## Unsupervised Learning with Permuted Data

Sergey Kirshner<br>SKIRSHNE@ICS.UCI.EDU<br>Sridevi Parise SPARISE@ICS.UCI.EDU<br>Padhraic Smyth<br>SMYTH@ICS.UCI.EDU

School of Information and Computer Science, University of California, Irvine, CA 92697-3425 USA


#### Abstract

We consider the problem of unsupervised learning from a matrix of data vectors where in each row the observed values are randomly permuted in an unknown fashion. Such problems arise naturally in areas such as computer vision and text modeling where measurements need not be in correspondence with the correct features. We provide a general theoretical characterization of the difficulty of "unscrambling" the values of the rows for such problems and relate the optimal error rate to the well-known concept of the Bayes classification error rate. For known parametric distributions we derive closed-form expressions for the optimal error rate that provide insight into what makes this problem difficult in practice. Finally, we show how the Expectation-Maximization procedure can be used to simultaneously estimate both a probabilistic model for the features as well as a distribution over the correspondence of the row values.


## 1. Introduction

There are a number of real-world machine learning problems that can be characterized as follows: there are $N$ objects of interest, where for each object we measure a number of features or attributes of the objects but we do not necessarily know the correspondence between the measurements for different objects. In this paper we focus specifically on the case where the feature values have been permuted in a random manner. Table 1 shows a simple example of this type of permutation problem. We would like to be able to learn the joint probability density of the original data on the left given only the permuted data and knowledge of the type of permutations that may have been

Table 1. An example of the permutation problem

| Original Data |  |  |  |  |  |  |  | Permuted Data |  |  |  | Permutations |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0.2 | 10 | 4 | 56 | 10 | 4 | 56 | 0.2 | $(2,3,4,1)$ |  |  |  |  |
| 0.1 | 12 | 4 | 62 | 62 | 0.1 | 4 | 12 | $(4,1,3,2)$ |  |  |  |  |
| 0.1 | 11 | 3 | 70 | 11 | 3 | 0.1 | 70 | $(2,3,1,4)$ |  |  |  |  |
| 0.3 | 14 | 4 | 61 | 61 | 0.3 | 14 | 4 | $(4,1,2,3)$ |  |  |  |  |

applied to the data (e.g., cyclic shifts). Two questions naturally arise: (a) how hard is this type of learning problem in general? and (b) what kinds of algorithms can we use to solve this problem in practice?
In considering the first problem, our intuition tells us that the "more different" the features in the original (unpermuted) table are then the "easier" the unscrambling problem may be. For example, in Table 1, the distributions of each individual feature in the table on the left appear quite different from each other, so that one hopes that given enough data one could eventually recover a model for the original data given only permuted data. In Section 2 we make this notion of learnability precise by introducing the notion of a Bayes-optimal permutation error rate. In Section 3 we show that under certain conditions this error rate is upper-bounded by an appropriately defined Bayesoptimal classification error rate, confirming the intuition that the ability to unmix the row-values should be related to the overlap of the column densities. In Section 4 we derive closed-form expressions for the permutation error rate for specific parametric models finding for example that negative correlation among Gaussian columns can make the unmixing problem significantly harder, while positive correlation can make it arbitrarily easy. In Section 5 we address the second question (how to simultaneously unlearn the mixing and estimate the original joint density) using an EM
framework, and provide various experimental results to illustrate how these learning algorithms work. Conclusions are presented in Section 6.
Our interest in this "learning from permuted data" problem is motivated by recent work in applying machine learning techniques to astronomical image data (Kirshner et al., 2003a). In this problem each image consists of three intensity "blobs" that represent a particular type of galactic structure of interest to astronomers. For each blob a vector of features can be extracted, such as mean intensity, ellipticity, and so forth. Astronomers can visually identify a central "core" blob, and right and left "lobe" blobs in each image. However, in the training data the groups of features are not associated with any labels that identify whether they are from the center, left, or right blob this information is hidden and must be estimated from the data. This is a version of the permutation problem described earlier. In prior work (Kirshner et al., 2003a) we developed an EM algorithm to solve this image analysis problem and focused on domain-specific aspects of the astronomy application.
Permutation problems frequently occur in computer vision where, for example, features in the form of landmarks are calculated for an object of interest in an image (such as a face) but the features are not necessarily in correspondence across different images (Gold et al., 1995). Similar problems also arise in language modeling and information extraction, e.g., in learning models documents, where different text fields can occur in different positions on the page (McCallum et al., 2000). A significant advance in recent years is the realization that many of these types of correspondence issues can be cast as machine learning problems, viewing the unknown correspondences as hidden variables that can be estimated from the data using techniques such as Expectation-Maximization (EM), e.g., in vision (Frey \& Jojic, 2003). Much of this prior work takes advantage of domain-specific information to help solve the correspondence problem, e.g., the use of prior knowledge of likely types of spatial deformations in images, or sequential constraints on text formation in information extraction. In contrast, in this paper we focus on a more abstract theoretical characterization of learning from permuted data. The primary novel contribution of this paper is the introduction and analysis of the notion of Bayes-optimal error rates for "unscrambling" the permutations, providing a lower bound on the performance of any unsupervised learning algorithm for this problem.

