

Solution path clustering with adaptive concave penalty

Yuliya Marchetti* and Qing Zhou^{†,‡}

*Department of Statistics, University of California, Los Angeles
8125 Math Sciences Bldg., Los Angeles, California 90095*

e-mail: yuliya@stat.ucla.edu; zhou@stat.ucla.edu

Abstract: Fast accumulation of large amounts of complex data has created a need for more sophisticated statistical methodologies to discover interesting patterns and better extract information from these data. The large scale of the data often results in challenging high-dimensional estimation problems where only a minority of the data shows specific grouping patterns. To address these emerging challenges, we develop a new clustering methodology that introduces the idea of a regularization path into unsupervised learning. A regularization path for a clustering problem is created by varying the degree of sparsity constraint that is imposed on the differences between objects via the minimax concave penalty with adaptive tuning parameters. Instead of providing a single solution represented by a cluster assignment for each object, the method produces a short sequence of solutions that determines not only the cluster assignment but also a corresponding number of clusters for each solution. The optimization of the penalized loss function is carried out through an MM algorithm with block coordinate descent. The advantages of this clustering algorithm compared to other existing methods are as follows: it does not require the input of the number of clusters; it is capable of simultaneously separating irrelevant or noisy observations that show no grouping pattern, which can greatly improve data interpretation; it is a general methodology that can be applied to many clustering problems. We test this method on various simulated datasets and on gene expression data, where it shows better or competitive performance compared against several clustering methods.

MSC 2010 subject classifications: Primary 62H30, 62J07; secondary 68T05.

Keywords and phrases: Clustering, sparsity, concave regularization, coordinate descent, MM algorithm.

Received April 2014.

Contents

1	Introduction	1570
2	Formulation	1573
3	Solution path clustering	1575
3.1	An MM algorithm	1575

*Supported by NSF graduate research fellowship and UCLA dissertation year fellowship.

[†]Supported in part by NSF grants DMS-1055286 and DMS-1308376.

[‡]Corresponding author.

3.2	Solution path construction	1578
3.3	The full SPC algorithm	1580
4	Simulation study	1581
4.1	Competing methods and cluster quality assessment	1581
4.2	Results for $n > p$	1584
4.3	Results for $n < p$	1587
4.4	Results for non-convex and non-spherical clusters	1589
4.5	Sensitivity to tuning parameters	1591
5	Solution selection	1591
6	Clustering gene expression data	1594
7	Discussion	1597
	Appendix	1598
	A.1 Proof of Lemma 1	1598
	A.2 Proof of Lemma 2	1599
	A.3 ARI	1599
	References	1599

1. Introduction

Cluster analysis allows us to group a collection of objects into subsets such that objects within a subset are similar to each other, while objects in different subsets are dissimilar from each other. Clustering is widely used in exploratory data analysis and has a great variety of applications ranging from biology, astrophysics to social sciences and psychology. Clustering is a first step in knowledge discovery where prior information is rarely available to a researcher. It is also instrumental for visualization of complex data and for partitioning a dataset into more homogeneous groups in which simpler models might be adequate. At the same time, clustering can enable discovery of unknown groups or associations, providing deeper insights into the data. For example, in biological research, clustering can help determine which genes are associated with particular cellular functions or phenotypes and can help isolate subclasses of diseases for targeted treatments.

There exists a great variety of clustering methods. A majority of them rely on minimizing some loss function, usually by an iterative procedure, like the well known k-means algorithm, while other methods recursively organize objects into trees, like the popular hierarchical clustering. Spectral clustering techniques are based on graph theory and matrix decomposition and are gaining popularity as being simple, accurate, and able to find non-convex clusters (von Luxburg, 2007; Hastie, Tibshirani and Friedman, 2009). Along with k-means and its variants, there also exist mixture-likelihood clustering approaches (Mclachlan, Bean and Peel, 2002; Fraley and Raftery, 2002; Yeung et al., 2001) that assume an underlying statistical model for the data and maximize the likelihood function with the EM algorithm or MCMC methods (Bensmail et al., 1997; Oh and Raftery, 2007). Yet other methods take modified or combined approaches, for example, self-organizing maps (Kohonen, 1990), dp-means (Kulis and Michael, 2011),

CLICK (Sharan, Maron-Katz and Shamir, 2003), gene shaving (Hastie et al., 2000), pclust (Wang, Neill and Miller, 2008), support vector clustering (Ben-Hur et al., 2001) to name just a few. Clustering literature is vast, and surveys of clustering methodologies are usually specialized. A recent comprehensive review of basic and new clustering methodologies is presented in Aggarwal and Reddy (2013), and more general developments and trends in clustering are covered in Jain (2010).

Increasingly large and complex datasets, such as those in gene expression analysis or data mining, have created the need for new efficient approaches to clustering. Data now often contain large amounts of both noisy observations and irrelevant variables. Most existing clustering methods do not address the problems of identifying noise and selecting meaningful variables. Only recently, researchers have shown that simultaneously accommodating for the presence of noisy or irrelevant observations can immensely improve clustering results and provide a better interpretation of the patterns in the data (Tseng and Wong, 2005; Thalamuthu et al., 2006). A number of new clustering algorithms have been proposed, the most popular of which include the resampling-based tight clustering method (Tseng and Wong, 2005), penalized weighted k-means (Tseng, 2007), and model-based clustering (Fraley and Raftery, 2002). Other methods that take noise into account include adap_Cluster (De Smet et al., 2002), k-clips (Maitra and Ramler, 2009), DWCN (Shen, Sun and Li, 2010), trimmed k-means (Garcia-Escudero et al., 2008) and other robust clustering algorithms (Soltanolkotabi, Elhamifar and Candes, 2013; Forero, Kekatos and Giannakis, 2012). Another challenging task in clustering in general is the specification of the number of clusters, which is required as an input for most of the existing methods. When the number of clusters is a required input, the obtained solution is prone to error, especially when a dataset is large and complex. There are a number of methods that suggest rules for choosing the number of clusters. For a recent overview, please refer to Fang and Wang (2012).

A different class of methodologies that have gained popularity for high-dimensional complex datasets is sparsity regularization techniques. Some examples of such penalization methods include the lasso (Tibshirani, 1996), the elastic net (Zou and Hastie, 2005), the group lasso (Yuan and Lin, 2006), the fused lasso (Tibshirani et al., 2005), SCAD (Fan and Li, 2001), SparseNet (Mazumder, Friedman and Hastie, 2011), and MC+ (Zhang, 2010). These methods are mostly used in linear and generalized linear models for identifying useful predictors among a large number of covariates. They introduce a penalty to a loss function to find sparse solutions in challenging problems. Efficient optimization methods exist that compute the entire regularization path or a particular solution for a penalized loss function, such as the least angle regression (Efron et al., 2004) and coordinate descent (Friedman et al., 2007; Friedman, Hastie and Tibshirani, 2010; Wu and Lange, 2008).

We propose a novel solution path clustering (SPC) method that can be applied to a wide range of data settings, including high dimensionality and the presence of noisy or irrelevant observations, as this method is able to isolate these observations in singleton or very small clusters. Our algorithm minimizes

a penalized quadratic loss function under the minimax concave penalty (MCP) (Zhang, 2010). The regularization allows us to obtain sparse solutions and to construct a solution path with a decreasing number of clusters, which eliminates the need to specify the number of clusters as an input parameter. The method minimizes a non-convex objective function via the majorization-minimization (MM) algorithm (Lange, 2004) coupled with block coordinate descent. We also develop adaptive data-driven strategies for selecting the penalty parameters along a solution path so that the SPC algorithm has in effect only one tuning parameter for initializing the path. Overall, SPC is a simple, easily implemented and relatively fast algorithm that has worked well in practice, although its convergence properties are to be established in future work.

Penalized estimation has been previously utilized in clustering mostly for variable selection, such as in gene expression analysis (Pan and Shen, 2007; Wang and Zhu, 2008; Xie, Pan and Shen, 2008; Zhou, Pan and Shen, 2009; Guo et al., 2010; Witten and Tibshirani, 2010; Sun and Wang, 2012). Unlike SPC, these methods assume a given number of clusters and select useful variables to partition objects. Very recently, a number of authors introduced penalized clustering methods that are similar to our method in that they also impose a penalty on the pairwise differences between cluster centers and can generate solution paths that do not require the specification of the number of clusters. Pelckmans et al. (2005), Hocking et al. (2011), Lindsten, Ohlsson and Ljung (2011) and Chi and Lange (2013) suggested minimizing an objective function where the penalty on the differences between the cluster centers is convex and, thus, the resulting algorithms are guaranteed to converge to a global minimizer. The potential severe biases in the cluster center estimates from these procedures are handled through penalty weights. The first three referenced papers focus primarily on the optimization of objective functions with convex penalties and do not discuss the choice of the penalty parameter, the detection of noise, performance on high-dimensional data, or the impact of the penalty weights on the clustering results. Chi and Lange (2013) have improved on these convex clustering methods and have provided a general unified algorithm for solving such problems. They have also noted that the solution paths obtained from convex clustering can be unsatisfactory if the weights are not selected properly.

Finally, Pan, Shen and Liu (2013) have proposed a penalized regression-based clustering method (PRclust) using a novel non-convex penalty on the pairwise differences in order to alleviate the possible bias of convex penalties. The authors have re-parametrized the objective function to ensure the convergence of the coordinate descent algorithm to a stationary point. PRclust, however, has not been shown to handle noisy and large high-dimensional datasets. In contrast to our adaptive selection of penalty parameters, Pan, Shen and Liu (2013) mainly focus on determining the number of clusters by searching over a pre-specified grid of three penalty parameters resulting from the re-parametrization, which might not be efficient for large complex datasets.

The remainder of this paper is organized as follows. Section 2 provides a general formulation of a clustering problem under a concave penalty. In Section 3 we develop our clustering algorithm with adaptive solution path construction.

In Section 4 we use simulated data to illustrate and compare the SPC algorithm to several clustering methods. In Section 5 we briefly discuss solution selection for SPC. In Section 6 we apply SPC to a gene expression dataset from mouse embryonic stem cells to show the performance on bigger real data. Finally, Section 7 contains further discussion and future research directions.

2. Formulation

Let $Y = (y_{im})_{n \times p}$ be an observed data matrix, where $y_i = (y_{i1}, \dots, y_{ip}) \in \mathbb{R}^p$ represents the i th object. Assuming that the underlying model for y_i , $i = 1, \dots, n$, is multivariate Gaussian with a mean parameter $\theta_i \in \mathbb{R}^p$ and a constant diagonal covariance matrix $\sigma^2 I_p$, we propose to cluster the n objects into an unknown number of clusters K by minimizing a penalized ℓ_2 loss function. This is achieved by the use of sparsity regularization on the difference between pairwise mean parameters $d(\theta_i, \theta_j) = \|\theta_i - \theta_j\|_2$. Our goal is then to minimize

$$\ell(\theta) = \sum_{i=1}^n \|y_i - \theta_i\|_2^2 + \lambda \sum_{i < j} \rho(\|\theta_i - \theta_j\|_2) \quad (1)$$

over $\theta = (\theta_1, \dots, \theta_n)$, where $\lambda > 0$ and $\rho(\cdot)$ is some penalty function. With a careful choice of $\rho(\cdot)$ we can achieve sparsity such that $\|\hat{\theta}_i - \hat{\theta}_j\|_2$ is arbitrarily small when λ is sufficiently large, where $(\hat{\theta}_1, \dots, \hat{\theta}_n)$ is the minimizer of (1). An important advantage of this formulation, for the situations when there is very little prior knowledge about the data, is that the number of clusters K does not need to be specified beforehand. This regularization also makes it possible to naturally separate noisy objects into singletons and to prevent them from erroneously merging into other clusters.

