

[Objectives](#)

Maximum Likelihood:

[Classification](#)

[Maximization](#)

[Mahalanobis](#)

Discrimination:

[2D Gaussian](#)

[Support Region](#)

[Decision Regions](#)

[Approach](#)

[Scatter](#)

[Example](#)

On-Line Resources:

[LDA Tutorial](#)

[ICA](#)

[Statistical Normalization](#)

[Pattern Recognition Applet](#)

[LNKNET Software](#)

- Objectives:
 - Review maximum likelihood classification
 - Appreciate the importance of weighted distance measures
 - Introduce the concept of discrimination
 - Understand under what conditions linear discriminant analysis is useful

This material can be found in most pattern recognition textbooks. This is the book we recommend:

- R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification* (Second Edition), Wiley Interscience, New York, New York, USA, ISBN: 0-471-05669-3, 2000.

and use in our pattern recognition course. The material in this lecture follows this textbook closely:

J. Deller, et. al., *Discrete-Time Processing of Speech Signals*, MacMillan Publishing Co., ISBN: 0-7803-5386-2, 2000.

Each of these sources contain references to the seminal publications in this area, including our all-time favorite:

- K. Fukunaga, *Introduction to Statistical Pattern Recognition*, MacMillan Publishing Company, San Diego, California, USA, ISBN: 0-1226-9851-7, 1990.



Introduction:

01: Organization
([html](#), [pdf](#))

Speech Signals:

02: Production
([html](#), [pdf](#))

03: Digital Models
([html](#), [pdf](#))

04: Perception
([html](#), [pdf](#))

05: Masking
([html](#), [pdf](#))

06: Phonetics and Phonology
([html](#), [pdf](#))

07: Syntax and Semantics
([html](#), [pdf](#))

Signal Processing:

08: Sampling
([html](#), [pdf](#))

09: Resampling
([html](#), [pdf](#))

10: Acoustic Transducers
([html](#), [pdf](#))

11: Temporal Analysis
([html](#), [pdf](#))

12: Frequency Domain Analysis
([html](#), [pdf](#))

13: Cepstral Analysis
([html](#), [pdf](#))

14: **Exam No. 1**
([html](#), [pdf](#))

15: Linear Prediction
([html](#), [pdf](#))

16: LP-Based Representations
([html](#), [pdf](#))

17: Spectral Normalization
([html](#), [pdf](#))

Parameterization:

18: Differentiation
([html](#), [pdf](#))

19: Principal Components
([html](#), [pdf](#))

20: Linear Discriminant Analysis
([html](#), [pdf](#))

Statistical Modeling:

21: Dynamic Programming
([html](#), [pdf](#))

ECE 8463: FUNDAMENTALS OF SPEECH RECOGNITION

Professor Joseph Picone
Department of Electrical and Computer Engineering
Mississippi State University

email: picone@isip.msstate.edu
phone/fax: 601-325-3149; office: 413 Simrall
URL: http://www.isip.msstate.edu/resources/courses/ece_8463

Modern speech understanding systems merge interdisciplinary technologies from Signal Processing, Pattern Recognition, Natural Language, and Linguistics into a unified statistical framework. These systems, which have applications in a wide range of signal processing problems, represent a revolution in Digital Signal Processing (DSP). Once a field dominated by vector-oriented processors and linear algebra-based mathematics, the current generation of DSP-based systems rely on sophisticated statistical models implemented using a complex software paradigm. Such systems are now capable of understanding continuous speech input for vocabularies of hundreds of thousands of words in operational environments.

In this course, we will explore the core components of modern statistically-based speech recognition systems. We will view speech recognition problem in terms of three tasks: signal modeling, network searching, and language understanding. We will conclude our discussion with an overview of state-of-the-art systems, and a review of available resources to support further research and technology development.

Tar files containing a compilation of all the notes are available. However, these files are large and will require a substantial amount of time to download. A tar file of the html version of the notes is available [here](#). These were generated using wget:

```
wget -np -k -m http://www.isip.msstate.edu/publications/courses/ece_8463/lectures/current
```

A pdf file containing the entire set of lecture notes is available [here](#). These were generated using Adobe Acrobat.

Questions or comments about the material presented here can be directed to help@isip.msstate.edu.

LECTURE 20: LINEAR DISCRIMINANT ANALYSIS

- Objectives:
 - Review maximum likelihood classification
 - Appreciate the importance of weighted distance measures
 - Introduce the concept of discrimination
 - Understand under what conditions linear discriminant analysis is useful

This material can be found in most pattern recognition textbooks. This is the book we recommend:

- R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification* (Second Edition), Wiley Interscience, New York, New York, USA, ISBN: 0-471-05669-3, 2000.

and use in our pattern recognition course. The material in this lecture follows this textbook closely:

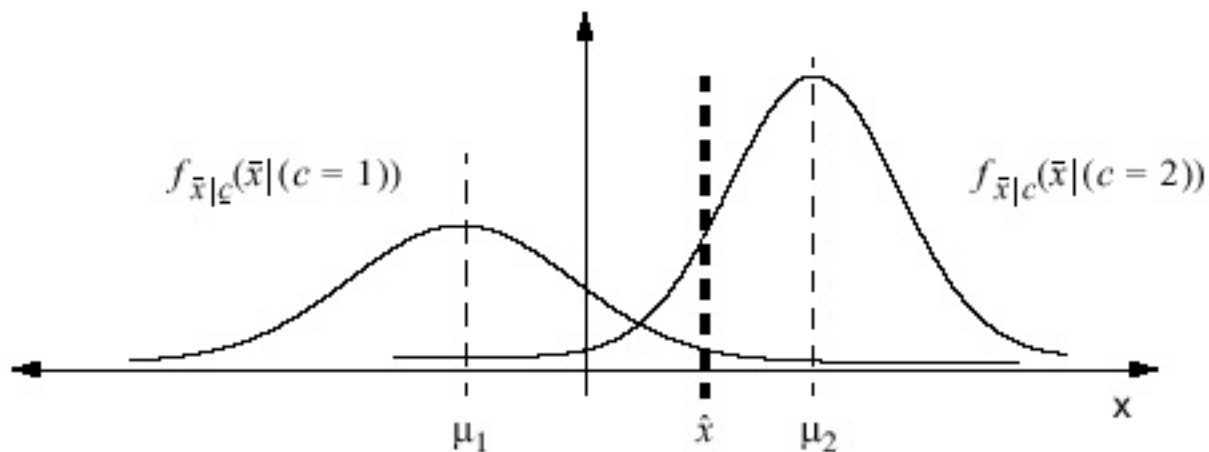
J. Deller, et. al., *Discrete-Time Processing of Speech Signals*, MacMillan Publishing Co., ISBN: 0-7803-5386-2, 2000.

Each of these sources contain references to the seminal publications in this area, including our all-time favorite:

- K. Fukunaga, *Introduction to Statistical Pattern Recognition*, MacMillan Publishing Company, San Diego, California, USA, ISBN: 0-1226-9851-7, 1990.

MAXIMUM LIKELIHOOD CLASSIFICATION

Consider the problem of assigning a measurement to one of two sets:



What is the best criterion for making a decision?

Ideally, we would select the class for which the conditional probability is highest:

$$c^* = \operatorname{argmax}_c P((c = \hat{c}) | (\bar{x} = \hat{\bar{x}}))$$

However, we can't estimate this probability directly from the training data. Hence, we consider:

$$c^* = \operatorname{argmax}_c P((\bar{x} = \hat{\bar{x}}) | (c = \hat{c}))$$

By definition

$$P((c = \hat{c}) | (\bar{x} = \hat{\bar{x}})) = \frac{P((c = \hat{c}), (\bar{x} = \hat{\bar{x}}))}{P(\bar{x} = \hat{\bar{x}})}$$

and

$$P((\bar{x} = \hat{\bar{x}}) | (c = \hat{c})) = \frac{P((c = \hat{c}), (\bar{x} = \hat{\bar{x}}))}{P(c = \hat{c})}$$

from which we have

$$P((c = \hat{c}) | (\bar{x} = \hat{\bar{x}})) = \frac{P((\bar{x} = \hat{\bar{x}}) | (c = \hat{c}))P(c = \hat{c})}{P(\bar{x} = \hat{\bar{x}})}$$

SPECIAL CASE: GAUSSIAN DISTRIBUTIONS

Clearly, the choice of c that maximizes the right side also maximizes the left side. Therefore,

$$\begin{aligned}c^* &= \operatorname{argmax}_c [P((\bar{x} = \hat{x}) | (c = \hat{c}))] \\ &= \operatorname{argmx}_c [P((\bar{x} = \hat{x}) | (c = \hat{c}))P(c = \hat{c})]\end{aligned}$$

if the class probabilities are equal,

$$c^* = \operatorname{argmx}_c [P((\bar{x} = \hat{x}) | (c = \hat{c}))]$$

A quantity *related* to the probability of an event which is used to make a decision about the occurrence of that event is often called a *likelihood measure*.

