

---

# Clustering Stable Instances of Euclidean $k$ -means

---

Abhratanu Dutta\*  
Northwestern University  
adutta@u.northwestern.edu

Aravindan Vijayaraghavan\*  
Northwestern University  
aravindv@northwestern.edu

Alex Wang†  
Carnegie Mellon University  
alexwang@u.northwestern.edu

## Abstract

The Euclidean  $k$ -means problem is arguably the most widely-studied clustering problem in machine learning. While the  $k$ -means objective is NP-hard in the worst-case, practitioners have enjoyed remarkable success in applying heuristics like Lloyd’s algorithm for this problem. To address this disconnect, we study the following question: *what properties of real-world instances will enable us to design efficient algorithms and prove guarantees for finding the optimal clustering?* We consider a natural notion called additive perturbation stability that we believe captures many practical instances of Euclidean  $k$ -means clustering. Stable instances have unique optimal  $k$ -means solutions that does not change even when each point is perturbed a little (in Euclidean distance). This captures the property that  $k$ -means optimal solution should be tolerant to measurement errors and uncertainty in the points. We design efficient algorithms that provably recover the optimal clustering for instances that are additive perturbation stable. When the instance has some additional separation, we can design a simple, efficient algorithm with provable guarantees that is also robust to outliers. We also complement these results by studying the amount of stability in real datasets, and demonstrating that our algorithm performs well on these benchmark datasets.

## 1 Introduction

One of the major challenges in the theory of clustering is to bridge the large disconnect between our theoretical and practical understanding of the complexity of clustering. While theory tells us that most common clustering objectives like  $k$ -means or  $k$ -median clustering problems are intractable in the worst case, many heuristics like Lloyd’s algorithm or  $k$ -means++ seem to be effective in practice. In fact, this has led to the “CDNM” thesis [11, 9]: “Clustering is difficult only when it does not matter”.

We try to address the following natural questions in this paper: *Why are real-world instances of clustering easy? Can we identify properties of real-world instances that make them tractable?*

We focus on the Euclidean  $k$ -means clustering problem where we are given  $n$  points  $X = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$ , and we need to find  $k$  centers  $\mu_1, \mu_2, \dots, \mu_k \in \mathbb{R}^d$  minimizing the objective  $\sum_{x \in X} \min_{i \in [k]} \|x - \mu_i\|^2$ . The  $k$ -means clustering problem is the most well-studied objective for clustering points in Euclidean space [18, 3]. The problem is NP-hard in the worst-case [14] even for  $k = 2$ , and a constant factor hardness of approximation is known for larger  $k$  [5].

---

\*Supported by the National Science Foundation (NSF) under Grant No. CCF-1637585.

†Part of the work was done while the author was at Northwestern University.

One way to model real-world instances of clustering problems is through *instance stability*, which is an implicit structural assumption about the instance. Practically interesting instances of  $k$ -means clustering problem often have a clear optimal clustering solution (usually the ground-truth clustering) that is stable: i.e., it remains optimal even under small perturbations of the instance. As argued in [7], clustering objectives like  $k$ -means are often just a proxy for recovering a ground-truth clustering that is close to the optimal solution. Instances in practice always have measurement errors, and optimizing the  $k$ -means objective is meaningful only when the optimal solution is stable to these perturbations.

This notion of stability was formalized independently in a pair of influential works [11, 7]. The predominant strand of work on instance stability assumes that the optimal solution is resilient to multiplicative perturbations of the distances [11]. For any  $\gamma \geq 1$ , a metric clustering instance  $(X, d)$  on point set  $X \subset \mathbb{R}^d$  and metric  $d : X \times X \rightarrow \mathbb{R}_+$  is said to be  $\gamma$ -factor stable iff the (unique) optimal clustering  $C_1, \dots, C_k$  of  $X$  remains the optimal solution for any instance  $(X, d')$  where any (subset) of the distances are increased by up to a  $\gamma$  factor i.e.,  $d(x, y) \leq d'(x, y) \leq \gamma d(x, y)$  for any  $x, y \in X$ . In a series of recent works [4, 8] culminating in [2], it was shown that 2-factor perturbation stable (i.e.,  $\gamma \geq 2$ ) instances of  $k$ -means can be solved in polynomial time.

Multiplicative perturbation stability represents an elegant, well-motivated formalism that captures robustness to measurement errors for clustering problems in general metric spaces ( $\gamma = 1.1$  captures relative errors of 10% in the distances). However, multiplicative perturbation stability has the following drawbacks in the case of Euclidean clustering problems:

- Measurement errors in Euclidean instances are better captured using additive perturbations. Uncertainty of  $\delta$  in the position of  $x, y$  leads to an additive error of  $\delta$  in  $\|x - y\|_2$ , irrespective of how large or small  $\|x - y\|_2$  is.
- The amount of stability  $\gamma$  needed to enable efficient algorithms (i.e.,  $\gamma \geq 2$ ) often imply strong structural conditions, that are unlikely to be satisfied by many real-world datasets. For instance,  $\gamma$ -factor perturbation stability implies that every point is a multiplicative factor of  $\gamma$  closer to its own center (say  $\mu_i$ ) than to any other cluster center  $\mu_j$ .
- Algorithms that are known to have provable guarantees under multiplicative perturbation stability are based on single-linkage or MST algorithms that are very non-robust by nature. In the presence of a few outliers or noise, any incorrect decision in the lower layers gets propagated up to the higher levels.

In this work, we consider a natural additive notion of stability for Euclidean instances, when the optimal  $k$ -means clustering solution does not change even where each point is moved by a small Euclidean distance of up to  $\delta$ . Moving each point by at most  $\delta$  corresponds to a small additive perturbation to the pairwise distances between the points<sup>3</sup>. Unlike multiplicative notions of perturbation stability [11, 4], this notion of additive perturbation is not scale invariant. Hence the normalization or scale of the perturbation is important.

Ackerman and Ben-David [1] initiated the study of additive perturbation stability when the distance between any pair of points can be changed by at most  $\delta = \varepsilon \text{diam}(X)$  with  $\text{diam}(X)$  being the *diameter* of the whole dataset. The algorithms take time  $n^{O(k/\varepsilon^2)} = n^{O(k \text{diam}^2(X)/\delta^2)}$  and correspond to polynomial time algorithms when  $k, 1/\varepsilon$  are constants. However, this dependence of  $k \text{diam}^2(X)/\delta^2$  in the exponent is not desirable since the diameter is a very non-robust quantity — the presence of one outlier (that is even far away from the decision boundary) can increase the diameter arbitrarily. Hence, these guarantees are useful mainly when the whole instance lies within a small ball and for a small number of clusters [1, 10]. Our notion of additive perturbation stability will use a different scale parameter that is closely related to the distance between the centers, instead of the diameter  $\text{diam}(X)$ . As we will see soon, our results for additive perturbation stability have no explicit dependence on the diameter, and allows instances to have potentially unbounded clusters (as in the case of far-way outliers). Further with some additional assumptions, we also obtain polynomial time algorithmic guarantees for large  $k$ .

---

<sup>3</sup>Note that not all additive perturbations to the distances can be captured by an appropriate movement of the points in the cluster. Hence the notion we consider in our paper is a weaker assumption on the instance.




Figure 1: a) **Left:** the figure shows an instance with  $k = 2$  satisfying  $\varepsilon$ -APS with  $D$  being separation between the means. The half-angle of the cone is  $\arctan(1/\varepsilon)$  and the distance between  $\mu_1$  and the apex of the cone ( $\Delta$ ) is at most  $D/2$ . b) **Right:** The figure shows a  $(\rho, \Delta, \varepsilon)$ -separated instance, with scale parameter  $\Delta$ . All the points lie inside the cones of half-angle  $\arctan(1/\varepsilon)$ , whose apexes are separated by a margin of at least  $\rho$ .

