

Utility metric for unsupervised feature selection

Amalia Villa^{Corresp., 1}, **Abhijith Mundanad Narayanan**¹, **Sabine Van Huffel**¹, **Alexander Bertrand**^{1,2}, **Carolina Varon**^{1,3,4}

¹ STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Department of Electrical Engineering (ESAT), KU Leuven, Leuven, Belgium

² Leuven.AI, KU Leuven Institute for AI, Leuven, Belgium

³ Circuits and Systems (CAS) group, Delft University of Technology, Delft, The Netherlands

⁴ e-Media Research Lab, Campus GroepT, KU Leuven, Leuven, Belgium

Corresponding Author: Amalia Villa

Email address: amalia.villagomez@kuleuven.be

Feature selection techniques are very useful approaches for dimensionality reduction in data analysis. They provide interpretable results by reducing the dimensions of the data to a subset of the original set of features. When the data lack annotations, unsupervised feature selectors are required for their analysis. Several algorithms for this aim exist in the literature, but despite their large applicability, they can be very inaccessible or cumbersome to use, mainly due to the need for tuning non-intuitive parameters and the high computational demands. In this work, a publicly available ready-to-use unsupervised feature selector is proposed, with comparable results to the state-of-the-art at a much lower computational cost. The suggested approach belongs to the methods known as spectral feature selectors. These methods generally consist of two stages: manifold learning and subset selection. In the first stage, the underlying structures in the high-dimensional data are extracted, while in the second stage a subset of the features is selected to replicate these structures. This paper suggests two contributions to this field, related to each of the stages involved. In the manifold learning stage, the effect of non-linearities in the data is explored, making use of a radial basis function (RBF) kernel, for which an alternative solution for the estimation of the kernel parameter is presented for cases with high-dimensional data. Additionally, the use of a backwards greedy approach based on the least-squares utility metric for the subset selection stage is proposed. The combination of these new ingredients results in the Utility metric for Unsupervised feature selection U2FS algorithm. The proposed U2FS algorithm succeeds in selecting the correct features in a simulation environment. In addition, the performance of the method on benchmark datasets is comparable to the state-of-the-art, while requiring less computational time. Moreover, unlike the state-of-the-art, U2FS does not require any tuning of parameters.

Utility metric for unsupervised feature selection

Amalia Villa¹, Abhijith Mundanad Narayanan¹, Sabine Van Huffel¹,
Alexander Bertrand^{1,2}, and Carolina Varon^{1,3,4}

¹STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics,
Department of Electrical Engineering (ESAT), KU Leuven, Leuven, Belgium

²Leuven.AI, KU Leuven Institute for AI, Leuven, Belgium

³Circuits and Systems (CAS) group, Delft University of Technology, Delft, The
Netherlands

⁴e-Media Research Lab, Campus GroepT, KU Leuven, Leuven, Belgium

Corresponding author:

Amalia Villa¹

Email address: amalia.villagomez@kuleuven.be

ABSTRACT

Feature selection techniques are very useful approaches for dimensionality reduction in data analysis. They provide interpretable results by reducing the dimensions of the data to a subset of the original set of features. When the data lack annotations, unsupervised feature selectors are required for their analysis. Several algorithms for this aim exist in the literature, but despite their large applicability, they can be very inaccessible or cumbersome to use, mainly due to the need for tuning non-intuitive parameters and the high computational demands.

In this work, a publicly available ready-to-use unsupervised feature selector is proposed, with comparable results to the state-of-the-art at a much lower computational cost. The suggested approach belongs to the methods known as spectral feature selectors. These methods generally consist of two stages: manifold learning and subset selection. In the first stage, the underlying structures in the high-dimensional data are extracted, while in the second stage a subset of the features is selected to replicate these structures.

This paper suggests two contributions to this field, related to each of the stages involved. In the manifold learning stage, the effect of non-linearities in the data is explored, making use of a radial basis function (RBF) kernel, for which an alternative solution for the estimation of the kernel parameter is presented for cases with high-dimensional data. Additionally, the use of a backwards greedy approach based on the least-squares utility metric for the subset selection stage is proposed.

The combination of these new ingredients results in the Utility metric for Unsupervised feature selection (U2FS) algorithm. The proposed U2FS algorithm succeeds in selecting the correct features in a simulation environment. In addition, the performance of the method on benchmark datasets is comparable to the state-of-the-art, while requiring less computational time. Moreover, unlike the state-of-the-art, U2FS does not require any tuning of parameters.

INTRODUCTION

Many applications of data science require the study of highly multi-dimensional data. A high number of dimensions implies a high computational cost as well as a large amount of memory required. Furthermore, this often leads to problems related to the curse of dimensionality (Verleysen and François, 2005) and thus, to irrelevant and redundant data for machine learning algorithms (Mairon, 2007). Therefore, it is crucial to perform dimensionality reduction before analyzing the data.

There are two types of dimensionality reduction techniques. So-called feature selection techniques directly select a subset of the original features. On the other hand, transformation techniques compute a new (smaller) set of features, each of which are derived from all features of the original set. Some examples of these are Principal Component Analysis (PCA) (Wold et al., 1987), Independent Component Analysis (ICA) (Jiang et al., 2006) or the Extended Sammon Projection (ESP) (Ahmad et al., 2019).

47 While these methods lead to a reduction in the number of dimensions, results are less interpretable, since
48 their direct relationship with the original set of features is lost.

49 In this work, the focus is on unsupervised feature selectors. Since these methods do not rely on the
50 availability of labels or annotations in the data, the information comes from the learning of the underlying
51 structure of the data. Despite this challenge, the generalization capabilities of these methods are typically
52 better than for supervised or semi-supervised methods (Guyon and Elisseeff, 2003). Within unsupervised
53 feature selectors, sparse learning based methods have gained attention in the last 20 years (Li et al., 2017).
54 These methods rely on graph theory and manifold learning to learn the underlying structures of the data
55 (Lunga et al., 2013), and they apply sparsity inducing techniques to perform subset selection. However, to
56 the best of our knowledge, none explores specifically the behavior of these methods with data presenting
57 non-linear relationships between the features (i.e., dimensions). While the graph definition step can make
58 use of kernels to tackle non-linearities, these can be heavily affected by the curse of dimensionality, since
59 they are often based on a distance metric (Aggarwal et al., 2001).

60 After the manifold learning stage, sparse regression is applied to score the relevance of the features
61 in the structures present in the graph. These formulations make use of sparsity-inducing regularization
62 techniques to provide the final subset of features selected, and thus, they are highly computationally
63 expensive. These methods are often referred to as structured sparsity-inducing feature selectors (SSFS),
64 or sparse learning based methods (Gui et al., 2016)(Li et al., 2017).

65 Despite the large amount of unsupervised SSFS algorithms described in the literature, these methods
66 are cumbersome to use for a novice user. This is not only due to the codes not being publicly available,
67 but also due to the algorithms requiring regularization parameters which are difficult to tune, in particular
68 in unsupervised settings.

69 In this work, an efficient unsupervised feature selector based on the utility metric (U2FS) is proposed.
70 U2FS is a ready-to-use, publicly available unsupervised sparsity-inducing feature selector designed to be
71 robust for data containing non-linearities. The code is available here: <https://github.com/avillago/u2fs>,
72 where all functions and example codes are published. The main contributions of this work are:

- 73 • The definition of a new method to automatically approximate the radial-basis function (RBF) kernel
74 parameter without the need for a user-defined tuning parameter. This method is used to tackle the
75 curse of dimensionality when embedding the data taking non-linearities into account.
- 76 • The suggestion of a backwards greedy approach for the stage of subset selection, based on the
77 utility metric for the least-squares problem. The utility metric was proposed in the framework of
78 supervised learning (Bertrand, 2018), and has been used for channel selection in applications such
79 as electroencephalography (EEG) (Narayanan and Bertrand, 2020), sensor networks (Szurley et al.,
80 2014), and microphone arrays (Szurley et al., 2012). Nevertheless, this is the first work in which
81 this type of approach is proposed for the sparsity-inducing stage of feature selection.
- 82 • Propose a non-parametric and efficient unsupervised SSFS algorithm. This work analyzes the
83 proposed method U2FS in terms of its complexity, and of its performance on simulated and
84 benchmark data. The goal is to reduce the computational cost while maintaining a comparable
85 performance with respect to the state-of-the-art. In order to prove this, U2FS is compared to three
86 related state-of-the-art algorithms in terms of accuracy of the features selected, and computational
87 complexity of the algorithm.

88 The rest of the paper is structured as follows. In Related Work, previous algorithms on SSFS are
89 summarized. In Methods, the proposed U2FS method is described: first the manifold learning stage,
90 together with the algorithm proposed for the selection of the kernel parameter; and further on, the utility
91 metric is discussed and adapted to feature selection. The experiments performed in simulations and
92 benchmark databases, as well as the results obtained are described in the Results and Discussion sections.
93 Finally, the last section provides some conclusions.

94 RELATED WORK

95 Sparsity-inducing feature selection methods have become widely used in unsupervised learning applica-
96 tions for high-dimensional data. This is due to two reasons. On the one hand, the use of manifold learning
97 guarantees the preservation of local structures present in the high-dimensional data. Additionally, its

98 combination with feature selection techniques not only reduces the dimensionality of the data, but also
99 guarantees interpretability.

