

Unsupervised Learning of Prototypes and Attribute Weights

Summary

We introduce a new approach, called Simultaneous Clustering and Attribute Discrimination (SCAD), that performs clustering and feature weighting *simultaneously* in an unsupervised manner. The proposed approach is computationally and implementation ally simple, and learns a different set of feature weights for each identified cluster. The cluster dependent feature weights offer two advantages. First, they guide the clustering process to partition the data set into more meaningful clusters. Second, they can be used in the subsequent steps of a learning system to improve its learning behavior. Moreover, SCAD can adapt to the variations that exist within a data set by categorizing it into distinct clusters, which are allowed to overlap because of the use of fuzzy membership degrees.

Based on the SCAD approach, we present two new clustering algorithms, called SCAD1 and SCAD2 respectively. These two algorithms achieve the same goal, however, they minimize different objective functions. SCAD1 borrows more from variance-based (or distance-based) approaches, while SCAD2 uses a regularization term in the objective function that results in a tunable feature discrimination parameter. Both SCAD1 and SCAD2 minimize one objective function for both the optimal prototype parameters and feature weights for each cluster. This optimization is done iteratively by dynamically updating the prototype parameters and the feature weights in each iteration. The optimization is done based on gradient descent, which makes the proposed clustering approach simple and fast. Since the objective functions of SCAD1 and SCAD2 are based on that of the Fuzzy C-Means (FCM), they inherit most of the advantages of FCM-type clustering algorithms, such as ease of computation, and simplicity.

Since the number of categories in a data set is not always known a priori, we present an extension of SCAD, which can determine the “optimal” number of clusters. The extension is based on Competitive Agglomeration, whereby the number of clusters is over specified, and adjacent clusters are allowed to compete for data points in a manner that causes clusters that lose in the competition to gradually become depleted and vanish.

We illustrate the performance of the proposed approach for the case of supervised and unsupervised learning. In the unsupervised learning application, color image segmentation, we show that SCAD can identify homogeneous regions within an image while learning the relevant descriptors of each region. In the supervised learning application, building a nearest prototype classifier, we use SCAD to model the class distribution by several clusters (i.e., sub-classes) and determine the optimal feature weights within each one of those clusters in an unsupervised fashion. In both applications, we compare our results with those generated by a similar clustering algorithm that treats all features as equally important.

Contributions

- The basic idea of using one objective function to identify clusters and their relevant feature weights have been published by the same authors as a short conference paper (“Simultaneous Clustering and Attribute Discrimination”, IEEE International Conference on Fuzzy Systems, San Antonio, 2000).
- The application of the above approach to clustering a collection of documents and identifying relevant keywords for each cluster of documents has been presented in a workshop (“Clustering with Adaptive Attribute Weighting for text documents”, Text Mining Workshop, held in conjunction with the second SIAM conf. on Data Mining, 2002).
- The major contributions of this manuscript include:
 1. A more detailed derivation and analysis of the algorithm (SCAD1)
 2. An alternative formulation of the objective function, and the derivation of a different algorithm that accomplishes a similar task (SCAD2).
 3. Extension of SCAD2 to deal with the case where the number of clusters is unknown (SCAD-CA).
 4. Unsupervised learning application: Color image segmentation.
 5. Supervised learning application: building a nearest prototype classifier.
 6. For both applications, we compare the performance of the proposed approach with those generated by a similar clustering algorithm that treats all features as equally important.

Unsupervised Learning of Prototypes and Attribute Weights

Hichem Frigui

Department of Electrical and Computer Engineering
The University of Memphis
206 Engineering Science Bldg., Memphis, TN 38152

phone: (901) 678-4948 fax: (901) 678-5469
hfrigui@memphis.edu

Olfa Nasraoui

Department of Electrical and Computer Engineering
The University of Memphis
206 Engineering Science Bldg., Memphis, TN 38152

phone: (901) 678-3251 fax: (901) 678-5469
onasraou@memphis.edu

Unsupervised Learning of Prototypes and Attribute Weights

Abstract

In this paper, we introduce new algorithms that perform clustering and feature weighting simultaneously and in an unsupervised manner. The proposed algorithms are computationally and implementation ally simple, and learn a different set of feature weights for each identified cluster. The cluster dependent feature weights offer two advantages. First, they guide the clustering process to partition the data set into more meaningful clusters. Second, they can be used in the subsequent steps of a learning system to improve its learning behavior. An extension of the algorithm to deal with an unknown number of clusters is also proposed. The extension is based on competitive agglomeration, whereby the number of clusters is over specified, and adjacent clusters are allowed to compete for data points in a manner that causes clusters which lose in the competition to gradually become depleted and vanish. We illustrate the performance of the proposed approach by using it to segment color images, and to build a nearest prototype classifier.

keywords: Feature weighting, fuzzy clustering, competitive agglomeration, image segmentation, nearest prototype classifier.

1 Introduction

The problem of selecting or weighting the best subset of attributes constitutes an important part of the design of good learning algorithms for real word tasks. The generalization performance of these algorithms can be significantly degraded if irrelevant attributes are used. As a result, several methods have been proposed for feature selection and weighting [1, 11, 16, 18]. In feature selection, the task's dimensionality is reduced by completely eliminating irrelevant features. This amounts to assigning binary relevance weights to the features (1 for relevant, and 0 for irrelevant). Feature weighting is an extension of the selection process where the features are assigned continuous weights which can be regarded as degrees of relevance. Continuous weighting provides a richer feature relevance representation. Therefore, it tends to outperform feature selection from an accuracy point of view in tasks where some features are useful but less important than others. Most feature weighting (and selection) methods assume that feature relevance is invariant over the task's domain, and hence learn a single set of weights for the entire data set. This assumption can impose unnecessary and pernicious constraints on the learning task when the data set is made of different categories or classes. If the data is already labeled (supervised learning), then it is possible to learn a different set of weights for each class. For instance, in [5], context-sensitive feature selection is introduced as a mechanism to improve lazy supervised learning techniques. However, this technique requires labeled data, and relates relevance of a given feature in terms of the feature values of other features.

On the other hand, if the data set is unlabeled (unsupervised learning), then no methods exist for assigning different relevance weights for the different categories of a data set prior to clustering. The classical approach in this case is to determine a single subset of features or feature weights for the entire unlabeled data set prior to clustering. However, by ignoring the existence of different sub-structures in the data set, which require different subsets of feature weights, the performance of any clustering procedure can be severely degraded. Hence, it is clear that the clustering and feature selection/weighting steps are coupled, and applying these steps in sequence can degrade the performance of the learning system. In fact, even if the data set is labeled according to several

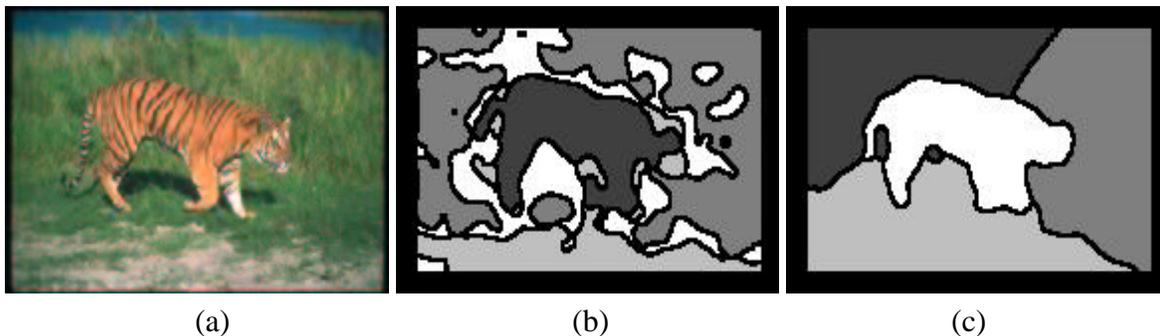


Figure 1. (a) Original color image, (b) results of the FCM with no position features, (c) results of the FCM with position features

known classes, it is preferable to model each complex (non-convex) class by several simple sub-classes or clusters, and to use a different set of feature weights for each cluster.

To illustrate the need for different sets of feature weights for different clusters, we consider the unsupervised learning problem of segmenting the color image shown in Fig. 1(a). In this problem, the pixels of the image must be categorized into meaningful clusters corresponding to different homogeneous regions of the image without any prior knowledge about its contents. In order to be categorized, the pixels are first mapped to feature vectors as explained in Section 5.1. Fig. 1(b) and (c) show the segmentation results with different feature subsets using the Fuzzy C-Means (FCM) algorithm [2], when the number of clusters is fixed to 4. The FCM, like other clustering algorithms, has no provision for cluster dependent feature weights. Any given feature must either be used (completely relevant) or ignored (irrelevant) for all clusters. Fig. 1(b) shows the segmentation results when the x and y position features are not used during clustering. It can be seen that some clusters suffer from a fragmentation phenomenon. That is, instead of being clean and compact, these clusters are scattered around the image. Fig. 1(c) shows the results when the x and y position features are used. As expected, the clusters obtained are more compact. However, except for the tiger cluster, the remaining clusters do not correspond to homogeneous regions in the original image. This is because, for example for the dirt region, the y - position feature is relevant while the x - position feature is irrelevant since this region forms a narrow strip that extends over the entire width of the image. For the grass region, both the x and y position features are irrelevant

because grass can be found in almost all areas of the image. We conclude that in this case, the x and y position features are both useful, and thus relevant for delineating clusters corresponding to the compact regions of the image such as the tiger. However, that is not the case for non-compact regions such as the background. The remaining features cannot be easily visualized as in the case of position. However, their relevance is also expected to vary across the different regions of the image because of the inherent differences in their color and texture properties. This simple example shows how feature relevance can vary widely within the domain of a data set.

