

This is a repository copy of *Identifying the most informative features using a structurally interacting elastic net*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/132582/>

Version: Accepted Version

Article:

Cui, Lixin, Bai, Lu, Zhang, Zhihong et al. (2 more authors) (2018) Identifying the most informative features using a structurally interacting elastic net. *Neurocomputing*. pp. 1-14. ISSN 0925-2312

<https://doi.org/10.1016/j.neucom.2018.06.081>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Identifying The Most Informative Features Using A Structurally Interacting Elastic Net

Lixin Cui^{1*}, Lu Bai^{1*}, Zhihong Zhang², Yue Wang², Edwin R. Hancock³

¹ *Central University of Finance and Economics, Beijing, China.*

² *Xiamen University, Xiamen, Fujian, China.*

³ *University of York, York, UK.*

Abstract

Feature selection can efficiently identify the most informative features with respect to the target feature used in training. However, state-of-the-art vector-based methods are unable to encapsulate the relationships between feature samples into the feature selection process, thus leading to significant information loss. To address this problem, we propose a new graph-based structurally interacting elastic net method for feature selection. Specifically, we commence by constructing feature graphs that can incorporate pairwise relationship between samples. With the feature graphs to hand, we propose a new information theoretic criterion to measure the joint relevance of different pairwise feature combinations with respect to the target feature graph representation. This measure is used to obtain a structural interaction matrix where the elements represent the proposed information theoretic measure between feature pairs. We then formulate a new optimization model through the combination of the structural interaction matrix and an elastic net regression model for the feature subset selection problem. This allows us to a) preserve the information of the original vectorial space, b) remedy the information loss of the original feature space caused by using graph representation, and c) promote a sparse solution and also encourage correlated features to be selected. Because the proposed optimization problem is non-convex, we develop an efficient alternating direction multiplier method (ADMM) to locate the optimal solutions. Extensive experiments on various datasets demonstrate the effectiveness of the proposed

*These authors contribute equally to this work and are co-first authors. Correspondence author: Lu Bai, email: bailucs@cufe.edu.cn.

method.

Keywords: Feature Selection; Graph; Interacting Elastic Net; Sparse; ADMM

1. Introduction

There has recently been a rapid growth in both the size and dimension of the data encountered in many real world applications of pattern recognition including image processing, bioinformatics, and financial analysis. Finding useful information and building effective prediction models from such data presents new challenges for machine learning and pattern recognition [1]. One way to overcome this problem is to develop efficient spectral methods including stochastic neighbour embedding [2], elastic embedding methods [3] and feature selection [4] methods to reduce the dimensionality of the data.

Feature selection aims to identify an optimal subset of the most informative features by removing irrelevant and redundant features [4]. One of the main advantages is that feature selection can improve the predictive accuracy and enhance comprehensibility of learning tasks [5]. Unlike feature extraction, feature selection can maintain the properties of the original features and has better interpretability. This is very important for understanding which features are most informative with respect to the target feature used in training. For instance, in Peer-to-Peer (P2P) lending analysis, it is important to understand which features of the P2P lending platforms (e.g., operation time, registered capital, and management team) affect the investors' decisions [6]. For medical diagnosis, it is crucial to know which characteristics of the patients (e.g., age, gender, and weight) affect the occurrence of a certain disease [7].

Because of these advantages, many efficient feature selection methods have been developed [5] [8]. Existing feature selection algorithms can be broadly categorized as filter and wrapper methods depending on whether the learning algorithm is used in the feature subset selection process [9]. Filter methods utilize the intrinsic properties of the data to build quantitative evaluation criteria [10]. By contrast, wrapper methods [11] evaluate the selected feature subsets based on the performance measures of the classifier including accuracy and precision. Although wrapper methods often perform

better than filter methods, they require significantly higher computational costs. In addition, in the presence of redundant features, wrappers tend to locate suboptimal subsets, and the characteristics of the selected subset are inevitably biased depending on the choice of the classifier [9]. Therefore, for high-dimensional data, filter methods are often preferred [12].

To construct effective filter methods, a good evaluation criterion is necessary. To date, many evaluation criteria have been used, including correlation [13], consistency [14], Fisher score [15], and mutual information (MI) [16]. For instance, MI measures the mutual dependence of two variables [16] and has been shown to have similar or better performance than other more sophisticated methods [17]. Due to its excellent performance, MI has received considerable attention for developing various information theoretic feature selection methods. Representative examples include 1) the Mutual Information-Based Feature Selection (MIFS) [18], 2) the MIFS method under the assumption of a uniform distribution of input variables (MIFS-U) [19], 3) the Maximum-Relevance Minimum-Redundancy criterion (MRMR) [20], and 4) the Normalized Feature-Feature Mutual Information method (NMIFS) [21]. Although the performance of MI-based feature selection methods have been demonstrated in many applications, they suffer from two widely acknowledged limitations. First, they require the number of selected features to be predetermined. Second, they adopt greedy search methods to identify the most informative feature subsets [22]. To overcome these shortcomings, Liu et al. [23] have proposed the adaptive MI-based feature selection method that automatically determines the number of most informative features, by maximizing the average pairwise informativeness. Zhang and Hancock [24] have developed a hypergraph based information theoretic feature selection method that can automatically determine the size of the most informative feature subset through dominant hypergraph clustering [25].

However, none of the aforementioned information theoretic feature selection methods can incorporate pairwise relationship between samples of each feature dimension. More specifically, assume a dataset with N features denoted as $\mathcal{X} = \{\mathbf{f}_1, \dots, \mathbf{f}_i, \dots, \mathbf{f}_N\}$, and each feature \mathbf{f}_i has M samples as $\mathbf{f}_i = (f_{i1}, \dots, f_{ia}, \dots, f_{ib}, \dots, f_{iM})^T$. Traditional information theoretic feature selection methods represent each feature \mathbf{f}_i as a vec-

tor, and thus ignore the relationship between pairwise samples f_{ia} and f_{ib} in \mathbf{f}_i . This
60 deficiency restricts the precision of the information theoretic measure between pairs
of features. To address this drawback, Cui et al. [26] have recently developed a new
feature selection method using graph-based features. Specifically, they transform each
feature vector into a graph structure that encapsulates pairwise relationship between
samples. The most relevant vectorial features are located by selecting the graph-based
65 features that are most similar to the graph-based target feature, in terms of the Jensen-
Shannon divergence measure between the graphs. To adaptively determine the most
relevant feature subset, Cui et al. [27] have further developed a new information theo-
retic feature selection method which a) encapsulates the relationship between sample
pairs for each feature dimension and b) automatically identifies the subset containing
70 the most informative and least redundant features by solving a quadratic programming
problem.

However, the aforementioned graph-based feature selection methods may lead to
significant information loss concerning the relationships between samples from the o-
riginal vector space. To illustrate this point, assume that two pairs of samples from
75 the same feature dimension \mathbf{f}_i are denoted as $\{f_{i1}, f_{i3}\}$ and $\{f_{i3}, f_{i5}\}$, respectively.
Following Cui et al. [27], we transform the feature vector \mathbf{f}_i into a graph-based repre-
sentation \mathbf{G}_i , which is a complete weighted graph. Each vertex v_a of \mathbf{G}_i represents
a corresponding sample f_{ia} in \mathbf{f}_i and each weighted edge (v_a, v_b) represents the rela-
tionship between sample pair f_{ia} and f_{ib} . If the Euclidean distances of the two pairs,
80 i.e., $\{f_{i1}, f_{i3}\}$ and $\{f_{i3}, f_{i5}\}$ are the same, the weights associated with the two pairs
of samples are also the same in the feature graph \mathbf{G}_i . However, these two pairs of
samples are located differently in the original vector space. This means that the graph-
based feature representation may lead to information loss. One exception is that the
vertex label is the associated sample value of the original features, i.e., the vertex label
85 is continuous. However, in this case, we need to measure the affinity between a pair of
graphs associated with continuous vertex labels and this results in significantly higher
computational complexity [28].

To summarize the above, it is fair to say that it still remains a challenge to develop
an effective information theoretic feature selection methods that can both encapsulate

90 pairwise relationship between samples of each feature dimension and avoid information loss from the original vector space.

On the other hand, sparse feature selection methods have received increasing attention [29]. By formulating feature selection as a regression model with an ordinary least square (OLS) term and a specifically designed sparsity inducing regularizer, the regression model can be efficiently represented by a linear combination of a set of the most active variables. The cardinality of the set of the selected variables is significantly smaller than the entire number of variables [30]. In other words, the regression model retains information concerning the original feature space and also allows us to adaptively select the most informative feature subset. Because of these advantages, many efficient regularization techniques including Lasso [31], Elastic Net [32], and Group Lasso [33] have been extensively studied for high-dimensional data feature selection. For instance, Zheng and Liu [7] have developed a Lasso operator to identify the most informative features for cancer classification, where Lasso enforces automatic feature selection by forcing at least some features to zero. Panagakis et al. [34] have developed a new similarity measure based on the matrix Elastic Net regularization to efficiently deal with highly correlated audio feature vectors. Marafino et al. [35] have proposed an efficient sparse feature selection method for biomedical text classification using the Elastic Net. More recently, Zhang et al. [30] have devised a new regularization term in the Lasso regression model to impose high order interactions between covariates and responses. The high-order relations among covariates are represented by a feature hypergraph and then used as a regularizer on the covariate coefficients to automatically select the most relevant features.