100 Sparsity-inducing feature selectors learn the structures present in the data via connectivity graphs
101 obtained in the high-dimensional space (Yan et al., 2006). The combination of manifold learning and
102 regularization techniques to impose sparsity, allows to select a subset of features from the original dataset
103 that are able to describe these structures in a smaller dimensional space.

104 These algorithms make use of sparsity-inducing regularization approaches to stress those features that
105 are more relevant for data separation. The sparsity of these approaches is controlled by different statistical
106 norms ($l_{r,p}$ -norms), which contribute to the generalization capability of the methods, adapting them to
107 binary or multi-class problems (Gui et al., 2016). One drawback of these sparse regression techniques is
108 that generally, they rely on optimization methods, which are computationally expensive.

109 The Laplacian Score (He et al., 2006) was the first method to perform spectral feature selection in an
110 unsupervised way. Based on the Laplacian obtained from the spectral embedding of the data, it obtains
111 a score based on locality preservation. SPEC (Zhao and Liu, 2007) is a framework that contains this
112 previous approach, but it additionally allows for both supervised or unsupervised learning, including
113 other similarity metrics, as well as other ranking functions. These approaches evaluate each feature
114 independently, without considering feature interactions. These interactions are, however, taken into
115 account in Multi-Cluster Feature Selection (MCFS) (Cai et al., 2010), where a multi-cluster approach
116 is defined based on the eigendecomposition of a similarity matrix. The subset selection is performed
117 applying an l_1 -norm regularizer to approximate the eigenvectors obtained from the spectral embedding
118 of the data inducing sparsity. In UDFS (Yang et al., 2011) the l_1 -norm regularizer is substituted by a
119 $l_{2,1}$ -norm to apply sample and feature-wise constraints, and a discriminative analysis is added in the
120 graph description. In NDFS (Li et al., 2012), the use of the $l_{2,1}$ -norm is preserved, but a non-negative
121 constraint is added to the spectral clustering stage. Additionally, this algorithm performs feature selection
122 and spectral clustering simultaneously.

123 The aforementioned algorithms perform manifold learning and subset selection in a sequential way.
124 However, other methods tackle these simultaneously, in order to adaptively change the similarity metric
125 or the selection criteria regarding the error obtained between the original data and the new representation.
126 Examples of these algorithms are JELSR (Hou et al., 2013), SOGFS (Nie et al., 2019), (R)JGSC (Zhu
127 et al., 2016) and DSRMR (Tang et al., 2018), and all make use of an $l_{2,1}$ -norm. Most recently, the
128 SAMM-FS algorithm was proposed (Zhang et al., 2019), where a combination of similarity measures
129 is used to build the similarity graph, and the $l_{2,0}$ -norm is used for regression. This group of algorithms
130 are currently the ones achieving the best results, at the cost of using complex optimization techniques
131 to adaptively tune both stages of the feature selection process. While this can lead to good results, it
132 comes with a high computation cost, which might hamper the tuning process, or might simply not be
133 worthy for some applications. SAMM-FS and SOGFS are the ones that more specifically suggest new
134 approaches to perform the embedding stage, by optimally creating the graph (Nie et al., 2019) or deriving
135 it from a combination of different similarity metrics (Zhang et al., 2019). Again, both approaches require
136 computationally expensive optimization techniques to select a subset of features.

137 In summary, even if SSFS methods are getting more sophisticated and accurate, this results in
138 algorithms becoming more complex in terms of computational time, and in the ease of use. The use of
139 advanced numerical optimization techniques to improve results makes algorithms more complex, and
140 requires regularization parameters which are not easy to tune. In this work, the combination of a new
141 approach to estimate the graph connectivity based on the RBF kernel, together with the use of the utility
142 metric for subset selection, results in an efficient SSFS algorithm, which is easy to use and with lower
143 complexity than the state-of-the-art. This efficient implementation is competitive with state-of-the-art
144 methods in terms of performance, while using a simpler strategy, which is faster to compute and easier to
145 use.

146 METHODS

147 This section describes the proposed U2FS algorithm, which focuses on selecting the relevant features
148 in an unsupervised way, at a relatively small computational cost. The method is divided in three parts.
149 Firstly, the suggested manifold learning approach is explained, where an embedding based on binary
150 weighting and the RBF kernel are used. Then a method to select the kernel parameter of the RBF kernel

151 is proposed, specially designed for high-dimensional data. Once the manifold learning stage is explained,
152 the Utility metric is proposed as a new approach for subset selection.

153 **Manifold learning considering non-linearities**

154 Given is a data matrix $\mathbf{X} \in \mathbb{R}^{N \times d}$, with $\mathbf{X} = [\mathbf{x}_1; \mathbf{x}_2; \dots; \mathbf{x}_N]$, $\mathbf{x}_i = [x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(d)}]$, $i = 1, \dots, N$, N the
155 number of data points, and d the number of features (i.e., dimensions) in the data. The aim is to learn the
156 structure hidden in the d -dimensional data and approximate it with only a subset of the original features.
157 In this paper, this structure will be identified by means of clustering, where the dataset is assumed to be
158 characterized by c clusters.

159 In spectral clustering, the clustering structure of this data can be obtained by studying the eigenvectors
160 derived from a Laplacian built from the original data (Von Luxburg (2007), Biggs et al. (1993)). The
161 data is represented using a graph $G = (\mathcal{V}, \mathcal{E})$. \mathcal{V} is the set of vertices \mathbf{v}_i , $i = 1, \dots, N$ where $\mathbf{v}_i = \mathbf{x}_i$.
162 $\mathcal{E} = \{e_{ij}\}$ with $i = 1, \dots, N$ $j = 1, \dots, N$ is the set of edges between the vertices where $\{e_{ij}\}$ denotes the
163 edge between vertices v_i and v_j . The weight of these edges is determined by the entries $w_{ij} \geq 0$ of a
164 similarity matrix \mathbf{W} . We define the graph as undirected. Therefore, the similarity matrix \mathbf{W} , is symmetric
165 (since $w_{ij} = w_{ji}$, with the diagonal set to $w_{ii} = 0$).

166 Typically, \mathbf{W} is computed after coding the pairwise distances between all N data points. There are
167 several ways of doing this, such as calculating the k -nearest neighbours (KNN) for each point, or choosing
168 the ε -neighbors below a certain distance (Belkin and Niyogi, 2002).

169 In this paper, two similarity matrices are adopted inspired by the work in (Cai et al., 2010), namely a
170 binary one and one based on an RBF kernel. The binary weighting is based on KNN, being $w_{ij} = 1$ if and
171 only if vertex i is within the K closest points to vertex j . Being a non-parametric approach, the binary
172 embedding allows to simply characterize the connectivity of the data.

173 Additionally, the use of the RBF kernel is considered, which is well suited for non-linearities and
174 allows to characterize complex and sparse structures (Von Luxburg, 2007). The RBF kernel is defined as
175 $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2\sigma^2)$. The selection of the kernel parameter σ is a long-standing challenge
176 in machine learning. For instance, in Cai et al. (2010), σ^2 is defined as the mean of all the distances
177 between the data points. Alternatively, a rule of thumb, uses the sum of the standard deviations of the data
178 along each dimension (Varon et al., 2015). However, the estimation of this parameter is highly influenced
179 by the amount of features or dimensions in the data, making it less robust to noise and irrelevant features.
180 In the next section, a new and better informed method to approximate the kernel parameter is proposed.

181 The graph G , defined by the similarity matrix \mathbf{W} , can be partitioned into multiple disjoint sets. Given
182 the focus on multi-cluster data of our approach, the k -Way Normalized Cut ($NCut$) Relaxation is used,
183 as proposed in Ng et al. (2002). In order to obtain this partition, the degree matrix \mathbf{D} of \mathbf{W} must be
184 calculated. \mathbf{D} is a diagonal matrix for which each element on the diagonal is calculated as $D_{ii} = \sum_j W_{i,j}$.
185 The normalized Laplacian \mathbf{L} is then obtained as $\mathbf{L} = \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$, as suggested in Von Luxburg (2007).
186 The vectors \mathbf{y} embedding the data in \mathbf{L} can be extracted from the eigenvalue problem (Chung and Graham,
187 1997):

$$\mathbf{L}\mathbf{y} = \lambda\mathbf{y} \quad (1)$$

188 Given the use of a normalized Laplacian for the data embedding, the vectors \mathbf{y} must be adjusted using
189 the degree matrix \mathbf{D} :

$$\alpha = \mathbf{D}^{1/2}\mathbf{y}, \quad (2)$$

190 which means that α is the solution of the generalized eigenvalue problem of the pair \mathbf{W} and \mathbf{D} . These
191 eigenvectors α are a new representation of the data, that gathers the most relevant information about the
192 structures appearing in the high-dimensional space. The c eigenvectors, corresponding to the c highest
193 eigenvalues (after excluding the largest one), can be used to characterize the data in a lower dimensional
194 space (Ng et al., 2002). Thus, the matrix $\mathbf{E} = [\alpha_1, \alpha_2, \dots, \alpha_c]$ containing column-wise the c selected
195 eigenvectors, will be the low-dimensional representation of the data to be mimicked using a subset of the
196 original features, as suggested in Cai et al. (2010).

197 **Kernel parameter approximation for high-dimensional data**

198 One of the most used similarity functions is the RBF kernel, which allows to explore non-linearities in
 199 the data. Nevertheless, the kernel parameter σ^2 must be selected correctly, to avoid overfitting or the
 200 allocation of all data points to the same cluster. This work proposes a new approach to approximate this
 201 kernel parameter, which will be denoted by $\hat{\sigma}^2$ when derived from our method. This method takes into
 202 account the curse of dimensionality and the potential irrelevant features or dimensions in the data.

