# Spectral Relaxation for K-means Clustering 

Hongyuan Zha \& Xiaofeng He<br>Dept. of Comp. Sci. \& Eng.<br>The Pennsylvania State University<br>University Park, PA 16802<br>\{zha,xhe\}@cse.psu.edu

Chris Ding \& Horst Simon<br>NERSC Division<br>Lawrence Berkeley National Lab.<br>UC Berkeley, Berkeley, CA 94720<br>\{chqding,hdsimon\}@lbl.gov

Ming Gu<br>Dept. of Mathematics<br>UC Berkeley, Berkeley, CA 95472<br>mgu@math.berkeley.edu


#### Abstract

The popular K-means clustering partitions a data set by minimizing a sum-of-squares cost function. A coordinate descend method is then used to find local minima. In this paper we show that the minimization can be reformulated as a trace maximization problem associated with the Gram matrix of the data vectors. Furthermore, we show that a relaxed version of the trace maximization problem possesses global optimal solutions which can be obtained by computing a partial eigendecomposition of the Gram matrix, and the cluster assignment for each data vectors can be found by computing a pivoted QR decomposition of the eigenvector matrix. As a by-product we also derive a lower bound for the minimum of the sum-of-squares cost function.


## 1 Introduction

K-means is a very popular method for general clustering [6]. In K-means clusters are represented by centers of mass of their members, and it can be shown that the K-means algorithm of alternating between assigning cluster membership for each data vector to the nearest cluster center and computing the center of each cluster as the centroid of its member data vectors is equivalent to finding the minimum of a sum-of-squares cost function using coordinate descend. Despite the popularity of Kmeans clustering, one of its major drawbacks is that the coordinate descend search method is prone to local minima. Much research has been done on computing refined initial points and adding explicit constraints to the sum-of-squares cost function for K-means clustering so that the search can converge to better local minimum [1, 2]. In this paper we tackle the problem from a different angle: we find an equivalent formulation of the sum-of-squares minimization as a trace maximization problem with special constraints; relaxing the constraints leads to a maximization problem
that possesses optimal global solutions. As a by-product we also have an easily computable lower bound for the minimum of the sum-of-squares cost function. Our work is inspired by $[9,3]$ where connection to Gram matrix and extension of Kmeans method to general Mercer kernels were investigated.

The rest of the paper is organized as follows: in section 2, we derive the equivalent trace maximization formulation and discuss its spectral relaxation. In section 3, we discuss how to assign cluster membership using pivoted QR decomposition, taking into account the special structure of the partial eigenvector matrix. Finally, in section 4 , we illustrate the performance of the clustering algorithms using document clustering as an example.
Notation. Throughout, $\|\cdot\|$ denotes the Euclidean norm of a vector. The trace of a matrix $A$, i.e., the sum of its diagonal elements, is denoted as trace $(A)$. The Frobenius norm of a matrix $\|A\|_{F}=\sqrt{\operatorname{trace}\left(A^{T} A\right)}$. $I_{n}$ denotes identity matrix of order $n$.

## 2 Spectral Relaxation

Given a set of $m$-dimensional data vectors $a_{i}, i=1, \ldots, n$, we form the $m$-by- $n$ data $\operatorname{matrix} A=\left[a_{1}, \ldots, a_{n}\right]$. A partition $\Pi$ of the date vectors can be written in the following form

$$
\begin{equation*}
A E=\left[A_{1}, \ldots, A_{k}\right], \quad A_{i}=\left[a_{1}^{(i)}, \ldots, a_{s_{i}}^{(i)}\right] \tag{1}
\end{equation*}
$$

where $E$ is a permutation matrix, and $A_{i}$ is $m$-by- $s_{i}$, i.e., the $i$ th cluster contains the data vectors in $A_{i}$. For a given partition $\Pi$ in (1), the associated sum-of-squares cost function is defined as