Although sparsity is desirable for designing effective feature selection algorithms, it is worth noting that most of the existing sparse feature selection methods seldom consider pairwise relationship between samples from each feature dimension. Intuitively, such structural information is important for improving the efficiency of feature selection methods. In addition, as opposed to the Elastic Net, the use of Lasso proves to be problematic when at least some features are highly correlated. In this case, Lasso selects one feature at random. As a result, given n training samples, Lasso can only select at most n features.

Motivated by the above discussion, we aim to overcome the limitations of existing feature selection methods by developing a novel structurally interacting elastic net feature selection method. The proposed method not only considers the structural relationships between feature samples but also remedies the information loss caused by the graph-based representation of features. In addition, we also explore how to ensure sparsity and promote a grouping effect among selected features via elastic net regularization.

1.1. Related Work

Feature selection has been widely studied in machine learning and pattern analysis. The topic of feature selection has been reviewed in a number of recent papers [36], [37] and [38]. In this section, we briefly state-of-the-art MI-based and sparse feature selection methods, which are related to our proposed method.

MI is often considered as an evaluation criterion to measure the relevance between features and the target labels due to its effectiveness at quantifying how much information is shared by two random variables. Because of this, MI has been extensively used for developing information theoretic feature selection methods. In the earlier reported work, Battiti [18] introduced a first order incremental search algorithm based on MI known as the MIFS criterion,

$$J_{MIFS} = I(f_i; C) - \beta \sum_{f_s \in S} I(f_s; f_i). \quad (1)$$

For a given set of existing selected features S , at each step MIFS locates the candidate feature f_i which maximizes the relevance to the class $I(f_i; C)$, instead of calculating the joint MI between the selected features and the class label C . The proportional term $\beta I(f_s; f_i)$ measures the overlap information between the candidate feature and existing features, and is used to regulate the feature selection process. The parameter β may significantly influence the features selected and needs to be carefully controlled. It is worth noting that because MIFS only considers features that have maximum MI with the output classes, it might treat features that have rich information about the output class as redundant, leading to a suboptimal subset. To overcome this drawback,

Kwak and Choi [19] improved MIFS by developing MIFS-U under the assumption of a uniform distribution for selected features. This uniform criterion is defined as

$$J_{MIFS-U} = I(f_i; C) - \beta \sum_{f_s \in S} \frac{I(f_s; C)}{H(f_s)} I(f_i; f_s), \quad (2)$$

where $H(f_s) = -\sum_{f_s \in S} P(f_s) \log P(f_s)$ is the entropy associated with the probability distribution for f_s . Instead of calculating $I(C; f_i|S)$ directly, only $I(S; f_i)$ and $I(C; f_i)$ are computed. The conditional MI (denoted as $I(C; f_i|S)$) between the class label C and the candidate feature f_i for a given feature subset S can be approximated as $I(C; f_i|S) = I(C; f_i) - I(S; f_i) - I(S; f_i|C)$.

Although MIFS-U gives better estimation than MIFS, the model parameters need to be carefully controlled to avoid bad results. To overcome this problem, Peng et al. [20] proposed a parameter-free method, referred to as MRMR, which is defined as

$$J_{MRMR} = I(f_i; C) - \frac{1}{|S|} \sum_{f_s \in S} I(f_i; f_s), \quad (3)$$

where $|S|$ is the cardinality of the selected feature set S . MRMR uses the average of the redundancy term to eliminate the difficulty of parameter selection of β with MIFS and MIFS-U methods. However, as a first-order incremental method, that sequentially selects one feature after another based on the evaluation function, MRMR presents similar limitations as MIFS and MIFS-U in the presence of many irrelevant or redundant features. This is because the conditional MI $I(C; f_i|S)$ between the target class C and the candidate feature f_i for a given subset of features S is ignored. To deal with this problem, Estevez et al. [21] developed the normalized NMIFS method to achieve a balance between the relevance and the redundancy term, defined as

$$J_{NMIFS} = I(f_i; C) - \frac{1}{|S|} \sum_{f_s \in S} \hat{I}(f_i; f_s), \quad (4)$$

where $\hat{I}(f_i; f_s) = \frac{I(f_i; f_s)}{\inf(H(f_s), H(f_i))}$ is the normalized MI.

On the other hand, sparse feature selection methods have recently attracted much attention. Typically we have a set of training data $\{(x_i, y_i), i = 1, \dots, N\}$, which is used to estimate the regression coefficients β . Each $x_i = (f_1^i, f_2^i, \dots, f_d^i)^T \in R^{d \times 1}$ is a predictive vector of feature measurements for the i th sample. To fit the linear regression

model, the most popular ordinary least square (OLS) method is adopted. OLS selects the coefficients $\beta = (\beta_1, \dots, \beta_d)^T$ by minimizing the residual sum of squares denoted as

$$\begin{aligned} \min_{\beta \in \mathbb{R}^d} \sum_{i=1}^N \|y_i - \sum_{j=1}^d \beta_j f_j^i\|_2^2 &= \min_{\beta \in \mathbb{R}^d} \|\mathbf{y}^T - \beta^t \mathbf{X}\|_2^2 \\ \text{s.t. } \sum_{i=1}^N \|\beta\|_0 &= k, \end{aligned} \quad (5)$$

where $\mathbf{y} \in \mathbb{R}^{N \times 1}$ is the label (response) vector, $\mathbf{X} \in \mathbb{R}^{d \times N}$ is the training dataset, k is a predetermined number of selected features. The minimisation of Eq.(5) has been proved to be a NP-hard optimization problem and is very difficult to be solved. In practice, we can relax the constraint equation by imposing a positive regularization parameter λ and adding it to the objective function, that is

$$\min_{\beta \in \mathbb{R}^d} \|\mathbf{y}^T - \beta^t \mathbf{X}\|_2^2 + \lambda \|\beta\|_0. \quad (6)$$

Unfortunately solving Eq.(6) is still challenging. Therefore, an alternative formulation using l_1 -norm regularization instead of l_0 -norm has been proposed for practical purposes. This corresponds to the regularized counterpart of the Lasso (least absolute shrinkage and selection operator) problem in statistical learning [31]. Lasso imposes a l_1 constraint on the regression coefficients, so that some of the regression coefficients in the regression model will shrink to zeros. Thus it can automatically select a set of the informative variables. Correspondingly, the feature selection problem with Lasso penalty is defined as

$$\min_{\beta \in \mathbb{R}^d} \|\mathbf{y}^T - \beta^t \mathbf{X}\|_2^2 + \lambda \|\beta\|_1, \quad (7)$$

145 where $\|\beta\|_1$ is the l_1 -norm of vector β , that is, $\|\beta\|_1 = \sum_{j=1}^d |\beta_j|$. The parameter $\lambda \geq 0$ controls the amount of regularization applied to the estimate. The larger λ , the larger the number of zeros in β . The nonzero components give the selected variables. After the optimal value of β is obtained, one can choose the feature indices corresponding to the top k largest values of the summation of the absolute values among each column.

Lasso requires the independence assumption of the input variables, however, in most real world data, features are often correlated. Therefore, in the presence of highly

correlated features, Lasso tends to select only one of these features at random, resulting in suboptimal performance (37). Moreover, the l_1 minimization algorithm is not stable when compared with l_2 minimization (30). For this reason, the elastic net (5) adds an additional l_2 regularization term into the Lasso objective function, which can be formulated as

$$\min_{\beta} \|\mathbf{y}^T - \beta^t \mathbf{X}\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2, \quad (8)$$

150 where $\lambda_1, \lambda_2 \geq 0$ are the tuning parameters.

Elastic Net can be seen as a linear combination of the Lasso and Ridge penalty. When $\lambda_1 = 0$, it becomes simple Ridge regression, when $\lambda_2 = 0$, it is equivalent to the Lasso penalty. Thus, Elastic Net enjoys a similar sparsity of representation of Lasso and also allows groups of correlated features to be selected. In the literature, it is reported that Elastic Net usually outperforms Lasso when the number of features is
155 much larger than the number of samples [32].

In summary, most of the exiting MI-based feature selection methods aim to develop an efficient quantitative evaluation criterion that simultaneously maximizes relevancy and minimizes redundancy. Unfortunately, it has been noted that such MI-based feature
160 selection methods have two common limitations. First, they tend to ignore pairwise relationship between samples of each feature dimension, which leads to significant information loss. Second, the majority of these methods cannot adaptively identify the most informative feature subset. Alternatively, sparse feature selection methods like Lasso and Elastic Net ensure parameter vector sparsity and allow the relevant features
165 to be adaptively selected. However, existing sparse feature selection methods also fail to encapsulate pairwise relationship between samples for each feature dimension. These drawbacks motivate us to develop a novel structurally interacting elastic net feature selection method to adaptively locate the most informative feature subset.