An appropriately chosen penalty should result in an estimator that satisfies the properties of unbiasedness, sparsity, and continuity (Fan and Li, 2001). To achieve these three properties and to specifically avoid excessive bias in the estimation of θ , which could lead to unsatisfactory cluster assignment, we propose to use the minimax concave penalty (MCP) developed by Zhang (2010),

$$\begin{aligned} \rho(t) &= \int_0^t \left(1 - \frac{x}{\delta\lambda}\right)_+ dx \\ &= \left(t - \frac{t^2}{2\lambda\delta}\right) I(t < \lambda\delta) + \left(\frac{\lambda\delta}{2}\right) I(t \geq \lambda\delta), \quad (t \geq 0), \end{aligned} \quad (2)$$

where $I(\cdot)$ is the indicator function. The MCP penalty $\rho(t)$ in (2) defines a family of penalty functions that are concave in $t \in [0, \infty)$, where $\lambda > 0$ controls the amount of regularization and $\delta > 0$ controls the degree of concavity. It has been noted that such non-convex penalties promote sparser models than the ℓ_1 penalty with the same or superior prediction accuracy in regression models (Zhang, 2010; Mazumder, Friedman and Hastie, 2011). In fact, MCP includes both the ℓ_1 penalty when $\delta \rightarrow \infty$ and the ℓ_0 penalty when $\delta \rightarrow 0+$, forming

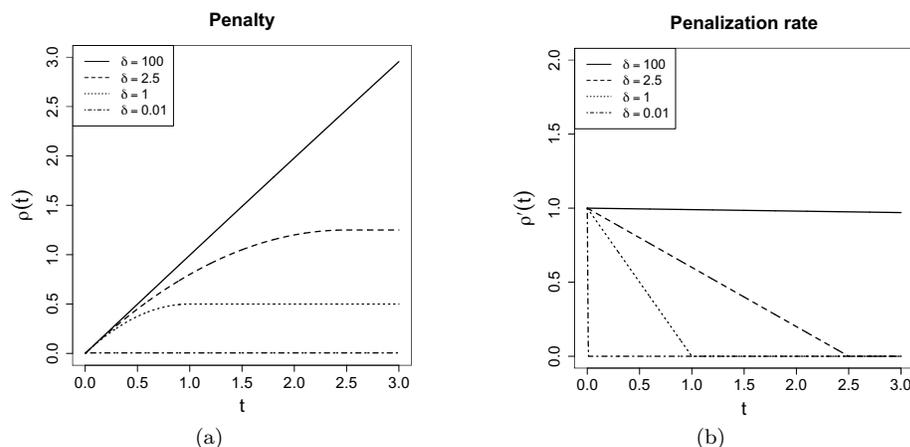


FIG 1. MCP in (a) and its derivative in (b) are plotted for different values of δ and $\lambda = 1$. It approaches the ℓ_1 penalty when δ is high (solid line).

a continuum between the two extremes. MCP is a simple differentiable penalty function with only two parameters and is designed to minimize maximum concavity. Compared to other non-convex penalties such as SCAD (Fan and Li, 2001) or the truncated Lasso penalty (Pan, Shen and Liu, 2013), it includes the explicit concavity parameter δ in its formulation that is easily separated from the penalization rate. Increasing the concavity through this parameter allows us to effectively control the bias of θ using a data-driven approach. The minimax concave penalty is demonstrated in Figure 1.

We illustrate this regularization in the penalized loss function (1) with a special case $n = 2$. Denote the sample mean of the two observations by $\bar{y} = \frac{1}{2}(y_1 + y_2)$. The objective function in this special case is

$$\ell(\theta_1, \theta_2) = \|y_1 - \theta_1\|_2^2 + \|y_2 - \theta_2\|_2^2 + \lambda\rho(\|\theta_1 - \theta_2\|_2). \quad (3)$$

Let $\gamma = (\theta_2 - \theta_1) \in \mathbb{R}^p$. Then, for any fixed γ , $\ell(\theta_1, \theta_2)$ is minimized at $(\theta_1, \theta_2) = (\bar{y} - \gamma/2, \bar{y} + \gamma/2)$, and thus, minimizing (3) reduces to

$$\begin{aligned} \min_{\gamma} \left[\ell(\gamma) &= \|y_1 - \bar{y} + \gamma/2\|_2^2 + \|y_2 - \bar{y} - \gamma/2\|_2^2 + \lambda\rho(\|\gamma\|_2) \right. \\ &= \left. \frac{1}{2}\|\gamma - (y_2 - y_1)\|_2^2 + \lambda\rho(\|\gamma\|_2) \right]. \end{aligned} \quad (4)$$

Figure 2 plots $\ell(\theta_1, \theta_2)$ and the corresponding $\ell(\gamma)$ for different combinations of (λ, δ) for the MCP in (2), and demonstrates how the choice of these two tuning parameters affects the estimation of the cluster centers. The light gray contours in Figure 2a represent the unpenalized ℓ_2 loss (with $\delta \rightarrow 0$ or $\lambda \rightarrow 0$), which is minimized at (y_1, y_2) and gives two different clusters. The black contours in

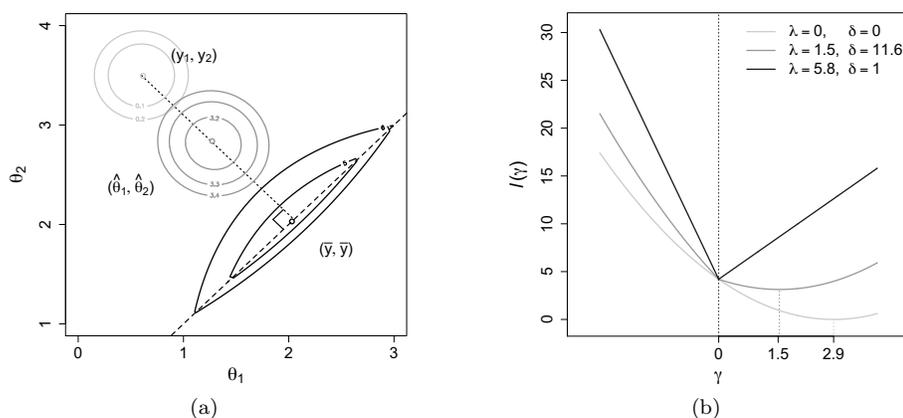


FIG 2. Demonstration of the regularization with $n = 2$ and $p = 1$ for three sets of (λ, δ) : black color for $(\lambda, \delta) = (5.8, 1)$, dark gray for $(\lambda, \delta) = (1.5, 11.6)$, and light gray for $(\lambda, \delta) = (0, 0)$. (a) The contours of $\ell(\theta_1, \theta_2)$ (3). The dark gray contours centered at $(\hat{\theta}_1, \hat{\theta}_2) = (1.3, 2.8)$ demonstrate the bias in the estimates of the cluster centers. (b) The penalized loss function $\ell(\gamma)$ (4) with minimizer indicated by a vertical dashed line.

the same figure depict the objective function (3) when the value of $\delta = 1$ is low enough and the value of $\lambda = 5.8$ is big enough to produce an unbiased estimate of one cluster center $\hat{\theta}_1 = \hat{\theta}_2 = \bar{y}$. The penalty on $|\theta_1 - \theta_2|$ forces the minimizer to move from (y_1, y_2) , when the loss function has no penalty, to $\hat{\theta}_1 = \hat{\theta}_2 = \bar{y}$, which lies on the dashed line $\theta_1 = \theta_2$. The dark gray contours plot the objective function for a larger value of $\delta = 11.6$ and a smaller value of $\lambda = 1.5$, which is minimized at $(\hat{\theta}_1, \hat{\theta}_2) = (1.3, 2.8)$. In this case, two clusters are obtained, and both centers are estimated with substantial bias. Therefore, a proper and data-driven choice of δ and λ is key to our method. This also shows that a penalty in the form $\|\theta_i - \theta_j\|_2$, which corresponds to $\delta \rightarrow \infty$, may not be appropriate for clustering.

Figure 2b plots the objective function $\ell(\gamma)$ with the same combinations of tuning parameters as those in Figure 2a. It can be seen from the figure that $\ell(\gamma)$ with $(\lambda, \delta) = (5.8, 1)$ is minimized when $\gamma = 0$, i.e. $\theta_1 = \theta_2 = \bar{y}$, and the unpenalized loss function in light gray color is minimized at $\gamma = (y_2 - y_1)$. The loss function with $(\lambda, \delta) = (1.5, 11.6)$ in dark gray color is minimized at $\gamma = 1.5$, between 0 and $(y_2 - y_1)$.

3. Solution path clustering

3.1. An MM algorithm

Minimization of a non-convex objective function is usually non-trivial. Motivated by the MM algorithm (Lange, 2004) we propose to majorize the penalty term of (1) by a linear function (Wu and Lange, 2008). We then minimize the

majorizing surrogate function by cyclic block coordinate descent. We initialize the algorithm assuming all objects form singleton clusters and gradually merge the objects into a decreasing number of clusters for an appropriately chosen sequence of parameters (δ, λ) , stopping when all the objects form one cluster. Correspondingly, once two objects are merged into a cluster, we do not consider splitting them in the later stages of the algorithm (Remark 3). Suppose that in the current solution the objects y_i are assigned to K clusters with centers μ_1, \dots, μ_K , where $\mu_k \in \mathbb{R}^p$. Let $C_k = \{i : \theta_i = \mu_k\}$ represent the k th current cluster and $N_k = |C_k|$ denote the size of this cluster. We may then rewrite the objective function (1) as

$$\ell_K(\mu) = \sum_{k=1}^K \sum_{i \in C_k} \|y_i - \mu_k\|_2^2 + \lambda \sum_{k < \ell} N_k N_\ell \rho(\|\mu_k - \mu_\ell\|_2), \quad (5)$$

where $\mu = (\mu_1, \dots, \mu_K)$. With a proper choice of (λ, δ) , minimizing $\ell_K(\mu)$ over μ will force some μ_k 's to be very close to each other, effectively merging these clusters into a bigger cluster in the next solution.

To minimize $\ell_K(\mu)$, we employ a blockwise MM step to cycle through μ_k . At each step, we fix $\mu_{[-k]} = (\mu_1, \dots, \mu_{k-1}, \mu_{k+1}, \dots, \mu_K)$ to its current value and majorize

$$\ell_K(\mu_k) \triangleq \sum_{i \in C_k} \|y_i - \mu_k\|_2^2 + \lambda N_k \sum_{\ell \neq k} N_\ell \rho(\|\mu_k - \mu_\ell\|_2). \quad (6)$$

Let $\mu_k^{(t)}$ be the value of μ_k before the current MM step, where t corresponds to the iteration number. By assumption $\mu_k^{(t)} \neq \mu_\ell$ for all ℓ , and we majorize $\rho(\|\mu_k - \mu_\ell\|_2)$ by

$$\begin{aligned} & \rho(\|\mu_k - \mu_\ell\|_2) \\ & \leq \rho(\|\mu_k^{(t)} - \mu_\ell\|_2) + \rho'(\|\mu_k^{(t)} - \mu_\ell\|_2) \left(\|\mu_k - \mu_\ell\|_2 - \|\mu_k^{(t)} - \mu_\ell\|_2 \right) \\ & \leq \rho(\|\mu_k^{(t)} - \mu_\ell\|_2) + \rho'(\|\mu_k^{(t)} - \mu_\ell\|_2) \left(\frac{\|\mu_k - \mu_\ell\|_2^2 - \|\mu_k^{(t)} - \mu_\ell\|_2^2}{2\|\mu_k^{(t)} - \mu_\ell\|_2} \right), \end{aligned} \quad (7)$$

due to the concavity of the functions $\rho(x)$ and \sqrt{x} for $x > 0$. When the majorization (7) is substituted into $\ell_K(\mu_k)$ (6) we obtain a quadratic surrogate function in μ_k for which the minimizer is

$$\mu_k^{(t+1)} = \frac{\bar{y}_k + \lambda \sum_{\ell \neq k} w_{k,\ell}^{(t)} \mu_\ell}{1 + \lambda \sum_{\ell \neq k} w_{k,\ell}^{(t)}}, \quad (8)$$

where $\bar{y}_k = \frac{1}{N_k} \sum_{i \in C_k} y_i$ and $w_{k,\ell}^{(t)}$ can be regarded as the weight for μ_ℓ :

$$w_{k,\ell}^{(t)} = \frac{N_\ell \rho'(\|\mu_k^{(t)} - \mu_\ell\|_2)}{2\|\mu_k^{(t)} - \mu_\ell\|_2} = \frac{N_\ell (1 - \|\mu_k^{(t)} - \mu_\ell\|_2 / \lambda \delta)_+}{2\|\mu_k^{(t)} - \mu_\ell\|_2}. \quad (9)$$

Algorithm 1 One iteration of the MM algorithm

```

1: for  $k = 1, \dots, K$  do
2:   majorization: compute weights  $w_{k,\ell}^{(t)}$  as in (9) for all  $\ell \neq k$ 
3:   minimization: update  $\mu_k^{(t+1)}$  as in (8)
4:   if  $\|\mu_k^{(t+1)} - \mu_\ell\|_2 < \xi$  for some  $\ell$  then
5:     set  $\mu_k^{(t+1)}$  and  $\mu_\ell$  to their weighted mean
6:   end if
7: end for

```

In effect, the derivative ρ' in (9) becomes the adaptive weight for the estimation of the cluster centers μ_k , which is similar to the convex clustering penalty weights in Hocking et al. (2011) and Chi and Lange (2013). The weight in (9), however, uses the distances between the cluster centers μ_k and varies with each iteration, whereas weights in Hocking et al. (2011) and Chi and Lange (2013) are based on the distances between the data points, which do not change throughout the estimation procedure and solution path. It can be seen from (8) and (9) that:

- When $\|\mu_k^{(t)} - \mu_\ell\|_2 \geq \lambda\delta$ for all $\ell \neq k$, then $w_{k,\ell}^{(t)} = 0$ and, thus, in the next iteration $\mu_k^{(t+1)} = \bar{y}_k$, the sample mean.
- When $\|\mu_k^{(t)} - \mu_\ell\|_2 \ll \lambda\delta$ so that $\lambda w_{k,\ell}^{(t)} \gg 1$ for a particular ℓ , then $\mu_k^{(t+1)} \approx \mu_\ell$ and the two clusters C_k and C_ℓ will merge.

A single iteration of the MM algorithm cycles through all K blocks as summarized in Algorithm 1. To account for data scaling we set $\xi = \frac{\epsilon}{\sqrt{p}} \sum_{m=1}^p \sigma_m$, where $\epsilon = 10^{-4}$ and σ_m is the standard deviation of the m th component of the data.