In light of our discussion, it is only natural to conclude that ideally, clustering and feature selection should be performed simultaneously in an unsupervised learning problem. In this paper, we propose a new approach, called Simultaneous Clustering and Attribute Discrimination (SCAD), that performs clustering and feature weighting *simultaneously*. When used in conjunction with a supervised or unsupervised learning system, SCAD offers several advantages. First, it uses continuous feature weighting, hence providing a much richer feature relevance representation than feature selection. Secondly, SCAD learns the feature relevance representation of each cluster independently and in an unsupervised manner. Moreover, SCAD can adapt to the variations that exist within a data set by categorizing it into distinct clusters, which are allowed to overlap because of the use of fuzzy membership degrees. Based on the SCAD approach, we will present two new clustering algorithms, called SCAD1 and SCAD2 respectively. These two algorithms achieve the same goal, however, they minimize different objective functions. Finally, because the number of categories in a data set is not always known a priori, we present an extension of SCAD that can determine the number of clusters using competitive agglomeration [6].

The rest of the paper is organized as follows. In section 2, we review existing approaches to feature selection/weighting, and prototype-based clustering algorithms. In section 3, we present our new algorithms: SCAD1 and SCAD2. In section 4, we extend SCAD2 to the case where the number of clusters is unknown. In section 5, we present two applications of the proposed algorithm. The first one involves the use of SCAD2 to segment color images in an unsupervised manner. The second application uses SCAD2 to build a nearest prototype classifier in a supervised

learning environment. Finally, section 6 contains the summary conclusions.

To illustrate the need for different sets of feature weights for different clusters, we consider the unsupervised learning problem of segmenting the color image shown in Fig. 1(a). In this problem, the pixels of the image must be categorized into meaningful clusters corresponding to different homogeneous regions of the image without any prior knowledge about its contents. In order to be categorized, the pixels are first mapped to feature vectors as explained in Section 5.1. Fig. 1(b) and (c) show the segmentation results with different feature subsets using the Fuzzy C-Means (FCM) algorithm [2], when the number of clusters is fixed to 4. The FCM, like other clustering algorithms, has no provision for cluster dependent feature weights. Any given feature must either be used (completely relevant) or ignored (irrelevant) for all clusters. Fig. 1(b) shows the segmentation results when the x and y position features are not used during clustering. It can be seen that some clusters suffer from a fragmentation phenomenon. That is, instead of being clean and compact, these clusters are scattered around the image. Fig. 1(c) shows the results when the x and y position features are used. As expected, the clusters obtained are more compact. However, except for the tiger cluster, the remaining clusters do not correspond to homogeneous regions in the original image. This is because, for example for the dirt region, the y - position feature is relevant while the x - position feature is irrelevant since this region forms a narrow strip that extends over the entire width of the image. For the grass region, both the x and y position features are irrelevant because grass can be found in almost all areas of the image. We conclude that in this case, the x and y position features are both useful, and thus relevant for delineating clusters corresponding to the compact regions of the image such as the tiger. However, that is not the case for non-compact regions such as the background. The remaining features cannot be easily visualized as in the case of position. However, their relevance is also expected to vary across the different regions of the image because of the inherent differences in their color and texture properties. This simple example shows how feature relevance can vary widely within the domain of a data set.

In light of our discussion, it is only natural to conclude that ideally, clustering and feature selection should be performed simultaneously in an unsupervised learning problem. In this paper, we

propose a new approach, called Simultaneous Clustering and Attribute Discrimination (SCAD), that performs clustering and feature weighting *simultaneously*. When used in conjunction with a supervised or unsupervised learning system, SCAD offers several advantages. First, it uses continuous feature weighting, hence providing a much richer feature relevance representation than feature selection. Secondly, SCAD learns the feature relevance representation of each cluster independently and in an unsupervised manner. Moreover, SCAD can adapt to the variations that exist within a data set by categorizing it into distinct clusters, which are allowed to overlap because of the use of fuzzy membership degrees. Based on the SCAD approach, we will present two new clustering algorithms, called SCAD1 and SCAD2 respectively. These two algorithms achieve the same goal, however, they minimize different objective functions. Finally, because the number of categories in a data set is not always known a priori, we present an extension of SCAD which can determine the number of clusters using competitive agglomeration [6].

The rest of the paper is organized as follows. In section 2, we review existing approaches to feature selection/weighting, and prototype-based clustering algorithms. In section 3, we present our new algorithms: SCAD1 and SCAD2. In section 4, we extend SCAD2 to the case where the number of clusters is unknown. In section 5, we present two applications of the proposed algorithm. The first one involves the use of SCAD2 to segment color images in an unsupervised manner. The second application uses SCAD2 to build a nearest prototype classifier in a supervised learning environment. Finally, section 6 contains the summary conclusions.

2 Background

2.1 Feature Selection and Feature Weighting

Feature selection and weighting techniques generally rely on a criterion function and a search strategy. The criterion function is used to decide whether a feature subset is better than another, while the search strategy determines feature subset candidates. Depending on the criterion function used, there are two types of feature selection/weighting methods: the *wrapper* and *filter* approaches [10]. The wrapper approach relies on feedback from the performance algorithm, such as classifier

accuracy, to learn feature weights or to decide whether a particular feature subset is superior to another. The filter approach optimizes a classifier independent criterion function. The wrapper approach tends to perform better, however it can cause over fitting [12]. Moreover, it should only be applied in combination with classifiers of low complexity to limit its computational cost.

Feature selection methods have exploited several search strategies. The most rudimentary strategy, exhaustive search, considers $2^n - 1$ (where n is the maximum number of features) possible feature subsets, and is impractical for large n . As a result, other optimized search strategies such as forward selection and backward selection [8] can be used. These approaches are computationally feasible. Unfortunately, they are only appropriate for binary weighting and are prone to yielding sub optimal solutions [10]. Other strategies include random mutation hill climbing [17], and parallel search [15].

Feature selection/weighting can either be applied to the entire data set to obtain a single set of features or weights, or to each class independently to obtain a different set of features or weights for each class. It is also possible to use the entire data set and the class labels to learn a different set of weights for each class [13, 17]. In general, class dependent feature selection/weighting is superior to the approach yielding a single set of features/weights. Unfortunately, class dependent feature selection/weighting cannot be applied when the data set is unlabeled.

2.2 Prototype-Based Clustering

Let $\mathcal{X} = \{\mathbf{x}_j \mid j = 1, \dots, N\}$ be a set of N feature vectors in an n -dimensional feature space. Let $\mathbf{B} = (\beta_1, \dots, \beta_c)$ represent a C -tuple of prototypes each of which characterizes one of the C clusters. Each β_i consists of a set of parameters. Let u_{ij} represent the grade of membership of feature point \mathbf{x}_j in cluster β_i . The $C \times N$ matrix $\mathbf{U} = [u_{ij}]$ is called a constrained fuzzy C -partition matrix if it satisfies the following conditions [2]

$$u_{ij} \in [0, 1] \quad \forall i; \quad 0 < \sum_{j=1}^N u_{ij} < N \quad \forall i, j; \quad \sum_{i=1}^C u_{ij} = 1 \quad \forall j. \quad (1)$$

The problem of fuzzily partitioning the feature vectors into C clusters can be formulated as the minimization of an objective function $J(\mathbf{B}, \mathbf{U}; \mathcal{X})$ of the form

$$J(\mathbf{B}, \mathbf{U}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^m d_{ij}^2, \quad (2)$$

subject to the constraint in (1). In the above equation, $m \in [1, \infty)$ is a weighting exponent called the fuzzifier, and d_{ij}^2 represents the distance from a feature point \mathbf{x}_j to the prototype β_i . Minimization of (2) with respect to \mathbf{U} subject to (1) gives us [2]

$$u_{ij} = \frac{1}{\sum_{k=1}^C \left(\frac{d_{kj}^2}{d_{ij}^2} \right)^{\frac{1}{m-1}}}. \quad (3)$$

Minimization of (2) with respect to \mathbf{B} varies according to the choice of the prototypes and the distance measure. For example, in the Fuzzy C-Means (FCM) algorithm, the Euclidean distance is used, and each of the prototypes is described by the cluster center \mathbf{c}_i , which may be updated in each iteration using [2]

$$\mathbf{c}_i = \frac{\sum_{j=1}^N u_{ij}^m \mathbf{x}_j}{\sum_{j=1}^N u_{ij}^m}. \quad (4)$$

2.3 Competitive Agglomeration

The objective function in (2), which is essentially the sum of (fuzzy) intra-cluster distances, has a monotonic tendency with respect to the number of clusters, C , and has the minimum value of zero when $C = N$. Therefore, it is not useful for the automatic determination of the ‘‘optimum’’ number of clusters, and C has to be specified *a priori*. The Competitive Agglomeration (CA) algorithm [7, 6] overcomes this drawback by adding a second regularization term to prevent over fitting the data set with too many prototypes. The CA algorithm starts by partitioning the data set into a large number of small clusters. As the algorithm progresses, adjacent clusters compete for data points, and clusters that lose the competition gradually vanish. The CA algorithm minimizes the following objective function

$$J_A(\mathbf{B}, \mathbf{U}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^2 d_{ij}^2 - \alpha \sum_{i=1}^C \left[\sum_{j=1}^N u_{ij} \right]^2, \quad \text{subject to } \sum_{i=1}^C u_{ij} = 1, \quad \forall j. \quad (5)$$