1.2. Contributions

170 As previously stated, the aim of this paper is to overcome the limitations of existing MI-based and sparse feature selection methods by developing a new structurally interacting elastic net feature selection algorithm. In summary, the contributions of

this work are threefold. First, we transform each vector feature into a graph-based feature representation, where each vertex represents a corresponding sample in each feature dimension and each weighted edge represents the pairwise relationship between samples from each feature dimension. We use the Euclidean distance to measure the pairwise relationship between samples. Similarly, we also construct a complete feature graph for the target feature. To measure the joint relevance of different pairwise feature combinations with respect to the target feature, we propose a new information theoretic criterion using the Jensen-Shannon divergence (JSD). Based on this criteria, we obtain a new interaction matrix which characterizes the informative relationships between feature pairs. Second, to a) incorporate the pairwise sample relationships in each feature dimension, b) remedy information loss in the original feature space, and c) adaptively select the most informative feature subset, we formulate the proposed graph-based feature selection method as an elastic net regression model. Specifically, the interaction matrix encapsulates high-dimensional structural relationships between feature samples and thus provides richer representation of structural interaction information between features. The ordinary least-square (OLS) term utilizes information from the original feature space, and thus remedies the information loss caused by representing features as graphs. In addition, the l_1 -norm regularizer ensures sparsity in the coefficients of variables and the l_2 -norm regularizer promotes a grouping effect to select correlated features. Third, an efficient alternating direction multiplier method (ADMM) is presented to solve the proposed elastic net optimization problem. Comprehensive experiments on eight standard machine learning datasets and two publically available datasets demonstrate the effectiveness of the proposed method.

1.3. Paper Outline

The remainder of the paper is organized as follows. Section 2 discusses the important concepts which will be used for the proposed feature selection method. Section 3 presents the proposed structurally interacting elastic net for feature selection. Section 4 provides our experimental evaluation. Finally, Section 5 concludes our work.

2. Preliminary Concepts

In this section, we review some preliminary concepts that will be used in this work. We commence by reviewing how to construct the feature graph to incorporate structural information for feature samples. Then we introduce the concept of Jensen-Shannon
 205 divergence, which will be used to calculate the similarity between feature graphs.

2.1. Construction of Feature Graphs

In this subsection, we introduce how to transform each vectorial feature into a complete weighted graph. The advantages of using the graph-based representation are twofold. First, graph structures have a stronger ability to encapsulate global topological
 210 information than vectors. Second, graphs can incorporate the relationships between samples of each original vector feature into the feature selection process, thus reducing information loss.

Given a dataset of N features denoted as $\mathcal{X} = \{\mathbf{f}_1, \dots, \mathbf{f}_i, \dots, \mathbf{f}_N\} \in \mathbb{R}^{M \times N}$, \mathbf{f}_i represents the i -th vectorial feature and has M samples $\mathbf{f}_i = (f_{i1}, \dots, f_{ia}, \dots, f_{ib}, \dots, f_{iM})^T$.
 215 We transform each feature \mathbf{f}_i into a graph-based feature representation $\mathbf{G}_i(V_i, E_i)$, where the vertex $v_{ia} \in V_i$ indicates the a -th sample f_{ia} of \mathbf{f}_i . Each pair of vertices v_{ia} and v_{ib} are connected by a weighted edge $(v_{ia}, v_{ib}) \in E_i$, and the dissimilarity weight $\omega(v_{ia}, v_{ib})$ of (v_{ia}, v_{ib}) is the Euclidean distance between f_{ia} and f_{ib} , i.e.,

$$\omega(v_{ia}, v_{ib}) = \sqrt{(f_{ia} - f_{ib})^2}. \quad (9)$$

Similarly, if the sample values of the target feature $\mathbf{Y} = (y_1, \dots, y_a, \dots, y_b, \dots, y_M)^T$ are continuous, its graph-based feature representation $\hat{\mathbf{G}}(\hat{V}, \hat{E})$ can be computed using Eq.(9) and the vertex \hat{v}_a represents the a -th sample y_a . However, for classification problems, the target features Y are the class labels and thus takes the discrete values $c \in \{1, 2, \dots, C\}$, i.e., the samples of each feature \mathbf{f}_i are assigned to the C different classes. In this case, we propose to compute the graph-based target feature $\hat{\mathbf{G}}_i(\hat{V}_i, \hat{E}_i)$ for each feature \mathbf{f}_i , where the dissimilarity weight $\omega(\hat{v}_{ia}, \hat{v}_{ib})$ of each edge $(\hat{v}_{ia}, \hat{v}_{ib}) \in \hat{E}_i$ is

$$\omega(\hat{v}_{ia}, \hat{v}_{ib}) = \sqrt{(\mu_{ia} - \mu_{ib})^2}, \quad (10)$$

where μ_{ia} is the mean value of all samples in \mathbf{f}_i from the same class c .

2.2. The Jensen-Shannon Divergence

In information theory, the JSD is a dissimilarity measure between probability distributions. Let two (discrete) probability distributions be $\mathcal{P} = (p_1, \dots, p_a, \dots, p_A)$ and $\mathcal{Q} = (q_1, \dots, q_b, \dots, q_B)$, then the JSD between \mathcal{P} and \mathcal{Q} is defined as

$$\text{JSD}(\mathcal{P}, \mathcal{Q}) = H_S\left(\frac{\mathcal{P} + \mathcal{Q}}{2}\right) - \frac{1}{2}H_S(\mathcal{P}) - \frac{1}{2}H_S(\mathcal{Q}), \quad (11)$$

where $H_S(\mathcal{P}) = \sum_{a=1}^A p_a \log p_a$ is the Shannon entropy of the probability distribution \mathcal{P} . In [39], the JSD has been used as a means of measuring the information theoretic dissimilarity between graphs associated with their probability distributions. In this work, we are concerned with the similarity between graph-based feature representations. Therefore, we adopt the negative exponential of $\text{JSD}(\mathcal{P}, \mathcal{Q})$ to compute the similarity measure I_S between probability distributions, i.e.,

$$I_S(\mathcal{P}, \mathcal{Q}) = \exp\{-\text{JSD}(\mathcal{P}, \mathcal{Q})\}. \quad (12)$$

3. The Proposed Feature Selection Method

In this section, we introduce our proposed structurally interacting elastic net feature selection method. We first detail the formulation of the structurally interacting elastic net model. To solve the optimization model, the alternating direction method of multiplier (ADMM) algorithm [40] is used to identify the most informative feature subset. Finally, we provide the convergence proof and complexity analysis for the method.

We propose to use the following information theoretic criterion to measure the joint relevance of different pairwise feature combinations with respect to target labels. For a set of N features $\mathbf{f}_1, \dots, \mathbf{f}_i, \dots, \mathbf{f}_j, \dots, \mathbf{f}_N$ and the associated continuous target feature \mathbf{Y} , the relevance degree of the feature pair $\{\mathbf{f}_i, \mathbf{f}_j\}$ is

$$W_{\mathbf{f}_i, \mathbf{f}_j} = \frac{I_S(\mathbf{G}_i, \hat{\mathbf{G}}) + I_S(\mathbf{G}_j, \hat{\mathbf{G}})}{I_S(\mathbf{G}_i, \mathbf{G}_j)}, \quad (13)$$

where \mathbf{G}_i and $\hat{\mathbf{G}}$ are the graph-based feature representations of \mathbf{f}_i and \mathbf{Y} , I_S is the JSD based information theoretic similarity measure defined in Eq.(12). The above relevance

measure consists of three terms. The first and second terms $I_S(\mathbf{G}_i, \hat{\mathbf{G}})$ and $I_S(\mathbf{G}_j, \hat{\mathbf{G}})$ are the relevance degrees of individual features \mathbf{f}_i and \mathbf{f}_j with respect to the target feature \mathbf{Y} , respectively. The third term $I_S(\mathbf{G}_i, \mathbf{G}_j)$ measures the relevance between the feature pair $\{\mathbf{f}_i, \mathbf{f}_j\}$. Therefore, $W_{\mathbf{f}_i, \mathbf{f}_j}$ is large if and only if both $I_S(\mathbf{G}_i, \hat{\mathbf{G}})$ and $I_S(\mathbf{G}_j, \hat{\mathbf{G}})$ are large (i.e., both \mathbf{f}_i and \mathbf{f}_j are informative themselves with respect to the target feature representation \mathbf{Y}) and $I_S(\mathbf{G}_i, \mathbf{G}_j)$ is small (i.e., \mathbf{f}_i and \mathbf{f}_j are not relevant).

For classification problems, the samples of the target feature \mathbf{Y} take the discrete value c and $c \in \{1, 2, \dots, C\}$. In this case, we compute the individual graph-based target feature representation $\hat{\mathbf{G}}_i$ for each feature \mathbf{f}_i , and the relevance measure defined in Eq.(13) can be written as

$$W_{\mathbf{f}_i, \mathbf{f}_j} = \frac{I_S(\mathbf{G}_i, \hat{\mathbf{G}}_i) + I_S(\mathbf{G}_j, \hat{\mathbf{G}}_j)}{I_S(\mathbf{G}_i, \mathbf{G}_j)}. \quad (14)$$

Similarly to Eq.(13), the three terms of Eq.(14) have the same corresponding theoretical significance.

Furthermore, based on the graph-based feature representations, we construct a feature informativeness matrix \mathbf{W} , where each element $W_{i,j} \in \mathbf{W}$ represents the information theoretic measure between a feature pair $\{\mathbf{f}_i, \mathbf{f}_j\}$ based on Eq.(13) (for \mathbf{Y} is continuous) or Eq.(14) (for \mathbf{Y} is discrete). Given the informativeness matrix \mathbf{W} and the d -dimensional feature indicator vector β , where β_i represents the coefficient for the i -th feature, we can identify the informative feature subset by solving the following maximization problem

$$\max f(\beta) = \max_{\beta \in \mathbb{R}^d} \beta^T \mathbf{W} \beta, \quad (15)$$

subject to $\beta \in \mathbb{R}^N, \beta \geq 0$. The solution vector $\beta = (\beta_1, \dots, \beta_d)^T$ to the above quadratic program is an N -dimensional vector. When $\beta_i > 0$, the i -th feature \mathbf{f}_i belongs to the most informative feature subset i.e., feature \mathbf{f}_i is selected if and only if the i -th component of β is positive ($i \in \{1, 2, \dots, d\}$).