Remark 1 (Relative sparsity). It should be noted that the minimizer of (5) in general does not have exactly identical pairs of μ_k 's for any finite λ , and thus only relative sparsity can be achieved so that some μ_k 's become very close to each other. This is due to the penalized loss function (5) itself, not because of the algorithm implementation. However, such relative sparsity is sufficient for practical applications, with a simple thresholding step like on lines 4-5 of Algorithm 1. We perform the thresholding step for every iteration of the MM algorithm to decrease the overall running time of SPC as K may become smaller for the next iteration.

Remark 2 (Convergence). Convergence for coordinate descent type algorithms has been established for the sum of a smooth function and a non-convex penalty under certain conditions that are in fact met by MCP (Mazumder, Friedman and Hastie, 2011). However, since the penalty term in (5) is non-separable, this result as well as that in Tseng and Yun (2009) cannot be applied. Another difficulty is that the objective function (6) is non-differentiable when $\mu_k = \mu_\ell$, which does not meet the assumption for the fixed points of the MM algorithm to coincide with the set of the stationary points of the objective function (Lange, 1995; Lange, Hunter and Yang, 2000). Consequently, there is no theoretical guarantee

that Algorithm 1 always converges to a stationary point. In our implementation, Algorithm 1 is repeated until it reaches the stopping criterion

$$\max_{1 \leq k \leq K} \|\mu_k^{(t+1)} - \mu_k^{(t)}\|_2 < \xi \quad (10)$$

or until 50 iterations. In practice, we have observed that Algorithm 1 almost always terminates in fewer than 50 iterations. See Sections 4.2 and 4.3 for more discussion with numerical results.

Remark 3 (Cluster splitting). In order to minimize (1), it is necessary that objects are allowed to be split or unfused from the clusters to which they were assigned. However, to save computation time, we have made a heuristic assumption that a cluster is never split so that μ_k 's are gradually merged into a decreasing number of clusters.

3.2. Solution path construction

The penalty function in (2) has two parameters λ and δ , the former controlling the amount of regularization and the latter determining the degree of concavity of the function. We would like to create some simple data-driven rules for selecting several combinations of the penalty parameters to produce a solution path for any clustering problem. We start the algorithm assuming that each individual observation y_i forms its own singleton cluster. We then gradually enforce sparsity in the differences between cluster centers by using an increasing sequence of λ while reducing the bias, if necessary, through a decreasing sequence of δ . Since the penalty function (2) depends on the distances between μ_k 's, the sequences of the tuning parameters can be guided by these distances, and the solution from the current combination of (δ, λ) can be used as a warm start for the next combination.

We define a decreasing sequence $\Delta = \{\delta_1, \dots, \delta_H\}$ and for each δ_h , $h = 1, \dots, H$, define an increasing sequence $\Lambda(\delta_h) = \{\lambda_1(\delta_h), \dots, \lambda_G(\delta_h)\}$. The sequence Δ is simply determined by

$$\delta_h = \delta_{h-1}\alpha, \quad (11)$$

for $h = 2, \dots, H$, where $\alpha \in (0, 1)$ is a constant. We discuss the choice of δ_1 later in this section when we talk about $\lambda_1(\delta_1)$. Each time the value of δ is decreased a new $\Lambda(\delta)$ is computed. The initial solution is then obtained with the lowest concavity. As mentioned previously, high values of δ decrease concavity of the penalty function and make it behave more like the ℓ_1 penalty, which could introduce considerable bias into the estimate of μ_k . Therefore, in order to determine whether the value of δ needs to be decreased, we define the bias-variance ratio (BVR) for each cluster C_k as

$$\text{BVR}_k = \begin{cases} \frac{\|\mu_k - \bar{y}_k\|_2^2}{\sum_{i \in C_k} \|y_i - \bar{y}_k\|_2^2 / (N_k - 1)}, & N_k > 1 \\ \frac{\|\mu_k - y_i\|_2^2}{(r_k/2)^2}, & N_k = 1, C_k = \{i\}, \end{cases} \quad (12)$$

where $r_k = \min_{\ell \neq k} \|y_i - \mu_\ell\|_2$ is the distance between y_i and the nearest cluster center. If $\text{BVR}_k > 1$ for any k , then we decrease δ as in (11). The idea behind the bias-variance ratio is that if a certain estimated cluster center μ_k moves beyond the range of the observations y_i in that cluster C_k , then the concavity is increased in order to reduce the bias that could lead to a bad solution.

We now address the choice of $\Lambda(\delta_h)$ determined by the lower and upper bounds of the sequence, $\lambda_1(\delta_h)$ and $\lambda_G(\delta_h)$. The values of λ in between are evenly spaced in log-scale. We state two lemmas for a simple case with $n = 2$ and then use these lemmas to motivate our choice of the lower and upper bounds for $\Lambda(\delta_h)$. The proofs for the lemmas are provided in the Appendix.

Lemma 1. *Assume $n = 2$ and that there are only two points (y_1, y_2) with distance $d = \|y_1 - y_2\|_2$. Let $\theta_i^{(t)}$ be the value of θ_i generated by the MM algorithm with $\theta_i^{(0)} = y_i, i = 1, 2$. Fix $\lambda\delta = \eta > 0$. For a given $\phi \in (0, 1)$, if $\eta > d$ and*

$$\lambda = \frac{2\phi\eta d}{(1 - \phi)(\eta - d)}, \tag{13}$$

then $\|\theta_1^{(1)} - \theta_2^{(0)}\|_2 = (1 - \phi)\|\theta_1^{(0)} - \theta_2^{(0)}\|_2$. If $\eta \leq d$, then $\theta_i^{(t)} = y_i$ for all $t \geq 1$ and $i = 1, 2$.

Lemma 2. *Assume $n = 2$ and that there are only two points (y_1, y_2) with distance $d = \|y_1 - y_2\|_2$. For a given $\delta > 0$, if*

$$\lambda \geq \left(1 + \frac{1}{\delta}\right) d, \tag{14}$$

then the global minimizer of (3) is given by $(\hat{\theta}_1, \hat{\theta}_2) = (\bar{y}, \bar{y})$.

We use Lemma 1 to determine δ_1 and the initial lower bound $\lambda_1(\delta_1)$. From this lemma one sees that η serves as a threshold: if $d \geq \eta$, then $\theta_i^{(t)}$ will not change and the points will not merge. Let $\eta_1 = \lambda_1(\delta_1)\delta_1$ and Q_β be the β -quantile of the nearest neighbor distances among y_i 's. We choose $\eta_1 = Q_\omega$, where $\omega \in (0, 1)$ can be regarded as the approximate proportion of data points that may merge in the initial solution. On the other hand, it follows from Lemma 1 that $\phi \in (0, 1)$ can be considered the minimization step size. By default, we set $\phi = 0.5$. Then, to use (13) to determine $\lambda_1(\delta_1)$ we need to specify d . We may choose d as the distance between a pair of points such that $d < \eta_1$. This can be achieved by simply setting $d = Q_\tau$, where $\tau \in (0, \omega)$. Plugging these choices of parameters into (13), we obtain the initial lower bound for λ as

$$\lambda_1(\delta_1) = \frac{2\phi Q_\omega Q_\tau}{(1 - \phi)(Q_\omega - Q_\tau)}. \tag{15}$$

The first value of the sequence Δ then follows directly from

$$\delta_1 = Q_\omega / \lambda_1(\delta_1). \tag{16}$$

Denote the maximum penalty by $z = \lambda \max_t \rho(t) = \frac{1}{2} \lambda^2 \delta$. In order to achieve gradual merging of objects into fewer clusters, both the threshold $\eta = \lambda\delta$ and

z should be non-decreasing. Suppose that δ_{h-1} is decreased to δ_h by the BVR criterion and $z_{h-1} = \frac{1}{2}[\lambda_{\tilde{G}}(\delta_{h-1})]^2 \delta_{h-1}$, where $\lambda_{\tilde{G}}(\delta_{h-1})$ is the value of λ before decrease. The lower bound for $\Lambda(\delta_h)$ is then

$$\lambda_1(\delta_h) = \left(\frac{2z_{h-1}}{\delta_h} \right)^{1/2} = \left(\frac{[\lambda_{\tilde{G}}(\delta_{h-1})]^2 \delta_{h-1}}{\delta_h} \right)^{1/2} = \alpha^{-1/2} \lambda_{\tilde{G}}(\delta_{h-1}). \quad (17)$$

Next, we directly apply Lemma 2 to define the upper bound $\lambda_G(\delta_h)$, $h = 1, \dots, H$. Given a collection of objects, we can conservatively choose d in (14) to be the maximum distance among all pairs of objects to obtain

$$\lambda_G(\delta_h) = \left(1 + \frac{1}{\delta_h} \right) \max_{i,j} \|y_i - y_j\|_2. \quad (18)$$

For a collection of $n > 2$ objects, the upper bound in (18) becomes only an approximation for the value of λ such that all objects merge into a single cluster. This value of λ does not necessarily guarantee that the objects will merge, but in practice, we have not encountered a situation when this approximation did not work. In a case when the value of the upper bound of λ is not sufficiently large, one can simply decrease the value of δ , recalculate the new sequence of λ and run the algorithm until all data points form a single cluster.

In summary, the construction of the solution path involves the specification of four parameters, α in (11), and ω , ϕ and τ in (15), all ranging between 0 and 1. In effect, the parameters τ and ϕ control the step size of the MM iteration, and consequently the length of the solution path. Based on our experience and as demonstrated by the sensitivity analysis in Section 4.5, the solution path is not much affected by the choice of ϕ and τ , except when very small values are used ($\phi, \tau \leq 0.01$), which can slow down the progression of the solution path. We recommend setting τ to be slightly smaller than ω in order to avoid unnecessary detail in the solution path. In general, we recommend setting $\alpha = 0.9$, $\phi = 0.5$ and $\tau = 0.9\omega$, and we use these default values throughout the paper.

The only tuning parameter that needs to be specified by the user is ω for the calculation of $\lambda_1(\delta_1)$ in (15) and δ_1 in (16). Since it stands for the approximate proportion of the nearest neighbors that may merge initially, the nature of the dataset might help determine its value. For instance, if the dataset is very noisy and clusters are not tight, ω could be set to a low value, and if the dataset is well separated into clusters and there is little noise, a high value of ω could be used. If chosen too high in cases where the data is very noisy, ω could force the algorithm to skip the correct solution by initially merging too many noisy observations. In practice, we also found that ω should be small in high-dimensional settings. In this paper we use $\omega = 0.1$ for the high-dimensional examples when $n < p$ and $\omega = 0.5$ when $n > p$.

3.3. The full SPC algorithm

We now combine the MM algorithm with the construction of the solution path to describe the full SPC algorithm. It is initialized assuming that each observation

Algorithm 2 Solution path clustering**Inputs**

required input: $Y = (y_{im})_{n \times p}$, $\omega \in (0, 1)$
 default input: $\tau = 0.9\omega$, $\phi = 0.5$, $\alpha = 0.9$, $G = \min(20, p)$
 initialization: $h = 1$, $K = n$, $\mu_k = y_k$, $k = 1, \dots, n$

```

1: repeat
2:   compute  $\delta_h$ ,  $\lambda_1(\delta_h)$ ,  $\lambda_G(\delta_h)$  and construct  $\Lambda(\delta_h)$  in logarithmic scale of size  $G$ 
3:   for  $g = 1, \dots, G$  do
4:     run the MM algorithm (Algorithm 1) until  $\max_k \|\mu_k^{(t+1)} - \mu_k^{(t)}\|_2 < \xi$  or  $t > 50$  to
       obtain  $K(g, h)$  and  $\{\hat{\mu}_k(h, g), k = 1, \dots, K(h, g)\}$ 
5:     for  $k = 1, \dots, K(h, g)$  do
6:       compute  $\text{BVR}_k$ 
7:       if  $\text{BVR}_k > 1$  then
8:          $h \leftarrow h + 1$  and go to line 2
9:       end if
10:    end for
11:  end for
12:   $h \leftarrow h + 1$ 
13: until  $K(h, g) = 1$ 

```

forms its own singleton cluster and is run until all observations merge into one cluster using a sequence of (δ, λ) . The solution obtained from a particular (δ, λ) is used as a warm start for the next solution. We do not require the input of the number of clusters K , and the algorithm typically yields a short path of 2–15 solutions.

For a particular $(\delta_h, \lambda_g(\delta_h))$, let $K(h, g)$ be the estimated number of clusters and $\hat{\mu}_k(h, g)$ for $k = 1, \dots, K(h, g)$ be the estimated cluster centers. Using these notations, the full SPC algorithm is provided in Algorithm 2.

4. Simulation study

4.1. Competing methods and cluster quality assessment

We illustrate the performance of SPC on simulated data and compare it to the k-means++ algorithm (Arthur and Vassilvitskii, 2007), the convex clustering method of Chi and Lange (2013) and several popular clustering algorithms that were developed to account for the noise in data. For the latter purpose we selected model-based clustering (mclust) (Fraley and Raftery, 2002), tight clustering (Tseng and Wong, 2005), and penalized weighted k-means (PWK-means) (Tseng, 2007).

