

A Framework for Local Supervised Dimensionality Reduction of High Dimensional Data

Charu C. Aggarwal
IBM T. J. Watson Research Center
charu@us.ibm.com

Abstract

High dimensional data presents a challenge to the classification problem because of the difficulty in modeling the precise relationship between the large number of feature variables and the class variable. In such cases, it may be desirable to reduce the information to a small number of dimensions in order to improve the accuracy and effectiveness of the classification process. While data reduction has been a well studied problem for the unsupervised domain, the technique has not been explored quite as extensively for the supervised case. Existing techniques which try to perform dimensionality reduction are too slow for practical use in the high dimensional case. These techniques try to find *global discriminants* in the data. However, the behavior of the data often varies considerably with data locality and different subspaces may show better discrimination in different localities. This is an even more challenging task than the global discrimination problem because of the additional issue of data localization. In this paper, we propose the novel idea of supervised subspace sampling in order to create a reduced representation of the data for classification applications in an efficient and effective way. The method exploits the natural distribution of the different classes in order to sample the best subspaces for class discrimination. Because of its sampling approach, the procedure is extremely fast and scales almost linearly both with data set size and dimensionality.

Keywords: classification, dimensionality reduction

1 Introduction

The classification problem is defined as follows: We have a set of records \mathcal{D} containing the training data, and each record is associated with a class label. The classification problem constructs a model which connects the training data to the class variables. The classification problem is a widely studied problem by the data mining, statistics and the machine learning communities [6, 8, 14]. In

this paper, we will explore the dimensionality reduction problem in the context of classification.

Dimensionality Reduction Methods have been widely studied in the unsupervised domain [9, 13, 15]. The idea in dimensionality reduction methods is to transform the data into a new orthonormal coordinate system in which the second order correlations are eliminated. In typical applications, the resulting axis-system has the property that the variance of the data along many of the new dimensions is very small [13]. These dimensions can then be eliminated, a process resulting in a compact representation of the data with some loss of representational accuracy.

Dimensionality reduction has been studied somewhat sparingly in the supervised domain. This is partially because the presence of class labels significantly complicates the reduction process. It is more important to find a new set of dimensions in which the new axis-system retains the discriminatory behavior of the data, whereas the maintenance of representational accuracy becomes secondary. Aside from the advantages of a compact representation after reduction, dimensionality reduction also serves the dual purpose of removing the irrelevant subspaces in the data. This improves the accuracy of the classifiers on the reduced data. Some techniques such as those discussed in [8] achieve this by repeated discriminant computation. This is extremely expensive for high dimensional databases.

Most data reduction methods in the supervised and unsupervised domain use *global data reduction*. In these techniques, a single axis system is constructed on which the entire data is projected. Such techniques assume uniformity of class behavior throughout the data set while computing the new representation. Recent research for the unsupervised domain has shown that different parts of the data show considerably different behavior. As a result, while global dimensionality reduction often fails to capture the important characteristics of the data in a small number of dimensions, local dimensionality reduction methods [2, 5] can often provide

a more effective solution. In these techniques, a different axis system is constructed for each data locality for more effective reduction.

For the supervised domain, the analogous intuition is that different parts of the data may show different patterns of discrimination. However, since even global reduction methods such as the Fisher method are so computationally intensive, the task of effective local reduction becomes even more intractable. In this paper, we will show that this task can actually be accomplished quite efficiently by using a sampling process in which the random process of subspace selection is biased by the underlying class distribution. We will also show that the reduction process is very useful for the classification problem itself, since it facilitates the development of some interesting decomposable classification algorithms. The overall result is a greatly improved classification process.

The technique of subspace sampling [1, 11] has recently been used to perform data reduction in the unsupervised version of the problem. In this paper, we propose an effective subspace sampling approach for supervised problems. The aim is to exploit the data distribution in such a way that the axis system of representation in each data locality is able to expose the class discrimination. Since the class discrimination can be modeled with a small number of dimensions in many parts of the data, the resulting representation is often more concise and effective for the classification process.

We will also show how existing classification algorithms can be enhanced by the local reduction approach described in this paper. Since our reduction approach decomposes the data set by optimizing the discrimination behavior in each segment, different classification techniques may vary in effectiveness on different parts of the data. This fact can be exploited in order to ensure that the particular classification model being used is best suited to its data locality. To the best of our knowledge, the *decomposable* method discussed in this paper which picks an optimal classifier depending upon locality specific properties of the compressed data is unique in its approach.

In order to facilitate further development of the ideas, we will introduce additional notations and definitions. We assume that the data set is denoted by \mathcal{D} . The number of points in the data set is denoted by N and the dimensionality by d . The full dimensional data space is denoted by \mathcal{U} . We define the l -dimensional hyperplane $\mathcal{H}(\bar{y}, \mathcal{E})$ by an anchor \bar{y} and a mutually orthogonal set of vectors $\mathcal{E} = \{\bar{e}_1 \dots \bar{e}_l\}$. The hyperplane passes through \bar{y} , and the vectors in \mathcal{E} form the basis system for its subspace. The projection of a point \bar{x} onto this hyperplane is denoted by $\mathcal{P}(\bar{x}, \bar{y}, \mathcal{E})$ and is the closest

approximation of \bar{x} , which lies on this hyperplane. In order to find the value of $\mathcal{P}(\bar{x}, \bar{y}, \mathcal{E})$, we use \bar{y} as the reference point for the computation. Specifically, we determine the projections of $\bar{x} - \bar{y}$ onto the hyperplane defined by $\{\bar{e}_1 \dots \bar{e}_l\}$. Then, we translate the resulting point by the reference point \bar{y} . Therefore, we have:

$$(1.1) \quad \mathcal{P}(\bar{x}, \bar{y}, \mathcal{E}) = \bar{y} + \sum_{i=1}^l [(\bar{x} - \bar{y}) \cdot \bar{e}_i] \bar{e}_i$$

A pictorial representation of $\bar{x}' = \mathcal{P}(\bar{x}, \bar{y}, \mathcal{E})$ is illustrated in Figure 1(a). The value of \bar{x}' can be represented in the orthonormal axis system for \mathcal{E} with the use of only l coordinates $((\bar{x} - \bar{y}) \cdot \bar{e}_1 \dots (\bar{x} - \bar{y}) \cdot \bar{e}_l)$. This results in an additional overhead of storing \bar{y} and \mathcal{E} . This storage overhead is however not significant, if it can be averaged over a large number of points stored on this hyperplane. While the error of approximating \bar{x} with $\mathcal{P}(\bar{x}, \bar{y}, \mathcal{E})$ is given by the euclidean distance between \bar{x} and $\mathcal{P}(\bar{x}, \bar{y}, \mathcal{E})$, this measure is secondary to the classification accuracy of the reduced data.