A decision rule that maximizes a likelihood is called a maximum likelihood decision.

In a case where the number of outcomes is not finite, we can use an analogous continuous distribution. It is common to assume a multivariate Gaussian distribution:

$$\begin{aligned}f_{\bar{x}|c}(x_1, \dots, x_N | c) &= f_{\bar{x}|c}(\hat{x} | \hat{c}) \\ &= \frac{1}{\sqrt{2\pi} |C_{\bar{x}|c}|} \exp \left\{ -\frac{1}{2} (\hat{x} - \bar{\mu}_{\bar{x}|c})^T C_{\bar{x}|c}^{-1} (\hat{x} - \bar{\mu}_{\bar{x}|c}) \right\}\end{aligned}$$

We can elect to maximize the log, $\ln[f_{\bar{x}|c}(\bar{x}|c)]$ rather than the likelihood (we refer to this as the log likelihood). This gives the decision rule:

$$c^* = \operatorname{argmin}_c \left[(\hat{x} - \bar{\mu}_{\bar{x}|c})^T C_{\bar{x}|c}^{-1} (\hat{x} - \bar{\mu}_{\hat{x}|c}) + \ln \left\{ |C_{\bar{x}|c}^{-1}| \right\} \right]$$

(Note that the maximization became a minimization.)

We can define a distance measure based on this as:

$$d_{ml}(\bar{x}, \bar{\mu}_{\bar{x}|c}) = (\hat{x} - \bar{\mu}_{\bar{x}|c})^T C_{\bar{x}|c}^{-1} (\hat{x} - \bar{\mu}_{\hat{x}|c}) + \ln \left\{ |C_{\bar{x}|c}^{-1}| \right\}$$

THE MAHALANOBIS DISTANCE

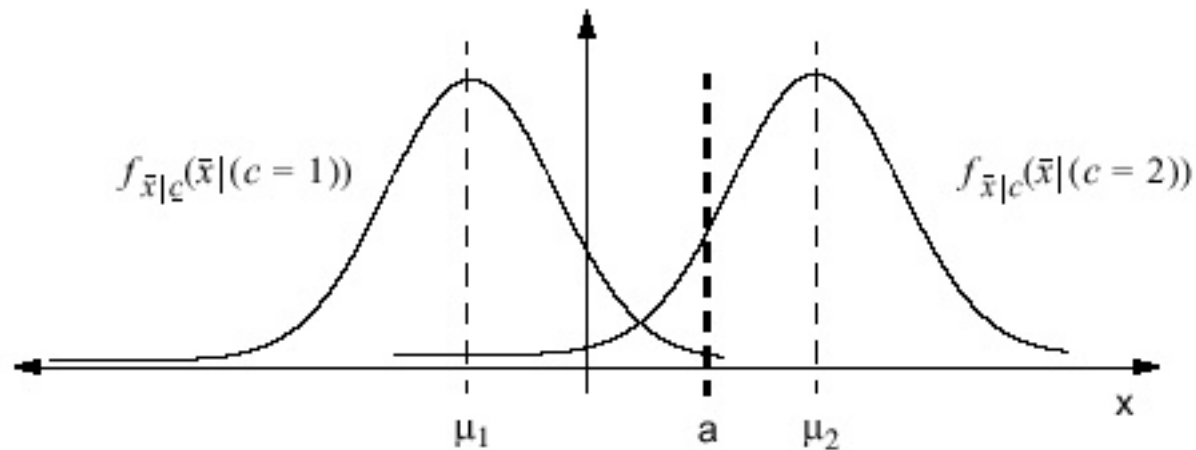
Note that the distance is conditioned on each class mean and covariance. This is why "generic" distance comparisons are a joke.

If the mean and covariance are the same across all classes, this expression simplifies to:

$$d_M(\hat{x}, \bar{\mu}_{\hat{x}|c}) = (\hat{x} - \bar{\mu}_{\hat{x}|c})^T C_{\hat{x}|c}^{-1} (\hat{x} - \bar{\mu}_{\hat{x}|c})$$

This is frequently called the *Mahalanobis distance*. But this is nothing more than a weighted Euclidean distance.

This result has a relatively simple geometric interpretation for the case of a single random variable with classes of equal variances:



The decision rule involves setting a threshold:

$$a = \left(\frac{\mu_1 + \mu_2}{2} \right) + \frac{\sigma^2}{\mu_1 - \mu_2} \ln \left(\frac{P(c=2)}{P(c=1)} \right)$$

and,

$$\begin{array}{ll} \text{if} & x < a & x \in (c=1) \\ \text{else} & & x \in (c=2) \end{array}$$

If the variances are not equal, the threshold shifts towards the distribution with the smaller variance.

What is an example of an application where the classes are not equiprobable?

TWO-DIMENSIONAL GAUSSIAN DISTRIBUTIONS

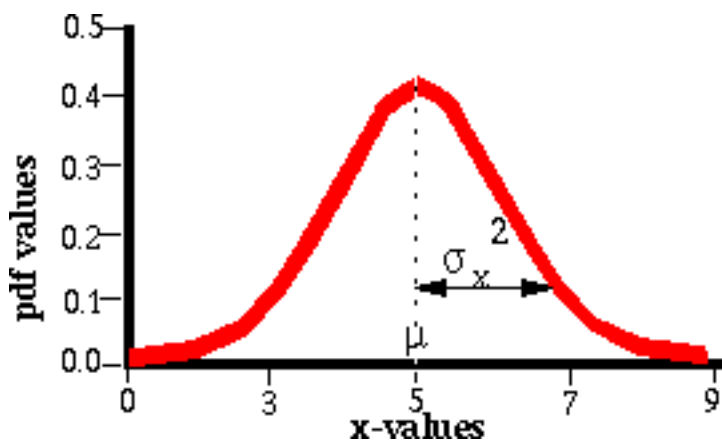
The equation of multivariate Gaussian distribution is as follows:

$$z = p(x,y) = \frac{1}{(2\pi)^{d/2} \sqrt{|C|}} e^{-\frac{1}{2}((x,y) - \mu) C^{-1} ((x,y) - \mu)^T}$$

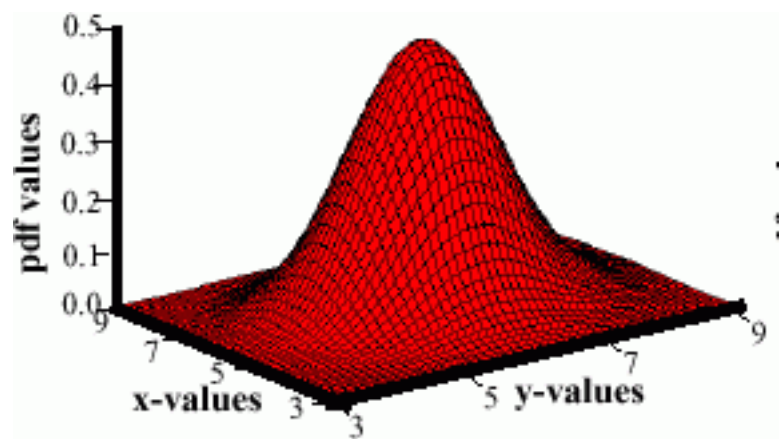
$$C = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{yx} & \sigma_y^2 \end{bmatrix}$$

where is the covariance matrix and d is the dimension of the distribution.

