

RESEARCH ARTICLE

Constrained Clustering: General Pairwise and Cardinality Constraints

ADEL BIBI¹, ALI ALQAHTANI^{2,3}, AND BERNARD GHANEM⁴¹Department of Engineering Science, University of Oxford, OX1 3PJ Oxford, U.K.²Department of Computer Science, King Khalid University, Abha 61421, Saudi Arabia³Center for Artificial Intelligence (CAI), King Khalid University, Abha 61421, Saudi Arabia⁴Visual Computing Center, King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia

Corresponding author: Adel Bibi (adel.bibi@eng.ox.ac.uk)


This work was supported in part by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research; and in part by the Deanship of Scientific Research, King Khalid University, Saudi Arabia, under Grant RGP1/357/43.

ABSTRACT In this work, we study constrained clustering, where constraints are utilized to guide the clustering process. In existing works, two categories of constraints have been widely explored, namely pairwise and cardinality constraints. Pairwise constraints enforce the cluster labels of two instances to be the same (must-link constraints) or different (cannot-link constraints). Cardinality constraints encourage cluster sizes to satisfy a user-specified distribution. However, most existing constrained clustering models can only utilize one category of constraints at a time. In this paper, we enforce the above two categories into a unified clustering model starting with the integer program formulation of the standard K-means. As these two categories provide useful information at different levels, utilizing both of them is expected to allow for better clustering performance. However, the optimization is difficult due to the binary and quadratic constraints in the proposed unified formulation. To alleviate this difficulty, we utilize two techniques: equivalently replacing the binary constraints by the intersection of two continuous constraints; the other is transforming the quadratic constraints into bi-linear constraints by introducing extra variables. Then we derive an equivalent continuous reformulation with simple constraints, which can be efficiently solved by Alternating Direction Method of Multipliers (ADMM) algorithm. Extensive experiments on both synthetic and real data demonstrate: 1) when utilizing a single category of constraint, the proposed model is superior to or competitive with state-of-the-art constrained clustering models, and 2) when utilizing both categories of constraints jointly, the proposed model shows better performance than the case of the single category. The experimental results show that the proposed method exploits the constraints to achieve perfect clustering performance with improved clustering to 2 – 5% in classical clustering metrics, e.g. Adjusted Random Index (ARI), Mirkin's Index (MI), and Huber's Index (HI), outperforming all compared-against methods across the board. Moreover, we show that our method is robust to initialization.

INDEX TERMS Constrained clustering, K-means, pairwise, cardinality constraints.

I. INTRODUCTION

Clustering is the task of partitioning data into different clusters, based on some specific cluster assumptions. For example, K-means and Gaussian mixture models (GMM) assume each cluster is sampled from a Gaussian distribution. In contrast, density-based clustering assumes that the densities of data points in different clusters should be

The associate editor coordinating the review of this manuscript and approving it for publication was Liangxiu Han .

different, such as Chameleon [1] and AITC [2], or clusters should be partitioned at low density regions [3]. However, if the adopted cluster assumption is not suited to the target dataset, this may result in a poor performance. To avoid such performance instability, prior knowledge or constraints on the data can be used to guide the clustering process. These constraints are independent of cluster assumptions, and they provide weak supervision to reflect user preferences. Thus, clustering with constraints, called *constrained clustering*, [4], [5], [6], is expected to give

better and more stable performance than unconstrained clustering.

Two main categories of constraints have been widely studied in the field of constrained clustering, namely pairwise and cardinality constraints. *Pairwise constraints* may arise from some form of perceived similarity between samples. For instance, the continuity property is a form of *Pairwise constraints* that suggests that neighbouring samples are likely to be clustered together and vice versa. Thus, *Pairwise constraints* include must-link and cannot-link constraints. Must-link constraints enforce that a set of pairs of instances should be in the same cluster, while cannot-link constraints enforce that they belong to different clusters. Thereafter, this category can be viewed as instance-level constraints. On the other hand, *Cardinality constraints* provide extra knowledge on the size distribution of all clusters. This sort of constraints become particularly necessary in clustering tasks of data that is high dimensional and sparse with many clusters to assign [7]. This often leads to solutions of empty clusters or unbalanced cluster assignments. Balancing constraints that lead to equal sized clusters are only a special case of *Cardinality constraints*. This category in general can be viewed as cluster-level constraints.

Many clustering methods have been proposed to utilize one of the two categories of constraints, such as the ones with pairwise constraints [4], [5], [6], [8], [9], [10], [11], [12], [13], and the ones with cardinality constraints [7], [14], [15], [16]. However, in some cases, one might want to enforce the continuity property among a set of points and the same time requiring to have solutions of balanced or user specified cluster sizes. In general, both constraints can be provided simultaneously, as they are derived from different sources. For example, pairwise constraints are usually obtained from an oracle query, while cardinality constraints can be obtained from experience or user preference. Moreover, they represent supervision at different levels. Each of them can provide particularly useful information that is not covered by the other. Thus, having both sets of constraints together in a clustering task should significantly improve the performance and to the best of our knowledge, there is no existing work that can seamlessly incorporate both categories jointly.

For existing constrained clustering methods that handle one constraint category, it is non-trivial to directly add the other. For example, embedding cardinality constraints into the COP-KMEANS [13], will lead to instability in performance where COP-KMEANS will often fail in finding a feasible solution. This is because COP-KMEANS is very sensitive to cluster initialization. Moreover, it is also not easy to embed the pairwise constraints into normalized/ratio-cut [16], which exploits balanced distribution constraints. In short, existing models are designed to exploit one category of constraints at a time.

We propose a unified model to incorporate both categories of constraints to guide the clustering process. Specifically, we start from the formulation of the standard K-mean method,

and formulate cardinality constraints into linear constraints and pairwise constraints into quadratic constraints. Then we obtain a discrete optimization problem with quadratic constraints, which is difficult to be solved by off-the-shelf optimization methods. Thus we propose to utilize two techniques. One is to equivalently replace the binary constraints by the intersection of two continuous constraints, which was firstly proposed in [17]. The other is to transform the quadratic constraints to bi-linear constraints by introducing extra variables. Our key contributions revolve around the new novel continuous reformulation for the K-Means problem allowing to incorporate both cardinality and pair-wise constraints. The reformulation is simple, flexible, and enjoys nice convergence properties with competitive performance. The contributions can be summarized in three folds.

- We embed both pairwise and cardinality constraints into one unified clustering model. To the best of our knowledge, this is the first attempt in the field of constrained clustering.
- We propose to equivalently transform the binary and quadratic constraints in the original problem to continuous and bi-linear constraints, to obtain a simple continuous reformulation.
- We conduct various experiments on several synthetic and real datasets comparing against 5 different algorithms. Our approach demonstrate competitive edge over all compared-against methods in the final clustering performance while respecting the imposed constraints.

II. RELATED WORK

Here, we briefly review existing clustering models that utilize pairwise or cardinality constraints.

A. PAIRWISE CONSTRAINTS

They were first introduced into clustering in [4] and [13]. In [4], a method called COP-COBWEB inserted the pairwise constraints into the clustering process of the incremental clustering method COBWEB [18], which utilizes four operators (i.e. add, new, merge, and split) to maximize the intra-cluster similarity and the inter-cluster dissimilarity. In each operator of COP-COBWEB, the given pairwise constraints are checked to ensure the satisfaction of all constraints. In [13], the method COP-KMEANS checks the pairwise constraints in each assignment step of K-means. Both COP-COBWEB and COP-KMEANS treat the pairwise constraints as hard constraints (i.e. all constraints must be satisfied), and the constraints are somewhat independent of the original objective. A common limitation of these two methods is that the processing order of instances influences clustering performance, and sometimes they may even fail to output a feasible partition.

To avoid this limitation, many methods treat pairwise constraints as soft constraints (i.e. a subset of these constraints could be violated) to develop more flexible approaches to embed constraints. For example, in constrained

complete-link (CCL) [5], pairwise constraints are used to modify the instance proximity computed in the original feature space. Then, standard complete-link clustering is applied using the modified proximity matrix. Penalized probabilistic clustering (PPC) [10] uses pairwise constraints as a prior term w.r.t. the cluster labels within the underlying GMM-based model. Clustering configurations not satisfying the constraints have a lower probability. Moreover, HMRF-KMEANS [19] embeds pairwise constraints as correlations between cluster labels in a hidden Markov random field (HMRF). A metric learning step is added into standard K-means to encourage gradual satisfaction of pairwise constraints. Other methods propagate pairwise constraints via instance similarity to obtain soft constraints, such as constrained spectral clustering [8] and HMRF-pc [12], [20].

B. CARDINALITY CONSTRAINTS

They are widely used to guide the clustering process. Balancing constraints are a special type that encourages all clusters to be balanced in size or in connecting weights. For example, normalized cuts [16] divides the standard cut cost (sum of edge weights connecting the two clusters) by the sum of edge weights between each cluster and all other instances. Hence, each cluster is encouraged to have similar edge weights connecting to other clusters. Similarly, ratio cut [21] normalizes the cut function by the size of each cluster to encourage similar sized clusters. Equi-sized Fuzzy c-means (FCM) [15] formulates the balancing constraints as equality constraints, where the size of each cluster equals to the average cluster size. More general cardinality constraints have also been explored. For example, a constrained K-means method [7] sets a lower bound on the cluster size, to avoid very small or empty clusters that occur in standard K-means. An extension of the Equi-sized FCM is proposed in [14], where the size of one single cluster is set to a specific size.

To the best of our knowledge, the only clustering frameworks that enable both sets of constraints (Cardinality and Pairwise) either target a very specific class of methods that suffer from the locality property [22], or are greedy heuristics that propagate constraints [23]. Clustering methods that suffer from the locality property result in clusters located partially or entirely outside the Voronoi cell of their centers [22]. However, popular methods like K-means, K-medians, and many others always satisfy the locality property by definition, thus, limiting the theoretical results of [22] to a smaller class of clustering methods. There has also been an attempt to use standard constraint propagation methods to enforce both classes of constraints [23]. However, this is done in a greedy heuristic fashion that may often fail in finding a feasible solution. Therefore, we believe that the combination of both pairwise and cardinality categories into a constraint generic and unified clustering model that can be systematically solved (ie using a flexible continuous optimization framework) has not been explored in any existing work.

