Simultaneous Unsupervised Learning of Disparate Clusterings

Prateek Jain, Raghu Meka and Inderjit S. Dhillon

Department of Computer Sciences, University of Texas

Austin, TX 78712-1188, USA

{pjain, raghu, inderjit}@cs.utexas.edu

Abstract

Most clustering algorithms produce a single clustering for a given data set even when the data can be clustered naturally in multiple ways. In this paper, we address the difficult problem of uncovering disparate clusterings from the data in a *totally unsupervised manner*. We propose two new approaches for this problem. In the first approach we aim to find good clusterings of the data that are also *decorrelated* with one another. To this end, we give a new and tractable characterization of decorrelation between clusterings, and present an objective function to capture it. We provide an iterative “decorrelated” k -means type algorithm to minimize this objective function. In the second approach, we model the data as a sum of mixtures and associate each mixture with a clustering. This approach leads us to the problem of learning a convolution of mixture distributions. Though the latter problem can be formulated as one of factorial learning [8, 13, 16], the existing formulations and methods do not perform well on many real high-dimensional data sets. We propose a new regularized factorial learning framework that is more suitable for capturing the notion of disparate clusterings in modern, high-dimensional data sets. The resulting algorithm does well in uncovering multiple clusterings, and is much improved over existing methods. We evaluate our methods on two real-world data sets - a music data set from the text mining domain, and a portrait data set from the computer vision domain. Our methods achieve a substantially higher accuracy than existing factorial learning as well as traditional clustering algorithms.

1 Introduction

Clustering data into groups based on similarity is often one of the most important steps in any data analysis application. Currently, most clustering algorithms partition the data into groups that are disjoint, while other algorithms extend this approach to probabilistic or overlapping clustering. However, in many important applications it is necessary to uncover disparate or alternative clusterings¹ in order to reflect

¹Throughout this paper, a *clustering* will refer to a set of disjoint clusters of the data.

Figure 1: Images of different persons in different poses. Each row has different persons in the same pose. Each column has the same person in different poses.

the different groupings inherent in the data. As an example, consider a set of pictures of different persons in different poses (see Figure 1). These images can be clustered by the identity of the person in the picture or by the pose of the person. Given such a dataset it would be desirable to recover two disparate clusterings of the data - one based on the identity of the person and the other based on their pose.

The above problem arises naturally for many other widely used datasets, for instance: news articles (can be clustered by the main topic, or by the news source), reviews of various musical albums (can be clustered by composers, or by other characteristics like genre of the album), and movies (can be clustered based on actors/actresses or genre).

Most existing methods to recover alternative clusterings use semi-supervision or side-information about one or more of the clusterings. Since clustering is generally the first step in data analysis, such information might not be available beforehand. For example, news articles change dynamically and it is infeasible to manually label them by topic and the source. Thus, completely unsupervised techniques to find disparate clusterings are immensely useful.

In this paper, we present two novel unsupervised ap-

proaches for discovering disparate clusterings in a given dataset. In the first approach we aim to find multiple clusterings of the data which satisfy two criteria: a) the clustering error of each individual clustering is small and b) different clusterings have small *correlation* between them. To this end, we present a new and computationally tractable characterization of *correlation* (or decorrelation) between different clusterings. We use this characterization to formulate a k -means type objective function which contains error terms for each individual clustering along with a regularization term corresponding to the correlation between clusterings. We provide a computationally efficient k -means type algorithm for minimizing this objective function.

In the second approach we model the problem of finding disparate clusterings as one of learning the component distributions when the given data is sampled from a convolution of mixture distributions. This formulation is appropriate when the different clusterings come from independent additive components of the data. The problem of learning a convolution of mixture distributions is closely related to factorial learning [8, 13, 16]. However, the methods of [8, 13, 16] are not suited for recovering multiple clusterings. The problem with applying factorial learning directly is that there are multiple solutions to the problem of learning a convolution of mixture distributions. Out of all such possible solutions, the desirable solutions are the ones that give maximally disparate clusterings. To address this problem we propose a regularized factorial learning model that intuitively captures the notion of decorrelation between clusterings and aims to estimate the parameters of the decorrelated model.

An important aspect of both our approaches is the notion of decorrelation between clusterings. The decorrelation measures that we propose quantify the “orthogonality” between the mean vectors corresponding to different clusterings. We show that the characterization of disparity between different clusterings by the “orthogonality” between the mean vectors of the respective cluster centers has a well-founded theoretical basis (see Section 3.3.1).

We evaluate our methods on synthetic and real-world datasets that have multiple disparate clusterings. We consider real-world datasets from two different domains - a music dataset from the text-mining domain and a portrait dataset from the computer vision domain. We compare our methods to two factorial learning algorithms, Co-operative Vector Quantization (CVQ)[8] and Multiple Cause Vector Quantization (MCVQ)[16]. We also compare against traditional single clustering algorithms like k -means and non-negative matrix approximation (NNMA)[15]. On all the datasets, both of our algorithms significantly outperform the factorial learning as well as the single clustering algorithms. The factorial learning methods work reasonably well on a few synthetic datasets which exactly satisfy their respective model assumptions. But they are not robust in the case where model

assumptions are even slightly violated. Because of this, their performance is poor on real-world datasets and other synthetic datasets. In comparison, our algorithms are more robust and perform significantly better on all the datasets. For the music dataset both our algorithms achieve around 20% improvement in accuracy over the factorial learning and single clustering algorithms (k -means and NNMA). Similarly, for the portrait dataset we achieve an improvement of 30% over the baseline algorithms.

2 Related Work

Most of the existing work for finding disparate clusterings has been in the semi-supervised setting. The semi-supervised clustering problem of finding a clustering consistent with a given set of constraints has been extensively studied ([21, 23, 2]). This approach has been applied to the problem of recovering multiple clusterings by providing appropriate constraints. Must-link and cannot-link constraints have been extensively used for semi-supervised clustering ([21, 22, 2]). Recently, Davidson et al.[4] proposed an efficient incremental algorithm for must-link and cannot-link constraints. An alternative approach to the problem is taken by [1, 10, 11] where it is assumed that a clustering of the data is given and the objective is to find a clustering different from the given one. Our work differs from the above approaches in that our methods for discovering the disparate clusterings are completely unsupervised.

A supervised approach to the related problem of learning hidden two-factor structures from the observed data was suggested in [20]. Their method, named Separable Mixture Model (SMM), models the data using a bilinear function of the factors and can also be used for obtaining two clusterings of the data. An advantage of our methods over SMM is that our methods are unsupervised compared to the supervised approach of SMM. Also, our model can be extended to more than two factors, whereas it is unclear how SMM could be extended to a data generated from more than two-factors.

Our second approach (“sum of parts” approach) is closely related to the factorial learning problem where each data point is assumed to be generated by combining multiple factors. Ghahramani[8] introduced a novel architecture named co-operative vector quantization (CVQ), in which a set of multiple vector quantizers (VQ) combine linearly to generate the input data. However, a drawback of CVQ is that it can have multiple solutions. Many of these solutions give poor results for the problem of discovering disparate clusterings, especially on our real-world applications. Also, CVQ can be seen as a special case of our model. Another recent model related to factorial learning is multiple cause vector quantization (MCVQ) (Ross and Zemel[16]). In MCVQ it is assumed that the dimensions of the data can be separated into several disjoint factors, which take on values independently of each other. The factors are then modeled using a

vector quantizer as in CVQ. However, MCVQ also faces the same drawbacks of CVQ - existence of multiple solutions - which leads to a poor performance for our application of discovering disparate clusterings.

