Using the left Gram matrix to cluster high-dimensional data

Shahina Rahman[a,1], Valen E. Johnson[a,1,*], Suhasini Subba Rao[a,1]

[a]*Department of Statistics, Texas A & M University*
[b]*3143 TAMU, College Station, TX 77843-3143*

Abstract

For high dimensional data, where P features for N objects ($P \gg N$) are represented in an $N \times P$ matrix $\mathbf{X}$, we describe a clustering algorithm based on the normalized left Gram matrix, $\mathbf{G} = \mathbf{X}\mathbf{X}^T/P$. Under certain regularity conditions, the rows in $\mathbf{G}$ that correspond to objects in the same cluster converge to the same mean vector. By clustering on the row means, the algorithm does not require preprocessing by dimension reduction or feature selection techniques, and does not require specification of tuning or hyperparameter values. Because it is based on the $N \times N$ matrix $\mathbf{G}$, it has lower computational cost than many methods based on clustering the feature matrix $\mathbf{X}$. When compared to 14 other clustering algorithms applied to 32 benchmarked microarray datasets, the proposed algorithm provided the most accurate estimate of the underlying cluster configuration more than twice as often as its closest competitors.

Keywords: Clustering, Gram Matrix, High-Dimensional Features, Unsupervised

1. Introduction

Despite their ubiquity in real applications, clustering of objects based on high dimensional features remains a challenging unsupervised learning task, made harder by the fact that the number of clusters is seldom known a priori. To detect clusters in high dimensional data, many clustering methods rely on preprocessing steps like feature selection or dimension reduction. Feature selection techniques are useful in finding clusters in sparse settings, where underlying clusters differ only by the values of a small number of features. However, the success of these techniques often depends on the astute selection of tuning parameters, which can be problematic in unsupervised

*Corresponding author
Email address: `vejohnson@exchange.tamu.edu` (Valen E. Johnson)

settings. Dimension reduction techniques, like PCA, SVD and matrix factorization, perform well

when underlying assumptions apply, but can fail to preserve cluster structure when they don't. [1].

Our goal in this article is to describe an algorithm to cluster N objects based on P features when $P \gg N$. When P features for N objects are represented in an $N \times P$ matrix $\mathbf{X}$, we base the algorithm on the $N \times N$ left Gram matrix, $\mathbf{X}\mathbf{X}^T$. Under certain regularity conditions, we show that this lower dimensional matrix preserves the cluster structure of the data. Standard clustering algorithms can then be applied to accurately identify the clusters at a reduced computational cost. The proposed algorithm does not require specification of tuning parameters or hyperparameters.

2. Notation and framework

We denote matrices by upper case bold letters (e.g., $\mathbf{X}$) and column vectors by lower case bold letters (e.g., $\mathbf{u}$). We use $\mathbb{1}_P$ to denote the P-dimensional vector of ones and $\mathbb{I}_A$ to denote the indicator function which equals 1 if A is true and 0 otherwise. We let $||\cdot||_2$ denote the Euclidean distance of a vector or Frobenius norm of a matrix, and $||\cdot||_1$ denote the absolute sum of the entries of a vector or matrix. We use $|\mathcal{S}|$ to denote the cardinality of the set $\mathcal{S}$. We write $i \simeq j$ if objects i and j are in the same cluster, and $i \not\simeq j$ otherwise.

Let K_0 be the (possibly unknown) number of true clusters. Let $\mathbf{f}_i = (x_{i,1}, \ldots, x_{i,P})$ denote the P-dimensional feature vector measured on the i^{th} object and stack the feature vectors of N objects into an $N \times P$ feature matrix $\mathbf{X}$. For a given K_0, define δ_i to be an integer in $\{1, \cdots, K_0\}$ that denotes the cluster membership of the i^{th} object. We denote $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_N)$ as the vector of cluster identifiers, where $Pr(\delta_i = a) = w_a$ with $\sum_{a=1}^{K_0} w_a = 1$. We assume that $\{\mathbf{f}_i\}_{i=1}^N$ are conditionally independent P dimensional random vectors with $\mathbb{E}(\mathbf{f}_i|\delta_i = a) = \boldsymbol{\mu}_a$, where $\boldsymbol{\mu}_a^{\text{T}} = (\mu_{a,1}, \ldots, \mu_{a,P})$ and finite covariance matrix $\text{Var}(\mathbf{f}_i|\delta_i = a) = \boldsymbol{\Sigma}_a$. We denote the vector of diagonal entries of $\boldsymbol{\Sigma}_a$ as $\boldsymbol{d}_a^{\text{T}} = (\sigma_1^a, \ldots, \sigma_P^a)$, and we define $\theta_{a,b} = \boldsymbol{\mu}_a^{\text{T}}\boldsymbol{\mu}_b/P$ for $a \neq b$ and $\theta_a = \boldsymbol{\mu}_a^{\text{T}}\boldsymbol{\mu}_a/P + \boldsymbol{d}_a^{\text{T}}\mathbb{1}_P/P$, where $1 \leq a, b \leq K_0$.

3. The clustering algorithm

The goal of our clustering algorithm is to partition N objects into K_0 clusters using a transformed left Gram matrix, rather than directly from the $\mathbf{X}$ matrix. We assume that K_0 is unknown. The Gram matrix or similarity matrix in general is a central quantity for many clustering algorithms. However, as far as we are aware the following transformation has not been previously exploited.

2

3.1. The left Gram matrix and it's properties

In the first step of the algorithm, we standardize each column of $\mathbf{X}$ to have mean 0 and standard deviation 1. Using the standardized $\mathbf{X}$, we construct the matrix $\mathbf{G}$ according to

$$\mathbf{G} = \mathbf{X}\mathbf{X}^{\mathrm{T}}/\mathbf{P}. \tag{1}$$

We denote the $(i,j)^{\mathrm{th}}$ entry of the $\mathbf{G}$ matrix by $g_{i,j}$.

If $\delta_i = a$ and $\delta_j = b$, $a, b \in 1, \ldots, K_0$, then the expectation of the entries of the $\mathbf{G}$ matrix is

$$\mathbb{E}(g_{i,j}|\delta_i = a, \delta_j = b) = \theta_{a,b} \qquad \text{and} \qquad \mathbb{E}(g_{i,i}|\delta_i = a) = \theta_a. \tag{2}$$

where, $\theta_{a,b} = \sum_{p=1}^{P} \mu_{a,p}\mu_{b,p}/P$ and $\theta_a = \sum_{p=1}^{P}(\mu_{a,p}^2 + \sigma_p^a)/P$. This pivotal property of the $\mathbf{G}$ matrix motivates us to detect the underlying clusters efficiently using the following two transformations.

