliefF	80.94	84.38	85.94	87.50	84.69	0.52	99.56	99.78	99.78	99.78	99.72	1.09
RFE	89.06	85.00	86.88	85.62	86.64	0.18	100	99.78	99.56	99.78	99.78	0.14
EC-FS	91.40	91.10	91.11	90.63	91.06	0.45	99.92	99.92	99.77	99.85	99.86	1.50
Method	Lymphoma						Prostate					
	# Features					Time	# Features					Time
	50	100	150	200	Average		50	100	150	200	Average	
Fisher-S	98.75	98.38	98.38	100	98.87	0.01	96.10	96.20	96.30	97.30	96.47	0.02
FSV	98.22	98.44	99.11	99.33	98.77	0.18	96.70	96.70	96.50	96.30	96.55	0.63
Inf-FS	98.12	98.75	98.75	99.38	98.75	7.61	96.80	96.90	97.10	96.70	96.87	26.85
LS	90.00	96.88	99.38	98.75	96.25	0.04	85.80	94.60	96.90	97.00	93.57	0.24
MI	97.50	98.75	99.38	99.38	98.75	0.59	96.00	96.90	96.00	96.20	96.27	1.01
ReliefF	96.80	97.00	98.80	98.80	97.85	0.74	92.72	93.46	93.62	93.85	93.41	2.68
RFE	96.00	98.00	98.80	99.00	97.95	0.02	93.40	96.40	97.10	96.32	95.80	0.3
EC-FS	99.40	99.20	99.60	99.20	99.20	1.50	96.28	96.90	96.80	98.10	97.02	2.81

The quality of a feature subset is measured by an estimate of the classification accuracy of a chosen classifier trained on the candidate subset. Stability of the ranking is an important aspect when the task is knowledge discovery. The rationale behind this fact is that the estimate of the quality of the candidate subsets usually depends on many the training/testing split of the data. Therefore different sequences of features may be returned from repeated runs of FS approaches. In such a case, it is important to define if these numerous different

subsets of features have approximately equal quality, otherwise presenting the user with only one subset may be misleading. We assessed the stability of the selected features using the Kuncheva index [21]. This stability measure represents the similarity between the set of rankings generated over the different splits of the dataset. The similarity between sequences of size N can be seen as the number of elements n they have in common (i.e. the size of their intersection). The Kuncheva index takes values in $[-1, 1]$, and the higher its value, the larger the number of commonly selected features in both sequences. The index is shown in Fig. 5, comparing our approach and the other methods. A valid alternative is the stability index based on Jensen-Shannon Divergence D_{JS} , proposed by [18], with a $[0, 1]$ range, where 0 indicates completely random rankings and 1 means stable rankings. Unlike Kuncheva measure, this metric is suitable for different algorithm outcomes: partial sublists (top- k lists) as well as the least studied partial ranked lists. Since in our case we work with full ranked lists, because all the feature selection algorithms considered in this study produce permutations of the original set of features, we preferred the widely used Kuncheva index. The proposed method shows, in most of the cases, a high stability whereas the highest performance is achieved.




Fig. 5. The Kuncheva stability indices for each method in comparison are presented. The figure reports the stability while varying the cardinality of the selected features from 10 to 200 on different benchmarks.

4.4 Exp. 2: Other Benchmarks

GINA has sparse input variables consisting of 970 features. It is a balanced data set with 49.2% instances belonging to the positive class. Results obtained on GINA indicate that the proposed approach overcomes the methods in comparison, and select the most useful features from a data set with high-complexity and dimensionality. MADELON is an artificial dataset, which was part of the NIPS 2003 feature selection challenge. It represents a two-class classification problem with continuous input variables. The difficulty is that the problem is multivariate and highly non-linear. Results are reported in Table 5. This gives a proof about the classification performance of our approach that is attained on the test sets of GINA and MADELON.

Table 5. Varying the cardinality of the selected features. (ROC) AUC (%) on different datasets by SVM classification. Performance obtained with the first 50, 100, 150, and 200 features.

FS challenge datasets												
Method	GINA - handwritten recognition						MADELON - artificial data					
	# Features						# Features					
	50	100	150	200	Average	Time	50	100	150	200	Average	Time
Fisher-S	89.8	89.4	90.2	90.4	89.9	0.05	61.9	63.0	62.3	64.0	62.5	0.02
FSV	81.9	83.7	82.0	83.6	82.7	138	59.9	60.6	61.0	61.0	60.7	732
Inf-FS	89.0	88.7	89.1	89.0	88.9	41	62.6	63.8	65.4	60.8	63.2	0.04
LS	82.2	82.4	83.4	83.2	82.7	1.30	62.8	62.9	63.3	64.7	63.4	8.13
MI	89.3	89.7	89.8	90.1	89.6	1.13	63.0	63.7	63.5	64.7	63.6	0.4
Relieff	77.9	76.3	77.3	76.9	77.2	0.12	62.9	63.1	63.2	64.9	63.5	10.41
RFE	82.2	82.4	83.4	83.2	82.7	6.60	55.0	61.2	57.1	60.2	56.5	50163
EC-FS	90.9	90.3	90.4	89.5	90.3	1.56	63.6	63.8	63.7	63.3	63.7	0.57

FS techniques definitely represent an important class of preprocessing tools, by eliminating uninformative features and strongly reducing the dimension of the problem space, it allows to achieve high performance, useful for practical purposes in those domains where high speed is required.