III. PROPOSED METHOD

Unlike previous methods that can only handle either pairwise or cardinality constraints, we show, in this section, a detailed derivation of our framework that embeds both constraints simultaneously. In fact, this formulation is flexible and generic enough to handle any other linear equality. In our framework, we adopt the K-means integer program formulation [24] expressed as follows:

$$\begin{aligned} \min_{\{x_{ij}\}_{i=1, j=1}^{n,k}} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \left\| \mathbf{s}_i - \frac{\sum_{p=1}^n x_{pj} \mathbf{s}_p}{\sum_{l=1}^n x_{lj}} \right\|_2^2 \\ \text{s.t.} & \sum_{j=1}^k x_{ij} = 1, \quad \forall i \quad x_{ij} \in \{0, 1\} \quad \forall i, j \end{aligned} \quad (1)$$

where $\mathbf{s}_p \in \mathbb{R}^d$ is the p^{th} data point to be clustered and k is the number of clusters. The variable x_{ij} defines the binary association between data point i and cluster j . The constraint $\sum_{j=1}^k x_{ij} = 1$ enforces data point i to belong to one and only one cluster. This constraint can be simply written as a matrix vector multiplication: $\Psi^T \mathbf{x} = \mathbf{1}_n$, where $\Psi^T \in \mathbb{R}^{n \times nk}$ is a binary matrix that has in each row a vector $\mathbf{1}_k^T$ that sums all the binary labels for a given data point while the rest are 0.

To simplify the fractional objective, we introduce variable w_{pj} , such that: $x_{pj} = w_{pj} \sum_{l=1}^n x_{lj}$. For ease of notation, we concatenate all the binary labels x_{ij} into one vector ordered by the data points one at a time as follows: $\mathbf{x}^T = [(x_{11} \dots x_{1k}) \dots (x_{n1} \dots x_{nk})]^T$. We also concatenate and reorder the w_{pj} values one cluster at a time: $\mathbf{w}^T = [(w_{11} \dots w_{n1}) \dots (w_{1k} \dots w_{nk})]^T$. A matrix $\mathbf{P} \in \mathbb{R}^{nk \times nk}$ is used to swap the order of the binary vectors from a cluster based order to a data point order and vice versa. Note that \mathbf{P} is a proper permutation matrix that is symmetric and it satisfies: $\mathbf{P}\mathbf{P}^T = \mathbf{I}_{nk}$. Thus, the compact form of unconstrained K-means can be re-written as follows:

$$\begin{aligned} \min_{\{x_{ij}\}_{i=1, j=1}^{n,k}, \{w_{pj}\}_{p=1, j=1}^{n,k}} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \left\| \mathbf{s}_i - \sum_{p=1}^n w_{pj} \mathbf{s}_p \right\|_2^2 \\ \text{s.t.} & \Psi^T \mathbf{x} = \mathbf{1}_n, \quad \mathbf{x} = \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x}, \quad \mathbf{x} \in \{0, 1\}^{nk} \end{aligned} \quad (2)$$

where $\mathbf{C} \in \mathbb{R}^{nk \times nk}$ sums the binary labels of each cluster and is defined as follows:

$$\begin{aligned} \mathbf{C}^T &= [\gamma_1 \ \gamma_2 \ \dots \ \gamma_n] \\ \gamma_1^T &= [1 \ \mathbf{0}_{k-1}^T \ 1 \ \dots], \quad \gamma_2^T = [0 \ 1 \ \mathbf{0}_{k-2}^T \ \dots] \end{aligned}$$

A. CARDINALITY CONSTRAINTS

They are enforced by a set of linear constraints that specify the cluster size as:

$$\sum_{i=1}^n x_{ij} = u_j \quad \forall j \Leftrightarrow \mathbf{Q}\mathbf{P}^T \mathbf{x} = \mathbf{u} \quad (3)$$

where u_j is the size of cluster j and $\mathbf{Q} \in \mathbb{R}^{k \times nk}$ sums the binary labels of each cluster for all data points.

B. MUST-LINK CONSTRAINTS

We define $\mathbf{E}_1, \mathbf{E}_2 \in \mathbb{R}^{kv \times nk}$ as selection matrices that choose the two sets of data points ($\mathbf{E}_1 \mathbf{x}$ and $\mathbf{E}_2 \mathbf{x}$) involved in the v must-link constraints. We show next that the set of all must-link constraints can be expressed with a single quadratic.

$$\mathbf{x}^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} = v \quad (4)$$

Proposition 1: For the binary association $\mathbf{x} \in \{0, 1\}^{nk}$ between n data points and k clusters, where $\Psi^\top \mathbf{x} = \mathbf{1}_n$, enforcing must-link constraints through $\mathbf{E}_1 \mathbf{x} = \mathbf{E}_2 \mathbf{x}$ is equivalent to enforcing a single quadratic $\mathbf{x}^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} = v$.

Proof:

$$\begin{aligned} \mathbf{E}_1 \mathbf{x} = \mathbf{E}_2 \mathbf{x} &\Leftrightarrow \|\mathbf{E}_1 \mathbf{x} - \mathbf{E}_2 \mathbf{x}\|_2^2 = 0 \\ &= \|\mathbf{E}_1 \mathbf{x}\|_2^2 + \|\mathbf{E}_2 \mathbf{x}\|_2^2 - 2\mathbf{x}^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} = 0 \\ &= 2v - 2\mathbf{x}^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} = 0 \quad \square \end{aligned}$$

The last equality ($\|\mathbf{E}_1 \mathbf{x}\|_2^2 = \|\mathbf{E}_2 \mathbf{x}\|_2^2 = v$) is true since \mathbf{x} is binary and that each data point is associated to only one cluster (i.e. $\Psi^\top \mathbf{x} = \mathbf{1}_n$).

This concludes that only one quadratic constraint can be used to enforce all must-link constraints.

C. CANNOT-LINK CONSTRAINTS

We define $\mathbf{E}_3, \mathbf{E}_4 \in \mathbb{R}^{ke \times nk}$ as selection matrices for the two sets of data points ($\mathbf{E}_3 \mathbf{x}$ and $\mathbf{E}_4 \mathbf{x}$) involved in the e cannot-link constraints. Similar to before, we show that the set of cannot-link constraints can be expressed with a single quadratic.

$$\mathbf{x}^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} = 0 \quad (5)$$

Proposition 2: For the binary association $\mathbf{x} \in \{0, 1\}^{nk}$ between n data points and k clusters, where $\Psi^\top \mathbf{x} = \mathbf{1}_n$, enforcing cannot-link constraints through the selection matrices $\mathbf{E}_3, \mathbf{E}_4$ is equivalent to enforcing the quadratic $\mathbf{x}^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} = 0$.

Proof:

Similar to Proposition 1 with $\|\mathbf{E}_3 \mathbf{x} + \mathbf{E}_4 \mathbf{x}\|_2^2 = 2e \quad \square$

Incorporating Eqs (3), (4) and (5) into (2) we obtain the following constrained K-means formulation:

$$\begin{aligned} \min_{\{x_{ij}\}_{i=1, j=1}^{n, k}, \{w_{pj}\}_{p=1, j=1}^{n, k}} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \left\| \mathbf{s}_i - \mathbf{S} \Lambda_j \mathbf{w} \right\|_2^2 \\ \text{s.t.} & \Psi^\top \mathbf{x} = \mathbf{1}_n, \quad \mathbf{x} \in \{0, 1\}^{nk}, \\ & \mathbf{Q} \mathbf{P}^\top \mathbf{x} = \mathbf{u} \\ & \mathbf{x} = \mathbf{P} \mathbf{w} \odot \mathbf{C} \mathbf{x}, \quad (\mathbf{E}_1 \mathbf{x})^\top \mathbf{E}_2 \mathbf{x} \\ & = v, \quad (\mathbf{E}_3 \mathbf{x})^\top \mathbf{E}_4 \mathbf{x} = 0 \quad (6) \end{aligned}$$

where $\mathbf{S} \in \mathbb{R}^{d \times n}$ contains all the data points in its columns and $\Lambda_j \in \mathbb{R}^{n \times nk}$ is zero everywhere except for the j^{th} block that is identity, i.e. $\Lambda_j = [\mathbf{0} \cdots \mathbf{I}_j \cdots \mathbf{0}]$.

D. ADMM SOLVER

Problem (6) is still difficult to solve due to the mixed binary and quadratic constraints. To handle these difficulties, (i) we first replace the binary constraints with an exact equivalent set that is the intersection of the ℓ_2 -sphere (defined by set S_2) and box constraints (defined by set S_b) following [17]. (ii) Moreover, by introducing the auxiliary variables ($\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3$, and \mathbf{z}_4), the quadratic constraints are now changed to bi-linear ones and separated from the binary constraints. Thus, the resultant problem is given as follows:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{w}, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \left\| \mathbf{s}_i - \mathbf{S} \Lambda_j \mathbf{w} \right\|_2^2 \\ \text{s.t.} & \Psi^\top \mathbf{x} = \mathbf{1}_n, \quad \mathbf{z}_1 \in S_b, \quad \mathbf{z}_2 \in S_2, \\ & \mathbf{Q} \mathbf{P}^\top \mathbf{x} = \mathbf{u} \mathbf{x} = \mathbf{P} \mathbf{w} \odot \mathbf{C} \mathbf{x}, \\ & (\mathbf{E}_1 \mathbf{z}_3)^\top \mathbf{E}_2 \mathbf{x} = v, \quad (\mathbf{E}_3 \mathbf{z}_4)^\top \mathbf{E}_4 \mathbf{x} = 0 \\ & \mathbf{x} = \mathbf{z}_1, \quad \mathbf{x} \\ & = \mathbf{z}_2, \quad \mathbf{x} = \mathbf{z}_3, \quad \mathbf{x} = \mathbf{z}_4 \quad (7) \end{aligned}$$

which can be solved using in the standard ADMM framework. Let $\mathcal{L}_{\rho_{1-9}}$ be the augmented Lagrangian function of problem (7). We define it as follows:

$$\begin{aligned} \mathcal{L}_{\rho_{1-9}}(\mathbf{x}, \mathbf{w}, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4, \mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4, \mathbf{y}_5, \mathbf{y}_6, \mathbf{y}_7, \mathbf{y}_8, \mathbf{y}_9) \\ := \sum_{j=1}^k \sum_{i=1}^n x_{ij} \left\| \mathbf{s}_i - \mathbf{S} \Lambda_j \mathbf{w} \right\|_2^2 + \mathbf{y}_1^\top (\Psi^\top \mathbf{x} - \mathbf{1}_n) \\ + \frac{\rho_1}{2} \|\Psi^\top \mathbf{x} - \mathbf{1}_n\|_2^2 \\ + \mathbb{I}_{\{\mathbf{z}_1 \in S_b\}} + \mathbf{y}_2^\top (\mathbf{x} - \mathbf{z}_1) + \frac{\rho_2}{2} \|\mathbf{x} - \mathbf{z}_1\|_2^2 \\ + \mathbb{I}_{\{\mathbf{z}_2 \in S_2\}} + \mathbf{y}_3^\top (\mathbf{x} - \mathbf{z}_2) \\ = \frac{\rho_3}{2} \|\mathbf{x} - \mathbf{z}_2\|_2^2 + \mathbf{y}_4^\top (\mathbf{Q} \mathbf{P}^\top \mathbf{x} - \mathbf{u}) \\ + \frac{\rho_4}{2} \|\mathbf{Q} \mathbf{P}^\top \mathbf{x} - \mathbf{u}\|_2^2 \\ = \mathbf{y}_5^\top (\mathbf{I} - \text{diag}(\mathbf{P} \mathbf{w}) \mathbf{C}) \mathbf{x} + \frac{\rho_5}{2} \|(\mathbf{I} - \text{diag}(\mathbf{P} \mathbf{w}) \mathbf{C}) \mathbf{x}\|_2^2 \\ = \mathbf{y}_6^\top (\mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v) + \frac{\rho_6}{2} \|\mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v\|_2^2 + \mathbf{y}_7^\top (\mathbf{x} - \mathbf{z}_3) \\ = \frac{\rho_7}{2} \|\mathbf{x} - \mathbf{z}_3\|_2^2 + \mathbf{y}_8^\top (\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x}) + \frac{\rho_8}{2} \|\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x}\|_2^2 \\ = \mathbf{y}_9^\top (\mathbf{x} - \mathbf{z}_4) + \frac{\rho_9}{2} \|\mathbf{x} - \mathbf{z}_4\|_2^2 \quad (8) \end{aligned}$$

where the \mathbf{y} variables are the Lagrange multipliers of the corresponding constraints, \mathbb{I} is the indicator function that penalizes infeasible \mathbf{z}_1 and \mathbf{z}_2 , and $\rho_{1-9} \geq 0$ are the penalty parameters. In our experiments, we set all the ρ coefficients to the same value. The iterative ADMM steps for problem (7) are described in Algorithm 1. ADMM updates are performed by optimizing for the set of primal variables one at a time, while keeping the rest of the primal and dual variables fixed. Then, the dual variables are updated using gradient ascent on the corresponding dual problem.