The problem of learning convolutions of distributions that forms the basis of our second approach has been considered in the statistics community - see for instance [6], [18], [17]. However, these methods deal with learning convolutions of simple distributions like binomial, Gaussian and Poisson, and do not consider mixtures of distributions. A fundamental problem with learning a convolution of Gaussians, as mentioned in [18], is that the problem is not well-defined - there exist many solutions to the learning problem. We face a similar problem in the M-step of our algorithm for learning the convolution of mixtures of Gaussians, where the maximum likelihood estimation has multiple solutions. We deal with this issue by regularizing the solution space in a way suitable for the purpose of recovering disparate clusterings so that the problem becomes well-posed.

We emphasize that though we state the problem of recovering disparate clusterings as one of learning independent components from the data, the problem we address is completely different from that of independent component analysis (ICA) [14]. ICA tries to separate a multivariate signal into independent additive univariate signals, whereas in our problem we try to decompose the signal into independent multivariate signals, each of which may have high correlation between its different dimensions.

For our experiments, we also evaluated various simple extensions of k -means such as the removal of important features of the first clustering to uncover the second clustering, and projection of the data onto the space orthogonal to the means of the first clustering. The later heuristic was motivated by principal gene shaving [12]. But, these approaches are ad-hoc and do not perform well in our experiments.

3 Methodology

For simplicity, we present our methods for uncovering two disparate clusterings from the data; our techniques can be generalized to uncover more than two clusterings. We propose the following approaches:

- **Decorrelated-kmeans approach:** In this approach we try to fit each clustering to the entire data, while requiring that different clusterings be decorrelated with each other. To this end, we introduce a novel measure for correlation between clusterings. This measure is motivated by the fact that if the representative vectors of two clusterings are orthogonal to one another, then the labellings generated by nearest neighbor assignments for these representative vectors are independent under some mild conditions (see Section (3.3.1)).
- **Sum of parts approach:** In this approach we model

the data as a sum of independent components, each of which is a mixture model. We then associate each component with a clustering. Further, as the distribution of the sum of two independent random variables is the convolution of the distributions (see [5]), we model the observed data as being sampled from a convolution of two mixtures. Thus, our approach leads us to the problem of learning a convolution of mixtures. Note that the individual components uncovered by this approach may not be good approximations to the data by themselves, but their sum is. This is in complete contrast to the first approach where we try to approximate the data individually by each component.

3.1 First Approach: Decorrelated-kmeans

Given a set of data points $Z = \{z_1, z_2, \dots, z_n\} \subseteq \mathbb{R}^m$, we aim to uncover two clusterings C^1 and C^2 . Specifically, we wish to partition the set Z into k_1 groups for the first clustering C^1 and k_2 groups for the second clustering C^2 . To achieve this task, we try to find *decorrelated* clusterings each of which approximates the data as a whole. We propose the following objective function:

$$(3.1) \quad G(\mu_{1\dots k_1}, \nu_{1\dots k_2}) = \sum_i \sum_{z \in C_i^1} \|z - \mu_i\|^2 + \sum_j \sum_{z \in C_j^2} \|z - \nu_j\|^2 + \lambda \sum_{i,j} (\beta_j^T \mu_i)^2 + \lambda \sum_{i,j} (\alpha_i^T \nu_j)^2,$$

where C_i^1 is cluster i of the first clustering, C_j^2 is cluster j of the second clustering, and $\lambda > 0$ is a regularization parameter. The vector μ_i is the *representative* vector of C_i^1 , ν_j is the *representative* vector of C_j^2 , α_i is the mean vector of C_i^1 and β_j is the mean vector of C_j^2 .

The first two terms of (3.1) correspond to a k -means type error term for the clusterings, with a crucial difference being that the “representative” vector of a cluster may not be its mean vector. The last two terms are regularization terms that measure the decorrelation between the two clusterings. In order to extend this formulation for $T \geq 2$ clusterings, we add k -means type error terms for each of the T clusterings. Furthermore, we add $T \times (T - 1)/2$ terms corresponding to the decorrelation between pairs of clusterings.

The decorrelation measure given above is motivated by the intuition that if the “representative” vectors of two clusterings are orthogonal to one another, then the labellings generated by nearest neighbor assignments for these vectors are independent. We provide a theoretical basis for the above intuition in Section 3.3.1. Also, an important advantage of the proposed decorrelation measure is that the objective function remains strictly and jointly convex in the μ_i ’s and ν_j ’s (assuming fixed C_i^1 ’s and C_j^2 ’s).

To minimize the objective function (3.1), we present

an iterative algorithm which we call Decorrelated-kmeans (Algorithm 1). We fix C^1 and C^2 to obtain μ_i 's and ν_j 's that minimize (3.1) and then assign each point z to C_i^1 such that $i = \arg\min_l \|z - \mu_l\|^2$ and to C_j^2 such that $j = \arg\min_l \|z - \nu_l\|^2$. We initialize one of the clusterings using k -means with $k = k_1$ and the other clustering randomly.

For computing the μ_i 's and ν_j 's, we need to minimize (3.1). The gradient of the objective function in (3.1) w.r.t μ_i is given by:

$$\frac{\partial G}{\partial \mu_i} = -2 \left(\sum_{z \in C_i^1} z \right) + 2 \left(\sum_j n_{ij} \right) \mu_i + 2\lambda \sum_j (\beta_j^T \mu_i) \beta_j,$$

where n_{ij} is the number of points that belong to C_i^1 and C_j^2 .

Now, $(\beta_j^T \mu_i) \beta_j = (\beta_j \beta_j^T) \mu_i$ and $\alpha_i = \frac{(\sum_{z \in C_i^1} z)}{\sum_j n_{ij}}$. Thus,

$$\frac{\partial G}{\partial \mu_i} = -2 \sum_j n_{ij} \alpha_i + 2 \sum_j n_{ij} \mu_i + 2\lambda \left(\sum_j \beta_j \beta_j^T \right) \mu_i.$$

Similarly,

$$\frac{\partial G}{\partial \nu_j} = -2 \sum_i n_{ij} \beta_j + 2 \sum_i n_{ij} \nu_j + 2\lambda \left(\sum_i \alpha_i \alpha_i^T \right) \nu_j.$$

Setting the gradients to zero gives us the following equations:

$$(3.2) \quad \mu_i = \left(I + \frac{\lambda}{\sum_j n_{ij}} \sum_j \beta_j \beta_j^T \right)^{-1} \alpha_i,$$

$$(3.3) \quad \nu_j = \left(I + \frac{\lambda}{\sum_i n_{ij}} \sum_i \alpha_i \alpha_i^T \right)^{-1} \beta_j.$$

Since the objective function (3.1) is strictly and jointly convex in both μ_i 's and ν_j 's, the above updates lead to a global minima of the objective function (3.1) for fixed C^1 and C^2 .