3.1. The Proposed Structurally Interacting Elastic Net

The proposed feature subset selection algorithm aims to incorporate structural information between pairwise features and simultaneously allow correlated features to

be selected, as well as promote a sparse solution. Therefore, we combine Eq.(15) and Eq.(8) to construct the associated structurally interacting elastic net for feature selection with the following mathematical form

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2 - \lambda_3 \beta^T \mathbf{W} \beta, \quad (16)$$

where λ_1 and λ_2 are the tuning parameters in the elastic net regression model, and λ_3 is the associated tuning parameter for the structural interaction matrix \mathbf{W} .

245 It can be seen that the first term in Eq.(16) remedies the information loss from the original feature space, while the second and third terms ensure the sparsity and grouping among selected features. The fourth term incorporates structural information concerning the relationships between feature samples. Because $\beta^T \mathbf{W} \beta$ is a non-convex constraint, the proposed method distinguishes itself from existing Lasso-type
 250 methods using convex optimization methods, which may become trapped in suboptimal solutions. Specifically, for the proposed model (16), we need to develop efficient algorithms to obtain the optimal solutions (denoted as β^*). A feature f_i is selected if and only if $\beta_i^* > 0$. Consequently, we can recover the number of features in the optimal feature subset according to the number of positive components of β^* .

255 3.2. Optimization Algorithm

To solve the non-convex problem (16), we develop an optimization algorithm using ADMM [40]. The ADMM approach is a powerful algorithm that is well suited to problems arising in machine learning. The basic principle of the ADMM approach is to decompose a hard optimization problem into a series of smaller ones, each of which
 260 is simpler to handle. It takes the form of a decomposition-coordination procedure, in which the solutions to small local subproblems are coordinated to find a solution to a large global problem. ADMM can be viewed as an attempt to blend the benefits of dual decomposition and augmented Lagrangian methods for constrained optimization. It turns out to be equivalent or closely related to many well known algorithms as well,
 265 such as Douglas-Rachford splitting from numerical analysis [41], proximal methods [42], and many others.

In ADMM form, problem (16) can be re-written as

$$\begin{aligned} \min_{\beta \in \mathbb{R}^d} \quad & \frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2 + \lambda_2 \|\beta\|_2^2 - \lambda_3 \beta^T \mathbf{W} \beta + \lambda_1 \|\gamma\|_1 \\ \text{s.t.} \quad & \beta - \gamma = 0, \end{aligned} \quad (17)$$

where γ is an auxiliary variable, which can be regarded as a proxy for vector β . By doing so, the objective function can be split into two separate parts associated with two different variables, i.e., β and γ , indicating that the hard constrained problem can be solved separately. As in the method of multipliers, we form the augmented Lagrangian function associated with the constrained problem (16) as follows

$$L_\rho(\beta, \gamma, z) = \frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2 + \lambda_2 \|\beta\|_2^2 - \lambda_3 \beta^T \mathbf{W} \beta + \lambda_1 \|\gamma\|_1 + \langle \beta - \gamma, z \rangle + \frac{\rho}{2} \|\beta - \gamma\|_2^2, \quad (18)$$

where $\langle \cdot, \cdot \rangle$ is an Euclidean inner product, z is a dual variable (i.e., the Lagrange multiplier) associated with the equality constraint $\beta = \gamma$, and ρ is a positive penalty parameter (step size for dual variable update). By introducing an additional variable γ and an additional constraint $\beta - \gamma = 0$, we have simplified the optimization problem (16) by decoupling the objective function into two parts that depend on two different variables. In other words, ADMM decomposes the minimization of $L_\rho(\beta, \gamma, z)$ into two simpler subproblems. Specifically, ADMM solves the original problem (16) by seeking for a saddle point of the augmented Lagrangian by iteratively minimizing $L_\rho(\beta, \gamma, z)$ over β , γ , and z . In ADMM, the variables β , γ , and z are updated in an alternating or sequential fashion, which accounts for the term alternating direction. This updating rule is shown below

- (1) $\beta^{k+1} = \arg \min_{\beta \in \mathbb{R}^d} L(\beta, \gamma^k, z^k)$, // β -minimization
- (2) $\gamma^{k+1} = \arg \min_{\gamma \in \mathbb{R}^d} L(\beta^{k+1}, \gamma, z^k)$, // γ -minimization
- (3) $z^{k+1} = z^k + \rho(\beta^{k+1} - \gamma^{k+1})$, // z -update

Given the above updating rule, we need to resolve each sub-problem iteratively until the termination criteria is satisfied. Using ADMM, we perform the following calculation steps at each iteration.

(a) Update β

In the $(k + 1)$ -th iteration, in order to update β^k , we need to solve the following

sub-problem, where the values of γ^k and z^k are fixed

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2 + \lambda_2 \|\beta\|_2^2 - \lambda_3 \beta^T \mathbf{W} \beta + \lambda_1 \|\gamma\|_1 + \langle \beta - \gamma^k, z^k \rangle + \frac{\rho}{2} \|\beta - \gamma^k\|_2^2. \quad (19)$$

Let the partial derivative with respect to β be equal to zero, we have

$$\frac{\partial [\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2 + \lambda_2 \|\beta\|_2^2 - \lambda_3 \beta^T \mathbf{W} \beta + \lambda_1 \|\gamma\|_1 + \langle \beta - \gamma^k, z^k \rangle + \frac{\rho}{2} \|\beta - \gamma^k\|_2^2]}{\partial \beta} = 0, \quad (20)$$

because

$$\begin{cases} \frac{\partial (\frac{1}{2} \|\mathbf{y}^T - \beta^T \mathbf{X}\|_2^2)}{\partial \beta} = -\mathbf{X} \mathbf{y} + \mathbf{X} \mathbf{X}^T \beta \\ \frac{\partial (\lambda_2 \|\beta\|_2^2)}{\partial \beta} = \lambda_2 \beta \\ \frac{\partial (-\lambda_3 \beta^T \mathbf{W} \beta)}{\partial \beta} = -2\lambda_3 \mathbf{W} \beta \\ \frac{\partial \langle \beta - \gamma^k, z^k \rangle}{\partial \beta} = z^k \\ \frac{\partial (\frac{\rho}{2} \|\beta - \gamma^k\|_2^2)}{\partial \beta} = \rho(\beta - \gamma^k), \end{cases} \quad (21)$$

we have

$$-\mathbf{X} \mathbf{y} + \mathbf{X} \mathbf{X}^T \beta + \lambda_2 \beta - 2\lambda_3 \mathbf{W} \beta + z^k + \rho(\beta - \gamma^k) = 0, \quad (22)$$

that is,

$$\beta^{k+1} = (\mathbf{X} \mathbf{X}^T + \lambda_2 \mathbf{I} - 2\lambda_3 \mathbf{W} + \rho \mathbf{I})^{-1} (\mathbf{X} \mathbf{y} - z^k + \rho \gamma^k). \quad (23)$$

285

(b)Update γ

Based on the results, assume β_i^{k+1} and z_i^k are fixed, for $i = 1, 2, \dots, d$, we update γ_i^{k+1} by solving the following sub-optimization problem

$$\min_{\gamma_i} \lambda_1 \sum_{i=1}^d \|\gamma_i\|_1 - \sum_{i=1}^d \langle \gamma_i, z_i^k \rangle + \frac{\rho}{2} \sum_{i=1}^d \|\beta_i^{k+1} - \gamma_i\|_2^2, \quad (24)$$

$$\frac{\partial [\min_{\gamma_i} \lambda_1 \sum_{i=1}^d \|\gamma_i\|_1 - \sum_{i=1}^d \langle \gamma_i, z_i^k \rangle + \frac{\rho}{2} \sum_{i=1}^d \|\beta_i^{k+1} - \gamma_i\|_2^2]}{\partial \gamma_i} = 0. \quad (25)$$

We therefore have the following results

$$\gamma_i^{k+1} = \begin{cases} \frac{1}{\rho} (z_i^k + \rho \beta_i^{k+1} - \lambda_1), & \text{if } z_i^k + \rho \beta_i^{k+1} > \lambda_1 \\ \frac{1}{\rho} (z_i^k + \rho \beta_i^{k+1} - \lambda_1), & \text{if } z_i^k + \rho \beta_i^{k+1} < -\lambda_1 \\ 0, & \text{if } z_i^k + \rho \beta_i^{k+1} \in [-\lambda_1, \lambda_1] \end{cases} \quad (26)$$

(c) **Update \mathbf{z}**

Then, assume β_i^{k+1} and γ_i^{k+1} are fixed, for $i = 1, 2, \dots, d$, we update z_i^{k+1} by using the following equation

$$z_i^{k+1} = z_i^k + \rho(\beta_i^{k+1} - \gamma_i^{k+1}). \quad (27)$$

Based on procedures (a), (b), and (c), we summarize the optimization algorithm below

Algorithm 1 The proposed ADMM algorithm for structurally interacting Elastic Net.