1D Gaussian distribution

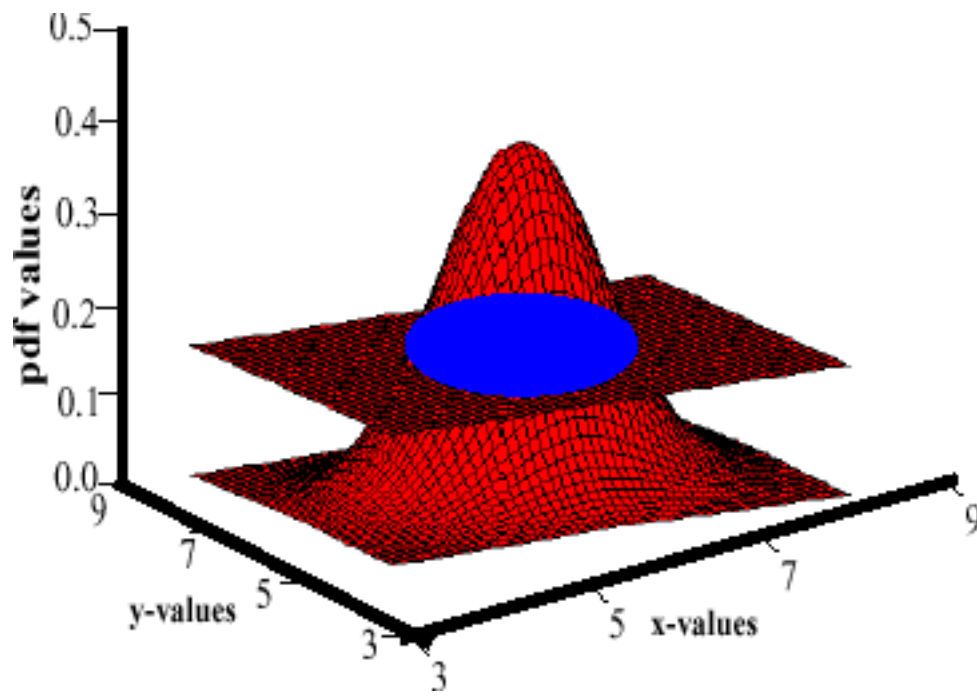


2D Gaussian distribution



SUPPORT REGION: A CONVENIENT VISUALIZATION TOOL

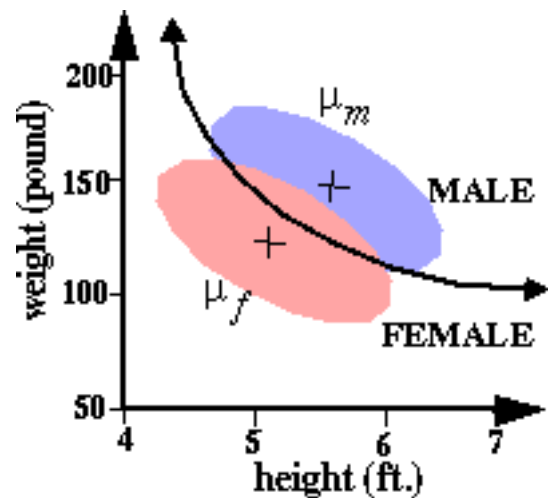
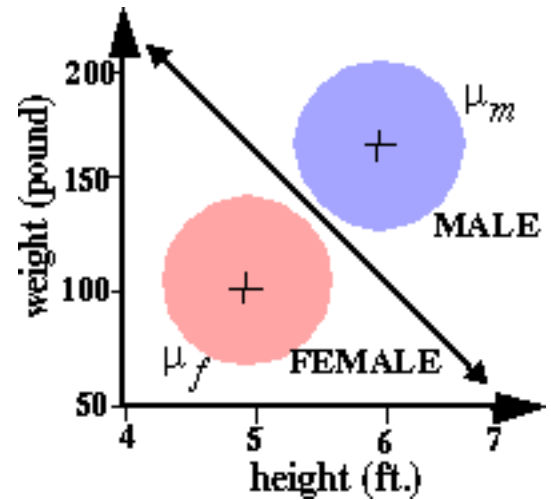
- A support region is the surface defined by the intersection of a Gaussian distribution with a plane.
- The shape of a support region is elliptical and depends on the covariance matrix of the original data.
- A convenient visualization tool in pattern recognition.



Some examples demonstrating the relationship between the covariance matrix and the 2D Gaussian distribution are shown below:

<p>Identity:</p> $C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	<p>A 3D plot of a Gaussian distribution with a circular support region. The vertical axis is 'pdf values' (0.0 to 0.5) and the horizontal axes are 'x-values' and 'y-values' (3 to 9).</p>	<p>A 2D plot of a Gaussian distribution with a circular support region. The horizontal axis is 'x-values' (3.0 to 9.0) and the vertical axis is 'y-values' (3.0 to 9.0).</p>
<p>Unequal Variances:</p> $C = \begin{bmatrix} 5 & 0 \\ 0 & 2 \end{bmatrix}$	<p>A 3D plot of a Gaussian distribution with an elongated support region. The vertical axis is 'pdf values' (0.0 to 0.5) and the horizontal axes are 'x-values' and 'y-values' (3 to 9).</p>	<p>A 2D plot of a Gaussian distribution with an elongated support region. The horizontal axis is 'x-values' (3.0 to 9.0) and the vertical axis is 'y-values' (3.0 to 9.0).</p>
<p>Nonzero off-diagonal:</p> $C = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$	<p>A 3D plot of a Gaussian distribution with a rotated support region. The vertical axis is 'pdf values' (0.0 to 0.5) and the horizontal axes are 'x-values' and 'y-values' (3 to 9).</p>	<p>A 2D plot of a Gaussian distribution with a rotated support region. The horizontal axis is 'x-values' (3.0 to 9.0) and the vertical axis is 'y-values' (3.0 to 9.0).</p>
<p>Unconstrained:</p> $C = \begin{bmatrix} 5 & 0.5 \\ 0.5 & 2 \end{bmatrix}$	<p>A 3D plot of a Gaussian distribution with a rotated and elongated support region. The vertical axis is 'pdf values' (0.0 to 0.5) and the horizontal axes are 'x-values' and 'y-values' (3 to 9).</p>	<p>A 2D plot of a Gaussian distribution with a rotated and elongated support region. The horizontal axis is 'x-values' (3.0 to 9.0) and the vertical axis is 'y-values' (3.0 to 9.0).</p>

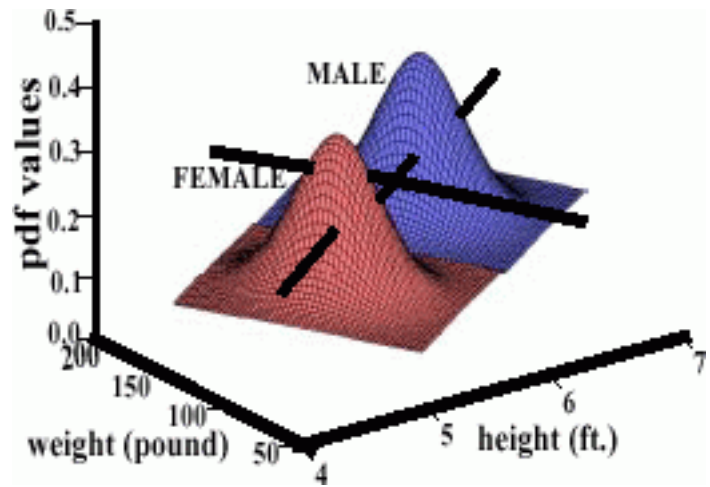
- **Case I: Distributions with equal variances**
- **Decision surface boundary is a line separating the two distributions (general case is a hyperplane).**
- **Case II: Distributions with unequal variances**
- **Direction of greatest variance is not the direction of discrimination.**



GOAL: MAXIMIZE SEPARABILITY

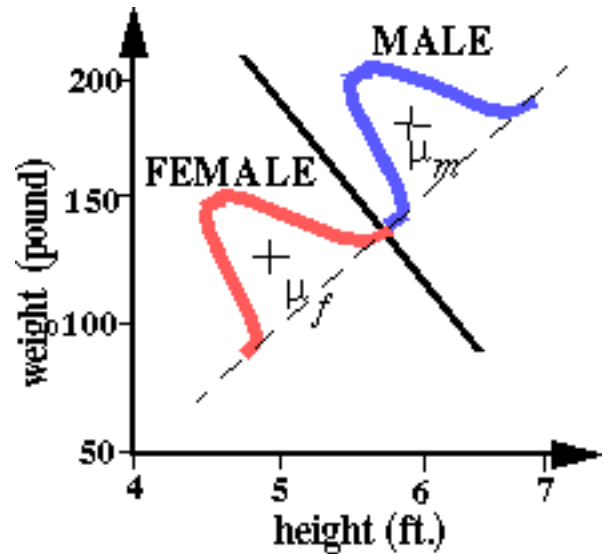
Principal Component Analysis:

Transform features to a new space in which the features are uncorrelated.



Linear Discriminant Analysis:

Projection of d -dimensional data onto a line; Dimensionality reduction by mapping L distributions to $(L-1)$ -dimensional subspace; maximize class separability.



OPTIMIZATION CRITERIA BASED ON SCATTER

- Within-class scatter defines the scatter of samples around their respective means.

$$S_w = \sum_{i=1}^L P_i E \left[(X - \mu_i)(X - \mu_i)^T \right]$$

- Between-class scatter defines scatter of expected vectors around the global mean.

$$S_b = \sum_{i=1}^L P_i \left[(\mu_i - \mu_0)(\mu_i - \mu_0)^T \right]$$

- Mixture-class scatter is the overall scatter obtained from the covariance of all samples:

$$S_m = E \left[(X - \mu_0)(X - \mu_0)^T \right] = S_w + S_b$$

where μ_0 is the overall mean.