203 As a rule of thumb, σ^2 is approximated as the sum of the standard deviation of the data along each
 204 dimension (Varon et al., 2015). This approximation grows with the number of features (i.e. dimensions)
 205 of the data, and thus, it is not able to capture its underlying structures in high-dimensional spaces.
 206 Nevertheless, this σ^2 is commonly used as an initialization value, around which a search is performed,
 207 considering some objective function (Alzate and Suykens, 2008; Varon et al., 2015).

208 The MCFS algorithm skips the search around an initialization of the σ^2 value by substituting the sum
 209 of the standard deviations by the mean of these (Cai et al., 2010). By doing so, the value of σ^2 does not
 210 overly grow. This estimation of σ^2 suggested in Cai et al. (2010) will be referred to as σ_0^2 . A drawback of
 211 this approximation in high-dimensional spaces is that it treats all dimensions as equally relevant for the
 212 final estimation of σ_0^2 , regardless of the amount of information that they actually contain.

213 The aim of the proposed approach is to provide a functional value of σ^2 that does not require
 214 any additional search, while being robust to high-dimensional data. Therefore, this work proposes an
 215 approximation technique based on two factors: the distances between the points, and the number of
 216 features or dimensions in the data.

217 The most commonly used distance metric is the euclidean distance. However, it is very sensitive to
 218 high-dimensional data, deriving unsubstantial distances when a high number of features is involved in the
 219 calculation (Aggarwal et al., 2001). In this work, the use of the Manhattan or taxicab distance (Reynolds,
 220 1980) is proposed, given its robustness when applied to high-dimensional data (Aggarwal et al., 2001).
 221 For each feature l , the Manhattan distance δ_l is calculated as:

$$\delta_l = \frac{1}{N} \sum_{i,j=1}^N |x_{il} - x_{jl}| \quad (3)$$

222 Additionally, in order to reduce the impact of irrelevant or redundant features, a system of weights
 223 is added to the approximation of $\hat{\sigma}^2$. The goal is to only take into account the distances associated to
 224 features that contain relevant information about the structure of the data. To calculate these weights, the
 225 probability density function (PDF) of each feature is compared with a Gaussian distribution. Higher
 226 weights are assigned to the features with less Gaussian behavior, i.e. those the PDF of which differs the
 227 most from a Gaussian distribution. By doing so, these will influence more the final $\hat{\sigma}^2$ value, since they
 228 allow a better separation of the structures present in the data.

229 Figure 1 shows a graphical representation of this estimation. The dataset in the example has 3
 230 dimensions or features: f_1 , f_2 and f_3 . f_1 and f_2 contain the main clustering information, as it can be
 231 observed in Figure 1a, while f_3 is a noisy version of f_1 , derived as $f_3 = f_1 + 1.5n$, where n is drawn from
 232 a normal distribution $\mathcal{N}(0, 1)$. Figures 1b, 1c and 1d show in a continuous black line the PDFs derived
 233 from the data, and in a grey dash line their fitted Gaussian, in dimensions f_1 , f_2 and f_3 respectively. This
 234 fitted Gaussian was derived using the Curve Fitting toolbox of MatlabTM. As it can be observed, the
 235 matching of a Gaussian with an irrelevant feature is almost perfect, while those features that contain more
 236 information, like f_1 and f_2 , deviate much more from a normal distribution.

237 Making use of these differences, an error, denoted ϕ_l , for each feature l , where $l = 1, \dots, d$, is
 238 calculated as:

$$\phi_l = \frac{1}{H} \sum_{i=1}^H (p_i - g_i)^2, \quad (4)$$

239 where H is the number of bins in which the range of the data is divided to estimate the PDF (p), and g
 240 is the fitted Gaussian. The number of bins in this work is set to 100 for standardization purposes. Equation
 241 (4) corresponds to the mean-squared error (MSE) between the PDF of the data over feature l and its fitted
 242 Gaussian. From these ϕ_l , the final weights b_l are calculated as:

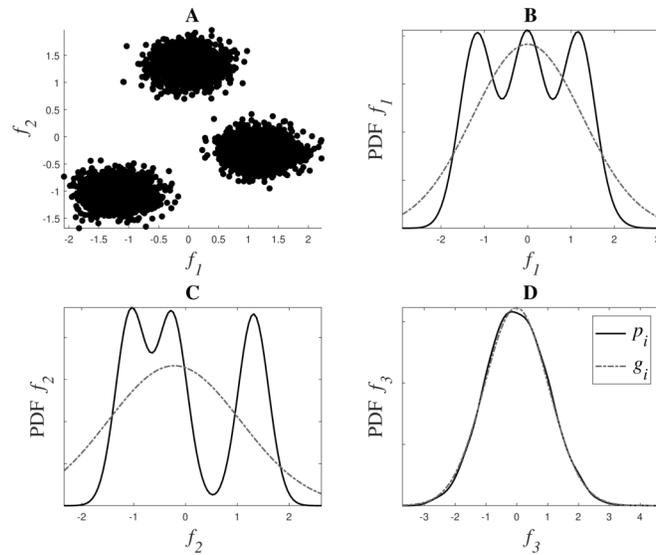


Figure 1. Weight system for relevance estimation. In Figure 1A, f_1 and f_2 can be seen. 1B, 1C and 1D show in black the PDFs p_i of f_1 , f_2 and f_3 respectively, and in grey dotted line their fitted Gaussian g_i .

$$b_l = \frac{\phi_l}{\sum_{l=1}^d \phi_l} \quad (5)$$

243 Therefore, combining (3) and (5), the proposed approximation, denoted $\hat{\sigma}^2$, is derived as:

$$\hat{\sigma}^2 = \sum_{l=1}^d b_l \delta_l, \quad (6)$$

244 which gathers the distances present in the most relevant features, giving less importance to the
 245 dimensions that do not contribute to describe the structure of the data. The complete algorithm to calculate
 246 $\hat{\sigma}^2$ is described in Algorithm 1.

Algorithm 1 Kernel parameter approximation for high-dimensional data.

Input: Data $\mathbf{X} \in \mathbb{R}^{N \times d}$.

Output: Sigma parameter $\hat{\sigma}^2$

- 1: Calculate the Manhattan distances between the datapoints using Equation (3): vector of distances per feature δ_l .
 - 2: Obtain the weights for each of the features using Equations (4) and (5): weights b_l .
 - 3: Calculate $\hat{\sigma}^2$ using Equation (6).
-

247 Utility metric for feature subset selection

248 In the manifold learning stage, a new representation \mathbf{E} of the data based on the eigenvectors was built,
 249 which described the main structures present in the original high-dimensional data. The goal is to select
 250 a subset of the features which best approximates the data in this new representation. In the literature,
 251 this feature selection problem is formulated using a graph-based loss function and a sparse regularizer
 252 of the coefficients is used to select a subset of features, as explained in Zhu et al. (2016). The main
 253 idea of these approaches is to regress the data to its low dimensional embedding along with some sparse
 254 regularization. The use of such regularization techniques reduces overfitting and achieves dimensionality
 255 reduction. This regression is generally formulated as a least squares (LS) problem, and in many of these

256 cases, the metric that is used for feature selection is the magnitude of their corresponding weights in
 257 the least squares solution (Cai et al., 2010; Gui et al., 2016). However, the optimized weights do not
 258 necessarily reflect the importance of the corresponding feature as it is scaling dependent and it does
 259 not properly take interactions across features into account (Bertrand, 2018). Instead, the importance
 260 of a feature can be quantified using the increase in least-squared error (LSE) if that feature was to be
 261 removed and the weights were re-optimized. This increase in LSE, called the ‘utility’ of the feature can
 262 be efficiently computed (Bertrand, 2018) and can be used as an informative metric for a greedy backwards
 263 feature selection procedure (Bertrand, 2018; Narayanan and Bertrand, 2020; Szurley et al., 2014), as an
 264 alternative for (group-)LASSO based techniques. Under some technical conditions, a greedy selection
 265 based on this utility metric can even be shown to lead to the optimal subset (Couvreur and Bresler, 2000).

After representing the dataset using the matrix $\mathbf{E} \in \mathbb{R}^{N \times c}$ containing the c eigenvectors, the following LS optimization problem finds the weights \mathbf{p} that best approximate the data \mathbf{X} in the c -dimensional representation in \mathbf{E} :

$$J = \min_{\mathbf{p}} \frac{1}{N} \|\mathbf{X}\mathbf{p} - \mathbf{E}\|_F^2 \quad (7)$$

266 where J is the cost or the LSE and $\|\cdot\|_F$ denotes the Frobenius norm.

If \mathbf{X} is a full rank matrix and if $N > d$, the LS solution $\hat{\mathbf{p}}$ of (7) is

$$\hat{\mathbf{p}} = \mathbf{R}_{\mathbf{X}\mathbf{X}}^{-1} \mathbf{R}_{\mathbf{X}\mathbf{E}}, \quad (8)$$

267 with $\mathbf{R}_{\mathbf{X}\mathbf{X}} = \frac{1}{N} \mathbf{X}^T \mathbf{X}$ and $\mathbf{R}_{\mathbf{X}\mathbf{E}} = \frac{1}{N} \mathbf{X}^T \mathbf{E}$.