### 1.1 Additive Perturbation Stability and Our Contributions

We consider a notion of additive stability where the points in the instance can be moved by at most  $\delta = \varepsilon D$ , where  $\varepsilon \in (0, 1)$  is a parameter, and  $D = \max_{i \neq j} D_{ij} = \max_{i \neq j} \|\mu_i - \mu_j\|_2$  is the maximum distance between pairs of means. Suppose  $X$  is a  $k$ -means clustering instance with optimal clustering  $C_1, C_2, \dots, C_k$ . We say that  $X$  is  $\varepsilon$ -APS (additive perturbation stable) iff every  $\delta = \varepsilon D$ -additive perturbation of  $X$  has  $C_1, C_2, \dots, C_k$  as an optimal clustering solution. (See Definition 2.3 for a formal definition). Note that there is no restriction on the diameter of the instance, or even the diameters of the individual clusters. Hence, our notion of additive perturbation stability allows the instance to be unbounded.

**Geometric property of  $\varepsilon$ -APS instances.** Clusters in the optimal solution of an  $\varepsilon$ -APS instance satisfy a natural geometric condition, that implies an “**angular separation**” between every pair of clusters.

**Proposition 1.1** (Geometric Implication of  $\varepsilon$ -APS). *Consider an  $\varepsilon$ -APS instance  $X$ , and let  $C_i, C_j$  be two clusters of the optimal solution. Any point  $x \in C_i$  lies in a cone whose axis is along the direction  $(\mu_i - \mu_j)$  with half-angle  $\arctan(1/\varepsilon)$ . Hence if  $u$  is the unit vector along  $\mu_i - \mu_j$  then*

$$\forall x \in C_i, \frac{|\langle x - \frac{\mu_i + \mu_j}{2}, u \rangle|}{\|x - \frac{\mu_i + \mu_j}{2}\|_2} > \frac{\varepsilon}{\sqrt{1 + \varepsilon^2}}. \quad (1)$$

For any  $j \in [k]$ , all the points in cluster  $C_i$  lie inside the cone with its axis along  $(\mu_i - \mu_j)$  as in Figure 1. The distance between  $\mu_i$  and the apex of the cone is  $\Delta = (\frac{1}{2} - \varepsilon)D$ . We will call  $\Delta$  the **scale parameter** of the clustering.

We believe that many clustering instances in practice satisfy  $\varepsilon$ -APS condition for reasonable constants  $\varepsilon$ . In fact, our experiments in Section 4 suggest that the above geometric condition is satisfied for reasonable values e.g.,  $\varepsilon \in (0.001, 0.2)$ .

While the points can be arbitrarily far away from their own means, the above angular separation (1) is crucial in proving the polynomial time guarantees for our algorithms. For instance, this implies that at least  $1/2$  of the points in a cluster  $C_i$  are within a Euclidean distance of at most  $O(\Delta/\varepsilon)$  from  $\mu_i$ . This geometric condition (1) of the dataset enables the design of a tractable algorithm for  $k = 2$  with provable guarantees. This algorithm is based on a modification of the perceptron algorithm in supervised learning, and is inspired by [13].

**Informal Theorem 1.2.** *For any fixed  $\varepsilon > 0$ , there exists an  $dn^{\text{poly}(1/\varepsilon)}$  time algorithm that correctly clusters all  $\varepsilon$ -APS 2-means instances.*

For  $k$ -means clustering, similar techniques can be used to learn the separating halfspace for each pair of clusters. But this incurs an exponential dependence on  $k^2$ , which renders this approach

inefficient for large  $k$ .<sup>4</sup> We now consider a natural strengthening of this assumption that allows us to get  $\text{poly}(n, d, k)$  guarantees.

**Angular Separation with additional Margin Separation.** We consider a natural strengthening of additive perturbation stability where there is an additional margin between any pair of clusters. This is reminiscent of margin assumptions in supervised learning of halfspaces and spectral clustering guarantees of Kumar and Kannan [15] (see Section 1.2). Consider a  $k$ -means clustering instance  $X$  having an optimal solution  $C_1, C_2, \dots, C_k$ . This instance is  $(\rho, \Delta, \varepsilon)$ -separated iff for each  $i \neq j \in [k]$ , the subinstance induced by  $C_i, C_j$  has parameter scale  $\Delta$ , and all points in the clusters  $C_i, C_j$  lie inside cones of half-angle  $\arctan(1/\varepsilon)$ , which are separated by a margin of at least  $\rho$ . This is implied by the stronger condition that the subinstance induced by  $C_i, C_j$  is  $\varepsilon$ -additive perturbation stable with scale parameter  $\Delta$  even when  $C_i$  and  $C_j$  are moved towards each other by  $\rho$ . Please see Figure 1 for an illustration. We formally define  $(\rho, \Delta, \varepsilon)$ -separated stable instances geometrically in Section 2.

**Informal Theorem 1.3** (Polytime algorithm for  $(\rho, \Delta, \varepsilon)$ -separated instances). *There is an algorithm running in time<sup>5</sup>  $\tilde{O}(n^2 kd)$  that given any instance  $X$  that is  $(\rho, \Delta, \varepsilon)$ -separated with  $\rho \geq \Omega(\Delta/\varepsilon^2)$  recovers the optimal clustering  $C_1, \dots, C_k$ .*

A formal statement of the theorem (with unequal sized clusters), and its proof are given in Section 3. We prove these polynomial time guarantees for a new, simple algorithm (Algorithm 3.1). The algorithm constructs a graph with one vertex for each point, and edges between points that within a distance of at most  $r$  (for an appropriate threshold  $r$ ). The algorithm then finds the  $k$ -largest connected components. It then uses the  $k$  empirical means of these  $k$  components to cluster all the points.

In addition to having provable guarantees, the algorithm also seems efficient in practice, and performs well on standard clustering datasets. Experiments that we conducted on some standard clustering datasets in UCI suggest that our algorithm manages to almost recover the ground truth, and achieves a  $k$ -means objective cost that is very comparable to Lloyd’s algorithm and  $k$ -means++ (see Section 4).

In fact, our algorithm can also be used to initialize the Lloyd’s algorithm: our guarantees show that when the instance is  $(\rho, \Delta, \varepsilon)$ -separated, one iteration of Lloyd’s algorithm already finds the optimal clustering. Experiments suggest that our algorithm finds initializers of smaller  $k$ -means cost compared to the initializers of  $k$ -means++ [3] and also recover the ground-truth to good accuracy (see Section 4 and Supplementary material for details).

**Robustness to Outliers.** Perturbation stability requires the optimal solution to remain completely unchanged under any valid perturbation. In practice, the stability of an instance may be dramatically reduced by a few outliers. We show provable guarantees for a slight modification of the algorithm, in the setting when an  $\eta$ -fraction of the points can be arbitrary outliers, and do not lie in the stable regions. More formally, we assume that we are given an instance  $X \cup Z$  where there is an (unknown) set of points  $Z$  with  $|Z| \leq \eta|X|$  such that  $X$  is a  $(\rho, \Delta, \varepsilon)$ -separated-stable instance. Here  $\eta n$  is assumed to be smaller than size of the smallest cluster by a constant factor. This is similar to robust perturbation resilience considered in [8, 16]. Our experiments in Section 4 indicate that the stability or separation can increase a lot after ignoring a few points close to the margin.

In what follows,  $w_{\max} = \max|C_i|/n$  and  $w_{\min} = \min|C_i|/n$  are the maximum and minimum weight of clusters, and  $\eta < w_{\min}$ .

**Theorem 1.4.** *Given  $X \cup Z$  where  $X$  satisfies  $(\rho, \Delta, \varepsilon)$ -separated with  $\eta < w_{\min}$  and*

$$\rho = \Omega\left(\frac{\Delta}{\varepsilon^2} \left(\frac{w_{\max} + \eta}{w_{\min} - \eta}\right)\right)$$

*and  $\eta < w_{\min}$ , there is a polynomial time algorithm running in time  $\tilde{O}(n^2 dk)$  that returns a clustering consistent with  $C_1, \dots, C_k$  on  $X$ .*

This robust algorithm is effectively the same as Algorithm 3.1 with one additional step that removes all low-degree vertices in the graph. This step removes bad outliers in  $Z$  without removing too many points from  $X$ .

<sup>4</sup>We remark that the results of [1] also incur an exponential dependence on  $k$ .

<sup>5</sup>The  $\tilde{O}$  hides an inverse Ackerman function of  $n$ .