It should be noted that the number of clusters C in (5) is dynamically updated in the CA algorithm. The first term in (5) controls the shape and size of the clusters and encourages partitions with many clusters, while the second term penalizes solutions with a large number of clusters and encourages the agglomeration of clusters. When both components are combined and α is chosen properly, the final partition will minimize the sum of intra-cluster distances, while partitioning the data set into the smallest possible number of clusters. It can be shown [6] that the membership update equation for (5) is given by

$$u_{ij} = u_{ij}^{FCM} + u_{ij}^{Bias}, \quad (6)$$

where

$$u_{ij}^{FCM} = \frac{1/d_{ij}^2}{\sum_{k=1}^C 1/d_{kj}^2}, \quad (7)$$

and

$$u_{ij}^{Bias} = \frac{\alpha}{d_{ij}^2} (N_i - \bar{N}_j). \quad (8)$$

In (8),

$$N_i = \sum_{j=1}^N u_{ij},$$

is the cardinality of cluster i , and

$$\bar{N}_j = \frac{\sum_{k=1}^C 1/d_{kj}^2 N_k}{\sum_{k=1}^C 1/d_{kj}^2}$$

is a weighted average of the cardinalities of all clusters. The first term in (6) is the membership term in the FCM algorithm [2] (see equation (3)) which takes into account only the relative distances of the feature point to all clusters. The second term is a signed bias term which allows good clusters to agglomerate and spurious clusters to disintegrate.

The value of α needs to be initially small to encourage the formation of small clusters. Then, it should be increased gradually to promote agglomeration. After a few iterations, when the number of clusters becomes close to the ‘‘optimum’’, the value of α should again decay slowly to allow the algorithm to converge. In [6], α is updated in every iteration (t) using

$$\alpha(t) = \eta(t) \frac{\sum_{i=1}^C \sum_{j=1}^N (u_{ij}^{(t-1)})^2 (d_{ij}^2)^{(t-1)}}{\sum_{i=1}^C \left[\sum_{j=1}^N u_{ij}^{(t-1)} \right]^2}, \quad (9)$$

where

$$\eta(t) = \begin{cases} \eta_0 e^{-|t_0-t|/\tau} & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases} \quad (10)$$

In (10), η_0 is the initial value, τ is the time constant, and t_0 is the iteration number at which η starts to decrease. The superscript $(t - 1)$ is used on u_{ij} and d_{ij}^2 to denote their values in iteration $(t - 1)$.

3 Simultaneous Clustering and Attribute Discrimination

The proposed approach to perform clustering and attribute discrimination is designed to search for the optimal prototype parameters, \mathbf{B} , and the optimal set of feature weights, \mathbf{V} , simultaneously. Each cluster i is allowed to have its own set of feature weights $\mathbf{V}_i = [v_{i1}, \dots, v_{in}]$. We present two versions of the SCAD algorithm: SCAD1 tries to balance between the two terms of a compound objective function in order to determine the optimal attribute relevance weights, while SCAD2 minimizes a single term criterion.

3.1 SCAD1

The SCAD1 algorithm minimizes the following objective function:

$$J_1(\mathbf{C}, \mathbf{U}, \mathbf{V}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N u_{ij}^m \sum_{k=1}^n v_{ik} d_{ijk}^2 + \sum_{i=1}^C \delta_i \sum_{k=1}^n v_{ik}^2, \quad (11)$$

subject to the constraint on fuzzy memberships, u_{ij} in (1), and

$$v_{ik} \in [0, 1] \forall i, k; \quad \text{and} \quad \sum_{k=1}^n v_{ik} = 1, \forall i. \quad (12)$$

In (11), v_{ik} represents the relevance weight of feature k in cluster i , and d_{ijk} is given by

$$d_{ijk} = |x_{jk} - c_{ik}|, \quad (13)$$

where x_{jk} is the k^{th} feature value of data point \mathbf{x}_j , and c_{ik} is the k^{th} component of the i^{th} cluster center vector. In other words, d_{ijk} is the projection of the displacement vector between feature point \mathbf{x}_j and the i^{th} class center (\mathbf{c}_i) along the k^{th} dimension.

The objective function in (11) has two components. The first component, which is similar to the FCM objective function [2], is the sum of feature-weighted Euclidean distances to the prototypes, additionally weighted by constrained memberships. This component allows us to obtain compact clusters. From a feature-relevance point of view, this term is minimized when, in each cluster, only one feature is completely relevant, while all other features are irrelevant. The second component in equation (11) is the sum of the squared feature weights. The global minimum of this component is achieved when all the features are equally weighted. When both components are combined as in (11), and the coefficients δ_i are chosen properly, the final partition will minimize the sum of intra-cluster weighted distances, where the weights are optimized for each cluster.

To optimize J_1 , with respect to \mathbf{V} , we use the Lagrange multiplier technique, and obtain

$$J_1(\mathbf{\Lambda}, \mathbf{V}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^m \sum_{k=1}^n v_{ik} d_{ijk}^2 + \sum_{i=1}^C \delta_i \sum_{k=1}^n v_{ik}^2 - \sum_{i=1}^C \lambda_i \left(\sum_{k=1}^n v_{ik} - 1 \right),$$

where $\mathbf{\Lambda} = [\lambda_1, \dots, \lambda_c]^t$. Since the rows of \mathbf{V} are independent of each other, we can reduce the above optimization problem to the following C independent problems:

$$J_{1i}(\lambda_i, \mathbf{V}_i) = \sum_{j=1}^N (u_{ij})^m \sum_{k=1}^n v_{ik} (x_{jk} - c_{ik})^2 + \delta_i \sum_{k=1}^n v_{ik}^2 - \lambda_i \left(\sum_{k=1}^n v_{ik} - 1 \right) \text{ for } i = 1, \dots, C,$$

where \mathbf{V}_i is the i^{th} row of \mathbf{V} . By setting the gradient of J_{1i} to zero, we obtain

$$\frac{\partial J_{1i}(\lambda_i, \mathbf{V}_i)}{\partial \lambda_i} = \left(\sum_{k=1}^n v_{ik} - 1 \right) = 0, \quad (14)$$

and

$$\frac{\partial J_{1i}(\lambda_i, \mathbf{V}_i)}{\partial v_{ik}} = \sum_{j=1}^N (u_{ij})^m d_{ijk}^2 + 2\delta_i v_{ik} - \lambda_i = 0. \quad (15)$$

Solving (14) and (15) for v_{ik} , we obtain

$$v_{ik} = \frac{1}{n} + \frac{1}{2\delta_i} \sum_{j=1}^N (u_{ij})^m \left[\frac{\|\mathbf{x}_j - \mathbf{c}_i\|^2}{n} - d_{ijk}^2 \right]. \quad (16)$$

The first term in (16), $(1/n)$, is the default value if all attributes are treated equally, and no discrimination is performed. The second term is a bias that can be either positive or negative. It is positive for compact features where the distance along the corresponding dimension is, on average, less

than the average distance values along all the dimensions. In other words, if an attribute is very compact, compared to the other attributes, for most of the points that belong to a given cluster (high u_{ij}), then it is considered to be very relevant for that cluster.

The choice of δ_i in equation (11) is important to the performance of SCAD1 since it reflects the importance of the second term relative to the first term. If δ_i is too small, then the first term dominates, and only one feature in cluster i will be maximally relevant and assigned a weight of one, while the remaining features get assigned zero weights. On the other hand, if δ_i is too large, then the second term will dominate, and all features in cluster i will be relevant, and assigned equal weights of $1/n$. Hence, the values of δ_i should be chosen such that both terms are of the same order of magnitude. This can be accomplished by computing δ_i in iteration, t , as follows

$$\delta_i^{(t)} = K \frac{\sum_{j=1}^N (u_{ij}^{(t-1)})^m \sum_{k=1}^n v_{ik}^{(t-1)} (d_{ijk}^{(t-1)})^2}{\sum_{k=1}^n (v_{ik}^{(t-1)})^2}. \quad (17)$$

In (17), K is a constant, and the superscript $(t - 1)$ is used on u_{ij} , v_{ik} , and d_{ijk} to denote their values in iteration $(t - 1)$.

It should be noted that depending on the values of δ_i , the feature relevance values v_{ik} may not be confined to $[0,1]$. This inequality constraint could be added to the objective function in 11, and Karush-Kuhn-Tucker (KKT) Theorem [14] could be used to derive the necessary conditions. A much simpler and practical heuristic, and almost as effective as the KKT optimization conditions, would be to simply set negative values to zero, and to clip values that are greater than one to one. However, if this constraint is violated very often, then it is an indication that the value of δ is too small, and that it should be increased (increase K).