3.2. The transformation on the left Gram matrix

Next, we modify the elements of $\mathbf{G}$ to form a matrix $\mathbf{M}$ by first appending the diagonal entries of $\mathbf{G}$ as an additional column to $\mathbf{M}$. In their place, we substitute the column-wise average from $\mathbf{G}$. The resulting $\mathbf{M} = \{m_{i,j}\}$ matrix is an $N \times (N+1)$ matrix with entries

$$m_{i,j} = \begin{cases} g_{i,j}, & \text{for } j \neq i = 1, \ldots, N \\ g_{i,i}, & \text{for } j = N+1 \\ \frac{1}{N-1}\sum_{\substack{j=1 \\ j \neq i}}^{N} g_{i,j}, & \text{for } j = i. \end{cases}$$

We denote the $(N+1)$ dimensional i^{th} row of $\mathbf{M}$ by $\mathbf{m}_i$.

To illustrate the above transformations, we consider an example with $N = 4$ objects where objects 1 and 2 belong to cluster 1, and objects 3 and 4 belong to cluster 2. The transformation from $\mathbf{G}$ to $\mathbf{M}$ is

$$\mathbf{G} = \begin{bmatrix} g_{1,1} & g_{1,2} & g_{1,3} & g_{1,4} \\ g_{2,1} & g_{2,2} & g_{2,3} & g_{2,4} \\ g_{3,1} & g_{3,2} & g_{3,3} & g_{3,4} \\ g_{4,1} & g_{4,2} & g_{4,3} & g_{4,4} \end{bmatrix} \Rightarrow \mathbf{M} = \begin{bmatrix} \frac{g_{2,1}+g_{3,1}+g_{4,1}}{3} & g_{1,2} & g_{1,3} & g_{1,4} & g_{1,1} \\ g_{2,1} & \frac{g_{1,2}+g_{3,2}+g_{4,2}}{3} & g_{2,3} & g_{2,4} & g_{2,2} \\ g_{3,1} & g_{3,2} & \frac{g_{1,3}+g_{2,3}+g_{4,3}}{3} & g_{3,4} & g_{3,3} \\ g_{4,1} & g_{4,2} & g_{4,3} & \frac{g_{1,4}+g_{2,4}+g_{3,4}}{3} & g_{4,4} \end{bmatrix}$$

The transformation from $\mathbf{G}$ to $\mathbf{M}$ matrix and equation 2 implies that the expected values of rows of $\mathbf{M}$ are equal for objects in the same cluster, with the exception of the diagonal elements. Hence $\mathbf{M}$ is useful in initializing our clustering algorithm when there is no prior information on the cluster identifiers, $\boldsymbol{\delta}$.

3.3. Updated modification of the Gram matrix

At the model selection stage, we propose further modification of $\mathbf{M}$ matrix to $\mathbf{M}^{\widehat{\boldsymbol{\delta}}}$ based on the knowledge of $\widehat{\boldsymbol{\delta}}$ obtained from clustering the rows of $\mathbf{M}$ matrix. With this modification, expectation of all the elements of the i^{th} and the j^{th} row of $\mathbf{M}^{\widehat{\boldsymbol{\delta}}}$ become equal for all the columns whenever the corresponding i^{th} and j^{th} objects belong to the same cluster. For a given estimate of cluster identifiers $\widehat{\boldsymbol{\delta}}$ from $\mathbf{M}$ matrix, we construct a matrix $\mathbf{M}^{\widehat{\boldsymbol{\delta}}}$ in which the diagonal elements of $\mathbf{M}^{\widehat{\boldsymbol{\delta}}}$ are updated with the column means restricted only to the rows of the remaining objects having the same cluster identifier , i.e.,

$$m_{i,i}^{\widehat{\boldsymbol{\delta}}} = \sum_{\substack{j=1\\j\neq i}}^{N} g_{j,i}.\mathbb{I}_{\{\widehat{\delta}_j=\widehat{\delta}_i\}} / \sum_{\substack{j=1\\j\neq i}}^{N} \mathbb{I}_{\{\widehat{\delta}_j=\widehat{\delta}_i\}}.$$

For the above example in 3.2 the transformation from $\mathbf{G}$ matrix to $\mathbf{M}^{\boldsymbol{\delta}}$ is

$$\mathbf{G} = \begin{bmatrix} g_{1,1} & g_{1,2} & g_{1,3} & g_{1,4} \\ g_{2,1} & g_{2,2} & g_{2,3} & g_{2,4} \\ g_{3,1} & g_{3,2} & g_{3,3} & g_{3,4} \\ g_{4,1} & g_{4,2} & g_{4,3} & g_{4,4} \end{bmatrix} \Rightarrow \mathbf{M}^{\boldsymbol{\delta}} = \begin{bmatrix} g_{2,1} & g_{1,2} & g_{1,3} & g_{1,4} & g_{1,1} \\ g_{2,1} & g_{1,2} & g_{2,3} & g_{2,4} & g_{2,2} \\ g_{3,1} & g_{3,2} & g_{4,3} & g_{3,4} & g_{3,3} \\ g_{4,1} & g_{4,2} & g_{4,3} & g_{3,4} & g_{4,4} \end{bmatrix}.$$

When $\widehat{\boldsymbol{\delta}} = \boldsymbol{\delta}$, the expected values of the rows of $\mathbf{M}^{\boldsymbol{\delta}}$ are described in the following lemma.

Lemma 3.1. *If $\delta_i = a$ (i.e., object i belongs to cluster a), then the expectation of the i^{th} row of $\mathbf{M}^{\boldsymbol{\delta}}$ is the following $(N+1)$ dimensional vector*

$$\mathbb{E}\left(\mathbf{m}_i^{\boldsymbol{\delta}}|\delta_i = a\right) = \boldsymbol{\theta}_a \equiv \left(\{\theta_{a,\delta_j}\}_{j=1}^{N} \quad ,\theta_a\right) \tag{3}$$

Let $\boldsymbol{\Theta}^{N\times N+1} = \mathbb{E}\left[\mathbf{M}^{\boldsymbol{\delta}}|\boldsymbol{\delta}\right]$, where $\boldsymbol{\Theta} = (\boldsymbol{\theta}_{\delta_1}, \cdots, \boldsymbol{\theta}_{\delta_N})^{\text{T}}$. Here $\boldsymbol{\Theta}$ represents the cluster means in the transformed space. The above transformation of $\mathbf{M}$ matrix makes the expected value of $\mathbf{m}_i^{\boldsymbol{\delta}}$ and $\mathbf{m}_j^{\boldsymbol{\delta}}$ equal whenever the i^{th} and j^{th} object belong to the same cluster. This facilitates various model selection criterion, including , the Bayesian information criterion (BIC) [2], to select the number of clusters by minimizing the discrepancy between $\mathbf{M}^{\boldsymbol{\delta}}$ and the assumed model.

Next, we assume the following condition holds on the distinct cluster means on the transformed space.