5 Reliability and Validity

In order to assess if the difference in performance is statistically significant, t-tests have been used for comparing the accuracies. Statistical tests are used to determine if the accuracies obtained with the proposed approach are significantly different from the others (whereas both the distribution of values were normal). The test for assessing whether the data come from normal distributions with unknown, but equal, variances is the *Lilliefors* test. Results have been obtained by comparing the results produced by each method over 100 trials (at each

trial corresponds a different split of the data). Given the two distributions x_p of the proposed method and x_c of the current competitor, of size 1×100 , a *two-sample t-test* has been applied obtaining a test decision for the null hypothesis that the data in vectors x_p and x_c comes from independent random samples from normal distributions with equal means and equal but unknown variances. Results (highlighted in Tables 4 and 5) show a statistical significant effect in performance (p-value < 0.05 , Lilliefors test $H = 0$).

6 Conclusion

In this paper we present the idea of solving feature selection via the Eigenvector centrality measure. We design a graph – where features are the nodes – weighted by a kernelized adjacency matrix, which draws upon the best-practice in feature selection while assigning scores according to how well features discriminate between classes. The method (supervised) estimates some indicators of centrality identifying the most important features within the graph. The results are remarkable: the proposed method has been extensively tested on 7 different datasets selected from different scenarios (i.e., object recognition, handwritten recognition, biological data, and synthetic testing datasets), in all the cases we achieve top performances against 7 competitors selected from recent literature in feature selection. Our approach is also robust and stable on different splits of the training data, it performs effectively in ranking high the most relevant features, and it has a very competitive complexity. This study also points to many future directions; focusing on the investigation of different implementations for parallel computing for big data analysis or focusing on the investigation of different relations among the features. Finally, we provide an open and portable library of feature selection algorithms, integrating the methods with uniform input and output formats to facilitate large scale performance evaluation. The *Feature Selection Library* (FSLib is available on *Matlab File Exchange* at <https://goo.gl/bvg1ha>) and interfaces are fully documented. The library integrates directly with MATLAB, a popular language for machine learning and pattern recognition research.

References

1. GINA digit recognition database. In: IEEE Conference International Joint Conference on Neural Networks (2007)
2. Alon, U., Barkai, N., Notterman, D.A., Gish, K., Ybarra, S., Mack, D., Levine, A.J.: Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. *PNAS* **96**(12), 6745–6750 (1999)
3. Bamber, D.: The area above the ordinal dominance graph and the area below the receiver operating characteristic graph. *J. Math. Psychol.* **12**(4), 387–415 (1975)
4. Battiti, R.: Using mutual information for selecting features in supervised neural net learning. *IEEE Trans. Neural Netw.* **5**(4), 537–550 (1994)