To minimize J_1 with respect to \mathbf{U} , we rewrite the objective function in (11) as

$$J_1(\mathbf{C}, \mathbf{U}, \mathbf{V}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N u_{ij}^m \tilde{d}_{ij}^2 + \sum_{i=1}^C \delta_i \sum_{k=1}^n v_{ik}^2, \quad (18)$$

where

$$\tilde{d}_{ij}^2 = \sum_{k=1}^n v_{ik} d_{ijk}^2 \quad (19)$$

is the weighted Euclidean distance. Since the second term in (18) does not depend on u_{ij} explicitly,

the update equation of the memberships is similar to that of the FCM (see equation (3)), i.e.,

$$u_{ij} = \frac{1}{\sum_{k=1}^C \left(\frac{\tilde{d}_{ij}^2}{\tilde{d}_{kj}^2} \right)^{\frac{1}{m-1}}}. \quad (20)$$

To minimize J_1 with respect to the centers, we fix \mathbf{U} and \mathbf{V} , and set the gradient to zero. We obtain

$$\frac{\partial J}{\partial c_{ik}} = -2 \sum_{j=1}^N (u_{ij})^m v_{ik} (x_{jk} - c_{ik}) = 0. \quad (21)$$

Solving (21) for c_{ik} , we obtain

$$c_{ik} = \frac{v_{ik} \sum_{j=1}^N (u_{ij})^m x_{jk}}{v_{ik} \sum_{j=1}^N (u_{ij})^m}. \quad (22)$$

There are two cases depending on the value of v_{ik} .

Case 1: $v_{ik} = 0$ In this case the k^{th} feature is completely irrelevant relative to the i^{th} cluster. Hence, regardless of the value of c_{ik} , the values of this feature will not contribute to the overall weighted distance computation. Therefore, in this situation, any arbitrary value can be chosen for c_{ik} . In practice, we set $c_{ik} = 0$.

Case 2: $v_{ik} \neq 0$ For the case when the k^{th} feature has some relevance to the i^{th} cluster, equation (22) reduces to

$$c_{ik} = \frac{\sum_{j=1}^N (u_{ij})^m x_{jk}}{\sum_{j=1}^N (u_{ij})^m}.$$

To summarize, the update equation for the centers is

$$c_{ik} = \begin{cases} 0 & \text{if } v_{ik} = 0, \\ \frac{\sum_{j=1}^N (u_{ij})^m x_{jk}}{\sum_{j=1}^N (u_{ij})^m} & \text{if } v_{ik} > 0 \end{cases} \quad (23)$$

The SCAD1 algorithm is summarized below.

Simultaneous Clustering and Attribute Discrimination: SCAD1

Fix the number of clusters C ;

Fix m , $m \in [1, \infty)$;

Initialize the centers;

Initialize the relevance weights to $1/n$;

Initialize the fuzzy partition matrix \mathbf{U} ;

REPEAT

 Compute d_{ijk}^2 for $1 \leq i \leq C$, $1 \leq j \leq N$, and $1 \leq k \leq n$;

 Update the relevance weights v_{ik} by using (16);

 Update the partition matrix $\mathbf{U}^{(k)}$ by using (20);

 Update the centers by using (23);

 Update δ_i by using (17);

UNTIL (centers stabilize);

3.2 SCAD2

The SCAD2 algorithm omits the second term from its objective function by incorporating a discriminant exponent, q , and minimizing

$$J_2(\mathbf{B}, \mathbf{U}, \mathbf{V}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^m \sum_{k=1}^n (v_{ik})^q d_{ijk}^2, \quad (24)$$

subject to (1), and

$$v_{ik} \in [0, 1] \forall i, k; \quad \text{and} \quad \sum_{k=1}^n v_{ik} = 1, \forall i. \quad (25)$$

To optimize J_2 with respect to \mathbf{V} , we use the Lagrange multiplier technique, and obtain

$$J_2(\mathbf{\Lambda}, \mathbf{V}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^m \sum_{k=1}^n (v_{ik})^q d_{ijk}^2 - \sum_{i=1}^C \lambda_i \left(\sum_{k=1}^n v_{ik} - 1 \right),$$

where $\mathbf{\Lambda} = [\lambda_1, \dots, \lambda_C]^t$ is a vector of Lagrange multipliers corresponding to the C constraints in (25). Since the rows of \mathbf{V} are independent of each other and d_{ijk}^2 is independent of \mathbf{V} , we can reduce the above optimization problem to the following C independent problems:

$$J_{2i}(\lambda_i, \mathbf{V}_i) = \sum_{j=1}^N (u_{ij})^m \sum_{k=1}^n (v_{ik})^q d_{ijk}^2 - \lambda_i \left(\sum_{k=1}^n v_{ik} - 1 \right) \quad \text{for } i = 1, \dots, C,$$

where \mathbf{V}_i is the i^{th} row of \mathbf{V} . By setting the gradient of J_i to zero, we obtain

$$\frac{\partial J_{2i}(\lambda_i, \mathbf{V}_i)}{\partial \lambda_i} = \left(\sum_{k=1}^n v_{ik} - 1 \right) = 0, \quad (26)$$

and

$$\frac{\partial J_i(\lambda_i, \mathbf{V}_i)}{\partial v_{it}} = q (v_{it})^{(q-1)} \sum_{j=1}^N (u_{ij})^m d_{ijt}^2 - \lambda_i = 0. \quad (27)$$

Equation (27) yields

$$v_{it} = \left[\frac{\lambda_i}{q \sum_{j=1}^N (u_{ij})^m d_{ijt}^2} \right]^{1/(q-1)}. \quad (28)$$

Substituting (28) back into (26), we obtain

$$\sum_{k=1}^n v_{ik} = \left[\lambda_i / q \right]^{1/(q-1)} \sum_{k=1}^n \left[\frac{1}{\sum_{j=1}^N (u_{ij})^m d_{ijk}^2} \right]^{1/(q-1)} = 1.$$

Thus,

$$\left[\lambda_i / q \right]^{1/(q-1)} = \frac{1}{\sum_{k=1}^n \left[\frac{1}{\sum_{j=1}^N (u_{ij})^m d_{ijk}^2} \right]^{1/(q-1)}}.$$

Substituting this expression back in (28), we obtain

$$v_{it} = \frac{\left[\frac{1}{\sum_{j=1}^N (u_{ij})^m d_{ijt}^2} \right]^{1/(q-1)}}{\sum_{k=1}^n \left[\frac{1}{\sum_{j=1}^N (u_{ij})^m d_{ijk}^2} \right]^{1/(q-1)}}. \quad (29)$$

Simplifying (28) and using k to represent the dimension, we obtain the following update equation

$$v_{ik} = \frac{1}{\sum_{t=1}^n \left(\tilde{D}_{ik} / \tilde{D}_{it} \right)^{1/(q-1)}}, \quad (30)$$

where

$$\tilde{D}_{ik} = \sum_{j=1}^N (u_{ij})^m d_{ijk}^2$$

can be interpreted as a measure of dispersion of the i^{th} cluster along the k^{th} dimension, and $\sum_{t=1}^n \tilde{D}_{it}$ can be viewed as the total dispersion of the i^{th} cluster (taking all dimensions into account). Hence v_{ik} is inversely related to the ratio of the dispersion along the k^{th} dimension to the total dispersion for the i^{th} cluster. This means that the more compact the i^{th} cluster is along the k^{th} dimension (smaller \tilde{D}_{ik}), the higher will the relevance weight, v_{ik} , be for the k^{th} feature.

Equation (30) may be likened to the estimation and use of a covariance matrix in an inner-product norm-induced metric in various clustering algorithms [9, 4]. In fact, if d_{ijk}^2 is defined to be the Euclidean distance between x_{jk} and c_{ik} , then \tilde{D}_{ik} becomes a measure of the fuzzy variance of

the i^{th} cluster along the k^{th} dimension, and SCAD2 becomes very similar to the Gustafson-Kessel (GK) algorithm [9] with a diagonal covariance matrix. However, the estimation of a covariance matrix relies on the assumption that the data has a multivariate Gaussian distribution. On the other hand, SCAD2 is free of any such assumptions when estimating the feature weights. This means that SCAD2 can incorporate more general dissimilarity measures.

The role of the attribute weight exponent, q , can be deduced from equation (30). It can be shown (see Appendix) that as q approaches 1, v_{ik} tends to take binary values, i. e.,

$$v_{ik} = \begin{cases} 1 & \text{if } \tilde{D}_{ik} = \min_{t=1}^n \tilde{D}_{it}, \\ 0 & \text{otherwise} \end{cases}$$

This case is analogous to the winner-take-all situation where the feature along which the i^{th} cluster is the most compact gets all the relevancy ($v_{ik} = 1$), while all other attributes get assigned zero relevance, and hence do not contribute to the distance or center computations. On the other hand, when q approaches infinity, it can easily be shown that

$$v_{ik} = 1/n.$$

This means that all attributes share the relevancy equally. This is equivalent to the situation where no feature selection/weighting takes place. As can be expected, for the case where q takes finite values in $(1, \infty)$, we obtain weights that provide a moderate level of feature discrimination. For this reason, we will refer to q as a “discrimination exponent”.

To minimize J_2 with respect to \mathbf{U} , we follow similar steps and obtain

$$u_{ij} = \frac{1}{\sum_{k=1}^C \left[\frac{\tilde{d}_{ij}^2}{d_{kj}^2} \right]^{1/(m-1)}}, \quad (31)$$

where

$$\tilde{d}_{ij}^2 = \sum_{k=1}^n (v_{ik})^q d_{ijk}^2. \quad (32)$$

Even though the discrimination exponent, q , is necessary in order to derive the update equation for the feature weights, v_{ik} , and is useful to control their degree of discrimination (see (30)), we have noticed that for the purpose of computing the fuzzy memberships, u_{ij} , it is preferable to set

this exponent to one, and simply use:

$$\tilde{d}_{ij}^2 = \sum_{k=1}^n v_{ik} d_{ijk}^2. \quad (33)$$

That is, we compute the distance measure as in SCAD1.