$$
\operatorname{ss}(\Pi)=\sum_{i=1}^{k} \sum_{s=1}^{s_{i}}\left\|a_{s}^{(i)}-m_{i}\right\|^{2}, \quad m_{i}=\sum_{s=1}^{s_{i}} a_{s}^{(i)} / s_{i}
$$

i.e., $m_{i}$ is the mean vector of the data vectors in cluster $i$. Let $e$ be a vector of appropriate dimension with all elements equal to one, it is easy to see that $m_{i}=A_{i} e / s_{i}$ and

$$
\mathrm{ss}_{i} \equiv \sum_{s=1}^{s_{i}}\left\|a_{s}^{(i)}-m_{i}\right\|^{2}=\left\|A_{i}-m_{i} e^{T}\right\|_{F}^{2}=\left\|A_{i}\left(I_{s_{i}}-e e^{T} / s_{i}\right)\right\|_{F}^{2}
$$

Notice that $I_{s_{i}}-e e^{T} / s_{i}$ is a projection matrix and $\left(I_{s_{i}}-e e^{T} / s_{i}\right)^{2}=I_{s_{i}}-e e^{T} / s_{i}$, it follows that

$$
\operatorname{ss}_{i}=\operatorname{trace}\left(A_{i}\left(I_{s_{i}}-e e^{T} / s_{i}\right) A_{i}^{T}\right)=\operatorname{trace}\left(\left(I_{s_{i}}-e e^{T} / s_{i}\right) A_{i}^{T} A_{i}\right)
$$

Therefore,

$$
\operatorname{ss}(\Pi)=\sum_{i=1}^{k} \mathrm{ss}_{i}=\sum_{i=1}^{k}\left(\operatorname{trace}\left(A_{i}^{T} A_{i}\right)-\left(\frac{e^{T}}{\sqrt{s_{i}}}\right) A_{i}^{T} A_{i}\left(\frac{e}{\sqrt{s_{i}}}\right)\right)
$$

Let the $n$-by- $k$ orthonormal matrix $X$ be

$$
X=\begin{align*}
& s_{1}  \tag{2}\\
& s_{2} \\
& \vdots \\
& s_{k}
\end{align*}\left(\begin{array}{llll}
e / \sqrt{s_{1}} & & & \\
& e / \sqrt{s_{2}} & & \\
& & \vdots & \\
& & & e / \sqrt{s_{k}}
\end{array}\right)
$$

The sum-of-squares cost function can now be written as

$$
\mathrm{ss}(\Pi)=\operatorname{trace}\left(A^{T} A\right)-\operatorname{trace}\left(X^{T} A^{T} A X\right)
$$

and its minimization is equivalent to

$$
\max \left\{\operatorname{trace}\left(X^{T} A^{T} A X\right) \mid X \text { of the form in }(2)\right\}
$$

Remark. Without loss of generality, let $E=I$ in (1). If we let $x_{i}$ be the cluster indicator vector, i.e.,

$$
x_{i}^{T}=[0, \ldots, 0, \underbrace{1, \ldots, 1}_{s_{i}}, 0, \ldots, 0] \text {. }
$$

Then it is easy to see that

$$
\operatorname{trace}\left(X^{T} A^{T} A X\right)=\sum_{i=1}^{k} \frac{x_{i}^{T} A^{T} A x_{i}}{x_{i}^{T} x_{i}}=\sum_{i=1}^{k} \frac{\left\|A x_{i}\right\|^{2}}{\left\|x_{i}\right\|^{2}}
$$

Using the partition in (1), the right-hand side of the above can be written as

$$
\sum_{i=1}^{k} s_{i}\left\|\frac{A_{i} e}{s_{i}}\right\|^{2}=\sum_{i=1}^{k} s_{i}\left\|m_{i}\right\|^{2}
$$

a weighted sum of the squared Euclidean norms of the mean vector of each clusters.
REmark. If we consider the elements of the Gram matrix $A^{T} A$ as measuring similarity between data vectors, then we have shown that Euclidean distance leads to Euclidean inner-product similarity. This inner-product can be replaced by a general Mercer kernel as is done in [9,3].
Ignoring the special structure of $X$ and let it be an arbitrary orthonormal matrix, we obtain a relaxed maximization problem

$$
\begin{equation*}
\max _{X^{T} X=I_{k}} \operatorname{trace}\left(X^{T} A^{T} A X\right) \tag{3}
\end{equation*}
$$