## 1.2 Comparisons to Other Related Work

Awasthi et al. [4] showed that  $\gamma$ -multiplicative perturbation stable instance also satisfied the notion of  $\gamma$ -center based stability (every point is a  $\gamma$ -factor closer to its center than to any other center) [4]. They showed that an algorithm based on the classic single linkage algorithm works under this weaker notion when  $\gamma \geq 3$ . This was subsequently improved by [8], and the best result along these lines [2] gives a polynomial time algorithm that works for  $\gamma \geq 2$ . A robust version of  $(\gamma, \eta)$ -perturbation resilience was explored for center-based clustering objectives [8]. As such, the notions of additive perturbation stability, and  $(\rho, \Delta, \varepsilon)$ -separated instances are incomparable to the various notions of multiplicative perturbation stability. Further as argued in [9], we believe that additive perturbation stability is more realistic for Euclidean clustering problems.

Ackerman and Ben-David [1] initiated the study of various deterministic assumptions for clustering instances. The measure of stability most related to this work is Center Perturbation (CP) clusterability (an instance is  $\delta$ -CP-clusterable if perturbing the centers by a distance of  $\delta$  does not increase the cost much). A subtle difference is their focus on obtaining solutions with small objective cost [1], while our goal is to recover the optimal clustering. However, the main qualitative difference is how the length scale is defined – this is crucial for additive perturbations. The run time of the algorithm in [1] is  $n^{\text{poly}(k, \text{diam}(X)/\delta)}$ , where the length scale of the perturbations is  $\text{diam}(X)$ , the diameter of the whole instance. Our notion of additive perturbations uses a much smaller length-scale of  $\Delta$  (essentially the inter-mean distance; see Prop. 1.1 for a geometric interpretation), and Theorem 1.2 gives a run-time guarantee of  $n^{\text{poly}(\Delta/\delta)}$  for  $k = 2$  (Theorem 1.2 is stated in terms of  $\varepsilon = \delta/\Delta$ ). By using the largest inter-mean distance instead of the diameter as the length scale, our algorithmic guarantees can also handle unbounded clusters with arbitrarily large diameters and outliers.

The exciting results of Kumar and Kannan [15] and Awasthi and Sheffet [6] also gave deterministic margin-separation condition, under which spectral clustering (PCA followed by  $k$ -means) finds the optimum clusters under deterministic conditions about the data. Suppose  $\sigma = \|X - C\|_{\text{op}}^2/n$  is the “spectral radius” of the dataset, where  $C$  is the matrix given by the centers. In the case of equal-sized clusters, the improved results of [6] proves approximate recovery of the optimal clustering if the margin  $\rho$  between the clusters along the line joining the centers satisfies  $\rho = \Omega(\sqrt{k}\sigma)$ . Our notion of margin  $\rho$  in  $(\rho, \Delta, \varepsilon)$ -separated instances is analogous to the margin separation notion used by the above results on spectral clustering [15, 6]. In particular, we require a margin of  $\rho = \Omega(\Delta/\varepsilon^2)$  where  $\Delta$  is our scale parameter, with no extra  $\sqrt{k}$  factor. However, we emphasize that the two margin conditions are incomparable, since the spectral radius  $\sigma$  is incomparable to the scale parameter  $\Delta$ .

We now illustrate the difference between these deterministic conditions by presenting a couple of examples. Consider an instance with  $n$  points drawn from a mixture of  $k$  Gaussians in  $d$  dimensions with identical diagonal covariance matrices with variance 1 in the first  $O(1)$  coordinates and roughly  $\frac{1}{d}$  in the others, and all the means lying in the subspace spanned by these first  $O(1)$  co-ordinates. In this setting, the results of [15, 6] require a margin separation of at least  $\sqrt{k \log n}$  between clusters. On the other hand, these instances satisfy our geometric conditions with  $\varepsilon = \Omega(1)$ ,  $\Delta \sim \sqrt{\log n}$  and therefore our algorithm only needs a margin separation of  $\rho \sqrt{\log n}$  (hence, saving a factor of  $\sqrt{k}$ ).<sup>6</sup> However, if the  $n$  points were drawn from a mixture of spherical Gaussians in high dimensions (with  $d \gg k$ ), then the margin condition required for [15, 6] is weaker.

## 2 Stability definitions and geometric properties

$X \subseteq \mathbb{R}^d$  will denote a  $k$ -means clustering instance and  $C_1, \dots, C_k$  will often refer to its optimal clustering. It is well-known that given a cluster  $C$  the value of  $\mu$  minimizing  $\sum_{x \in C} \|x - \mu\|^2$  is given by  $\mu = \frac{1}{|C|} \sum_{x \in C} x$ , the mean of the points in the set. We give more preliminaries about the  $k$ -means problem in the Supplementary Material.

### 2.1 Balance Parameter

We define an instance parameter,  $\beta$ , capturing how balanced a given instance’s clusters are.

<sup>6</sup>Further, while algorithms for learning GMM models may work here, adding some outliers far from the decision boundary will cause many of these algorithms to fail, while our algorithm is robust to such outliers.




Figure 2: An example of the family of perturbations considered by Lemma 2.4. Here  $v$  is in the upwards direction. If  $a$  is to the right of the diagonal solid line, then  $a'$  will be to the right of the slanted dashed line and will lie on the wrong side of the separating hyperplane.

**Definition 2.1** (Balance parameter). Given an instance  $X$  with optimal clustering  $C_1, \dots, C_k$ , we say  $X$  satisfies balance parameter  $\beta \geq 1$  if for all  $i \neq j$ ,  $\beta|C_i| > |C_j|$ .

We will write  $\beta$  in place of  $\beta(X)$  when the instance is clear from context.

## 2.2 Additive perturbation stability

**Definition 2.2** ( $\varepsilon$ -additive perturbation). Let  $X = \{x_1, \dots, x_n\}$  be a  $k$ -means clustering instance with optimal clustering  $C_1, C_2, \dots, C_k$  whose means are given by  $\mu_1, \mu_2, \dots, \mu_k$ . Let  $D = \max_{i,j} \|\mu_i - \mu_j\|$ . We say that the instance  $X' = \{x'_1, \dots, x'_n\}$  is an  $\varepsilon$ -additive perturbation of  $X$  if for all  $i$ ,  $\|x'_i - x_i\| \leq \varepsilon D$ .

**Definition 2.3** ( $\varepsilon$ -additive perturbation stability). Let  $X$  be a  $k$ -means clustering instance with optimal clustering  $C_1, C_2, \dots, C_k$ . We say that  $X$  is  $\varepsilon$ -additive perturbation stable (APS) if every  $\varepsilon$ -additive perturbation of  $X$  has unique optimal clustering given by  $C_1, C_2, \dots, C_k$ .

Intuitively, the difficulty of the clustering task increases as the stability parameter  $\varepsilon$  decreases. For example, when  $\varepsilon = 0$  the set of  $\varepsilon$ -APS instances contains any instance with a unique solution. In the following we will only consider  $\varepsilon > 0$ .

## 2.3 Geometric implication of $\varepsilon$ -APS

Let  $X$  be an  $\varepsilon$ -APS  $k$ -means clustering instance such that each cluster has at least 4 points. Fix  $i \neq j$  and consider a pair of clusters  $C_i, C_j$  with means  $\mu_i, \mu_j$  and define the following notations.

- Let  $D_{i,j} = \|\mu_i - \mu_j\|$  be the distance between  $\mu_i$  and  $\mu_j$  and let  $D = \max_{i,j'} \|\mu_i - \mu_{j'}\|$  be the maximum distance between any pair of means.
- Let  $u = \frac{\mu_i - \mu_j}{\|\mu_i - \mu_j\|}$  be the unit vector in the intermean direction. Let  $V = u^\perp$  be the space orthogonal to  $u$ . For  $x \in \mathbb{R}^d$ , let  $x_{(u)}$  and  $x_{(V)}$  be the projections  $x$  onto  $u$  and  $V$ .
- Let  $p = \frac{\mu_i + \mu_j}{2}$  be the midpoint between  $\mu_i$  and  $\mu_j$ .