Algorithm 1 ADMM for Solving Problem (7)

Input : Set $\mathbf{S} \in \mathbb{R}^{d \times n}$. Set $\rho_{1-9}, \mathbf{y}_{1-5,7,9} = \mathbf{0}$,
 $y_{6,8} = 0, \mathbf{x}_{\text{kmeans}},$
 $\mathbf{w} = \mathbf{P}^\top \mathbf{x} \odot \text{diag}^{-1}(\mathbf{C}\mathbf{x})\mathbf{1}_{nk}.$

Output: \mathbf{x}

while not converged **do**

update: \mathbf{x} by solving Eq (9). **update:** \mathbf{w} by
 solving Eq (10). **update:** \mathbf{z}_{1-4} via
 Eqs (11,12, 13,14). **update:** $\mathbf{y}_{1-5,7,9}, y_{6,8}$ via
 Eqs (15).

end

We next show the final updates for each subproblem, but the exact derivations are found in the **supplementary material**.

a: UPDATE \mathbf{x}

We need to solve the following linear system using the conjugate gradient method.

$$\begin{aligned} & \left(\rho_1 \Psi \Psi^\top + (\rho_2 + \rho_3 + \rho_7 + \rho_9) \mathbf{I}_{nk} + \rho_4 \mathbf{P} \mathbf{Q}^\top \mathbf{Q} \mathbf{P}^\top \right. \\ & \left. + \rho_5 \left(\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C} \right)^\top \left(\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C} \right) \right. \\ & \left. + \rho_6 \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_3 \mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 + \rho_8 \mathbf{E}_4^\top \mathbf{E}_4 \mathbf{z}_4 \mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \right) \mathbf{x} = \\ & - \left(\text{vect}(\mathbf{B}) + \Psi \mathbf{y}_1 + \mathbf{y}_2 + \mathbf{y}_3 - \rho_1 \Psi \mathbf{1}_n - \rho_2 \mathbf{z}_1 \right. \\ & - \rho_3 \mathbf{z}_2 - \mathbf{C}^\top \text{diag}(\mathbf{P}\mathbf{w}) \mathbf{y}_5 + y_6 \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_4 - \rho_6 v \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_3 + y_7 \\ & \left. - \rho_7 \mathbf{z}_3 + y_8 \mathbf{E}_4^\top \mathbf{E}_3 \mathbf{z}_4 + y_9 - \rho_9 \mathbf{z}_4 + \mathbf{P} \mathbf{Q}^\top \mathbf{y}_4 - \rho_4 \mathbf{P} \mathbf{Q}^\top \mathbf{u} \right) \end{aligned} \quad (9)$$

where $\mathbf{B}(i, j) = \|\mathbf{s}_i - \mathbf{S} \Lambda_j \mathbf{w}\|_2^2$, and $\text{vect}(\mathbf{B})$ is simply a columnwise vectorization of the matrix \mathbf{B} .

b: UPDATE \mathbf{w}

We need to solve the following linear system using the conjugate gradient method.

$$\begin{aligned} & \left[\sum_j^k \sum_i^n 2x_{ij} \Lambda_j^\top \mathbf{S}^\top \mathbf{S} \Lambda_j + \rho_5 \mathbf{P}^\top \text{diag}(\mathbf{C}\mathbf{x} \odot \mathbf{C}\mathbf{x}) \mathbf{P} \right] \mathbf{w} \\ & = \sum_j^k \sum_i^n 2x_{ij} \Lambda_j^\top \mathbf{S}^\top \mathbf{s}_i + \mathbf{P}^\top \mathbf{C}\mathbf{x} \odot \mathbf{y}_5 + \rho_5 \mathbf{P}^\top \mathbf{C}\mathbf{x} \odot \mathbf{x} \end{aligned} \quad (10)$$

c: UPDATE \mathbf{z}_1

Here, we need to perform a simple projection onto the box: $S_b = \{\mathbf{a} : \mathbf{0} \leq \mathbf{a} \leq \mathbf{1}\}$. This projection is an elementwise clamping between 0 and +1.

$$\mathbf{z}_1 = \mathbf{P}_{S_b} \left(\mathbf{x} + \frac{\mathbf{y}_2}{\rho_2} \right) = \min \left(\max \left(\mathbf{x} + \frac{\mathbf{y}_2}{\rho_2}, \mathbf{0} \right), \mathbf{1} \right) \quad (11)$$

d: UPDATE \mathbf{z}_2

We need to perform a simple projection onto the ℓ_2 -sphere: $S_2 = \{\mathbf{a} \in \mathbb{R}^{nk} : \|\mathbf{a} - \frac{1}{2} \mathbf{1}\|_2^2 = \frac{nk}{4}\}$. This involves an elementwise shift and ℓ_2 vector normalization.

$$\mathbf{z}_2 = \mathbf{P}_{S_2} \left(\mathbf{x} + \frac{\mathbf{y}_3}{\rho_3} \right) = \frac{\sqrt{nk}}{2} \frac{\left(\mathbf{x} + \frac{\mathbf{y}_3}{\rho_3} \right) - \frac{1}{2} \mathbf{1}}{\left\| \left(\mathbf{x} + \frac{\mathbf{y}_3}{\rho_3} \right) - \frac{1}{2} \mathbf{1} \right\|_2} + \frac{1}{2} \mathbf{1} \quad (12)$$

e: UPDATE \mathbf{z}_3

We need to solve the following linear system using the conjugate gradient method.

$$\begin{aligned} & \left[\rho_6 \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} \mathbf{x}^\top \mathbf{E}_2^\top \mathbf{E}_1 + \rho_7 \mathbf{I}_{nk} \right] \mathbf{z}_3 = \mathbf{y}_7 + \rho_7 \mathbf{x} - y_6 \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} \\ & = \rho_6 v \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} \end{aligned} \quad (13)$$

f: UPDATE \mathbf{z}_4

We need to solve the following linear system using the conjugate gradient method.

$$\left[\rho_8 \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \mathbf{x}^\top \mathbf{E}_4^\top \mathbf{E}_3 + \rho_9 \mathbf{I}_{nk} \right] \mathbf{z}_4 = \mathbf{y}_9 + \rho_9 \mathbf{x} - y_8 \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \quad (14)$$

g: UPDATE $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4, \mathbf{y}_5, y_6, \mathbf{y}_7, y_8, \mathbf{y}_9$

Lastly, we need to perform dual ascent on the dual variables as follows:

$$\begin{aligned} \mathbf{y}_1 & \leftarrow \mathbf{y}_1 + \rho_1 \left(\Psi^\top \mathbf{x} - \mathbf{1}_n \right), \mathbf{y}_2 \leftarrow \mathbf{y}_2 + \rho_2 \left(\mathbf{x} - \mathbf{z}_2 \right) \\ \mathbf{y}_3 & \leftarrow \mathbf{y}_3 + \rho_3 \left(\mathbf{x} - \mathbf{z}_3 \right), \mathbf{y}_4 \leftarrow \mathbf{y}_4 + \rho_4 \left(\mathbf{Q} \mathbf{P}^\top \mathbf{x} - \mathbf{u} \right) \\ \mathbf{y}_5 & \leftarrow \mathbf{y}_5 + \rho_5 \left(\mathbf{x} - \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x} \right), \\ \mathbf{y}_6 & \leftarrow \mathbf{y}_6 + \rho_6 \left(\mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right) \\ \mathbf{y}_7 & \leftarrow \mathbf{y}_7 + \rho_7 \left(\mathbf{x} - \mathbf{z}_3 \right), \mathbf{y}_8 \leftarrow \mathbf{y}_8 + \rho_8 \left(\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right) \\ \mathbf{y}_9 & \leftarrow \mathbf{y}_9 + \rho_9 \left(\mathbf{x} - \mathbf{z}_4 \right) \end{aligned} \quad (15)$$

The ADMM iterations are run until convergence (i.e. when the standard deviation between the last 10 objective values is $\leq 10^{-5}$). Upon convergence, all the primal variables (\mathbf{x} and \mathbf{z}_{1-4}) converge to the same feasible binary vector. Despite that the problem is non-convex, we show empirically in the experiments' section and in the **supplementary material** that the performance using is very stable.

IV. EXPERIMENTS

In this section, we conduct extensive experiments to motivate and evaluate our proposed clustering method, both on synthetic and real datasets. We also compare our method against other constrained clustering methods on well-known benchmarks, thus, demonstrating superior performance and flexibility, as well as, superior gain that can be achieved when both categories of constraints (cardinality and pairwise) are combined in our framework.

A. DATASETS AND IMPLEMENTATION DETAILS

The datasets used in this section vary from synthetic to real. As for the synthetic ones, we construct two datasets, one is cluster balanced (denoted as *Balanced*) and the other is imbalanced (denoted as *ImBalanced*) as shown in Figure 1. Each dataset comprises 700 data points with 2 clusters. In *Balanced*, each cluster has exactly 350 data points, while in *ImBalanced* one cluster has 600 data points while the other contains 100. As for the real datasets, we make use of various popular UCI datasets [7], e.g. iris, wine, glass, ionosphere, Hepatitis, Hepatitis1 and Breast Cancer Wis-D. These are the most popular UCI datasets used for clustering purposes [8], [13]. Following convention, data points are normalized to have a value in $[-1, +1]$. For Hepatitis and Hepatitis1, we remove all points with missing or none categorical features. Table 1 lists the details of all UCI datasets used in the experiments.