5. Bólon-Canedo, V., Sánchez-Maroo, N., Alonso-Betanzos, A.: Recent advances and emerging challenges of feature selection in the context of big data. *Knowl.-Based Syst.* **86**, 33–45 (2015)
6. Bonacich, P.: Power and centrality: a family of measures. *Am. J. Sociol.* **92**(5), 1170–1182 (1987)
7. Bradley, P.S., Mangasarian, O.L.: Feature selection via concave minimization and support vector machines. In: *Conference International Conference on Machine Learning (ICML)* (1998)
8. Duch, W., Wiecek, T., Biesiada, J., Blachnik, M.: Comparison of feature ranking methods based on information entropy. In: *IJCNN*, vol. 2. IEEE (2004)
9. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results (2007)
10. Garrison, W.L.: Connectivity of the interstate highway system. *Pap. Reg. Sci.* **6**(1), 121–137 (1960)
11. Golub, T.R.: Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* **286**(5439), 531–537 (1999)
12. Grinblat, G.L., Izetta, J., Granitto, P.M.: SVM based feature selection: why are we using the dual? In: *Conference Ibero-American Conference on AI* (2010)
13. Gu, Q., Li, Z., Han, J.: Generalized fisher score for feature selection. In: *Computing Research Repository (CoRR)* (2012)
14. Guyon, I.: *Feature Extraction: Foundations and Applications*, vol. 207. Springer Science & Business Media, Berlin (2006)
15. Guyon, I., Gunn, S., Ben-Hur, A., Dror, G.: Result analysis of the nips 2003 feature selection challenge. In: *NIPS*, pp. 545–552 (2004)
16. Guyon, I., Li, J., Mader, T., Pletscher, P.A., Schneider, G., Uhr, M.: Competitive baseline methods set new standards for the NIPS 2003 feature selection benchmark. *PRL* **28**(12), 1438–1444 (2007)
17. Guyon, I., Weston, J., Barnhill, S., Vapnik, V.: Gene selection for cancer classification using support vector machines. *Mach. Learn. J.* **46**(1), 389–422 (2002)
18. Guzmán-Martínez, R., Alaiz-Rodríguez, R.: Feature selection stability assessment based on the Jensen-Shannon divergence. In: Gunopulos, D., Hofmann, T., Malerba, D., Vazirgiannis, M. (eds.) *ECML PKDD 2011. LNCS*, vol. 6911, pp. 597–612. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-23780-5_48](https://doi.org/10.1007/978-3-642-23780-5_48)
19. He, X., Cai, D., Niyogi, P.: Laplacian score for feature selection. In: *Advances in Neural Information Processing Systems*, vol. 18 (2005)
20. Kang, U., Papadimitriou, S., Sun, J., Tong, H.: Centralities in large networks: algorithms and observations. In: *Proceedings of the 2011 SIAM International Conference on Data Mining. Society for Industrial and Applied Mathematics*, pp. 119–130 (2011)
21. Kuncheva, L.I.: A stability index for feature selection. In: *Proceedings of the 25th Conference on Proceedings of the 25th IASTED International Multi-Conference: Artificial Intelligence and Applications, AIAP 2007*, pp. 390–395. ACTA Press, Anaheim (2007)
22. Lehoucq, R.B., Sorensen, D.C., Yang, C.: *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, vol. 6. SIAM, Philadelphia (1998)
23. Lerman, K., Ghosh, R., Kang, J.H.: Centrality metric for dynamic networks. In: *Proceedings of the Eighth Workshop on Mining and Learning with Graphs, MLG 2010*, pp. 70–77. ACM, New York (2010)
24. Liu, H., Motoda, H. (eds.): *Computational Methods of Feature Selection*. CRC Press, Boca Raton (2007)

25. Meyer, C.D. (ed.): Matrix Analysis and Applied Linear Algebra. Society for Industrial and Applied Mathematics, Philadelphia (2000)
26. Obertino, S., Roffo, G., Granziera, C., Menegaz, G.: Infinite feature selection on shore-based biomarkers reveals connectivity modulation after stroke. In: 2016 International Workshop on Pattern Recognition in Neuroimaging (PRNI), pp. 1–4, June 2016
27. Peng, H., Long, F., Ding, C.: Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)* **27**(8), 1226–1238 (2005)
28. Pitts, F.R.: A graph theoretic approach to historical geography. *Prof. Geogr.* **17**(5), 15–20 (1965)
29. Rawat, A., Saha, S., Ghrera, S.P.: Time efficient ranking system on map reduce framework. In: 2015 Third International Conference on Image Information Processing (ICIIP), pp. 496–501 (2015)
30. Roffo, G., Melzi, S.: Online feature selection for visual tracking. In: International Conference the British Machine Vision Conference (BMVC), September 2016
31. Roffo, G., Melzi, S., Cristani, M.: Infinite feature selection. In: IEEE Conference International Conference on Computer Vision (ICCV) (2015)
32. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556 (2014)
33. Singh, D., Febbo, P.G., Ross, K., Jackson, D.G., Manola, J., Ladd, C., Tamayo, P., Renshaw, A.A., D’Amico, A.V., Richie, J.P., Lander, E.S., Loda, M., Kantoff, P.W., Golub, T.R., Sellers, W.R.: Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell* **1**(2), 203–209 (2002)
34. Wu, D.D., Deng, X., Li, Y.: Safety and emergency systems engineering mapreduce based betweenness approximation engineering in large scale graph. *Syst. Eng. Procedia* **5**, 162–167 (2012)
35. Zaffalon, M., Hutter, M.: Robust feature selection using distributions of mutual information. In: Conference International Conference on Uncertainty in Artificial Intelligence (UAI) (2002)
36. Zhang, Z., Hancock, E.R.: A graph-based approach to feature selection. In: Jiang, X., Ferrer, M., Torsello, A. (eds.) *GbRPR 2011*. LNCS, vol. 6658, pp. 205–214. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-20844-7_21](https://doi.org/10.1007/978-3-642-20844-7_21)

New Frontiers in Mining Complex Patterns
5th International Workshop, NFMCP 2016, Held in
Conjunction with ECML-PKDD 2016, Riva del Garda, Italy,
September 19, 2016, Revised Selected Papers
Appice, A.; Ceci, M.; Loglisci, C.; Masciari, E.; Ras, Z.W.
(Eds.)
2017, XIV, 263 p. 66 illus., Softcover
ISBN: 978-3-319-61460-1