Assumption 3.1. *For $a \neq b \in \{1,\ldots,K_0\}$, there exists $\eta > 0$ such that for all $P > 0$*

$$||\boldsymbol{\theta}_a - \boldsymbol{\theta}_b||_2 > \eta. \tag{4}$$

4

Such a condition is required to ensure that distinct clusters can be identified. It requires that the proportion of informative feature vectors measured on objects does not converge to 0 as P grows.

3.4. Clustering strategy based on the rows of $\mathbf{M}$

To estimate the underlying number of clusters K_0 and the cluster indicator vector $\boldsymbol{\delta}$ we maximize a quasi mixture likelihood for each possible value of $K = 1, \ldots, K_{max}$. Here, K_{max} is an user-defined upper bound on K_0, which may equal N if a prior bound is not known. Because the elements of $\mathbf{M}$ and $\mathbf{M}^{\boldsymbol{\delta}}$ represent an average of P pairwise products, under certain regularity conditions the rows of $\mathbf{M}$ and $\mathbf{M}^{\boldsymbol{\delta}}$ converge to a multivariate normal distribution. We therefore maximize a quasi mixture likelihood function of the form $L_K(\mathbf{M}) = \log \prod_{i=1}^{N} \sum_{k=1}^{K} w_k \phi\left(\mathbf{m}_i; \boldsymbol{\theta}_k, \boldsymbol{\Gamma}_k\right)$, where

$$\phi\left(\mathbf{m}_i; \boldsymbol{\theta}_k, \boldsymbol{\Gamma}_k\right) = \frac{1}{(2\pi)^{N/2}\det(\boldsymbol{\Gamma}_k)^{1/2}} \times \exp\left\{-\frac{1}{2}\left(\mathbf{m}_i - \boldsymbol{\theta}_k\right)^{\mathrm{T}} \boldsymbol{\Gamma}_k^{-1}\left(\mathbf{m}_i - \boldsymbol{\theta}_k\right)\right\}.$$

Here the mixing weights w_k satisfy $w_k \geq 0$ and $\sum_{k=1}^{K} w_k = 1$. Let $\widehat{L}_K(\mathbf{M}) = L_K(\mathbf{M}; \widehat{w}_K, \widehat{\boldsymbol{\theta}}_K, \widehat{\boldsymbol{\Gamma}}_K)$ denote the maximized quasi log-likelihood for an assumed value of K. The EM algorithm [3] is used to obtain maximum likelihood estimates for the parameters of the mixture model and the latent cluster identifiers $\widehat{\boldsymbol{\delta}}$ given K.

We estimate the number of clusters by maximizing the Bayesian Information Criterion (BIC) [2], which can be expressed as

$$\mathrm{BIC}_K = 2\widehat{L}_K(\mathbf{M}^{\widehat{\boldsymbol{\delta}}}) - \nu_K \log N,$$

where ν_K is the number of estimated parameters in the likelihood function $L_K(\cdot)$.

3.5. Implementation details

The EM algorithm is only guaranteed to arrive at a local optimum of the mixture likelihood criterion [4]. As a consequence, the choice of starting values for $\boldsymbol{\delta}$ is important. Our experience and previous studies on multivariate mixture models by [5], [6], [7] suggests that initial estimates of $\boldsymbol{\delta}$ obtained from agglomerative hierarchical cluster analysis on the rows of $\mathbf{M}$ for a given K generally provide effective initialization. Pseudo code for implementing the clustering algorithm based on the modified Gram matrix, $\mathbf{M}$ and $\mathbf{M}^{\boldsymbol{\delta}}$ is described in Algorithm 1.

Algorithm 1: GMCluster

Input: $\mathbf{X} \in \mathbb{R}^{N \times P}$ and $K_{\max}$.

Output: $\widehat{K}$; Cluster identifier, $\widehat{\boldsymbol{\delta}} = (\widehat{\delta}_1, \ldots, \widehat{\delta}_N)$.

1. **G-step:** Construct the $N \times N$ similarity matrix, $\mathbf{G} = \mathbf{X}\mathbf{X}^{\mathbf{T}}/\mathbf{P}$.

2. **M-step:** Rearrange the diagonal elements of $\mathbf{G}$ to construct the $N \times (N+1)$ matrix $\mathbf{M}$.

3. **for** $K = 1$ **to** $K_{\max}$:

    a. **Initialize:** Use agglomerative hierarchical clustering on $\mathbf{M}$ to initialize the cluster identifier $\widehat{\boldsymbol{\delta}}$ having K clusters.

    b. **Repeat**

        i. $\boxed{\textbf{Mstep}:}$ **for** $k = 1$ **to** K :

        $n_k = \sum_{i=1}^{N} \mathbb{I}_{\{\widehat{\delta}_i = k\}}; \quad \widehat{w}_k = n_k/N; \quad \widehat{\boldsymbol{\theta}}_k = \sum_{i=1}^{N} \mathbf{m}_i \cdot \mathbb{I}_{\{\widehat{\delta}_i = k\}}/n_k;$

        $\widehat{\boldsymbol{\Gamma}}_k = \sum_{i=1}^{N} \left(\mathbf{m}_i - \widehat{\boldsymbol{\theta}}_k\right)\left(\mathbf{m}_i - \widehat{\boldsymbol{\theta}}_k\right)^{\mathbf{T}} \cdot \mathbb{I}_{\{\widehat{\delta}_i = k\}}/n_k.$

        ii. $\boxed{\textbf{Estep}:}$ **for** $i = 1$ **to** N: Set $\widehat{\delta}_i = k$ for $\arg\max_{k=1}^{K} \widehat{w}_k \phi\left(\mathbf{m}_i; \widehat{\boldsymbol{\theta}}_k, \widehat{\boldsymbol{\Gamma}}_k\right) = k.$

        until until $\widehat{\boldsymbol{\delta}}$ doesn't change.

    c. Given $\widehat{\boldsymbol{\delta}}$, calculate BIC_K based on $\mathbf{M}^{\widehat{\delta}}$.

4. **Return** K and $\widehat{\boldsymbol{\delta}}$ that maximize BIC_K

---

4. Convergence of the $\mathbf{M}^{\delta}$ matrix

In the following lemma we obtain a rate of convergence for the mean squared error difference between $\mathbf{M}^{\widehat{\delta}}$ and $\boldsymbol{\Theta}$. We require the following assumptions on the features. Recall that $\boldsymbol{\Sigma}_a = \mathrm{Var}(\boldsymbol{x}_i|\delta_i = a)$, and let $\tau_P = \sup_{1 \leq k \leq K_0} \|\boldsymbol{\Sigma}_k\|_1^{1/2}$. For each feature $p = 1, \ldots, P$, let $y_{i,p} = x_{i,p} - \mu_{\delta_i, p}$ and define the P-dimensional vector $\mathbf{z}_i^{\mathrm{T}} = \left(y_{i,1}^2, \ldots, y_{i,P}^2\right)$. Let $\boldsymbol{\Upsilon}_a = \mathrm{Var}(\mathbf{z}_i|\delta_i = a)$ and $\kappa_P = \sup_{1 \leq k \leq K_0} \|\boldsymbol{\Upsilon}_k\|_1^{1/2}$.