This paper is organized as follows. In the next section, we will introduce the supervised subspace sampling technique and discuss some of its properties. Section 3 will discuss the application of the method to the classification problem. We will also discuss how the performance of the classification system can be considerably optimized by using the decomposition created by the supervised subspace sampling technique. The empirical results are discussed in section 4. Finally, we present the conclusions and summary in section 5.

1.1 Contributions of this paper The paper discusses a highly effective and scalable approach to the problem of supervised data reduction. While unsupervised data reduction has been well studied in the literature, the supervised problem is extremely difficult in practice because of the need to use the class distributions in the reduction process. The available discriminant methods are highly computationally intensive even on memory resident databases. In contrast, the technique discussed in this paper provides a significantly more effective reduction process, while exhibiting linear scalability with data set size and dimensionality because of its sampling approach. The process of sampling subspaces which are optimized to particular data localities results in a technique in which each segment of the data is more suited to the classification task. Furthermore, the unique decomposition created by the reduction process facilitates the creation of optimized decomposable approaches to the classification problem. Thus, the overall approach not only provides savings in terms of data compression, but also a greatly improved classification process.




Figure 1: Illustration of Localized Sampling




Figure 2: Subspace Tree (Example)

2 Supervised Subspace Sampling

An interesting approach for dimensionality reduction of the unsupervised version of the problem is that of random projections [11, 12]. In this class of techniques, we repeatedly sample spherically symmetric random directions in order to determine an optimum hyperplane on which the data is projected. These methods can also be extended to the classification problem by projection of the data onto random subspaces and measuring the discrimination of such spaces. However, such a direct extension of the random projection technique [11, 12] may often turn out to be ineffective in practice, since it is blind both to the data and class distributions of the points. We will try to explain this point by using 1-dimensional projections of 2-dimensional data. Consider the data set illustrated in Figure 1(b) in which we have illustrated two kinds of projections. In the left figure, the data *space* is sampled in order to find a 1-dimensional line along which the projection is performed. In data space sampling, random projections are chosen in a spherically symmetric fashion irrespective of the data distribution. The reduced data in this 1-dimensional representation is simply the projection of the data points onto the line. We note that such a pro-

jection neither follows the basic pattern in the data, nor does it provide a direction along which the class distributions are well discriminated. For high dimensional cases, such a projection may be poor at distinguishing the different classes even after repeated subspace sampling. In the other case of Figure 1(b), we have sampled the *points* in order to create a random projection. The sampled subspace is defined as the $(l - 1)$ -dimensional hyperplane containing l (locally proximate) points (of different classes) from the data.¹ The reason for picking points from different classes is that we would like the resulting subspace to represent the discrimination behavior of different classes more effectively. At the same time these points should be picked carefully only from a local segment of the data in order to ensure that the class discrimination is determined by data locality. For example, in Figure 1(b), the 1-dimensional line obtained by sampling two points of different classes picks the direction of greater discrimination most effectively than the space sampled projection in the same figure.

While it is intuitively clear that point sampling is more effective than space sampling for variance preservation, the advantages are limited when the data distribution varies considerably with locality. For example, in Figure 1(c), even the optimal 1-dimensional random projection cannot represent all points without losing a substantial amount of class discrimination. In fact, there is no 1-dimensional line along which a projection can effectively separate out the classes. In Figure 1(c), we have used the random projection technique locally in conjunction with data partitioning. In this technique, each data point is projected on the closest of a number of point sampled hyperplanes. In this case, it is evident that the data points are often well discriminated along each of the sampled hyperplanes, while they may be

¹The actual methodology of choosing the points is discussed at a later stage. At this point, we are concerned only with choosing the points in such a way that the chosen subspace is naturally biased by the original data distribution.

poorly discriminated at a global level. This is because the nature of the class distribution varies considerably with data locality, and a global dimensionality reduction method cannot reduce the representation below the original set of two dimensions. On the other hand, the 1-dimensional representation of the data created by local projection of the data points along the sampled lines represents the class discrimination very well.

It should be noted that the improvements of the localized subspace sampling technique come at the additional storage costs of the different hyperplanes. This limits the number of hyperplanes which can be retained from the sampling process, and requires us to make judicious choices in picking these hyperplanes. A second important issue is that even the implicit dimensionalities of the different data localities may be different. Therefore, we need a mechanism by which the sampling process is able to effectively choose hyperplanes of the lowest possible dimensionality for each data locality. This is an issue which we will discuss after developing some additional notational machinery:

DEFINITION 1. *Let $P = (\overline{x_1} \dots \overline{x_{l+1}})$ be a set of $(l + 1)$ linearly independent points. The representative hyperplane $\mathcal{R}(P)$ of P is defined as the l -dimensional hyperplane which passes through each of these $(l + 1)$ points.*

The hyperplane $R(P)$ can also be represented with the use of any point \overline{y} on the hyperplane, and an orthonormal set of vectors $\mathcal{E} = \{\overline{e_1} \dots \overline{e_l}\}$, which lie on the hyperplane. We shall call $(\overline{y}, \mathcal{E})$ the *axis* representation of the hyperplane, whereas the set P is referred to as the *point* representation. Thus, $\mathcal{R}(P)$ (point representation) is the same as $\mathcal{H}(\overline{y}, \mathcal{E})$ (axis representation). We note that there can be infinitely many point or axis representations of the same hyperplane. The axis representation is more useful for performing distance computations of the hyperplane from individual points in the database, whereas the point representation has advantages in storage efficiency in the context of a hierarchical arrangement of subspaces. We will discuss this issue in a later section.