3.1.1 Computing the updates efficiently: Computing the updates given by (3.2) and (3.3) requires computing the inverse of an $m \times m$ matrix, where m is the dimensionality of the data. Thus updating all the μ_i 's and ν_j 's directly would seem to require $O(k_1 m^3 + k_2 m^3)$ operations, which is cubic in the dimensionality of the data. We now give a substantially faster way to compute the updates in time linear in the dimensionality. Using the Sherman-Morrison-Woodbury formula (see [9]) for the inverse in (3.2), we get

$$(I + \xi_i V V^T)^{-1} = I - \xi_i V (I + \xi_i V^T V)^{-1} V^T,$$

where $\xi_i = \frac{\lambda}{\sum_j n_{ij}}$ and $V = [\beta_1, \dots, \beta_{k_2}]$. Using the eigenvalue decomposition $V^T V = Q \Sigma Q^T$ we see that

$$(I + \xi_i V^T V)^{-1} = Q (I + \xi_i \Sigma)^{-1} Q^T.$$

Algorithm 1 Decorrelated-kmeans (Dec-kmeans)

Input: Data $\mathcal{Z} = \{z_1, z_2, \dots, z_n\}$

k_1 : Number of clusters in first clustering (C^1)

k_2 : Number of clusters in second clustering (C^2)

λ : regularization parameter

Output: C^1, C^2 : Two different clusterings

1. $C^1 \leftarrow k\text{-means}(\mathcal{Z}), C^2 \leftarrow \text{Random assignment}$

2. **repeat**

2.1. $\alpha_i \leftarrow \text{ComputeMean}(C_i^1)$, for all $1 \leq i \leq k_1$

2.2. $\beta_j \leftarrow \text{ComputeMean}(C_j^2)$, for all $1 \leq j \leq k_2$

2.3. Update μ_i and ν_j for all i, j using (3.4), (3.5)

2.4. $\forall z, C_i^1 \leftarrow C_i^1 \cup \{z\}$,
if $i = \arg\min_l \|z - \mu_l\|^2$.

2.5. $\forall z, C_j^2 \leftarrow C_j^2 \cup \{z\}$,
if $j = \arg\min_l \|z - \nu_l\|^2$.

4. **until** convergence

return C^1, C^2

Since $V^T V$ is a $k_2 \times k_2$ matrix its eigenvalue decomposition can be computed in $O(k_2^3)$ time. Also, as $(I + \xi_i \Sigma)^{-1}$ is a diagonal matrix, calculating its inverse requires just $O(k_2)$ operations. The updates for μ_i 's can now be rewritten as,

$$(3.4) \quad \mu_i = (I - \xi_i V Q (I + \xi_i \Sigma)^{-1} Q^T V^T) \alpha_i.$$

Similarly, the updates for ν_j 's can now be written as,

$$(3.5) \quad \nu_j = (I - \zeta_j M U (I + \zeta_j \Lambda)^{-1} U^T M^T) \beta_j,$$

where $\zeta_j = \frac{\lambda}{\sum_i n_{ij}}$, $M = [\alpha_1, \dots, \alpha_{k_1}]$, and $U \Lambda U^T$ is the eigenvalue decomposition of $M^T M$.

Using these updates reduces the computational complexity of computing all the μ_i 's and ν_j 's to $O(m k_1^2 + m k_2^2 + k_1^3 + k_2^3)$. If $m > k = \max(k_1, k_2)$, which is typically the case, the above bound becomes $O(m k^2)$.

3.1.2 Determining λ : The regularization parameter λ plays an important role in the Decorrelated-kmeans algorithm. It determines the tradeoff between minimizing the individual clustering error of each clustering (first two terms in (3.1)) and finding decorrelated cluster centers for the different clusterings (last two terms in (3.1)). Empirically, we observe that the clustering accuracies are good when $\lambda \in [100, 10000]$, which is a large range. But, a different scaling of the data can change this range for λ . Hence, we determine λ using a simple heuristic. Note that for small values of λ , the Decorrelated-kmeans algorithm finds approximately the same clusters for both the clusterings. While for

high value of λ it tries to find clusterings which are orthogonal to each other, even though both the clusterings may not fit the data well. Thus, a suitable λ balances out both the objectives and hence generally there is a sudden change in objective function value. Based on this intuition we form a heuristic to determine λ : start with a large λ and find different clusterings of the data while decreasing λ , and select a λ for which the drop in the objective function is the highest. Note that different variants of the heuristic can be used depending on the data and domain knowledge. E.g. if the data is large, then a subset of the data can be used for finding clusterings or if the data is noisy then a more robust measure like average change in objective function should be preferred over the maximum change measure for selecting λ .

3.2 Second Approach: Sum of Parts

In this section, we describe our “sum of parts” approach. Let $\mathcal{Z} = \{z_1, \dots, z_n\}$ be the observed m -dimensional data sampled from a random variable Z . We model Z as a sum $X + Y$, where X, Y are independent random variables and are drawn from mixtures of distributions. Specifically,

$$p_X = \sum_{i=1}^{k_1} a_i p_{X_i}, \quad p_Y = \sum_{j=1}^{k_2} b_j p_{Y_j}.$$

The problem of learning independent components can now be stated as: Given data sampled according to Z , recover the parameters of the probability distributions p_{X_i}, p_{Y_j} along with the mixing weights a_i, b_j .

As $Z = X + Y$, the probability density function of Z is the convolution of p_X and p_Y [5, Section A.4.11]. Thus,

$$p_Z(z) = (p_X * p_Y)(z) = \sum_{i=1}^{k_1} \sum_{j=1}^{k_2} (a_i b_j) \cdot (p_{X_i} * p_{Y_j})(z), \quad (3.6)$$

where $f_1 * f_2(z) = \int_{\mathbb{R}^m} f_1(x) \cdot f_2(z - x) dx$ denotes the convolution of f_1 and f_2 .

From (3.6) it follows that when the distributions p_{X_i} and p_{Y_j} belong to a family of distributions closed under convolution, Z can be viewed as a mixture of $k_1 \times k_2$ distributions. However, the problem of learning the components X and Y from Z is harder than that of simply learning the parameters of a mixture model, as along with learning the $k_1 \times k_2$ component distributions one must also be able to factor them out. In the following section, we give a generalized Expectation Maximization (EM) algorithm for learning the parameters of the component mixtures when the base distributions are spherical multi-variate Gaussians. Our techniques can be extended to more general distributions like non-spherical Gaussians and potentially to other families closed under convolution.

3.2.1 Learning the convolution of a mixture of Gaussians

Let the components X and Y be mixtures of spherical Gaussians, i.e., $p_X = \sum_{i=1}^{k_1} a_i \mathcal{N}(\mu_i, \sigma^2)$ and $p_Y = \sum_{j=1}^{k_2} b_j \mathcal{N}(\nu_j, \sigma^2)$. As in our first approach we initialize the EM algorithm (Algorithm 2) by k -means for the first clustering and a random assignment for the second clustering. We initialize μ_i^0 's and ν_j^0 's to be the means of the first and second clusterings respectively. To initialize σ we use a heuristic presented in [3],

$$\sigma = \frac{1}{\sqrt{2m}} \min \left(\min_{i \neq j} \|\mu_i^0 - \mu_j^0\|, \min_{i \neq j} \|\nu_i^0 - \nu_j^0\| \right).$$

E-step:

Let $p_{ij}^t(z)$ denote the conditional probability that z comes from the Gaussian $p_{X_i} * p_{Y_j}$ given the current parameters. As our main objective is to cluster the data, we use hard assignments in the E-step to ease the computations involved. The E-step in this case will be:

$$(3.7) \quad p_{ij}^{t+1}(z) = \begin{cases} 1, & \text{if } (i, j) = \\ & \text{argmax}_{(r,s)} \{a_r^t b_s^t \cdot \mathcal{N}(\mu_r^t + \nu_s^t, 2(\sigma^t)^2)(z)\} \\ 0, & \text{otherwise.} \end{cases}$$

Note that, to uncover T different clusterings from the data, $O(k^T)$ computational operations are required for each data point in the E-step. Ghahramani[8] suggested various approximation methods to reduce the time complexity of this estimation, and the same can be applied to our setting as well. In our implementation, we use Gibbs sampling for approximating the distribution of labels, $p_{ij}^{t+1}(z)$, when the parameters of the base distributions are fixed.