TABLE 1. Lists the total number of points, clusters and features of all UCI datasets used in the experiments.

Datasets	#Points	#Features	#Clusters
wine	178	13	3
iris	150	4	3
glass	214	9	7
ionosphere	351	33	2
Hepatitis	142	14	2
Hepatitis1	80	19	2
Breast Cancer Wis-D	569	30	2

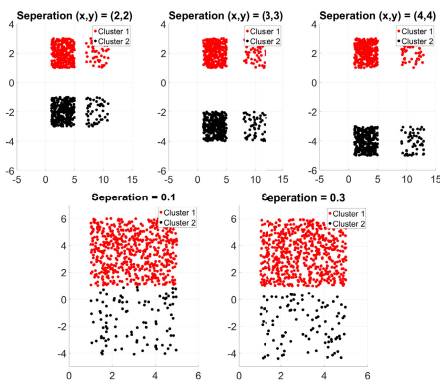


FIGURE 1. The *Balanced* and *ImBalanced* in the two consecutive rows respectively. They comprise two clusters (red/black) with an increasing separation between clusters.

As for the implementation details, none of the selection matrices used in the proposed framework (i.e. $E_1, E_2, E_3, E_4, P, C, Q, \Psi, \Lambda_j$) are actually constructed. Only element indexing within vectors is used, thus, keeping the necessary computation cost minimal. For ease, all ρ_i parameters have the same value and updated similarly. We find that setting all ρ_i parameters to 20 and by increasing it every 5 iterations by 10% for all real datasets achieves the fastest convergence. Moreover, we initialize all the optimization variables using zero vectors, while x is initialized to random (i.e. random

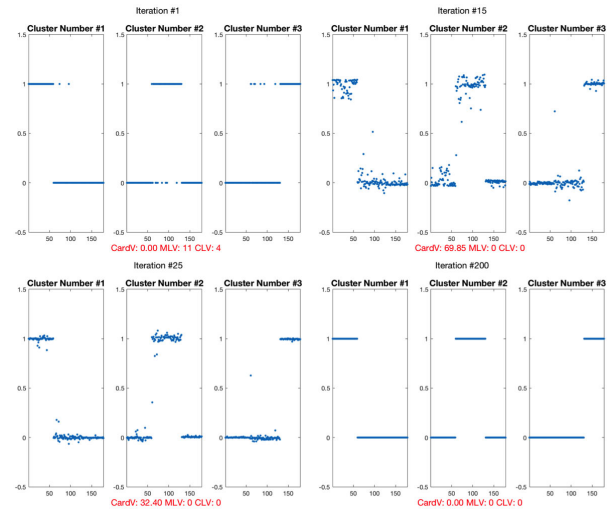


FIGURE 2. Convergence of the solution x using ℓ_p Km-Mix with random binary initialization that satisfy the cardinality constraints on the Wine dataset.

TABLE 2. Comparison between K-means and ℓ_p Km on real UCI datasets using K-means objective value, *ARI*(%), *MI* and *HI*(%) along with the standard deviation.

Datasets	K-mean (Obj Value)	ℓ_p Km (Obj Value)
wine	195.91 ± 0.05	195.86 ± 0.07
iris	29.894 ± 4.84	28.282 ± 0.05
glass	81.277 ± 5.49	85.944 ± 1.307

	K-means (<i>ARI</i>)	ℓ_p Km (<i>ARI</i> (%))
wine	84.765 ± 0.92	86.779 ± 0.10
iris	67.831 ± 8.79	72.41 ± 0.90
glass	17.454 ± 1.31	14.96 ± 0.91

	K-means (<i>MI</i>)	ℓ_p Km (<i>MI</i>)
wine	0.068 ± 0.004	0.059 ± 0.001
iris	0.145 ± 0.048	0.122 ± 0.004
glass	0.338 ± 0.019	0.319 ± 0.019

	K-means (<i>HI</i>)	ℓ_p Km (<i>HI</i> (%))
wine	86.423 ± 0.82	88.225 ± 0.001
iris	70.969 ± 9.51	75.590 ± 0.008
glass	32.378 ± 3.57	36.134 ± 0.038

assignment of data points to clusters) if the comparison is against K-means. When comparing against other clustering methods, we use the same K-means initialization as other methods. In all comparisons, w is initialized to a feasible point as given in Algorithm 1. MATLAB is used to implement our method. The most expensive operation in our framework is the x and w updates, which involve solving an $n \times k$ linear system. This is the bottleneck of our framework causing it to have a computational complexity $\mathcal{O}(n^3k^3)$ per iteration. In the final experiment, we report the runtime of our framework on different sized datasets with a variety of constraint choices.

As for the evaluation metric, we adopt the 3 most common criteria used in the clustering community to compare different clustering methods, namely the Adjusted Random Index (*ARI*)(\nearrow), Mirkin’s Index (*MI*)(\searrow) and Hubert’s Index (*HI*)(\nearrow) which calculate a measure of agreement between two partitions of a dataset [25], [26]. The symbol \nearrow indicate that

TABLE 3. Comparison between K-means and ℓ_p Km-Car on synthetic balanced and Imbalanced synthetic datasets using ARI(%), MI and HI(%).

Datasets	K-mean (ARI)	ℓ_p Km-Car (ARI)
Balanced (x,y)=(2,2)	59.97 \pm 51.68	100 \pm 0
Balanced (x,y)=(3,3)	89.99 \pm 31.64	100 \pm 0
Balanced (x,y)=(4,4)	79.99 \pm 42.19	100 \pm 0
Imbalanced y=0.1	34.67 \pm 0	99.46 \pm 0.37
Imbalanced y=0.3	32.26 \pm 0	100 \pm 0

	K-means (MI)	ℓ_p Km-Car (MI)
Balanced (x,y)=(2,2)	0.2003 \pm 0.2586	0 \pm 0
Balanced (x,y)=(3,3)	0.0501 \pm 0.1583	0 \pm 0
Balanced (x,y)=(4,4)	0.1001 \pm 0.2111	0 \pm 0
Imbalanced y=0.1	0.3082 \pm 0	0.002 \pm 0.0014
Imbalanced y=0.3	0.3222 \pm 0	0 \pm 0

	K-means (HI)	ℓ_p Km-Car (HI)
Balanced (x,y)=(2,2)	59.94 \pm 5.171	100 \pm 0
Balanced (x,y)=(3,3)	89.99 \pm 31.67	100 \pm 0
Balanced (x,y)=(4,4)	79.97 \pm 42.22	100 \pm 0
Imbalanced y=0.1	38.35 \pm 0	99.6 \pm 0.28
Imbalanced y=0.3	35.57 \pm 0	100 \pm 0

the higher the number the better performance and vice versa for \searrow . In all experiments, clustering is repeated 10 times with different initializations and we report the average and standard deviation of the metric used in comparison.

B. COMPARING DIFFERENT CONSTRAINT DESIGN CHOICES

We apply our proposed method, ℓ_p Km, on the same clustering task with several choices of constraints: no constraints, only cardinality constraints, only pairwise constraints and both types jointly. We refer to each as ℓ_p Km, ℓ_p Km-Car, ℓ_p Km-Pair and ℓ_p Km-Mix respectively.

a: AN AUXILIARY EXPERIMENT

Despite that we do not provide a proof for the convergence of the non-quadratic objective in Eq. 1, as it is proven for the quadratic case in [17], we find the performance very stable where we demonstrate it empirically. For instance, we run ℓ_p Km-Mix that enforces cardinality, 20 must-link and 20 cannot-link constraints. In figure 2, we plot the three pieces of the solution label vector \mathbf{x} at four different ADMM iterations (1, 15, 25, and 200). In the first iteration, the initial clustering is random however satisfying the cardinality constraints, so it is binary but it does not lead to a good objective. As ADMM progresses, the continuous solution \mathbf{x} becomes more and more binary, until it converges to a feasible binary solution where the three clusters are disjoint satisfying all constraints. Moreover, we also report the number of cardinality (CardV), must-link (MLV), and cannot-link (CLV) violations at each of these iterations. These violations gradually decrease until convergence occurs, when no violations persist. We find this stable performance across all datasets as will be presented in later sections. Further detailed experiments can be found in the **supplementary material**.

Traditional K-Means Versus ℓ_p Km: First, we start by comparing our vanilla constrained free version clustering method ℓ_p Km against K-means. We show that ℓ_p Km method can

in fact attain very similar, if not better, performance than traditional K-means (builtin MATLAB function). This is clearly because both methods use the same objective value and that ℓ_p Km does converge to good solutions. Experiments are conducted on some of the UCI datasets [7] (wine, iris and glass). Table 2 reports the K-means objective value, ARI(%), MI and HI(%) metrics for both methods.

b: TRADITIONAL K-MEANS VERSUS ℓ_p Km-Car.

Here, we demonstrate that our framework coupled with only cardinality constraints outperforms traditional unconstrained K-means on a variety of synthetic data. This highlights the importance of having this prior information available and harnessing it in the clustering process. In these experiments, the cardinality constraints are generated from ground truth labels. To show that cardinality does in fact help clustering performance, we apply ℓ_p Km-Car on the two synthetic datasets (Balanced and ImBalanced) and report their ARI, MI and HI results in Table 3.

For the Balanced dataset, the separation between the four groups of points increases. In fact, K-means tends to cluster points together such that each cluster has a similar variance as other clusters. Consequently, K-means clusters the high-density points of the Balanced dataset together and groups the remaining less dense points into another cluster. In comparison, our framework exploits the cardinality constraints to achieve perfect clustering performance. Similarly, the ImBalanced dataset contains two imbalanced clusters with very different densities, where the separation between them is increased. In this case, K-means often mixes data points between clusters, since the cardinality constraints are not used. On the other hand, ℓ_p Km-Car can almost perfectly predict the ground truth clustering labels. Interestingly, the variance of our results is much lower than that of K-means even though they both use the same clustering initialization. This indicates that the cardinality constraints afford our method robustness to the initialization.

To the best of our knowledge, all previous work that handles generic cardinality constraints do not have readily available code for comparison. Therefore, we only compare our method with traditional unconstrained K-means, so as to demonstrate the effectiveness of adding cardinality constraints to an unconstrained clustering method.

c: PAIRWISE CONSTRAINED CLUSTERING METHODS VERSUS ℓ_p Km-Pair

Here, we compare our ℓ_p Km-Pair method against several pairwise constrained methods from the literature, namely Constrained Clustering [13] (COP-KMEANS), Spectral Clustering [8], Penalized probabilistic Clustering (PPC) [10] and CCL [11]. All pairwise constraints were randomly generated. Among these methods, only COP-KMEANS [13] and ℓ_p Km-Pair exactly enforce the constraints, while the others incorporate them as soft pairwise constraints in their clustering framework. Consequently, Spectral Clustering, PPC and CCL may result in clustering violations.