- Optimizing criteria used to obtain the transforms is a combination of within-class scatter, between-class scatter and overall scatter.

$$criterion = \frac{tr(S_b)}{tr(S_w)} \qquad criterion = S_b - \mu(S_w - c)$$

$$criterion = inv(S_w) \times S_b \qquad criterion = inv(S_m) \times S_b$$

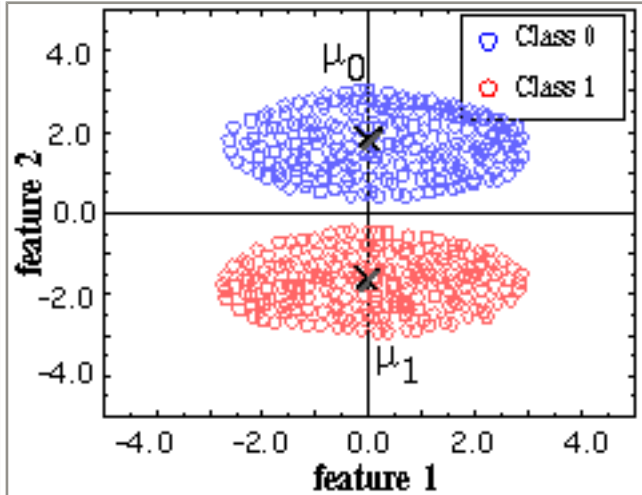
- Transformation matrix is given by:

$$T = \Phi$$

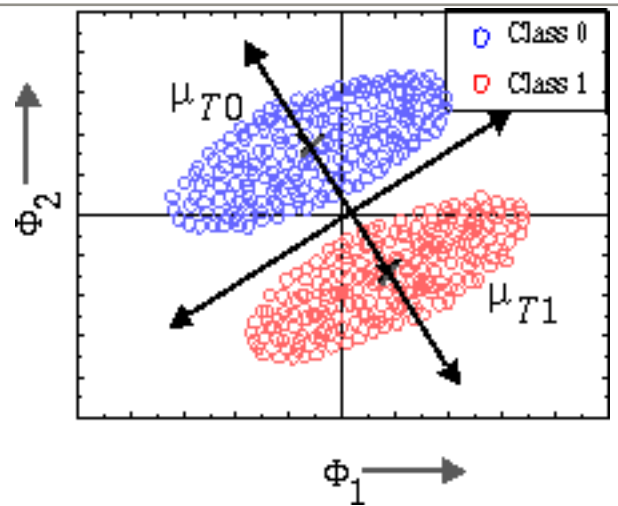
Φ

where Φ are the eigenvectors corresponding to non-zero eigenvalues.

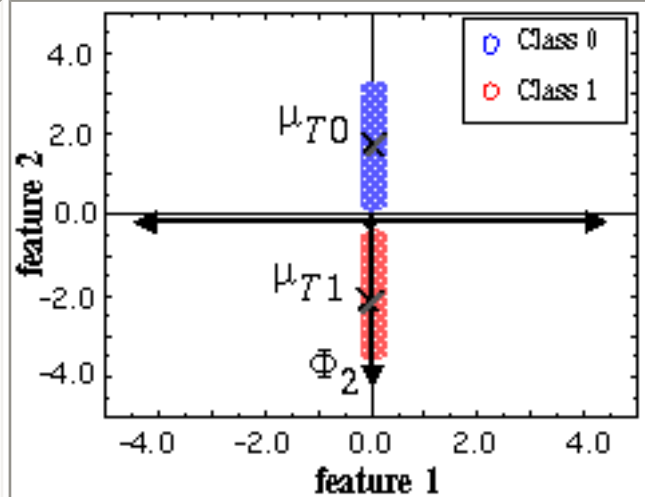
Original Data



Class-Independent PCA



Class-Independent LDA



LINEAR DISCRIMINANT ANALYSIS - A BRIEF TUTORIAL

S. Balakrishnama, A. Ganapathiraju

Institute for Signal and Information Processing
Department of Electrical and Computer Engineering
Mississippi State University
Box 9571, 216 Simrall, Hardy Rd.
Mississippi State, Mississippi 39762
Tel: 601-325-8335, Fax: 601-325-3149
Email: {balakris, ganapath}@isip.msstate.edu



1. INTRODUCTION

There are many possible techniques for classification of data. Principle Component Analysis (PCA) and Linear Discriminant Analysis (LDA) are two commonly used techniques for data classification and dimensionality reduction. Linear Discriminant Analysis easily handles the case where the within-class frequencies are unequal and their performances has been examined on randomly generated test data. This method maximizes the ratio of between-class variance to the within-class variance in any particular data set thereby guaranteeing maximal separability. The use of Linear Discriminant Analysis for data classification is applied to classification problem in speech recognition. We decided to implement an algorithm for LDA in hopes of providing better classification compared to Principle Components Analysis. The prime difference between LDA and PCA is that PCA does more of feature classification and LDA does data classification. In PCA, the shape and location of the original data sets changes when transformed to a different space whereas LDA doesn't change the location but only tries to provide more class separability and draw a decision region between the given classes. This method also helps to better understand the distribution of the feature data. Figure 1 will be used as an example to explain and illustrate the theory of LDA.

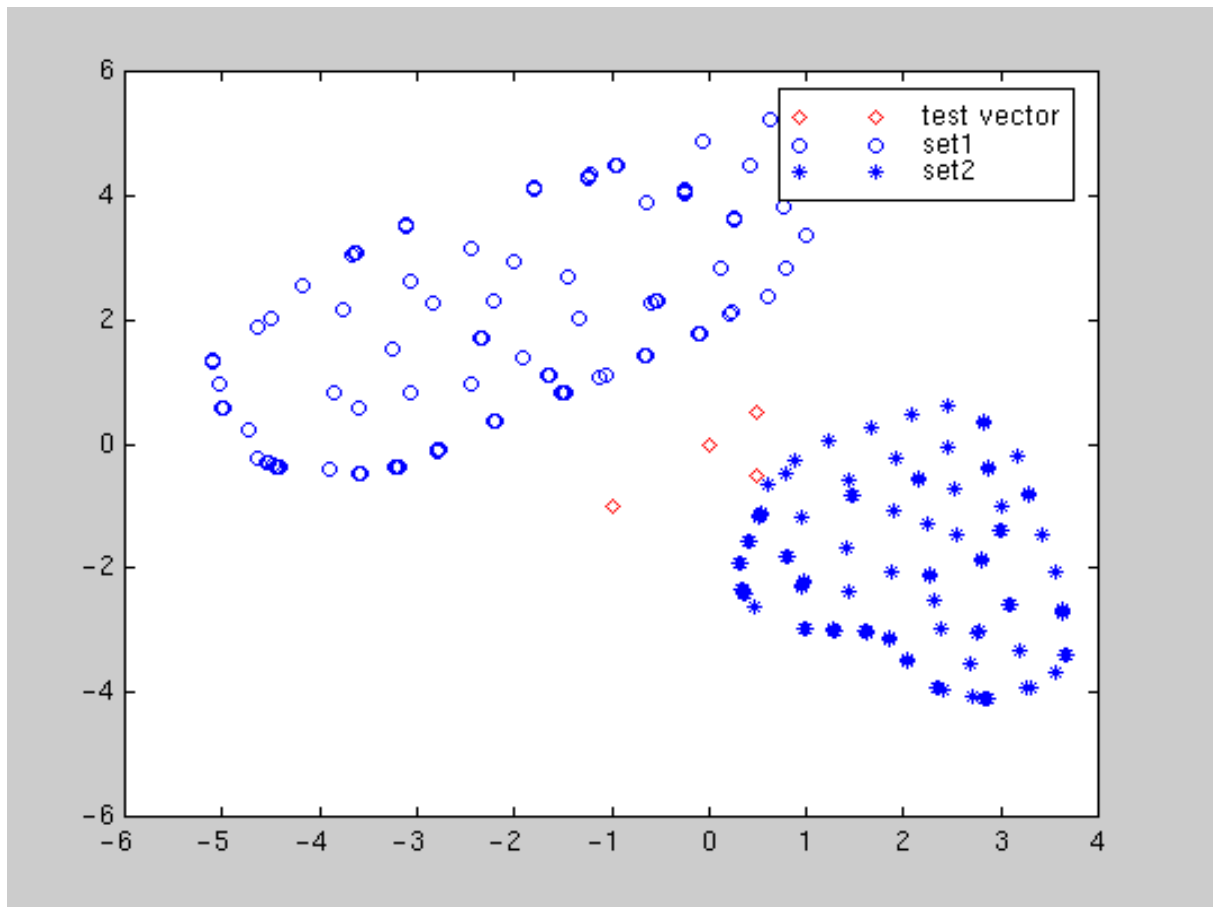


Figure 1. Figure showing data sets and test vectors in original

2. DIFFERENT APPROACHES TO LDA

Data sets can be transformed and test vectors can be classified in the transformed space by two different approaches.