The majority of authors use the adjusted rand index (ARI) (Hubert and Arabie, 1985) to compare the clustering results across different methods when the true cluster assignment is known. ARI calculates the similarity of a clustering result to the underlying true clustering assignment. It takes a maximum value of 1 if the clustering solution is identical to the true structure and takes a value close to 0 if the clustering result is obtained from random partitioning. The

TABLE 1
Contingency table and notation for the calculation of (ARI_c, ARI_n)

Cluster	v_1	...	v_R	v_{R+1}	Sum
\hat{u}_1	n_{11}	...	n_{1R}	$n_{1(R+1)}$	$n_{1\bullet}$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
\hat{u}_K	n_{K1}	...	n_{KR}	$n_{K(R+1)}$	$n_{K\bullet}$
\hat{u}_{K+1}	$n_{(K+1)1}$...	$n_{(K+1)R}$	$n_{(K+1)(R+1)}$	$n_{(K+1)\bullet}$
Sum	$n_{\bullet 1}$...	$n_{\bullet R}$	$n_{\bullet(R+1)}$	n

exact definition of the ARI can be found in the [Appendix](#). For the simulation study we also use ARI in order to be consistent with the prevailing method of cluster quality assessment.

Since most of the competing methods will detect noise, i.e. the data points that do not belong to any well-defined clusters, we calculate two ARI scores (ARI_c, ARI_n) for each of them. Suppose the true partition and an estimated partition are $C = \{C_1, \dots, C_R, C_{R+1}\}$ and $\hat{C} = \{\hat{C}_1, \dots, \hat{C}_K, \hat{C}_{K+1}\}$, where C_r and \hat{C}_k contain the indices of the data points assigned to true and estimated clusters for $r = 1, \dots, R$ and $k = 1, \dots, K$, respectively, and C_{R+1} and \hat{C}_{K+1} are the indices assigned to noise. Denote the respective cluster labels as $v = \{v_1, \dots, v_R, v_{R+1}\}$ and $\hat{u} = \{\hat{u}_1, \dots, \hat{u}_K, \hat{u}_{K+1}\}$, where v_{R+1} and \hat{u}_{K+1} indicate the labels for the true and estimated noise, respectively. See Table 1 for the full contingency table of the counts $n_{kr} = |\hat{C}_k \cap C_r|$. Our ARI scores are calculated based on parts of the counts in this table.

ARI_c accounts for data points that are identified as belonging to estimated clusters \hat{C}_k , $k = 1, \dots, K$, and is calculated as in (22) in the [Appendix](#) using only the first K rows of Table 1, $\{n_{kr} : 1 \leq k \leq K, 1 \leq r \leq R + 1\}$, with the corresponding row and column sums. In effect, ARI_c provides a quality assessment of the identified clusters \hat{C}_k , i.e. misclassification of clustered data and the amount of noise in the estimated clusters.

ARI_n indicates how sensitive a method is in identifying noise and whether any clustered data point is misclassified as noise. It is based on all the data points except the noise in the estimated clusters, which is accounted for in ARI_c . We collapse Table 1 to a 2×2 table with counts $n_{cc}^* = \sum_{k=1}^K \sum_{r=1}^R n_{kr}$, $n_{nn}^* = n_{(K+1)(R+1)}$, $n_{nc}^* = \sum_{r=1}^R n_{(K+1)r}$, and $n_{cn}^* = \sum_{k=1}^K n_{k(R+1)}$. We set $n_{cn}^* = 0$ since we account for these data points in ARI_c . The total number of data points to be considered for ARI_n is thus $n^* = n_{cc}^* + n_{nn}^* + n_{nc}^*$. Again, we use (22) in the [Appendix](#) to calculate ARI_n by plugging in n_{cc}^* , n_{nn}^* , n_{nc}^* , and $n_{cn}^* = 0$. In general, if n_{nn}^* is large and n_{nc}^* is small, ARI_n will be close to 1.

PWK-means, tight clustering, and k-means++, which is a classical k-means method combined with a randomized seeding procedure to select the starting centers, require, as an input, an estimated number of clusters, and mclust requires an input of a range of the number of clusters. In most cases, we provided the comparison algorithms with ideal input parameters, which will likely result in optimal performance for these methods. In addition to the number of clusters,

the competing methods, except k-means++, have other tuning parameters that we mention below.

PWK-means also requires the input of a penalty parameter λ since this method imposes a penalty on the number of noisy data points. We used the suggested prediction-based resampling method (Tibshirani and Walther, 2005) to find λ , calculated the prediction strength criterion for an increasing sequence of λ , and selected the value of λ corresponding to the highest prediction strength computed.

Tight clustering has several tuning parameters, but most of them are recommended to stay at their default values, which we follow for the simulated data. Along with the user-specified target number of clusters k_{target} , tight clustering relies on a starting number of clusters $k_0 > k_{\text{target}}$. The tight clustering algorithm is then applied to a decreasing sequence, decremented by 1, starting with k_0 and ending with k_{target} . The authors recommend $k_0 \geq k_{\text{target}} + 5$, however, a too large k_0 results in smaller clusters and many of the clustered data identified as noise. Conversely, a small k_0 can result in a smaller size of the estimated noise and more noise assigned to clusters. We set $k_0 = k_{\text{target}} + 5$, the smallest recommended value.

One additional tuning parameter of mclust is the reciprocal of the hypervolume V of the data region, and the authors note that the method is sensitive to this value. The default method of calculating the hypervolume is $V = \prod_{j=1}^p (\max_i \{y_{ij}\} - \min_i \{y_{ij}\})$, which is used for the simulated data. Another input into mclust is the estimated categorization of each data point as clustered data or noise. Once the categorization is provided, mclust applies hierarchical clustering to the identified clustered data to get a good initialization for the EM algorithm that generates a final clustering result. We have used the recommended K th-nearest neighbor cleaning method (Byers and Raftery, 1998) to obtain the initial categorization into clusters and noise. It is suggested that the value of K should be the size of the smallest cluster to be detected, but if K is selected too high, the noise identification might not perform well. We chose $K = 5$ for the nearest neighbor cleaning such that it is smaller than the average cluster size.

Finally, convex clustering requires the specification of three tuning parameters, which are the Gaussian kernel weight ϕ , the number of nearest neighbors k , and an increasing sequence of the regularization parameter γ for obtaining a solution path. We chose the parameter ϕ such that the penalty weights $w_{ij} = \iota_{\{i,j\}}^k \exp(-\phi \|y_i - y_j\|_2^2)$, where $\iota_{\{i,j\}}^k$ is 1 if y_j is among y_i 's k -nearest neighbors and 0 otherwise, are on average around 0.25. This rule forced the parameter ϕ to be very close to 0 for our simulated data, but any larger values for ϕ resulted in unsatisfactory clustering. We also used $k = 5$ and an increasing sequence of $\gamma \in [0, 50]$ of size 20 at even intervals, following an example in the documentation for the R package `cvxclust`.

To compare the performance, we created four different clustering scenarios: 1) well-separated clusters, 2) overlapping clusters, 3) well-separated clusters with added noise, 4) overlapping clusters with added noise. All clusters in the examples in Sections 4.2 and 4.3 are spherical, with equal variance. The cluster

centers and the noise points were generated from a uniform distribution on $[-5, 5]^p$. All noise was generated outside of the radius of the clusters, where the radius is the largest distance from the cluster center to the data points in that cluster. Overlapping clusters were generated such that 15–20% of the data points in a pair of clusters are located within the radiuses of both clusters. For each scenario, we simulated 20 datasets with $n > p$ and $n < p$. To demonstrate the ability of SPC to identify noise and for the purposes of comparison with other methods, we simply regard all estimated clusters of size $N_k \leq 3$ as noise. The same cut off is also used for defining noise from k-means++ and convex clustering results.

4.2. Results for $n > p$

The simulated datasets for both well-separated and overlapping scenarios when $n > p$ are of size $n = 400$, dimension $p = 20$, and with $K = 10$ clusters. For each of the two scenarios with noise, 200 uniformly distributed noise points were added to the clustered data. We chose $\omega = 0.5$ for all the scenarios assuming that about half of the nearest neighbors should merge. As the output from each dataset, we obtained a solution path of 7–12 solutions, each containing the number of clusters K with size $N_k > 3$, estimated cluster centers $\hat{\mu}_k$, and cluster assignments C_k .

The ARI_c and ARI_n scores for the comparison with other methods are presented in Figure 3. In addition to the true number of clusters $K = 10$, we supplied the competing methods with the number of clusters (of size $N_k > 3$) along the SPC solution path to demonstrate their performance when the number of clusters is misspecified. We report the ARI scores averaged over 20 datasets for $K = 10$ and for different ranges of K , e.g. 6 – 9, since each of the SPC solution paths for the 20 datasets might contain a different number of clusters. It must be noted that $\text{ARI}_n = 0$ when no noise is detected, which is misleading for the scenarios without noisy data. Thus, we use a special score, $S_n = 1 - n_{nc}^*/n$, for these scenarios, such that if no noise is identified by a method, then $S_n = 1$. SPC can clearly outperform k-means++, especially in the scenarios with noise, due to the fact that k-means++ is not designed to separate noisy observations. SPC performs similarly or slightly better in most scenarios compared with tight clustering and PWK-means. For $K = 10$, mclust outperforms all the other methods in all scenarios and separates the overlapping clusters and noise well. The pre-classification of the noise and hierarchical clustering of the remaining data provide mclust with an excellent initialization. The spherical nature of the clusters and uniformly generated noise also perfectly match the model assumptions of mclust. For $K = 10$ or a slightly smaller value, the performance of SPC is very comparable to that of mclust for noisy data scenarios, although our method makes much weaker data generation assumptions and does not use any specific initialization.

All the methods recover well-separated clusters with accuracy whether noise is present or not, however, tight clustering tends to considerably underperform when no noise is present due to the fact that it is targeted specifically for noisy

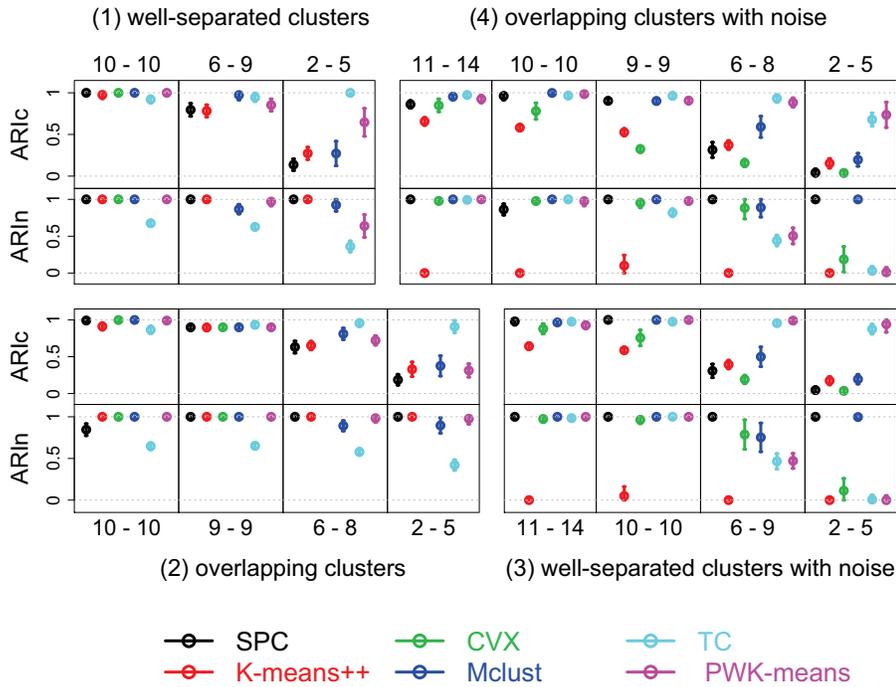


FIG 3. ARI_c and ARI_n for comparison methods for all four scenarios when $n > p$. Each scenario is represented by a panel consisting of a block of plots. Each colored point indicates the average ARI_c or ARI_n across 20 datasets and the error bar shows its 95% confidence interval. ARI scores are between 1 (best) and 0 (worst) as labeled on the y-axis. The range of numbers above or below each plot refers to the number of clusters of size $N_k > 3$ found by SPC. Different colors correspond to different competing methods, as shown in the color legend with convex clustering and tight clustering abbreviated as CVX and TC, respectively. For scenarios (1) and (2) ARI_n reports $S_n = 1 - n_{nc}^*/n$.

data. When clusters are not well-separated, SPC tends to merge overlapping clusters but identifies noise very well (high ARI_n). With the progression toward greater sparsity (smaller K), SPC adds more and more noise into the clusters or creates bigger clusters from noise, which results in a low ARI_c score. On the other hand, the remaining noise is identified accurately, which is reflected in the high ARI_n scores. In contrast, when the number of clusters is misspecified and $K < 10$, tight clustering and PWK-means have the tendency to leave out clustered data as noise, reflected by low ARI_n scores, while mclust merges the closest clusters together. Convex clustering performs well with overlapping clusters, however, it tends to produce less satisfactory results when any noise is present, compared to SPC. In the scenarios with noise, aside from k-means++, convex clustering adds the most noise into the clusters.