TABLE 4. Comparison of several pairwise constrained clustering methods against ℓ_p Km-Pair using *ARI*(%), as well as, must-link (MLV) and cannot-link (CLV) violations in the constraints. The cells indicated with x imply that the underlying method does not attain a feasible solution after 1000 runs.

	Constraints	COP-KMEANS			Spectral Clustering			PPC			CCL			ℓ_p Km-Pair		
		ARI(%)	MLV	CLV	ARI(%)	MLV	CLV	ARI(%)	MLV	CLV	ARI(%)	MLV	CLV	ARI(%)	MLV	CLV
wine	20ml, 20cl	87.14 ± 1.06	0	0	91.07 ± 2	0	0	88.19 ± 0	1	0	74.52 ± 0	0	0	91.67 ± 0	0	0
	40ml, 40cl	84.48 ± 0.71	0	0	96.67 ± 0	0	0	89.68 ± 1.64	1	1.5	53.58 ± 0	3	4	88.04 ± 0	0	0
	60ml, 60cl	91.22 ± 0	0	0	94.87 ± 0	2	0	89.74 ± 0.04	4.8	2.9	82.79 ± 0	7	1	98.32 ± 0	0	0
	80ml, 80cl	93.09 ± 0	0	0	96.67 ± 0	1	0	87.73 ± 0.74	6	3	55 ± 0	11	10	94.87 ± 0	0	0
	100ml, 100cl	96.51 ± 0	0	0	94.87 ± 0	3	2	88.33 ± 0.14	4.4	2.1	90.34 ± 0	1	5	96.51 ± 0	0	0
iris	20ml, 20cl	69.57 ± 0.61	0	0	80.27 ± 0.64	1.1	2.2	90.39 ± 0	1	0	68.28 ± 0	0	2	72.87 ± 0	0	0
	40ml, 40cl	74.37 ± 0	0	0	90.38 ± 0	1	1	92.22 ± 0	1	2	45.65 ± 0	3	11	72.76 ± 0	0	0
	60ml, 60cl	x	x	x	69.00 ± 0.59	0	19.5	97.99 ± 0	0	1	54.86 ± 0	2	0	100 ± 0	0	0
	80ml, 80cl	x	x	x	68.44 ± 0	0	20	97.99 ± 0	0	1	54.17 ± 0	2	0	100 ± 0	0	0
	100ml, 100cl	x	x	x	68.44 ± 0	0	20	97.99 ± 0	0	1	52.52 ± 0	4	1	100 ± 0	0	0

TABLE 5. Comparison of several pairwise constrained clustering methods against ℓ_p Km-Pair using *MI* and *HI*(%). The cells indicated with x imply that the underlying method does not attain a feasible solution after 1000 runs.

	Constraints	COP-KMEANS		Spectral Clustering		PPC		CCL		ℓ_p Km-Pair	
		MI	HI(%)	MI	HI(%)	MI	HI(%)	MI	HI(%)	MI	HI(%)
wine	20ml, 20cl	0.06 ± 0.01	88.55 ± 0.95	0.04 ± 0.01	92.03 ± 1.8	0.05 ± 0	89.48 ± 0	0.11 ± 0	77.32 ± 0	0.04 ± 0	92.57 ± 0
	40ml, 40cl	0.07 ± 0	86.18 ± 0.62	0.02 ± 0	97.03 ± 0	0.05 ± 0.02	90.80 ± 1.46	0.21 ± 0	57.37 ± 0	0.05 ± 0	89.34 ± 0
	60ml, 60cl	0.04 ± 0	92.15 ± 0	0.02 ± 0	95.42 ± 0	0.05 ± 0	90.85 ± 0.04	0.08 ± 0	84.69 ± 0	0.01 ± 0	98.50 ± 0
	80ml, 80cl	0.03 ± 0	93.83 ± 0	0.02 ± 0	97.03 ± 0	0.06 ± 0	89.07 ± 0.66	0.20 ± 0	59.63 ± 0	0.02 ± 0	95.42 ± 0
	100ml, 100cl	0.02 ± 0	96.88 ± 0	0.02 ± 0	95.42 ± 0	0.05 ± 0	89.60 ± 0.13	0.04 ± 0	91.39 ± 0	0.02 ± 0	96.88 ± 0
iris	20ml, 20cl	0.14 ± 0.01	72.92 ± 0.49	0.09 ± 0.036	82.41 ± 5.11	0.04 ± 0	91.50 ± 0	0.14 ± 0	71.36 ± 0	0.12 ± 0	75.95 ± 0
	40ml, 40cl	0.11 ± 0	77.18 ± 0.75	0.04 ± 0	91.50 ± 0	0.03 ± 0	93.13 ± 0	0.26 ± 0	49.01 ± 0	0.12 ± 0	75.95 ± 0
	60ml, 60cl	x	x	0.14 ± 0	71.91 ± 0.58	0.01 ± 0	98.23 ± 0	0.22 ± 0	55.38 ± 0	0 ± 0	100 ± 0
	80ml, 80cl	x	x	0.14 ± 0	71.36 ± 0	0.01 ± 0	98.23 ± 0	0.22 ± 0	55.69 ± 0	0 ± 0	100 ± 0
	100ml, 100cl	x	x	0.14 ± 0	71.36 ± 0	0.01 ± 0	98.23 ± 0	0.23 ± 0	53.18 ± 0	0 ± 0	100 ± 0




FIGURE 3. Effect of increasing must-link and cannot-link constraints separately, as compared to unconstrained K-means.

TABLE 6. Comparison of several pairwise constrained clustering methods against ℓ_p Km-Pair using *MI* and *HI*(%). The cells indicated with x imply that the underlying method does not attain a feasible solution after 1000 runs.

Datasets	ℓ_p Km-Car				ℓ_p Km-Pair				ℓ_p Km-Mix			
	ARI(%)	MI	HI(%)	Time	ARI(%)	MI	HI(%)	Time	ARI(%)	MI(%)	HI	Time
ionosphere	8.49 ± 0	0.46 ± 0	9.03 ± 0	1.50 sec	35.65 ± 0.67	0.321 ± 0.003	35.68 ± 0.57	5.00 sec	80.42 ± 4.35	0.097 ± 0.002	80.54 ± 4.32	6.57 sec
Hepatitis	17.12 ± 1.5	0.36 ± 0.01	29.07 ± 1.28	0.55 sec	38.17 ± 4.01	0.303 ± 0.021	39.49 ± 4.33	0.99 sec	46.25 ± 15.74	0.230 ± 0.067	54 ± 13.47	1.09 sec
Hepatitis1	40.98 ± 7.48	0.24 ± 0.03	52.86 ± 5.29	0.91 sec	26.97 ± 12.26	0.358 ± 0.065	28.50 ± 12.96	0.70 sec	77.27 ± 10.01	0.091 ± 0.040	81.85 ± 7.99	1.26 sec
Breast Cancer Wis-D	74.93 ± 0	0.125 ± 0	75.03 ± 0	3.36 sec	74.93 ± 0	0.128 ± 0	74.42 ± 0	3.27 sec	74.93 ± 0	0.125 ± 0	75.03 ± 0	9.27 sec

However, due to the heuristic nature of COPKMEANS, it may lead to a situation where depending on the initialization no feasible solution is attained. We run all five methods on two UCI datasets (wine, iris) and ensure that all methods receive the same randomly generated pairwise constraints. Tables 4 and 5 report the performance of these methods on all discussed metrics. For each experiment, we also report the number of must-link (ml) and cannot-link (cl) constraints, as well as, the number of must-link violations (MLV) and the cannot-link violations (CLV). It is clear that ℓ_p Km-Pair outperforms all other methods, while satisfying all the constraints.

d: (K-MEANS VERSUS ℓ_p Km-Mix.

Here, we demonstrate the main motivation behind our flexible framework, namely its ability to incorporate both cardinality and pairwise constraints simultaneously in the clustering optimization. Firstly, and following previous work [13], we demonstrate that increasing the number of pairwise constraints (either must-link or cannot-link) with the same cardinality constraints consistently improves performance. We conduct this experiment on two different datasets: one synthetic (ImBalanced $y=0, 1$) and one real (wine). Figure 3 compares our method against traditional K-means in such setup. Obviously, K-means does not benefit from the

constraints while ours consistently improves in performance. Secondly, we compare all three variants of our framework, i.e. cardinality only constraints (ℓ_p Km-Car), pairwise only constraints (ℓ_p Km-Pair) and both (ℓ_p Km-Mix), on several UCI datasets (ionosphere, Hepatitis, Hepatitis1 and Breast Cancer Wis-D). The number of must-link and cannot-link constraints were equal for each dataset and they were set proportional to the dataset size to (20,25,20,100) respectively. Results in Table 6 show that our method performs increasingly and significantly better, when more constraint categories are used simultaneously. This improvement reaches as high as 40% in ARI for some datasets. We also report in table 6 the runtime for all 3 different variants on all 4 datasets. The time vary depending on the dataset size and the number of clusters from 0.5 – 10 seconds. We did not compare against other methods here, since there is no existing work that combines both categories of constraints in a unified framework and extending the pairwise constrained methods to cardinality constraints is not trivial.

V. CONCLUSION

We proposed a new flexible framework to handle both pairwise and cardinality constraints for K-means clustering. The resulting integer program is transformed into an equivalent continuous reformulation where pairwise constraints are incorporated as quadratic constraints. The resultant problem is solved using ADMM. Extensive experiments have been conducted on both synthetic and real datasets to demonstrate the competitive performance of our method under different constraint choices and that the proposed method achieves state-of-art performance when both types of constraints are used simultaneously. As a future work, we seek to adapt our work to deep clustering approaches [27], [28], [29], which jointly learn feature representations and cluster assignments. That is to say, the ADMM updates will involve another minimization step that trains the encoder for the feature learning procedure while the other steps involve on solving both the assignment and imposing the constraints. This shall have applications in biology, medicine, finance, and animation, as in many such applications, where pseudo labels in terms of assignments are provided as constraints, e.g., two patients are labeled to have a similar syndrome, feature learning will allow for further improvements in clustering performance.

APPENDIX REFORMULATION OF CONSTRAINED K-MEANS

We start with our formulation of constrained K-means in Equation (1) below (or Equation (4) in the manuscript).