Assumption 4.1. *As $P \to \infty$, we assume that $\kappa_P/P \to 0$ and $\tau_P/P \to 0$.*

We call the features weakly dependent if κ_P and τ_P are $O(P^{1/2})$. This condition holds if the feature vectors are independent and have bounded variance.

Lemma 4.1. *Suppose Assumption 2 holds. Let $\mu_{\sup} = \sup_{a,p} |\mu_{a,p}|$ and $\sigma_{\sup} = \sup_{a,p} \sqrt{\sigma_p^a}$ and let $\mu_{\sup}, \sigma_{\sup} < \infty$. Then*

$$\mathbb{E} \left\|\mathbf{M}^{\delta} - \boldsymbol{\Theta}\right\|_2^2 \leq \Delta_P^2$$

6

where $\Delta_P = \frac{1}{P} \left[N \{ (N-1)\tau_P^2 (2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + (\kappa_P + 2\tau_P \mu_{\mathrm{sup}})^2 \} \right]^{1/2}$.

The proof is in the supplementary material.

This lemma suggests that for a correctly specified cluster configuration in which $\widehat{\boldsymbol{\delta}} = \boldsymbol{\delta}$, $\mathbf{M}^{\boldsymbol{\delta}}$ converges to $\boldsymbol{\Theta}$, the transformed cluster mean, in the order of $O(P^{-1/2})$, if the features are weakly dependent. For $K \leq K_0$, the identifiability condition (4)) guarantees that this sum-of-squares is bounded away from 0. For $K > K_0$, the BIC penalty is sufficiently large to prevent sub-clusters from a given cluster from forming since the decrease in the sum-of-squares accumulated from such a split cannot offset a fixed penalty that is greater than $\log(N)$. Thus, standard clustering algorithms are likely to be able to identify the correct cluster identifiers, provided that the conditions stated above are satisfied and P is sufficiently large.

5. Application of the proposed method to bench-marked gene data sets

Shah and Koltun [8] provided a recent comparison of several popular clustering methods for 32 gene expression data sets based on adjusted mutual information (AMI) [9]. The cluster configurations of these data sets are well studied, validated and are available in DataLink [10].

To evaluate our method, GMcluster (GMC), we compared its performance to the clustering methods considered in [8], adding 7 additional state-of-the-art algorithms to the comparison. Overall, we compared GMC with 12 clustering algorithms that require pre-specification of the number of clusters and 6 clustering algorithms that do not. As in [8], for the methods that required pre-specification of the number of clusters we used those given in [11]. Details of the methods considered and their parameter settings are provided in the supplementary section.

Figure 1 summarizes the results of these comparisons. (For brevity, we have displayed the AMI results of only the 15 algorithms that yielded the best AMI value for at least 1 out of 32 data sets in Table B.2 of the appendix.) The barplot depicts the number of times each method yielded the highest accuracy based on AMI. As the figure shows, GMC achieved the highest AMI for 11 data sets, while the next best algorithm, RCC, provided the highest AMI in 5 data sets. AMI values acheived by each clustering algorithm for each data set are provided in the Appendix.

We also compared the computational performance of GMC to the six algorithms that, like GMC, also estimated the number of clusters. Because the runtime complexity of the GMC algorithm is linear in P, it was generally faster than the other algorithms (Figure 1b). Indeed, except for the

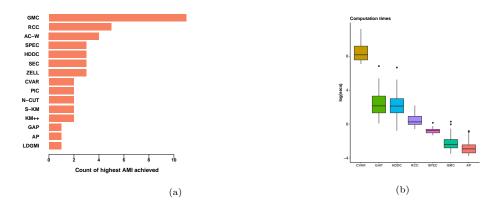(a)   (b)

Figure 1: (a) Barplot depicts the frequency of obtaining the highest accuracy (AMI) by each of the 14 clustering algorithms across 32 gene expression data sets. (b) Boxplots display the distribution of computation times for the algorithms that do not require any predefined number of clusters (K_0).

AP algorithm (which only achieved the best AMI in one test data set), then GMC was faster by a factor of at least 7 than all other algorithms that we tested. All comparisons were performed on a workstation with an Intel(R) Core(TM) i7-3770 CPU clocked at 3.40GHz with 8.00 GB RAM. The data sets and algorithms that produced these results are available at Github Link.

6. Discussion

Our clustering algorithm offers a simple and efficient method for clustering objects based on high-dimensional feature vectors. A software implementation of our proposed algorithm GMcluster is available in the R package, *RJcluster* [12] [1]. The overall complexity of the current algorithm is $O(N^2P)$. Application to several genomic data sets suggest that the proposed algorithm provides a useful method for clustering objects when $N \ll P$. In ongoing work, we are developing scalable techniques that will facilitate the use of our proposed algorithm in large N and large P settings. In future work, we will explore the utility of the proposed transformation on other similarity measures for meaningful clustering of unstructured data, like images or text documents.

---

[1] Our proposed algorithm has been implemented as RJcluster in CRAN as a R package

References

References

[1] W.-C. Chang, On using principal components before separating a mixture of two multivariate normal distributions, Journal of the Royal Statistical Society: Series C (Applied Statistics) 32 (3) (1983) 267–275.

[2] G. Schwarz, Estimating the dimension of a model, Ann. Statist. 6 (2) (1978) 461–464.

[3] A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the em algorithm, Journal of the royal statistical society. Series B (methodological) (1977) 1–38.

[4] G. McLachlan, D. Peel, Finite mixture models, John Wiley & Sons, 2004.

[5] L. Scrucca, M. Fop, T. B. Murphy, A. E. Raftery, mclust 5: clustering, classification and density estimation using Gaussian finite mixture models, The R Journal 8 (1) (2016) 289–317. URL https://doi.org/10.32614/RJ-2016-021

[6] C. Fraley, A. E. Raftery, Model-based clustering, discriminant analysis, and density estimation, Journal of the American statistical Association 97 (458) (2002) 611–631.

[7] A. E. Raftery, N. Dean, Variable selection for model-based clustering, Journal of the American Statistical Association 101 (473) (2006) 168–178.

[8] S. A. Shah, V. Koltun, Robust continuous clustering, Proceedings of the National Academy of Sciences 114 (37) (2017) 9814–9819.

[9] T. M. Cover, J. A. Thomas, Entropy, relative entropy and mutual information, Elements of information theory 2 (1) (1991) 12–13.