M-step:

In the M-step, we use the clusterings updated in the E-step (specified by p_{ij}^{t+1} 's) to estimate the parameters of the distributions. Formally, we maximize the log-likelihood:

$$\begin{aligned} & (\mu_{1 \dots k_1}^{t+1}, \nu_{1 \dots k_2}^{t+1}, \sigma^{t+1}, a_{1 \dots k_1}^{t+1}, b_{1 \dots k_2}^{t+1}) = \\ & \text{argmax}_{\substack{\mu_{1 \dots k_1}, \nu_{1 \dots k_2}, \sigma, \\ a_{1 \dots k_1}, b_{1 \dots k_2}}} \sum_{i,j,z} p_{ij}^{t+1}(z) \log(a_i b_j \mathcal{N}(\mu_i + \nu_j, \sigma)(z)). \end{aligned}$$

The mixture weights and variance σ can be easily computed by differentiating w.r.t. a_i 's, b_j 's, σ and setting the derivatives to zero. This gives us the following expressions:

$$(3.8) \quad a_i^{t+1} = \frac{1}{n} \sum_j \sum_z p_{ij}^{t+1}(z),$$

$$(3.9) \quad b_j^{t+1} = \frac{1}{n} \sum_i \sum_z p_{ij}^{t+1}(z),$$

$$(3.10) \quad (\sigma^{t+1})^2 = \frac{1}{2mn} \sum_{i,j,z} p_{ij}^{t+1}(z) \|z - \mu_i^t - \nu_j^t\|^2.$$

Computing the means to maximize the log-likelihood is more involved and it reduces to minimizing the following objective function:

$$(3.11) \quad \min_{\mu_{1\dots k_1}, \nu_{1\dots k_2}} F(\mu_{1\dots k_1}, \nu_{1\dots k_2}) = \sum_{i,j,z} p_{ij}^{t+1}(z) \|z - \mu_i - \nu_j\|^2.$$

Note that there exist multiple solutions for the above equation; since we can translate the means μ_i 's by a fixed vector w and the means ν_j 's by $-w$ to get another set of solutions. As discussed in Section 2, the CVQ [8] algorithm also suffers from the same problem of multiple solutions. Out of all the solutions to (3.11), the solutions which give maximally disparate clusterings are more desirable. To obtain such solutions we regularize the μ_i 's and ν_j 's to have small correlation with each other. To this end we introduce a regularization term to make the μ_i 's and ν_j 's orthogonal to one another. This correlation measure is similar to the measure discussed in the previous Decorrelated-kmeans approach (Section 3.1). Formally, we minimize the following objective function:

$$(3.12) \quad \tilde{F}(\mu_{1\dots k_1}, \nu_{1\dots k_2}) = \sum_{i,j,z} p_{ij}^{t+1}(z) \|z - \mu_i - \nu_j\|^2 + \lambda \sum_{i,j} (\mu_i^T \nu_j)^2,$$

where $\lambda > 0$ is a regularization parameter and can be selected using a heuristic similar to the one described in Section 3.1.2.

Observe that the above objective is not jointly convex in μ_i and ν_j but is strictly convex in μ_i for fixed ν_j 's and vice-versa. To minimize $\tilde{F}$, we use the block coordinate descent algorithm ([24]) where we fix ν_j 's to minimize μ_i and vice-versa. By differentiating (3.12) w.r.t. μ_i and ν_j and setting the derivatives to zero we get,

$$\begin{aligned} \left(I + \frac{\lambda \sum_j \nu_j \nu_j^T}{\sum_j n_{ij}} \right) \mu_i + \frac{\sum_j n_{ij} \nu_j}{\sum_j n_{ij}} &= \alpha_i, \\ \left(I + \frac{\lambda \sum_i \mu_i \mu_i^T}{\sum_i n_{ij}} \right) \nu_j + \frac{\sum_i n_{ij} \mu_i}{\sum_i n_{ij}} &= \beta_j, \end{aligned}$$

where $n_{ij} = \sum_z p_{ij}^{t+1}(z)$ is the number of data-points that belong to cluster i of the first clustering and cluster j of the second clustering, α_i denotes the mean of all points that are assigned to cluster i in the first clustering and β_j denotes the mean of points assigned to cluster j in the second clustering, i.e. ,

$$(3.13) \quad \alpha_i = \frac{1}{na_i} \sum_j \sum_z z p_{ij}^{t+1}(z),$$

$$(3.14) \quad \beta_j = \frac{1}{nb_j} \sum_i \sum_z z p_{ij}^{t+1}(z).$$

To solve for μ_i and ν_j in the above equations we use an alternative minimization scheme - we iteratively update the μ_i and ν_j as follows:

$$(3.15) \quad \mu_i = \left(I + \frac{\lambda \sum_j \nu_j \nu_j^T}{\sum_j n_{ij}} \right)^{-1} \left(\alpha_i - \frac{\sum_j n_{ij} \nu_j}{\sum_j n_{ij}} \right)$$

$$(3.16) \quad \nu_j = \left(I + \frac{\lambda \sum_i \mu_i \mu_i^T}{\sum_i n_{ij}} \right)^{-1} \left(\beta_j - \frac{\sum_i n_{ij} \mu_i}{\sum_i n_{ij}} \right).$$

For initialization we set ν_j to be β_j for each j . Below we prove that this scheme converges to a local minima of (3.12).

LEMMA 3.1. *The updates for μ_i and ν_j given by (3.15) converge to a local minimum of the regularized objective function given by (3.12).*

Proof. As the updates (3.15) minimize the objective function at each iteration, the updates converge to a fixed point[24]. Also, the objective function (3.12) is strictly-convex in μ_i for fixed ν_j 's and vice-versa. Thus, any fixed point of (3.12) is also a local minimum. It now follows that our updates converge to a local minimum of the objective function. $\square$

THEOREM 3.1. *Algorithm 2 monotonically decreases the objective function:*

$$(3.17) \quad F = \sum_{i,j,z} p_{ij}^{t+1}(z) \frac{\|z - \mu_i - \nu_j\|^2}{2\sigma^2} + \lambda \sum_{i,j} (\mu_i^T \nu_j)^2$$

Proof. Let F_t be the objective function value at the start of t -th iteration, F_t^E be the objective function value after the E-step of t -th iteration and $F_t^M = F_{t+1}$ be the objective function after M -step of t -th iteration. The E-step assigns new labels according to 3.7, which is equivalent to minimizing:

$$\sum_{i,j,z} p_{ij}^{t+1}(z) \frac{\|z - \mu_i - \nu_j\|^2}{2\sigma^2},$$

with μ_i and ν_j being fixed.