## 2. Probabilistic Generative Models

### 2.1. Notation

Let $\mathbf{x}=\left(\vec{x}_{1}, \ldots, \vec{x}_{i}, \ldots, \vec{x}_{c}\right)$ be composed of $c$ featurevectors $\vec{x}_{i}$, with each vector $\vec{x}_{i}$ taking values from the same $d$-dimensional set S . Thus, $\mathbf{x}$ has dimension $c \times d$ and takes values in the set $S^{c}$.
Example 1. Let each $\vec{x}_{i}$ be a one-dimensional realvalued feature-vector. In this case S is the real line $(S=\mathbb{R})$ with $d=1$, each $\vec{x}_{i}$ is a scalar, and $\mathrm{S}^{c}=\mathbb{R}^{c}$.
Example 2. Consider the case where $A=\{\mathrm{a}, \mathrm{b}, \ldots, \mathrm{z}\}$ is a set of letters in an alphabet, and let $S=A \times \mathbb{R}$. In this case each feature-vector takes values as pairs of a letter and a real number, i.e. $d=2$, and the space $\mathrm{S}^{c}$ is $2 c$-dimensional.

We define $p(\mathbf{x})$ as a probability density (distribution) function over the set $S^{c}$. For example, in Example 1 above $p(\mathbf{x})$ could be a $c$-dimensional multivariate Gaussian density.

### 2.2. A Generative Model for Permuted Data

Our generative model consists of two parts:
In the first part we generate samples from $p(\mathbf{x})$ in a standard manner - throughout this paper we assume independent and identically distributed random samples. In this manner we can generate a data matrix of $N$ rows and $c$ columns, where each column has dimension $d$.

The second part of the generative model randomly applies a permutation $\rho$ to each row of the data matrix in the following manner. Let $\mathcal{P}=\left\{\rho_{1}, \ldots, \rho_{m}\right\}$ be a set of permutations defined on $(1, \ldots, c)$. For example, $\mathcal{P}$ could be the set of all $c$ cyclic shifts, e.g., $c=3, \rho_{1}=(1,2,3), \rho_{2}=(2,3,1)$, and $\rho_{3}=(3,1,2)$. For each row $\mathbf{x}$ of the data matrix a permutation $\rho \in \mathcal{P}$ is randomly selected according to a probability distribution $p(\rho)$ over $\mathcal{P}$. The components of $\mathbf{x}$ are then permuted according to the selected $\rho$ to obtain a permuted vector taking values in the same set $S^{c}$.
The size of $\mathcal{P},|\mathcal{P}|$, is denoted by $m$. Note that if all possible permutations are allowed then $m=c$ !. Unless stated otherwise, in this paper we will generally assume that all permutations in $\mathcal{P}$ are equally likely, i.e. $p\left(\rho_{j}\right)=\frac{1}{m}, j=1, \ldots, m$.

### 2.3. Probability Densities for Permuted Data

It is useful to express the probability density of a permuted vector, call it $q(\mathbf{x})$, as a function of (a) the density of the original data rows $p(\mathbf{x})$, and (b) the distribution over permutations $p(\rho)$. In the remainder

Table 2. Probability distributions for Example 4.

| $\mathbf{x}$ | $p(\mathbf{x})$ | $q(\mathbf{x})$ |
| :---: | :---: | :---: |
| $(0,0)$ | 0.28 | 0.28 |
| $(0,1)$ | 0.42 | 0.27 |
| $(1,0)$ | 0.12 | 0.27 |
| $(1,1)$ | 0.18 | 0.18 |

of the paper whenever the symbol $q$ is used it is implicitly assumed that the argument $\mathbf{x}$ has been permuted. Note that $q(\mathbf{x})$ can be expressed as a finite mixture over all $m$ possible permutations that could have led to the generation of $\mathbf{x}$ :

$$
\begin{align*}
q(\mathbf{x}) & =\sum_{j=1}^{m} q\left(\mathbf{x} \mid \rho_{j}\right) p\left(\rho_{j}\right) \\
& =\sum_{j=1}^{m} q\left(\vec{x}_{1}, \ldots, \vec{x}_{c} \mid \rho_{j}\right) p\left(\rho_{j}\right) \\
& =\sum_{j=1}^{m} p\left(\rho_{j}^{-1}\left(\vec{x}_{1}\right), \ldots, \rho_{j}^{-1}\left(\vec{x}_{c}\right)\right) p\left(\rho_{j}\right) \\
& =\sum_{j=1}^{m} p\left(\vec{x}_{\rho_{j}^{-1}(1)}, \ldots, \vec{x}_{\rho_{j}^{-1}(c)}\right) p\left(\rho_{j}\right) \tag{1}
\end{align*}
$$