268 The goal of this feature selection method is to select the subset of $s (< d)$ features that best represents
 269 \mathbf{E} . This feature selection problem can be reduced to the selection of the best $s (< d)$ columns of \mathbf{X} which
 270 minimize (7). However, this is inherently a combinatorial problem and is computationally unfeasible to
 271 solve. Nevertheless, several greedy and approximative methods have been proposed (Gui et al., 2016; Nie
 272 et al., 2019; Narayanan and Bertrand, 2020). In the current work, the use of the utility metric for subset
 273 selection is proposed to select these best s columns.

274 The utility of a feature l of \mathbf{X} , in an LS problem like (7), is defined as the increase in the LSE J when
 275 the column corresponding to the l -th feature in \mathbf{X} is removed from the problem and the new optimal
 276 weight matrix, $\hat{\mathbf{p}}_{-l}$, is re-computed similar to (8). Consider the new LSE after the removal of feature l
 277 and the re-computation of the weight matrix $\hat{\mathbf{p}}_{-l}$ to be J_{-l} , defined as:

$$J_{-l} = \frac{1}{N} \|\mathbf{X}_{-l} \hat{\mathbf{p}}_{-l} - \mathbf{E}\|_F^2 \quad (9)$$

278 where \mathbf{X}_{-l} denotes the matrix \mathbf{X} with the column corresponding to l -th feature removed. Then
 279 according to the definition, the utility of feature l , U_l is:

$$U_l = J_{-l} - J \quad (10)$$

280 A straightforward computation of U_l would be computationally heavy due to the fact that the compu-
 281 tation of $\hat{\mathbf{p}}_l$ requires a matrix inversion of $\mathbf{X}_{-l} \mathbf{X}_{-l}^T$, which has to be repeated for each feature l .

282 However, it can be shown that the utility of the l -th feature of \mathbf{X} in (10) can be computed efficiently
 283 without the explicit recomputation of $\hat{\mathbf{p}}_{-l}$ by using the following expression (Bertrand, 2018):

$$U_l = \frac{1}{q_l} \|\bar{\mathbf{p}}_l\|_2, \quad (11)$$

284 where q_l is the l -th diagonal element of $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{-1}$ and p_l is the l -th row in $\hat{\mathbf{p}}$, corresponding to the l -th
 285 feature. The mathematical proof of (11) can be found in Bertrand (2018). Note that $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{-1}$ is already known
 286 from the computation of $\hat{\mathbf{p}}$ such that no additional matrix inversion is required.

287 However, since the data matrix \mathbf{X} can contain redundant features or features that are linear combi-
 288 nations of each other in its columns, it cannot be guaranteed that the matrix \mathbf{X} in (7) is full-rank. In
 289 this case, the removal of a redundant column from \mathbf{X} will not lead to an increase in the LS cost of (7).

290 Moreover, $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{-1}$, used to find the solution of (7) in (8), will not exist in this case since the matrix \mathbf{X} is rank
 291 deficient. A similar problem appears if $N < d$, which can happen in case of very high-dimensional data.
 292 To overcome this problem, the definition of utility generalized to a minimum l_2 -norm selection (Bertrand,
 293 2018) is used in this work. This approach eliminates the feature yielding the smallest increase in the
 294 l_2 -norm of the weight matrix when the column corresponding to that feature were to be removed and the
 295 weight matrix would be re-optimized. Moreover, minimizing the l_2 -norm of the weights further reduces
 296 the risk of overfitting.

297 This generalization is achieved by first adding an l_2 -norm penalty β to the cost function that is
 298 minimized in (7):

$$J = \min_{\mathbf{p}} \frac{1}{2} \|\mathbf{X}\mathbf{p} - \mathbf{E}\|_F^2 + \beta \|\mathbf{p}\|_2^2 \quad (12)$$

299 where $0 < \beta \leq \mu$ with μ equal to the smallest non-zero eigenvalue of $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ in order to ensure that the
 300 bias added due to the penalty term in (12) is negligible. The minimizer of (12) is:

$$\hat{\mathbf{p}} = \mathbf{R}_{\mathbf{X}\mathbf{X}\beta}^{-1} \mathbf{R}_{\mathbf{X}\mathbf{E}} = (\mathbf{R}_{\mathbf{X}\mathbf{X}} + \beta \mathbf{I})^{-1} \mathbf{R}_{\mathbf{X}\mathbf{E}} \quad (13)$$

301 It is noted that (13) reduces to $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{\dagger} \mathbf{R}_{\mathbf{X}\mathbf{E}}$ when $\beta \rightarrow 0$, where $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{\dagger}$ denotes the Moore-Penrose pseudo-
 302 inverse. This solution corresponds to the minimum norm solution of (7) when \mathbf{X} contains linearly
 303 dependent columns or rows. The utility U_l of the l -th column in \mathbf{X} based on (12) is (Bertrand, 2018):

$$\begin{aligned} U_l &= (\|\mathbf{X}_{-l}\hat{\mathbf{p}}_{-l} - \mathbf{E}\|_2^2 - \|\mathbf{X}\hat{\mathbf{p}} - \mathbf{E}\|_2^2) \\ &\quad + \beta (\|\hat{\mathbf{p}}_{-l}\|_2^2 - \|\hat{\mathbf{p}}\|_2^2) \\ &= (J_{-l} - J) + \beta (\|\hat{\mathbf{p}}_{-l}\|_2^2 - \|\hat{\mathbf{p}}\|_2^2) \end{aligned} \quad (14)$$

304 Note that if column l in \mathbf{X} is linearly independent from the other columns, (14) closely approximates
 305 to the original utility definition in (10) as the first term dominates over the second. However, if column l
 306 is linearly dependent, the first term vanishes and the second term will dominate. In this case, the utility
 307 quantifies the increase in l_2 -norm after removing the l -th feature.

308 To select the best s features of \mathbf{X} , a greedy selection based on the iterative elimination of the features
 309 with the least utility is carried out. After the elimination of each feature, a re-estimation of the weights $\hat{\mathbf{p}}$
 310 is carried out and the process of elimination is repeated, until s features remain.

311 Note that the value of β depends on the smallest non-zero eigenvalue of $\mathbf{R}_{\mathbf{X}\mathbf{X}}$. Since $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ has to be
 312 recomputed every time when a feature is removed, also its eigenvalues change along the way. In practice,
 313 the value of β is selected only once and fixed for the remainder of the algorithm, as smaller than the
 314 smallest non-zero eigenvalue of $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ before any of the features are eliminated (Narayanan and Bertrand,
 315 2020). This value of β will be smaller than all the non-zero eigenvalues of any principal submatrix of
 316 $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ using the Cauchy's interlace theorem (Hwang, 2004).

317 The summary of the utility subset selection is described in Algorithm 2. Algorithm 3 outlines the
 318 complete U2FS algorithm proposed in this paper.

Algorithm 2 Utility metric algorithm for subset selection.

Input: Data \mathbf{X} , Eigenvectors \mathbf{E} , Number of features s to select

Output: s features selected

- 1: Calculate $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ and $\mathbf{R}_{\mathbf{X}\mathbf{E}}$ as described in Equation (8).
 - 2: Calculate β as the smallest non-zero eigenvalue of $\mathbf{R}_{\mathbf{X}\mathbf{X}}$
 - 3: **while** Number of features remaining is $> s$ **do**
 - 4: Compute $\mathbf{R}_{\mathbf{X}\mathbf{X}\beta}^{-1}$ and $\hat{\mathbf{p}}$ as described in (13).
 - 5: Calculate the utility of the remaining features using (11)
 - 6: Remove the feature f_l with the lowest utility.
 - 7: Update $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ and $\mathbf{R}_{\mathbf{X}\mathbf{E}}$ by removing the rows and columns related to that feature f_l .
 - 8: **end while**
-

Algorithm 3 Unsupervised feature selector based on the utility metric (U2FS).

Input: Data \mathbf{X} , Number of clusters c , Number of features s to select

Output: s features selected

- 1: Construct the similarity graph \mathbf{W} as described in Section selecting one of the weightings:
 - Binary
 - RBF kernel, using σ_0^2
 - RBF kernel, using $\hat{\sigma}^2$ based on Algorithm 1
 - 2: Calculate the normalized Laplacian \mathbf{L} and the eigenvectors α derived from Equation (2). Keep the c eigenvectors corresponding to the highest eigenvalues, excluding the first one.
 - 3: Apply the backward greedy utility algorithm 2.
 - 4: Return the s features remaining from the backward greedy utility approach.
-

319 As it has been stated before, one of the most remarkable aspects of the U2FS algorithm is the use of
 320 a greedy technique to solve the subset selection problem. The use of this type of method reduces the
 321 computational cost of the algorithm. This can be confirmed analyzing the computational complexity of
 322 U2FS, where the most demanding steps are the eigendecomposition of the Laplacian matrix (step 2 of
 323 Algorithm 3), which has a cost of $O(N^3)$ (Tsironis et al., 2013), and the subset selection stage in step 3 of
 324 Algorithm 3. Contrary to the state-of-the-art, the complexity of U2FS being a greedy method depends on
 325 the number of features to select. The most computationally expensive step of the subset selection in U2FS
 326 is the calculation of the matrix $\mathbf{R}_{\mathbf{X}\mathbf{X}}^{-1}$, which has a computational cost of $O(d^3)$. In addition, this matrix
 327 needs to be updated $d - s$ times. This update can be done efficiently using a recursive updating equation
 328 from Bertrand (2018) with a cost of $O(t^2)$, with t the number of features remaining in the dataset, i.e.
 329 $t = d - s$. Since $t < d$, the cost for performing $d - s$ iterations will be $O((d - s)d^2)$, which depends on
 330 the number of features s to be selected. Note that the cost of computing the least squares solution $\hat{\mathbf{p}}_{-l}$ for
 331 each l in (14) is eliminated using the efficient equation (11), bringing down the cost for computing the
 332 utility from $O(t^4)$ to $O(t)$ in each iteration. This vanishes with respect to the $O(d^3)$ term (remember that
 333 $t < d$). Therefore, the total asymptotic complexity of U2FS is $O(N^3 + d^3)$.