To minimize J_2 with respect to the centers, we fix \mathbf{U} and \mathbf{V} , and set the gradient to zero. We obtain

$$\frac{\partial J}{\partial c_{ik}} = -2 \sum_{j=1}^N (u_{ij})^m (v_{ik})^q (x_{jk} - c_{ik}) = 0. \quad (34)$$

Equation (34) is similar to (21), and its solution yields the same equation to update the centers as in SCAD1.

The SCAD2 algorithm is summarized below.

Simultaneous Clustering and Attribute Discrimination: SCAD2

Fix the number of clusters C ;

Fix the fuzzifier m , $m \in [1, \infty)$;

Fix the discrimination exponent q , $q \in [1, \infty)$;

Initialize the centers and the fuzzy partition matrix \mathbf{U} ;

Initialize all the relevance weights to $1/n$;

REPEAT

Compute d_{ijk}^2 for $1 \leq i \leq C$, $1 \leq j \leq N$, and $1 \leq k \leq n$;

Update the relevance weights matrix \mathbf{V} by using (30);

Compute \tilde{d}_{ij}^2 by using (33);

Update the partition matrix \mathbf{U} by using (31);

Update the centers by using (23);

UNTIL (*centers stabilize*);

We now illustrate the performance of SCAD1 and SCAD2 on a simple data set. Table 1 contains the coordinates of 20 points of two synthetic Gaussian clusters with $(\mu_1, \Sigma_1) = ([0, 0]^T, \mathbf{I}_2)$ and $(\mu_2, \Sigma_2) = ([5, 5]^T, \mathbf{I}_2)$. In this example, m , K (for SCAD1), and q (for SCAD2) were set to 2.0, and the centers and the fuzzy partition matrix were initialized by running the FCM algorithm for 2 iterations. SCAD1 converged after 4 iterations, and SCAD2 converged after 5 iterations giving the results displayed in Table 2 and Table 3 respectively. Since both features are relevant, SCAD1 and SCAD2 assigned high weights for both features, and the estimated center coordinates are close

to those of the actual centers. To demonstrate the ability of the proposed algorithms to cluster and identify relevant features, we increase the number of features to four by adding two irrelevant features to each cluster which are highlighted in Table 4. The first two features of the first cluster are uniformly distributed in the intervals $[0, 20]$ and $[0, 10]$ respectively. Features two and four of the second cluster are uniformly distributed in the intervals $[0, 10]$ and $[0, 5]$ respectively. A traditional feature selection algorithm can only discriminate against the second feature since it is irrelevant for both clusters. Clustering the remaining three features will not provide a compact description of each cluster. SCAD1 converged after 10 iterations, and SCAD2 converged after 5 iterations, and their results are displayed in Table 5 and Table 6 respectively. The first feature of the first cluster is correctly identified as irrelevant by both algorithms. The second feature of the first cluster is identified as irrelevant by SCAD2 (relevance=0.05), but was assigned a larger weight by SCAD1 (relevance=0.23) because it has a relatively smaller dynamic range. Feature four of the second cluster was not identified as totally irrelevant by both algorithms. This is because it has a dynamic range which is close to that of the actual features, and therefore this feature will be treated as almost equally important. Notice, however, that this feature was judged by SCAD2 to be more irrelevant compared to the remaining features than SCAD1.

We have used SCAD1 and SCAD2 to cluster several other data sets, and we have observed that they have similar behavior. Both algorithms succeed in identifying the relevant and irrelevant features, and assign similar weights. In the rest of this paper, we will extend only SCAD2 to the case of an unknown number of clusters, and illustrate its performance.

4 Simultaneous Clustering and Attribute Discrimination: Unknown Number of Clusters

The objective functions in (5) and (24) complement each other, and can easily be combined into one objective function. The algorithm that minimizes the resulting objective function would inherit the advantages of the CA and the SCAD2 algorithms. This algorithm, called SCAD2-CA,

Table 1. Two 2-dimensional Gaussian clusters

Cluster # 1		Cluster # 2	
x_1	x_2	x_1	x_2
-0.33	1.11	4.66	6.11
-2.02	-0.73	2.97	4.26
-0.33	0.72	4.66	5.72
-0.25	0.04	4.74	5.04
-1.08	-0.37	3.91	4.62
0.15	-0.36	5.15	4.63
-1.22	0.11	3.77	5.11
1.80	1.43	6.80	6.43
-1.48	-0.70	3.51	4.29
-0.87	1.02	4.12	6.02
-0.21	-0.45	4.78	4.54
-0.28	1.06	4.71	6.06
0.45	0.16	5.45	5.16
-2.29	1.98	2.70	6.98
0.84	-0.68	5.84	4.31
1.49	1.61	6.49	6.61
-0.23	0.31	4.76	5.31
-0.46	-0.82	4.53	4.17
-1.58	-1.09	3.41	3.90
0.72	1.27	5.72	6.27

Table 2. Results of SCAD1 on the data set in Table 1

	Cluster # 1		Cluster # 2	
Features	x_1	x_2	x_1	x_2
Centers	-0.4	0.24	4.65	5.27
Relevance Weights	0.49	0.51	0.48	0.52

Table 3. Results of SCAD2 on the data set in Table 1

	Cluster # 1		Cluster # 2	
Features	x_1	x_2	x_1	x_2
Centers	-0.37	0.27	4.64	5.28
Relevance Weights	0.43	0.57	0.43	0.57

minimizes the following objective function

$$J_3(\mathbf{B}, \mathbf{U}, \mathbf{V}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^2 \sum_{k=1}^n (v_{ik})^q d_{ijk}^2 - \alpha \sum_{i=1}^C \left[\sum_{j=1}^N u_{ij} \right]^2, \quad (35)$$

Table 4. Two 4-dimensional clusters

Cluster # 1				Cluster # 2			
x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4
19.00	2.09	-0.33	1.11	4.66	2.13	6.11	0.28
4.62	3.79	-2.02	-0.73	2.97	6.43	4.26	1.76
12.13	7.83	-0.33	0.72	4.66	3.20	5.72	4.06
9.71	6.80	-0.25	0.04	4.74	9.60	5.04	0.04
17.82	4.61	-1.08	-0.37	3.91	7.26	4.62	0.69
15.24	5.67	0.15	-0.36	5.15	4.11	4.63	1.01
9.12	7.94	-1.22	0.11	3.77	7.44	5.11	0.99
0.37	0.59	1.80	1.43	6.80	2.67	6.43	3.01
16.42	6.02	-1.48	-0.70	3.51	4.39	4.29	1.36
8.89	0.50	-0.87	1.02	4.12	9.33	6.02	0.99
12.30	4.15	-0.21	-0.45	4.78	6.83	4.54	0.07
15.83	3.05	-0.28	1.06	4.71	2.12	6.06	3.73
18.43	8.74	0.45	0.16	5.45	8.39	5.16	2.22
14.76	0.15	-2.29	1.98	2.74	6.28	6.98	4.65
3.52	4.98	0.84	-0.68	5.84	1.33	4.31	2.33
8.11	7.67	1.49	1.61	6.49	2.07	6.61	2.09
18.70	9.70	-0.23	0.31	4.76	6.07	5.31	4.23
18.33	9.90	-0.46	-0.82	4.53	6.29	4.17	2.62
8.20	7.88	-1.58	-1.09	3.41	3.70	3.90	1.01
17.87	4.38	0.72	1.27	5.72	5.75	6.27	3.36

Table 5. Results of SCAD1 on the data set in Table 4

	Cluster # 1				Cluster # 2			
Features	x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4
Centers	13.06	5.56	-0.32	0.22	4.67	5.17	5.19	2.08
Relevance Weights	0.00	0.23	0.38	0.40	0.28	0.16	0.29	0.27

Table 6. Results of SCAD2 on the data set in Table 4

	Cluster # 1				Cluster # 2			
Features	x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4
Centers	12.72	5.39	-0.40	0.26	4.62	5.26	5.26	2.03
Relevance Weights	0.02	0.05	0.40	0.53	0.32	0.06	0.42	0.20

subject to (25), and $\sum_{i=1}^C u_{ij} = 1, \forall j$. Since the second term in (35) does not depend on v_{ik} , minimization of (35) with respect to \mathbf{V} yields the same update equation for v_{ik} as in SCAD2 (see equation (30)).

To minimize (35) with respect to \mathbf{U} , we rewrite it as

$$J(\mathbf{B}, \mathbf{U}, \mathbf{V}; \mathcal{X}) = \sum_{i=1}^C \sum_{j=1}^N (u_{ij})^2 \tilde{d}_{ij}^2 - \alpha \sum_{i=1}^C \left[\sum_{j=1}^N u_{ij} \right]^2,$$

where \tilde{d}_{ij}^2 is given by (32). Since \tilde{d}_{ij}^2 does not depend directly on u_{ij} , it can be shown that the update equation for u_{ij} is the same as the one for the CA algorithm if d_{ij}^2 is replaced with \tilde{d}_{ij}^2 in equations (7) and (8). Similarly, the agglomeration constant, α , can still be updated as in the CA algorithm using equation (9) and \tilde{d}_{ij}^2 instead of d_{ij}^2 .