2.1 The Supervised Subspace Tree The Supervised Subspace Tree is a conceptual organization of subspaces used in the data reduction process. This conceptual organization imposes a hierarchical arrangement of the subspaces of different dimensionalities. Each such subspace provides effective data discrimination which is specific to a particular locality of the data. Since the subspaces have different data dimensionalities, this results in a variable dimensionality decomposition of the data. The nodes at level- m in the subspace tree correspond to m -dimensional subspaces. The root node

corresponds to the null subspace. Thus, the dimensionality of the hyperplane for any node in the tree is determined by its depth. The subspace at a node is hierarchically related to that of its immediate parent. Each subspace other than the null subspace at the root is a 1-dimensional extension of its parent hyperplane. This 1-dimensional extension is obtained by adding a sampled data point to the representative set of the parent hyperplane. In order to elucidate the concept of a subspace tree, we will use an example. In Figure 2, we have illustrated a hierarchically arranged set of subspaces. The figure contains a two-level tree structure which corresponds to 1- and 2-dimensional subspaces. For each level-1 node in the tree, we store two points which correspond to the 1-dimensional line for that node. For each lower level node, we store an additional data point which increases the dimensionality of its parent subspace by 1. Therefore, a level- m node has a representative set of cardinality $(m + 1)$. For example, in the case of Figure 2, the node A in the subspace tree (with representative set $\{i_1, i_2\}$) corresponds to the 1-dimensional line defined by $\{i_1, i_2\}$. This node is extended to a 2-dimensional hyperplane in two possible ways corresponding to the nodes C and D . In each case, an extra point needs to be added to the representative set for creating the 1-dimensional extension. In order to extend to the 2-dimensional hyperplane for node C , we use the point i_5 , whereas in order to extend to the hyperplane for node D , we use the point i_6 . Note from Figure 2(a) that the intersection of the 2-dimensional hyperplanes C and D is the 1-dimensional line A . Thus, each node in the subspace tree corresponds to a hyperplane which is defined by its representative set drawn from the database \mathcal{D} . The representative set for a given hyperplane is obtained by adding one point to the representative set of its immediate parent. The subspace tree is formally defined as follows:

DEFINITION 2. *The subspace tree is a hierarchical arrangement of subspaces with the following properties: (1) Nodes at level- m correspond to m -dimensional hyperplanes (2) Nodes at level- $(m + 1)$ correspond to 1-dimensional extensions of their parent hyperplanes at level- m . (3) The point representative set of a level- $(m + 1)$ node is obtained by adding a sampled data point to the representative set of its m -dimensional parent subspace.*

The data points in \mathcal{D} are partitioned among the different nodes of the subspace tree. We note that since the hyperplane is a subspace of the full dimensional space, it has a lower dimensional axis system in terms of which the coordinates of \overline{x} are represented. Since the dimensionality of a hyperplane depends directly on the

distance of the node to the root, higher levels of the tree provide greater advantages in the reduction process.

2.2 Subspace Tree Construction Each node of the subspace tree corresponds to a hyperplane defined by the sequence of representative points sampled, starting from the root up to that node. The terms hyperplane and node are therefore used interchangeably through this paper.

In order to measure the quality of a given node N for the classification process, a *discrimination index* $\beta(N)$ is maintained along with each node. This discrimination index $\beta(N)$ always lies between 0 and 1, and is a measure of how well the different classes are separated out in the data set for node N . A value of 1 indicates perfect discrimination among the classes, whereas a value of 0 indicates very poor discrimination. We will discuss the methodology for computation of the discriminant in a later section.

The input to the subspace sampling algorithm for tree construction is the compression tolerance parameter ϵ , the data set \mathcal{D} , the maximum number of nodes L , a discrimination tolerance γ_1 , and a discrimination target γ_2 . The value of γ_2 is always larger than γ_1 . Each of these discrimination thresholds lie between 0 and 1. Intuitively, these discrimination thresholds impose a minimum and maximum threshold on the quality of class separation classes in the individual nodes. Correspondingly, each node N is classified into one of three types which is recorded by a variable called the *Status*(\cdot) vector: (1) *Default Node*: In this case, the discrimination index $\beta(N)$ lies between γ_1 and γ_2 . The value of *Status*(N) is set to 0. (2) *Discriminative Node*: Such a node is good for the classification process. The discrimination index $\beta(N)$ is larger than γ_2 in this case, the variable *Status*(N) is set to 2. (3) *Forbidden Node*: Such a node is bad for the classification process. In this case, the discriminant index $\beta(N)$ is smaller than γ_1 . The value of *Status*(N) is set to 1.

A top-down algorithm is used to construct the nodes of the subspace tree, and the data set \mathcal{D} is partitioned along this hierarchy in order to maximize the localized discrimination during the dimensionality reduction process. A *discrimination index* $\beta(N)$ is maintained with each node N in the tree. At each stage of the algorithm, every node N in the subspace tree has a set of *descendent assignments* $\mathcal{T}(N) \subseteq \mathcal{D}$ from the database \mathcal{D} . These are the data points which will be assigned to one of the descendants of node N during the tree construction process, but not to node N itself. In addition, each node also has a set of *direct assignments* $\mathcal{Q}(N)$, which are data points that are reduced onto node N . A data point becomes a direct assignment of node N , when

one of the following two properties is satisfied: (1) The data point is at most a distance of ϵ from the hyperplane corresponding to node N and the discrimination factor $\beta(N)$ for the node is larger than γ_1 . (2) The discrimination factor $\beta(N)$ for the node is larger than γ_2 . All assignments of the node which are not direct assignments automatically become descendent assignments. In each iteration, the descendent assignments $\mathcal{T}(N)$ of the nodes at a given level of the tree are partitioned further into at most k_{max} children of node N . This partitioning is based on the distance of the data points to the hyperplanes corresponding to the k_{max} children of N . Specifically, each data point is assigned to the hyperplane from which it has the least distance. The assigned points are then classified either as a descendent or direct assignments depending upon the distance from the hyperplane and corresponding discrimination index. As noted earlier, the latter value determines whether a node is a default node, a discriminative node, or forbidden node. Forbidden nodes do not have any direct assignments and therefore all data points at forbidden nodes automatically become descendent assignments. On the other hand, in the case of a discriminative node, the reverse is true and all points become direct assignments. For the case of default nodes, a point becomes a direct assignment only if it is at a distance of at most ϵ from the corresponding hyperplane. This process continues until each data point becomes the direct assignment of some node, or is identified in the anomaly set. The overall algorithm for subspace tree construction is illustrated in Figure 3.