where $\rho_{j}^{-1}$ is the unique inverse permutation for $\rho_{j}$.
Example 3. Let $c=3$, and $\mathcal{P}=\left\{\rho_{1}, \rho_{2}, \rho_{3}\right\}$ be a set of cyclic shifts: $\rho_{1}=(1,2,3), \rho_{2}=(2,3,1)$, and $\rho_{3}=$ $(3,1,2)$. If $\mathbf{x}=\left(\vec{x}_{1}, \vec{x}_{2}, \vec{x}_{3}\right)$ is a permuted vector, it could have been obtained from one of three possible permutations, as reflected by the mixture model:

$$
\begin{aligned}
q(\mathbf{x})= & q\left(\vec{x}_{1}, \vec{x}_{2}, \vec{x}_{3} \mid \rho_{1}\right) p\left(\rho_{1}\right) \\
& +q\left(\vec{x}_{1}, \vec{x}_{2}, \vec{x}_{3} \mid \rho_{2}\right) p\left(\rho_{2}\right) \\
& +q\left(\vec{x}_{1}, \vec{x}_{2}, \vec{x}_{3} \mid \rho_{3}\right) p\left(\rho_{3}\right) \\
= & p\left(\vec{x}_{1}, \vec{x}_{2}, \vec{x}_{3}\right) p\left(\rho_{1}\right)+p\left(\vec{x}_{3}, \vec{x}_{1}, \vec{x}_{2}\right) p\left(\rho_{2}\right) \\
& +p\left(\vec{x}_{2}, \vec{x}_{3}, \vec{x}_{1}\right) p\left(\rho_{3}\right) .
\end{aligned}
$$

An important point is that $p$ and $q$ are not the same distribution although both are defined over $\mathrm{S}^{c}$.
Example 4. Let $\mathrm{S}=\{0,1\}$ with $c=2$. Let $\mathcal{P}=$ $\left\{\rho_{1}, \rho_{2}\right\}$ where $\rho_{1}=(1,2)$ and $\rho_{2}=(2,1)$. Let $p(\mathbf{x})$ be as defined in the Table 2. Assume $p\left(\rho_{1}\right)=p\left(\rho_{2}\right)=$ 0.5 . The resulting $q$ distribution on permuted vectors $\mathbf{x}$ is also listed in the Table 2. $p \neq q$ since, for example,

$$
\begin{aligned}
q(0,1) & =\sum_{j=1}^{2} p\left(0_{\rho_{j}^{-1}(1)}, 1_{\rho_{j}^{-1}(2)}\right) p\left(\rho_{j}\right) \\
& =p(0,1) \times 0.5+p(1,0) \times 0.5 \\
& =0.42 \times 0.5+0.12 \times 0.5 \\
& =0.27 \neq 0.42=p(0,1)
\end{aligned}
$$

### 2.4. Inference and Learning

There are two problems of direct interest. In the first problem, we assume that $p(\mathbf{x})$ is known, and that the set of permutations $\mathcal{P}$ and their probabilities $p(\rho)$ are also known. Then, given a permuted vector $\mathbf{x}$, we can calculate

$$
q\left(\rho_{j} \mid \mathbf{x}\right)=\frac{q\left(\mathbf{x} \mid \rho_{j}\right) p\left(\rho_{j}\right)}{\sum_{i} q\left(\mathbf{x} \mid \rho_{i}\right) p\left(\rho_{i}\right)}, \quad j=1, \ldots, m
$$

using the mixture model in Equation 1 and Bayes rule. This allows us to identify (for example) the most likely permutation, $\arg \max _{j} q\left(\rho_{j} \mid \mathbf{x}\right)$. This decision rule is Bayes-optimal in that no other decision rule can achieve a lower average error in terms of identifying the permutations. Of interest here is the probability that we make an error (on average) using this decision rule, i.e., what is the optimal error rate achievable in terms of unscrambling the row values. Here an "error" occurs whenever the most likely permutation is not the same as the true permutation that generated x. We will refer to this error rate as the Bayes-optimal permutation error rate, defined as

$$
E_{P}^{\star}=\int_{\mathrm{S}^{c}} q(\mathbf{x}) \times\left(1-\max _{j} q\left(\rho_{j} \mid \mathbf{x}\right)\right) \mathrm{d} \mathbf{x}
$$