A simple perturbation that we can use will move all points in  $C_i$  and  $C_j$  along the direction  $\mu_i - \mu_j$  by a  $\delta$  amount, while another perturbation moves these points along  $\mu_j - \mu_i$ ; these will allow us to conclude that  $\exists$  margin of size  $2\delta$ . To establish Proposition 1.1, we will choose a clever  $\varepsilon$ -perturbation that allows us to show that clusters must live in *cone* regions (see figure 1 left). This perturbation chooses two clusters and moves their means in opposite directions orthogonal to  $u$  while moving a single point towards the other cluster (see figure 2). The following lemma establishes Proposition 1.1.

**Lemma 2.4.** For any  $x \in C_i \cup C_j$ ,  $\|(x - p)_{(V)}\| < \frac{1}{\varepsilon} (\|(x - p)_{(u)}\| - \varepsilon D_{i,j})$ .

*Proof.* Let  $v \in V$  be a unit vector perpendicular to  $u$ . Without loss of generality, let  $a, b, c, d \in C_i$  be distinct. Note that  $D_{i,j} \leq D$  and consider the  $\varepsilon$ -additive perturbation given by

$$X' = \{a - \delta u, b + \delta u, c - \delta v, d - \delta v\} \cup \{x - \frac{\delta}{2}v \mid x \in C_i \setminus \{a, b, c, d\}\} \cup \{x + \frac{\delta}{2}v \mid x \in C_j\}$$

and  $X \setminus \{C_i \cup C_j\}$  where  $\delta = \varepsilon D_{i,j}$  (see figure 2). By assumption,  $\{C_i, C_j\}$  remains the optimal clustering of  $C_i \cup C_j$ . We have constructed  $X'$  such that the new means are at  $\mu'_i = \mu_i - \frac{\varepsilon D_{i,j}}{2} v$  and  $\mu'_j = \mu_j + \frac{\varepsilon D_{i,j}}{2} v$ , and the midpoint between the means is  $p' = p$ . The halfspace containing  $\mu'_i$  given by the linear separator between  $\mu'_i$  and  $\mu'_j$  is  $\langle x - p', \mu'_i - \mu'_j \rangle > 0$ . Hence, as  $a'$  is classified correctly by the  $\varepsilon$ -APS assumption,

$$\begin{aligned} \langle a' - p', \mu'_i - \mu'_j \rangle &= \langle a - p - \varepsilon D_{i,j} u, D_{i,j} u - \varepsilon D_{i,j} v \rangle \\ &= D_{i,j} (\langle a - p, u \rangle - \varepsilon \langle a - p, v \rangle - \varepsilon D_{i,j}) > 0 \end{aligned}$$

Then noting that  $\langle u, a - p \rangle$  is positive, we have that  $\langle a - p, v \rangle < \frac{1}{\varepsilon} (\|(a - p)_{(u)}\| - \varepsilon D_{i,j})$ .  $\square$

Note that this property follows from perturbations which only affect two clusters at a time. Our results follow from this weaker notion.

## 2.4 $(\rho, \Delta, \varepsilon)$ -separation

Motivated by Lemma 2.4, we define a geometric condition where the angular separation and margin separation are parametrized separately. This notion of separation is implied by a stronger stability assumption where any pair of clusters is  $\varepsilon$ -APS with scale parameter  $\Delta$  even after being moved towards each other a distance of  $\rho$ .

We say that a pair of clusters is  $(\rho, \Delta, \varepsilon)$ -separated if their points lie in cones with axes along the intermean direction, half-angle  $\arctan(1/\varepsilon)$ , and apexes at distance  $\Delta$  from their means and at least  $\rho$  from each other (see figure 1 right). Formally, we require the following.

**Definition 2.5** (Pairwise  $(\rho, \Delta, \varepsilon)$ -separation). Given a pair of clusters  $C_i, C_j$  with means  $\mu_i, \mu_j$ , let  $u = \frac{\mu_i - \mu_j}{\|\mu_i - \mu_j\|}$  be the unit vector in the intermean direction and let  $p = (\mu_i + \mu_j)/2$ . We say that  $C_i$  and  $C_j$  are  $(\rho, \Delta, \varepsilon)$ -separated if  $D_{i,j} \geq \rho + 2\Delta$  and for all  $x \in C_i \cup C_j$ ,

$$\|(x - p)_{(v)}\| \leq \frac{1}{\varepsilon} (\|(x - p)_{(u)}\| - (D_{i,j}/2 - \Delta)).$$

**Definition 2.6**  $(\rho, \Delta, \varepsilon)$ -separation. We say that an instance  $X$  is  $(\rho, \Delta, \varepsilon)$ -separated if every pair of clusters in the optimal clustering is  $(\rho, \Delta, \varepsilon)$ -separated.

## 3 $k$ -means clustering for general $k$

We assume that our instance has balance parameter  $\beta$ . Our algorithm takes in as input the set of points  $X$  and  $k$ , and outputs a clustering of all the points.

### Algorithm 3.1.

**Input:**  $X = \{x_1, \dots, x_n\}, k$ .

- 1: **for all** pairs  $a, b$  of distinct points in  $\{x_i\}$  **do**
- 2:     Let  $r = \|a - b\|$  be our guess for  $\rho$
- 3:     **procedure** INITIALIZE
- 4:         Create graph  $G$  on vertex set  $\{x_1, \dots, x_n\}$  where  $x_i$  and  $x_j$  have an edge iff  $\|x_i - x_j\| < r$
- 5:         Let  $a_1, \dots, a_k \in \mathbb{R}^d$  where  $a_i$  is the mean of the  $i$ th largest connected component of  $G$
- 6:     **procedure** ASSIGN
- 7:         Let  $C_1, \dots, C_k$  be the clusters obtained by assigning each point in  $X$  to the closest  $a_i$
- 8:         Calculate the  $k$ -means objective of  $C_1, \dots, C_k$
- 9:     Return clustering with smallest  $k$ -means objective found above

---

**Theorem 3.2.** *Algorithm 3.1 recovers  $C_1, \dots, C_k$  for any  $(\rho, \Delta, \varepsilon)$ -separated instance with  $\rho = \Omega\left(\frac{\Delta}{\varepsilon^2} + \frac{\beta\Delta}{\varepsilon}\right)$  and the running time is  $\tilde{O}(n^2 kd)$ .*

We maintain the connected components and their centers via a union-find data structure and keep it updated as we increase  $\rho$  and add edges to the dynamic graph. Since we go over  $n^2$  possible choices of  $\rho$  and each pass takes  $O(kd)$  time, the algorithm runs in  $\tilde{O}(n^2 kd)$ .

The rest of the section is devoted to proving Theorem 3.2. Define the following regions of  $\mathbb{R}^d$  for every pair  $i, j$ . Given  $i, j$ , let  $C_i, C_j$  be the corresponding clusters with means  $\mu_i, \mu_j$ . Let  $u = \frac{\mu_i - \mu_j}{\|\mu_i - \mu_j\|}$  be the unit vector in the inter-mean direction and  $p = \frac{\mu_i + \mu_j}{2}$  be the point between the two means. We first define formally  $S_{i,j}^{(\text{cone})}$  which was described in the introduction (the feasible region) and two other regions of the clusters that will be useful in the analysis (see Figure 1b). We observe that  $C_i$  belongs to the intersection of all the cones  $S_{i,j}^{(\text{cone})}$  over  $j \neq i$ .

**Definition 3.3.**

- $S_{i,j}^{(\text{cone})} = \{x \in \mathbb{R}^d \mid \|(x - (\mu_i - \Delta u))_{(V)}\| \leq \frac{1}{\varepsilon} \langle x - (\mu_i - \Delta u), u \rangle\},$
- $S_{i,j}^{(\text{nice})} = \{x \in S_{i,j}^{(\text{cone})} \mid \langle x - \mu_i, u \rangle \leq 0\},$
- $S_i^{(\text{good})} = \bigcap_{j \neq i} S_{i,j}^{(\text{nice})}.$

The nice area of  $i$  with respect to  $j$  i.e.  $S_{i,j}^{(\text{nice})}$ , is defined as all points in the cap of  $S_{i,j}^{(\text{cone})}$  “above”  $\mu_i$ . The good area of a cluster  $S_i^{(\text{good})}$  is the intersection of its nice areas with respect to all other clusters.