Class-dependent transformation: This type of approach involves maximizing the ratio of between class variance to within class variance. The main objective is to maximize this ratio so that adequate class separability is obtained. The class-specific type approach involves using two optimizing criteria for transforming the data sets independently.

Class-independent transformation: This approach involves maximizing the ratio of overall variance to within class variance. This approach uses only one optimizing criterion to transform the data sets and hence all data points irrespective of their class identity are transformed using this transform. In this type of LDA, each class is considered as a separate class against all other classes.

3. MATHEMATICAL OPERATIONS

In this section, the mathematical operations involved in using LDA will be analyzed the aid of sample set in Figure 1. For ease of understanding, this concept is applied to a two-class problem. Each data set has 100 2-D data points. Note that the mathematical formulation of this classification strategy parallels the Matlab implementation associated with this work.

1. Formulate the data sets and the test sets, which are to be classified in the original space. The given data sets and the test vectors are formulated, a graphical plot of the data sets and test vectors for the example considered in original space is shown in Figure 1. For ease of understanding let us represent the data sets as a matrix consisting of features in the form given below:

$$set1 = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ \dots & \dots \\ \dots & \dots \\ a_{m1} & a_{m2} \end{bmatrix} \quad set2 = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ \dots & \dots \\ \dots & \dots \\ b_{m1} & b_{m2} \end{bmatrix} \quad (1)$$

2. Compute the mean of each data set and mean of entire data set. Let μ_1 and μ_2 be the mean of set 1 and set 2 respectively and μ_3 be mean of entire data, which is obtained by merging set 1 and set 2, is given by Equation 1.

$$\mu_3 = p_1 \times \mu_1 + p_2 \times \mu_2 \quad (2)$$

where p_1 and p_2 are the apriori probabilities of the classes. In the case of this simple two class problem, the probability factor is assumed to be 0.5.

3. In LDA, within-class and between-class scatter are used to formulate criteria for class separability. Within-class scatter is the expected covariance of each of the classes. The scatter measures are computed using Equations 3 and 4.

$$S_w = \sum_j p_j \times (cov_j) \quad (3)$$

Therefore, for the two-class problem,

$$S_w = 0.5 \times cov_1 + 0.5 \times cov_2 \quad (4)$$

All the covariance matrices are symmetric. Let cov_1 and cov_2 be the covariance of set 1 and set 2 respectively. Covariance matrix is computed using the following equation.

$$cov_j = (\mathbf{x}_j - \mu_j)(\mathbf{x}_j - \mu_j)^T \quad (5)$$

The between-class scatter is computed using the following equation.

$$S_b = \sum_j (\mu_j - \mu_3) \times (\mu_j - \mu_3)^T \quad (6)$$

Note that S_b can be thought of as the covariance of data set whose members are the mean vectors of each class. As defined earlier, the optimizing criterion in LDA is the ratio of between-class scatter to the within-class scatter. The solution obtained by maximizing this criterion defines the axes of the transformed space. However for the class-dependent transform the optimizing criterion is computed using equations and (5). It should be noted that if the LDA is a class dependent type, for L -class L separate optimizing criterion are required for each class. The optimizing factors in case of class dependent type are computed as

$$criterion_j = inv(cov_j) \times S_b \quad (7)$$

For the class independent transform, the optimizing criterion is computed as

$$criterion = inv(S_w) \times S_b \quad (8)$$

4. By definition, an eigen vector of a transformation represents a 1-D invariant subspace of the vector space in which the transformation is applied. A set of these eigen vectors whose corresponding eigen values are non-zero are all linearly independent and are invariant under the transformation. Thus any vector space can be represented in terms of linear combinations of the eigen vectors. A linear dependency between features is indicated by a

zero eigen value. To obtain a non-redundant set of features all eigen vectors corresponding to non-zero eigen values only are considered and the ones corresponding to zero eigen values are neglected. In the case of LDA, the transformations are found as the eigen vector matrix of the different criteria defined in Equations 7 and 8.

5. For any L -class problem we would always have $L-1$ non-zero eigen values. This is attributed to the constraints on the mean vectors of the classes in Equation 2. The eigen vectors corresponding to non-zero eigen values for the definition of the transformation.

For our 2-class example, Figures 2 and 3 show the direction of the significant eigen vector along which there is maximum discrimination information. Having obtained the transformation matrices, we transform the data sets using the single LDA transform or the class specific transforms whichever the case may be. From the figures it can be observed that, transforming the entire data set to one axis provides definite boundaries to classify the data. The decision region in the transformed space is a solid line separating the transformed data sets thus

For the class dependent LDA,

$$transformed_set_j = transform_j^T \times set_j \quad (9)$$

For the class independent LDA,

$$transformed_set = transform_spec^T \times data_set^T \quad (10)$$

Similarly the test vectors are transformed and are classified using the euclidean distance of the test vectors from each class mean.

The two Figures 4 and 5 clearly illustrate the theory of Linear Discriminant Analysis applied to a 2-class problem. The original data sets are shown and the same data sets after transformation are also illustrated. It is quite clear from these figures that transformation provides a boundary for proper classification. In this example the classes were properly defined but cases where there is overlap between classes, obtaining a decision region in original space will be very difficult and in such cases transformation proves to be very essential. Transformation along largest eigen vector axis is the best transformation.

Figures 6 and 7, are interesting in that they show how the linear transformation process can be viewed as projecting data points onto the maximally discriminating axes represented by the eigen vectors.

6. Once the transformations are completed using the LDA transforms, Euclidean distance or RMS distance is used to classify data points. Euclidean distance is computed using Equation 11 where μ_{ntrans} is the mean of the transformed data set, n is the class index and \mathbf{x} is the test vector. Thus for n classes, n euclidean distances are obtained for each test point.