Thus, the first term of the objective function (3.17) is decreased by the E-step while the second term remains fixed. Hence, $F_t \geq F_t^E$. Using Lemma 3.1, $F_t^E \geq F_t^M$, as only μ_i 's and ν_j 's are variables with p_{ij} fixed (σ can be absorbed in λ). Thus, $F_t \geq F_{t+1}$. $\square$

3.2.2 Computing the updates efficiently: Using techniques similar to Section (3.1.1), the update for μ_i can be written as:

$$(3.18) \quad \mu_i = \left(I - \xi_i V Q (I + \xi_i \Sigma)^{-1} Q^T V^T \right) \left(\alpha_i - \frac{\sum_j n_{ij} \nu_j}{\sum_j n_{ij}} \right),$$

Algorithm 2 Convolutional-EM (Conv-EM)

Input: Data $\mathcal{Z} = \{z_1, z_2, \dots, z_n\}$ k_1 : Number of clusters in first clustering (C^1) k_2 : Number of clusters in second clustering (C^2) λ : regularization parameter**Output:** C^1, C^2 : Two different clusterings1. $C^1 \leftarrow k\text{-means}(\mathcal{Z}), C^2 \leftarrow \text{Random assignment}$ 2. $\mu_i \leftarrow \text{ComputeMean}(C_i^1), \nu_j \leftarrow \text{ComputeMean}(C_j^2)$ 3. $a_i = \frac{1}{k_1}, b_j = \frac{1}{k_2}$ 4. **repeat****E Step:**4.1. For each z , assign $p_{ij}(z)$ using (3.7).**M Step:**4.2. Assign a_i, b_j and σ using (3.8), (3.9), (3.10).4.3. Assign α_i and β_j using (3.13), (3.14).4.4. $\nu_j \leftarrow \beta_j$ 4.5. **repeat until convergence**• Update μ_i using (3.18).• Update ν_j using (3.19).5. **until** convergence6. $C_i^1 = \{z | p_{ij}(z) = 1, \forall j\}, C_j^2 = \{z | p_{ij}(z) = 1, \forall j\}$ **return** C^1, C^2

where, $\xi_i = \frac{\lambda}{\sum_j n_{ij}}$ and $V = [\nu_1, \dots, \nu_{k_2}]$ and $Q\Sigma Q^T$ is

the eigenvalue decomposition of $V^T V$.

Similarly, the update for ν_j can be written as, (3.19)

$$\nu_j = \left(I - \zeta_j M U (I + \zeta_j \Lambda)^{-1} U^T M^T \right) (\beta_j - \frac{\sum_i n_{ij} \mu_i}{\sum_i n_{ij}}),$$

where, $\zeta_j = \frac{\lambda}{\sum_i n_{ij}}$, $M = [\mu_1, \dots, \mu_{k_1}]$ and $M^T M = U \Lambda U^T$.

As in Section (3.1.1), the above updates reduce the computational complexity of computing all the μ_i 's and ν_j 's from $O(k_1 m^3 + k_2 m^3)$ to $O(m(k_1^2 + k_2^2))$.

3.3 Discussion

3.3.1 Decorrelation measure: Now we motivate the decorrelation measures used in equations (3.1) and (3.12). For this, we will need the following lemma about multivariate Gaussians. In the lemma, let $Z \in \mathbb{R}^m$ denote a random variable with spherical Gaussian distribution. For a subspace S of $\mathbb{R}^m$ let $P_S : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be the orthogonal projection operator onto the subspace S .

LEMMA 3.2. *Let $S_1, S_2 \subseteq \mathbb{R}^m$ be two subspaces such that $S_1 \cap S_2 = \{0\}$ and let $Z_1 = P_{S_1}(Z)$, $Z_2 = P_{S_2}(Z)$ be the random variables obtained by projecting Z onto S_1, S_2 respectively. Then, the random variables Z_1 and Z_2 are*

independent if and only if the subspaces S_1 and S_2 are orthogonal.

Proof. $\Leftarrow$ If S_1 and S_2 are orthogonal, then for $u_1 \in S_1$ and $u_2 \in S_2$, $Pr[Z = u_1 + u_2] = Pr[Z_1 = u_1 \wedge Z_2 = u_2]$. Further, since Z has a spherical Gaussian distribution so do Z_1 and Z_2 . The independence of Z_1 and Z_2 follows easily from the above observations.

$\Rightarrow$ Let the random variables Z_1 and Z_2 be independent. Note that without loss of generality we can assume that Z has mean 0 (as else we can translate Z). Furthermore, we can also assume that the support of Z is contained in $S_1 + S_2$. This is because, $P_{S_1} = P_{S_1} \circ P_{S_1 + S_2}$ and $P_{S_1 + S_2}(Z)$ is also distributed as a spherical multivariate Gaussian. For the rest of the proof we will suppose that $S_1 + S_2 = \text{support}(Z) = \mathbb{R}^m$ and that Z has mean 0 .

Using Lemma (3.3) (given below), for $u \in \mathbb{R}^m$, we have

$$Pr[Z = u] = Pr[P_{S_1}(Z) = P_{S_1}(u) \wedge P_{S_2}(Z) = P_{S_2}(u)].$$

As Z_1 and Z_2 are independent the above can be rewritten as

$$Pr[Z = u] = Pr[P_{S_1}(Z) = P_{S_1}(u)] \cdot Pr[P_{S_2}(Z) = P_{S_2}(u)].$$

Now, since the projection of a spherical multivariate Gaussian is also a spherical multivariate Gaussian, substituting probability density formulae in the above equation we get the following:

$$\frac{1}{(2\pi)^{m/2}} e^{-\frac{1}{2}\|u\|^2} = \frac{1}{(2\pi)^{m_1/2}} e^{-\frac{1}{2}\|u_1\|^2} \cdot \frac{1}{(2\pi)^{m_2/2}} e^{-\frac{1}{2}\|u_2\|^2},$$

where, m_1, m_2 denote the dimensions of S_1, S_2 respectively and $u_1 = P_{S_1}(u), u_2 = P_{S_2}(u)$. Noting that $m = m_1 + m_2$ (since $S_1 \cap S_2 = \{0\}$) the above equation can be simplified to

$$\|u\|^2 = \|P_{S_1}(u)\|^2 + \|P_{S_2}(u)\|^2.$$

As the above equation holds for all u it also holds in particular for $u \in S_1$. Now, for $u \in S_1$ we have $P_{S_1}(u) = u$, thus we get

$$\forall u \in S_1, P_{S_2}(u) = 0.$$

That is, all the vectors in S_1 are projected onto zero by P_{S_2} . Using elementary linear algebra the above condition can easily be shown to be equivalent to S_1 and S_2 being orthogonal. $\square$

To complete the proof above we need the following Lemma.

LEMMA 3.3. *Let S_1, S_2 be subspaces of $\mathbb{R}^m$ such that $S_1 \cap S_2 = \{0\}$. Then, for all $x \in S_1$, and $y \in S_2$, there exists a unique $u \in S_1 + S_2$ such that $P_{S_1}(u) = x$ and $P_{S_2}(u) = y$.*

Proof. Let $x \in S_1$ and $y \in S_2$. Also, let P_1, P_2 be the projection matrices for the projection operators P_{S_1} and P_{S_2}

respectively. We first formulate the hypothesis that $S_1 \cap S_2 = \{0\}$ in terms of the matrices P_1, P_2 by showing that $I - P_1 P_2$ and $I - P_2 P_1$ are invertible. Suppose on the contrary that $I - P_1 P_2$ is not invertible. Then, for some non-zero z we must have, $(I - P_1 P_2)z = 0$, i.e., $z = P_1 P_2 z$. Recall that for a projection matrix P into a subspace S we always have $\|Pu\| \leq \|u\|$ with equality if and only if $u \in S$. Thus, we have

$$\|z\| = \|P_1 P_2 z\| \leq \|P_2 z\| \leq \|z\|.$$

Therefore, $z = P_1 P_2 z = P_2 z$, which is possible only if $z \in S_1$ and $z \in S_2$. As this contradicts the assumption that $S_1 \cap S_2 = \{0\}$, we must have $I - P_1 P_2$ to be invertible. Similarly, we can also show that $I - P_2 P_1$ is invertible.