A levelwise algorithm is used during the tree construction phase. The reason for this levelwise approach is that the database operations during the construction of a given level of nodes can be consolidated into a single database pass. The actual construction of the m th level is achieved by sampling one representative point for each of the k_{max} children of the level- $(m-1)$ nodes. This representative point is added in order to create the corresponding 1-dimensional extension. These k_{max} representative points are sampled from the local segment $\mathcal{T}(N)$ of the database. Picking these representative points at a node N is tricky, since we would like to ensure that they satisfy the following two properties: (1) The points represent the behavior of localized regions in the data. (2) The points are sufficiently representative of the different classes in that data locality. In order to achieve this, we would like to ensure that the set $\mathcal{R}(N)$ represents as many different classes as possible. Since the representative extensions at a node N are sampled from the local segment $\mathcal{T}(N)$ only, the first property is satisfied. In order to satisfy the second property, we choose the class from which the data points are sampled

Algorithm *SampleSubspaceTree*(CompressionTolerance: ϵ , MaximumTreeDegree: k_{max} , Database: \mathcal{D} , Node Limit: L , Discrimination Tolerance: γ_1 , Discrimination Target: γ_2)

begin

Sample $2 * k_{max} * sampfactor$ points from \mathcal{D} and pair up points randomly to create

$k_{max} * sampfactor$ 1-dim. point representative hyperplanes (lines) denoted by \mathcal{S} ;

$(\mathcal{S}, \beta(S_1) \dots \beta(S_{k_{max}})) = SelectSubspaces(\mathcal{S}, k_{max})$;

$(\mathcal{T}(S_1), \dots, \mathcal{T}(S_{k_{max}}), \mathcal{Q}(S_1), \dots, \mathcal{Q}(S_{k_{max}})) = PartitionData(\mathcal{D}, \mathcal{S})$;

for $i = 1$ **to** k_{max} **do** **if** $\beta(S_i) \geq \gamma_2$ **then** { Discriminative Node } $\mathcal{Q}(S_i) = \mathcal{Q}(S_i) \cup \mathcal{T}(S_i)$; $\mathcal{T}(S_i) = \phi$; $Status(S_i) = 2$;

else if $\beta(S_i) < \gamma_1$ **then** { Forbidden Node } $\mathcal{T}(S_i) = \mathcal{T}(S_i) \cup \mathcal{Q}(S_i)$; $\mathcal{Q}(S_i) = \phi$; $Status(S_i) = 1$;

else $Status(S_i) = 0$;

$\mathcal{S} = DeleteNodes(S_1 \dots S_{k_{max}}, min-thresh)$; { \mathcal{L}_m is the set of level- m nodes }

$m = 1$; $\mathcal{L}_1 = \mathcal{S}$; { Each hyperplane (line) in \mathcal{S} is the child of *Root* };

while ($\mathcal{L}_m \neq \{\}$) and (less than L nodes have been generated) **do** **begin**

for each non-null level- m node $R \in \mathcal{L}_m$ **do** **begin**

Sample $k_{max} * sampfactor$ points from $\mathcal{T}(R)$;

Extend the node R by each of these $k_{max} * sampfactor$ points (in turn) to create the $k_{max} * sampfactor$ corresponding $(m + 1)$ -dimensional hyperplanes denoted by \mathcal{S} ;

$(\mathcal{S}, \beta(S_1) \dots \beta(S_{k_{max}})) = SelectSubspaces(\mathcal{S}, k_{max})$;

$(\mathcal{T}(S_1), \dots, \mathcal{T}(S_{k_{max}}), \mathcal{Q}(S_1), \dots, \mathcal{Q}(S_{k_{max}})) = PartitionData(\mathcal{T}(R), \mathcal{S})$;

for $i = 1$ **to** k_{max} **do** **if** $\beta(S_i) \geq \gamma_2$ **then** { Discriminative Node } $\mathcal{Q}(S_i) = \mathcal{Q}(S_i) \cup \mathcal{T}(S_i)$; $\mathcal{T}(S_i) = \phi$; $Status(S_i) = 2$;

else if $\beta(S_i) < \gamma_1$ **then** { Forbidden Node } $\mathcal{T}(S_i) = \mathcal{T}(S_i) \cup \mathcal{Q}(S_i)$; $\mathcal{Q}(S_i) = \phi$; $Status(S_i) = 1$;

else $Status(S_i) = 0$;

$\mathcal{S} = DeleteNodes(S_1 \dots S_{k_{max}}, min-thresh)$; { Thus \mathcal{S} contains at most k_{max} children of R }

$\mathcal{L}_{m+1} = \mathcal{L}_{m+1} \cup \mathcal{S}$; $m = m + 1$;

end;

end;

end

Figure 3: Supervised Subspace Tree Construction

as follows: Let $f_1^R \dots f_k^R$ be the fractional class distributions in $\mathcal{R}(N)$ and $f_1^T \dots f_k^T$ be the fractional class distributions in $\mathcal{T}(N)$. We sample a point belonging to the class i in $\mathcal{T}(N)$ which always belongs to the class with the least value of f_i^R / f_i^T . Thus, this process picks the point from the class which is most under-represented in $\mathcal{R}(N)$ relative to $\mathcal{T}(N)$.

A total of $k_{max} * sampfactor$ points (belonging to the selected class) are picked for extension of the nodes from level- $(m - 1)$ to level- m . Thus, a total of $k_{max} * sampfactor$ m -dimensional hyperplanes can be generated by combining the representative set $\mathcal{R}(N)$ of node N with each of these sampled points. The purpose of oversampling by a factor of $sampfactor$ is to increase the effectiveness of the final children subspaces which are picked. The larger the value of $sampfactor$, the better the sampled subspaces, but the greater the computational requirement. Next, the procedure *SelectSubspaces* picks k_{max} hyperplanes out of these $k_{max} * sampfactor$ possibilities so that the different classes are as well separated as possible. The first task is to partition the $k_{max} * sampfactor$ hyperplanes into $sampfactor$ sets of k_{max} hyperplanes. We will pick one of these partitions depending upon the quality of the assignment. In order to achieve this, the distance of the data point \bar{x} to each of the $k_{max} * sampfactor$

hyperplanes is determined. For each of the $sampfactor$ sets of hyperplanes, we assign the data point \bar{x} to the closest hyperplane from that partition. This results in a total of $sampfactor$ possible assignments of the data points. The quality of the assignment depends upon how well the different classes are discriminated from one another in the resulting localized data sets. The *SelectSubspaces* procedure quantifies this separation in terms of the discrimination index $\beta(\cdot)$ for each of these nodes. The average discrimination index for each of the $sampfactor$ sets of nodes is calculated. The set of k_{max} hyperplanes with the smallest average discrimination index is chosen for the purpose of reduction. These hyperplanes are returned as the set $\mathcal{S} = (S_1 \dots S_{k_{max}})$. In addition, the discrimination index of each of these nodes is returned as $(\beta_1 \dots \beta_{k_{max}})$.