334 RESULTS

335 The aim of the following experiments is to evaluate the U2FS algorithm based on multiple criteria. With
 336 the focus on the new estimation of the embedding proposed, the proposed RBF kernel approach using
 337 the estimated $\hat{\sigma}^2$ is compared to the σ_0^2 parameter proposed in Cai et al. (2010), and to the binary KNN
 338 graph commonly used in Gui et al. (2016). On the other hand, the utility metric for subset selection is
 339 compared to other sparsity-inducing techniques, based on l_p -norm regularizations. In these experiments,
 340 this is evaluated using the l_1 -norm. The outline of the different combinations considered in this work
 341 summarized in Table 1. The last method, RBF $_{\hat{\sigma}^2}$ + Utility, would be the one referred to as U2FS,
 342 combining the novelties suggested in this work.

Table 1. Methods compared in the experiments

	Similarity measure	Subset selection
KNN$_{Bin}$ + l_1 - norm	KNN + binary weighting	l_1 -norm
RBF$_{\sigma_0^2}$ + l_1 - norm	RBF kernel, σ_0^2	l_1 -norm
KNN$_{Bin}$ + Utility	KNN + binary weighting	Utility metric
RBF$_{\sigma_0^2}$ + Utility	RBF kernel, σ_0^2	Utility metric
RBF$_{\hat{\sigma}^2}$ + Utility	RBF kernel, $\hat{\sigma}^2$	Utility metric

343 These novelties are evaluated in two different scenarios, namely a simulation study, and in the
 344 application of the methods on benchmark datasets. In particular for the latter, the methods are not only
 345 evaluated in terms of accuracy, but also regarding computational cost. Additionally, U2FS is compared
 346 with 3 representative state-of-the-art algorithms from the general family of unsupervised sparsity-inducing
 347 feature selection algorithms:

- 348 • **MCFS**(Cai et al., 2010)¹. This algorithm served as inspiration to create U2FS, and therefore,
 349 it is added to the set of comparison algorithms as baseline reference. MCFS performs spectral
 350 embedding and l_1 -norm regularization sequentially, and which served as inspiration to create U2FS.
- 351 • **NDFS**(Li et al., 2012)², which performs nonnegative spectral analysis with $l_{2,1}$ -norm regularization.
 352 This algorithm is added to the experiments since it is an improvement of MCFS, while being the first
 353 algorithm simultaneously adapting both stages of manifold learning and subset selection. Therefore,
 354 NDFS represents the transition to these adaptive optimization-based feature selection algorithms.
- 355 • **RJGSC**(Zhu et al., 2016) optimally derives the embedding of the data by adapting the results with
 356 $l_{2,1}$ -norm regularization. This algorithm is taken as a reference for the large class of adaptive
 357 sparsity-inducing feature selection algorithms, which are much more complex than U2FS, since
 358 they apply optimization to recursively adapt the embedding and feature selection stages of the
 359 methods. RJGSC was already compared to several feature selectors in Zhu et al. (2016), and
 360 therefore, it is taken here as upper-bound threshold in performance.

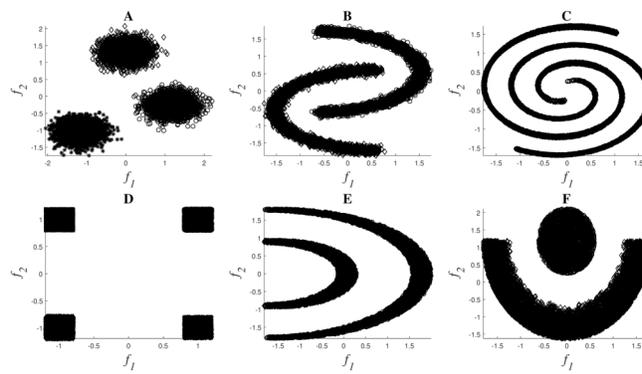


Figure 2. Toy examples used for simulations: Clouds (A), Moons (B), Spirals (C), Corners (D), Half-Kernel (E), Crescent Moon (F).

361 Simulations

362 A set of nonlinear toy examples typically used in clustering problems are proposed to test the different
 363 feature selection methods. In these experiments, the goal was to verify the correct selection of the original
 364 set of features. Figure 2 shows the toy examples considered³, which are described by features f_1 and f_2 ,
 365 and the final description of the datasets can be seen in Table 2.

366 All these problems are balanced, except for the last dataset Cres-Moon, for which the data is divided
 367 25% to 75% between the two clusters. Five extra features in addition to the original f_1 and f_2 were added
 368 to each of the datasets in order to include redundant or irrelevant information:

- 369 • f'_1 and f'_2 : random values extracted from two Pearson distributions characterized by the same
 370 higher-order statistics as f_1 and f_2 respectively.
- 371 • f'_3 and f'_4 : Original f_1 and f_2 contaminated with Gaussian noise ($v\mathcal{N}(0, 1)$), with $v = 1.5$.
- 372 • f'_5 : Constant feature of value 0.

373 The first step in the preprocessing of the features was to standardize the data using z-score to reduce
 374 the impact of differences in scaling and noise. In order to confirm the robustness of the feature selection
 375 techniques, the methods were applied using 10-fold cross-validation on the standardized data. For each
 376 fold a training set was selected using m -medoids, setting m to 2000 and using the centers of the clusters
 377 found as training samples. By doing so, the generalization ability of the methods can be guaranteed

¹<http://www.cad.zju.edu.cn/home/dengcai/Data/MCFS.html>

²<http://www.cs.cmu.edu/~yiyang/Publications.html>

³The codes used to generate these datasets are available in <https://github.com/avillago/u2fs>

Table 2. Description of the toy example datasets.

	# samples	# classes
Clouds	9000	3
Moons	10000	2
Spirals	10000	2
Corners	10000	4
Half-Kernel	10000	2
Crescent-Moon	10000	2

378 (Varon et al., 2015). On each of the 10 training sets, the features were selected applying the 5 methods
 379 mentioned in Table 1. For each of the methods, the number of clusters c was introduced as the number of
 380 classes presented in Table 2. Since these experiments aim to evaluate the correct selection of the features,
 381 and the original features f_1 and f_2 are known, the number of features s to be selected was set to 2.

382 Regarding the parameter settings within the embedding methods, the binary was obtained setting k in
 383 the k NN approach to 5. For the RBF kernel embedding, σ_0^2 was set to the mean of the standard deviation
 384 along each dimension, as done in Cai et al. (2010). When using $\hat{\sigma}^2$, its value was obtained by applying
 385 the method described in Algorithm 1.

386 In terms of subset selection approaches, the method based on the l_1 - norm automatically sets the
 387 value of the regularization parameter required for the LARS implementation, as described in (Deng Cai,
 388 Chiyuan Zhang, 2020). For the utility metric, β was automatically set to the smallest non-zero eigenvalue
 389 of the matrix $\mathbf{R}_{\mathbf{X}\mathbf{X}}$ as described in Algorithm 2.

390 The performance of the algorithm is evaluated comparing the original set of features f_1 and f_2 to those
 391 selected by the algorithm. In these experiments, the evaluation of the selection results is binary: either the
 392 feature set selected is correct or not, regardless of the additional features f'_i , for $i = 1, 2, \dots, 5$, selected.

393 In Table 3 the most common results obtained in the 10 folds are shown. The utility-based approaches
 394 always obtained the same results for all 10 folds of the experiments. On the contrary, the l_1 - norm
 395 methods provided different results for different folds of the experiment. For these cases, Table 3 shows
 396 the most common feature pair for each experiment, occurring at least 3 times.

Table 3. Results feature selection for toy examples

Method	Utility metric			l_1 - norm	
	KNN _{Bin}	RBF _{σ_0^2}	RBF _{$\hat{\sigma}^2$}	KNN _{Bin}	RBF _{σ_0^2}
Clouds	f_1, f_2	f'_1, f'_4	f_1, f_2	f'_1, f'_2	f'_1, f'_2
Moons	f_1, f_2	f'_3, f'_4	f_1, f_2	f'_1, f'_3	f'_1, f'_3
Spirals	f_1, f_2	f_1, f_2	f_1, f_2	f_2, f'_2	f_2, f'_2
Corners	f_1, f_2	f'_1, f'_2	f_1, f_2	f_2, f'_2	f_2, f'_2
Half-Kernel	f_1, f_2	f_2, f'_3	f_1, f_2	f_1, f'_3	f_1, f'_3
Cres-Moon	f_1, f_2	f_1, f'_4	f_1, f_2	f_2, f'_1	f_2, f'_2

397 As shown in Table 3, the methods that always obtain the adequate set of features are based on utility,
 398 both with the binary weighting and with the RBF kernel and the suggested $\hat{\sigma}^2$. Since these results were
 399 obtained for the 10 folds, they confirm both the robustness and the consistency of the U2FS algorithm.