It suffices to prove the following two main lemmas. Lemma 3.4 states that the ASSIGN subroutine correctly clusters all points given an initialization satisfying certain properties. Lemma 3.5 states that the initialization returned by the INITIALIZE subroutine satisfies these properties when we guess  $r = \rho$  correctly. As  $\rho$  is only used as a threshold on edge lengths, testing the distances between all pairs of data points i.e.  $\{\|a - b\| : a, b \in X\}$  suffices.

**Lemma 3.4.** *For a  $(\rho, \Delta, \varepsilon)$ -separated instance with  $\rho = \Omega(\Delta/\varepsilon^2)$ , the ASSIGN subroutine recovers  $C_1, C_2, \dots, C_k$  correctly when initialized with  $k$  points  $\{a_1, a_2, \dots, a_k\}$  where  $a_i \in S_i^{(\text{good})}$ .*

**Lemma 3.5.** *For an  $(\rho, \Delta, \varepsilon)$ -separated instance with balance parameter  $\beta$  and  $\rho = \Omega(\beta\Delta/\varepsilon)$ , the INITIALIZE subroutine outputs one point each from  $\{S_i^{(\text{good})} : i \in [k]\}$  when  $r = \rho$ .*

To prove Lemma 3.5 we define a region of each cluster  $S_i^{(\text{core})}$ , the core, such that most (at least  $\beta/(1 + \beta)$  fraction) of the points in  $C_i$  will belong to the connected component containing  $S_i^{(\text{core})}$ . Hence, any large connected component (in particular, the  $k$  largest ones) must contain the core of one of the clusters. Meanwhile, the margin ensures points across clusters are not connected. Further, since  $S_i^{(\text{core})}$  accounts for most points in  $C_i$ , the angular separation ensures that the empirical mean of the connected component is in  $S_i^{(\text{good})}$ .

## 4 Experimental results

We evaluate Algorithm 3.1 on multiple real world datasets and compare its performance to the performance of  $k$ -means++, and also check how well these datasets satisfy our geometric conditions. See supplementary results for details about ground truth recovery.

**Datasets.** Experiments were run on unnormalized and normalized versions of four labeled datasets from the UCI Machine Learning Repository: Wine ( $n = 178, k = 3, d = 13$ ), Iris ( $n = 150, k = 3, d = 4$ ), Banknote Authentication ( $n = 1372, k = 2, d = 5$ ), and Letter Recognition ( $n = 20,000, k = 26, d = 16$ ). Normalization was used to scale each feature to unit range.

**Performance** We ran Algorithm 3.1 on the datasets. The cost of the returned solution for each of the normalized and unnormalized versions of the datasets is recorded in Table 1 column 2. Our guarantees show that under  $(\rho, \Delta, \varepsilon)$ -separation for appropriate values of  $\rho$  (see section 3), the algorithm will find the optimal clustering after a single iteration of Lloyd’s algorithm. Even when  $\rho$  does not satisfy our requirement, we can use our algorithm as an initialization heuristic for Lloyd’s algorithm. We compare our initialization with the  $k$ -means++ initialization heuristic ( $D^2$  weighting). In Table 1, this is compared to the smallest initialization cost of 1000 trials of  $k$ -means++ on each of the datasets, the solution found by Lloyd’s algorithm using our initialization and the smallest  $k$ -means cost of 100 trials of Lloyd’s algorithm using a  $k$ -mean++ initialization.

**Separation in real data sets.** As the ground truth clusterings in our datasets are not in general linearly separable, we consider the clusters given by Lloyd’s algorithm initialized with the ground



Table 1: Comparison of  $k$ -means cost for Alg 3.1 and  $k$ -means++

Dataset	Alg 3.1	$k$ -means++	Alg 3.1 with Lloyd's	$k$ -means++ with Lloyd's
Wine	2.376e+06	2.426e+06	2.371e+06	2.371e+06
Wine (normalized)	48.99	65.50	48.99	48.95
Iris	81.04	86.45	78.95	78.94
Iris (normalized)	7.035	7.676	6.998	6.998
Banknote Auth.	44808.9	49959.9	44049.4	44049.4
Banknote (norm.)	138.4	155.7	138.1	138.1
Letter Recognition	744707	921643	629407	611268
Letter Rec. (norm.)	3367.8	4092.1	2767.5	2742.3

Table 2: Values of  $(\rho, \varepsilon, \Delta)$  satisfied by  $(1 - \eta)$ -fraction of points

Dataset	$\eta$	$\varepsilon$	minimum $\rho/\Delta$	average $\rho/\Delta$	maximum $\rho/\Delta$
Wine	0.1	0.1	0.566	1.5	3.05
		0.01	0.609	1.53	3.07
Iris	0.1	0.1	0.398	4.35	7.7
		0.01	0.496	5.04	9.06
Banknote Auth.	0.1	0.1	0.264	0.264	0.264
		0.01	0.398	0.398	0.398
Letter Recognition	0.1	0.1	0.018	2.19	7.11
		0.01	0.378	3.07	11.4

truth solutions. *Values of  $\varepsilon$  for Lemma 2.4.* We calculate the maximum value of  $\varepsilon$  such that a given pair of clusters satisfies the geometric condition in Proposition 1.1. The results are collected in the Supplementary material in Table 3. We see that the average value of  $\varepsilon$  lies approximately in the range  $(0.01, 0.1)$ .

*Values of  $(\rho, \Delta, \varepsilon)$ -separation.* We attempt to measure the values of  $\rho$ ,  $\Delta$ , and  $\varepsilon$  in the datasets. For  $\eta = 0.05, 0.1$ ,  $\varepsilon = 0.1, 0.01$ , and a pair of clusters  $C_i, C_j$ , we calculate  $\rho$  as the maximum margin separation a pair of axis-aligned cones with half-angle  $\arctan(1/\varepsilon)$  can have while capturing a  $(1 - \eta)$ -fraction of all points. For some datasets and values for  $\eta$  and  $\varepsilon$ , there may not be any such value of  $\rho$ , in this case we leave the row blank. The results for the unnormalized datasets with  $\eta = 0.1$  are collected in Table 2. (See Supplementary material for the full table).

## 5 Conclusion and Future Directions

We studied a natural notion of additive perturbation stability, that we believe captures many real-world instances of Euclidean  $k$ -means clustering. We first gave a polynomial time algorithm when  $k = 2$ . For large  $k$ , under an additional margin assumption, we gave a fast algorithm of independent interest that provably recovers the optimal clustering under these assumptions (in fact the algorithm is also robust to noise and outliers). An appealing aspect of this algorithm that it is not tailored towards the model; our experiments indicate that this algorithm works well in practice even when the assumptions do not hold. Our results with the margin assumption hence gives an algorithm that (A) has provable guarantees (under reasonable assumptions) (B) is efficient and practical (C) is robust to errors.

While the margin assumption seems like a realistic assumption qualitatively, we believe that the exact condition we assume is not optimal. An interesting open question is understanding whether such a margin is necessary for designing tractable algorithms for large  $k$ . We conjecture that for higher  $k$ , clustering remains hard even with  $\varepsilon$  additive perturbation resilience (without any additional margin assumption). Improving the margin condition or proving lower bounds on the amount of additive stability required are interesting future directions.