Once the hyperplanes which form the optimal extensions of node N have been determined, each point in $\mathcal{T}(N)$ is re-assigned to one of the children of N by the use of the procedure *PartitionData*. Specifically, each point in $\mathcal{T}(N)$ is assigned to the child node to which it is the least distance. Furthermore, the assigned nodes are classified into direct assignments $\mathcal{T}(S_i)$ and descendent assignments $\mathcal{Q}(S_i)$. Initially, the *PartitionData* procedure returns the direct assignments $\mathcal{T}(S_i)$ and descendent assignments $\mathcal{Q}(S_i)$ using only the distance of the

data points from the corresponding hyperplanes. Specifically, the *PartitionData* procedure returns a point as a direct assignment, if it is at a distance of at most ϵ from the corresponding hyperplane. Otherwise, it returns the data point as a descendent assignment. After application of the *PartitionData* procedure, we further re-adjust the direct and descendent assignments using the discrimination levels $\beta(S_i)$ of each child node S_i . If the node S_i is a forbidden node, then the direct assignment set $\mathcal{Q}(S_i)$ is reset to null, and all points become descendent assignments. Therefore $\mathcal{Q}(S_i)$ is added to $\mathcal{T}(S_i)$. The reverse is true when the node is a discriminative node. In that case, all points become direct assignments and $\mathcal{T}(S_i)$ is set to null.

Nodes which have too few points assigned to them are not useful for the dimensionality reduction process. Such nodes are deleted by the procedure *DeleteNodes*. The corresponding data points are considered exceptions which are stored separately by the algorithm.

The first iteration of the algorithm ($m = 1$) is special in which we sample $2 * k_{max} * sampfactor$ points in order to create the initial set of $k_{max} * sampfactor$ lines. Thus, the only difference is that twice the number of points need to be sampled in order to create the 1-dimensional hyperplanes used by the algorithm. The procedure for selection of these points is exactly similar to that of the general case. The other procedures such as selection and deletion of subspaces and data partitioning are also the same as in the general case.

In order to ease conceptual abstraction, we have presented the *PartitionData* and *SelectSubspaces* procedures separately for each node. In the actual implementation, this procedure is executed simultaneously for all nodes at a given level in one scan. Similarly, the process of picking the best hyperplanes for all nodes at a given level is executed simultaneously in a single scan of the data.

The process of levelwise tree construction continues until no node in the current level can be extended any further, or the maximum limit L for the number of nodes has been reached. We would like this limit L to be determined by main memory limitations, since it will be found useful in applying it effectively during the classification process. For our implementation, we used a conservative limit of only $L = 10,000$ nodes, which was well within current main memory limitations for even 1000-dimensional data sets.

Each of the procedures *SelectSubspaces* and *PartitionData* require the computation of distances of data points \bar{x} to the representative hyperplanes. In order to perform these distance computations, the axis representations of the hyperplanes need to be determined. A hyperplane node N at level- m is only implicitly de-

finied by the $(m + 1)$ data points $\{\bar{z}_1 \dots \bar{z}_{m+1}\}$ stored at the nodes along the path from the root to N . The next tricky issue is to compute the axis representation $(\bar{y}, \mathcal{E} = \{\bar{e}_1 \dots \bar{e}_m\})$ of the points $\{\bar{z}_1 \dots \bar{z}_{m+1}\}$ efficiently in a way that can be replicated exactly at the time of data reconstruction. This is especially important, since there can be an infinite number of axis representations of the same hyperplane, but the projection coordinates are computed only with respect to a particular axis-representation. The corresponding representation $(\bar{y}, \mathcal{E} = \{\bar{e}_1 \dots \bar{e}_m\})$ is computed as follows:

We first set $\bar{y} = \bar{z}_1$ and $\bar{e}_1 = (\bar{z}_2 - \bar{z}_1) / \|\bar{z}_2 - \bar{z}_1\|$. Next, we iteratively compute \bar{e}_i from $\bar{e}_1 \dots \bar{e}_{i-1}$ as follows:

$$(2.2) \quad \bar{e}_i = \frac{\bar{z}_{i+1} - \bar{z}_1 - \sum_{j=1}^{i-1} [(\bar{z}_{i+1} - \bar{z}_1) \cdot \bar{e}_j] \bar{e}_j}{\|\bar{z}_{i+1} - \bar{z}_1 - \sum_{j=1}^{i-1} [(\bar{z}_{i+1} - \bar{z}_1) \cdot \bar{e}_j] \bar{e}_j\|}$$

Equation 2.2 is essentially the iteration for the Gram-Schmidt orthogonalization process [10]. The following observation is a direct consequence of this fact:

OBSERVATION 2.1. *The set (\bar{z}_1, \mathcal{E}) generated by Equation 2.2 is an axis representation of the hyperplane $\mathcal{R}(\bar{z}_1 \dots \bar{z}_{m+1})$.*

Many axis representations can be generated using Equation 2.2 for the same hyperplane $\mathcal{R}(\{\bar{z}_1 \dots \bar{z}_{m+1}\})$ depending upon the ordering of $\{\bar{z}_1 \dots \bar{z}_{m+1}\}$. Since we need to convert from point representations to axis representations in a consistent way for both data reduction and reconstruction, this ordering needs to be fixed in advance. For the purpose of this paper, we will assume that the point ordering is always the same as one in which it was sampled during the top-down tree construction process. This leads to representative points sampled at higher levels of the tree to be ordered first, and points at lower levels to be ordered last. The only ambiguity is for the level-1 nodes at which 2 points are stored instead of one. In that case, the record which is lexicographically smaller is ordered earlier. We shall refer to this particular convention for axis representation as the *path-ordered axis representation*.

2.3 Computation of Discrimination Index Several methods can be used in order to calculate the discrimination index of the reduced data at a given node. The most popularly used method is the Fisher's linear discriminant [7] of the data points in the node. This discriminant minimizes the ratio of the intra-class distance to the inter-class distance. Our approach however uses a nearest neighbor discriminant. In this technique, we find the nearest neighbor to each data point in the dimensionality reduced database, and calculate the fraction of data points which share the same class as their

nearest neighbor. This fraction is reported as the discrimination index.

2.4 Storage of Compressed Representation The storage algorithm needs to store the tree structure and the set of points associated with each node of the tree. Each of these components are stored as follows:

(1) The Subspace Tree: If the axis-systems are explicitly stored at each node, the storage requirements for the tree structure could be considerable. This is because the axis representation (\bar{y}, \mathcal{E}) for an m -dimensional node requires $m+1$ orthonormal vectors (including the origin \bar{y} for the corresponding hyperplane). However, it turns out that we do not need to store the axis representations explicitly, For each level- m node, we only maintain the additional data point which increases the dimensionality of the corresponding subspace by one. In addition, we need to maintain the identity of the node, its immediate parent and the status of the node (corresponding to whether it is forbidden, discriminative, or a default node). Thus, a total of $(d+3)$ values are required for each node. For the (at most k_{max}) level-1 nodes, the storage requires $(2 \cdot d + 3)$ values since we need to maintain the *two* points which define the sampled line in lexicographic ordering.