It turns out the above trace maximization problem has a closed-form solution.
Theorem. (Ky Fan) Let $H$ be a symmetric matrix with eigenvalues

$$
\lambda_{1} \geq \lambda_{2} \geq \cdots \geq \lambda_{n}
$$

and the corresponding eigenvectors $U=\left[u_{1}, \ldots, u_{n}\right]$. Then

$$
\lambda_{1}+\cdots \lambda_{k}=\max _{X^{T} X=I_{k}} \operatorname{trace}\left(X^{T} H X\right)
$$

Moreover, the optimal $X^{*}$ is given by $X^{*}=\left[u_{1}, \ldots, u_{k}\right] Q$ with $Q$ an arbitrary orthogonal matrix.
It follows from the above theorem that we need to compute the largest $k$ eigenvectors of the Gram matrix $A^{T} A$. As a by-product, we have

$$
\begin{equation*}
\min _{\Pi} \mathrm{ss}(\Pi) \geq \operatorname{trace}\left(A^{T} A\right)-\max _{X^{T} X=I_{k}} \operatorname{trace}\left(X^{T} A^{T} A X\right)=\sum_{i=k+1}^{\min \{m, n\}} \sigma_{i}^{2}(A) \tag{4}
\end{equation*}
$$

where $\sigma_{i}(A)$ is the $i$ largest singular value of $A$. This gives a lower bound for the minimum of the sum-of-squares cost function.

Remark. It is easy to see from the above derivation that we can replace $A$ with $A-a e^{T}$, where $a$ is an arbitrary vector. Then we have the following lower bound

$$
\min _{\Pi} \operatorname{ss}(\Pi) \geq \max _{a} \sum_{i=k+1}^{\min \{m, n\}} \sigma_{i}^{2}\left(A-a e^{T}\right)
$$

REmark. One might also try the following approach: notice that

$$
\left\|A_{i}-m_{i} e^{T}\right\|_{F}^{2}=\frac{1}{2 s_{i}} \sum_{a_{j} \in A_{i}} \sum_{a_{j^{\prime}} \in A_{i}}\left\|a_{j}-a_{j^{\prime}}\right\|^{2}
$$

Let $W=\left(\left\|a_{i}-a_{j}\right\|^{2}\right)_{i, j=1}^{n}$, and and $x_{i}=\left[x_{i j}\right]_{j=1}^{n}$ with

$$
x_{i j}= \begin{cases}1 & \text { if } a_{j} \in A_{i} \\ 0 & \text { otherwise }\end{cases}
$$

Then

$$
\operatorname{ss}(\Pi)=\frac{1}{2} \sum_{i=1}^{k} \frac{x_{i}^{T} W x_{i}}{x_{i}^{T} x_{i}} \geq \frac{1}{2} \min _{Z^{T} Z=I_{k}} Z^{T} W Z=\frac{1}{2} \sum_{i=n-k+1}^{n} \lambda_{i}(W) .
$$

Unfortunately, some of the smallest eigenvalues of $W$ can be negative.
Let $X_{k}$ be the $n$-by- $k$ matrix consisting of the $k$ largest eigenvectors of $A^{T} A$. Each row of $X_{k}$ corresponds to a data vector, and the above process can be considered as transforming the original data vectors which live in a $m$-dimensional space to new data vectors which now live in a $k$-dimensional space. One might be attempted to compute the cluster assignment by applying the ordinary K-means method to those data vectors in the reduced dimension space. In the next section, we discuss an alternative that takes into account the structure of the eigenvector matrix $X_{k}$ [5].
REMARK. The similarity of the projection process to principal component analysis is deceiving: the goal here is not to reconstruct the data matrix using a low-rank approximation but rather to capture its cluster structure.