Figure 4 shows an example of the cluster assignment for the overlapping scenario with noise for different methods. Figures 4a–4b show cluster assignments

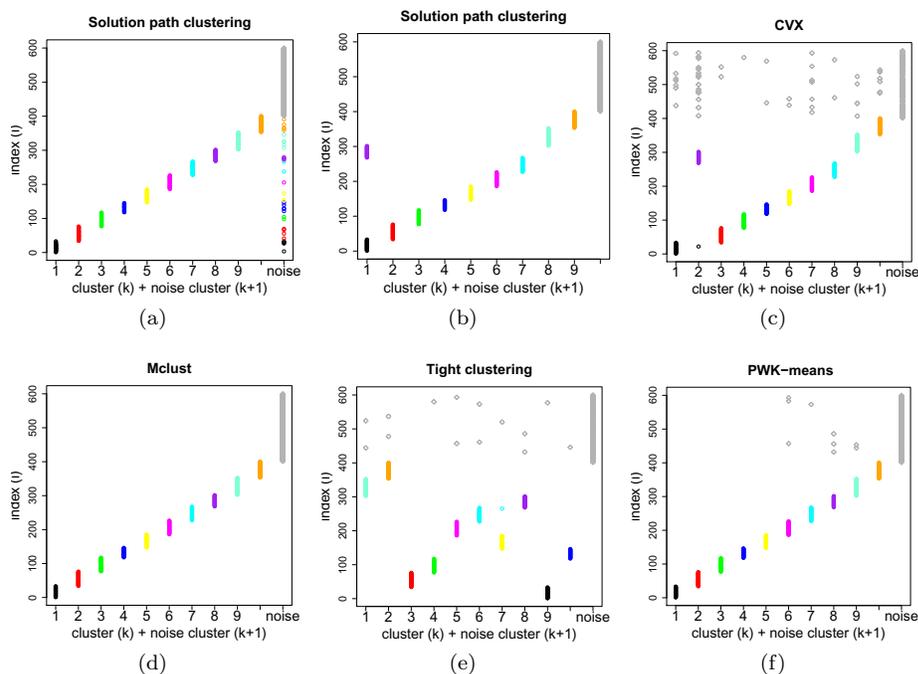


FIG 4. Cluster assignment plots for comparison methods. The y-axis displays the index i of each data point and the x-axis shows K estimated clusters and one additional cluster for the noise. The color indicates the true cluster assignment with true noise in gray, such that if a clustered data point is identified as noise it will appear as noise on the x-axis but will have a color associated with it. If noise is misclassified then it will appear in the clusters but in gray color. (a)–(b) Cluster assignments for solution path clustering for $K = 10$ and $K = 9$, corresponding to two consecutive solutions in a solution path. (c)–(f) Cluster assignments for $K = 10$ for convex clustering (CVX), mclust, tight clustering, and PWK-means.

for two consecutive solutions with $K = 10$ and $K = 9$ for SPC. The solution for $K = 10$ leaves out a number of clustered data points as noise. The solution with $K = 9$ separates the noise perfectly, but merges two overlapping clusters. Figures 4c–4f show the results of competing methods with $K = 10$. One sees that mclust (Figure 4d) separates noise well and misclassifies only one noisy point, while tight clustering (Figure 4e) and PWK-means (Figure 4f) add noise to the clusters. Convex clustering (Figure 4c) has the most noise added to the clusters.

Finally, we provide some empirical evidence which suggests that the SPC algorithm may reach the stopping criterion (10) quickly with the simulated data. Let T^* be the number of iterations before meeting the stopping criterion (10). Figure 5a shows the histogram of T^* for the combined four scenarios and confirms that the SPC algorithm terminated within 35 iterations per solution for all the datasets. We note that T^* generally decreases along a solution path due to warm starts and the reduction in K .

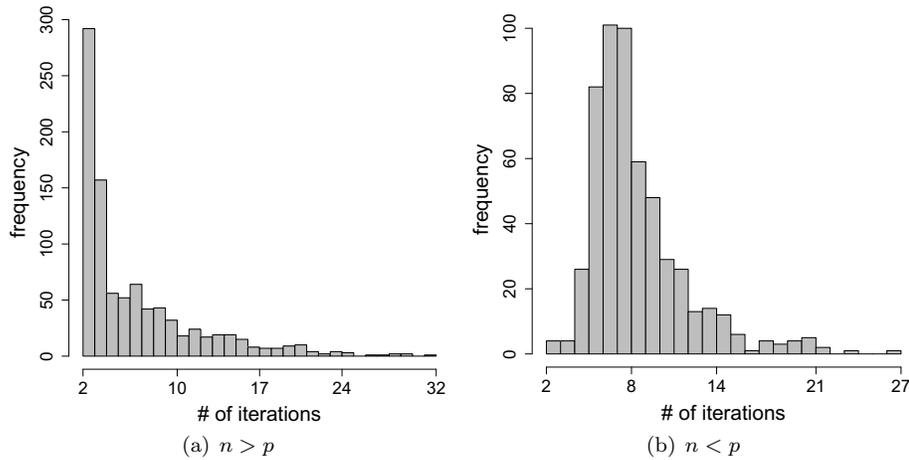


FIG 5. Histograms of T^* for the combined four scenarios for the simulated data.

4.3. Results for $n < p$

For $n < p$ we simulated datasets with $n = 100$ clustered data points, $p = 200$, and $K = 10$. For each dataset, 50 uniformly distributed noise points were added to the clustered data. Similarly to the scenarios with $n > p$, we consider clusters of size $N_k \leq 3$ as noise so that the results can be compared across the methods.

With high-dimensional data we chose $\omega = 0.1$ in order to create a longer and a more detailed solution path. By setting a smaller ω , fewer observations merge into clusters in the initial stages so that the majority of data points are left as singletons or very small clusters. As the sparsity is increased, more observations form new clusters. Consequently, even though the total number of clusters (K_{total}) decreases, the number of clusters of size $N_k > 3$ (K_{clust}) increases along the solution path. As shown in the results for SPC in Figure 6, ARI_c is always high because almost no noisy points merge into clusters and there are few misclassified clustered data points. On the other hand, ARI_n increases as more clusters are formed so that fewer data points are regarded as noise. One sees that when K_{clust} is close to the true number of clusters ($K = 10$), satisfactory results are obtained in terms of both cluster assignment and noise detection.

When the number of dimensions p is higher than the number of observations n , SPC shows a better performance than all the other methods in most scenarios. It performs at least as well as or slightly better than tight clustering. We could not obtain solutions for mclust due to error messages when the initial noise categorization was provided for the EM initialization and, thus, we used the solutions without the pre-specification of noise, which resulted in all the noise included in the clusters. This clearly demonstrates that mclust is sensitive to pre-specification of noise. K-means++ recognized clustered data quite

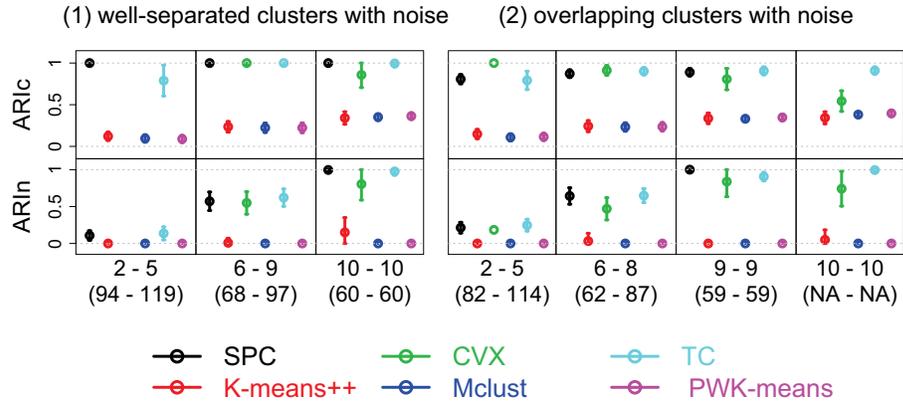


FIG 6. ARI_c and ARI_{in} for comparison methods for two scenarios with noise when $n < p$. The two scenarios are presented in two panels of plots in the same format as those in Figure 3. The x-axis indicates K_{clust} ranges and K_{total} ranges in brackets. NA on the x-axis in the last column in scenario (2) reflects that SPC did not obtain solutions with $K = 10$ for this scenario.

well but grouped the noise with the clusters. The performance of mclust and k-means++ was very similar to that of PWK-means, which also did not recognize noise with penalty parameter λ selected by prediction-based resampling. This comparison shows that identification of noise in high-dimensional space is very challenging and the uniform assumption for the noise may not be appropriate because the volume of the data range becomes huge when p is large.

Convex clustering showed slightly worse ARI scores in the $n < p$ scenarios, adding noise to the clusters as well as leaving out the clustered data as noise, which indicates that high dimensionality might create challenges for this method. The distances between observations in high-dimensional data become very large, forcing the parameter ϕ to be very close to 0 in order to produce reasonably large weights w_{ij} , which corresponds to uniform w_{ij} and introduces more bias into the estimation of the cluster centers, possibly leading to unsatisfactory solutions. The k -nearest-neighbor approach could mitigate this problem somewhat but, as mentioned in qualitative comparisons section in Chi and Lange (2013), can still lead to data points not being agglomerated correctly.

As in the previous section, we demonstrate in Figure 5b the empirical evidence that the SPC algorithm reached the stopping criterion (10) for the simulated data with $n < p$. From the figure we see that T^* is somewhere between 3 and 28 for each solution. In the high-dimensional case T^* is larger on average (about 8 iterations vs. about 4 for $n > p$).

Table 2 summarizes an example of the running time for all the methods except kmeans++, which is by far the fastest method. The run time for SPC is based on the longest solution path for each scenario, usually 10–14 solutions. The matching run times for mclust, tight clustering and PWK-means are calculated based on the unique cluster counts and those for convex clustering are based

TABLE 2
 Summary of run times (in seconds) for SPC and comparison methods

	$n > p$					$n < p$				
	SPC	CVX	mclust	TC	PWK	SPC	CVX	mclust	TC	PWK
(1)	4.56	1.87	1.36	9.42	19.07	4.26	3.63	0.92	812.75	17.60
(2)	3.88	3.76	1.34	12.61	19.26	2.68	3.36	0.95	369.51	14.29
(3)	11.95	19.14	1.43	10.32	39.40	7.38	1.69	0.89	11.16	4.29
(4)	11.43	24.50	1.56	10.36	41.29	9.10	0.76	1.08	11.92	4.64

NOTE: (1) well-separated clusters, (2) overlapping clusters, (3) well-separated clusters with noise, (4) overlapping clusters with noise.

on its corresponding full solution path of length 20. If the number of clusters K_{clust} (of size $N_k > 3$) is duplicated in a solution path, these duplicates are not included in the run times of mclust, tight clustering and PWK-means. The running time for PWK-means includes penalty parameter search with prediction-based resampling. SPC is currently implemented in R, and tight clustering and PWK-means are written in C, mclust in Fortran, and convex clustering in R and Fortran. It seems that SPC's run times compare well to the other methods and that it has a good potential in terms of speed, especially after implementation in a faster language.

4.4. Results for non-convex and non-spherical clusters

We now show the performance of SPC on non-convex and non-spherical simulated clusters. For the non-convex case we generated $n = 400$ observations in $p = 2$ dimensions grouped in $K = 4$ clusters, where data in each cluster were generated from a normal distribution with a high correlation. Three out of the four clusters have a negative or positive correlation of around 0.9 and the other one has a slightly lower correlation of 0.5. After the data points were simulated, non-convexity was introduced to each cluster as shown in Figure 7 (true model). The cluster size was varied to be $N_1 \approx 250$, $N_2 \approx 100$, $N_3 \approx 30$, and $N_4 \approx 20$. Finally, 20 noise data points were added similarly to the scenarios in the previous sections. We then applied all the methods except tight clustering, which did not converge, to this simulated dataset.

It can be seen from Figure 7 that SPC can perform well in relatively complex settings when the clusters are non-convex, the noise or outliers are present and when the cluster sizes are different. All the comparison methods tend to split one of the two bigger clusters (black and green in the true model plot), while SPC does not. Mclust and PWK-means perform well with noise, however, they do not recognize the smallest cluster with the lowest correlation (blue). Mclust also assigns the ends of the non-convex clusters to noise, showing its sensitivity to the assumption of normality. Convex clustering does not seem to be robust to outliers and noise, due to the design of its penalty, specifically the weights, even though it can obtain very good results for non-noisy non-convex data for exactly the same reason.