9

[10] M. de Souto, I. Costa, D. S. de Araujo, T. B. Ludermir, A. Schliep, Clustering cancer gene expression data: a comparative study, BMC Bioinformatics 9 (1) (2008) 497.

[11] M. De Souto, I. Costa, D. de Araujo, T. Ludermir, a schliep, a.(2008),'clustering cancer gene expression data: a comparative study', BMC Bioinformatics 9.

[12] S. Rahman, V. E. Johnson, S. Subba Rao, R. Shudde, RJcluster: A Fast Clustering Algorithm for High Dimensional Data Based on the Gram Matrix Decomposition, r package version 3.2.4 (2021).
URL https://CRAN.R-project.org/package=RJcluster

[13] D. Arthur, S. Vassilvitskii, k-means++: The advantages of careful seeding, in: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, 2007, pp. 1027–1035.

[14] D. Comaniciu, P. Meer, Mean shift: A robust approach toward feature space analysis, IEEE Transactions on pattern analysis and machine intelligence 24 (5) (2002) 603–619.

[15] J. Shi, J. Malik, Normalized cuts and image segmentation, IEEE Transactions on pattern analysis and machine intelligence 22 (8) (2000) 888–905.

[16] D. Zhao, X. Tang, Cyclizing clusters via zeta function of a graph, in: Advances in Neural Information Processing Systems, 2009, pp. 1953–1960.

[17] F. Nie, D. Xu, I. W. Tsang, C. Zhang, Spectral embedded clustering., in: IJCAI, 2009, pp. 1181–1186.

[18] Y. Yang, D. Xu, F. Nie, S. Yan, Y. Zhuang, Image clustering using local discriminant models and global integration, IEEE Transactions on Image Processing 19 (10) (2010) 2761–2773.

[19] W. Zhang, D. Zhao, X. Wang, Agglomerative clustering via maximum incremental path integral, Pattern Recognition 46 (11) (2013) 3056–3065.

[20] D. M. Witten, R. Tibshirani, A framework for feature selection in clustering, Journal of the American Statistical Association 105 (490) (2010) 713–726.

[21] E. Elhamifar, R. Vidal, Sparse subspace clustering: Algorithm, theory, and applications, IEEE transactions on pattern analysis and machine intelligence 35 (11) (2013) 2765–2781.

[22] B. J. Frey, D. Dueck, Clustering by passing messages between data points, science 315 (5814) (2007) 972–976.

[23] R. Tibshirani, G. Walther, T. Hastie, Estimating the number of clusters in a data set via the gap statistic, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63 (2) (2001) 411–423.

[24] L. Kaufman, P. J. Rousseeuw, Partitioning around medoids (program pam), Finding groups in data: an introduction to cluster analysis 344 (1990) 68–125.

[25] C. Bouveyron, S. Girard, C. Schmid, High-dimensional data clustering, Computational statistics & data analysis 52 (1) (2007) 502–519.

[26] C. R. John, D. Watson, Spectrum: Fast Adaptive Spectral Clustering for Single and Multi-View Data, r package version 1.1 (2020).
URL https://CRAN.R-project.org/package=Spectrum

Appendix A. Proof of Lemma 4.1

The proof of Lemma 4.1 follows from the proof of the following lemmas.

Lemma Appendix A.1. Lemma A.1 *Recall that $g_{i,j} = \sum_{p=1}^{P} x_{i,p} x_{j,p}/P$ with $\mathbb{E}(x_{i,p}|\delta_i) = \mu_{\delta_i,p}$ and $\mathbb{E}(\mathbf{f}_i|\delta_i) = \boldsymbol{\mu}_{\delta_i}$ and $Var(\mathbf{f}_i|\delta_i) = \boldsymbol{\Sigma}_{\delta_i}$. Also recall that for clusters $a, b \in \{1, \ldots, K_0\}$, in equation 2 we defined $\theta_{a,b} = \boldsymbol{\mu}_a^{\mathrm{T}} \boldsymbol{\mu}_b/P$ and $\theta_a = \boldsymbol{\mu}_a^{\mathrm{T}} \boldsymbol{\mu}_a/P + \boldsymbol{d}_a^{\mathrm{T}} \mathbb{1}_P/P$, where $\boldsymbol{d}_a^{\mathrm{T}} = (\sigma_1^a, \ldots, \sigma_P^a)$ are the diagonal entries of $\boldsymbol{\Sigma}_a$. Then for $a, b \in \{1, \ldots, K_0\}$, $\mathbb{E}[g_{i,j}|\delta_i = a, \delta_j = b] = \theta_{a,b}$ and $\mathbb{E}[g_{i,i}|\delta_i] = \theta_a$.*

Proof:

Suppose Assumption 4.1 holds. We define a P-dimensional vector $\mathbf{y}_i = (y_{i,1}, \ldots, y_{i,P})$ with components $y_{i,p} = x_{i,p} - \mu_{\delta_i,p}$. It follows that $\mathbb{E}[y_{i,p}|\delta_i = a] = 0$ and $\mathbb{E}[y_{i,p}^2|\delta_i = a] = \sigma_p^a$. For $a \neq b \in \{1, \ldots, K_0\}$, If $i \neq j$, $\delta_i = a$, and $\delta_j = b$, then

$$g_{i,j} - \theta_{a,b} = \frac{1}{P} \sum_{p=1}^{P} y_{i,p} y_{j,p} + \frac{1}{P} \sum_{p=1}^{P} \mu_{a,p} y_{j,p} + \frac{1}{P} \sum_{p=1}^{P} \mu_{b,p} y_{i,p},$$

and for $i = j$

$$g_{i,i} - \theta_a = \frac{1}{P} \sum_{p=1}^{P} (y_{i,p}^2 - \sigma_p^a) + \frac{2}{P} \sum_{p=1}^{P} \mu_{a,p} y_{i,p}.$$

11

Using the above and evaluating the conditional expectations $\mathbb{E}[g_{i,j}|\delta_i, \delta_j]$ and $\mathbb{E}[g_{i,i}|\delta_i]$ proves

$$\mathbb{E}(g_{i,j}|\delta_i = a, \delta_j = b) = \theta_{a,b} \text{ and } \mathbb{E}(g_{i,i}|\delta_i = a) = \theta_a. \tag{A.1}$$

Lemma Appendix A.2. Lemma A.2 *Suppose Assumption 2 holds. Recall in section 4 for the P-dimensional vector* $\mathbf{y}_i^{\mathrm{T}}$*, we defined* $\mathbf{\Upsilon}_k = Var(\mathbf{y}_i|\delta_i = k)$ *and* $\kappa_P = \sup_{1 \le k \le K_0} \|\mathbf{\Upsilon}_k\|_1^{1/2}$ *and* $\tau_P = \sup_{1 \le k \le K_0} \|\mathbf{\Sigma}_k\|_1^{1/2}$*. Let* $\mu_{\mathrm{sup}} = \sup_{k,p} |\mu_{k,p}|$ *and* $\sigma_{\mathrm{sup}} = \sup_{k,p} \sqrt{\sigma_p^k}$*. Then for* $i \ne j$

$$\left(\mathbb{E}|g_{i,j} - \theta_{\delta_i,\delta_j}|^2\right)^{1/2} \le \frac{\tau_P}{P}(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}}) \tag{A.2}$$

and

$$\left(\mathbb{E}|g_{i,i} - \theta_{\delta_i}|^2\right)^{1/2} \le \frac{1}{P}(\kappa_P + 2\tau_P\mu_{\mathrm{sup}}). \tag{A.3}$$