## 3 Cluster Assignment Using Pivoted QR Decomposition

Without loss of generality, let us assume that the best partition of the data vectors in $A$ that minimizes $\mathrm{ss}(\Pi)$ is given by $A=\left[A_{1}, \ldots, A_{k}\right]$, each submatrix $A_{i}$ corresponding to a cluster. Now write the Gram matrix of $A$ as

$$
A^{T} A=\left[\begin{array}{cccc}
A_{1}^{T} A_{1} & 0 & \cdots & 0 \\
0 & A_{2}^{T} A_{2} & \cdots & 0 \\
\vdots & \vdots & \cdots & \vdots \\
0 & 0 & \cdots & A_{k}^{T} A_{k}
\end{array}\right]+E \equiv B+E
$$

If the overlaps among the clusters represented by the submatrices $A_{i}$ are small, then the norm of $E$ will be small as compare with the block diagonal matrix $B$ in the above equation. Let the largest eigenvector of $A_{i}^{T} A_{i}$ be $y_{i}$, and

$$
A_{i}^{T} A_{i} y_{i}=\mu_{i} y_{i}, \quad\left\|y_{i}\right\|=1, \quad i=1, \ldots, k
$$

then the columns of the matrix

$$
Y_{k}=\begin{aligned}
& s_{1} \\
& s_{2} \\
& \vdots \\
& s_{k}
\end{aligned}\left(\begin{array}{llll}
y_{1} & & & \\
& y_{2} & & \\
& & \vdots & \\
& & & y_{k}
\end{array}\right)
$$

span an invariant subspace of $B$. Let the eigenvalues and eigenvectors of $A^{T} A$ be

$$
\lambda_{1} \geq \lambda_{2} \geq \ldots \geq \lambda_{n}, \quad A^{T} A x_{i}=\lambda_{i} x_{i}, \quad i=1, \ldots, n
$$

Assume that there is a gap between the two eigenvalue sets $\left\{\mu_{1}, \ldots \mu_{k}\right\}$ and $\left\{\lambda_{k+1}, \ldots \lambda_{n}\right\}$, i.e.,

$$
0<\delta=\min \left\{\left|\mu_{i}-\lambda_{j}\right| \mid i=1, \ldots, k, j=k+1, \ldots, n\right\}
$$

Then Davis-Kahan $\sin (\Theta)$ theorem states that $\left\|Y_{k}^{T}\left[x_{k+1}, \ldots, x_{n}\right]\right\| \leq\|E\| / \delta[11$, Theorem 3.4]. After some manipulation, it can be shown that

$$
X_{k} \equiv\left[x_{1}, \ldots, x_{k}\right]=Y_{k} V+O(\|E\|)
$$

where $V$ is an $k$-by- $k$ orthogonal matrix. Ignoring the $O(\|E\|)$ term, we see that

$$
X_{k}^{T}=[\underbrace{y_{11} v_{1}, \ldots, y_{1 s_{1}} v_{1}}_{\text {cluster } 1}, \ldots, \underbrace{y_{k 1} v_{k}, \ldots, y_{k s_{k}} v_{k}}_{\text {cluster } k}]
$$

where we have used $y_{i}^{T}=\left[y_{i 1}, \ldots, y_{i s_{i}}\right]$, and $V^{T}=\left[v_{1}, \ldots, v_{k}\right]$. A key observation is that all the $v_{i}$ are orthogonal to each other: once we have selected a $v_{i}$, we can jump to other clusters by looking at the orthogonal complement of $v_{i}$. Also notice that $\left\|y_{i}\right\|=1$, so the elements of $y_{i}$ can not be all small. A robust implementation of the above idea can be obtained as follows: we pick a column of $X_{k}{ }^{T}$ which has the largest norm, say, it belongs to cluster $i$, we orthogonalize the rest of the columns of $X_{k}^{T}$ against this column. For the columns belonging to cluster $i$ the residual vector will have small norm, and for the other columns the residual vectors will tend to be not small. We then pick another vector with the largest residual norm, and orthogonalize the other residual vectors against this residual vector. The process can be carried out $k$ steps, and it turns out to be exactly QR decomposition with column pivoting applied to $X_{k}{ }^{T}$ [4], i.e., we find a permutation matrix $P$ such that