To further test the relative performance between SPC and the competing methods we applied them in higher dimensions. We generated $K = 4$ similarly

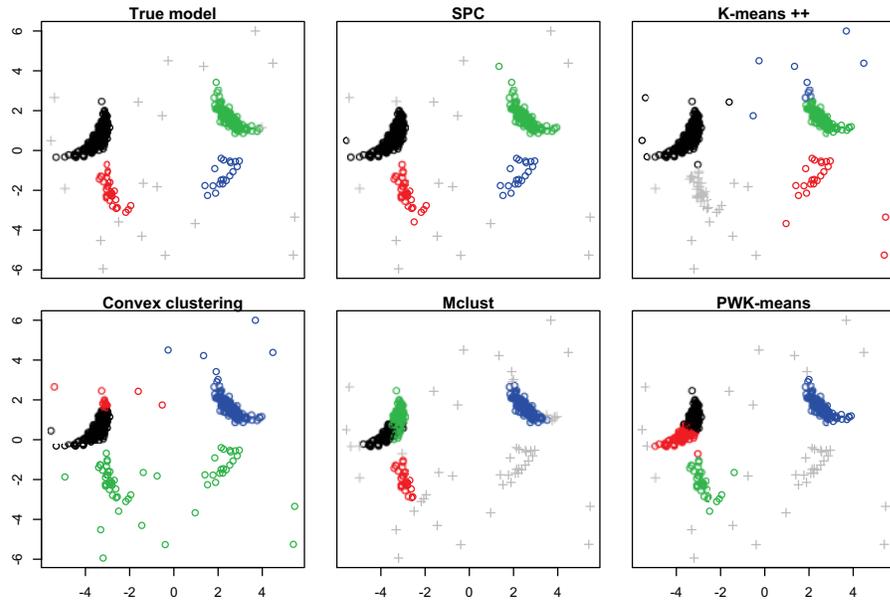


FIG 7. Cluster assignment results for non-convex clusters. Colored points indicate data in clusters and gray points indicate noise.

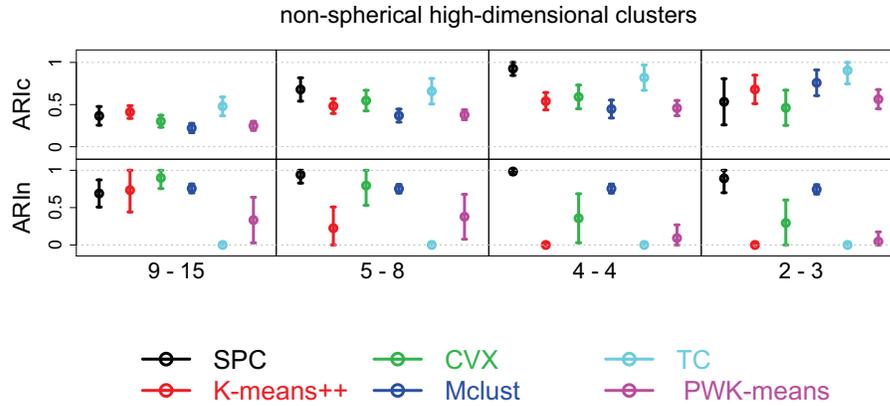


FIG 8. ARI_c and ARI_n for comparison methods for non-spherical simulated clusters with $p = 20$ and $K = 4$ true clusters, in the same format as the panels in Figure 3.

correlated clusters in $p = 20$ dimensions with the same sizes and added 50 noise points. We did not, however, introduce any non-convexity into the clusters. The resulting ARI scores, averaged over 20 randomly generated datasets, are presented in Figure 8. SPC has the highest ARI_c and ARI_n scores for the true

TABLE 3
Solution path length and (ARI_c, ARI_n) scores for different tuning parameters ϕ and τ

τ	$\phi = 0.05$	$\phi = 0.1$	$\phi = 0.3$	$\phi = 0.5$	$\phi = 0.7$
0.05	6 (0.91, 0.99)	9 (0.90, 1.00)	10 (0.91, 1.00)	10 (0.91, 1.00)	11 (0.91, 0.99)
0.15	8 (0.90, 1.00)	13 (0.90, 1.00)	10 (0.91, 1.00)	11 (0.91, 0.99)	11 (0.91, 0.99)
0.25	10 (0.90, 1.00)	10 (0.90, 1.00)	10 (0.91, 1.00)	11 (0.91, 0.99)	11 (0.92, 0.99)
0.35	11 (0.90, 1.00)	10 (0.91, 1.00)	11 (0.91, 0.99)	11 (0.91, 1.00)	11 (0.91, 0.99)
0.45	10 (0.91, 0.99)	10 (0.91, 1.00)	11 (0.91, 1.00)	11 (0.92, 0.99)	11 (0.92, 0.99)

number of clusters $K = 4$ and also performs well when the number of clusters is misspecified. This confirms the usefulness of SPC for clustering high-dimensional non-spherical data.

4.5. Sensitivity to tuning parameters

In this section, we further comment on the sensitivity of the solution path to the tuning parameters ω , ϕ , τ in (15), and α in (11). As mentioned in the previous section, the parameter ω could have a considerable impact by skipping the best solution in the initial step if set too high. This can be avoided by setting this parameter to a low value at the expense of a more detailed solution path with a longer running time. We have found, however, that the rest of the tuning parameters do not impact the solution path in any major way. Setting a different value for α will automatically trigger a corresponding change in the value of λ , which might affect the speed of the algorithm. To further demonstrate the sensitivity to τ and ϕ , we have run SPC on 20 randomly generated datasets with different combinations of these parameters and averaged the length of the solution path with ARI_c and ARI_n for each combination. Table 3 is based on the overlapping case with noise for $n > p$ and illustrates that the cluster assignment and noise detection are not sensitive at all to these two parameters. The length of the solution path, except for the smallest values for both τ and ϕ , is also very stable. The greatest difference between the smaller and larger values of these tuning parameters was the speed of each solution due to the different minimization step size.

5. Solution selection

SPC does not require the specification of the number of clusters, however, it produces a solution path. Unlike the entire clustering path of hierarchical clustering or arbitrary regularization paths of convex clustering (Chi and Lange, 2013; Hocking et al., 2011) or PRclust (Pan, Shen and Liu, 2013), a SPC solution path includes only a limited number of solutions obtained using adaptive data-driven approach. This makes it easier for a user to explore different possible clustering assignments. However, it is still very useful in practice to be able to select a particular member along the solution path, especially for large datasets.

Resampling-based methods such as the gap statistic (Tibshirani, Walther and Hastie, 2001) or the clustering instability method in Fang and Wang (2012),

are computationally intensive, while Bayesian information criterion and cross-validation have been shown to select models that are too complex compared to the true model in sparse linear regression. These problems are especially relevant for our method that aims at high-dimensional and large data. We, therefore, adopt the empirical approach in Fu and Zhou (2013) to demonstrate that solution selection for SPC is possible in a simple and fast way.

The approach of Fu and Zhou (2013) is based on the fact that as sparsity decreases, i.e. the number of clusters increases, the unpenalized log-likelihood of the data will increase. The increase in the number of clusters, then, is justifiable only by a significant increase in the unpenalized log-likelihood. We should choose a solution after which an increase in the number of clusters will not correspond to a large increase in the unpenalized log-likelihood. To determine this, we sort the solution path according to an increasing number of clusters K_{total} and, for each solution s , $s = 1, \dots, S$, we calculate the difference ratio for two adjacent solutions:

$$dr^{(s,s+1)} = \frac{L(K^{(s+1)}) - L(K^{(s)})}{K^{(s+1)} - K^{(s)}}, \quad (19)$$

provided that $K^{(s+1)} - K^{(s)} \geq 1$, where

$$L(K^{(s)}) = \sum_{i=1}^n \log \left[\sum_{k=1}^{K^{(s)}} \pi_k^{(s)} \phi(y_i; \mu_k^{(s)}, \Sigma_k^{(s)}) \right] \quad (20)$$

is the mixture Gaussian log-likelihood function and $K^{(s)}$ is the total number of estimated clusters K_{total} for that solution. The cluster centers $\mu_k^{(s)}$ and the cluster proportions $\pi_k^{(s)}$ are estimated given the cluster assignments for each solution s . We set the covariance $\Sigma_k^{(s)}$ to be the identity matrix, in line with the implicit assumption behind the use of ℓ_2 loss in (1) and (5). The likelihood for a singleton cluster in this case is well-defined and in effect amounts to $1/n$. We choose the solution indexed by

$$K^* = \max \left\{ K^{(s)} : dr^{(s,s+1)} \geq a \times \max \left(dr^{(1,2)}, \dots, dr^{(S-1,S)} \right) \right\}, \quad (21)$$

with $a = 0.05$ as suggested by Fu and Zhou (2013). We have plotted the difference ratio (19) and the unpenalized log-likelihood (20) for two simulated datasets in Figure 9. In Figure 9a the selected solution according to (21) is the solution with the true number of clusters $K^* = K_{\text{total}} = K_{\text{clust}} = 10$. An example from the scenario with noise in Figure 9b is more ambiguous and in this particular case we would barely choose a solution with $K_{\text{total}} = 214$ and $K_{\text{clust}} = 10$, which does not correspond to the highest ARI score and leaves 4 of the clustered points as noise. Generally, however, the selected solutions have relatively high ARI scores, as reported in Table 4. In this table, we demonstrate the averages for ARI_c and ARI_n scores for the selected solutions as well as the averages for the best ARI scores (in terms of the sum of the two ARI scores)

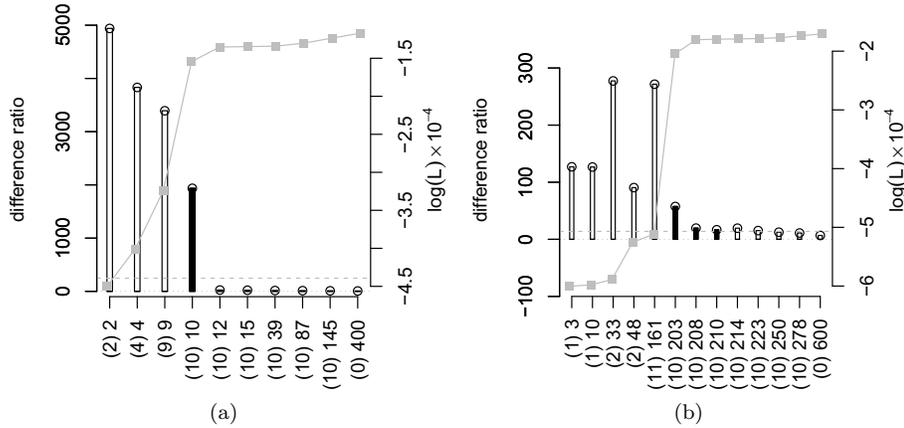


FIG 9. The difference ratio (vertical bars) and the unpenalized log-likelihood (gray solid line) as a function of the number of clusters in a solution path. The dashed line represents the cutoff with $\alpha = 0.05$ and the black color highlights the difference ratio for the solutions with the highest ARI scores. The number of clusters on the x-axis is represented by the total number of clusters K_{total} and the number of clusters of size $N_k > 3$ in brackets (K_{clust}). (a) An example from the $n > p$ scenario with well-separated clusters and no noise. (b) An example from the $n > p$ scenario with well-separated clusters and noise.

along each solution path of the corresponding scenario. One sees that the ARI scores of a selected solution are close to the best along the solution path. The table also reports the average selected K_{clust} and its range, showing that the number of clusters given by a selected solution is close to the true one. All averages and ranges are taken over 20 simulated datasets.

It is generally hard to determine whether singleton or very small clusters are truly noise, and thus, solution selection is more challenging when noisy data points are present, which is reflected in selected average ARI_n being slightly lower than the best ARI_n scores in Table 4. For overlapping scenario with no noise, we tended to select the solutions with $K_{clust} = 9$ clusters, while the best ARI scores were given by solutions with $K_{clust} = 10$ (with a few clustered data

TABLE 4
Solution selection summary for $n > p$ scenarios

	ARI_c		ARI_n		K_{clust}	
	select	best	select	best	mean	range
(1)	1.000	1.000	1.000	1.000	10	10-10
(2)	0.899	0.935	1.000	0.986	9	9-9
(3)	0.986	1.000	0.979	1.000	10	10-13
(4)	0.940	0.918	0.900	0.987	10	9-10

NOTE: (1) well-separated clusters, (2) overlapping clusters, (3) well-separated clusters with noise, (4) overlapping clusters with noise. ARI_c and ARI_n are averaged over 20 datasets for each scenario for the selected solutions (select) and for the largest ARI scores in each solution path (best).

points misclassified as noise). On the contrary, for the overlapping scenario with noise, most of the solutions selected were given by $K_{\text{clust}} = 10$, which resulted in a higher than best ARI_c but lower than best ARI_n . Altogether this method is very fast and appears to select good solutions for SPC, but more experiments are needed, especially for various possible assumptions on Σ_k .

6. Clustering gene expression data

To further test its performance in a real data environment we applied SPC to a gene expression dataset (Zhou et al., 2007). The full dataset consists of over 45,000 genes across 16 different experimental conditions that were created to study the regulation of these genes by Oct4, a transcription factor (TF) important for the self-renewal and maintenance of mouse embryonic stem cells. The 16 gene expression profiles include 3 generated from undifferentiated cells which are naturally high in Oct4 expression (conditions 1–3), 5 from early differentiated cells with high Oct4 expression (conditions 4–8), and 8 with low Oct4 expression (conditions 9–16). The study in Zhou et al. (2007) has identified and referenced 1,325 Oct4-high genes (Oct4+) and 1,440 Oct4-low genes (Oct4–) out of all the genes, which we regard as two true separated clusters for validation. We then added 3,000 extra randomly selected genes with various levels of coefficient of variation and randomly permuted the expression vector of each gene to obtain noise data points. Thus, we analyzed a final dataset with $n = 5,765$ observations (genes), $p = 16$ dimensions, and $K = 2$ distinct clusters of Oct4+ and Oct4– genes. Following common practice, we normalized the expression data of each gene to have zero mean and unit standard deviation.