with the superscript $\star$ referring to "Bayes-optimal" and subscript $P$ referring to "permutation".
In the second problem, we are given the set of permuted vectors $D=\left\{\mathbf{x}^{1}, \ldots, \mathbf{x}^{N}\right\}$, a set of permutations $\mathcal{P}$, and an assumed functional form for $p(\mathbf{x})$. Our task in this case is to estimate the parameters of $p(\mathbf{x})$, the probabilities $p(\rho)$, and for each row $\mathbf{x}$ to estimate the probability that permutation $\rho_{j}$ was used to generate it. This problem is discussed in Section 5.
The two problems above are intimately related. The error rate of any learning algorithm in problem 2 (in terms of identifying permutations) will be lowerbounded by the Bayes-optimal permutation error rate $E_{P}^{\star}$ as defined in problem 1 . Thus, $E_{P}^{\star}$ is a fundamental characteristic of the difficulty of learning in the presence of permutations and it is of direct interest to study it. In what follows we first show in Section 3 that $E_{P}^{\star}$ is it itself upper-bounded (under certain assumptions) by a well-known characteristic of density overlap (the Bayes-optimal classification error rate), and we then in Section 4 derive closed form expressions for $E_{P}^{\star}$ for specific simple forms for $p(\mathbf{x})$.

## 3. Analysis of Bayes-Optimal Error Rates

Recall that $\mathbf{x}=\left(\vec{x}_{1}, \ldots, \vec{x}_{c}\right)$. We can define a marginal distribution for each feature-vector $\vec{x}_{i}, i=1, \ldots, c$ as

$$
\begin{aligned}
p\left(\vec{x}_{i}\right)= & \int_{\mathrm{S}} \ldots \int_{\mathrm{S}} p\left(\vec{x}_{1}, \ldots, \vec{x}_{i-1}, \vec{x}, \vec{x}_{i+1}, \ldots, \vec{x}_{c}\right) \\
& \mathrm{d} \vec{x}_{1} \ldots \mathrm{~d} \vec{x}_{i-1} \mathrm{~d} \vec{x}_{i+1} \ldots \mathrm{~d} \vec{x}_{c}
\end{aligned}
$$

i.e., the marginal density for $\vec{x}_{i}$ defined on the set S . Each of the $c$ features has a similar marginal density on the same set S . We will use $p\left(\vec{x} \mid C_{i}\right)=p\left(\vec{x}_{i}\right)$ to denote the marginal density of $\vec{x}_{i}$ on the set S .

We now have $c$ different densities defined on S , which in turn defines a finite mixture $p_{M}(\vec{x})$ on S :

$$
p_{M}(\vec{x})=\sum_{i=1}^{c} p\left(\vec{x} \mid C_{i}\right) \times p\left(C_{i}\right)
$$

where $p\left(C_{i}\right)=\frac{1}{c}$, since all marginals have equal weight in the process of defining the mixture. In the space $S$ consider a classification problem with $c$ classes, where, given a measurement $\vec{x} \in \mathrm{~S}$, we infer the most likely feature-vector $\vec{x}_{j}$ that it originated from, $j=1, \ldots, c$. The Bayes-optimal classification rate for this problem is defined as

$$
E_{C}^{\star}=\int_{\mathrm{S}} p_{M}(\vec{x}) \times\left(1-\max _{j} p\left(C_{j} \mid \vec{x}\right)\right) \mathrm{d} \vec{x}
$$

Intuitively, $E_{C}^{\star}$ is the error rate obtained if we were given vectors $\vec{x}_{i}$ one at a time, and asked to identify which of the $c$ "columns" they originated from, based on knowing each of the $p\left(\vec{x} \mid C_{i}\right)$ densities, $i=1, \ldots, c$. $E_{C}^{\star}$ is proportional to the overlap of the individual feature densities $p\left(\vec{x}_{i}\right)$ in the space S . For example, for the data on the left in Table 1 we would expect the overlap of the 4 densities, as reflected by $E_{C}^{\star}$, to be quite small. Furthermore, we would expect intuitively that the permutation error rate $E_{P}^{\star}$ should also be low in this case, and more generally that it should be related to $E_{C}^{\star}$ in some manner. In what follows below we quantify this intuition. Specifically, under certain choices of $\mathcal{P}$, we show that $E_{P}^{\star}$ is upper-bounded by $E_{C}^{\star}$.
Definition 1. For $\mathcal{P}$, let $k$ be a key index if $\left(\rho_{1}(k), \ldots, \rho_{m}(k)\right)$ is a permutation of $(1, \ldots, c)$.

Note that $\mathcal{P}$ having a key implies $m=|\mathcal{P}|=c$. The set of all cyclic shifts for example has a key.
Example 5. For $\mathcal{P}=\left\{\rho_{1}, \rho_{2}\right\}$ with $\rho_{1}=(1,2)$ and $\rho_{2}=(2,1)$, both indices 1 and 2 are keys.
Example 6. A set of permutations $\mathcal{P}=$ $\left\{\rho_{1}, \rho_{2}, \rho_{3}, \rho_{4}\right\}$ with $\rho_{1}=(1,2,3,4), \rho_{2}=(2,1,3,4)$,
$\rho_{3}=(1,2,4,3)$, and $\rho_{4}=(2,1,4,3)$ does not have a key.
Theorem 1. If a set of permutations $\mathcal{P}$ for a permutation problem has a key, and if each permutation is equally likely, then

$$
E_{p}^{*} \leq E_{c}^{*}
$$

The proof is omitted in this version of the paper due to lack of space, but can be found online in a longer version of the paper (Kirshner et al., 2003b). Note that if the set of permutations $\mathcal{P}$ does not have a key, $E_{P}^{\star}$ may in fact be larger than $E_{C}^{\star}$.