$$
X_{k}^{T} P=Q R=Q\left[R_{11}, R_{12}\right]
$$

where $Q$ is a $k$-by- $k$ orthogonal matrix, and $R_{11}$ is a $k$-by- $k$ upper triangular matrix. We then compute the matrix

$$
\hat{R}=R_{11}^{-1}\left[R_{11}, R_{12}\right] P^{T}=\left[I_{k}, R_{11}^{-1} R_{12}\right] P^{T}
$$

Then the cluster membership of each data vector is determined by the row index of the largest element in absolute value of the corresponding column of $\hat{R}$.

Remark. Sometimes it may be advantageous to include more than $k$ eigenvectors to form $X_{s}{ }^{T}$ with $s>k$. We can still use QR decomposition with column pivoting to select $k$ columns of $X_{s}{ }^{T}$ to form an $s$-by- $k$ matrix, say $X$. Then for each column $z$ of $X_{s}{ }^{T}$ we compute the least squares solution of $t^{*}=\operatorname{argmin}_{t \in \mathcal{R}^{k}}\|z-X t\|$. Then the cluster membership of $z$ is determined by the row index of the largest element in absolute value of $t^{*}$.

## 4 Experimental Results

In this section we present our experimental results on clustering a dataset of newsgroup articles submitted to 20 newsgroups. ${ }^{1}$ This dataset contains about 20,000 articles (email messages) evenly divided among the 20 newsgroups. We list the names of the newsgroups together with the associated group labels.

[^0]

Figure 1: Clustering accuracy for five newsgroups NG2/NG9/NG10/NG15/NG18: p-QR vs. p-Kmeans (left) and p-Kmeans vs. Kmeans (right)

```
NG1: alt.atheism NG2: comp.graphics
NG3: comp.os.ms-windows.misc NG4: comp.sys.ibm.pc.hardware
NG5: comp.sys.mac.hardware NG6: comp.windows.x
NG7:misc.forsale NG8: rec.autos
NG9:rec.motorcycles NG10: rec.sport.baseball
NG11:rec.sport.hockey NG12: sci.crypt
NG13:sci.electronics NG14: sci.med
NG15:sci.space NG16: soc.religion.christian
NG17:talk.politics.guns NG18: talk.politics.mideast
NG19:talk.politics.misc NG20: talk.religion.misc
```

We used the bow toolkit to construct the term-document matrix for this dataset, specifically we use the tokenization option so that the UseNet headers are stripped, and we also applied stemming [8]. The following three preprocessing steps are done: 1) we apply the usual tf.idf weighting scheme; 2) we delete words that appear too few times; 3) we normalized each document vector to have unit Euclidean length.
We tested three clustering algorithms: 1) $\mathrm{p}-\mathrm{QR}$, this refers to the algorithm using the eigenvector matrix followed by pivoted QR decomposition for cluster membership assignment; 2) p-Kmeans, we compute the eigenvector matrix, and then apply K-means on the rows of the eigenvector matrix; 3) K-means, this is K-means directly applied to the original data vectors. For both K-means methods, we start with a set of cluster centers chosen randomly from the (projected) data vectors, and we aslo make sure that the same random set is used for both for comparison. To assess the quality of a clustering algorithm, we take advantage of the fact that the newsgroup data are already labeled and we measure the performance by the accuracy of the clustering algorithm against the document category labels [10]. In particular, for a $k$ cluster case, we compute a $k$-by- $k$ confusion matrix $C=\left[c_{i j}\right]$ with $c_{i j}$ the number of documents in cluster $i$ that belongs to newsgroup category $j$. It is actually quite subtle to compute the accuracy using the confusion matrix because we do not know which cluster matches which newsgroup category. An optimal way is to solve the following maximization problem