We ran SPC with $\omega = 0.5$ and the resulting solution path is presented in Table 5. SPC has clearly identified the largest two clusters of Oct4+ and Oct4– genes which are shown in Figure 10a–10b. It also suggested an additional smaller cluster of size $N_k = 43$ shown in Figure 10c. We do not address in this paper the problem of choosing a proper threshold for the size of a cluster. For demonstration purposes we consider only $K = 2$ largest clusters and regard the rest

TABLE 5
Solution path for the gene expression data

	δ	λ	# of clusters of size			# of genes (Oct4+ / Oct4– / noise)		
			[1K, n]	[30, 1K]	(1, 30)	Cluster 1	Cluster 2	Other
1	0.3	6.2	2	1*	2954	1313/ 0 / 3	0/ 1392/ 1	12/ 48/ 2996
2	0.3	6.7	2	1*	2868	1313/ 0 / 3	0/ 1392/ 1	12/ 48/ 2996
3	0.3	7.3	2	1*	2617	1313/ 0 / 3	0/ 1392/ 1	12/ 48/ 2996
4	0.3	8.0	2	1*	2000	1313/ 0 / 3	0/ 1392/ 1	12/ 48/ 2996
5	0.3	8.7	3	1	837	1315/ 0 / 4	0/ 1435/ 9	10/ 5 / 2987
6	0.3	9.5	1	0	234	1325/ 1440/ 2658	0/ 0 / 8	0 / 0 / 334
7	0.3	10.4	1	0	80	1325/ 1440/ 2910	0/ 0 / 3	0 / 0 / 87
8	0.3	11.3	1	0	46	1325/ 1440/ 2954	0/ 0 / 1	0 / 0 / 45
9	0.3	12.4	1	0	4	1325/ 1440/ 2996	0/ 0 / 1	0 / 0 / 3
10	0.3	13.5	1	0	0	1325/ 1440/ 3000	0/ 0 / 0	0 / 0 / 0

NOTE: The asterisks * identify the small cluster of size $N_k = 43$ (Oct4–) in the solution path.

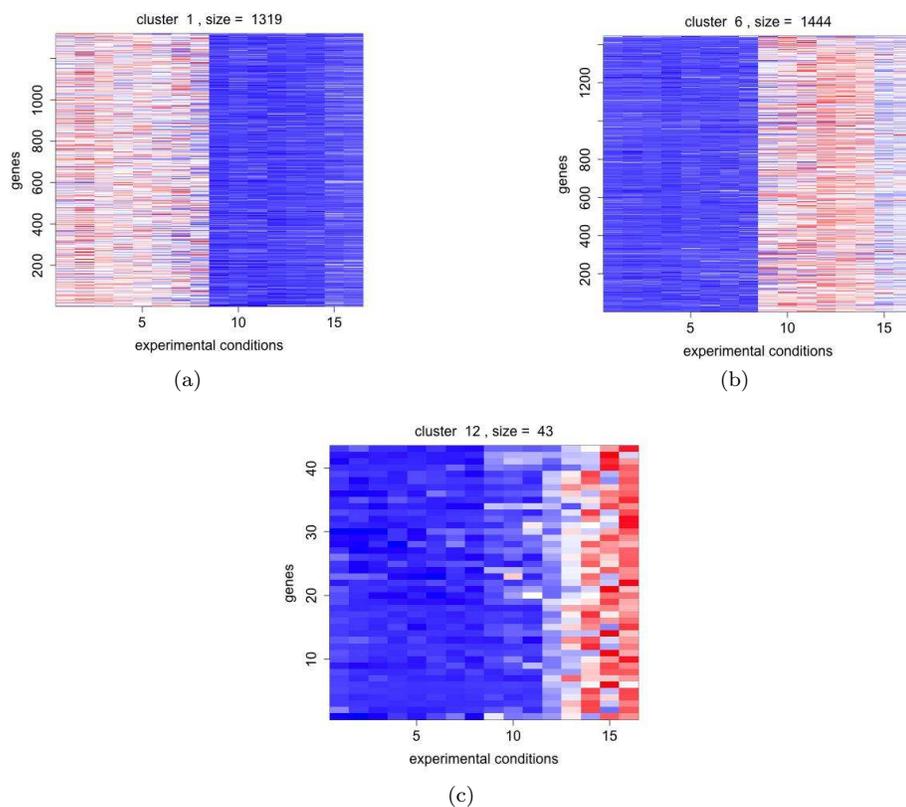


FIG 10. Heatmaps of the three largest clusters identified by SPC. Red color indicates high expression and blue color indicates low expression.

of the observations as noise, including the small cluster in Figure 10c. The left side of Table 5 shows the number of clusters of different size per each solution and the right side shows the numbers of Oct4+, Oct4-, and random genes in the two largest clusters (Cluster 1 and Cluster 2) and those in the rest of the genes (Other). The differences among the solutions 1–4 are only due to the random genes forming mostly very small clusters of size $N_k < 30$, and there are only 4 random genes assigned to the Oct4 clusters. It can be seen from the last column that a small number of Oct4 genes were left out as random, but most of these genes (43 out of 48 Oct4- genes) formed the small Oct4- cluster of Figure 10c. In solution 5 the small Oct4- cluster is merged with the main Oct4- cluster and another large cluster of size $> 1,000$ is formed from many random genes. In solution 6 the majority of the genes merge into a single cluster and in the subsequent solutions the remaining random genes (outliers) are gradually added to this cluster. All the solutions in Table 5 were obtained in fewer than 41 iterations.

TABLE 6
Comparison of different clustering methods for the Oct4-sorted expression data

	# of genes (Oct4+ / Oct4- / noise)														
	Cluster 1			Cluster 2			Other								
SPC	1315	/	0	/	3	0	/	1435	/	9	10	/	5	/	2987
mclust	1302	/	0	/	4	0	/	1412	/	5	23	/	28	/	2991
tight clust	1323	/	0	/	108	0	/	1402	/	64	2	/	38	/	2828
PWK	1250	/	0	/	11	0	/	1375	/	13	75	/	65	/	2976

Table 6 compares the result of SPC with that of mclust, tight clustering, and PWK-means. For best results, we ran tight clustering with $K = 2$ and $k_0 = K + 1$ since any larger values of k_0 produced inferior results, with large amounts of Oct4 genes added to the random category. For mclust, we overrode the default value for the hypervolume V and applied the calculation from Section 4.1. For the nearest neighbor cleaning $K = 5$ was used (the resulting partitioning into Oct4 clusters and random genes was not sensitive for $K < 40$, while if $K \geq 40$ was used, no noise would be identified in this data). Finally, for PWK-means we performed the penalty parameter λ search as with the simulated data. However, we had to start the sequence with a value greater than 1, otherwise the chosen value of λ would be less than 1 by prediction-based resampling and would result in most of the genes classified as random. Overall, we felt that all the other methods required more user fine tuning and rerunning for this dataset than SPC. It should be noted that the variation in the results of tight clustering across different runs was negligible, while the results of the other three methods were deterministic.

It can be seen from Table 6 that none of the methods misclassified genes between the Oct4+ and Oct4- sets in this well-separated data, however, generally SPC added the fewest random genes into the Oct4 clusters and left out the fewest Oct4 genes as random. SPC (solution 5 in Table 5) performed better compared to tight clustering, which tended to add more random genes into Oct4 clusters, and PWK-means, which left out more Oct4 genes as random. SPC identified a larger number of Oct4 genes compared to mclust. Even when comparing mclust to solutions 1–4 of SPC, it can be seen that SPC's overestimation of the number of random genes is attributed to the small Oct4- cluster in Figure 10c that can arguably be classified as a separate cluster. Indeed, compared with those genes in the big Oct4- cluster in Figure 10b, the genes in Figure 10c have lower expression levels in conditions 9 to 12 and higher expression levels in conditions 13 to 16, corresponding to cells starting differentiation.

Gene ontology (GO) term enrichment analysis further confirms the specific biological roles of the genes in this small cluster. Using the full set of the 1,440 Oct4- genes as the background set, we found 64 enriched GO terms from the process ontology with $P < 1.7 \times 10^{-4}$ and three from the function ontology with $P < 1 \times 10^{-4}$, both having a false discovery rate (FDR) of less than 0.005% as strong evidence that these genes are not clustered by chance. From Table 7 we see that this small cluster is highly enriched for genes involved in nerve development and organismal process and for TFs with sequence-specific

TABLE 7
 Top GO terms enriched in the small Oct4- cluster

GO term	% in cluster	% in full set	P-value
nervous system development	43.6	12.9	1.41×10^{-6}
multicellular organismal process	71.8	34.8	1.99×10^{-6}
cranial nerve development	10.3	0.3	2.27×10^{-6}
sequence-specific DNA binding	25.6	5.6	2.76×10^{-5}
sequence-specific DNA binding TF activity	28.2	6.8	2.86×10^{-5}
nucleic acid binding TF activity	28.2	6.9	3.16×10^{-5}

Tabulated are the top three GO terms from the process (top) and the function (bottom) ontology.

DNA binding activity, compared to the corresponding proportions in the full Oct4- set. This analysis demonstrates a strong performance of SPC on a larger gene expression dataset and the possibility of discovering smaller, biologically meaningful clusters by further examining the solution path. Moreover, this shows an example where a detailed solution path is useful even when strong prior knowledge for the number of clusters is available. With the expected number of clusters $K = 2$, all the other methods failed to discover this small cluster. The results of mclust, tight clustering and PWK-means identified solely the two big clusters and only differed in how many random genes were added to the clusters and how many Oct4 genes were left out as random.

7. Discussion

SPC is a new general clustering method that provides a small set of solutions, each including a cluster assignment and a number of clusters, in a wide variety of situations including high-dimensional settings when model-based clustering, for example, may have difficulty in providing a solution. SPC can be applied to datasets of different complexities, whether they are small with only a few outliers or large with a high proportion of irrelevant observations. This method searches for the best solution given a certain degree of sparsity penalty, which is determined adaptively by the data. The irrelevant observations are simultaneously identified as singletons or very small clusters. SPC does not require extensive fine tuning and has only one required input tuning parameter, which could have some impact on the resulting set of solutions. However, even this parameter can be set at the lowest possible value, with the only drawback of a longer solution path and running time. Most importantly, SPC does not require the knowledge of the number of clusters and provides its own estimate, depending on the amount of sparsity imposed.

Solution path clustering is effective in separating noise from clustered data, however, it tends to merge overlapping clusters due to its dependence on distances between cluster centers. In the later stages of a solution path, when sparsity is higher, it also tends to add more noise into clusters as the larger distances, usually characterized by noisy observations or outliers, receive more

penalization. Overall, SPC have provided satisfactory clustering results across all the simulated data scenarios and the gene expression data.

Penalized estimation methods have been widely applied in linear model analysis and are emerging in clustering. The idea of a solution path or surface along with model selection methodology has been successfully utilized for fast search of sparse models. We extended the regularization path or surface and sparsity approach to unsupervised learning, where the gradual increase of sparsity is guided by the data and the bias of the cluster centers is explicitly adjusted to avoid bad solutions. All these features make SPC different from the majority of the existing penalization frameworks, which either encourage sparsity in cluster means or use a convex penalty for sparsity.

There are several questions that need to be addressed by future research. First, we have not discussed what can be considered a cluster and what can be considered noise. If a cluster consists of just one observation, it could safely be assumed that this singleton cluster is an outlier, however, a decision needs to be made, based on the data, with respect to clusters of size $N_k > 1$. This is easier to determine with smaller datasets, where the cut off $N_k > 1$ is usually sufficient, however, it will not be clear with larger and more complex data that might need higher cluster size cutoff values. We would like to address this issue in the future in a more principled manner. In this paper, we also only cursorily discuss the methods for selecting the best candidate solution from a solution path and we do not address the asymptotic properties of the cluster center estimates. These and other algorithmic improvements are interesting topics for future work.

Appendix

A.1. Proof of Lemma 1

Fix $\theta_2^{(t)}$ and one minimization step of the MM algorithm as in (8) gives

$$\theta_1^{(t+1)} - \theta_2^{(t)} = \frac{y_1 + \lambda w_{1,2}^{(t)} \theta_2^{(t)}}{1 + \lambda w_{1,2}^{(t)}} - \theta_2^{(t)}.$$

Equating this to $(1 - \phi)(\theta_1^{(t)} - \theta_2^{(t)})$ and simply re-arranging the terms, we get

$$\lambda = \frac{2\eta\|(\theta_1^{(t)} - y_1) + \phi(\theta_2^{(t)} - \theta_1^{(t)})\|_2}{\left(\eta - \|\theta_1^{(t)} - \theta_2^{(t)}\|_2\right)_+ (1 - \phi)}.$$

For the initial step of the MM algorithm $\theta_1^{(0)} = y_1$ and $\theta_2^{(0)} = y_2$, and thus

$$\lambda = \frac{2\phi\eta\|y_1 - y_2\|_2}{(1 - \phi)(\eta - \|y_1 - y_2\|_2)},$$

if $\eta > d = \|y_1 - y_2\|_2$. If $\eta \leq d$, then $w_{1,2}^{(t-1)} = 0$ and thus $\theta_i^{(t)} = y_i$ for all $t \geq 1$.