Input: $\mathbf{X}, \mathbf{y}, \beta^0, z^0, \lambda_1, \lambda_2, \lambda_3, \rho$

Step1: While (not converged), **do**

Step2: Update β^{k+1} according to Eq.(23)

Step3: Update $\gamma_i^{k+1}, i = 1, 2, \dots, d$ according to Eq.(26)

Step4: Update $\beta_i^{k+1}, i = 1, 2, \dots, d$ according to Eq.(27)

End While

Output: β^* .

290 **3.3. Complexity Analysis and Convergence Proof**

In this subsection, we provide an analyses of the properties of the proposed structurally interacting elastic net method. We commence by presenting the computational complexity which is followed by a convergence analysis.

3.3.1. Analysis of Computational Complexity

295 Let N be the number of features, M the number of samples, and K the required number of iterations to converge. At each iteration, the computational complexity for updating β according to Eq.(23) is $O(N^2M)$. Additionally, the computational costs for updating γ in Eq.(26) and z in Eq.(27) are both $O(N)$. Therefore, the overall time complexity of the ADMM algorithm is calculated as $\max\{O(N^2MK), O(MK)\}$.

300 **3.3.2. Convergence Proof**

To theoretically prove the convergence of the proposed ADMM algorithm, we present the following analysis.

Theorem 1. Assume the iterative sequences generated by the ADMM algorithm are denoted as $\{\beta^k\}$, $\{\gamma^k\}$, and $\{z^k\}$, respectively. Suppose as k tends to infinity, the sequence $\{z^k\}$ converges to a point z' , that is, $\lim_{k \rightarrow \infty} z^k = z'$. Following this, every limit point (β', γ') of the iteration sequence $\{\beta^k, \gamma^k\}$, together with z' , satisfy the necessary first order conditions of the problem (16), that is

(1) Primal feasibility, i.e., $\beta' - \gamma' = 0$.

(2) Dual feasibility, i.e., $\nabla f(\beta') + z' = 0$ and $0 \in \partial g(\gamma') - z'$, where ∂ denotes the sub-differential operator.

We can easily prove **Theorem 1** by following a proof similar to that of Proposition 3 in Magnússon[43]. We can conveniently draw the conclusion from **Theorem 1** that, in general, the ADMM algorithm converges to a local optimum solution to problem (16), that is, $(\beta', \gamma', z') = (\beta^*, \gamma^*, z^*)$.

4. Experimental Results and Discussion

In this section, we conduct a series of experiments to demonstrate the performance of the proposed structurally interacting elastic net feature selection method (InElasticNet). A comprehensive experimental study on two types of datasets is conducted to validate its effectiveness and make comparison with several state-of-the-art feature selection methods.

4.1. Experiments on Standard Machine Learning Datasets

Two categories of public datasets are used in our experiment, including eight widely used machine learning (ML) datasets and two public available datasets. The ML datasets are the USPS handwritten digit data set [44], Isolet speech data set and Pie data set from the UCI Machine Learning Repository [45], YaleB face data set [46], Lymphoma and Leukemia datasets [47], BASEHOCK and RELATHE. Note that the last two datasets are both large in feature dimension and sample size. Detailed information for these data sets are summarized in Table 1.

To evaluate the discriminative capabilities of the information captured by our method, we compare the classification results obtained using the selected features from our pro-

Table 1: Statistics of data sets used in the experiments

Name	Feature Dimension	Sample Number	Class Number
USPS	256	9298	10
Isolet1	617	1560	26
Pie	1024	11554	68
YaleB	1024	2414	38
Lymphoma	4026	96	9
Leukemia	7129	73	2
BASEHOCK	4862	1993	2
RELATHE	4322	1427	2

posed method with several state-of-the-art feature selection methods including Lasso [31], ULasso [48], Fused Lasso [49], Elastic Net [32], Group Lasso [33], InLasso [30], and one graph-based feature selection methods, namely, GF-RW [27].

a) Lasso [31]: As a typical sparse feature selection method, Lasso performs feature selection through the l_1 -norm, where features corresponding to zero coefficients in the parameter vector will be discarded.

b) ULasso [48]: The uncorrelated Lasso (ULasso) aims to conduct variable de-correlation and variable selection simultaneously, so that the variables selected are uncorrelated as much as possible.

c) Fused Lasso [49]: The fused lasso enforces sparsity in both the coefficients and their successive differences. It is desirable for applications with features ordered in some meaningful way.

d) Group Lasso [33]: The group Lasso is known to enforce the sparsity on variables at an inter-group level, where variables from different groups are competing to survive.

e) Elastic Net [32]: In statistics, the elastic net is a regularized regression method that linearly combines the l_1 and l_2 penalties of the Lasso and Ridge methods. This ensures democracy among groups of correlated groups and allows selection of the relevant groups while simultaneously promoting sparse solutions for feature selection.

f) InLasso [30]: Is a Lasso-type regression model which incorporates high-order feature interactions, InLasso can effectively evaluate whether a feature is redundant or interactive based on a neighborhood dependency measure. This method can avoid discarding some valuable features arising in individual feature combinations.


g) GF-RW [27]: Is a graph-based feature selection method which incorporates pairwise relationship between samples of each feature dimension.

Table 2: Comparison of classification accuracy


Dataset	Lasso	ULasso	FusedLasso	ElasticNet	GroupLasso	InLasso	GF-RW	InElasticNet
USPS	68.17% ±2.61	66.19% ±1.61	68.09% ±1.40	70.47% ±1.83	60.71% ±1.28	84.91% ±1.05	66.99 ±0.17	82.83% ±0.11
Isolet1	71.53% ±3.70	73.49% ±3.29	70.49% ±3.19	72.80% ±3.93	66.97% ±3.23	79.79% ±2.95	75.53 ±0.39	78.82% ±0.33
Pie	85.11% ±0.99	86.18% ±0.97	79.77% ±1.05	85.14% ±0.95	75.67% ±0.96	87.48% ±0.87	75.21 ±0.91	84.73% ±0.10
YaleB	31.62% ±2.77	34.93% ±2.74	29.52% ±2.52	36.75% ±3.90	28.44% ±2.71	54.42% ±2.67	92.86 ±0.49	78.94% ±0.25
Leukemia	70.43% ±1.75	76.71% ±1.63	92.00% ±1.26	94.29% ±1.05	83.71% ±1.53	98.29% ±0.33	97.14 ±1.67	98.29% ±1.73
Lymphoma	84.22% ±1.19	82.78% ±1.30	79.34% ±1.40	82.67% ±1.31	78.33% ±1.49	88.67% ±0.95	93.23 ±0.97	90.10% ±1.03
BASEHOCK	61.23% ±0.37	61.04% ±0.39	74.53% ±0.28	89.09% ±0.35	66.34% ±0.35	83.77% ±0.33	67.25 ±0.38	90.25% ±0.32
RELATHE	74.58% ±0.35	75.42% ±0.49	74.57% ±0.44	73.39% ±0.42	76.07% ±0.41	76.04% ±0.24	76.36 ±0.40	78.87% ±0.23
AVG	68.35%	69.60%	71.04%	75.57%	67.03%	81.62%	80.54%	85.36%

355 Generally, we adopt a 10-fold cross-validation method associated with C-SVM to
 evaluate the classification accuracy. To be specific, we first partition the entire sample
 randomly into 10 subsets (each subset with roughly equal size) and then we choose one
 subset for testing and use the remaining 9 subsets for training. We repeat this procedure
 for 10 times. The final accuracy is computed by averaging of the accuracies from all
 360 10 experiments and we also compute the associated standard error.

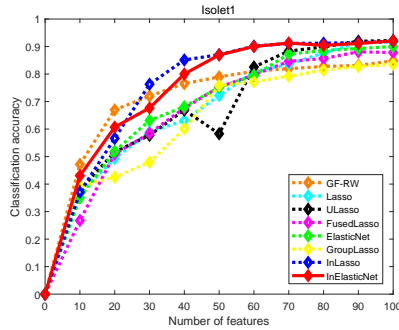
Fig. 1 shows the classification accuracy versus the number of selected features on
 the datasets for different methods. It is clear from the figure that the proposed method
 InElasticNet is, by and large, superior to the following alternative feature selection
 methods including Lasso, ULasso, Elastic Net, Fused Lasso, and Group Lasso on all
 365 datasets. When compared with InLasso, it is clear that our method significantly outper-
 forms InLasso on the YaleB dataset and also outperforms InLasso for BASEHOCK and
 RELATHE, which are challenging datasets both large in feature dimension and sample
 size. For the remaining datasets, our method is competitive to InLasso. In addition,
 our method significantly outperforms GF-RW method on USPS, Pie and BASEHOCK,
 370 and is superior than GF-RW on Isolet1, Leukemia, and RELATHE datasets. As Fig. 1
 shows, when the number of selected features is too small, the advantage of our method
 is not clear. However, when the number of total features in the selected subset increases
 to a certain number, the InElasticNet method performs much better than the alterna-




(a) YaleB dataset




(b) USPS dataset




(c) Isolet1 dataset




(d) Lymphoma dataset




(e) Pie dataset



(f) Leukemia dataset



(g) BASEHOCK dataset



(h) RELATHE dataset

Figure 1: Accuracy rate vs. the number of selected features on 8 benchmark machine learning datasets

tive methods. The results verify that the proposed structurally interacting Elastic Net
375 can identify more informative feature subsets than the state-of-the-art feature selection
methods. Although the total numbers of features in the selected feature subsets are
on average slightly larger than those obtained via Lasso, it is still small as compared
to the total number of features in the datasets. This is because our method can select
correlated features and also encourage a sparse solution, whereas Lasso tends to select
380 only one feature from a group of correlated features, which may decrease its accuracy.