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{w}, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4} \sum_{j=1}^k \sum_{i=1}^n x_{ij} \|\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}\|_2^2 \\ \text{s.t.} \quad & \Psi^\top \mathbf{x} = \mathbf{1}_n, \mathbf{z}_1 \in S_b, \mathbf{z}_2 \in S_2, \mathbf{Q}\mathbf{P}^\top \mathbf{x} = \mathbf{u}, \\ & \mathbf{x} = \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x}, \\ & (\mathbf{E}_1 \mathbf{z}_3)^\top \mathbf{E}_2 \mathbf{x} = v, (\mathbf{E}_3 \mathbf{z}_4)^\top \mathbf{E}_4 \mathbf{x} = 0, \\ & \mathbf{x} = \mathbf{z}_1, \mathbf{x} = \mathbf{z}_2, \mathbf{x} = \mathbf{z}_3, \mathbf{x} = \mathbf{z}_4. \end{aligned}$$

We define $\mathbf{X} \in \mathbb{R}^{k \times n}$ and therefore $\mathbf{x} = \text{vect}(\mathbf{X})$. The vect operator simply vectorizes the matrix one column at a time (i.e. one data point at a time). The order of concatenation is reverse for the label vector: $\mathbf{w} = \text{vect}(\mathbf{W})$, where $\mathbf{W} \in \mathbb{R}^{n \times k}$ (i.e. the vectorization is done one cluster at a time). Therefore, it is important to point out that the order of the binary labels \mathbf{x} and that of \mathbf{w} is swapped. Reordering these vectors based on data points or clusters is controlled by the permutation matrix $\mathbf{P} \in \mathbb{R}^{nk \times nk}$. Therefore, $\mathbf{P}\mathbf{x} = \mathbf{P}^\top \mathbf{x} = \text{vect}(\mathbf{X}^\top)$ and $\mathbf{P}\mathbf{w} = \mathbf{P}^\top \mathbf{w} = \text{vect}(\mathbf{W}^\top)$. Throughout the main manuscript and the appendix, we use \mathbf{P} and \mathbf{P}^\top to change the order of \mathbf{w} to the same order as \mathbf{x} and vice versa. Matrix $\mathbf{S} \in \mathbb{R}^{d \times n}$ in Equation (1) simply concatenates all the data points in its columns, where $\Lambda_j \in \mathbb{R}^{n \times nk}$ is zero everywhere except for the j^{th} block that is identity, i.e. $\Lambda_j = [\mathbf{0}, \dots, \mathbf{I}_k^j, \dots]$.

As for $\Psi^\top \in \mathbb{R}^{n \times nk}$, it is a binary matrix that has in each row a vector $\mathbf{1}_k^\top$ that sums all the binary labels for a given data point while the rest are zeros. We write this matrix in blockwise form as follows:

$$\Psi^\top = \begin{bmatrix} \mathbf{1}_k^\top & \mathbf{0}_k & \cdots \\ \mathbf{0}_k^\top & \mathbf{1}_k^\top & \mathbf{0}_k \\ \vdots & \vdots & \vdots \\ \mathbf{0}_k & \cdots & \mathbf{1}_k \end{bmatrix}$$

As for $\mathbf{Q} \in \mathbb{R}^{k \times nk}$, it sums all the binary labels of each cluster at a time for all the data points and its blockwise matrix form is given as follows:

$$\mathbf{Q} = \begin{bmatrix} \mathbf{1}_n^\top & \mathbf{0}_{nk-n}^\top & \cdots \\ \mathbf{0}_n^\top & \mathbf{1}_n^\top & \mathbf{0}_{nk-n}^\top & \cdots \\ \cdots & \cdots & \cdots & \cdots \end{bmatrix}$$

As for $\mathbf{C} \in \mathbb{R}^{nk \times nk}$, it sums the binary labels for each cluster at a time and its blockwise matrix form is as follows:

$$\mathbf{C} = \begin{bmatrix} \mathbf{I}_k & \cdots & \mathbf{I}_k \\ \vdots & \cdots & \vdots \\ \mathbf{I}_k & \cdots & \mathbf{I}_k \end{bmatrix}$$

As for the box and ℓ_2 -sphere constraints (which intersects with the binary vector space), we define two sets: $S_b := \{\mathbf{a} : \mathbf{0} \leq \mathbf{a} \leq \mathbf{1}\}$ and $S_2 := \{\mathbf{a} \in \mathbb{R}^n : \|\mathbf{a} - \frac{1}{2}\mathbf{1}\|_2^2 = \frac{n}{2}\}$, respectively. It is shown in [17] that: $\{0, 1\}^n = S_b \cap S_2$.

Lastly, $\mathbf{E}_1, \mathbf{E}_2 \in \mathbb{R}^{kv \times nk}$ and $\mathbf{E}_3, \mathbf{E}_4 \in \mathbb{R}^{ke \times nk}$ are selection matrices for the must-link and cannot-link constraints respectively. They select the data points that are involved in both types of constraints.

Applying ADMM: Following the conventional treatment of an optimization problem using ADMM, we first formulate the augmented Lagrangian function is given as follows:

$$\begin{aligned} & \mathcal{L}(\mathbf{x}, \mathbf{w}, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4, \mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4, \mathbf{y}_5, \mathbf{y}_6, \mathbf{y}_7, \mathbf{y}_8, \mathbf{y}_9) \\ & := \sum_{j=1}^k \sum_{i=1}^n x_{ij} \|\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}\|_2^2 + \mathbf{y}_1^\top (\Psi^\top \mathbf{x} - \mathbf{1}_n) \\ & \quad + \frac{\rho_1}{2} \|\Psi^\top \mathbf{x} - \mathbf{1}_n\|_2^2 \end{aligned}$$

$$\begin{aligned}
 & + \mathbb{I}_{\{z_1 \in S_b\}} + \mathbf{y}_2^\top (\mathbf{x} - \mathbf{z}_1) + \frac{\rho_2}{2} \|\mathbf{x} - \mathbf{z}_1\|_2^2 + \mathbb{I}_{\{z_2 \in S_2\}} \\
 & + \mathbf{y}_3^\top (\mathbf{x} - \mathbf{z}_2) \\
 & + \frac{\rho_3}{2} \|\mathbf{x} - \mathbf{z}_2\|_2^2 + \mathbf{y}_4^\top (\mathbf{Q}\mathbf{P}^\top \mathbf{x} - \mathbf{u}) + \frac{\rho_4}{2} \|\mathbf{Q}\mathbf{P}^\top \mathbf{x} - \mathbf{u}\|_2^2 + \\
 & \mathbf{y}_5^\top (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x} + \frac{\rho_5}{2} \|(\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x}\|_2^2 \\
 & = y_6 \left(\mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right) + \frac{\rho_6}{2} \left\| \mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right\|_2^2 + \mathbf{y}_7^\top \\
 & (\mathbf{x} - \mathbf{z}_3) \\
 & + \frac{\rho_7}{2} \|\mathbf{x} - \mathbf{z}_3\|_2^2 + y_8 \left(\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right) + \frac{\rho_8}{2} \left\| \mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right\|_2^2 \\
 & = \mathbf{y}_9^\top (\mathbf{x} - \mathbf{z}_4) + \frac{\rho_9}{2} \|\mathbf{x} - \mathbf{z}_4\|_2^2.
 \end{aligned}$$

ADMM updates steps tend to update each primal variable (\mathbf{x} , \mathbf{w} , and \mathbf{z}_{1-4} sequentially, while keeping the rest of these variables and the dual variables \mathbf{y}_{1-5} , $\mathbf{y}_{6,8}$, and $\mathbf{y}_{6,9}$) set to their most recent values. After the primal variables are updated, the dual variables are updated via a single gradient ascent step. Next, we detail each update step and the underlying optimization sub-problem that needs to be solved.

Update \mathbf{x} :

$$\begin{aligned}
 \mathbf{x} \leftarrow \arg \min_{\mathbf{x}} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \|\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}\|_2^2 + \mathbf{y}_1^\top (\Psi^\top \mathbf{x} - \mathbf{1}_n) \\
 & + \frac{\rho_1}{2} \|\Psi^\top \mathbf{x} - \mathbf{1}_n\|_2^2 + \mathbf{y}_2^\top (\mathbf{x} - \mathbf{z}_1) + \frac{\rho_2}{2} \|\mathbf{x} - \mathbf{z}_1\|_2^2 \\
 & + \mathbf{y}_3^\top (\mathbf{x} - \mathbf{z}_2) + \\
 & \frac{\rho_3}{2} \|\mathbf{x} - \mathbf{z}_2\|_2^2 + \mathbf{y}_4^\top (\mathbf{Q}\mathbf{P}^\top \mathbf{x} - \mathbf{u}) + \frac{\rho_4}{2} \|\mathbf{Q}\mathbf{P}^\top \mathbf{x} - \mathbf{u}\|_2^2 \\
 & + \mathbf{y}_5^\top (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x} + \frac{\rho_5}{2} \|(\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x}\|_2^2 \\
 & + y_6 \left(\mathbf{z}_4^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right) + \frac{\rho_6}{2} \left\| \mathbf{z}_4^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right\|_2^2 \\
 & + \mathbf{y}_7^\top (\mathbf{x} - \mathbf{z}_4) \\
 & + \frac{\rho_7}{2} \|\mathbf{x} - \mathbf{z}_4\|_2^2 + y_8 \left(\mathbf{z}_5^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right) \\
 & + \frac{\rho_8}{2} \left\| \mathbf{z}_5^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right\|_2^2 + \mathbf{y}_9^\top (\mathbf{x} - \mathbf{z}_5) \\
 & + \frac{\rho_9}{2} \|\mathbf{x} - \mathbf{z}_5\|_2^2
 \end{aligned}$$

The aforementioned problem is strongly convex quadratic in \mathbf{x} . Therefore, a stationary point is necessary and sufficient for optimality. By equating the gradient to zero, we get:

$$\begin{aligned}
 & \left(\rho_1 \Psi \Psi^\top + (\rho_2 + \rho_3 + \rho_7 + \rho_9) \mathbf{I}_{nk} + \rho_4 \mathbf{P}\mathbf{Q}^\top \mathbf{Q}\mathbf{P}^\top \right. \\
 & + \rho_5 (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})^\top (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C}) \\
 & + \rho_6 \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_3 \mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 + \rho_8 \mathbf{E}_4^\top \mathbf{E}_3 \mathbf{z}_4 \mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \left. \right) \mathbf{x} = \\
 & - (\text{vect}(\mathbf{B}) + \Psi \mathbf{y}_1 + \mathbf{y}_2 + \mathbf{y}_3 - \rho_1 \Psi \mathbf{1}_n - \rho_2 \mathbf{z}_1 - \\
 & \rho_3 \mathbf{z}_2 - \mathbf{C}^\top \text{diag}(\mathbf{P}\mathbf{w})\mathbf{y}_5 + y_6 \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_4 - \rho_6 v \mathbf{E}_2^\top \mathbf{E}_1 \mathbf{z}_3 + \mathbf{y}_7 \\
 & - \rho_7 \mathbf{z}_3 + y_8 \mathbf{E}_4^\top \mathbf{E}_3 \mathbf{z}_4 + \mathbf{y}_9 - \rho_9 \mathbf{z}_4 + \mathbf{P}\mathbf{Q}^\top \mathbf{y}_4 - \rho_4 \mathbf{P}\mathbf{Q}^\top \mathbf{u}),
 \end{aligned}$$

where $\mathbf{B}(i, j) = \|\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}\|_2^2$.