The theorem shows that under certain assumptions, the permutation problem is easier than a corresponding version of the classification problem. This generally agrees with our intuition since in the classification version of the problem we are classifying feature values one at a time in terms of which column they are thought to have originated from, whereas in the permutation version of the problem we are simultaneously classifying $c$ values together and have the additional information available that the $c$ values must all be assigned to different classes.

## 4. Analysis of Permutation Error Rates

In this section we derive closed-form expressions for the Bayes-optimal permutation error rate $E_{P}^{\star}$ for specific functional forms for $p(\mathbf{x})$ and we use these expressions to show how changes in the parameters of $p(\mathbf{x})$ can make learning and inference either harder or easier.

### 4.1. Gaussian Features

We begin with the case of Gaussian features, since the Gaussian model is both amenable to analysis and widely used in practice.

### 4.1.1. Case 1: Two Independent Features

Consider the case when $c=2, d=1$ and $\mathrm{S}=\mathbb{R}$. Let $p(\mathbf{x})$ be a Gaussian with covariance matrix of the form $\sigma^{2} I$ where $I$ is the identity matrix (the features are independent Gaussians with equal variances). Thus,

$$
p(\mathbf{x})=\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Sigma})
$$

where,

$$
\boldsymbol{\mu}=\left[\begin{array}{l}
\mu_{1} \\
\mu_{2}
\end{array}\right] \text { and } \boldsymbol{\Sigma}=\left(\begin{array}{cc}
\sigma^{2} & 0 \\
0 & \sigma^{2}
\end{array}\right)
$$

We have $m=2$ with $\rho_{1}=(1,2)$ and $\rho_{2}=(2,1)$.

$$
q\left(\mathbf{x} \mid \rho_{1}\right)=p\left(x_{1}, x_{2}\right)=\mathcal{N}\left(\left(x_{1}, x_{2}\right) \mid \boldsymbol{\mu}, \mathbf{\Sigma}\right)
$$



Figure 1. Simulated data (and covariance ellipses) from $q\left(\mathbf{x} \mid \rho_{1}\right)$ (x's) and $q\left(\mathbf{x} \mid \rho_{2}\right)$ (o's) for the correlated Gaussian case with equal variances (case 2). The optimal decision boundary corresponds to the line $x_{1}=x_{2}$.
and

$$
q\left(\mathbf{x} \mid \rho_{2}\right)=p\left(x_{2}, x_{1}\right)=\mathcal{N}\left(\left(x_{1}, x_{2}\right) \mid \tilde{\boldsymbol{\mu}}, \boldsymbol{\Sigma}\right)
$$

where

$$
\tilde{\boldsymbol{\mu}}=\left[\begin{array}{l}
\mu_{2} \\
\mu_{1}
\end{array}\right]
$$

It is straightforward to show that

$$
\begin{equation*}
E_{P}^{*}=\frac{1}{\sqrt{2 \pi}} \int_{\frac{r}{2}}^{\infty} e^{-u^{2} / 2} \mathrm{~d} u \tag{2}
\end{equation*}
$$

where

$$
r^{2}=(\boldsymbol{\mu}-\tilde{\boldsymbol{\mu}})^{T} \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}-\tilde{\boldsymbol{\mu}})
$$

Therefore, given the functional forms for $q\left(\mathbf{x} \mid \rho_{1}\right)$ and $q\left(\mathbf{x} \mid \rho_{2}\right)$,

$$
E_{P}^{*}=\frac{1}{\sqrt{2 \pi}} \int_{\frac{\left|\mu_{1}-\mu_{2}\right|}{\sqrt{2} \sigma}}^{\infty} e^{-u^{2} / 2} \mathrm{~d} u
$$

The quantity $\left|\mu_{1}-\mu_{2}\right| / \sigma$ is a measure of the overlap of the two Gaussians: as overlap increases $E_{P}^{*}$ decreases, and vice-versa. This is exactly the same qualitative behavior as one gets with the Bayes-optimal classification error rate $E_{C}^{*}$ for this problem, except that the range of integration is different. Specifically (using the results for $E_{C}^{*}$ in (Duda et al., 2000)) we have

$$
E_{C}^{*}-E_{P}^{*}=\frac{1}{\sqrt{2 \pi}} \int_{\frac{\left|\mu_{1}-\mu_{2}\right|}{2 \sigma}}^{\frac{\left|\mu_{1}-\mu_{2}\right|}{\sqrt{2} \sigma}} e^{-u^{2} / 2} \mathrm{~d} u
$$

In the cases of both maximal overlap $\left(\left|\mu_{1}-\mu_{2}\right| / \sigma\right.$ is very large) and minimal overlap (the overlap expression is very small) the difference in the two error rates is very small. The difference between the two types of error is maximized when $\frac{\left|\mu_{2}-\mu_{1}\right|}{2 \sigma}=\sqrt{\ln 2}$.