Now, to prove the lemma we need to show that there exists a unique $u \in S_1 + S_2$ such that ' $x = P_1 u$ and $y = P_2 u$ '. Since, $S_1 \cap S_2 = \{0\}$, solving the above system of equations is equivalent to solving for $v \in S_1$, and $w \in S_2$ such that

$$x = P_1(v + w), \quad y = P_2(v + w).$$

Manipulating the above equations, we get:

$$(I - P_1 P_2)v = x - P_1 y, \quad (I - P_2 P_1)w = y - P_2 x.$$

The existence and uniqueness of v, w follow from the fact that $I - P_1 P_2$ and $I - P_2 P_1$ are invertible.

We now give the motivation for our decorrelation measures. Let $\mu_1, \dots, \mu_{k_1}$ and $\nu_1, \dots, \nu_{k_2}$ be vectors in $\mathbb{R}^m$ such that μ_i and ν_j are orthogonal for all i, j . Let S_1 be the space spanned by μ_i 's and S_2 be the space spanned by ν_j 's. Define the "nearest-neighbor" random variables, $NN_1(Z), NN_2(Z)$ as follows:

$$(3.20) \quad NN_1(Z) = \operatorname{argmin}\{\|Z - \mu_i\| : 1 \leq i \leq k_1\}, \\ NN_2(Z) = \operatorname{argmin}\{\|Z - \nu_j\| : 1 \leq j \leq k_2\}.$$

Then as S_1 and S_2 are orthogonal to each other, it follows from Lemma (3.2) that when Z is a spherical multivariate Gaussian, the random variables $NN_1(Z)$ and $NN_2(Z)$ are independent. Similarly, it can be shown that when Z is a spherical multivariate Gaussian, the random variables $\overline{NN}_1$, and $\overline{NN}_2$ defined by,

$$(3.21) \quad (\overline{NN}_1(Z), \overline{NN}_2(Z)) = \operatorname{argmin}_{(i,j)}\{\|Z - \mu_i - \nu_j\|\},$$

are independent. Note that in equations (3.1), (3.12) we use inner products involving the mean vectors of different clusterings as the correlation measure. Thus, minimizing the correlation measure ideally leads to the mean vectors of different clusterings being orthogonal. Also, observe that we use nearest neighbor assignments of the form (3.20), (3.21) in our algorithms in Decorrelated-kmeans and Convolutional-EM. Thus, the decorrelation measures specified in equations (3.1) and (3.12) intuitively correspond to the labellings of the clusterings being independent.

Figure 2: Representative vectors obtained by Dec-kmeans and the parts obtained by Conv-EM. The bold line represents the separating hyperplane for the first clustering, while the dotted line represents the separating hyperplane for the second clustering. Conv-EM produced mean vectors $\{\mu_1, \mu_2, \nu_1, \nu_2\}$ and subsequently each of the four parts are obtained by $\mu_i + \nu_j$ ($i \in \{1, 2\}, j \in \{1, 2\}$).

3.3.2 Decorrelated-kmeans vs Convolutional-EM:

Decorrelated-kmeans (Algorithm 1) has a three-fold advantage over the "sum of the parts" approach (Algorithm 2):

- Computing the E-step exactly in the "sum of the parts" approach requires $O(k^T)$ computation for each data point, where T is the number of alternative clusterings. On the other hand, in Decorrelated-kmeans, each label assignment step requires just $O(kT)$ computations as the error terms for different clusterings are independent in (3.1). Thus, Decorrelated-kmeans is more scalable than Convolutional-EM with respect to the number of alternative clusterings.
- The M-step in the "sum of the parts" approach solves a non-convex problem and requires an iterative procedure to reach a local minimum. In the Decorrelated-kmeans approach, computing the representative vectors (the equivalent of M-step) requires solving a convex problem and the optimal solution can be written down in closed form. Hence, estimation of representative vectors is more accurate and efficient for Decorrelated-kmeans.
- Decorrelated-kmeans is a discriminative approach, while Convolutional-EM is a generative model based approach. Thus, the model assumptions are more stringent for the latter approach. This is observed empirically also, where Decorrelated-kmeans works well for all the datasets, but Convolutional-EM suffers on one

of the real-life datasets.

On the flip side, there is no natural interpretation of the “representative” vectors given by Decorrelated-kmeans. On the other hand, the means given by Convolutional-EM can naturally be interpreted as giving a part-based representation of the data. This argument is illustrated by Figure 2. The representative vectors obtained from Decorrelated-kmeans partition the data into two clusters accurately. But, they don’t give any intuitive characterization of the data. In contrast, Convolutional-EM is able to recover the four clusters in the data generated by the addition of two mixtures of Gaussians.

4 Experiments

We now provide experimental results on synthetic as well as real-world datasets to show the applicability of our methods. For real-world datasets we consider a music dataset from the text-mining domain and a portrait dataset from the computer-vision domain. We compare our methods against the factorial learning algorithms Co-operative Vector Quantization (CVQ)[8] and Multiple Cause Vector Quantization (MCVQ)[16]. We also compare against single-clustering algorithms such as k -means and NNMA. We will refer to the methods of Sections 3.1, 3.2 as Dec-kmeans (for Decorrelated-kmeans) and Conv-EM (for Convolutional-EM) respectively.

We also compare our methods against two simple heuristics:

1. *Feature Removal (FR)*: In this approach, we first cluster the data using k -means. Then, we remove the coordinates that have the most *correlation* with the labels in the obtained clustering. Next, we cluster the data again using the remaining features to obtain the alternative clustering. The correlation between a feature and the labels is taken to be proportional to the total weight of the mean vectors for the feature and inversely proportional to the entropy of the particular feature in the mean vectors. Formally:

$$C(i) = \frac{\sum_j \mu_j^i}{-\sum_j \left(\frac{\mu_j^i}{\sum_t \mu_t^i} \log \frac{\mu_j^i}{\sum_t \mu_t^i} \right)},$$

where μ_j^i is the i -th dimension of the j -th cluster.

2. *Orthogonal Projection (OP)*: This heuristic is motivated by principal gene shaving[12]. The heuristic proceeds by projecting the data onto the subspace orthogonal to the means of the first clustering and uses the projected data for computing the second clustering.

- (a) Cluster the data using a suitable method of clustering.

- (b) Let the means of the obtained clustering be $\mathbf{m}_1, \dots, \mathbf{m}_k$. Project the input matrix X onto the space orthogonal to the one spanned by the means $\mathbf{m}_1, \dots, \mathbf{m}_k$ to get X' .
- (c) Cluster the columns of X' to obtain a new set of labels, and compute the cluster means $\tilde{\mathbf{m}}_1, \dots, \tilde{\mathbf{m}}_k$.
- (d) Repeat steps (b),(c) with means $\tilde{\mathbf{m}}_1, \dots, \tilde{\mathbf{m}}_k$.
- (e) Until convergence, repeat steps (a)-(d).