Furthermore,

$$\mathbb{E}\left\|\mathbf{m}_i^{\boldsymbol{\delta}} - \boldsymbol{\theta}_{\delta_i}\right\|_2^2 \le \frac{1}{P^2}\left[(N-1)\tau_P^2(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + (\kappa_P + 2\tau_P\mu_{\mathrm{sup}})^2\right], \tag{A.4}$$

and

$$\mathbb{E}\left\|\mathbf{M}^{\boldsymbol{\delta}} - \boldsymbol{\Theta}\right\|_2^2 \le \frac{N}{P^2}\left[(N-1)\tau_P^2(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + (\kappa_P + 2\tau_P\mu_{\mathrm{sup}})^2\right]. \tag{A.5}$$

Proof:

We first prove (A.2), for the case $i \ne j$. We use (A.1) to give the bound

$$(\mathbb{E}|g_{i,j} - \theta_{\delta_i,\delta_j}|^2)^{1/2} \le \left[\mathbb{E}\left(\frac{1}{P}\sum_{p=1}^P y_{i,p}y_{j,p}\right)^2\right]^{1/2} + \left[\mathbb{E}\left(\frac{1}{P}\sum_{p=1}^P \mu_{\delta_i,p}y_{j,p}\right)^2\right]^{1/2} +$$
$$\left[\mathbb{E}\left(\frac{1}{P}\sum_{p=1}^P \mu_{\delta_j,p}y_{i,p}\right)^2\right]^{1/2}.$$

Recall that $\mathbb{E}(A^2) = \mathbb{E}[\mathbb{E}(A^2|\delta)]$ and if $\mathbb{E}[A|\delta] = 0$, then $\mathbb{E}(A^2) = \mathbb{E}[\mathrm{var}(A|\delta)]$. This implies

$$(\mathbb{E}|g_{i,j} - \theta_{\delta_i\delta_j}|^2)^{1/2} \le A_{1,P} + A_{2,P} + A_{3,P}, \tag{A.6}$$

12

where

$$A_{1,P} = \left(\mathbb{E} \left[\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} y_{i,p} y_{j,p} | \delta_i, \delta_j \right) \right] \right)^{1/2}$$

$$A_{2,P} = \left(\mathbb{E} \left[\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} \mu_{\delta_j,p} y_{i,p} | \delta_i, \delta_j \right) \right] \right)^{1/2}$$

$$\text{and} \quad A_{3,P} = \left(\mathbb{E} \left[\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} \mu_{\delta_i,p} y_{j,p} | \delta_i, \delta_j \right) \right] \right)^{1/2}.$$

We now bound each of the terms $A_{1,P}, A_{2,P}$ and $A_{3,P}$. To bound $A_{1,P}$ we use the following decomposition:

$$\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} y_{i,p} y_{j,p} | \delta_i = a, \delta_j = b \right)$$

$$= \frac{1}{P^2} \sum_{p_1,p_2=1}^{P} \operatorname{cov} \left(y_{i,p_1} y_{j,p_1}, y_{i,p_2} y_{j,p_2} | \delta_i = a, \delta_j = b \right)$$

$$= \frac{1}{P^2} \sum_{p_1,p_2=1}^{P} \operatorname{cov}[y_{i,p_1}, y_{i,p_2} | \delta_i = a] \operatorname{cov} \left(y_{j,p_1}, y_{j,p_2} | \delta_j = b \right)$$

$$\leq \sup_{a,p} \sigma_{a,p} \frac{1}{P^2} \sup_{a} \sum_{p_1,p_2=1}^{P} |\operatorname{cov} \left(y_{i,p_1}, y_{i,p_2} | \delta_i = a \right)| \leq \frac{\tau_P^2}{P^2} \sigma_{\sup}^2.$$

It follows that

$$A_{1,P} \leq \left(\mathbb{E} \left[\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} y_{i,p} y_{j,p} | \delta_i, \delta_j \right) \right] \right)^{1/2} \leq \frac{\tau_P}{P} \sigma_{\sup}.$$

Using a similar argument to bound the conditional variance inside $A_{2,P}$, we have

$$\operatorname{var} \left(\frac{1}{P} \sum_{p=1}^{P} \mu_{i,p} y_{j,p} | \delta_i = a, \delta_j = b \right) \leq \mu_{\sup}^2 \frac{1}{P^2} \sum_{p_1,p_2=1}^{P} |\operatorname{cov}(y_{i,p_1}, y_{i,p_2})| \leq \frac{1}{P^2} \mu_{\sup}^2 \tau_P^2.$$

This leads to

$$A_{2,P} \leq \frac{\tau_P}{P} \mu_{\sup},$$

and by a similar argument to $A_{3,P} \leq \frac{\tau_P}{P} \mu_{\sup}$. Substituting these bounds into (A.6) we obtain

$$(\mathbb{E}|g_{i,j} - \theta_{\delta_i,\delta_j}|^2)^{1/2} \leq \frac{\tau_P}{P} (2\mu_{\sup} + \sigma_{\sup}),$$

13

thus proving (A.2). We next bound $(\mathbb{E}|g_{i,i} - \theta_{\delta_i}|^2)^{1/2}$. We use (A.1) to give

$$(\mathbb{E}|g_{i,i} - \theta_{\delta_i}|^2)^{1/2} \quad \leq \quad B_{1,P} + B_{2,P} + B_{3,P}, \tag{A.7}$$

where

$$B_{1,P} \quad = \quad \left(\mathbb{E}\left[\mathrm{var}\left(\frac{1}{P} \sum_{p=1}^{P} y_{i,p}^2 | \delta_i \right) \right] \right)^{1/2}$$

$$B_{2,P} \quad = \quad \left(\mathbb{E}\left[\mathrm{var}\left(\frac{1}{P} \sum_{p=1}^{P} \mu_{\delta_i,p} y_{i,p} | \delta_i \right) \right] \right)^{1/2}$$

$$\text{and} \quad B_{3,P} \quad = \quad \left(\mathbb{E}\left[\mathrm{var}\left(\frac{1}{P} \sum_{p=1}^{P} \mu_{\delta_i,p} y_{i,p} | \delta_i \right) \right] \right)^{1/2}.$$