## References

- [1] Margareta Ackerman and Shai Ben-David. Clusterability: A theoretical study. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5, pages 1–8. PMLR, 2009.
- [2] Haris Angelidakis, Konstantin Makarychev, and Yury Makarychev. Algorithms for stable and perturbation-resilient problems. In *Symposium on Theory of Computing (STOC)*, 2017.
- [3] David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 1027–1035, 2007.
- [4] Pranjal Awasthi, Avrim Blum, and Or Sheffet. Center-based clustering under perturbation stability. *Information Processing Letters*, 112(1–2):49 – 54, 2012.
- [5] Pranjal Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop. The hardness of approximation of euclidean k-means. In *Symposium on Computational Geometry*, pages 754–767, 2015.
- [6] Pranjal Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 37–49, 2012.
- [7] Maria-Florina Balcan, Avrim Blum, and Anupam Gupta. Approximate clustering without the approximation. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 1068–1077, 2009.
- [8] Maria Florina Balcan and Yingyu Liang. Clustering under perturbation resilience. *SIAM Journal on Computing*, 45(1):102–155, 2016.
- [9] Shai Ben-David. Computational feasibility of clustering under clusterability assumptions. *CoRR*, abs/1501.00437, 2015.
- [10] Shalev Ben-David and Lev Reyzin. Data stability in clustering: A closer look. *Theoretical Computer Science*, 558:51 – 61, 2014. Algorithmic Learning Theory.
- [11] Yonatan Bilu and Nathan Linial. Are stable instances easy? In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 332–341, 2010.
- [12] Hans-Dieter Block. The perceptron: A model for brain functioning. *Reviews of Modern Physics*, 34(1):123–135, 1962.
- [13] Avrim Blum and John Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *Proceedings of Symposium on Discrete Algorithms (SODA)*, 2002.
- [14] Sanjoy Dasgupta. *The hardness of k-means clustering*. Department of Computer Science and Engineering, University of California, San Diego, 2008.
- [15] Amit Kumar and Ravindran Kannan. Clustering with spectral norm and the k-means algorithm. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 299–308. IEEE, 2010.
- [16] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Bilu-linial stable instances of max cut. *Proc. 22nd Symposium on Discrete Algorithms (SODA)*, 2014.
- [17] A.B.J Novikoff. On convergence proofs on perceptrons. *Proceedings of the Symposium on the Mathematical Theory of Automata*, XII(1):615–622, 1962.
- [18] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, New York, NY, USA, 1st edition, 2011.

## Appendix A Introduction to the $k$ -means clustering problem

In the  $k$ -means clustering problem, we are given  $n$  points  $\{x_1, \dots, x_n\}$  in  $\mathbb{R}^d$  and need to find  $k$  centers  $\mu_1, \mu_2, \dots, \mu_k \in \mathbb{R}^d$  minimizing

$$\sum_x \min_{i \in [k]} \|x - \mu_i\|^2.$$

Choosing the centers  $\mu_1, \mu_2, \dots, \mu_k$  fixes the clustering  $C_1, C_2, \dots, C_k$  given by  $C_i = \{x \mid \forall j, \|x - \mu_i\| \leq \|x - \mu_j\|\}$ . We can rewrite the objective as

$$\sum_{i \in [k]} \sum_{x \in C_i} \|x - \mu_i\|^2.$$

On the other hand, given a cluster  $C_i$  the value of  $\mu$  minimizing  $\sum_{x \in C_i} \|x - \mu\|^2$  is given by  $\mu = \frac{1}{|C_i|} \sum_{x \in C_i} x$ , the mean of the points in the set. Thus, we can reformulate the problem as finding a clustering  $C_1, C_2, \dots, C_k$  of  $\{x_i\}$  minimizing

$$\sum_{i \in [k]} \sum_{y \in C_i} \left\| y - \left( \frac{1}{|C_i|} \sum_{x \in C_i} x \right) \right\|^2.$$

## Appendix B Exact poly-time algorithm for $\varepsilon$ -APS 2-means clustering

The following algorithm is inspired by work in [ ] showing that the perceptron algorithm runs in poly-time with high probability in the smoothed analysis setting.

### B.1 Review of perceptron algorithm

Suppose  $y_1, \dots, y_n$  is a sequence of labeled  $\{+1, -1\}$ -samples consistent with a linear threshold function, i.e., there exists vector  $w^*$  such that the labeling function  $\ell(y_i)$  is consistent with  $\text{sgn}(\langle w^*, y_i \rangle)$ . At time  $t = 0$ , the perceptron algorithm sets  $w_0 = 0$ . At each subsequent time step, the algorithm sees sample  $y_t$ , outputs  $\text{sgn}(\langle w_{t-1}, y_t \rangle)$  as its guess for  $\ell(y_t)$ , sees the true label  $\ell(y_t)$ , and updates  $w_t$ . On a correct guess,  $w_t = w_{t-1}$ , and on a mistake  $w_t = w_{t-1} + \ell(y_t)y_t / \|y_t\|$ .

The following well-known theorem [12, 17] bounds the number of total mistakes the perceptron algorithm can make in terms of its angular margin.

**Theorem B.1.** *The number of mistakes made by the perceptron algorithm is bounded above by  $(1/\gamma)^2$  for*

$$\gamma = \min_{i \in [n]} \frac{|\langle w^*, y_i \rangle|}{\|w^*\| \|y_i\|}.$$

For a universe  $U$  of elements and a function  $f : U \rightarrow \mathbb{Z}_{\geq 0}$ , we will denote by  $(U, f)$  the multiset where  $u \in U$  appears in the multiset  $f(u)$ -many times. The size of a multiset is  $\sum_{u \in U} f(u)$ .

**Lemma B.2.** *There exists a multiset  $M = (\{y_1, \dots, y_n\}, f)$  of size at most  $(1/\gamma)^2$  such that  $\sum_{y \in M} \ell(y) \frac{y}{\|y\|}$  correctly classifies all of  $\{y_1, \dots, y_n\}$ .*

*Proof.* Let  $r = (1/\gamma)^2 + 1$ . Consider the performance of the perceptron algorithm on  $r$  consecutive runs of the  $y_1, \dots, y_n$ , i.e., let the input be

$$\underbrace{y_1, \dots, y_n}_{1 \text{ run}}, y_1, \dots, y_n, \dots, y_1, \dots, y_n \underbrace{\phantom{y_1, \dots, y_n}}_{r \text{ runs}}.$$

If no mistakes are made on a given run, no mistakes will be made on any subsequent run. Thus if a mistake is made on a given run, at least one mistake must have been made on every previous run. Suppose the perceptron algorithm makes a mistake on the  $r$ th run, then the algorithm must have made at least  $(1/\gamma)^2 + 1$  mistakes, a contradiction. Hence the direction of  $w$  after  $r$  runs correctly classifies all of  $\{y_1, \dots, y_n\}$ . The value of  $w$  is  $\sum_{i \in [n]} f(y_i) \ell(y_i) \frac{y_i}{\|y_i\|}$  where  $f(y_i)$  is the number of times  $y_i$  was misclassified.  $\square$

## B.2 The algorithm

Fix the following notation for this section: let  $I = \{x_1, \dots, x_n\}$  for  $x_i \in \mathbb{R}^d$  be an  $\varepsilon$ -APS 2-means clustering instance such that each cluster has at least 4 points. Let  $D = \|\mu_1 - \mu_2\|$ ,  $u = \frac{\mu_1 - \mu_2}{\|\mu_1 - \mu_2\|}$ ,  $p = \frac{\mu_1 + \mu_2}{2}$ .

### Algorithm B.3.

---

**Input:**  $X = \{x_1, \dots, x_n\}, \varepsilon$

- 1: **for all** pairs  $a, b$  of distinct points in  $\{x_i\}$  **do**
- 2:     Let  $\delta = \|a - b\|$
- 3:     Let  $X_{a,b} = \{y_1, \dots, y_n\}$  be an instance given by  $y_i = (x_i - a, \delta) \in \mathbb{R}^{d+1}$
- 4:     **for all** multisets  $M$  of size at most  $c_1^{-2}\varepsilon^{-8}$  and assignments  $\ell : M \rightarrow \{\pm 1\}$  **do**
- 5:         Let  $w = \sum_{y \in M} \ell(y) \frac{y}{\|y\|}$
- 6:         Calculate  $k$ -means cost of  $C_1 = \{x_i \mid \langle w, y_i \rangle \geq 0\}, C_2 = \{x_i \mid \langle w, y_i \rangle < 0\}$ .
- 7: Return clustering with smallest  $k$ -means objective found above

---

## B.3 Geometric conditions

The following two lemmas follow directly from Lemma 2.4.

**Lemma B.4.** *For any  $x \in C_i \cup C_j$ ,*

$$|\langle u, x - p \rangle| > \varepsilon D_{i,j}.$$

*In particular, for  $x \in C_i$ ,  $\langle u, x - p \rangle > \varepsilon D_{i,j}$  and for  $x \in C_j$ ,  $\langle u, x - p \rangle < -\varepsilon D_{i,j}$ .*

**Lemma B.5.** *For any  $x \in C_i \cup C_j$ ,*

$$\frac{|\langle u, x - p \rangle|}{\|x - p\|} > \sqrt{\frac{\varepsilon^2}{1 + \varepsilon^2}}.$$

The following statement says that the nice area of any cluster  $C_i$  with respect to another cluster  $C_j$  is nonempty.