To make a detailed comparison, we report the mean classification accuracies and
corresponding variances (i.e., $\text{MEAN} \pm \text{STD}$) obtained via various methods on each da-
ta set with different number of features selected by using the C-SVM classifier in Ta-
ble 2. The mean classification accuracy is obtained by averaging the accuracy achieved
385 via C-SVM using the largest number of features indicated in the corresponding sub-
figures of Fig. 1 for each data set. For instance, for the YaleB dataset, we use the top
10, 20, ..., 50 features selected by each algorithm. The boldfaced value of each row cor-
responds to the highest accuracy obtained by the different methods for the underlying
dataset. Our proposed method, i.e., InElasticNet improves the classification accuracy
390 by 24.52% for the YaleB dataset and 1.43% for the Lymphoma dataset, respectively.
For the Leukemia dataset, InElasticNet performs equally to InLasso, and is better than
the alternative methods. As for Isolet1, USPS, and Pie, InElasticNet can obtain better
classification accuracy than Lasso, ULasso, FusedLasso, Elastic Net, and Group Lasso,
and is competitive to InLasso, which retains the highest classification accuracy. Addi-
395 tionally, for the two large datasets Basehock and Relathe, InElasticNet outperforms all
the competitors.

The bottom row of Table 2 displays the average classification accuracy for each
algorithm over the eight datasets. It shows that our proposed method, i.e., InElasticNet
improves the classification accuracy by 17.11%(Lasso), 15.57%(ULasso), 15.75%(Fus-
400 edLasso), 11.93%(ElasticNet), 19.98%(GroupLasso), 3.36%(InLasso), and 4.82%(GF-
RW), respectively, compared to the averaged classification accuracy of all alternative
methods on the eight datasets. In addition, it is worth noting that the standard errors for
the proposed InElasticNet method are smaller than the competing methods for almost
all datasets, except for Leukemia. This indicates that InElasticNet is more stable than

Table 3: Win/Tie/Lost matrix for the feature selection methods used in the experiments.

Methods	Lasso	ULasso	FusedLasso	ElasticNet	GroupLasso	InLasso	GF-RW	InElasticNet	Total
Lasso		0/6/2	3/3/2	0/4/4	5/1/2	0/1/7	1/2/5	0/1/7	9/18/29
ULasso	2/6/0		4/2/2	1/4/3	5/1/2	0/2/6	1/2/5	0/1/7	13/18/25
FusedLasso	2/3/2	2/2/4		0/1/7	5/3/0	0/1/7	2/2/4	0/0/8	11/12/33
ElasticNet	4/4/0	3/4/1	7/1/0		7/1/0	1/1/6	3/0/5	0/2/6	25/13/18
GroupLasso	2/1/5	2/1/5	0/3/5	0/1/7		0/1/7	0/3/5	0/0/8	4/10/42
InLasso	7/1/0	6/2/0	7/1/0	6/1/1	7/1/0		4/2/2	2/3/3	39/11/6
GF-RW	5/2/1	5/2/1	4/2/2	5/0/3	5/3/0	2/2/4		2/1/5	28/12/16
InElasticNet	7/1/0	7/1/0	8/0/0	6/2/0	8/0/0	3/3/2	5/1/2		44/ 8/ 4

Table 4: The best result of all methods and the corresponding size of selected feature subset

Dataset	USPS	Isolet1	Pie	YaleB	Leukemia	Lymphoma	BASEHOCK	RELATHE
Lasso	86.30%(50)	91.67%(100)	94.48%(70)	46.64%(50)	82.86%(120)	94.44%(160)	66.33%(200)	86.00%(200)
ULasso	83.24%(50)	92.18%(100)	94.57%(70)	47.43%(50)	82.86%(200)	91.11%(200)	67.30%(200)	84.62%(200)
FusedLasso	87.40%(50)	88.08%(90)	93.53%(70)	55.89%(50)	98.57%(20)	94.44%(160)	84.62%(200)	86.50%(200)
ElasticNet	87.43%(50)	90.00%(100)	86.94%(70)	48.09%(50)	98.57%(180)	90.00%(120)	91.76%(200)	77.95%(180)
GroupLasso	83.93%(50)	83.53%(100)	92.35%(70)	45.02%(50)	91.43%(180)	91.11%(200)	73.33%(200)	74.16%(200)
InLasso	93.94%(50)	91.92%(100)	96.58%(70)	71.20%(50)	100%(80)	95.56%(140)	86.58%(200)	80.70%(180)
GF-RW	85.79%(50)	84.80%(100)	90.64%(70)	98.38%(50)	98.57%(20)	95.56%(160)	74.22%(200)	76.36%(180)
InElasticNet	94.10%(50)	92.23%(100)	96.81%(70)	94.62%(50)	100%(120)	95.56%(160)	92.75%(200)	83.66%(180)

405 the competing methods.

Table 3 presents the Win/Tie/Lost matrix for the feature selection methods used in the experiments. The (i, j) th element of the matrix represents the number of datasets where the method corresponding to the i th row has won/tied/lost against the method corresponding to the j th column. A tie is defined as a dataset on which difference in classification accuracy between two methods is not statistically significant. The last column of Table 3 shows the total number of wins/ties/lost for a given method, and the best performing method is highlighted in bold. InElasticNet has the largest total number of wins and the smallest total number of lost. This clearly indicates that the proposed InElasticNet method performs significantly better than the alternative feature selection methods.

415 Table 4 shows the best results for each competing method together with their corresponding number of features in the selected subset. In the table, the best classification accuracy is shown which is followed by the optimal number of features selected in brackets. From this table, it is clear that the proposed method achieves the highest classification accuracy using same number of features in the selected subset as the alternative methods. This implies that our proposed method tremendously outperforms


all the competing methods and has more discriminative power.

The ROC curves of the most competitive methods (ElasticNet, GF-RW, InLasso and InElasticNet) on the eight datasets are plotted in Figure 2. From this figure, we
425 observe that except for the YaleB dataset, the proposed InElasticNet method achieves superior performance to the competitors on all datasets. In summary, the aforementioned experimental results demonstrate that our proposed feature selection method outperforms the alternative methods on the standard ML datasets.


4.2. Experiments on Two Real World Datasets

430 Apart from the ML datasets, two publically available datasets including i) a Peer-to-Peer (P2P) dataset collected from the P2P lending sector in China and ii) a healthcare dataset collected from a well-known medical platform is used to validate the effectiveness of the proposed feature selection approach.


Since the launch of the first P2P lending platform in 2007, the P2P lending industry has developed rapidly and the market is enormous. Specifically, the total number of
435 operational P2P lending platforms nationwide has reached 2,448 with an accumulative loan amount of 20 trillion yuan by the end of year 2016. Along with this rapid development, the P2P lending industry has also experienced some serious problems with rising defaults and weak risk control. Therefore, it is of great significance to develop
440 an effective decision aid for the credit risk analysis of P2P platforms. However, the P2P lending data are often high-dimensional, highly correlated and unstable. Therefore, it presents a challenge for traditional statistical pattern recognition and machine learning techniques. The aim is to effectively analyze the data and identify which factors influence the performance of the lending platforms, or the default probability of
445 the borrowers, etc. To realize these goals, the sample relationships of the P2P data that encapsulates significant information should be incorporated into the feature selection process. However, the majority of existing feature selection methods ignore the sample relationships and may cause significant information loss. By contrast, our proposed structurally interacting feature selection approach is able to encapsulate the sample
450 relationships of P2P data and overcome these shortcomings.




(a) YaleB dataset




(b) USPS dataset




(c) Isolet1 dataset




(d) Lymphoma dataset




(e) Pie dataset



(f) Leukemia dataset



(g) Basehock dataset



(h) Realthe dataset

Figure 2: ROC Curves for Different Datasets.

The P2P dataset is collected from a reputable P2P lending portal in China¹, which tracks the industry. The dataset consists of the most popular 200 platforms (i.e., 200 samples) until Aug 2014. For each platform, we collect 19 features including 1) transaction volume, 2) total turnover, 3) average annualized interest rate, 4) the total number
455 of borrowers, 5) the total number of investors, 6) the online time, which refers to the foundation year of the platform, 7) the operation time, i.e., number of months since the foundation of the platform, 8) registered capital, 9) weighted turnover, 10) average term of loan, 11) average full mark time, i.e., tender period of a loan raised to the required full capital, 12) average amount borrowed, i.e., average loan amount of each success-
460 ful borrower, 13) average amount invested, which is the average investment amount of each successful investor, 14) loan dispersion, i.e., the ratio of the repayment amount to the total capital, 15) investment dispersion, the ratio of the invested amount to the total capital, 16) average times of borrowing, 17) average times of investment, 18) loan balance, and 19) popularity.