Using the same methods used to bound $A_{2,P}$ and $A_{3,P}$, it is straightforward to show that $B_{2,P}, B_{3,P} \leq \tau_P \mu_{\mathrm{sup}}/P$. To bound $B_{1,P}$ we note that

$$\mathrm{var}\left(\frac{1}{P} \sum_{p=1}^{P} y_{i,p}^2 | \delta_i = a \right) \quad = \quad \frac{1}{P^2} \sum_{p_1,p_2=1}^{P} \mathrm{cov}[y_{i,p_1}^2, y_{i,p_2}^2 | \delta_i = a] \leq P^{-2}\kappa_P^2,$$

which follows from Assumption 4.1. Thus $B_{1,P} \leq P^{-1}\kappa_P$. Substituting into (A.7) gives

$$\left(\mathbb{E}|g_{i,i} - \theta_{\delta_i})^2| \right)^{1/2} \quad \leq \quad \frac{1}{P}(\kappa_P + 2\tau_P\mu_{\mathrm{sup}}),$$

thus proving (A.3).

To prove (A.4), we apply (A.2) and (A.3), leading to

$$\mathbb{E}\left\| \mathbf{m}_i^{\boldsymbol{\delta}} - \boldsymbol{\theta}_{\delta_i} \right\|_2^2 \quad = \quad \sum_{i,j=1,i\neq j}^{N} \mathbb{E}(g_{i,j} - \theta_{\delta_i,\delta_j})^2 + \sum_{i=1}^{N} \mathbb{E}(g_{i,i} - \theta_{\delta_i})^2$$

$$\leq \quad (N-1)\frac{1}{P^2}\tau_P^2(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + \frac{1}{P^2}(\kappa_P + 2\tau_P\mu_{\mathrm{sup}})^2$$

$$\leq \quad \frac{1}{P^2}\left[(N-1)\tau_P^2(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + (\kappa_P + 2\tau_P\mu_{\mathrm{sup}})^2 \right],$$

proving (A.4). Finally, to prove (A.5) we note that

$$\mathbb{E}\left\| \mathbf{M}^{\boldsymbol{\delta}} - \boldsymbol{\Theta} \right\|_2^2 = \sum_{i=1}^{N} \mathbb{E}\left\| \mathbf{m}_i^{\boldsymbol{\delta}} - \boldsymbol{\theta}_{\delta_i} \right\|_2^2.$$

By substituting (A.4) above we have

$$\mathbb{E}\left\| \mathbf{M}^{\boldsymbol{\delta}} - \boldsymbol{\Theta} \right\|_2^2 \leq \frac{N}{P^2}\left[(N-1)\tau_P^2(2\mu_{\mathrm{sup}} + \sigma_{\mathrm{sup}})^2 + (\kappa_P + 2\tau_P\mu_{\mathrm{sup}})^2 \right],$$

which proves (A.5) and hence Lemma 4.1. $\qquad\square$

14

Appendix B. Description of other clustering algorithms

We divided the competing methods into two categories. We consider 12 algorithms that *require* the pre-specification of the number of clusters K_0 : k-means++ (**km++**) [13], Gaussian mixture models (**GMM**), fuzzy clustering (**fuzzy**), mean-shift clustering (**MS**) [14], agglomerative hierarchical clustering with ward linkage (**AC-W**), normalized cuts (**N-Cuts**) [15], Zeta l-links (**Zell**) [16], spectral embedded clustering (**SEC**) [17], clustering using local discriminant models and global integration (**LDMGI**) [18], path integral clustering (**PIC**) [19], sparse k-means (**sp-km**) [20] and sparse subspace clustering (**SSC**) [21]. We consider 6 other clustering methods that *do not require* the pre-specification of K_0 : affinity propagation (**AP**) [22], a robust graph continuous clustering (**RCC**) [8], GAP statistics with partitioning around medoids (**GAP**) ([23], [24]), a model based clustering with variable selection (**Cvarsl**) [7], a high dimensional data clustering (**HDDC**) [25] and a graph clustering based on tensors (**SPEC**) [26]. The hyperparameter settings used for these methods are provided in Table B.1.

Figure 1 provides a summary of the number of times each algorithm achieved the highest AMI. This figure does not display results for **GMM**, **MS**, **fuzzy**, and **SSC** algorithms because none of these algorithms provided the highest AMI for any data set.

Computation times

We compared the execution time of our proposed method, GMcluster algorithm to other algorithms which *do not require* the prespecification of the number of clusters. The overall distribution of the computation times taken by these algorithms across 32 datasets is displayed in Figure 1b. Execution times are displayed in log seconds. The **AP**, **HDDC**, **RCC** and **SPEC** algorithms are much more computationally efficient than **GAP** and **Cvarsl** methods. Excluding the AP algorithm, which provided the best AMI in only one case, the GMcluster algorithm was atleast 7 times faster than all of the remaining algorithms.

Hyperparameter settings of other clustering methods

We used the same hyperparameters settings recommended in [8] for the following clustering algorithms: **KM++**, **AC-W**, **N-CUT**, **ZELL**, **SEC**, **LDMGI**, **PIC** and **RCC** that. Table B.1 provides the hyperparameter settings and software used to obtain the results from the other methods.

Table B.1: Hyperparameters and software used for other methods in the comparative study.

Methods	hyper-parameters	Values	Software
AP	iter.max	100	apcluster
	s	negDistMat(r=2)	(R package)
CVAR	search	headlong	clustvarsel
	direction	forward	(R package)
	parallel	T	
	iter.max	100	
GAP	maximum number of clusters	20	cluster
(with partitioning around mediod)	d.power	2	(R package)
	bootstrap samples	max(100,n)	
	metric	Euclidean	cluster
	iter.max	100	
HDDC	max number of clusters	20	HDclassif
	model	"ALL"	(R package)
	threshold	0.2(default)	
	criterion	bic(default)	
	$d_{\max}$	100(default)	
S-KM	iter.max	100	sparcl
	wbounds grid	[1, 10](default)	(R package)
	nperm	100	
SPEC	method	2(default: multimodal eigen gap)	Spectrum
	kernel-type	density(default)	(R package)
	maxk	20	
	Nearest-Neighbor	7(default)	
	iter.max	100	
GMC	Cmax	20	RJcluster
	iter.max	100	

Data transformations

We took the logarithmic transformation of all data that contained only positive values. Several data sets were already preprocessed and centered and were therefore not log-transformed further. These data sets included Alizadeh-v1,v2,v3, Bittner, Garber, Lapointe-v1, Liang, Risinger, Singh-v1, Tomlins-v1 and West. For the remaining data sets, after logarithm transformation we standardized by centering on the median and scaling by the standard deviation. Further details regarding the transformations applied to each data set can be found in the "scaling.R" folder in the RJclust folder provided in the github repository https://github.com/srahman-24/GMclust.