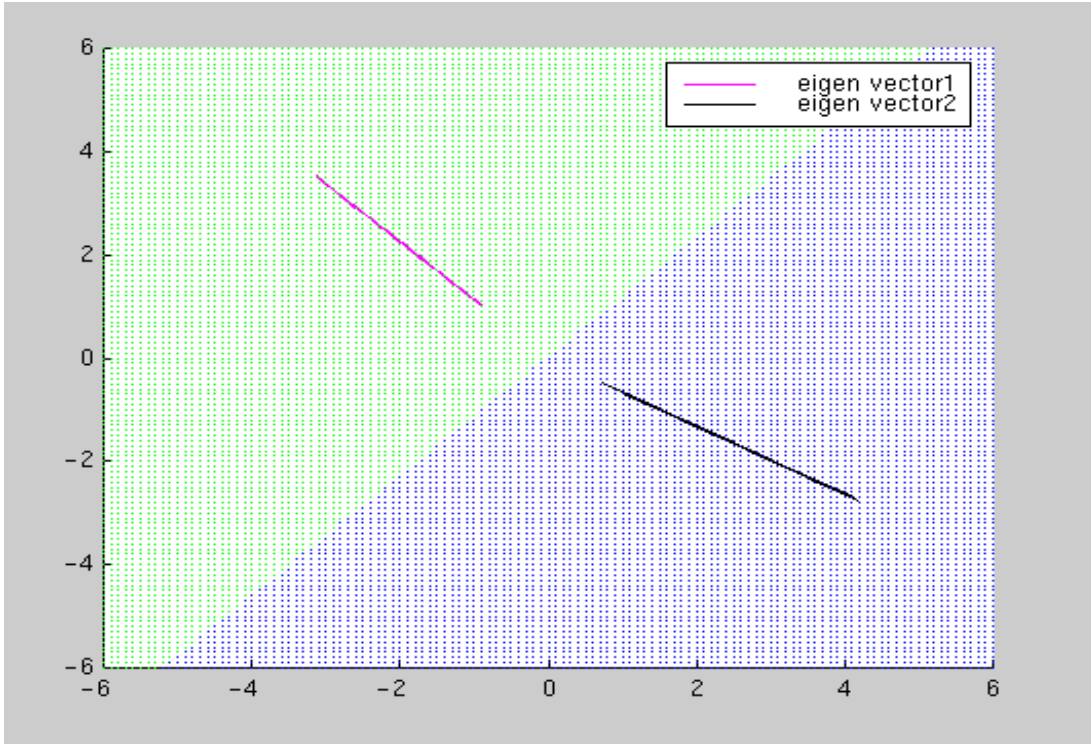


Figure 2. Figure for eigen vector direction in class dependent type

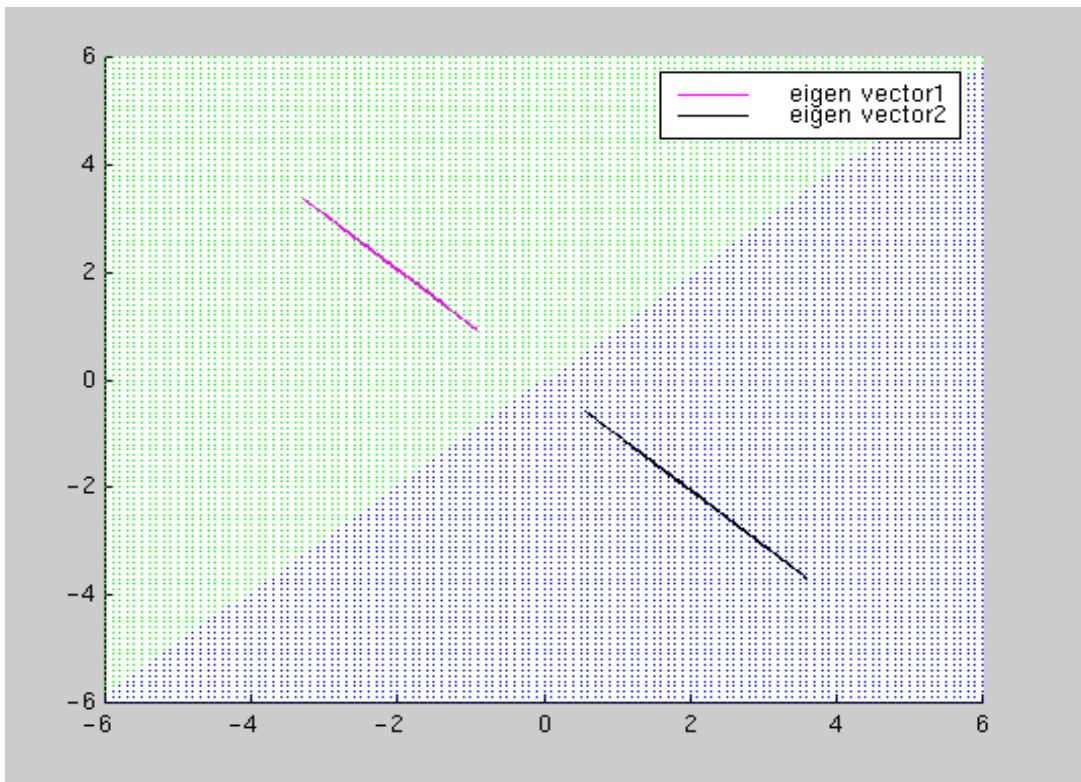


Figure 3. Figure for eigen vector direction in class independent type

$$dist_n = (transform_n_spec)^T \times \mathbf{x} - \mu_{ntrans} \quad (11)$$

7. The smallest Euclidean distance among the n distances classifies the test vector as belonging to class n .

4. CONCLUSIONS

We have presented the theory and implementation of LDA as a classification technique. Throughout the tutorial we have used a 2-class problem as an exemplar. Two approaches to LDA, namely, class independent and class dependent, have been explained. The choice of the type of LDA depends on the data set and the goals of the classification problem. If generalization is of importance, the class independent transformation is preferred. However, if good discrimination is what is aimed for, the class dependent type should be the first choice. As part of our future work, we plan to work on a Java-based demonstration which could be used to visualize LDA based transformations on user defined data sets and also help the user appreciate the difference between the various classification techniques.

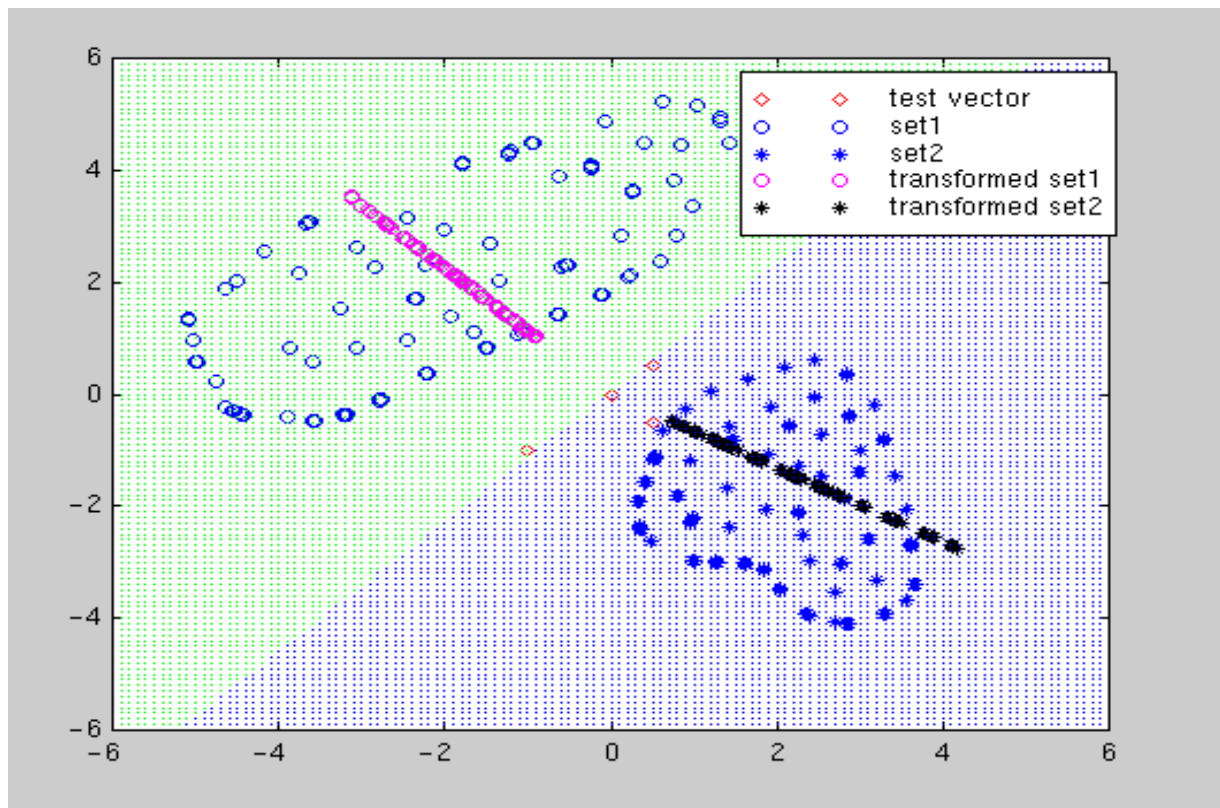


Figure 4. Data sets in original space and transformed space along with the transformation axis for class dependent LDA of a 2-class problem

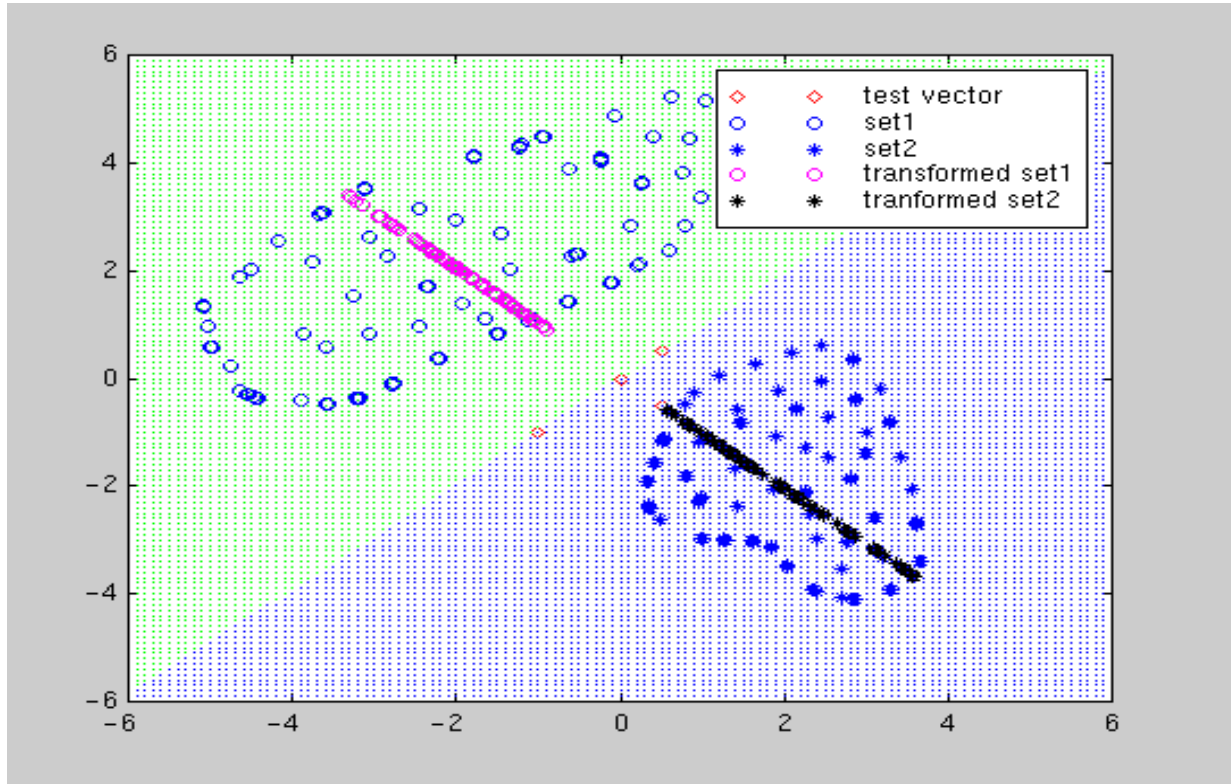


Figure 5. Data sets in original space and transformed space for class independent type of LDA of a 2-class problem