The subspaces in the tree structure are thus only *implicitly* defined by the sequence of points from the root to that node. Thus, by using this method, we require only one vector for an m -dimensional node rather than $O(m)$ vectors. Since most nodes in the tree are at lower levels, such savings add up considerably over the entire tree structure. The reason for this storage efficiency is the implicit representation of the subspace tree. This results in the reuse of the vector stored at a given node for all descendants of that node.

(2) The Reduced Database: Each data point is either an exception or is associated with one node in the tree. We store the identity of the node for which it is a direct assignment. In addition, we maintain the coordinates of the data point for the axis representation (\bar{y}, \mathcal{E}) of this hyperplane in accordance with Equation 1.1. The projection coordinates of \bar{x} on (\bar{y}, \mathcal{E}) are given by $(c_1 \dots c_m) = \{\bar{e}_1 \cdot (\bar{x} - \bar{y}) \dots \bar{e}_m \cdot (\bar{x} - \bar{y})\}$. The class labels were stored separately.

2.5 Reconstruction Algorithm Since the subspace tree is represented in an implicit format, the axis representations of the nodes need to be reconstructed. It is possible to do this efficiently because of the use of the path-ordered axis convention for representation. Since there are a large number of nodes, it would seem that the initial phase could be quite expensive to perform for each node. However, it turns out that because of the use

Algorithm *SelectClassifierModels*(Subspace Tree: ST)
begin

for each node N in the subspace tree ST **do**
begin
Divide the data set $\mathcal{T}(N)$ into two parts
 $\mathcal{T}_1(N)$ and $\mathcal{T}_2(N)$ with ratio $r : 1$;
Use classification training algorithms $\mathcal{A}_1 \dots \mathcal{A}_m$
on $\mathcal{T}_1(N)$ to create models $\mathcal{M}_1(N) \dots \mathcal{M}_m(N)$;
Compute accuracy of models $\mathcal{M}_1(N) \dots \mathcal{M}_m(N)$
on $\mathcal{T}_2(N)$;
Pick model with highest classification accuracy on $\mathcal{T}_2(N)$
and denote as $\mathcal{CL}(N)$;

end;
end

Figure 4: Node Specific Classification

of the path-ordered convention for axis-representations, this can be achieved in a time complexity which requires the computation of only one axis per node. The trick is to construct the axis representations of the nodes in the tree in a top-down fashion. This is because the Equation 2.2 computes the axis representation $\{\bar{e}_1 \dots \bar{e}_i\}$ of a node by using the axis representation $\{\bar{e}_1 \dots \bar{e}_{i-1}\}$ of its parent and the point \bar{z} stored at that node. (For the nodes at level-1, lexicographic ordering of the representative points is assumed.) It is easy to verify that this order of processing results in the path-ordered axis representation.

Once the axis representations of the nodes have been constructed, it is simple to perform the necessary axis transformations which represent the reconstructed database in terms of the original attributes. Recall that for each database point \bar{x} , the identity of the corresponding node is also stored along with it. Let (\bar{y}, \mathcal{E}) be the corresponding hyperplane and $(c_1 \dots c_m) = \{\bar{e}_1 \cdot (\bar{x} - \bar{y}) \dots \bar{e}_m \cdot (\bar{x} - \bar{y})\}$ be the coordinates of \bar{x} along this m -dimensional axis representation. Then, as evident from Equation 1.1, the reconstructed point \bar{x}' is given by $\bar{x}' = \bar{y} + \sum_{i=1}^m [c_i] \bar{e}_i$

3 Effective Classification by Locality Specific Decomposition

It turns out that the locality specific decomposition approach of the dimensionality reduction problem not only provides a reduced representation of the data, but can also be leveraged for effective classification. This is because the dimensionality reduction process decomposes the data into a number of different parts each of which have a unique and distinctive class distribution. The use of different training phases on different nodes may be quite effective in these cases. In fact, the fundamental variation in the dimensionalities and data characteristics of the different nodes makes

Algorithm *DecomposableClassify*(Test Instance: $\overline{x_T}$, SubspaceTree: \mathcal{ST} , Classifier Models: $\mathcal{CL}(\cdot)$)

begin

Perform a hierarchical traversal of the tree \mathcal{ST} so as to find the highest level node N_T which is not a forbidden node and is either a default node within a distance ϵ from x_T or is a discriminative node;

if no such node N_T exists (outlier node)

then report majority class of outlier points during subspace tree generation;

else use classifier model $\mathcal{CL}(N_T)$ in order to classify the test instance $\overline{x_T}$;

return class label of $\overline{x_T}$;

end

Figure 5: Locally Decomposable Classification

each segment amenable to different kinds of classifiers.

For each data set at a given node N of the tree, we apply a classifier whose identity is dependent only upon the results of a training process applied to the data at that particular node. We denote the classifier model at a given node N by $\mathcal{CL}(N)$. Thus, for one node a nearest neighbor classifier might be used, whereas for another node an association rule classifier may be used. We will discuss the details of the process of determination of the classifier model $\mathcal{CL}(N)$ slightly later.

Since the reduced data is contained only in nodes which are not forbidden, only those nodes are relevant to the classification process. For each such node N in the subspace tree \mathcal{ST} , we use the corresponding data $\mathcal{T}(N)$ in order to decide on the classification algorithm which best suits that particular set of records. In the first step, we divide the data $\mathcal{T}(N)$ into two parts $\mathcal{T}_1(N)$ and $\mathcal{T}_2(N)$ in the ratio $r : 1$. The set $\mathcal{T}_1(N)$ is used to build training different classification algorithms $\mathcal{A}_1 \dots \mathcal{A}_m$ on the particular locality of the data reduced to hyperplane N . The corresponding models are denoted by $\mathcal{M}_1 \dots \mathcal{M}_m$. The reduced representation of the data at node N is used for training purposes. Once the different algorithms have been trained using $\mathcal{T}_1(N)$, then the best algorithm is determined by using testing on $\mathcal{T}_2(N)$. The classification accuracy of each of the models $\mathcal{M}_1 \dots \mathcal{M}_m$ on $\mathcal{T}_2(N)$ is computed. The model \mathcal{M}_s with maximum classification accuracy on $\mathcal{T}_2(N)$ is determined. This is the classification model $\mathcal{CL}(N)$ used at node N . The overall algorithm is described in the procedure *SelectClassifierModels* of Figure 4.