400 Benchmark datasets

401 Additionally, the proposed methods were evaluated using 6 well-known benchmark databases. The
 402 databases considered represent image (USPS, ORL, COIL20), audio (ISOLET) and text data (PCMAC,
 403 BASEHOCK)⁴, proposing examples with more samples than features, and vice versa. The description of
 404 these databases is detailed in Table 4. All these datasets are balanced, except USPS.

405 In these datasets, the relevant features are unknown. Therefore, the common practice in the literature
 406 to evaluate feature selectors consists of applying the algorithms, taking from 10 to 80% of the original

⁴All datasets downloaded from <http://featureselection.asu.edu/datasets.php>

Table 4. Description of the benchmark databases

	Data Type	Samples	Features	Classes
USPS	Images	9298	256	10
Isolet	Audio	1560	617	26
ORL	Images	400	1024	40
COIL20	Images	1440	1024	20
PCMAC	Text	1943	3289	2
BASEHOCK	Text	1993	4862	2

407 set of features, and evaluating the accuracy of a classifier when trained and evaluated with the selected
 408 feature set (Zhu et al., 2016). The classifier used for this aim in other papers is k -Nearest Neighbors
 409 (KNN), setting the number of neighbors to 5.

410 These accuracy results are computed using 10-fold cross-validation to confirm the generalization
 411 capabilities of the algorithm. By setting m to 90% of the number of samples available in each benchmark
 412 dataset, m -medoids is used to select the m centroids of the clusters and use them as training set. Feature
 413 selection and the training of the KNN classifier are performed in these 9 folds of the standardized data,
 414 and the accuracy of the KNN is evaluated in the remaining 10% for testing. Exclusively for USPS, given
 415 the size of the dataset, 2000 samples were used for training and the remaining data was used for testing.
 416 These 2000 samples were also selected using m -medoids. Since PCMAC and BASEHOCK consist of
 417 binary data, these datasets were not standardized.

418 The parameters required for the binary and RBF embeddings, as well as β for the utility algorithm,
 419 are automatically set as detailed in section .

420 Figure 3 shows the median accuracy obtained for each of the 5 methods. The shadows along the
 421 lines correspond to the 25 and 75 percentile of the 10 folds. As a reference, the accuracy of the classifier
 422 without using feature selection is shown in black for each of the datasets. Additionally, Figure 4 shows the
 423 computation time for both the utility metric and the l_1 - norm applied on a binary weighting embedding.
 424 In this manner, the subset selection techniques can be evaluated regardless of the code efficiency of the
 425 embedding stage. Similarly to Figure 3, the computation time plots show in bold the median running time
 426 for each of the subset selection techniques, and the 25 and 75 percentiles around it obtained from the
 427 10-fold cross-validation.

428 The difference in the trends of the l_1 - norm and utility in terms of computation time is due to their
 429 formulation. Feature selection based on l_1 - norm regularization, solved using the LARS algorithm in
 430 this case, requires the same computation time regardless of the number of features aimed to select. All
 431 features are evaluated together, and later on, an MCFS score obtained from the regression problem is
 432 assigned to them (Cai et al., 2010). The features with the higher scores are the ones selected. On the other
 433 hand, since the utility metric is applied in a backward greedy trend, the computation times change for
 434 different number of features selected. The lower the number of features selected compared to the original
 435 set, the higher the computation time. This is aligned with the computational complexity of the algorithm,
 436 described in Section . In spite of this, it can be seen that even the highest computation time for utility is
 437 lower than the time taken using l_1 - norm regularization. The experiments were performed with 2x Intel
 438 Xeon E5-2640 @ 2.5 GHz processors and 64GB of working memory.

439 Finally, the experiments in benchmark databases are extended to compare U2FS to other key algorithms
 440 in the state-of-the-art. As it was mentioned at the beginning of this section, the selected algorithms are
 441 MCFS, NDFS, and RJGSC, which represent, respectively, the precursor of U2FS, an improved version of
 442 MCFS, and an example from the class of adaptive algorithms which recursively optimize the objective
 443 function proposed. NDFS and RJGSC require the tuning of their regularization parameters, for which
 444 the indications in their corresponding articles were followed. For NDFS, the value of γ was set to 10^8 ,
 445 and α and β were selected from the values $\{10^{-6}, 10^{-4}, \dots, 10^6\}$ applying grid search. The matrix F was
 446 initialized with the results of spectral clustering using all the features. For RJGSC, the results described
 447 in Zhu et al. (2016) for the BASEHOCK and PCMAC datasets are taken as a reference. In MCFS, the
 448 embedding is done using KNN and binary weighting, and the l_1 - norm is used for subset selection.
 449 U2FS, on the other hand, results from the combination of the RBF kernel with $\hat{\sigma}^2$ and the utility metric.
 450 Table 5 summarizes the results by showing the KNN accuracy (ACC) for 10% of the features used, and
 451 the maximum ACC achieved among the percentages of features considered, for the BASEHOCK and

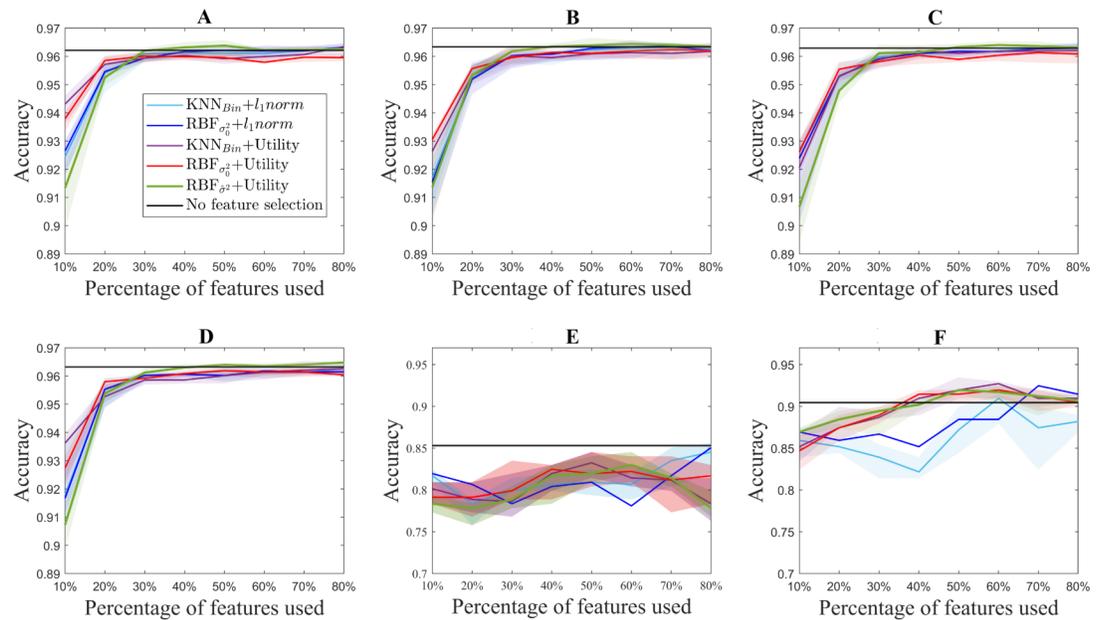


Figure 3. Accuracy results for the benchmark databases, for selecting from 10 to 80% of the original number of features. The thick lines represent the median accuracy of the 10-fold cross-validation, and the shadows, the 25 and 75 percentile. USPS (Figure 3A), Isolet (B), ORL (C), COIL20 (D), AR10 (E), PIE10 (F).

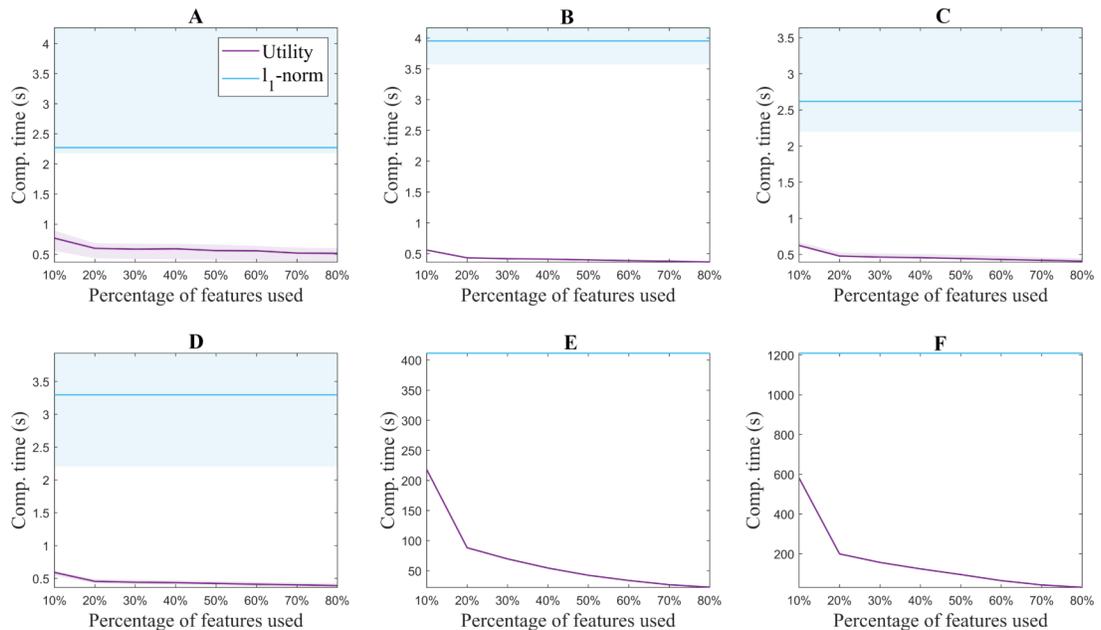


Figure 4. Computation time for extracting from 10 to 80% of the original number of features for each of the benchmark databases. USPS (Figure 4A), Isolet (B), ORL (C), COIL20 (D), AR10 (E), PIE10 (F).