5. SOFTWARE

All Matlab code written for this project is available for public from our website at www.isip.msstate.edu

6. REFERENCES

- [1] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, Academic Press, San Diego, California, 1990.
- [2] S. Axler, *Linear Algebra Done Right*, Springer-Verlag New York Inc., New York, New York, 1995.

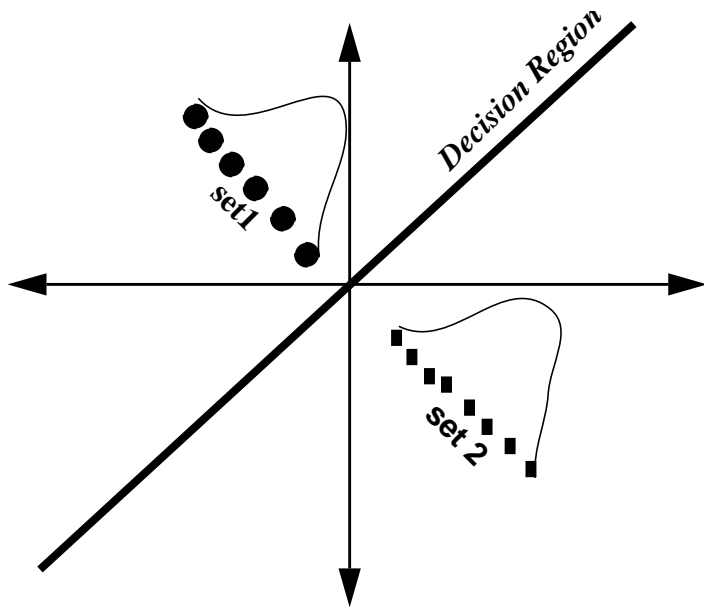


Figure 6. Figure showing histogram plot of transformed data with decision region in class independent type and the amount of class separability obtained in transformed space

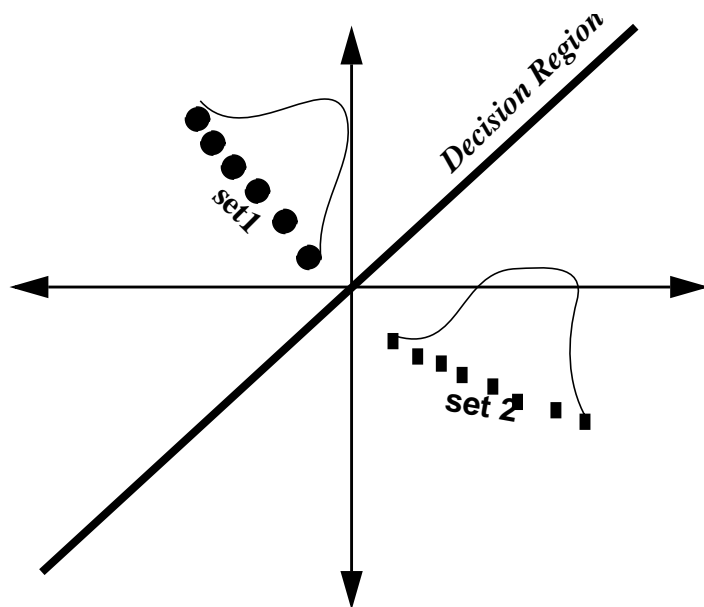
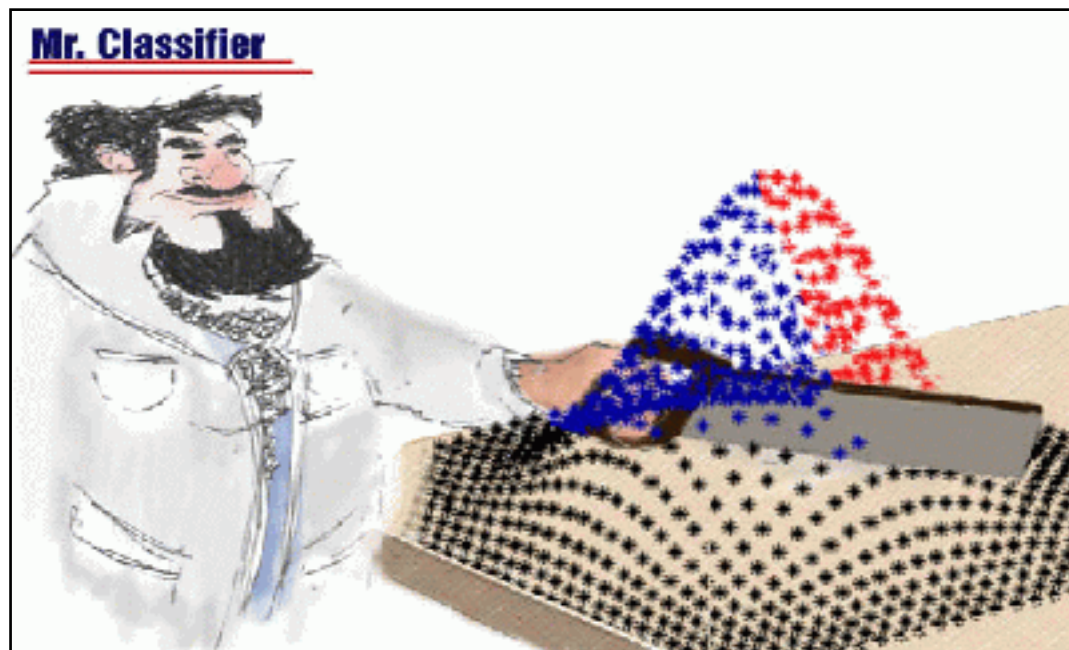


Figure 7. Histogram plot of transformed data with decision region in class dependent type and the amount of class separability obtained in transformed space



STATISTICAL NORMALIZATION FUNCTIONS FOR SIGNAL PROCESSING PROBLEMS



Suresh Balakrishnama

Candidate for Master of Science in Electrical Engineering
Institute for Signal and Information Processing
Department of Electrical and Computer Engineering
Mississippi State University, Mississippi State, MS 39762
Email: balakris@isip.msstate.edu

http://www.isip.msstate.edu/publications/seminars/masters_oral/1999/lda/index.html

Main:

[Title](#)
[Abstract](#)

Personal:

[Program](#)
[Achievements](#)

Introduction:

[Problem](#)
[Dist. Measure](#)
[Approach](#)
[Variants](#)

Multivariate Gaussian:

[Definition](#)
[Why Supp. Region?](#)
[Support Region](#)
[Identity](#)
[Unequal Variances](#)
[Off-diagonal](#)
[Unconstrained](#)

PCA:

[Math](#)
[Eigen Params](#)
[PCA Succeeds](#)
[PCA Fails](#)

LDA:

[Approach](#)
[Math](#)
[Opt. Criterion](#)
[Class-dep](#)
[LDA Vs PCA](#)

Examples:

[Uniform Ellipse](#)
[Parallel Ellipse](#)
[Class-dep](#)
[Max. Variance](#)

Application:

[Overview](#)
[Database](#)
[Algorithm](#)
[Evaluation](#)
[Best Features](#)
[Best System](#)
[Why LDA Failed?](#)

Conclusions:

[Summary](#)
[References](#)

PATTERN RECOGNITION

"All applets on our web site now require the java plugin to run. This is necessary so we can bring state-of-the-art features to you which are not currently supported by browser vendors. You can find the plugin at <http://java.sun.com/products/plugin>. For additional information or suggestions please contact help@isip.msstate.edu"
No JDK 1.2 support for APPLET!!

- [Source Code](#): Download the source code for this applet.
- [Tutorial](#): Learn how to use this applet.