5 Application1: Color Image Segmentation

We illustrate the ability of SCAD2 to perform clustering and learn cluster-dependent feature weighting by using it to segment real color images. We start by mapping each pixel in the original image to an 8-dimensional feature vector consisting of color, texture, and position features. These are the same features used by Carson et al. in their Blobword content-based image retrieval system [3]. In the following, we will first give a brief description of the features, then show the results of using SCAD2 to segment several images.

5.1 Feature Data Extraction

Texture Features First, the image $I(x, y)$ is convolved with Gaussian smoothing kernels $G_\sigma(x, y)$ of several scales, σ , as follows:

$$M_\sigma(x, y) = G_\sigma(x, y) * (\nabla I(x, y))(\nabla I(x, y))^t.$$

Then, the following three features are computed at each pixel location [3].

1. Polarity:

$$p = \frac{|E_+ - E_-|}{E_+ + E_-},$$

where E_+ and E_- represent the number of gradient vectors in the window $G_\sigma(x, y)$ that are on the positive and negative sides of the dominant orientation, respectively.

For each pixel, an optimal scale value is selected such that it corresponds to the value

where polarity stabilizes with respect to scale. Let p^* be the polarity at the selected scale.

2. Anisotropy:

$$a = 1 - \lambda_2/\lambda_1,$$

where λ_1 and λ_2 are the eigenvalues of $M_\sigma(x, y)$ at the selected scale.

3. Normalized texture contrast:

$$c = 2(\sqrt{\lambda_1 + \lambda_2})^3.$$

Since the anisotropy and polarity are meaningless in regions of low contrast, these two features are scaled by the contrast value to yield the texture feature vector $[ac, p^*c, c]$.

Color Features The three color features are the L*a*b* coordinates of the color image computed after smoothing the image with a Gaussian kernel at the selected optimal scale. Note that smoothing is performed to avoid over-segmentation of regions due to local color variations.

Position Features The (x, y) coordinates of each pixel are used to reduce over-segmentation of some regions, and to obtain smoother regions. In [3], the authors reported that including position yielded better segmentation results than excluding it.

5.2 Feature Data Clustering

The first color image to be segmented is shown in Fig. 2. This is the same image used to illustrate the need for adaptive relevance weights in Section 1 (See Fig. 1). As discussed in the introduction, the FCM succeeds in delineating the compact cluster corresponding to the tiger, but fails to correctly delineate the different regions in the background. This is because the FCM is unable to assign cluster dependent feature relevance weights. When the number of clusters is also fixed to 4, SCAD2 succeeds in simultaneously segmenting the image into 4 meaningful clusters as shown in Fig. 2(b), and in determining appropriate relevance weights for each cluster, as displayed in Table 7. Since a total of eight features are used in this experiment, the resulting feature weights

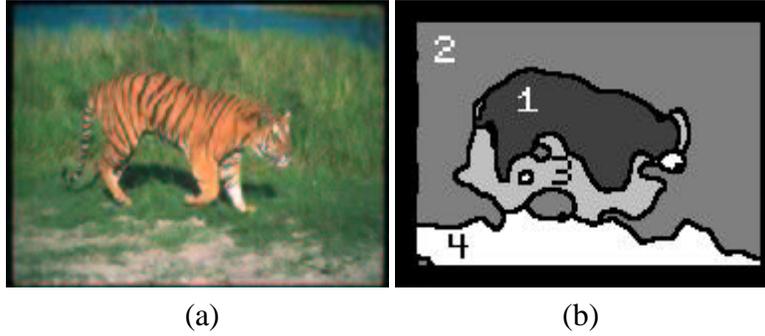


Figure 2. (a) Original color image, (b) results of SCAD2

should be compared to $1/8 \approx 0.12$, for the purpose of judging their relevancy. As can be seen, the x -position feature is found to be irrelevant ($v = 0.014$) for the “dirt” cluster (cluster #4 in Fig.2(b)), while the y -position ($v = 0.102$) is deemed relevant for this cluster. This is because the “dirt” region forms a narrow strip that extends over the entire width of the image. For the “grass” cluster (cluster #2 in Fig.2(b)), both the x and y features are found to be irrelevant because this region is scattered around the entire image. Similarly, the color and texture features receive unequal weights depending on the cluster. This is because the regions are not all characterized by the same features. For instance, the normalized texture contrast for the tiger region (cluster #1) was determined to be practically irrelevant ($v = 0.039$) because this feature has a large dynamic range. In fact, the value of this feature varies from small (in the uniformly textured areas) to large. Also, for the tiger cluster, the first texture feature, ac , was found to be practically irrelevant ($v = 0.029$). This means that the anisotropy feature ($a = 1 - \lambda_2/\lambda_1$) is irrelevant. This is because the orientation of the stripes varies significantly in some parts of the tiger cluster. In other words, some locations of this cluster have a dominant orientation ($\lambda_1 \gg \lambda_2$), while others do not.

Table 7. Feature relevance weights of the segmented objects in Fig. 1(b)

Clusters	Color Features			Texture Features			Position features	
	C_1	C_2	C_3	T_1	T_2	T_3	x	y
Cluster # 1	0.194	0.107	0.207	0.029	0.208	0.039	0.074	0.141
Cluster # 2	0.142	0.156	0.038	0.301	0.232	0.108	0.009	0.013
Cluster # 3	0.117	0.073	0.417	0.131	0.046	0.084	0.042	0.091
Cluster # 4	0.048	0.069	0.295	0.207	0.093	0.173	0.014	0.102

In the next experiment, we compare the performance of CA and SCAD2-CA to illustrate the importance of the simultaneous clustering and feature weighting mechanism when the number of clusters is unknown. The six color images in Fig. 3(a) are to be segmented into an unknown number of clusters. Both CA and SCAD2-CA were initialized by first dividing the image into a grid of 16 subimages. Then, the average of each feature within the i^{th} subimage is computed to yield the corresponding initial feature value of the i^{th} prototype. For the agglomeration coefficient, we set $(\eta_0, \tau, k_0) = (1.0, 10, 20)$. Fig. 3(b), displays the segmentation results using CA, and Fig. 3(c) displays the segmentation results using SCAD2-CA. The identified regions (i.e. clusters) are enumerated in Fig. 3(c), and their feature relevance weights are displayed in Fig. 3(d). As can be seen, the feature relevance weights differ not only from one feature to another, but also from one cluster to another.

As in the previous experiment, we notice that the SCAD2-CA performs better than CA, and it is able to yield meaningful clusters corresponding to both the objects in the foreground and the different areas of the background. This is because unlike CA, SCAD2-CA is capable of assigning distinct feature weights to different clusters. For example, the position features x and y are useful to delineate clusters corresponding to compact regions such as the animals in the several pictures. Hence, both algorithms are able to delineate these clusters. However, the x - position feature is irrelevant for most of the areas in the background such as the sky or the grass regions. Without an unsupervised feature weighting mechanism, the CA is unable to recognize this selective irrelevance. The same phenomenon can also be observed if an object in the foreground is too large. For example, the plane image in the last row of Fig. 3(b) gets over-segmented by CA into 2 clusters. In contrast, SCAD2-CA was capable of segmenting all the foreground and background regions in all the image in Fig. 3 into an appropriate number of clusters. These previous examples clearly demonstrate the effectiveness of SCAD2-CA in color image segmentation.