452 PCMAC datasets.

Table 5. Comparison of classification accuracy (ACC) with the state-of-the-art for PCMAC and BASEHOCK datasets.

Dataset	Method	ACC at 10% features	% features at Max ACC	Max ACC
PCMAC	U2FS	0.785	60%	0.83
	MCFS	0.67	20%	0.697
	NDFS	0.73	40%	0.83
	RJGSC	0.805	60%	0.83
BASEHOCK	U2FS	0.87	50%	0.925
	MCFS	0.815	80%	0.84
	NDFS	0.76	20%	0.794
	RJGSC	0.902	80%	0.917

453

454 DISCUSSION

455 The results obtained in the experiments suggest that the proposed U2FS algorithm obtains comparable
 456 results to the state-of-the-art in all the applications suggested, taking less computational time. Nevertheless,
 457 the performance of the utility metric for feature selection varies for the different experiments presented
 458 and requires a detailed analysis.

459 From Table 3, in Section , it can be concluded that the utility metric is able to select the correct
 460 features in an artificially contaminated dataset. Both the binary embedding and the RBF kernel with $\hat{\sigma}^2$
 461 select the original set of features for the 10 folds of the experiment. The stability in the results also applies
 462 for the RBF embedding with σ_0^2 , which always selected the same feature pair for all 10 folds even though
 463 they are only correct for the spirals problem.

464 Therefore, considering the stability of the results, it can be concluded that the proposed approach is
 465 more robust in the selection of results than that based on the $l_1 - norm$.

466 On the other hand, when considering the suitability of the features selected, two observations can be
 467 made. First of all, it can be seen that the lack of consistency in the $l_1 - norm$ approaches discards the
 468 selection of the correct set of features. Moreover, the wrong results obtained with both $l_1 - norm$ and
 469 utility methods for the RBF embedding using σ_0^2 reveal the drawback of applying this approximation
 470 of σ_0^2 in presence of redundant or irrelevant features. Since this value is calculated as the mean of the
 471 standard deviation of all the dimensions in the data, this measure can be strongly affected by irrelevant
 472 data, that could be very noisy and enlarge this sigma, leading to the allocation of all the samples to a
 473 mega-cluster.

474 While the use of the proposed approximation for $\hat{\sigma}^2$ achieves better results than σ_0^2 , these are
 475 comparable to the ones obtained with the KNN binary embedding when using the utility metric. The use
 476 of KNN to build graphs is a well-known practice, very robust for dense clusters, as it is the case in these
 477 examples. The definition of a specific field where each of the embeddings would be superior is beyond
 478 the scope of this paper. However, the excellence of both methods when combined with the proposed
 479 subset selection method only confirms the robustness of the utility metric, irrespective of the embedding
 480 considered.

481 For standardization purposes, the performance of the method was evaluated in benchmark databases.
 482 As it can be observed, in terms of the accuracy obtained for each experiment, U2FS achieves comparable
 483 results to the $l_1 - norm$ methods for most of the datasets considered, despite its condition of greedy
 484 method.

485 In spite of this, some differences in performance can be observed in the different datasets. The different
 486 ranking of the methods, as well as the accuracy obtained for each of the databases can be explained taking
 487 into account the type of data under study and the ratio between samples and dimensions.

488 With regard to the type of data represented by each test, it can be observed that for the ISOLET dataset,
 489 containing sound information, two groups of results are distinguishable. The group of the utility metric

490 results outperforms those derived from the l_1 - norm, which only reach comparable results for 60% of the
491 features selected. These two groups of results are caused by the subset selection method applied, and not
492 for the embedding, among which the differences are not remarkable.

493 In a similar way, for the case of the image datasets USPS, ORL and COIL20, the results derived
494 from utility are slightly better than those coming from the l_1 - norm. In these datasets, similarly to the
495 performance observed in ISOLET, accuracy increases with the number of features selected.

496 Regarding the differences between the proposed embeddings, it can be observed that the results
497 obtained are comparable for all of them. Nonetheless, Figure 3 shows that there is a slight improvement
498 in the aforementioned datasets for the RBF kernel with $\hat{\sigma}^2$, but the results are still comparable to those
499 obtained with other embeddings. Moreover, this similarity in the binary and RBF results holds for the
500 l_1 - norm methods, for which the accuracy results almost overlap in Figure 3. This can be explained by
501 the relation between the features considered. Since for these datasets the samples correspond to pixels,
502 and the features to the color codes, a simple neighboring method such as the binary weighting is able to
503 code the connectivity of pixels of similar colors.

504 The text datasets, PCMAC and BASEHOCK, are the ones that show bigger differences between the
505 results obtained with utility and those obtained with the l_1 - norm. This can be explained by the amount
506 of zeros present in the data, with which the utility metric is able to cope slightly better. The sparsity of the
507 data leads to more error in the l_1 - norm results, since more features end up having the same MCFS score,
508 and among those, the order for selection comes at random. The results obtained with the utility metric
509 are more stable, in particular for the BASEHOCK dataset. For this dataset, U2FS even outperforms the
510 results without feature selection if at least 40% of the features are kept.

511 In all the datasets proposed, the results obtained with the l_1 - norm show greater variability, i.e. larger
512 percentiles. This is aligned with the results obtained in the simulations. The results for the l_1 - norm
513 are not necessarily reproducible in different runs, since the algorithm is too sensitive to the training set
514 selected. The variability of the utility methods is greater for the approaches based on the RBF kernel.
515 This is due to the selection of the σ^2 parameter, which also depends on the training set. The tuning of this
516 parameter is still very sensitive to high-dimensional and large-scale data, posing a continuous challenge
517 for the machine learning community (Yin and Yin, 2016; Tharwat et al., 2017).

518 Despite it being a greedy method, the utility metric proves to be applicable to feature selection
519 approaches and to strongly outperform the l_1 - norm in terms of computational time, without significant
520 reduction in accuracy. U2FS proves to be effective both in cases with more samples than features and
521 vice versa. The reduction in computation time is clear, for all the benchmark databases described, and is
522 particularly attractive for high-dimensional datasets. Altogether, our feature selection approach U2FS,
523 based on the utility metric, and with the binary or the RBF kernel with $\hat{\sigma}^2$ is recommended due to its fast
524 performance and its interpretability.

525 Additionally, the performance of U2FS is comparable to the state-of-the-art, as shown in Table 5. In
526 this table, the performance of U2FS (RBF kernel and $\hat{\sigma}^2$, with the utility metric) is compared to that of
527 MCFS, NDFS and RJGSC. For MCFS, it can be seen that, as expected, U2FS appears as an improvement
528 of this algorithm, achieving better results for both datasets. For NDFS, the results are slightly worse than
529 for U2FS, most probably due to problems in the tuning of regularization parameters. Given the consistent
530 good results for different datasets of RJGSC when compared against the state-of-the-art, and its condition
531 of simultaneously adapting the spectral embedding and subset selection stages, this algorithm is taken
532 as example of the most complex SSFS algorithms (SAMM-FS, SOGFS or DSRMR). These algorithms
533 perform manifold learning and feature selection simultaneously, iteratively adapting both steps to achieve
534 optimal results.

535 It is clear that in terms of accuracy, both for 10% of the features and for the maximal value of achieved,
536 U2FS obtains similar results to RJGSC, while at the same time having a much smaller computational
537 complexity. Furthermore, while RJGSC requires the manual tuning of extra parameters, similarly to other
538 algorithms in the state-of-the-art, U2FS tunes its parameters automatically. Hence, the application of the
539 method is straightforward for the users. The stages of higher complexity in U2FS, previously defined
540 as $O(N^3 + d^3)$, are shared by most of the algorithms in the state-of-the-art. However, on top of these
541 eigendecompositions and matrix inversions, the algorithms in the literature require a number of iterations
542 in the optimization process that U2FS avoids. Additionally, U2FS is the only algorithm for which the
543 computation time scales linearly with the amount of features selected.

544 The current state-of-the-art of unsupervised spectral feature selectors applies the stages of manifold

545 learning and subset selection simultaneously, which can lead to optimal results. In a field that gets more
546 and more complex and goes far from applicability, U2FS is presented as a quick solution for a sequential
547 implementation of both stages of SSFS algorithms, yet achieving comparable results to the state-of-the-art.
548 Being a greedy method, the utility metric cannot be applied simultaneously to the manifold learning and
549 subset selection stages. However, other sequential algorithms from the state-of-the-art could consider
550 the use of utility for subset selection, instead of the current sparsity-inducing techniques. One of the
551 most direct applications could be the substitution of group-LASSO for group-utility, in order to perform
552 selections of groups of features as proposed by Bertrand (2018). This can be of interest in cases where the
553 relations between features are known, such as in channel selection (Narayanan and Bertrand, 2020) or in
554 multi-modal applications (Zhao et al., 2015).