All applets on our web site now require a Java plug-in. This is necessary so we can bring state-of-the-art Java features to you which are not currently supported by browser vendors such as Netscape. You can find the appropriate plug-in at <http://java.sun.com/products/plugin>. We have generated a [list of steps](#) necessary for installing the plug-in in a Unix environment. For additional information or help with your installation please contact help@isip.msstate.edu.

[Up](#) | [Home](#) | [Site Map](#) | [What's New](#) | [Projects](#) | [Publications](#)
[Speech](#) | [Administration](#) | [About Us](#) | [Search](#) | [Contact](#)

Please direct questions or comments to help@isip.msstate.edu



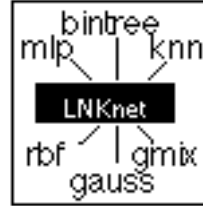
Programs
Communication Systems

- Information Systems Technology

Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington, MA 02420-9108
Phone (781) 981-5500

webmaster@ll.mit.edu

Information Systems Technology Group



LNKnet Pattern Classification Software

LNKnet is a software package developed at MIT Lincoln Laboratory which integrates more than 20 neural network, statistical, and machine learning classification, clustering, and feature selection algorithms into a modular software package.

Documentation includes a short [quick-start guide](#), a long 200-page [user's guide](#), and UNIX man pages for the many LNKnet commands.

Downloading LNKnet

LNKnet is public domain software available at no cost with a lenient [open-distribution license](#).

LNKnet software was originally developed under Sun Microsystem's Solaris 2.5.1 (SunOS 5.5.1) UNIX operating system and then ported to Solaris 2.6 (SunOS 5.6) under SUN Open Windows. It was recently ported to Red Hat Linux, and a test version is available under Linux. Source code is available both for individual programs (e.g. mlp, gauss, knn) and for the graphical user interface.

The following three gzipped tar files contain executables, man pages, documentation (the Quick-Start Guide and the LNKnet User's Guide), and sample data sets for Solaris 2.5.1 (SunOS 5.5.1), Solaris 2.6 (SunOS 5.6), and Red Hat Linux:

[LNKnet executables for Solaris 2.5.1 with Sample Data and Man Pages - gzipped tar file \(36 Mbytes\)](#)

[LNKnet executables for Solaris 2.6 with Sample Data and Man Pages - gzipped tar file \(37 Mbytes\)](#)

[LNKnet executables for Red Hat Linux with Sample Data and Man Pages - gzipped tar file \(39 Mbytes\)](#)

LNKnet now uses the GNU Autoconf tool to simplify compiling LNKnet on different host platforms. The following gzipped tar file uses Autoconf and contains source code for all LNKnet programs, man pages, documentation, and sample data sets:

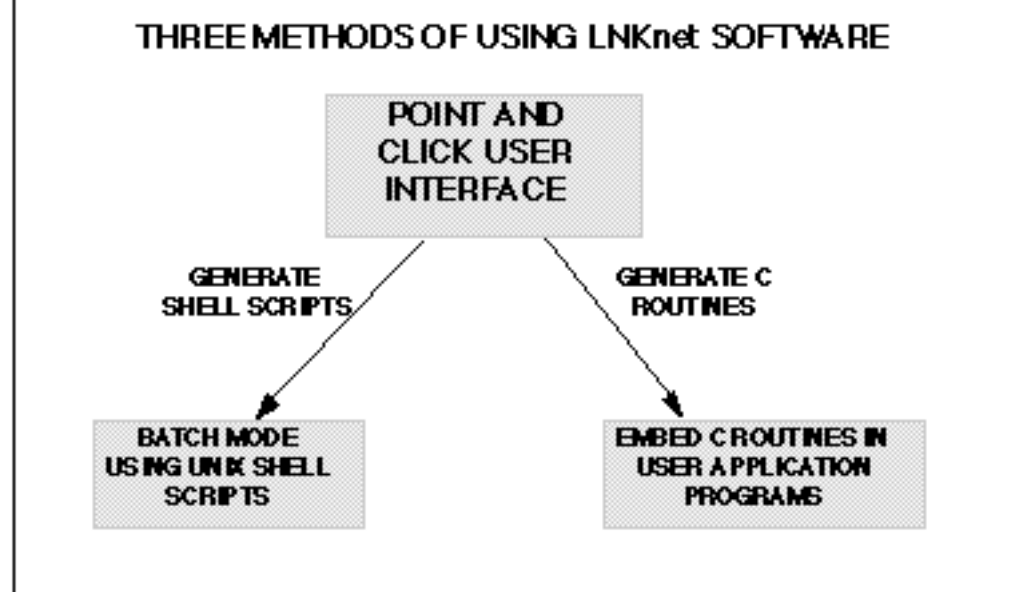
[Source Code - gzipped tar file \(8 Mbytes\)](#)

The following two pdf files contain the Quick-Start Guide and the LNKnet User's Guide:

[Quick-Start Guide - pdf file \(0.1 Mbytes\)](#)

[Users-Guide - pdf file \(5.2 Mbytes\)](#)

Methods Used to Access LNKnet Classifiers

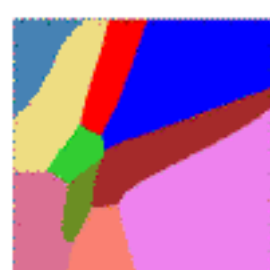


The above image shows three different methods are used to access LNKnet pattern classification software. The first and most popular approach is to take advantage of the high-level graphical user interface (GUI) to run experiments. This GUI allows a non-programmer who knows at least a little about pattern classification to run experiments by pressing buttons and making selections using a mouse. A second approach, used to create automatic batch jobs, is to run LNKnet classification driver programs from the UNIX command line or in shell scripts. A third popular approach, used by C-programmers to embed LNKnet classifiers in application programs, is to use the LNKnet GUI to automatically produce C source code which implements a trained classifier. This C-code can be copied into an application program and used with little knowledge concerning details of the classifier being used. These three interfaces have been used successfully for speech recognition, talker identification, medical risk prediction, chemical analysis, and other problems at MIT Lincoln Laboratory and many other government labs and Universities. The GUI has substantially reduced the time required to run classification experiments and explore new data sets.

Algorithms Included in LNKnet

	Supervised Training	Combine Unsupervised-Supervised Training	Unsupervised Training (Clustering)
Neural Network Algorithms	Back-Propagation (MLP) Hypersphere Classifier Perceptron Committee	Radial Basis Function (RBF) Incremental RBF Learning Vector Quantizer (LVQ) Nearest Cluster Classifier Fuzzy ARTMap Classifier	
Statistical and Machine Learning Algorithms	Gaussian Linear Discriminant Gaussian Quadratic K-Nearest Neighbor Binary Decision Tree Parzen Window Histogram	Gaussian Mixture Classifier - Diagonal/Full Covariance - Tied/Per-Class Centers	K-Means Clustering EM Clustering Leader Clustering Random Clustering
Feature Selection Algorithms	Linear Discriminant (LDA) Forward/Backward Search		Principal Components (PCA)

The above table shows the most commonly used LNKnet neural network, statistical, and machine learning classification and clustering algorithms and feature selection algorithms. Classification algorithms include: multi-layer sigmoid nets (multi-layer perceptrons) with squared-error, cross-entropy, maximum likelihood, and classification-figure-of-merit cost functions; hypersphere classifiers (sometimes called RCE classifiers); unimodal Gaussian classifier with full/diagonal and grand/per-class covariance matrices (linear discriminant and quadratic discriminant classifiers); k-nearest neighbor classifier, condensed k-nearest neighbor classifier; binary decision tree (similar to CART); radial basis function classifier, incremental adaptive center and variance radial basis function classifier with four cost functions; four versions of learning vector quantizer classifiers; nearest cluster classifier, perceptron convergence procedure; Parzen widow classifier, histogram classifier, and Gaussian mixture classifier with full/ diagonal and tied per-class covariance matrices. Clustering algorithms include: leader clustering (similar to ART); k-means clustering with binary splitting initialization; and Gaussian mixture clustering with binary splitting initialization. Feature reduction algorithms include principal components analysis and canonical linear discriminant analysis. Feature selection algorithms available with all classifiers include forward and backward search feature selection. In addition, n-fold cross validation is available with all classifiers and C source code is produced which replicates any trained classifier.



LNKnet Graphics and Plots

LNKnet produces a wide variety of plots and tables. The above plot is a decision region plot for a vowel classification problem. Plots produced include 2-D decision region plots, output profile and histogram plots, structure plots describing the computing architecture, and plots illustrating how the error rate and other cost functions vary with time. In addition, after an experiment is run, outputs include confusion matrices, tables of the error rate for each class, and overall rates.

Questions and Comments

email: KUKOLICH@SST.LL.MIT.EDU