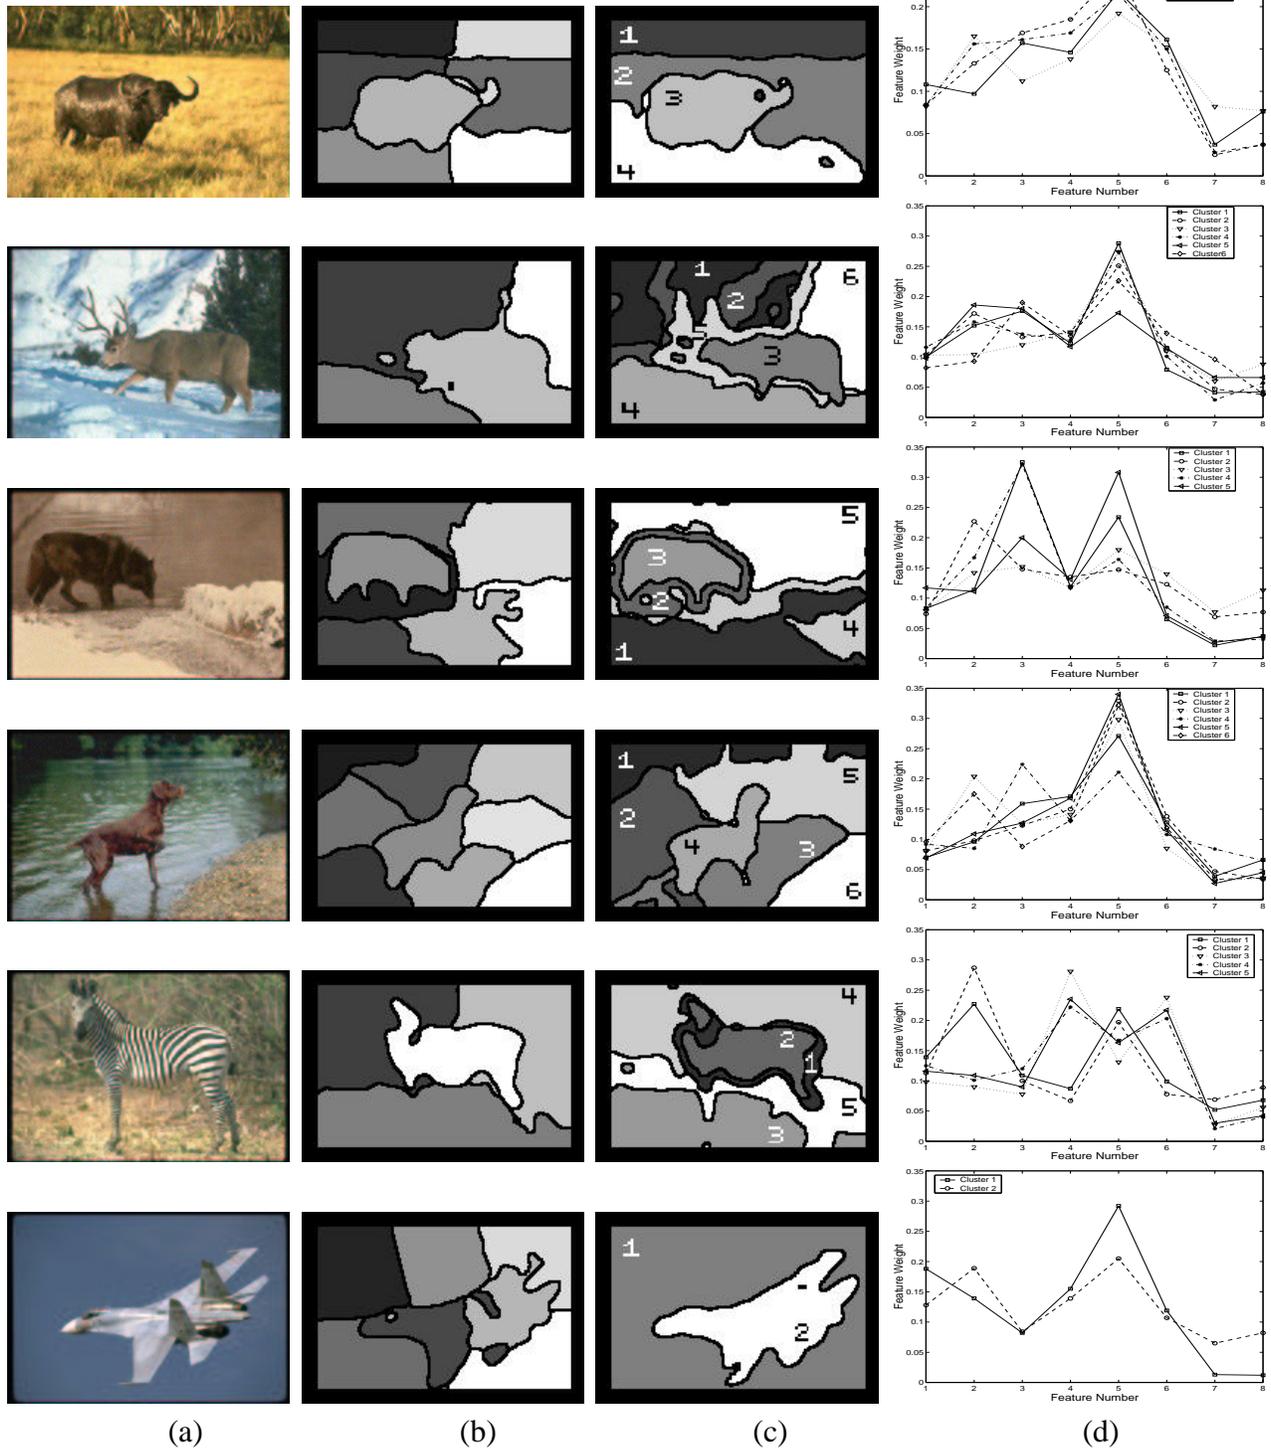


Figure 3. Image segmentation with CA and SCAD2-CA: (a) Original color images, (b) CA segmentation results, (c) SCAD2-CA segmentation results, and (d) Feature relevance weights for the regions labeled in (c)

5.3 Impact on Content-Based Image Retrieval Systems

Color image segmentation is an important step of many Content-Based Image Retrieval (CBIR) systems. When used for this task, SCAD2 and SCAD2-CA offer the following advantages.

Good segmentation As our results showed, the features are generally not equally relevant for all objects/scenes in an image. Hence, learning cluster-dependent relevance weights with SCAD2 will yield superior segmentation results. This translates into higher performance both in the indexing and retrieval phases of the CBIR system.

Compact and meaningful representation With SCAD2, each segmented region can automatically be represented by a different set of relevance weights. Therefore, there is no need for any user intervention to specify the feature relevance weights as in the case of most CBIR systems. This means that the user does not need to be concerned about the internal representation of the color image. This is particularly helpful for the “general” (non-specialized) user.

More efficient retrieval For the retrieval step, the segmented objects in the query image must be compared to the pre-segmented objects which are stored in the database. If all features are used indiscriminately, then the irrelevant features which can have unpredictable values, may have a contribution to the distance measure that is so large that even those images which are visually similar to the query image will not be retrieved. The cluster-dependent relevance weights provide an implicit mechanism for local similarity measures that take color, texture, and spatial context into account when comparing regions.

6 Application 2: Supervised Classification

In this application, we use SCAD2-CA for supervised classification. In this problem, the class labels for a set of (training) data are given, and the goal is to find the inter-class boundaries. Subsequently, any new (testing) data point can be classified into one of the known classes based on its position relative to the class boundaries. We will use SCAD2-CA for the construction of

the class boundaries by building a Multiple Prototype Classifier (MPC). We call this classifier SCAD2-MPC.

The MPC is a variation of the K-Nearest Neighbor (K-NN) approach which offers the advantage of reduced computational and storage costs. Instead of retaining all training instances in primary memory, only a few typical instances are stored and used to determine the class of a new instance. A smaller number of prototypes can achieve comparable or superior predictive accuracy because noisy instances will not be used as prototypes, and limiting the number of prototypes may avoid overfitting the training data. The number of prototypes that is required depends on the complexity of the distribution of each class. Most existing MPC algorithms fix the number of prototypes for each class, and treat all features as being equally important. Unfortunately, fixing the number of prototypes limits the complexity of the model of the problem being addressed. Moreover, equal weights make the classifier sensitive to the distance function. Existing feature weighting algorithms can only assign one set of feature weights for each class, and ignore the fact that different local regions (clusters) of complex classes are best characterized by different subsets of feature weights. In the following, we show that classification accuracy can be improved if we use SCAD2-CA to learn the optimal number of prototypes, and the optimal set of feature weights for each prototype.

Four standard data sets¹ were used to evaluate the performance of the SCAD2-MPC classifier. These are the Iris data set, the Wisconsin Breast Cancer data set, the Pima Indians Diabetes data set, and the Heart Disease data set. The Iris data set has 4 features and 3 classes, namely, "setosa", "versicolor", and "virginica". Each class has 50 samples. The Wisconsin Breast Cancer data set has 9 features and 2 classes, namely, benign and malignant. These 2 classes have 444 and 239 samples respectively. The Pima Indian Diabetes data set has 8 features and 2 classes. The first class, negative test, has 500 samples and the second class, positive test, has 268 samples. The Heart Disease data set has 13 features and 2 classes. The first class, absence of heart disease, has 150 samples, and the second class, presence of heart disease, has 120 samples. Since these data sets are not divided into training and testing sets, we use a 25% Jackknifing procedure. In other

¹Theses data sets are available via anonymous ftp: <ftp://ics.uci.edu>

words, the training data set is created by leaving out every fourth data point. These left out data points are later used for testing. This process is repeated 4 times, each time leaving out a different set of points. The reported results are averages of the four jackknifing experiments.

During training, the CA and SCAD2-CA algorithms used the same agglomeration values for all the data sets, namely, $(\eta_0, \tau, k_0) = (1.0, 10, 20)$. Since the data sets have different cardinalities and dimensions, we use different values of C_{max} (initial number of clusters) for the different classes of the data sets. This value is set automatically, depending on the cardinality (N) and the dimensionality (n) of the class being clustered, to $C_{max} = N/2n$.

Table 8. Average percentage of correct classification of 2 classifiers

Data Set	CA-NP		SCAD2-CA	
	Train	Test	Train	Test
Iris	95.11	92.67	95.11	94.00
Wisconsin Breast Cancer	95.90	95.46	96.83	96.78
Pima Indians Diabetes	74.04	71.74	76.30	74.87
Heart Disease	81.85	81.48	85.19	85.19

Table 8 shows the classification rates obtained using the CA-based classifier (CA-MPC), and the SCAD2-CA based classifier (SCAD2-MPC) for the four data sets used in this experiment. It can be seen that SCAD2-MPC results in better rates of correct classification for all data sets. Most importantly, the improvement is particularly significant when the features are assigned different weights depending on the clusters. This confirms the importance of the unsupervised feature weighting mechanism of SCAD2. For instance, the improvement in the correct classification rate of the “wisconsin breast cancer” data set is small, as expected, because the feature weights do not vary widely as can be seen in Table 9. On the other hand, a significant improvement can be observed for the “heart disease” data set because the feature weights determined by SCAD2-CA vary widely as shown in Table 10.

It should be noted that in addition to computing cluster-dependent feature weights, SCAD2-CA

determines the number of sub-classes or prototypes within each class in an unsupervised manner. For instance, Table 9 shows that for the “wisconsin breast cancer” data set, class 1 is probably compact since SCAD2-CA finds a single prototype in all four jackknives. On the other hand, class 2 may have a more complex distribution since it gets split into two clusters in all but one jackknife.