Data Availability and Software

The datasets used for the comparisons are available at <u>DataLink</u>. We executed all algorithms on a workstation with an Intel(R) Core(TM) i7-3770 CPU clocked at 3.40GHz with 8.00 GB RAM. The datasets and algorithms that produced the results are available at <u>https://github.com/srahman-24</u>.

AMI calculations

17

Table B.2: Adjusted Mutual Information for 15 clustering algorithms over 32 gene expression datasets. For each dataset, the maximum achieved AMI is highlighted in bold.

Datasets	N	P	K_0 is **known**								K_0 is **unknown**						
			km++	sp-km	AC-W	N-Cuts	Zell	SEC	LDGMI	PIC	AP	GAP	HDDC	SPEC	RCC	Cvarsl	GMC
Alizadeh-v1	42	1097	0.340	-0.015	0.101	0.096	0.250	0.238	0.123	0.033	0.211	0.000	0.133	0.157	0.426	-0.006	**0.515**
Alizadeh-v2	62	2095	0.568	0.872	0.922	0.922	0.922	0.922	0.738	0.922	0.563	**1.000**	0.571	0.753	**1.000**	0.533	**1.000**
Alizadeh-v3	62	2095	0.586	0.689	0.616	0.601	0.702	0.574	0.582	0.625	0.540	0.678	0.548	0.609	**0.792**	0.295	**0.792**
Armstrong-v1	72	1083	0.372	0.370	0.308	0.372	0.308	0.323	0.355	0.308	0.381	0.475	0.461	0.617	0.546	0.302	**0.637**
Armstrong-v2	72	2196	**0.891**	0.375	0.746	0.83	0.802	**0.891**	0.509	0.802	0.586	0.525	0.000	0.693	0.838	0.513	0.661
Bhattacharjee	203	1545	0.444	0.296	**0.601**	0.563	0.496	0.570	0.378	0.378	0.377	0.518	0.000	0.505	0.600	0.173	0.453
Bittner	38	2203	-0.012	0.195	0.002	0.042	0.115	-0.002	0.014	0.115	0.243	0.000	0.288	0.013	0.156	-0.020	**0.341**
Bredel	28	1072	0.297	0.000	0.384	0.203	0.278	0.259	0.295	0.278	0.139	0.035	0.227	0.356	**0.466**	-0.002	0.265
Chowdary	104	184	0.764	0.595	**0.859**	**0.859**	**0.859**	**0.859**	**0.859**	**0.859**	0.443	0.000	0.625	0.575	0.393	0.000	0.585
Dyrskjot	40	1205	0.507	**0.755**	0.474	0.303	0.269	0.389	0.385	0.177	0.558	0.348	0.607	0.629	0.383	0.292	0.742
Garber	66	4555	0.242	0.026	0.210	0.204	0.246	0.200	0.191	0.246	**0.274**	0.096	0.164	0.137	0.173	0.175	0.130
Golub-v1	72	1870	0.688	0.701	**0.831**	0.650	0.615	0.615	0.615	0.615	0.430	0.044	0.478	0.137	0.490	0.628	0.420
Golub-v2	72	1870	0.680	0.617	**0.737**	0.693	0.689	0.703	0.600	0.689	0.516	0.000	0.478	0.352	0.597	0.139	0.538
Gordon	181	1628	0.651	**0.937**	0.483	0.681	-0.005	0.791	0.669	0.664	0.304	0.435	0.000	**0.937**	0.343	0.140	0.499
Laiho	37	2204	0.007	0.062	-0.007	0.030	0.073	-0.007	0.093	0.044	0.061	0.000	**0.220**	0.036	0.000	0.091	0.185
Lapointe-v1	69	1627	0.088	0.012	0.151	0.179	0.151	0.088	0.149	0.151	0.162	0.034	0.165	0.012	0.156	**0.180**	0.181
Lapointe-v2	110	2498	0.008	0.097	0.033	0.153	0.147	0.028	0.118	0.171	0.210	0.199	0.000	-0.006	**0.239**	0.133	0.172
Liang	37	1413	0.301	0.301	0.301	0.301	0.301	0.301	0.301	0.301	0.481	0.243	**0.523**	0.301	0.419	0.296	0.481
Nutt-v1	50	1379	0.171	0.311	0.159	0.156	0.109	0.086	0.078	0.113	0.116	0.000	0.443	0.215	0.129	0.112	**0.459**
Nutt-v2	28	1072	-0.025	0.000	-0.024	-0.025	-0.031	-0.025	-0.027	-0.030	-0.027	0.035	0.152	**0.250**	-0.029	-0.002	**0.250**
Nutt-v3	22	1154	0.063	0.000	0.004	0.080	0.059	0.080	0.174	0.059	-0.002	0.000	0.589	0.511	0.000	0.225	**0.752**
Pomeroy-v1	34	859	0.012	-0.032	-0.020	-0.006	-0.020	0.008	-0.026	-0.020	0.061	-0.007	**0.589**	-0.014	0.140	0.056	0.067
Pomeroy-v2	42	1381	0.502	0.576	0.591	**0.617**	0.568	0.577	0.602	0.568	0.586	0.376	0.513	0.544	0.582	0.564	0.492
Ramaswamy	190	1365	0.618	0.401	0.623	0.651	0.618	0.620	0.663	0.639	0.592	0.336	0.000	0.547	**0.676**	0.182	0.450
Risinger	42	1773	0.210	0.162	0.297	0.223	0.201	0.258	0.153	0.201	0.309	0.000	0.399	0.308	0.248	0.087	**0.428**
Shipp-v1	77	800	**0.264**	0.035	0.208	0.132	-0.002	0.168	0.203	-0.002	0.113	0.079	0.101	0.069	0.124	0.013	0.065
Singh	102	341	0.048	0.037	0.019	0.033	-0.003	0.069	-0.003	0.066	0.079	0.066	0.000	0.029	0.034	0.083	**0.159**
Su	174	1573	0.666	0.672	0.662	0.738	0.687	0.650	0.667	0.660	0.657	0.589	0.000	**0.824**	0.702	0.288	0.738
Tomlins-v1	104	2317	0.396	0.382	0.454	0.409	**0.647**	0.469	0.419	0.590	0.374	0.423	0.000	0.485	0.513	0.165	0.413
Tomlins-v2	92	1290	0.368	0.222	0.215	0.292	0.226	0.383	0.354	0.311	0.340	0.000	0.000	0.468	0.373	**0.470**	0.288
West	49	1200	0.489	0.403	0.489	0.442	**0.515**	0.489	0.442	**0.515**	0.258	0.00	0.459	0.412	0.391	0.016	0.322
Yeohv2	248	2528	0.385	0.002	0.383	0.479	0.530	**0.550**	0.337	0.442	0.405	0.000	0.057	0.172	0.465	-0.001	0.135