555 CONCLUSION

556 This work presents a new method for unsupervised feature selection based on manifold learning and
557 sparse regression. The main contribution of this paper is the formulation of the utility metric in the field
558 of spectral feature selection, substituting other sparse regression methods that require more computational
559 resources. This method, being a backward greedy approach, has been proven to obtain comparable
560 results to the state-of-the-art methods with analogous embedding approaches, yet at considerably reduced
561 computational load. The method shows consistently good results in different applications, from images
562 to text and sound data; and it is broadly applicable to problems of any size: using more features than
563 samples or vice versa.

564 Furthermore, aiming to show the applicability of U2FS to data presenting non-linearities, the proposed
565 approach has been evaluated in simulated data, considering both a binary and an RBF kernel embedding.
566 Given the sensitivity of the RBF kernel to high-dimensional spaces, a new approximation of the RBF
567 kernel parameter was proposed, which does not require further tuning around the value obtained. The
568 proposed approximation outperforms the rule-of-thumb widely used in the literature in most of the
569 scenarios presented. Nevertheless, in terms of feature selection, the utility metric is robust against the
570 embedding.

571 U2FS is proposed as a non-parametric efficient algorithm, which does not require any manual tuning
572 or special knowledge from the user. Its simplicity, robustness and accuracy open a new path for structure
573 sparsity-inducing feature selection methods, which can benefit from this quick and efficient technique.

574 REFERENCES

- 575 Aggarwal, C. C., Hinneburg, A., and Keim, D. A. (2001). On the surprising behavior of distance metrics
576 in high dimensional space. In *International conference on database theory*, pages 420–434. Springer.
- 577 Ahmad, M., Khan, A. M., Mazzara, M., Distefano, S., Ali, A., and Tufail, A. (2019). Extended sammon
578 projection and wavelet kernel extreme learning machine for gait-based legitimate user identification. In
579 *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, pages 1216–1219.
- 580 Alzate, C. and Suykens, J. A. (2008). Multiway spectral clustering with out-of-sample extensions through
581 weighted kernel pca. *IEEE transactions on pattern analysis and machine intelligence*, 32(2):335–347.
- 582 Belkin, M. and Niyogi, P. (2002). Laplacian eigenmaps and spectral techniques for embedding and
583 clustering. In *Advances in neural information processing systems*, pages 585–591.
- 584 Bertrand, A. (2018). Utility metrics for assessment and subset selection of input variables for linear
585 estimation [tips & tricks]. *IEEE Signal Processing Magazine*, 35(6):93–99.
- 586 Biggs, N., Biggs, N. L., and Norman, B. (1993). *Algebraic graph theory*, volume 67. Cambridge university
587 press.
- 588 Cai, D., Zhang, C., and He, X. (2010). Unsupervised feature selection for multi-cluster data. In
589 *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data
590 mining*, pages 333–342.
- 591 Chung, F. R. and Graham, F. C. (1997). *Spectral graph theory*. Number 92. American Mathematical Soc.
- 592 Coudreau, C. and Bresler, Y. (2000). On the optimality of the backward greedy algorithm for the subset
593 selection problem. *SIAM Journal on Matrix Analysis and Applications*, 21(3):797–808.
- 594 Deng Cai, Chiyuan Zhang, X. H. (2020). Supervised/Unsupervised/Semi-supervised Feature Selection
595 for Multi-Cluster/Class Data.

- 596 Gui, J., Sun, Z., Ji, S., Tao, D., and Tan, T. (2016). Feature selection based on structured sparsity: A
597 comprehensive study. *IEEE transactions on neural networks and learning systems*, 28(7):1490–1507.
- 598 Guyon, I. and Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of machine*
599 *learning research*, 3(Mar):1157–1182.
- 600 He, X., Cai, D., and Niyogi, P. (2006). Laplacian score for feature selection. In *Advances in neural*
601 *information processing systems*, pages 507–514.
- 602 Hou, C., Nie, F., Li, X., Yi, D., and Wu, Y. (2013). Joint embedding learning and sparse regression: A
603 framework for unsupervised feature selection. *IEEE Transactions on Cybernetics*, 44(6):793–804.
- 604 Hwang, S.-G. (2004). Cauchy's interlace theorem for eigenvalues of hermitian matrices. *The American*
605 *Mathematical Monthly*, 111(2):157–159.
- 606 Jiang, X., Zhang, L., Zhao, Q., and Albayrak, S. (2006). Ecg arrhythmias recognition system based
607 on independent component analysis feature extraction. In *TENCON 2006-2006 IEEE Region 10*
608 *Conference*, pages 1–4. IEEE.
- 609 Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R. P., Tang, J., and Liu, H. (2017). Feature selection:
610 A data perspective. *ACM Computing Surveys (CSUR)*, 50(6):1–45.
- 611 Li, Z., Yang, Y., Liu, J., Zhou, X., and Lu, H. (2012). Unsupervised feature selection using nonnegative
612 spectral analysis. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.
- 613 Lunga, D., Prasad, S., Crawford, M. M., and Ersoy, O. (2013). Manifold-learning-based feature extraction
614 for classification of hyperspectral data: A review of advances in manifold learning. *IEEE Signal*
615 *Processing Magazine*, 31(1):55–66.
- 616 Maindonald, J. (2007). Pattern recognition and machine learning. *Journal of Statistical Software*, 17(b05).
- 617 Narayanan, A. M. and Bertrand, A. (2020). Analysis of miniaturization effects and channel selection
618 strategies for eeg sensor networks with application to auditory attention detection. *IEEE Transactions*
619 *on Biomedical Engineering*, 67(1):234–244.
- 620 Ng, A. Y., Jordan, M. I., and Weiss, Y. (2002). On spectral clustering: Analysis and an algorithm. In
621 *Advances in neural information processing systems*, pages 849–856.
- 622 Nie, F., Zhu, W., and Li, X. (2019). Structured graph optimization for unsupervised feature selection.
623 *IEEE Transactions on Knowledge and Data Engineering*.
- 624 Reynolds, B. E. (1980). Taxicab geometry. *Pi Mu Epsilon Journal*, 7(2):77–88.
- 625 Szurley, J., Bertrand, A., and Moonen, M. (2012). Efficient computation of microphone utility in a
626 wireless acoustic sensor network with multi-channel wiener filter based noise reduction. In *2012 IEEE*
627 *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2657–2660.
628 IEEE.
- 629 Szurley, J., Bertrand, A., Ruckebusch, P., Moerman, I., and Moonen, M. (2014). Greedy distributed node
630 selection for node-specific signal estimation in wireless sensor networks. *Signal Processing*, 94:57–73.
- 631 Tang, C., Liu, X., Li, M., Wang, P., Chen, J., Wang, L., and Li, W. (2018). Robust unsupervised
632 feature selection via dual self-representation and manifold regularization. *Knowledge-Based Systems*,
633 145:109–120.
- 634 Tharwat, A., Hassanien, A. E., and Elnaghi, B. E. (2017). A ba-based algorithm for parameter optimization
635 of support vector machine. *Pattern Recognition Letters*, 93:13–22.
- 636 Tsironis, S., Sozio, M., Vazirgiannis, M., and Polte, L. (2013). Accurate spectral clustering for
637 community detection in mapreduce. In *Advances in Neural Information Processing Systems (NIPS)*
638 *Workshops*. Citeseer.
- 639 Varon, C., Alzate, C., and Suykens, J. A. (2015). Noise level estimation for model selection in kernel pca
640 denoising. *IEEE transactions on neural networks and learning systems*, 26(11):2650–2663.
- 641 Verleysen, M. and François, D. (2005). The curse of dimensionality in data mining and time series
642 prediction. In *International Work-Conference on Artificial Neural Networks*, pages 758–770. Springer.
- 643 Von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416.
- 644 Wold, S., Esbensen, K., and Geladi, P. (1987). Principal component analysis. *Chemometrics and intelligent*
645 *laboratory systems*, 2(1-3):37–52.
- 646 Yan, S., Xu, D., Zhang, B., Zhang, H.-J., Yang, Q., and Lin, S. (2006). Graph embedding and extensions:
647 A general framework for dimensionality reduction. *IEEE transactions on pattern analysis and machine*
648 *intelligence*, 29(1):40–51.
- 649 Yang, Y., Shen, H. T., Ma, Z., Huang, Z., and Zhou, X. (2011). L_{2, 1}-norm regularized discriminative
650 feature selection for unsupervised. In *Twenty-Second International Joint Conference on Artificial*

651 *Intelligence.*

652 Yin, S. and Yin, J. (2016). Tuning kernel parameters for svm based on expected square distance ratio.
653 *Information Sciences*, 370:92–102.

654 Zhang, R., Nie, F., Wang, Y., and Li, X. (2019). Unsupervised feature selection via adaptive multimeasure
655 fusion. *IEEE transactions on neural networks and learning systems*, 30(9):2886–2892.

656 Zhao, L., Hu, Q., and Wang, W. (2015). Heterogeneous feature selection with multi-modal deep neural
657 networks and sparse group lasso. *IEEE Transactions on Multimedia*, 17(11):1936–1948.

658 Zhao, Z. and Liu, H. (2007). Spectral feature selection for supervised and unsupervised learning. In
659 *Proceedings of the 24th international conference on Machine learning*, pages 1151–1157.

660 Zhu, X., Li, X., Zhang, S., Ju, C., and Wu, X. (2016). Robust joint graph sparse coding for unsupervised
661 spectral feature selection. *IEEE transactions on neural networks and learning systems*, 28(6):1263–
662 1275.