Table 9. Feature relevance weights of the “wis” data set

Jackknife	Class	Prototype	Attribute weights								
			1	2	3	4	5	6	7	8	9
1	1	1	0.05	0.14	0.11	0.10	0.12	0.08	0.09	0.11	0.21
	2	1	0.12	0.13	0.13	0.09	0.12	0.09	0.14	0.09	0.09
		2	0.11	0.11	0.12	0.08	0.12	0.09	0.13	0.08	0.16
2	1	1	0.05	0.11	0.11	0.13	0.11	0.07	0.09	0.11	0.21
	2	1	0.13	0.11	0.12	0.09	0.13	0.09	0.14	0.08	0.12
3	1	1	0.05	0.11	0.09	0.10	0.12	0.08	0.09	0.10	0.25
	2	1	0.11	0.11	0.12	0.08	0.13	0.08	0.13	0.08	0.15
		2	0.14	0.13	0.12	0.09	0.13	0.08	0.14	0.09	0.09
4	1	1	0.05	0.11	0.10	0.10	0.10	0.08	0.08	0.09	0.29
	2	1	0.11	0.11	0.12	0.08	0.12	0.09	0.13	0.08	0.15
		2	0.12	0.12	0.13	0.09	0.13	0.09	0.14	0.09	0.09

Table 10. Feature relevance weights of the “heart” data set

Jackknife	Class	Prototype	Attribute weights												
			1	2	3	4	5	6	7	8	9	10	11	12	13
1	1	1	0.09	0.03	0.05	0.12	0.13	0.04	0.03	0.13	0.04	0.19	0.05	0.07	0.03
	2	1	0.12	0.04	0.06	0.11	0.21	0.04	0.03	0.12	0.03	0.09	0.06	0.05	0.04
2	1	1	0.09	0.03	0.05	0.12	0.17	0.04	0.03	0.13	0.04	0.15	0.05	0.08	0.03
	2	1	0.12	0.04	0.07	0.11	0.22	0.04	0.03	0.11	0.03	0.10	0.06	0.05	0.03
3	1	1	0.09	0.03	0.05	0.13	0.16	0.04	0.03	0.13	0.04	0.15	0.05	0.08	0.04
	2	1	0.12	0.04	0.07	0.12	0.21	0.04	0.03	0.12	0.03	0.09	0.06	0.05	0.03
4	1	1	0.09	0.02	0.05	0.12	0.18	0.04	0.03	0.13	0.04	0.13	0.05	0.09	0.04
	2	1	0.12	0.04	0.07	0.11	0.22	0.06	0.03	0.11	0.03	0.09	0.06	0.05	0.03

7 Conclusion

In this paper, we presented a new approach to perform clustering and feature weighting simultaneously. When used as part of an unsupervised learning system, SCAD1 and SCAD2 can

categorize the unlabeled data while determining the best feature weights within each cluster. They can also enhance supervised learning by modeling complex classes by several clusters and determining the optimal feature weights within each one of those clusters in an unsupervised fashion. In addition to the above advantages, SCAD2-CA can determine the “optimal” number of clusters automatically. Our experimental results on both synthetic and real examples showed that SCAD2 and SCAD2-CA could successfully cluster data and determine cluster dependent feature weights simultaneously. Our results demonstrated that the proposed paradigms outperform similar clustering algorithms specifically because of their ability to determine cluster-dependent feature weights. Both SCAD1 and SCAD2 minimize one objective function for both the optimal prototype parameters and feature weights for each cluster. This optimization is done iteratively by dynamically updating the prototype parameters and the feature weights in each iteration. The optimization is done based on gradient descent, which makes the proposed clustering approach simple and fast.

The proposed approach to simultaneous clustering and attribute discrimination (especially SCAD2) may be likened to existing clustering algorithms that estimate and use a covariance matrix. However, unlike these algorithms, our approach does not assume that the data has a multivariate Gaussian distribution, and can incorporate a variety of distance measures.

Since the objective functions of SCAD1 and SCAD2 are based on that of the FCM, they inherit most of the advantages of FCM-type clustering algorithms, such as ease of computation, and simplicity. In the future, we plan to extend SCAD1 and SCAD2 to take into account possible correlation between features, as well as robustness to noise.

Appendix

theorem 1

$$\lim_{q \rightarrow 1^+} v_{ik} = \begin{cases} 1 & \text{if } \tilde{D}_{ik} = \min_{t=1}^n \tilde{D}_{it}, \\ 0 & \text{otherwise.} \end{cases}$$

Proof

The relevance weight of feature k in cluster i is defined as

$$v_{ik} = \frac{1}{\sum_{t=1}^n (\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)}}, \quad (36)$$

Since the exponent $\frac{1}{q-1} \rightarrow \infty$ as $q \rightarrow 1^+$, the individual terms in the denominator have the following tendency

$$\lim_{q \rightarrow 1^+} (\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)} = \begin{cases} \infty & \text{if } \tilde{D}_{ik} > \tilde{D}_{it}, \\ 0 & \text{if } \tilde{D}_{ik} < \tilde{D}_{it}, \\ 1 & \text{if } \tilde{D}_{ik} = \tilde{D}_{it}. \end{cases}$$

Therefore, two cases arise:

- **Case1:**

$$\exists t \neq k \mid \tilde{D}_{ik} > \tilde{D}_{it} \implies \tilde{D}_{ik} \neq \min_{t=1}^n \tilde{D}_{it}.$$

The denominator in (36) becomes infinite, and

$$\lim_{q \rightarrow 1^+} v_{ik} = 0.$$

- **Case2:**

$$\nexists t \neq k \mid \tilde{D}_{ik} > \tilde{D}_{it} \implies \tilde{D}_{ik} = \min_{t=1}^n \tilde{D}_{it}.$$

hence,

$$\begin{aligned} v_{ik} &= \frac{1}{(\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)} + \sum_{t=1, t \neq k}^n (\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)}} \\ &= \frac{1}{1 + \sum_{t=1, t \neq k}^n (\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)}}. \end{aligned}$$

Since $\lim_{q \rightarrow 1^+} (\tilde{D}_{ik}/\tilde{D}_{it})^{1/(q-1)} = 0$ for all $t \neq k$,

$$\lim_{q \rightarrow 1^+} v_{ik} = 1.$$

Therefore, we conclude that

$$\lim_{q \rightarrow 1^+} v_{ik} = \begin{cases} 1 & \text{if } \tilde{D}_{ik} = \min_{t=1}^n \tilde{D}_{it} , \\ 0 & \text{otherwise.} \end{cases}$$

References

- [1] H. Almuallim and T. G. Dietterich. Learning with many irrelevant features. In *Ninth National Conference on artificial intelligence*, pages 547–552, 1991.
- [2] J. C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York, 1981.
- [3] C. Carson, S. Belongie, H. Greenspan, and J. Malik. Color- and texture-based image segmentation using em and its application to image querying and classification. Submitted to: *IEEE Trans. on Pattern Analysis and Machine Intelligence*.
- [4] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society Series B*, 39(1):1–38, 1977.
- [5] P. Domingos. Contextsensitive feature selection for lazy learners, 1997.
- [6] H. Frigui and R. Krishnapuram. Clustering by competitive agglomeration. *Pattern Recognition*, 30(7):1223–1232, 1997.
- [7] H. Frigui and R. Krishnapuram. A robust competitive clustering algorithm with applications in computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):450–465, 1999.
- [8] K. Fukunaga. *Statistical Pattern Recognition*. Academic Press, San Diego California, 1990.
- [9] E. E. Gustafson and W. C. Kessel. Fuzzy clustering with a fuzzy covariance matrix. In *IEEE CDC*, pages 761–766, San Diego, California, 1979.
- [10] G. John, R. Kohavi, and K. Pfleger. Irrelevant features and the subset selection problem. In *Eleventh International Machine Learning Conference*, pages 121–129, 1994.

- [11] K. Kira and L. A. Rendell. The feature selection problem: Traditional methods and a new algorithm. In *Tenth National Conference on artificial intelligence*, pages 129–134, 1992.
- [12] R. Kohavi and D. Sommerfield. Feature subset selection using the wrapper model: Overfitting and dynamic search space topology. In *First International Conference on Knowledge Discovery and Data Mining*, pages 192–197, 1995.
- [13] I. Kononenko. Estimation attributes: Analysis and extensions of relief. In *European Conference on machine learning*, pages 171–182, 1994.
- [14] H.W. Kuhn and A.W. Tucker. Nonlinear programming. In J. Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492. 1950.
- [15] M. S. Lee and A. W. Moore. Efficient algorithms for minimizing cross validation error. In W. Cohne and H. Hirsh, editor, *Machine Learning: Proceedings of the eleventh international conference*, pages 190–198. 1994.
- [16] L. A. Rendell and K. Kira. A practical approach to feature selection. In *International Conference on machine learning*, pages 249–256, 1992.
- [17] D. Skalak. Prototype and feature selection by sampling and random mutation hill climbing algorithms. In *Eleventh International Machine Learning Conference (ICML-94)*, pages 293–301, 1994.
- [18] D. Wettschereck, D. W. Aha, and T. Mohri. A review and empirical evaluation of feature weighting methods for a class of lazy learning algorithms. *Artificial Intelligence Review*, 11:273–314, 1997.