4.1 Implementation Details: All the methods have been implemented in MATLAB. The implementation of MCVQ was obtained from the authors of [16]. Lee and Seung’s algorithm[15] is used for NNMA. Experiments were performed on a Linux machine with a 2.4 GHz Pentium IV processor and 1 GB main memory. For the real-world datasets, we report results in terms of accuracy with the true labels. As the number of clusters can be high in the synthetic datasets, we report results in terms of normalized mutual information (NMI) [19]. For all the experiments, accuracy/NMI is averaged over 100 runs.

4.2 Synthetic Datasets: For our experiments we generate synthetic datasets as a sum of independent components. Let $\mathcal{X}$ and $\mathcal{Y}$ be samples drawn from two independent mixtures of multivariate Gaussians. To evaluate our methods in various settings, we generate the final dataset $\mathcal{Z}$ by combining $\mathcal{X}$ and $\mathcal{Y}$ in three different ways. By viewing $\mathcal{X}$ and $\mathcal{Y}$ as the *components* of the datasets, and clustering based on these components we get two different clusterings of the data.

1. *Concatenated dataset*: This dataset is produced by simply concatenating the features of $\mathcal{X}$ and $\mathcal{Y}$, i.e., $\mathcal{Z} = \begin{bmatrix} \mathcal{X} \\ \mathcal{Y} \end{bmatrix}$.
2. *Partial overlap dataset*: In this dataset we allow a few of the features of $\mathcal{X}$ and $\mathcal{Y}$ to overlap. Specifically, let $\mathcal{X} = \begin{bmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 \end{bmatrix}$ and $\mathcal{Y} = \begin{bmatrix} \mathcal{Y}_1 \\ \mathcal{Y}_2 \end{bmatrix}$, where $\mathcal{X}_1, \mathcal{X}_2, \mathcal{Y}_1$ and $\mathcal{Y}_2$ all have the same dimensionality. Then, we form $\mathcal{Z} = \begin{bmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 + \mathcal{Y}_1 \\ \mathcal{Y}_2 \end{bmatrix}$.
3. *Sum dataset*: In this dataset, all of the features of $\mathcal{X}$ and $\mathcal{Y}$ overlap, i.e., $\mathcal{Z} = \mathcal{X} + \mathcal{Y}$.

In our experiments the dimensionality of $\mathcal{X}$ and $\mathcal{Y}$ was set to 30 and there were 3000 data points. We label each x_i and y_j according to the Gaussian from which they were sampled. Thus, each z is associated with two true-labels. Both our methods produce two disparate clusterings, and

Figure 3: NMI achieved by various methods on the Concatenated Dataset. Top figure shows NMI for the first clustering and bottom figure shows NMI for the second clustering. Overall, Dec-kmeans achieves the highest NMI, while MCVQ also performs well on this dataset.

we associate each clustering with a unique true-labeling and report NMI with respect to that true-labeling. We use the same procedure for CVQ and MCVQ. For k -means and NNMA², which produce just one clustering, we report the NMI of the clustering with respect to the true-labellings.

Figure 3 compares the NMI achieved by various methods on the Concatenated dataset. It can be seen from the figure that Conv-EM and Dec-kmeans outperform k -means and NNMA for both the clusterings, achieving an average improvement of 50 – 60% in NMI. Similarly, both Conv-EM and Dec-kmeans achieve significantly higher NMI than CVQ. Note that the Concatenated dataset satisfies MCVQ’s assumption that each dimension of the data is generated from one of the two factors. This is empirically confirmed by the results, as MCVQ not only outperforms CVQ but also performs competitively with Conv-EM and Dec-kmeans.

Figure 4 compares the NMI for various methods on the Overlap dataset. Clearly, Conv-EM and Dec-kmeans perform better than both CVQ and MCVQ. Note that when the number of clusters is small, NMI of MCVQ with respect to both the clusterings drops to around 0.6. This is probably because the overlap dataset does not satisfy the model as-

Figure 4: NMI achieved by various methods on the Overlap Dataset. Top figure shows NMI for the first clustering and bottom figure shows NMI for the second clustering. Dec-kmeans and Conv-EM achieves similar NMI. Both achieve higher NMI than MCVQ or CVQ.

sumptions of MCVQ.

Figure 5 shows the NMI achieved by various methods on the Sum dataset. For this dataset also, both Conv-EM and Dec-kmeans perform comparably and both the methods achieve significantly higher NMI than other methods. Interestingly, NMI for MCVQ is even lower than the single-clustering algorithms (k -means and NNMA). This could be because the modeling assumption of MCVQ – each dimension in the data is generated by exactly one factor – is completely violated in the Sum dataset. Note that, although CVQ is designed to model the Sum datasets, it performs poorly compared to Conv-EM and Dec-kmeans. This trend can be attributed to the fact that due to the lack of regularization CVQ selects one of the many possible solutions to its optimization problem, which may or may not correspond to good disparate clusterings.

Also note that Conv-EM does not perform significantly better than Dec-kmeans, even though the datasets fit the Conv-EM model well. This is probably because of the non-convex nature of the optimization problem for the M-step in Conv-EM, due to which the maximum likelihood estimation gets stuck in a local minimum.

²As NNMA is useful for the non-negative data only, we made the data non-negative by choosing the means sufficiently far away from origin.

4.3 Real-World Datasets

Figure 5: NMI achieved by various methods on the Sum Dataset. Top figure shows NMI for the first clustering and bottom figure shows NMI for the second clustering. Dec-kmeans and Conv-EM achieves similar NMI. NMI achieved by MCVQ is very low.

4.3.1 Music Dataset: The music dataset is a collection of 270 documents, with each document being a review of a classical music piece taken from *amazon.com*. Each music piece is composed by one of Beethoven, Mozart or Mendelssohn and is in one of symphony, sonata or concerto forms. Thus, the documents can be clustered based on the composer or the genre of the musical piece. For the experiments a term-document matrix was formed with dimensionality 258 after stop word removal and stemming.

Table 1 shows that although all the methods are able to recover the true clustering for composer, most of the algorithms perform poorly for the genre based clustering. In particular, *k*-means and NNMA perform very poorly for the clustering based on genre as they produce just one clustering which has high NMI with the clustering based on composers. Note that in this dataset, the sets of features (words) that determine clustering with respect to composer and genre respectively are almost disjoint. Hence, methods like Feature Removal and Orthogonal Projection, which try to identify disjoint sets of features for the clusterings work fairly well. But, both MCVQ and CVQ algorithm achieve very low accuracy as they do not try to find *decorrelated* clusterings. Both our methods outperform the baseline algorithms.