### 4.1.2. Case 2: Two Correlated Gaussian Features

Next consider a generalization of Case 1 where the features are no longer assumed to be independent but are allowed to have non-zero correlation $\nu$ :

$$
p(\mathbf{x})=\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Sigma})
$$

where,

$$
\boldsymbol{\mu}=\left[\begin{array}{l}
\mu_{1} \\
\mu_{2}
\end{array}\right] \text { and } \boldsymbol{\Sigma}=\left(\begin{array}{cc}
\sigma^{2} & \nu \\
\nu & \sigma^{2}
\end{array}\right),-\sigma^{2}<\nu<\sigma^{2}
$$

$q\left(\mathbf{x} \mid \rho_{1}\right)$ and $q\left(\mathbf{x} \mid \rho_{2}\right)$ are defined as in Case 1 , but $\boldsymbol{\Sigma}$ now has a covariance term $\nu$ in the off-diagonal positions. Using Equation 2 again, we get

$$
E_{P}^{*}=\frac{1}{\sqrt{2 \pi}} \int_{\frac{\left|\mu_{1}-\mu_{2}\right|}{\sqrt{2} \sqrt{\sigma^{2}-\nu}}}^{\infty} e^{-u^{2} / 2} \mathrm{~d} u
$$

Thus, as in the independent case, $E_{P}^{*}$ decreases as $\mid \mu_{1}-$ $\mu_{2} \mid$ increases and vice-versa.
As $\sigma \rightarrow \nu$ the lower limit of the integral approaches $\infty$ and $E_{P}^{*}$ approaches zero. Thus, even though the two Gaussians could be heavily overlapped, as the correlation approaches 1 we can identify permuted pairs of values with accuracy approaching 1, in contrast to the Bayes-optimal classification error rate for the same problem which is defined based on classifying each value separately and cannot take advantage of the correlation information. This is a case where the permutation error rate can approach 0 even in cases where the classification error rate (proportional to the overlap of feature densities) can be as high as 0.5 .


Figure 2. Simulated data (and covariance ellipses) from $q\left(\mathbf{x} \mid \rho_{1}\right)$ (x's) and $q\left(\mathbf{x} \mid \rho_{2}\right)$ (o's) for the correlated Gaussian case with unequal variances (case 3). The optimal decision boundary corresponds to the line $x_{1}=x_{2}$.


Figure 3. Empirical values of the error rates for Case 3

Interestingly, negative correlation has the opposite effect in that as the correlation coefficient becomes more negative, $E_{P}^{*}$ increases and approaches $E_{C}^{*}$. Intuitively, negative correlation makes the problem harder by effectively leading to more overlap between the two densities. To see this visually, in Figure 1 we plot simulated data from $q(\mathbf{x})$ for the case of very negative correlation, zero correlation, and very positive correlation.

### 4.1.3. Case 3: Unequal Variances

Consider further the case where the Gaussian features have unequal variances $\sigma_{1}^{2}$ and $\sigma_{2}^{2}$ and covariance $\nu$. Since $q\left(\mathbf{x} \mid \rho_{1}\right)$ and $q\left(\mathbf{x} \mid \rho_{2}\right)$ have unequal covariance ma-
trices, there is no closed-form expression for $E_{P}^{*}$ or $E_{C}^{*}$ as before. Nevertheless, we examine the variation of $E_{P}^{*}$ as a function of $\sigma_{1}, \sigma_{2}$ and $\nu$ via simulations. Figure 2 shows some 2-D plots for various values of $\nu$, keeping $\sigma_{1}$ and $\sigma_{2}$ fixed. Variance inequality changes the nature of the overlap between $q\left(\mathbf{x} \mid \rho_{1}\right)$ and $q\left(\mathbf{x} \mid \rho_{2}\right)$ as compared to the equal variance case in Figure 1.

We can also calculate empirical error rates (in terms of classifying permutations) using the true model and varying the true parameters (note that these are empirical estimates of $E_{P}^{*}$ by definition). Figure 3 shows the variation of these empirical error rates with $\nu$ for different values of the variances keeping the ratio $\sigma_{1} \sigma_{2}$ constant. $E_{P}^{*}$ depends on both $\nu$ as well as the difference $\left|\sigma_{1}^{2}-\sigma_{2}^{2}\right|$, and the variation of the error rate with $\nu$ is not monotonic.