A.2. Proof of Lemma 2

Let $\gamma = \theta_2 - \theta_1$. Recall that $\ell(\theta_1, \theta_2)$ is minimized at $(\bar{y} - \gamma/2, \bar{y} + \gamma/2)$ for any γ and thus, minimizing (3) reduces to minimizing $\ell(\gamma)$ defined in (4). First assume that $\gamma = c(y_2 - y_1)$ for $c \in \mathbb{R}$ and define

$$f(c) \triangleq \ell(c(y_2 - y_1)) = \frac{1}{2}d^2(c-1)^2 + \lambda\rho(d \cdot |c|),$$

where $d = \|y_2 - y_1\|_2$. It is easy to see that $f(c) > f(0)$ for all $c < 0$ and $f(c) > f(1)$ for all $c > 1$. Therefore, f must be minimized globally at $c \in [0, 1]$. Then a sufficient condition for $f(c)$ to be minimized at $c = 0$, which gives $\gamma = 0$, is

$$\frac{\partial f(c)}{\partial c} = d^2(c-1) + \left(\lambda - \frac{dc}{\delta}\right)_+ d > 0,$$

for $0 < c \leq 1$. If (14) holds, the above inequality is equivalent to

$$\lambda > \left[(1-c) + \frac{c}{\delta}\right] d,$$

which is obviously implied by (14), noting that $c \in (0, 1]$. Lastly, for any nonzero γ' that is not collinear with $y_2 - y_1$, there is a $\gamma = c(y_2 - y_1)$ such that $\|\gamma\|_2 = \|\gamma'\|_2$ and $\|\gamma - (y_2 - y_1)\|_2 < \|\gamma' - (y_2 - y_1)\|_2$ by triangle inequality, and thus $\ell(\gamma) < \ell(\gamma')$. Therefore, the global minimizer of (3) is $(\hat{\theta}_1, \hat{\theta}_2) = (\bar{y}, \bar{y})$ if (14) holds.

A.3. ARI

Let $h(k) = \binom{k}{2}$. Given a contingency table $(n_{ij})_{I \times J}$ with entries n_{ij} , row sums $n_{i\bullet} = \sum_j n_{ij}$, column sums $n_{\bullet j} = \sum_i n_{ij}$, and a total sum of entries $n = \sum_{i,j} n_{ij}$, the ARI is defined by

$$\text{ARI} = \frac{\sum_{i=1}^I \sum_{j=1}^J h(n_{ij}) - \left[\sum_{i=1}^I h(n_{i\bullet}) \sum_{j=1}^J h(n_{\bullet j}) \right] / h(n)}{\frac{1}{2} \left[\sum_{i=1}^I h(n_{i\bullet}) + \sum_{j=1}^J h(n_{\bullet j}) \right] - \left[\sum_{i=1}^I h(n_{i\bullet}) \sum_{j=1}^J h(n_{\bullet j}) \right] / h(n)}. \quad (22)$$

References

- AGGARWAL, C. C. and REDDY, C. K. (2013). *Data Clustering: Algorithms and Applications*. Chapman & Hall/CRC Data Mining and Knowledge Discovery Series. Chapman and Hall/CRC.
- ARTHUR, D. and VASSILVITSKII, S. (2007). K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* 1027–1035. Society for Industrial and Applied Mathematics. [MR2485254](#)

- BEN-HUR, A., HORN, D., SIEGELMANN, H. T. and VAPNIK, V. (2001). Support vector clustering. *J. Mach. Learn. Res.* **2** 125–137.
- BENSMAIL, H., CELEUX, G., RAFTERY, A. E. and ROBERT, C. (1997). Inference in model-based cluster analysis. *Statist. Comput.* **7** 1–10.
- BYERS, S. and RAFTERY, A. E. (1998). Nearest-neighbor clutter removal for estimating features in spatial point processes. *J. Amer. Statist. Assoc.* **93** 577–584.
- CHI, E. C. and LANGE, K. (2013). Splitting methods for convex clustering. Preprint, available at [arXiv:1304.0499](https://arxiv.org/abs/1304.0499) [stat.ML].
- DE SMET, F., MATHYS, J., MACHAL, K., THIJIS, G., DE MOOR, B. and MOREAU, Y. (2002). Adaptive quality-based clustering of gene expression profiles. *Bioinformatics* **18** 735–746.
- EFRON, B., HASTIE, T., JOHNSTONE, I. and TIBSHIRANI, R. (2004). Least angle regression (with discussion). *Ann. Statist.* **32** 407–499. [MR2060166](#)
- FAN, J. and LI, R. (2001). Variable selection via non-concave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- FANG, Y. and WANG, J. (2012). Selection of the number of clusters via the bootstrap method. *Comput. Stat. Data Anal.* **56** 468–477. [MR2853747](#)
- FORERO, P., KEKATOS, V. and GIANNAKIS, G. B. (2012). Robust clustering using outlier-sparsity regularization. *IEEE Transactions on Signal Processing* **60** 4163–4177. [MR2960487](#)
- FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.* **97** 611–631. [MR1951635](#)
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Statist. Softw.* **33** 1–22.
- FRIEDMAN, J., HASTIE, T., HOFLING, H. and TIBSHIRANI, R. (2007). Pathwise coordinate optimization. *Ann. Appl. Statist.* **1** 302–332. [MR2415737](#)
- FU, F. and ZHOU, Q. (2013). Learning sparse causal Gaussian networks with experimental intervention: regularization and coordinate descent. *J. Amer. Statist. Assoc.* **108** 288–300. [MR3174620](#)
- GARCIA-ESCUADERO, L. A., GORDALIZA, A., MATRAN, C. and MAYO-ISCAR, A. (2008). A general trimming approach to robust cluster analysis. *Ann. Statist.* **36** 1324–1345. [MR2418659](#)
- GUO, J., LEVINA, E., MICHAILIDIS, G. and ZHU, J. (2010). Pairwise variable selection for high-dimensional model-based clustering. *Biometrics* **66** 793–804. [MR2758215](#)
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2 ed., Springer. [MR2722294](#)
- HASTIE, T., TIBSHIRANI, R., EISEN, M. B., ALIZADEH, A., LEVY, R., STAUDT, L., CHAN, W., BOTSTEIN, D. and BROWN, P. (2000). Gene shaving as a method for identifying distinct sets of genes with similar expression patterns. *Genome Biol.* **1** 0003.1–0003.21.

- HOCKING, T. D., JOULIN, A., BACH, F. and VERT, J.-P. (2011). Clusterpath: an algorithm for clustering using convex fusion penalties. In *28th International Conference on Machine Learning*.
- HUBERT, J. and ARABIE, P. (1985). Comparing partitions. *J. Classif.* **2** 193–218.
- JAIN, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Lett.* **31** 651–666.
- KOHONEN, T. (1990). The self-organizing map. *Proceedings of the IEEE* **78** 1464–1479.
- KULIS, B. and MICHAEL, I. J. (2011). Revisiting k-means: new algorithms via bayesian nonparametrics. Preprint, available at [arXiv:1111.0352](https://arxiv.org/abs/1111.0352).
- LANGE, K. (1995). A gradient algorithm locally equivalent to the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **57** 425–437. [MR1323348](#)
- LANGE, K. (2004). *Optimization*, 1 ed., Springer, New York, NY. [MR2072899](#)
- LANGE, K., HUNTER, D. R. and YANG, I. (2000). Optimization transfer using surrogate objective functions. *J. Comput. Graph. Statist.* **9** 1–20. [MR1819865](#)
- LINDSTEN, F., OHLSSON, H. and LJUNG, L. (2011). Just relax and come clustering! a convexication of k-means clustering Technical Report, Linköping University, Department of Electrical Engineering, Automatic Control.
- MAITRA, R. and RAMLER, I. P. (2009). Clustering in the presence of scatter. *Biometrics* **65** 341–352. [MR2751457](#)
- MAZUMDER, R., FRIEDMAN, J. and HASTIE, T. (2011). SparseNet: coordinate descent with non-convex penalties. *J. Amer. Statist. Assoc.* **106** 1125–1138. [MR2894769](#)
- MCLACHLAN, G. J., BEAN, R. W. and PEEL, D. (2002). A mixture model-based approach to the clustering of microarray expression data. *Bioinformatics* **18** 413–422.
- OH, M. S. and RAFTERY, A. E. (2007). Model-based clustering with dissimilarities: a bayesian approach. *J. Comput. Graph. Statist.* **16** 559–585. [MR2351080](#)
- PAN, W. and SHEN, X. (2007). Penalized model-based clustering with application to variable selection. *J. Mach. Learn. Res.* **8** 1145–1164.
- PAN, W., SHEN, X. and LIU, B. (2013). Cluster analysis: unsupervised learning via supervised learning with a non-convex penalty. *J. Mach. Learn. Res.* **14** 1865–1889. [MR3104498](#)
- PELCKMANS, K., DE BRABANTER, J., SUYKENS, J. and DE MOOR, B. (2005). Convex clustering shrinkage. In *PASCAL Workshop on Statistics and Optimization of Clustering Workshop*.
- SHARAN, R., MARON-KATZ, A. and SHAMIR, R. (2003). CLICK and EXPANDER: a system for clustering and visualizing gene expression data. *Bioinformatics* **19** 1787–1799.
- SHEN, Y., SUN, W. and LI, K. C. (2010). Dynamically weighted clustering with noise set. *Bioinformatics* **26** 341–347.
- SOLTANOLKOTABI, M., ELHAMIFAR, E. and CANDES, E. J. (2013). Robust subspace clustering. Preprint, available at [arXiv:1301.2603](https://arxiv.org/abs/1301.2603). [MR3210983](#)

- SUN, W. and WANG, J. (2012). Regularized k-means clustering of high-dimensional data and its asymptotic consistency. *Elec. J. Stat.* **6** 148–167. [MR2879675](#)
- THALAMUTHU, A., MUKHOPADHYAY, I., ZHENG, X. and TSENG, G. C. (2006). Evaluation and comparison of gene clustering methods in microarray analysis. *Bioinformatics* **22** 2405–2412.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- TIBSHIRANI, R., WALTHER, G. and HASTIE, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *J. Roy. Statist. Soc. Ser. B* **63** 411–423. [MR1841503](#)
- TIBSHIRANI, R. and WALTHER, G. (2005). Cluster validation by prediction strength. *J. Comput. Graph. Statist.* **14** 511–528. [MR2170199](#)
- TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. and KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *J. Roy. Statist. Soc. Ser. B* **67** 91–108. [MR2136641](#)
- TSENG, G. C. (2007). Penalized and weighted k-means for clustering with scattered objects and prior information in high-throughput biological data. *Bioinformatics* **23** 2247–2255.
- TSENG, G. C. and WONG, W. H. (2005). Tight clustering: a resampling-based approach for identifying stable and tight patterns in data. *Biometrics* **61** 10–16. [MR2129196](#)
- TSENG, P. and YUN, S. (2009). A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming* **117** 387–423. [MR2421312](#)
- VON LUXBURG, U. (2007). A tutorial on spectral clustering. *Stat. Comput.* **17** 395–416. [MR2409803](#)
- WANG, H., NEILL, J. and MILLER, F. (2008). Nonparametric clustering of functional data. *Stat. Interface* **1** 47–62. [MR2425344](#)
- WANG, S. and ZHU, J. (2008). Variable selection for model-based high-dimensional clustering and its applications to microarray data. *Biometrics* **64** 440–448. [MR2432414](#)
- WITTEN, D. and TIBSHIRANI, R. (2010). A framework for feature selection in clustering. *J. Amer. Statist. Assoc.* **105** 713–726. [MR2724855](#)
- WU, T. T. and LANGE, K. (2008). Coordinate descent algorithms for lasso penalized regression. *Ann. Appl. Statist.* **2** 224–244. [MR2415601](#)
- XIE, B., PAN, W. and SHEN, X. (2008). Variable selection in penalized model-based clustering via regularization on grouped parameters. *Biometrics* **64** 921–930. [MR2526644](#)
- YEUNG, K. Y., FRALEY, C., MURUA, A., RAFTERY, A. E. and RUZZO, W. L. (2001). Model-based clustering and data transformations for gene expression data. *Bioinformatics* **17** 977–987.
- YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *J. Roy. Statist. Soc. Ser. B* **68** 49–67. [MR2212574](#)
- ZHANG, C. H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. [MR2604701](#)

- ZHOU, H., PAN, W. and SHEN, X. (2009). Penalized model-based clustering with unconstrained covariance matrices. *Elec. J. Stat.* **3** 1473–1496. [MR2578834](#)
- ZHOU, Q., CHIPPERFIELD, H., MELTON, D. A. and WONG, W. H. (2007). A gene regulatory network in mouse embryonic stem cells. *PNAS* **104** 16438–16443.
- ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. Roy. Statist. Soc. Ser. B* **67** 301–320. [MR2137327](#)