Update \mathbf{w} :

$$\begin{aligned}
 \mathbf{w} \leftarrow \arg \min_{\mathbf{w}} & \sum_{j=1}^k \sum_{i=1}^n x_{ij} \|\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}\|_2^2 + \\
 & \mathbf{y}_5^\top (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x} + \frac{\rho_5}{2} \|(\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x}\|_2^2.
 \end{aligned}$$

Similarly to the \mathbf{x} -update, the problem is strongly convex quadratic and finding a stationary point is necessary and sufficient for a global solution. Thus, the gradient is given by:

$$\begin{aligned}
 & - \sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} (\mathbf{S}\Lambda_j)^T (\mathbf{s}_i - \mathbf{S}\Lambda_j \mathbf{w}) - \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5 \\
 & - \rho_5 \left(\mathbf{P}^T \text{diag}(\mathbf{C}\mathbf{x}) \right) (\mathbf{I} - \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C})\mathbf{x} = 0.
 \end{aligned}$$

Then, we have

$$\begin{aligned}
 & - \sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{s}_i + \sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{S}\Lambda_j \mathbf{w} \\
 & - \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5 \\
 & - \rho_5 \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{x} + \rho_5 \mathbf{P}^T \text{diag}(\mathbf{C}\mathbf{x}) \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C}\mathbf{x} = 0.
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 & \left[\sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{S}\Lambda_j \mathbf{w} + \rho_5 \mathbf{P}^T \text{diag}(\mathbf{C}\mathbf{x}) \text{diag}(\mathbf{P}\mathbf{w})\mathbf{C}\mathbf{x} \right] \\
 & = \sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{s}_i + \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5 + \rho_5 \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{x}
 \end{aligned}$$

Finally, we have,

$$\begin{aligned}
 & \left[\sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{S}\Lambda_j + \rho_5 \mathbf{P}^T \text{diag}(\mathbf{C}\mathbf{x} \odot \mathbf{C}\mathbf{x})\mathbf{P} \right] \mathbf{w} \\
 & = \sum_{j=1}^k \sum_{i=1}^n 2 x_{ij} \Lambda_j^T \mathbf{S}^T \mathbf{s}_i + \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5 + \rho_5 \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{x}
 \end{aligned}$$

In this derivation, we use the fact that $\nabla_{\mathbf{w}} (\mathbf{y}_5^T \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x}) = \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5$. To see why this is the case, note the following identities: $\mathbf{a} \odot \mathbf{b} = \mathbf{b} \odot \mathbf{a} = \text{diag}(\mathbf{a})\mathbf{b} = \text{diag}(\mathbf{b})\mathbf{a}$.

Therefore,

$$\begin{aligned}
 \nabla_{\mathbf{w}} \left(\mathbf{y}_5^T \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x} \right) & = \nabla_{\mathbf{w}} \left(\mathbf{y}_5^T \text{diag}(\mathbf{C}\mathbf{x})\mathbf{P}\mathbf{w} \right) \\
 & = \mathbf{P}^T \text{diag}(\mathbf{C}\mathbf{x})\mathbf{y}_5 \\
 & = \mathbf{P}^T \mathbf{C}\mathbf{x} \odot \mathbf{y}_5
 \end{aligned}$$

Update \mathbf{z}_1 :

$$\begin{aligned}
 \mathbf{z}_1 \leftarrow \arg \min_{\mathbf{z}_1 \in S_b} & \mathbf{y}_2^\top (\mathbf{x} - \mathbf{z}_1) + \frac{\rho_2}{2} \|\mathbf{x} - \mathbf{z}_1\|_2^2 \\
 \mathbf{z}_1 \leftarrow \arg \min_{\mathbf{z}_1 \in S_b} & \left\| \mathbf{z}_1 - \left(\mathbf{x} + \frac{\mathbf{y}_2}{\rho_2} \right) \right\|_2^2 \\
 \mathbf{z}_1 = & \mathbf{P}_{S_b} \left(\mathbf{x} + \frac{\mathbf{y}_2}{\rho_2} \right)
 \end{aligned}$$

Here, we need to perform a simple projection onto the box set S_b . The projection $\mathbf{P}_{S_b}(\cdot)$ is an elementwise clamping between 0 and +1 that is, $\mathbf{P}_{S_b}(a) = \min(\max(a, 0), 1)$ for a scalar value a .

Update \mathbf{z}_2 :

$$\begin{aligned} \mathbf{z}_2 &\leftarrow \arg \min_{\mathbf{z}_2 \in S_2} \mathbf{y}_3^\top (\mathbf{x} - \mathbf{z}_2) + \frac{\rho_3}{2} \|\mathbf{x} - \mathbf{z}_2\|_2^2 \\ \mathbf{z}_2 &\leftarrow \arg \min_{\mathbf{z}_2 \in S_2} \left\| \mathbf{z}_2 - \left(\mathbf{x} + \frac{\mathbf{y}_3}{\rho_3} \right) \right\|_2^2 \\ \mathbf{z}_2 &= \mathbf{P}_{S_2} \left(\mathbf{x} + \frac{\mathbf{y}_3}{\rho_3} \right) \end{aligned}$$

We need to perform a simple projection onto the ℓ_2 - sphere: $S_2 = \left\{ \mathbf{a} \in \mathbb{R}^n : \left\| \mathbf{a} - \frac{1}{2} \mathbf{1} \right\|_2^2 = \frac{n}{4} \right\}$. The projection $\mathbf{P}_{S_2}(\cdot)$ involves an elementwise shift and ℓ_2 vector normalization and thus is given as $\mathbf{P}_{S_2}(\mathbf{a}) = \frac{\sqrt{n}}{2} \left(\frac{\mathbf{a} - \frac{1}{2} \mathbf{1}}{\left\| \mathbf{a} - \frac{1}{2} \mathbf{1} \right\|_2} \right) + \frac{1}{2} \mathbf{1}$, for any vector $\mathbf{a} \in \mathbb{R}^n$.

Update \mathbf{z}_3 :

$$\begin{aligned} \mathbf{z}_3 &\leftarrow \arg \min_{\mathbf{z}_3} y_6 (\mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x}) \\ &\quad + \frac{\rho_6}{2} \left\| \mathbf{z}_3^\top \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} - v \right\|_2^2 - \mathbf{y}_7^\top \mathbf{z}_3 + \\ &\quad \frac{\rho_7}{2} \|\mathbf{x} - \mathbf{z}_3\|_2^2 \end{aligned}$$

The problem is strongly convex quadratic in \mathbf{z}_4 , so we obtain the unique global minimizer by equating the gradient to zero as follows

$$\left[\rho_6 \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} \mathbf{x}^\top \mathbf{E}_2^\top \mathbf{E}_1 + \rho_7 \mathbf{I}_{nk} \right] \mathbf{z}_4 = \mathbf{y}_7 + \rho_7 \mathbf{x} - y_6 \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x} + \rho_6 v \mathbf{E}_1^\top \mathbf{E}_2 \mathbf{x}$$

Update \mathbf{z}_4 :

$$\begin{aligned} \mathbf{z}_4 &\leftarrow \arg \min_{\mathbf{z}_4} y_8 (\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x}) + \frac{\rho_8}{2} \left\| \mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right\|_2^2 - \mathbf{y}_9^\top \mathbf{z}_4 + \\ &\quad \frac{\rho_9}{2} \|\mathbf{x} - \mathbf{z}_4\|_2^2 \end{aligned}$$

The problem is also strongly convex quadratic in \mathbf{z}_5 , so we obtain the unique global minimizer by equating the gradient to zero as follows

$$\left[\rho_8 \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \mathbf{x}^\top \mathbf{E}_4^\top \mathbf{E}_3 + \rho_9 \mathbf{I}_{nk} \right] \mathbf{z}_5 = \mathbf{y}_9 + \rho_9 \mathbf{x} - y_8 \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x}$$

Update $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \mathbf{y}_4, \mathbf{y}_5, \mathbf{y}_6, \mathbf{y}_7, \mathbf{y}_8, \mathbf{y}_9$: Lastly, we need to perform dual gradient ascent to update the dual variables as follows:

$$\begin{aligned} \mathbf{y}_1 &\leftarrow \mathbf{y}_1 + \rho_1 (\Psi^\top \mathbf{x} - \mathbf{1}_n), & \mathbf{y}_2 &\leftarrow \mathbf{y}_2 + \rho_2 (\mathbf{x} - \mathbf{z}_2) \\ \mathbf{y}_3 &\leftarrow \mathbf{y}_3 + \rho_3 (\mathbf{x} - \mathbf{z}_3), & \mathbf{y}_4 &\leftarrow \mathbf{y}_4 + \rho_4 (\mathbf{Q}\mathbf{P}^\top \mathbf{x} - \mathbf{u}) \\ \mathbf{y}_5 &\leftarrow \mathbf{y}_5 + \rho_5 (\mathbf{x} - \mathbf{P}\mathbf{w} \odot \mathbf{C}\mathbf{x}), & \mathbf{y}_6 &\leftarrow \mathbf{y}_6 + \rho_6 \left(\mathbf{z}_3^\top \mathbf{E}_1^\top \right. \\ \mathbf{y}_7 &\leftarrow \mathbf{y}_7 + \rho_7 (\mathbf{x} - \mathbf{z}_3), & \mathbf{y}_8 &\leftarrow \mathbf{y}_8 + \rho_8 \left(\mathbf{z}_4^\top \mathbf{E}_3^\top \mathbf{E}_4 \mathbf{x} \right) \\ \mathbf{y}_9 &\leftarrow \mathbf{y}_9 + \rho_9 (\mathbf{x} - \mathbf{z}_4) \end{aligned}$$

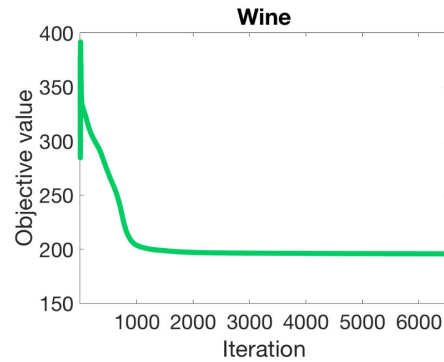


FIGURE 4. Convergence of the K-means objective value using FCKm with random initialization on the Wine dataset. Note the decreasing nature of the objective and its smooth convergence to the solution.

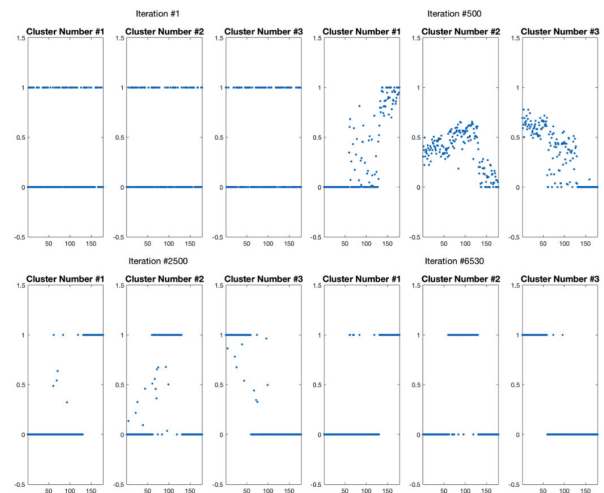


FIGURE 5. Convergence of the solution \mathbf{x} using FCKm with random initialization on the Wine dataset.