465 We evaluate the performance of the proposed feature selection approach with respect to continuous target features. Specifically, we use the proposed method to perform credit risk evaluation of the P2P lending platforms. As it is difficult to obtain sufficient data for the platforms which encountered a problem, we use the annualized average interest rate as an indicator of the credit risk of the P2P lending platforms. In
470 finance, interest rate is the amount charged, expressed as a percentage of principal, by a lender to a borrower for the use of assets. When the borrower is a low-risk party, they will usually be charged a low interest rate. On the other hand, if the borrower is considered high risk, the interest rate charged will be higher. Likewise, a higher annualized average interest rate of the P2P lending platforms often indicates a greater
475 likelihood of default, i.e., higher credit risk of the platforms. Identifying the features most relevant to the interest rate can help investors effectively manage the credit risks involved in P2P lending. Therefore, in our experiment, we set the average annualized interest rate as the target feature which takes continuous values. Our aim is to identify the most informative subset of features for the credit risk of the P2P platforms by us-

¹See the website <http://www.wdzt.com/> for more details

Table 5: Comparison of three methods for the P2P dataset

Ranking	ElasticNet	InLasso	InElasticNet
1#	Loan dispersion	Average amount invested	Average amount invested
2#	Investment dispersion	Average times of investment	Average times of investment
3#	Popularity	Online time	Online time
4#	Operation time	Total number of investors	Investment dispersion
5#	Average times of borrowing	Average term of loan	Loan balance
6#	Online time	Average times of borrowing	Popularity
7#	Total number of borrowers	Average amount borrowed	Total turnover
8#	Loan balance	Investment dispersion	Weighted turnover
9#	Transaction volume	Loan dispersion	Transaction volume
10#	Weighted turnover	Total number of borrowers	Average times of borrowing

480 ing the proposed feature selection method. To further strengthen our findings, we also compare the proposed feature selection method with two alternative methods including Elastic Net [32] and Interacted Lasso [30].

Table 5 presents a comparison of the results obtained using the competing methods. For each method, we display the top 10 features in terms of relevance to the average
485 annualized interest rate. It is worth noting that all three methods can identify some similar influential factors but differ from each other in the remaining factors. For instance, both InLasso and InElasticNet rank **the average amount invested, average times of investment, and online time** as the top three most influential factors. This is reasonable because a longer online time indicates that the P2P platform is in operation
490 for a relatively longer period of time, and is less risky. Moreover, a larger average amount invested and a higher level of the average times of investment indicate a higher preference of the investors for the P2P lending platform due to a higher degree of security. In addition, both methods consider **investment dispersion** as a relevant feature but with different rankings, i.e., the 4th for InElasticNet and the 8th for InLasso.
495 This is reasonable because investment dispersion is highly correlated with the average amount invested and the average times of investment. Therefore, when InElasticNet ranks these factors high, it also tends to rank investment dispersion high. This implies that our proposed method can encourage a grouping effect for highly correlated features. This is further demonstrated by the fact that **the rankings of popularity (6th), total turnover(7th), and weighted turnover (8th)** are close to each other. Moreover,
500 these three factors are also correlated to each other. Whereas for InLasso, when groups of correlated features exist, it can only select one feature from the group. Therefore,

InLasso may fail to recognize some highly relevant features.

When compared to ElasticNet, it is worth noting that although both Elastic Net and
505 the proposed method (InElasticNet) can identify **online time, popularity, weighted
turnover, transaction volume, average times of borrowing, and investment disper-
sion** as influential factors, their rankings are quite different. This meets our expecta-
tions because both ElasticNet and InElasticNet can promote a grouping effect. Howeve-
er, as InElasticNet utilizes the structural information between pairwise feature samples,
510 the results obtained are more encouraging. For instance, ElasticNet ranks **loan disper-
sion (1st)** and **investment dispersion (2nd)** as the most influential factors whereas
InElasticNet ranks **the average amount invested** as the highest and **the average times
of investment** as the second highest. Unfortunately, a higher level of loan dispersion
and investment dispersion does not necessarily correspond to a safer P2P platform with
515 a lower annualized interest rate. By contrast, a larger average amount invested and a
higher level of the average times of investment often indicate a higher preference of the
investors for the P2P lending platform due to a higher degree of security and a lower
level of annualized interest rate. These results demonstrate the effectiveness of the pro-
posed method for identifying the most influential factors for credit risk of P2P lending
520 platforms.

The healthcare dataset is collected from a well-known medical platform in China²,
which presents evaluation of doctors. The dataset consists of 2363 doctors (i.e., 2363
samples). For each doctor, we collect 13 features including 1) patients, i.e., total num-
ber of patients treated, 2) title, i.e., title of the doctor, 3) grade, i.e., recommended grade
525 of the doctor provided by the website, 4) notes, i.e., total number of notes of thanks
posted by the patients, 5) gifts, i.e., total number of gifts received, 6) outpatients, i.e.,
total number of outpatients of the doctor, 7) city, i.e., city of the doctor, 8) appointments,
i.e., total number of appointments received from the patients, 9) visits, i.e., total num-
ber of visits of the doctor’s personal website, 10) contribution value, i.e., contribution
530 value of the doctor, 11) posts, total number of posts published by the patients about the
doctor, 12) votes, i.e., total number of votes received from the patients for a doctor, and

²See the website <http://www.haodf.com/> for more details

Table 6: Comparison of three methods for the healthcare dataset

Ranking	ElasticNet	InLasso	InElasticNet
1#	City	Title	Title
2#	Title	Grade	City
3#	Grade	City	Grade
4#	Notes	Votes	Votes
5#	Votes	Contribution value	Outpatients
6#	Appointments	Visits	Contribution value

13) registration fee for the doctor.

We use the proposed feature selection method to evaluate the doctors. The registration fee is treated as the continuous target and we aim to identify which features are the most informative ones with respect to the target. Like for the P2P lending analysis, we also compare the proposed feature selection method with two alternative methods including Elastic Net [32] and Interacted Lasso [30].

Table 6 presents a comparison of the results obtained using the three methods. For each method, we display the top 6 features in terms of relevance to the registration fee. It is worth noting that all three methods can identify **title of the doctor**, **city located** and **grade of the doctor** as the top three most influential factors. However, the rankings of these factors are different. For instance, both InLasso and InElasticNet rank **title of the doctor** the first, but ElasticNet ranks this factor as second. Compared with InLasso, our method considers **city** as the second most influential factor whereas InLasso ranks **Grade** as the second. We believe the results obtained via our method is more reasonable because a doctor with a higher title and in bigger cities are more expensive. Although grade is also relevant to the registration fee of the doctor, it is not an objective evaluation criteria. In addition, both InLasso and InElasticNet consider **the number of votes** as the fourth highest influential factor. This is also reasonable because a greater number of votes received from the patients indicates a higher reputation of the doctor. An interesting finding is that only our method can identify **outpatients** as the top ranking features whereas the two competing methods consider **appointments** and **visits** as the most influential features. We believe outpatients is a more relevant feature to the registration fee of the doctor because a higher number of outpatients often indicate that more patients are willing to pay more to be treated by the doctor. In addition, outpatients and votes are closely related to each other, and only the proposed

method can select the highly correlated features.

5. Conclusion

The main goal of feature selection is to automatically identify a subset of the most
560 informative features that is small in size but high in classification accuracy. To realize
this goal, in this paper, we have developed a new structurally interacting elastic net
feature selection method. The major contributions of this paper are threefold. First,
the proposed method can encapsulate structural relationships between feature samples
into the feature selection process by representing features as graphs and samples as
565 graph vertices. Accordingly, the informativeness matrix obtained is used to construct
an optimization model to identify the features with maximum relevancy and minimum
redundancy to the target feature. Second, to remedy the information loss caused by
using graph-based feature representations, we formulate the feature selection problem
using an elastic net regression model and solve this model using ADMM. This allows
570 us to a) incorporate information from the original feature space, b) reduce the number
of features to a small size and c) promote grouping effects. The experimental results on
real datasets show that our method outperforms several well-known feature selection
methods.

We plan to extend our method in a number of ways. First, in this paper, the con-
575 structed feature graphs are complete weighted graphs. However, in real world appli-
cations, not all connections may be dominant and useful. In other words, the com-
plete weighted graphs may contain some noise. Therefore, one may want to define
a sparser graph. Second, in our previous work [50], we have developed a number
of quantum Jensen-Shannon kernels using both the continuous-time and discrete-time
580 quantum walks. It would be interesting to extend the proposed feature selection method
using the classical Jensen-Shannon divergence to that using its quantum counterpart.
Finally, the proposed feature selection method only considers the relationships between
pairwise features, i.e., it only evaluates the two-order relationships between features.
Our future work will extend the proposed method into a high-order feature selection
585 method by establishing higher order relationships between features.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant no. 61602535 and 61503422), the Open Project Program of the National Laboratory of Pattern Recognition (NLPR), and the program for innovation research in Central
590 University of Finance and Economics.