### 4.2. Categorical Data

Consider the simple case when $c=2$ and $d=1$, i.e., $\mathbf{x}$ consists of two scalar features. Assume that the features are discrete and can take one of $V$ values. Let $m=2$ with $\rho_{1}=(1,2)$ and $\rho_{2}=(2,1)$, and both permutations are assumed to be equally likely. We have,

$$
\begin{aligned}
E_{P}^{\star} & =\sum_{\mathbf{x}} q(\mathbf{x})\left(1-\max _{i} q\left(\rho_{i} \mid \mathbf{x}\right)\right) \\
& =\sum_{\mathbf{x}} q(\mathbf{x})-\sum_{\mathbf{x}} q(\mathbf{x}) \max _{i} q\left(\rho_{i} \mid \mathbf{x}\right) \\
& =1-\frac{1}{2} \sum_{\mathbf{x}} \max _{i} q\left(\mathbf{x} \mid \rho_{i}\right)
\end{aligned}
$$

$$
\begin{aligned}
& =1-\frac{1}{2} \sum_{x_{1}, x_{2}} \max \left\{p\left(x_{1}, x_{2}\right), p\left(x_{2}, x_{1}\right)\right\} \\
& =\frac{1}{2}-\frac{1}{4} \sum_{x_{1}, x_{2}}\left|p\left(x_{1}, x_{2}\right)-p\left(x_{2}, x_{1}\right)\right|
\end{aligned}
$$

Thus, $E_{P}^{\star}$ is a function of the quantity $\sum_{x_{1}, x_{2}}\left|p\left(x_{1}, x_{2}\right)-p\left(x_{2}, x_{1}\right)\right|$, which we can call the permutation distance between $\rho_{1}$ and $\rho_{2}$. $E_{P}^{\star}$ decreases linearly with this distance, reflecting the fact that the more dissimilar the probabilities of each permuted pair of values are from the probabilities of the unpermuted pairs, the more $E_{P}^{\star}$ decreases.

## 5. Learning from Permuted Data

In this section we briefly comment on the problem of unsupervised learning with permuted data. Space limitations do not permit a complete treatment of the topic: the goal here is to illustrate that learning with permutations can be achieved in a practical sense, and to demonstrate that the Bayes-optimal permutation error rate provides an absolute lower bound on the error rate of practical learning algorithms. We note in passing that there can be identifiability issues with mixtures of permutations that appear to be quite similar to those for standard finite mixtures, e.g., certain classes of distributions on categorical data cannot be identified uniquely.

### 5.1. EM Algorithms for Permuted Data

Consider a "matrix" data set $D=\left\{\mathbf{x}^{1}, \ldots, \mathbf{x}^{N}\right\}$ with $N$ rows and $c$ columns (each column being $d$ dimensional) where we assume that $D$ was generated using the generative model described in Section 3. Assume that we know the set of permutations $\mathcal{P}$ and the functional form (but not the parameters) of $p(\mathbf{x})$. If we knew the permutations that generated each data vector $\mathbf{x}^{i}$, then presumably the problem of estimating the parameters of $p(\mathbf{x})$ using the "unscrambled" data would be straightforward. This suggests the use of the EM framework for this problem treating the $m$ possible permutations as "hidden" information. Letting $\Theta$ be the unknown parameters of $p(\mathbf{x})$, the log-likelihood can be defined as

$$
\begin{aligned}
l(\Theta) & =\ln q(D \mid \Theta)=\sum_{i=1}^{N} \ln q\left(\mathbf{x}^{i} \mid \Theta\right) \\
& =\sum_{i=1}^{N} \ln \sum_{j=1}^{m} p\left(\rho_{j}\right) p\left(\rho_{j}^{-1}\left(\mathbf{x}^{i}\right)\right)
\end{aligned}
$$

After $\Theta$ has been initialized in some fashion, the parameters are changed iteratively, guaranteeing a non-
decreasing log-likelihood at the end of each iteration. In the E-step, the probability of each permutation is estimated for each data vector given the current $\Theta$. In the M-step, new values for $\Theta$ are chosen to maximize the expected log-likelihood of the data with respect to the distribution over permutations as estimated in the E-step. As an example, the $p\left(\rho_{j}\right)$ terms can always be updated analytically as follows:

$$
\hat{p}\left(\rho_{j}\right)=\frac{1}{N} \sum_{i=1}^{N} q\left(\rho_{j} \mid \mathbf{x}^{i}, \Theta\right)
$$

where here (unlike the analysis earlier in the paper) the probabilities of different permutations need not be equal and can be learned from the data.