Once the models for each of the nodes have been constructed, they can be used for classifying individual test instances. For a given test instance, we first decide the identity of the node that it belongs to. In order to find which node a test instance T belongs to, we use the same rules that are used to assign points to

Data Set	Records	Attributes
Forest Cover	581012	54
Keyword1	67121	239
Keyword2	68134	239
C35.I6.D100K.P100	100000	100
C40.I6.D100K.P100	100000	100

Table 1: Characteristics of the Data Sets

nodes in the construction of the subspace tree from the original database \mathcal{D} . Therefore, a hierarchical traversal is performed on the tree structure using the same rules as utilized by the tree construction algorithm in defining direct assignments. Thus, for each data point $\overline{x_T}$ which is to be classified, we perform a hierarchical traversal of the tree starting at the root. The traversal always picks that branch of the tree to which x_T is closest, until it either reaches a node L at which one of the following conditions is satisfied:

- (1) The node L is a discriminative node.
- (2) The node L is neither discriminative nor forbidden, and the corresponding hyperplane is within the specified error tolerance of ϵ .

We note that this choice of tree traversal and node selection directly mirrors the process of direct and descendent assignment of points to nodes during the tree construction process. In some cases, such a node may not be found if the point x_T is an outlier. In that case, the majority class from the set of outlier points found during tree construction is reported as the class label.

In the event that a node is indeed found which satisfied one of the above conditions, we denote it by \mathcal{N}_T . The classifier model $\mathcal{CL}(\mathcal{N}_T)$ is then used for the classification of the test instance $\overline{x_T}$ and the corresponding class label is reported as the final result of the classification process. The overall algorithm is described in the procedure *DecomposableClassify* of Figure 5.

4 Empirical Results

The system was tested on an AIX 4.1.4 system with a speed of 300 MHz and 100 MB of main memory. The data was stored on a 2 GB SCSI drive. The supervised subspace sampling algorithm was tested for the following measures: (1) Effectiveness of data reduction with respect to linear discriminant analysis. (2) Efficiency of data reduction process. (3) Effectiveness of the supervised subspace sampling method as an approach for decomposable classification.

A number of synthetic and real data sets from a were utilized in order to test the effectiveness of the reduction and classification process. The characteristics of the data sets are illustrated in Table 1. The Forest Cover Data set was available from the UCI KDD archive. All attributes and records of the forest cover data set were used. In this case, binary attributes were treated in a similar way to numerical attributes for the classification and dimensionality reduction process. The keyword data sets were derived from web pages in the *Yahoo!* taxonomy. The records were generated by finding the frequency of 239 most discriminative keywords in web pages drawn from the *Yahoo!* taxonomy. The classes correspond to the highest level categories in the *Yahoo!* taxonomy. These keywords were determined using the gini index value of each feature. The 239 features with highest gini index were used. Two data sets were generated corresponding to the commercial and non-commercial sections in the *Yahoo!* taxonomy. We denote these data sets by Keyword1 and Keyword2 respectively.

In order to test the algorithmic effectiveness further, we used synthetic data sets. These are also useful for scalability testing, since it is possible to show clear trends with varying data size and dimensionality. For this purpose, we used the market basket data sets proposed in [3]. In order to create different classes, we generated two different instantiations of the data set Tx.Iy.Dz (according to the notations of [3]) and created a two class data set which was a mixture of these two data sets in equal proportions. The only difference from the generation methodology in [3] is that a subset of w items were used instead of the standard 1000 items used in [3]. We refer to this data set as Cx.Iy.D(2z).Pw. Since the data set Tx.Iy.Dz contains z records, the data set Cx.Iy.D(2z).Pw contains $2 \cdot z$ records. Two data sets C35.I6.D100K.P100 and C40.I6.D100K.P100 were created using this methodology.

In order to test the effectiveness of the data reduction process, we used the Fisher’s discriminant method as a comparative baseline. This method was implemented as follows:

- (1) The first dimension was found by finding the Fisher’s direction [7] which maximized the discrimination between the class variables.
- (2) The data was then projected onto the remaining $d-1$ dimensions defined by the hyperplane orthonormal to the first direction.
- (3) The next dimension was again determined by finding the Fisher’s direction which maximized the total discrimination in the remaining $d-1$ dimensions.
- (4) The process was repeated iteratively until a new

orthonormal axis system was determined.

- (5) The most discriminative dimensions in the data were retained by using a threshold on the Fisher’s discriminant value. This threshold was determined by finding the mean μ and variance σ^2 of the Fisher’s index for the newly determined dimensions. The threshold was then set at $\mu + 2 \cdot \sigma$.

In order to test the effectiveness, we applied two different classification algorithms on the data sets. The specific classifiers tested were the C4.5 and the nearest neighbor algorithms. The following different approaches were tested:

- (1) Utilizing individual classifiers on the reduced data for global dimensionality reduction.
- (2) Utilizing individual classifiers for local dimensionality reduction. In this case, the same classifier was used for every local segment. Thus, the approach benefits from the use of different training models to the different segments which are quite heterogeneous because of the nature of the supervised partitioning process.
- (3) Utilizing the decomposable classification process on the data. This approach benefits not only from the use of different training models, but also from the different classifiers on the different nodes.

In Table 2, we have illustrated the effectiveness of the different classifiers on the data sets. The first two columns report the accuracy of the nearest neighbor and C4.5 classifiers on the full dimensional data sets. The next two columns report the accuracy, when the data was reduced with Fisher’s discriminant. The following two columns report the accuracy by using local training on each of the nodes with a particular classifier. The final column reports the accuracy of the decomposable classification process. We make the following observations from the results:

- (1) Neither of the two classifiers performed as effectively on the full dimensional data as it did on the reduced sets. This is quite natural, since the reduction procedure was able to remove the noise which was irrelevant to the classification process.
- (2) Both the classifiers showed superior classification accuracy when reduced with the subspace sampling approach as compared to the Fisher’s method.
- (3) The decomposable classifier always performed more effectively than any combination of classifier and reduction technique.