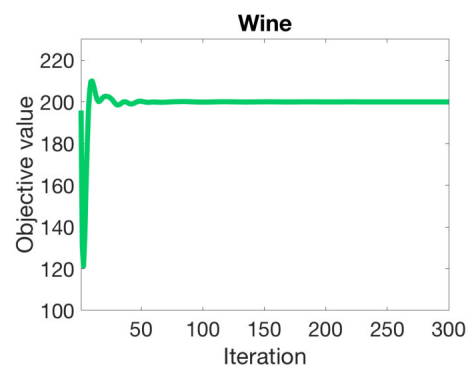


FIGURE 6. Convergence of the K-means objective value using FCKm-Mix with a random initialization on Wine dataset.

AUXILIARY RESULTS

We present some additional experimental results that augment the discussion made in the manuscript. Primarily, we provide empirical evidence that our FCKm method and its constrained variants converge to binary solutions that satisfy different constraints (pairwise and cardinality).

Convergence for Unconstrained Clustering: In Figure 4, we plot the K-means objective value at each ADMM iteration for an unconstrained clustering task with three clusters. Note that the initialization to this problem is a random threeway clustering. The objective decreases monotonically (after the first 2-3 iterations) and converges to a minimum value in approximately 6000 iterations. The optimization is stable with no perturbations at the onset of convergence.

In Figure 5, we plot the cluster vector for each of the three clusters being optimized at different iterations (1, 500, 2500, and 6530), i.e. we plot the three pieces of the label vector x . In the first iteration, the initial clustering is done randomly, so it is binary but it does not lead to a good objective. As ADMM progresses, the continuous solution x becomes more and more binary, until it converges to a feasible binary solution where the three clusters are completely disjoint.

Convergence for Constrained Clustering: In Figure 6, we plot the K-means objective value at each ADMM iteration for a constrained clustering task with three clusters. In this task, we enforce cardinality, must-link, and cannotlink constraints onto the optimization. The initialization is taken to be a random assignment between the three disjoint clusters. In this case, the objective tends to be monotonically increasing after the first few iterations. This might seem counter-intuitive, since we are trying to minimize the objective. However, it must be noted that the continuous solution vector x in each ADMM iteration tends not to be feasible with respect to the enforced constraints. In other words, these constraints are being enforced more and more as the ADMM process proceeds, which forces the tradeoff between objective and feasibility. But, similar to the unconstrained case, the variation in objective is smooth and no perturbations are exhibited when ADMM begins to converge.

Moreover, Figure 4 plots the solution vectors for each cluster at four different iterations (1, 15, 25, 200). A similar behavior to the unconstrained case is encountered, where disjoint binary solution vectors are converged to. However, the notable difference is that we also report the number of cardinality (CardV), must-link (MLV), and cannotlink (CLV) violations at each of these iterations. We see that these violations gradually decrease until convergence occurs, when no violations persist.

ACKNOWLEDGMENT

Adel Bibi was with the King Abdullah University of Science and Technology at the time of this work.

REFERENCES

- [1] G. Karypis, E.-H. Han, and V. Kumar, "Chameleon: Hierarchical clustering using dynamic modeling," *Computer*, vol. 32, no. 8, pp. 68–75, Aug. 1999.
- [2] B. Wu and B. Hu, "Density and neighbor adaptive information theoretic clustering," in *Proc. Int. Joint Conf. Neural Netw.*, Jul. 2011, pp. 230–237.
- [3] O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," in *Proc. Int. Workshop Artif. Intell. Statist. (AISTATS)*, Barbados, Caribbean, Jan. 2005, pp. 57–64.
- [4] K. Wagstaff and C. Cardie, "Clustering with instance-level constraints," in *Proc. AAAI/IAAI*, 2000, pp. 577–584.
- [5] D. Klein, S. D. Kamvar, and C. D. Manning, "From instance-level constraints to space-level constraints: Making the most of prior knowledge in data clustering," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2002, pp. 307–314.
- [6] A. Ng, M. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 14, 2001, pp. 1–15.
- [7] P. Bradley, K. Bennett, and A. Demiriz, "Constrained k -means clustering," Microsoft Research, Redmond, WA, USA, Tech. Rep. MSR-TR-2000-65, 2000.
- [8] Z. Lu and H. H. Ip, "Constrained spectral clustering via exhaustive and efficient constraint propagation," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2010, pp. 1–14.
- [9] U. Von Luxburg, "A tutorial on spectral clustering," *Statist. Comput.*, vol. 17, pp. 395–416, Aug. 2007.
- [10] Z. Lu and T. K. Leen, "Penalized probabilistic clustering," *Neural Comput.*, vol. 19, no. 6, pp. 1528–1567, 2007.
- [11] B.-R. Dai, C. R. Lin, and M. S. Chen, "On the techniques for data clustering with numerical constraints," in *Proc. SIAM Int. Conf. Data Mining*, 2003, pp. 264–268.
- [12] B. Wu, Y. Zhang, B.-G. Hu, and Q. Ji, "Constrained clustering and its application to face clustering in videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 3507–3514.
- [13] K. Wagstaff, C. Cardie, S. Rogers, and S. Schrödl, "Constrained k -means clustering with background knowledge," in *Proc. ICML*, vol. 1, 2001, pp. 577–584.
- [14] F. Hoppner and F. Klawonn, "Clustering with size constraints," in *Computational Intelligence Paradigms*. Berlin, Germany: Springer, 2008, pp. 167–180.
- [15] F. Klawonn and F. Hoppner, "Equi-sized, homogeneous partitioning," in *Proc. Int. Conf. Knowl.-Based Intell. Inf. Eng. Syst.* Cham, Switzerland: Springer, 2006, pp. 70–77.
- [16] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, Aug. 2000.
- [17] B. Wu and B. Ghanem, " ℓ_p -box ADMM: A versatile framework for integer programming," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 7, pp. 1695–1708, Jul. 2019.
- [18] D. H. Fisher, "Knowledge acquisition via incremental conceptual clustering," *Mach. Learn.*, vol. 2, no. 2, pp. 139–172, 1987.
- [19] S. Basu, M. Bilenko, and R. J. Mooney, "A probabilistic framework for semi-supervised clustering," in *Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2004, pp. 59–68.
- [20] B. Wu, S. Lyu, B.-G. Hu, and Q. Ji, "Simultaneous clustering and tracklet linking for multi-face tracking in videos," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 2856–2863.
- [21] I. S. Dhillon, Y. Guan, and B. Kulis, *A Unified View of Kernel K-Means, Spectral Clustering and Graph Cuts*. Princeton, NJ, USA: Citeseer, 2004.
- [22] H. Ding and J. Xu, "A unified framework for clustering constrained data without locality property," in *Proc. 26th Annu. ACM-SIAM Symp. Discrete Algorithms*, Oct. 2015, pp. 808–852.
- [23] T. B. H. Dao, K. C. Duong, and C. Vrain, "A declarative framework for constrained clustering," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*. Cham, Switzerland: Springer, 2013, pp. 419–434.
- [24] J. Peng and Y. Wei, "Approximating K-means-type clustering via semidefinite programming," *SIAM J. Optim.*, vol. 18, no. 1, pp. 186–205, Jan. 2007.
- [25] L. Hubert and P. Arabie, "Comparing partitions," *J. Classification*, vol. 2, no. 1, pp. 193–218, Jan. 1985, doi: [10.1007/BF01908075](https://doi.org/10.1007/BF01908075).
- [26] M. Meilă, "Comparing clusterings—An information based distance," *J. Multivariate Anal.*, vol. 98, no. 5, pp. 873–895, May 2007.
- [27] A. Alqahtani, X. Xie, J. Deng, and M. W. Jones, "A deep convolutional auto-encoder with embedded clustering," in *Proc. 25th IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2018, pp. 4058–4062.

- [28] A. Alqahtani, X. Xie, J. Deng, and M. W. Jones, "Learning discriminatory deep clustering models," in *Proc. Int. Conf. Comput. Anal. Images Patterns*, 2019, pp. 224–233.
- [29] A. Alqahtani, M. Ali, X. Xie, and M. W. Jones, "Deep time-series clustering: A review," *Electronics*, vol. 10, no. 23, p. 3001, Dec. 2021.



ADEL BIBI received the B.Sc. degree (Hons.) in electrical engineering from Kuwait University, in 2014, and the M.Sc. and Ph.D. degrees from the King Abdullah University of Science and Technology (KAUST), in 2016 and 2020, respectively. He is currently a Senior Research Fellow in machine learning and computer vision at the Department of Engineering Science, University of Oxford. He is also a Junior Research Fellow (JRF) of the Kellogg College. Prior to that, he was a Postdoctoral Researcher at the Kellogg College, for a year, since October 2020. He has published more than 20 papers in CVPRs, ECCVs, ICCVs, and ICLR, some which were selected as orals and spotlights. He has been recognized as an Outstanding Reviewer for CVPR 2018, CVPR 2019, ICCV 2019, and ICLR 2022, and won the Best Paper Award at the Optimization and Big Data Conference at the KAUST. He is going to serve as an Area Chair for AAAI 2023.



ALI ALQAHTANI received the Ph.D. degree in computer science from Swansea University, Swansea, U.K., in 2021. He is currently an Assistant Professor with the Department of Computer Science, King Khalid University, Abha, Saudi Arabia. He has published several refereed conference and journal publications. His research interests include various aspects of pattern recognition, deep learning, and machine intelligence and their applications to real-world problems.



BERNARD GHANEM received the bachelor's degree from the American University of Beirut (AUB), in 2005, and the M.S./Ph.D. degrees from the University of Illinois at Urbana–Champaign (UIUC), in 2010. He is currently a Professor with the Electrical and Computer Engineering Program, King Abdullah University of Science and Technology (KAUST). He has coauthored more than 100 peer-reviewed conferences and journal papers in his field as well as three issued patents.

His research interests include computer vision and machine learning with emphasis on topics in video understanding, 3D recognition, and theoretical foundations of deep learning. His work has received several awards and honors, including several Best Workshop Paper Awards in CVPR and ECCV, a Google Faculty Research Award, in 2015 (1st in MENA for Machine Perception), and a Abdul Hameed Shoman Arab Researchers Award for Big Data and Machine Learning, in 2020. He has served as an area chair (AC) for the leading computer vision and machine learning conferences. He serves as an Associate Editor for IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE.

• • •