**Lemma B.6.** *For any  $i \neq j$ ,  $S_{i,j}^{(\text{nice})}$  is nonempty.*

*Proof.* We have that  $\frac{1}{|C_i|} \sum_{x \in C_i} (x - p) = \mu_i - p = (\Delta_{i,j}/2)u$ . Then

$$\frac{1}{|C_i|} \sum_{x \in C_i} \langle u, x - p \rangle = \Delta_{i,j}/2$$

where each term in the summation is positive. Thus there is some  $x \in C_i$  such that  $|\langle u, x - p \rangle| \leq \Delta_{i,j}/2$ .  $\square$

Note that Lemmas B.4 and B.6 together imply that we *can not* have an instance with  $\varepsilon \geq 1/2$ .

**Lemma B.7.** *There is no  $\varepsilon$ -APS  $k$ -means clustering instance for  $\varepsilon \geq 1/2$ .*

## B.4 Proof of correctness of Algorithm B.3

Note that each new instance  $X_{a,b}$  has labeling consistent with some linear threshold function. We have that  $\ell(y_i) = \ell(x_i) = \text{sgn}(\langle u, x_i - p \rangle) = \text{sgn}(\langle u, x_i - a \rangle + \langle u, a - p \rangle)$ . Then taking  $w^* = (u, \langle u, a - p \rangle / \delta)$ , we have that  $\ell(y_i) = \text{sgn}(\langle w^*, y_i \rangle)$ .

**Lemma B.8.** *Let  $a \in S_{1,2}^{(\text{nice})}, b \in S_{2,1}^{(\text{nice})}$ . Then,*

$$(2\varepsilon)D < \|a - b\| < \left( \sqrt{\frac{1 + \varepsilon^2}{\varepsilon^2}} \right) D.$$

*Proof.* For the first inequality,  $\|a - b\| \geq |\langle u, a - b \rangle| = |\langle u, a - p \rangle - \langle u, b - p \rangle|$ . Then by Lemma B.4,  $\|a - b\| > 2\varepsilon D$ .

For the second inequality,  $\|a - b\| \leq \|a - p\| + \|p - b\|$ . As  $a \in S_{1,2}^{(\text{nice})}$ , we have  $|\langle u, a - p \rangle| \leq \Delta/2$ , hence by Lemma B.5,  $\|a - p\| < \sqrt{(1 + \varepsilon^2)/\varepsilon^2} D/2$ . Similarly,  $\|b - p\| < \sqrt{(1 + \varepsilon^2)/\varepsilon^2} D/2$ .  $\square$

The following lemma states that the angular separation  $\gamma$  in  $I_{a,b}$  is large when  $a \in S_{1,2}^{(\text{nice})}, b \in S_{2,1}^{(\text{nice})}$ .

**Lemma B.9.** *There exists constant  $c_1$  such that for any instance  $I_{a,b} = \{y_i\}$  corresponding to  $a \in S_{1,2}^{(\text{nice})}, b \in S_{2,1}^{(\text{nice})}$ ,*

$$\gamma = \min_{i \in [n]} \frac{|\langle w^*, y_i \rangle|}{\|w^*\| \|y_i\|} > c_1 \varepsilon^4.$$

*Proof.* We bound each term in the minimization individually. Let  $i \in [n]$ , then

$$\frac{|\langle w^*, y_i \rangle|}{\|w^*\| \|y_i\|} = \frac{|\langle u, x_i - p \rangle|}{\sqrt{1 + \left(\frac{\langle u, a - p \rangle}{\delta}\right)^2} \sqrt{\delta^2 + \|x_i - a\|^2}}.$$

We first observe the following facts.

- From Lemma B.5,  $|\langle u, x_i - p \rangle| > \sqrt{\frac{\varepsilon^2}{1 + \varepsilon^2}} \|x_i - p\| > \varepsilon \|x_i - p\|$
- As  $a$  is in  $S_{1,2}^{(\text{nice})}$ ,  $|\langle u, a - p \rangle| \leq \frac{D}{2} < \frac{\delta}{4\varepsilon}$
- As  $a$  is in  $S_{1,2}^{(\text{nice})}$ , by Lemma B.5,  $\|x_i - a\|^2 \leq 2\|x_i - p\|^2 + 2\|p - a\|^2 < 2\|x_i - p\|^2 + \frac{1}{2} \frac{1 + \varepsilon^2}{\varepsilon^2} D^2$
- From Lemma B.8,  $\delta^2 < \frac{1 + \varepsilon^2}{\varepsilon^2} D^2$ ,
- From Lemma B.4,  $\|x_i - p\| > \varepsilon D$

Making each of the substitutions above,

$$\begin{aligned} \frac{|\langle w^*, y_i \rangle|}{\|w^*\| \|y_i\|} &> \varepsilon \frac{\|x_i - p\|}{\sqrt{1 + \frac{1}{16\varepsilon^2}} \sqrt{2\|x_i - p\|^2 + \frac{3}{2} \frac{1 + \varepsilon^2}{\varepsilon^2} D^2}} \\ &> \varepsilon \frac{1}{\sqrt{1 + \frac{1}{16\varepsilon^2}} \sqrt{2 + \frac{3}{2} \frac{1 + \varepsilon^2}{\varepsilon^2} \left(\frac{D}{\|x_i - p\|}\right)^2}} \\ &> \varepsilon \frac{1}{\sqrt{1 + \frac{1}{16\varepsilon^2}} \sqrt{2 + \frac{3}{2\varepsilon^2} + \frac{3}{2\varepsilon^4}}}. \end{aligned}$$

Then, completing both squares,

$$\begin{aligned} \frac{|\langle w^*, y_i \rangle|}{\|w^*\| \|y_i\|} &> \varepsilon \frac{1}{\left(1 + \frac{1/4}{\varepsilon}\right) \left(\sqrt{2} + \frac{\sqrt{3/2}}{\varepsilon^2}\right)} \\ &= \varepsilon^4 \frac{1}{(\varepsilon + 1/4) \left(\sqrt{2}\varepsilon^2 + \sqrt{3/2}\right)} \end{aligned}$$

As  $\varepsilon < 1/2$  by Lemma B.7, we can bound the fraction below by some constant  $c_1 \approx 0.845$ .  $\square$

Finally, we prove Theorem 1.2 restated below.

**Theorem.** *There exists a universal constant  $c \geq 1$  such that for any fixed  $\varepsilon > 0$ , there exists an  $n^{O((1/\varepsilon)^c)}d$  time algorithm that correctly clusters all  $\varepsilon$ -APS 2-means instances.*

*Proof.* The correctness of Algorithm B.3 for all  $\varepsilon$ -APS 2-means clustering instances in which each cluster has at least 4 points follows from Lemmas B.6, B.2, and B.9. On the other hand, the optimal 2-means clustering where one of the clusters has at most 3 points can be calculated in  $O(n^4d)$  time. An algorithm that returns the better of these two solutions then correctly clusters all  $\varepsilon$ -APS 2-means instances.  $\square$

## Appendix C $k$ -means clustering for general $k$

We assume that our instance has balance parameter  $\beta$ . Our algorithm takes in as input the set of points  $X$  and  $k$ , and outputs a clustering of all the points.

### Algorithm C.1.

**Input:**  $X = \{x_1, \dots, x_n\}, k$ .

- 1: **for all** pairs  $a, b$  of distinct points in  $\{x_i\}$  **do**
- 2:     Let  $r = \|a - b\|$  be our guess for  $\rho$
- 3:     **procedure** INITIALIZE
- 4:         Create graph  $G$  on vertex set  $\{x_1, \dots, x_n\}$  where  $x_i$  and  $x_j$  have an edge iff  $\|x_i - x_j\| < r$
- 5:         Let  $a_1, \dots, a_k \in \mathbb{R}^d$  where  $a_i$  is the mean of the  $i$ th largest connected component of  $G$
- 6:     **procedure** ASSIGN
- 7:         Let  $C_1, \dots, C_k$  be the clusters obtained by assigning each point in  $X$  to the closest  $a_i$
- 8:         Calculate the  $k$ -means objective of  $C_1, \dots, C_k$
- 9:     Return clustering with smallest  $k$ -means objective found above

The main theorem of this section states that the above poly-time algorithm exactly recovers optimal clusters  $C_1, C_2, \dots, C_k$  for instances with  $\rho$  large enough in terms of  $\varepsilon$ ,  $\Delta$ , and  $\beta$ .