An important factor to be kept in mind is that the nearest neighbor and C4.5 classifiers did not show

Data Set	NN (Full)	C4.5 (Full)	NN (Fish.)	C4.5 (Fish.)	NN (Subsp.)	C4.5 (Subsp.)	Decomposable Classifier
Forest Cover	63.3%	60.7%	63.1%	61.3%	68.5%	67.3%	71.2%
Keyword1	53.2%	47.1%	56.7%	54.5%	65.3%	63.4%	69.3%
Keyword2	50.1%	43.4%	52.1%	51.3%	61.5%	60.7%	64.5%
C35.I6.D100K.P100	83.4%	84.3%	82.5%	84.7%	86.8%	86.5%	89.5%
C40.I6.D100K.P100	74.2%	75.1%	73.3%	76.7%	80.3%	80.9%	84.4%

Table 2: Effectiveness of Classifiers on Different Data Sets




Figure 6: Efficiency results

consistent performance across different data sets. In some data sets, the C4.5 classifier was better, whereas in others the nearest neighbor technique was better. Furthermore, the application of a particular kind of data reduction method affected the relative behavior of the two classifiers. This is undesirable from the perspective of a classification task, since this makes it more difficult to determine the best possible classification algorithm for each particular data set. On the other hand, the decomposable classification results *consistently outperformed all combinations of classifier and data reduction methods*.

A second area of examination was the reduction factor provided by each of the methods on the data. The *reduction factor* was defined as the ratio of the reduced data size to the original data size. For the subspace sampling method, this reduced data contained both the subspace tree and the database records. In Table 3, we have illustrated the corresponding reduction factors for each of the methods. It is clear that the supervised subspace sampling approach provides much higher reduction factors than the Fisher Discriminant. We have already shown that the classification accuracy is much better when the reduced data representation is generated using supervised subspace sampling. The fact that the more compact representation of the subspace sampling technique provides better accuracy indicates

Data Set	Reduc. Factor (Fisher)	Reduc. Factor (Subsp)
Forest Cover	0.204	0.11
Keyword1	0.167	0.091
Keyword2	0.163	0.087
C35.I6.D100K.P100	0.19	0.081
C40.I6.D100K.P100	0.22	0.094

Table 3: Reduction Ratios of Diff. Methods

that it is more effective at picking those subspaces which are most discriminatory for the classification process.

We used synthetic data sets for illustrating scalability trends with varying data size and dimensionality. Both the Fisher method and the subspace sampling method were applied to generations of the data of varying dimensionality corresponding to the synthetic data sets C35.I6.D100K.Px and C40.I6.D100K.Px. This provided an idea of the scalability of the algorithms with increasing data dimensionality. The results are illustrated in Figures 6(a) and 6(b). It is clear that the subspace sampling method was always more efficient than the Fisher data reduction method, and the performance gap increased with data dimensionality. This is because the subspace sampling method requires simple sampling computations which scale almost linearly




Figure 7: Efficiency vs. Data Size (C40.I6.Dx.P100)

with dimensionality. On the other hand, the running times of the Fisher’s discriminant method rapidly increase with dimensionality because of the costly computation of finding optimal axis directions. We also tested the two methods for increasing scalability with data size. In this case, we used samples of varying sizes of the data sets C35.I6.D100K.P100 and C40.I6.D100K.P100. The results are illustrated in Figures 6(c) and 7 respectively. It is clear that both methods scaled linearly with data size, though the subspace sampling technique consistently outperformed the Fisher method.

5 Conclusions and Summary

In this paper we have proposed an effective local dimensionality data reduction method in the supervised domain. Most current dimensionality reduction methods such as SVD are designed only for the unsupervised domain. The aim of dimensionality reduction method in the supervised domain is to create a new axis-system so that the discriminatory characteristics of the data are retained, and the classification accuracy is improved. Methods such as linear discriminant analysis turn out to be computationally intensive in practice, and often cannot be efficiently applied to large data sets. Furthermore, the global approach of these techniques often make the methods ineffective in practice. The supervised subspace sampling approach uses local data reduction in which the reduction of a data point depends upon the class distributions in its locality. As a result, the supervised subspace sampling technique allows a natural decomposition of the data so that the implicit dimensionality of each data locality is minimized. This improves the effectiveness of the reduction process, while retaining the efficiency of a sampling based technique. The reduction process improves the accuracy of the classification process because of removal of the irrelevant axis directions in the data. In addition, the supervised reduction process naturally facilitates the construction of a decomposable classifier which is able to provide much better classification accuracy than a global data

reduction process. Thus, the improved efficiency, compression quality and classification accuracy of this reduction method make it an attractive approach for a number of real data domains.

References

- [1] C. C. Aggarwal, *Hierarchical Subspace Sampling: A Unified Framework for High Dimensional Reduction, Selectivity Estimation, and Nearest Neighbor Search*, ACM SIGMOD Conference, (2002), pp. 452–463.
- [2] C. C. Aggarwal, and P. S. Yu, *Finding Generalized Projected Clusters in High Dimensional Spaces*, ACM SIGMOD Conference, (2000), pp. 70–81.
- [3] R. Agrawal, and R. Srikant, *Fast Algorithms for Mining Association Rules in Large Databases*, VLDB Conference, (1994), pp. 487–499.
- [4] L. Brieman, *Bagging Predictors*, Machine Learning, 24 (1996), pp. 123–140.
- [5] K. Chakrabarti, and S. Mehrotra, *Local Dimensionality Reduction: A New Approach to Indexing High Dimensional Spaces*, VLDB Conference, (2000), pp. 89–100.
- [6] S. Chakrabarti, S. Roy, and M. V. Soundalgekar, *Fast and Accurate Text Classification via Multiple Linear Discriminant Projections*, VLDB Conference, (2002), pp. 658–669.
- [7] T. Cooke. *Two variations on Fisher’s linear discriminant for pattern recognition* PAMI, 24(2) (2002), pp. 268–273.
- [8] R. Duda, and P. Hart, *Pattern Classification and Scene Analysis*, Wiley, NY (1973).
- [9] C. Faloutsos, and K.-I. Lin, *FastMap: A Fast Algorithm for Indexing, Data-Mining and Visualization of Traditional and Multimedia Datasets*, ACM SIGMOD Conference, (1995), pp. 163–174.
- [10] K. Hoffman, and R. Kunze, *Linear Algebra*, Prentice Hall, NJ, (1998).
- [11] P. Indyk, and R. Motwani, *Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality*, ACM STOC Proceedings, (1998), pp. 604–613.
- [12] W. Johnson, and J. Lindenstrauss, *Extensions of Lipschitz mapping into a Hilbert space*, Conference in modern analysis and probability, (1984), pp. 189–206.
- [13] I. T. Jolliffe, *Principal Component Analysis*, Springer-Verlag, New York, (1986).
- [14] J. R. Quinlan, *C4.5: Programs for Machine Learning*, Morgan Kaufmann, (1993).
- [15] K. V. Ravi Kanth, D. Agrawal, and A. Singh, *Dimensionality Reduction for Similarity Searching in Dynamic Databases*, SIGMOD Conference, (1998), pp. 166–176.