$$
\max \{\operatorname{trace}(C P) \mid P \text { is a permutation matrix }\},
$$

and divide the maximum by the total number of documents to get the accuracy. This is equivalent to finding perfect matching a complete weighted bipartite graph, one can use Kuhn-Munkres algorithm [7]. In all our experiments, we used a greedy algorithm to compute a sub-optimal solution.

Table 1: Comparison of $\mathrm{p}-\mathrm{QR}$, p -Kmeans, and K-means for two-way clustering

| Newsgroups | p-QR | p-Kmeans | K-means |
| :--- | ---: | ---: | ---: |
| NG1/NG2 | $89.29 \pm 7.51 \%$ | $89.62 \pm 6.90 \%$ | $76.25 \pm 13.06 \%$ |
| NG2/NG3 | $62.37 \pm 8.39 \%$ | $63.84 \pm 8.74 \%$ | $61.62 \pm 8.03 \%$ |
| NG8/NG9 | $75.88 \pm 8.88 \%$ | $77.64 \pm 9.00 \%$ | $65.65 \pm 9.26 \%$ |
| NG10/NG11 | $73.32 \pm 9.08 \%$ | $74.86 \pm 8.89 \%$ | $62.04 \pm 8.61 \%$ |
| NG1/NG15 | $73.32 \pm 9.08 \%$ | $74.86 \pm 8.89 \%$ | $62.04 \pm 8.61 \%$ |
| NG18/NG19 | $63.86 \pm 6.09 \%$ | $64.04 \pm 7.23 \%$ | $63.66 \pm 8.48 \%$ |

Table 2: Comparison of $\mathrm{p}-\mathrm{QR}, \mathrm{p}$-Kmeans, and K -means for multi-way clustering

| Newsgroups | p-QR | p-Kmeans | K-means |
| :--- | ---: | ---: | ---: |
| NG2/NG3/NG4/NG5/NG6 (50) | $40.36 \pm 5.17 \%$ | $41.15 \pm 5.73 \%$ | $35.77 \pm 5.19 \%$ |
| NG2/NG3/NG4/NG5/NG6 (100) | $41.67 \pm 5.06 \%$ | $42.53 \pm 5.02 \%$ | $37.20 \pm 4.39 \%$ |
| NG2/NG9/NG10/NG15/NG18 (50) | $777.83 \pm 9.26 \%$ | $70.13 \pm 11.6 \% \%$ | $58.10 \pm 9.60 \%$ |
| NG2/NG9/NG10/NG15/NG18 (100) | $79.91 \pm 9.90 \%$ | $75.56 \pm 10.63 \%$ | $66.37 \pm 10.89 \%$ |
| NG1/NG5/NG7//NG88GG11/ | (50) | $60.21 \pm 4.88 \%$ | $58.18 \pm 4.41 \%$ |
| NG12/NG13/NG14/NG15/NG17 | $40.18 \pm 4.64 \%$ |  |  |
| NG1/NG5/NG7/NG8/NG11/ | $(100)$ | $65.08 \pm 5.14 \%$ | $58.99 \pm 5.22 \%$ |
| NG12/NG13/NG14/NG15/NG17 | $48.33 \pm 5.64 \%$ |  |  |