### 5.2. Learning from Gaussian Data

We simulated data with $\mathrm{S}=\mathbb{R}, c=2$, using the setup in Section 4.1.1. In the first experiment, we performed an empirical analysis of how the permutation error rate of models learned with EM depends on the number of training examples $N$. For this we set $\mu_{1}=1, \mu_{2}=-1, \sigma^{2}=16$ which yields a Bayesoptimal permutation error rate of roughly 0.36 - thus, we know in advance that none of our learned models can have a lower error rate than this. 10 different training data sets were generated for each of the following sizes: $N=\{10,20,50, \ldots, 5000,10000\}$ and for each training data set the best fitting (maximum likelihood) model was chosen from 10 random restarts of EM. Each of these best-fit models were then evaluated on a large independent test data set $\left(N=2 \times 10^{6}\right.$ data points). The plot in Figure 4 shows that, as expected, the error rate of the predicted model approaches the Bayes-optimal error rate as the number of examples increases. For this particular problem, once the number of data points is on the order of 1000 or greater, EM is performing optimally in terms of identifying the permutations.

### 5.3. Learning with Rotated Triangles

For illustrative purposes we simulated 200 triangles from a distribution over angles (corresponding to $p(\mathbf{x}))$, and then rotated and reflected the triangles in a manner corresponding to a set of random permutations. The learning problem is to learn back the distribution which generated the triangles and put the triangles in geometric correspondence. For this problem, $c=3$ ( 3 angles), $\mathrm{S}=\mathbb{R}$, and $\mathcal{P}=\left\{\rho_{1}, \ldots, \rho_{6}\right\}$ (all six possible permutations of $(1,2,3)$ ). We set $p(\rho)$ as uniform. $p(\mathbf{x})$ is defined as $p\left(x_{1}, x_{2}, x_{3}\right) \propto$ $\mathcal{N}\left(x_{1} \mid \mu_{1}, \sigma_{1}^{2}\right) \times \mathcal{N}\left(x_{2} \mid \mu_{2}, \sigma_{2}^{2}\right)$ if $x_{1}+x_{2}+x_{3}=\pi$, and 0 otherwise. For $\mu_{1}, \mu_{2}>0$ such that $\mu_{1}+\mu_{2}<\pi$,


Figure 4. Error rates of models learned from permuted data as a function of the number of training examples.
and with small $\sigma_{1}^{2}$ and $\sigma_{2}^{2}$, this distribution generates triangles with angles $x_{1}, x_{2}, x_{3}$. Figure 5 demonstrates how EM learns both the underlying density model for angle generation and a distribution over rotations and reflections for each triangle. The rows represent different iterations of EM and the leftmost column is the learned density model as represented by the "mean triangle" at each iteration. The columns represent the 6 possible permutations for one of the simulated triangles in the training data, and the numbers in each row are the probability distribution $p\left(\rho_{j} \mid \mathbf{x}\right), 1 \leq j \leq 6$ for a specific iteration of EM. Starting from a random triangle model (upper left corner) and considerable uncertainty about the likely permutation (row 1), EM gradually learns both the correct "mean shape" and identifies the most likely orientation for this particular triangle (row 5).

## 6. Conclusions

We analyzed the problem of unsupervised learning in the presence of unknown permutations of feature values using the notion of a Bayes-optimal permutation error rate $E_{P}^{\star}$. We derived a general bound on $E_{P}^{\star}$ as well as closed-form expressions for specific learning problems and found (for example) that negative and positive correlation among the feature variables can lead to very different learning problems. The paper concluded with a brief empirical illustration of how EM can be used to perform unsupervised learning from permuted data. A related applications-oriented paper (Kirshner et al., 2003a) shows how this framework can be usefully applied to orienting images of galaxies in a real-world classification application in astronomy. There are several possible extensions of this work


Figure 5. Illustration of EM learning with triangle data.
including further analysis of the relationship between $E_{C}^{*}$ and $E_{P}^{*}$, computational techniques to handle large numbers of permutations ( $m=c$ ! for large $c$ ), and analysis of learning algorithms for more general transformations than permutations.

## References

Duda, R. O., Hart, P. E., \& Stork, D. G. (2000). Pattern Classification. New York: John Wiley \& Sons. Second edition.

Frey, B. J., \& Jojic, N. (2003). Transformationinvariant clustering using the EM algorithm. IEEE Transactions on Pattern Analysis and Machine Intelligence, 25, 1-17.

Gold, S., Lu, C. P., Rangarajan, A., Pappu, S., \& Mjolsness, E. (1995). New algorithms for 2D and 3D point matching: Pose estimation and correspondence. Advances in Neural Information Processing Systems (pp. 957-964). The MIT Press.

Kirshner, S., Cadez, I. V., Smyth, P., \& Kamath, C. (2003a). Learning to classify galaxy shapes using the EM algorithm. Advances in Neural Information Processing Systems 15. MIT Press.

Kirshner, S., Parise, S., \& Smyth, P. (2003b). Unsupervised learning with permuted data (Technical Report ICS TR-03-03). University of California Irvine. http://www.datalab.uci.edu.

McCallum, A., Nigam, K., Rennie, J., \& Seymore, K. (2000). Automating the construction of internet portals with machine learning. Information Retrieval, 3, 127-163.