References

- [1] J. Han, Z. Sun, H. Hao, Selecting feature subset with sparsity and low redundancy for unsupervised learning, *Knowl.-Based Syst.* 86 (2015) 210–223.
- [2] J. A. Lee, D. H. Peluffo-Ordóñez, M. Verleysen, Multi-scale similarities in stochastic neighbour embedding: Reducing dimensionality while preserving both
595 local and global structure, *Neurocomputing* 169 (2015) 246–261.
- [3] M. Á. Carreira-Perpiñán, The elastic embedding algorithm for dimensionality reduction, in: *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, June 21-24, 2010, Haifa, Israel, 2010, pp. 167–174.
- 600 [4] K. Benabdeslem, M. Hindawi, Efficient semi-supervised feature selection: Constraint, relevance, and redundancy, *IEEE Trans. Knowl. Data Eng.* 26 (5) (2014) 1131–1143.
- [5] C. Wang, M. Shao, Q. He, Y. Qian, Y. Qi, Feature subset selection based on fuzzy neighborhood rough sets, *Knowl.-Based Syst.* 111 (2016) 173–179.
- 605 [6] M. Malekipirbazari, V. Aksakalli, Risk assessment in social lending via random forests, *Expert Syst. Appl.* 42 (10) (2015) 4621–4631.
- [7] S. Zheng, W. Liu, An experimental comparison of gene selection by lasso and dantzig selector for cancer classification, *Comp. in Bio. and Med.* 41 (11) (2011) 1033–1040.
- 610 [8] C. Hou, F. Nie, X. Li, D. Yi, Y. Wu, Joint embedding learning and sparse regression: A framework for unsupervised feature selection, *IEEE Trans. Cybernetics* 44 (6) (2014) 793–804.

- [9] T. Naghibi, S. Hoffmann, B. Pfister, A semidefinite programming based search strategy for feature selection with mutual information measure, *IEEE Trans. Pattern Anal. Mach. Intell.* 37 (8) (2015) 1529–1541.
- [10] M. E. ElAlami, A filter model for feature subset selection based on genetic algorithm, *Knowl.-Based Syst.* 22 (5) (2009) 356–362.
- [11] A. Wang, N. An, G. Chen, L. Li, G. Alterovitz, Accelerating wrapper-based feature selection with k-nearest-neighbor, *Knowl.-Based Syst.* 83 (2015) 81–91.
- [12] W. Qian, W. Shu, Mutual information criterion for feature selection from incomplete data, *Neurocomputing* 168 (2015) 210–220.
- [13] M. A. Hall, Correlation-based feature selection for discrete and numeric class machine learning, in: *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford University, Stanford, CA, USA, June 29 - July 2, 2000, 2000, pp. 359–366.
- [14] M. Dash, H. Liu, Consistency-based search in feature selection, *Artif. Intell.* 151 (1-2) (2003) 155–176.
- [15] X. He, D. Cai, P. Niyogi, Laplacian score for feature selection, in: *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*, 2005, pp. 507–514.
- [16] G. Herman, B. Zhang, Y. Wang, G. Ye, F. Chen, Mutual information-based method for selecting informative feature sets, *Pattern Recognition* 46 (12) (2013) 3315–3327.
- [17] F. Fleuret, Fast binary feature selection with conditional mutual information, *Journal of Machine Learning Research* 5 (2004) 1531–1555.
- [18] R. Battiti, Using mutual information for selecting features in supervised neural net learning, *IEEE Trans. Neural Networks* 5 (4) (1994) 537–550.

- [19] N. Kwak, C. Choi, Input feature selection by mutual information based on parzen
640 window, *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (12) (2002) 1667–1671.
- [20] H. Peng, F. Long, C. H. Q. Ding, Feature selection based on mutual information:
Criteria of max-dependency, max-relevance, and min-redundancy, *IEEE Trans.*
Pattern Anal. Mach. Intell. 27 (8) (2005) 1226–1238.
- [21] P. A. Estévez, M. Tesmer, C. A. Perez, J. M. Zurada, Normalized mutual infor-
645 mation feature selection, *IEEE Trans. Neural Networks* 20 (2) (2009) 189–201.
- [22] G. Brown, A new perspective for information theoretic feature selection, in: *Pro-*
ceedings of AISTATS, 2009, pp. 49–56.
- [23] S. Liu, H. Liu, L. J. Latecki, S. Yan, C. Xu, H. Lu, Size adaptive selection of most
informative features, in: *Proceedings of the Twenty-Fifth AAAI Conference on*
650 *Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11,*
2011, 2011.
- [24] Z. Zhang, E. R. Hancock, Hypergraph based information-theoretic feature selec-
tion, *Pattern Recognition Letters* 33 (15) (2012) 1991–1999.
- [25] M. Pavan, M. Pelillo, Dominant sets and pairwise clustering, *IEEE Trans. Pattern*
655 *Anal. Mach. Intell.* 29 (1) (2007) 167–172.
- [26] L. Cui, L. Bai, Y. Wang, X. Bai, Z. Zhang, E. R. Hancock, P2P lending analysis
using the most relevant graph-based features, in: *Proceedings of S+SSPR 2016,*
2016, pp. 3–14.
- [27] L. Cui, Y. Jiao, L. Bai, L. Rossi, E. R. Hancock, Adaptive feature selection based
660 on the most informative graph-based features, in: *Graph-Based Representation-*
s in Pattern Recognition - 11th IAPR-TC-15 International Workshop, GbRPR
2017, Anacapri, Italy, May 16-18, 2017, *Proceedings, 2017*, pp. 276–287.
- [28] A. Feragen, N. Kasenburg, J. Petersen, M. de Bruijne, K. M. Borgwardt, Scalable
kernels for graphs with continuous attributes, in: *Advances in Neural Information*

- 665 Processing Systems 26: 27th Annual Conference on Neural Information Process-
ing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake
Tahoe, Nevada, United States., 2013, pp. 216–224.
- [29] H. Yan, J. Yang, Sparse discriminative feature selection, *Pattern Recognition*
48 (5) (2015) 1827–1835.
- 670 [30] Z. Zhang, Y. Tian, L. Bai, J. Xiahou, E. R. Hancock, High-order covariate inter-
acted lasso for feature selection, *Pattern Recognition Letters* 87 (2017) 139–146.
- [31] R. Tibshirani, Regression shrinkage and selection via the lasso, *Journal of the
Royal Statistical Society, Series B* 58 (1) (1996) 267–288.
- [32] H. Zou, T. Hastie, Regularization and variable selection via the elastic net, *Journal*
675 *of the Royal Statistical Society* 67 (5) (2005) 301–320.
- [33] S. Ma, X. Song, J. Huang, Supervised group lasso with applications to microarray
data analysis, *BMC Bioinformatics* 8.
- [34] Y. Panagakis, C. Kotropoulos, Elastic net subspace clustering applied to pop/rock
music structure analysis, *Pattern Recognition Letters* 38 (2014) 46–53.
- 680 [35] B. J. Marafino, W. J. Boscardin, R. A. Dudley, Efficient and sparse feature selec-
tion for biomedical text classification via the elastic net: Application to ICU risk
stratification from nursing notes, *Journal of Biomedical Informatics* 54 (2015)
114–120.
- [36] G. Brown, A. C. Pock, M. Zhao, M. Luján, Conditional likelihood maximisa-
685 tion: A unifying framework for information theoretic feature selection, *Journal of
Machine Learning Research* 13 (2012) 27–66.
- [37] J. R. Vergara, P. A. Estévez, A review of feature selection methods based on
mutual information, *Neural Computing and Applications* 24 (1) (2014) 175–186.
- [38] V. Bolón-Canedo, N. Sánchez-Marroño, A. Alonso-Betanzos, Recent advances
690 and emerging challenges of feature selection in the context of big data, *Knowl.-
Based Syst.* 86 (2015) 33–45.

- [39] L. Bai, L. Rossi, H. Bunke, E. R. Hancock, Attributed graph kernels using the jensen-tsallis q-differences, in: Proceedings of ECML-PKDD, 2014, pp. 99–114.
- [40] S. P. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, *Foundations and Trends in Machine Learning* 3 (1) (2011) 1–122.
- [41] P. Giselsson, S. Boyd, Linear convergence and metric selection for douglas-rachford splitting and ADMM, *IEEE Trans. Automat. Contr.* 62 (2) (2017) 532–544.
- [42] G. de Carvalho Bento, O. P. Ferreira, J. G. Melo, Iteration-complexity of gradient, subgradient and proximal point methods on riemannian manifolds, *J. Optimization Theory and Applications* 173 (2) (2017) 548–562.
- [43] S. Magnússon, P. C. Weeraddana, M. G. Rabbat, C. Fischione, On the convergence of alternating direction lagrangian methods for nonconvex structured optimization problems, *IEEE Trans. Control of Network Systems* 3 (3) (2016) 296–309.
- [44] J. J. Hull, A database for handwritten text recognition research, *IEEE Trans. Pattern Anal. Mach. Intell.* 16 (5) (1994) 550–554.
- [45] A. Frank, A. Asuncion, Uci machine learning repository.
- [46] A. S. Georghiades, P. N. Belhumeur, D. J. Kriegman, From few to many: Illumination cone models for face recognition under variable lighting and pose, *IEEE Trans. Pattern Anal. Mach. Intell.* 23 (6) (2001) 643–660.
- [47] N. X. Vinh, S. Zhou, J. Chan, J. Bailey, Can high-order dependencies improve mutual information based feature selection?, *Pattern Recognition* 53 (2016) 46–58.
- [48] S. Chen, C. H. Q. Ding, B. Luo, Y. Xie, Uncorrelated lasso, in: Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA., 2013.

- 720
- [49] R. Tibshirani, M. A. Saunders, S. Rosset, K. Knight, Sparsity and smoothness via the fused lasso, *Journal of the Royal Statistical Society Series B (Statistical Methodology)* 67 (1) (2005) 91–108.
- [50] L. Bai, L. Rossi, A. Torsello, E. R. Hancock, A quantum jensen-shannon graph kernel for unattributed graphs, *Pattern Recognition* 48 (2) (2015) 344–355.