Example 1. In this example, we look at binary clustering. We choose 50 random document vectors each from two newsgroups. We tested 100 runs for each pair of newsgroups, and list the means and standard deviations in Table 1. The two clustering algorithms $\mathrm{p}-\mathrm{QR}$ and p -Kmeans are comparable to each other, and both are better and sometimes substantially better than K-means.
Example 2. In this example, we consider $k$-way clustering with $k=5$ and $k=10$. Three newsgroup sets are chosen with 50 and 100 random samples from each newsgroup as indicated in the parenthesis. Again 100 runs are used for each tests and the means and standard deviations are listed in Table 2. Moreover, in Figure 1, we also plot the accuracy for the 100 runs for the test NG2/NG9/NG10/NG15/NG18 (50). Both p-QR and p-Kmeans perform better than Kmeans. For newsgroup sets with small overlaps, $\mathrm{p}-\mathrm{QR}$ performs better than p -Kmeans. This might be explained by the fact that $\mathrm{p}-\mathrm{QR}$ explores the special structure of the eigenvector matrix and is therefore more efficient. As a less thorough comparison with the information bottleneck method used in [10], there for 15 runs of NG2/NG9/NG10/NG15/NG18 (100) mean accuracy $56.67 \%$ with maximum accuracy $67.00 \%$ is obtained. For 15 runs of the 10 newsgroup set with 50 samples, mean accuracy $35.00 \%$ with maximum accuracy about $40.00 \%$ is obtained.

Example 3. We compare the lower bound given in (4). We only list a typical sample from NG2/NG9/NG10/NG15/NG18 (50). The column with "NG labels" indicates clustering using the newsgroup labels and by definition has $100 \%$ accuracy. It is quite clear that the newsgroup categories are not completely captured by the sum-of-squares cost function because p-QR and "NG labels" both have higher accuracy but also larger sum-of-squares values. Interestingly, it seems that p-QR captures some of this information of the newsgroup categories.

|  | p-QR | p-Kmeans | K-means | NG labels | lower bound |
| :---: | :---: | :---: | :---: | :---: | :---: |
| accuracy | $86.80 \%$ | $83.60 \%$ | $57.60 \%$ | $100 \%$ | N/A |
| $\mathrm{ss}(\Pi)$ | 224.1110 | 223.8966 | 228.8416 | 224.4040 | 219.0266 |

## Acknowledgments

This work was supported in part by NSF grant CCR-9901986 and by Department of Energy through an LBL LDRD fund.

## References

[1] P. S. Bradley and Usama M. Fayyad. (1998). Refining Initial Points for K-Means Clustering. Proc. 15th International Conf. on Machine Learning, 91-99.
[2] P. S. Bradley, K. Bennett and A. Demiritz. Constrained K-means Clustering. Microsoft Research, MSR-TR-2000-65, 2000.
[3] M. Girolani. (2001). Mercer Kernel Based Clustering in Feature Space. To appear in IEEE Transactions on Neural Networks.
[4] G. Golub and C. Van Loan. (1996). Matrix Computations. Johns Hopkins University Press, 3rd Edition.
[5] Ming Gu, Hongyuan Zha, Chris Ding, Xiaofeng He and Horst Simon. (2001). Spectral Embedding for K-Way Graph Clustering. Technical Report, Department of Computer Science and Engineering, CSE-01-007, Pennsylvania State University.
[6] J.A. Hartigan and M.A. Wong. (1979). A K-means Clustering Algorithm. Applied Statistics, 28:100-108.
[7] L. Lovasz and M.D. Plummer. (1986) Matching Theory. Amsterdam: North Holland.
[8] A. McCallum. Bow: A toolkit for statistical language modeling, text retrieval, classification and clustering. http://www.cs.cmu.edu/ mccallum/bow.
[9] B. Schölkopf, A. Smola and K.R. Müller. (1998). Nonlinear Component Analysis as a Kernel Eigenvalue Problem. Neural Computation, 10: 1299-1219.
[10] N. Slonim and N. Tishby. (2000). Document clustering using word clusters via the information bottleneck method. Proceedings of SIGIR-2000.
[11] G.W. Stewart and J.G. Sun. (1990). Matrix Perturbation Theory. Academic Press, San Diego, CA.


[^0]:    ${ }^{1}$ The newsgroup dataset together with the bow toolkit for processing it can be downloaded from http://www.cs.cmu.edu/afs/cs/project/theo-11/www/naive-bayes.html.