Table 1: Accuracy achieved by various methods on the Music dataset, which is a collection of text documents. Dec-kmeans performs the best on this dataset. CVQ and MCVQ perform very poorly compared to Conv-EM

Method\Type	Composer	Genre
NNMA	1.00	0.40
<i>k</i> -means	0.89	0.41
Feature Removal	0.97	0.64
Orthogonal Projection	0.99	0.66
CVQ	0.97	0.57
MCVQ	0.91	0.53
Conv-EM	1.00	0.65
Dec-kmeans	1.00	0.69

4.3.2 Portrait Dataset: The Portrait dataset consists of 324 images obtained from Yale Face Dataset B [7]. Each image in the dataset is a portrait of one of three people, in one of three poses in different backgrounds. The dimensionality of each image is 64×64 . As a first step we reduce the dimensionality of the data to 300 by using principal component analysis. As in the music dataset, the current dataset can be clustered in two natural ways - by the person in the picture or the pose. Table 2 shows that both *k*-means and NNMA perform poorly with respect to both the clusterings. This shows that in the datasets where there is more than one natural clustering, traditional clustering algorithms could fail to find even one good clustering. Hence, it can be beneficial to use alternative clustering methods even if one is interested in obtaining a single clustering.

Our hypothesis is that unlike the music dataset, there are no dominant features for any of the clusterings in this dataset. This hypothesis can be justified by observing the poor accuracies of methods like Feature Removal and Orthogonal Projection. Conv-EM outperforms baseline algorithms CVQ and MCVQ significantly, but interestingly Dec-kmeans achieves an even higher accuracy of 84% and 78% for the two clusterings.

5 Conclusions and Future Work

We address the difficult problem of uncovering disparate clusterings from the data in a totally unsupervised setting. We present two novel approaches for the problem - a decorrelated *k*-means approach and a sum of parts approach. In the first approach, we introduce a new regularization for *k*-means to uncover decorrelated clusterings. We provide theoretical justification for using the proposed decorrelation measure. The sum of parts approach leads us to the interesting problem of learning a convolution of mixture models and we present a regularized EM algorithm for learning a convolution of mixtures of spherical Gaussians. We address the problem of identifiability for learning a convolution of

Table 2: Accuracy achieved by various methods on the Portrait dataset, which is a collection of images. Dec-kmeans outperforms all other methods by a significant margin. Conv-EM achieves better accuracy than all other methods, especially CVQ and MCVQ.

Method\ Type	Person	Pose
NNMA	0.51	0.49
k -means	0.66	0.56
Feature Removal	0.56	0.48
Orthogonal Projection	0.66	0.70
CVQ	0.53	0.51
MCVQ	0.64	0.51
Conv-EM	0.69	0.72
Dec-kmeans	0.84	0.78

mixtures of Gaussians by using a regularization geared for providing disparate clusterings. We demonstrate the effectiveness and robustness of our algorithms on synthetic and real-world datasets. On each of these datasets, we significantly improve upon the accuracy achieved by existing factorial learning methods such as CVQ and MCVQ. Our methods also outperform the traditional clustering algorithms like k -means and NNMA.

For future work, it would be of interest to study the problem of learning a convolution of mixtures for a more general class of distributions and look for other settings where learning convolutions of mixtures could be useful. We also plan to further investigate the properties of the decorrelation measure, especially for more general distributions of the data. A problem that we do not address in this paper is model selection - choosing the number of clusters and the number of clusterings. Good heuristics for choosing these parameters would be very useful. Also, a more detailed comparison of Conv-EM and Dec-kmeans would be useful; it would be interesting to understand the comparable performance of Conv-EM and Dec-kmeans for the synthetic datasets which fit the convolution model well.

Acknowledgements: This research was supported by NSF grant CCF-0431257, NSF Career Award ACI-0093404, and NSF-ITR award IIS-0325116.

References

- [1] BAE, E. AND BAILEY, J., *COALA: A novel approach for the extraction of an alternate clustering of high quality and high dissimilarity*, in ICDM, 2006, pp. 53–62.
- [2] BILENKO, M., BASU, S., AND MOONEY, R.J., *Integrating constraints and metric learning in semi-supervised clustering*, in ICML, 2004, pp. 81–88.
- [3] BISHOP, C. M., *Neural networks for pattern recognition*, Oxford University Press, Oxford, UK, 1996.

- [4] I. DAVIDSON, S. S. RAVI, AND M. ESTER, *Efficient incremental constrained clustering*, in KDD, 2007, pp. 240–249.
- [5] R. O. DUDA, P. E. HART, AND D. G. STORK, *Pattern Classification*, Wiley-Interscience, November 2000.
- [6] W. R. GAFFEY, *A consistent estimator of a component of a convolution*, The Annals of Mathematical Statistics, 30 (1959), pp. 198–205.
- [7] A. GEORGHIADES, P. BELHUMEUR, AND D. KRIEGMAN, *From few to many: Illumination cone models for face recognition under variable lighting and pose*, IEEE Trans. Pattern Anal. Mach. Intelligence, 23 (2001), pp. 643–660.
- [8] Z. GHAHRAMANI, *Factorial learning and the EM algorithm*, in NIPS, 1994, pp. 617–624.
- [9] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins Univ. Press, second ed., 1989.
- [10] D. GONDEK AND T. HOFMANN, *Non-redundant data clustering*, in ICDM, 2004, pp. 75–82.
- [11] D. GONDEK, S. VAITHYANATHAN, AND A. GARG, *Clustering with model-level constraints*, in SDM, 2005, pp. 126–138.
- [12] T. HASTIE, R. TIBSHIRANI, A. EISEN, R. LEVY, L. STAUDT, D. CHAN, AND P. BROWN, *Gene shaving as a method for identifying distinct sets of genes with similar expression patterns*, Genome Biology, (2000).
- [13] G. E. HINTON AND R. S. ZEMEL, *Autoencoders, minimum description length and Helmholtz free energy*, in NIPS, 1993, pp. 3–10.
- [14] A. HYVÄRINEN AND E. OJA, *Independent component analysis: algorithms and applications*, Neural Networks, 13 (2000), pp. 411–430.
- [15] D. D. LEE AND H. S. SEUNG, *Algorithms for non-negative matrix factorization*, in NIPS, 2000, pp. 556–562.
- [16] D. A. ROSS AND R. S. ZEMEL, *Learning parts-based representations of data*, JMLR, 7 (2006), pp. 2369–2397.
- [17] SAMANIEGO, F. J. AND JONES, L. E., *Maximum likelihood estimation for a class of multinomial distributions arising in reliability*, Journal of the Royal Statistical Society. Series B (Methodological), 43 (1981), pp. 46–52.
- [18] SCLOVE, S. L. AND VAN RYZIN, J., *Estimating the parameters of a convolution*, Journal of the Royal Statistical Society. Series B (Methodological), 31 (1969), pp. 181–191.
- [19] A. STREHL AND J. GHOSH, *Cluster ensembles — a knowledge reuse framework for combining multiple partitions*, JMLR, 3 (2002), pp. 583–617.
- [20] J. B. TENENBAUM AND W. T. FREEMAN, *Separating style and content with bilinear models*, Neural Computation, 12 (2000), pp. 1247–1283.
- [21] K. WAGSTAFF AND C. CARDIE, *Clustering with instance-level constraints*, in ICML, 2000, pp. 1103–1110.
- [22] K. WAGSTAFF, C. CARDIE, S. ROGERS, AND S. SCHRÖDL, *Constrained k -means clustering with background knowledge*, in ICML, 2001, pp. 577–584.
- [23] E. P. XING, A. Y. NG, M. I. JORDAN, AND S. RUSSELL, *Distance metric learning, with application to clustering with side-information*, in NIPS, vol. 15, 2003, pp. 521–528.
- [24] W. I. ZANGWILL, *Nonlinear Programming: A Unified Approach*, Englewood Cliffs: Prentice-Hall, 1969.