**Theorem C.2.** *Algorithm C.1 recovers  $C_1, \dots, C_k$  for any  $(\rho, \Delta, \varepsilon)$ -separated instance with  $\rho = \Omega\left(\frac{\Delta}{\varepsilon^2} + \frac{\beta\Delta}{\varepsilon}\right)$ .*

The rest of the section is devoted to proving Theorem C.2. Define the following regions of  $\mathbb{R}^d$  for every pair  $i, j$ . Given  $i, j$ , let  $C_i, C_j$  be the corresponding clusters with means  $\mu_i, \mu_j$ . Let  $u = \frac{\mu_i - \mu_j}{\|\mu_i - \mu_j\|}$  be the unit vector in the inter-mean direction and  $p = \frac{\mu_i + \mu_j}{2}$  be the point between the two means. We first define formally  $S_{i,j}^{(\text{cone})}$  which was described in the introduction (the feasible region) and two other regions of the clusters that will be useful in the analysis (see Figure 1b). We observe that  $C_i$  belongs to the intersection of all the cones  $S_{i,j}^{(\text{cone})}$  over  $j \neq i$ .

**Definition C.3.** Let  $S_{i,j}^{(\text{cone})} = \{x \in \mathbb{R}^d \mid \|(x - (\mu_i - \Delta u))_{(V)}\| < \frac{1}{\varepsilon} \langle x - (\mu_i - \Delta u), u \rangle\}$ . The nice area of  $i$  with respect to  $j$  is defined as all points in the cap of  $S_{i,j}^{(\text{cone})}$  “above”  $\mu_i$  i.e.,  $S_{i,j}^{(\text{nice})} = \{x \in S_{i,j}^{(\text{cone})} \mid \langle x - \mu_i, u \rangle \leq 0\}$ . Finally, the good area of  $C_i$  is the intersection of its nice areas i.e.,  $S_i^{(\text{good})} = \bigcap_{j \neq i} S_{i,j}^{(\text{nice})}$ .

It suffices to prove the following two main lemmas. Lemma C.4 states that the ASSIGN subroutine correctly clusters all points given an initialization satisfying certain properties. Lemma C.5 states that the initialization returned by the INITIALIZE subroutine satisfies these properties when we guess  $r = \rho$  correctly. As  $\rho$  is only used as a threshold on edge lengths, testing the distances between all pairs of data points i.e.  $\{\|a - b\| : a, b \in X\}$  suffices.

**Lemma C.4.** *For a  $(\rho, \Delta, \varepsilon)$ -separated instance with  $\rho = \Omega(\Delta/\varepsilon^2)$ , the ASSIGN subroutine recovers  $C_1, C_2, \dots, C_k$  correctly when initialized with  $k$  points  $\{a_1, a_2, \dots, a_k\}$  where  $a_i \in S_i^{(\text{good})}$ .*

**Lemma C.5.** *For a  $(\rho, \Delta, \varepsilon)$ -separated instance with balance parameter  $\beta$  and  $\rho = \Omega(\beta\Delta/\varepsilon)$ , the INITIALIZE subroutine outputs one point each from  $\{S_i^{(\text{good})} : i \in [k]\}$  when  $r = \rho$ .*

**Proof of Lemma C.4.** We will show that for any  $a_i \in S_{i,j}^{(\text{nice})}$ ,  $a_j \in S_{j,i}^{(\text{nice})}$ , and  $x \in C_i$ ,  $x$  is closer to  $a_i$  than to  $a_j$ . The following simple lemma describes the properties of the perpendicular bisector between  $a_i$  and  $a_j$ . All these statements follow from the definitions of the nice regions and the angular separation (proof in Supplementary material).

**Lemma C.6.** Suppose  $\rho \geq \Delta/\varepsilon^2$ . Then, for  $a_i \in S_{i,j}^{(\text{nice})}$  and  $a_j \in S_{j,i}^{(\text{nice})}$ , we have

$$(1) \|(a_i - a_j)_{(u)}\| > \frac{\|(a_i - a_j)_{(V)}\|}{\varepsilon}. \quad (2) \left\langle \frac{a_i + a_j}{2}, u \right\rangle \leq \frac{\Delta}{2}. \quad (3) \left\| \left( \frac{a_i + a_j}{2} \right)_{(V)} \right\| \leq \Delta/\varepsilon.$$

*Proof.*

1. We have  $\|(a_i - a_j)_{(V)}\| \leq 2\Delta/\varepsilon$ . On the other hand,  $\rho \leq \|(a_i - a_j)_{(u)}\|$ . Thus the inequality holds for  $\rho > 2\Delta/\varepsilon^2$ .
2.  $\langle a_i + a_j, u \rangle = \langle a_i, u \rangle + \langle a_j, u \rangle \leq D_{i,j}/2 + (-D_{i,j}/2 + \Delta) = \Delta$ . Multiplying by  $1/2$  gives the desired inequality.
3.  $\|(a_i + a_j)_{(V)}\| \leq \|(a_i)_{(V)}\| + \|(a_j)_{(V)}\| < 2\Delta/\varepsilon$ . Multiplying by  $1/2$  gives the desired inequality.

□

We have all the properties we require from points in the nice region and we prove Lemma C.4 next. To do so we first rewrite the condition  $\|x - a_i\| < \|x - a_j\|$  as  $\frac{1}{2}\langle x - (a_i + a_j), a_i - a_j \rangle > 0$ . Then we write each vector in terms of their projection on  $u$  and  $V$  and use Lemma C.6 to bound each of the terms.

*Proof of Lemma C.4.* It suffices to show that for any  $a_i \in S_{i,j}^{(\text{nice})}$ ,  $a_j \in S_{j,i}^{(\text{nice})}$ , and  $x \in C_i$ ,  $\|x - a_i\| < \|x - a_j\|$ . Then by Lemma C.6 above,

$$\begin{aligned} \left\langle x - \frac{a_i + a_j}{2}, a_i - a_j \right\rangle &= \langle x_{(u)}, (a_i - a_j)_{(u)} \rangle + \langle x_{(V)}, (a_i - a_j)_{(V)} \rangle \\ &\quad - \frac{1}{2} \left\langle ((a_i + a_j))_{(u)}, (a_i - a_j)_{(u)} \right\rangle - \frac{1}{2} \left\langle ((a_i + a_j))_{(V)}, (a_i - a_j)_{(V)} \right\rangle \\ &> \|x_{(u)}\| \|(a_i - a_j)_{(u)}\| - \frac{1}{\varepsilon} \left( \|x_{(u)}\| - \frac{\rho}{2} \right) \varepsilon \|(a_i - a_j)_{(u)}\| \\ &\quad - \frac{\Delta}{2} \|(a_i - a_j)_{(u)}\| - \frac{\Delta}{\varepsilon} \|(a_i - a_j)_{(u)}\| \\ &= \frac{\rho}{2} - \frac{3}{2} \Delta > 0 \end{aligned}$$

where the inequality follows because of equality on the first term and Cauchy-Schwarz on the rest. So, for all  $a_i \in S_{i,j}^{(\text{nice})}$ ,  $a_j \in S_{j,i}^{(\text{nice})}$ , and  $x \in C_i$ ,  $x$  is closer to  $a_i$  than  $a_j$ . □

**Proof of Lemma C.5.** We now show that INITIALIZE finds one point in each of the  $k$  good areas. We first start by defining the “core region” of each cluster.

**Definition C.7** ( $S^{(\text{core})}$ ). Let  $S_i^{(\text{core})} = \{x \in \mathbb{R}^d \mid \|x - \mu_i\| < \Delta/\varepsilon\}$ .

The core region is defined in such a way that most (at least  $\beta/(1 + \beta)$  fraction) of the points in  $C_i$  will belong to the connected component containing  $S_i^{(\text{core})}$ . Hence, any large connected component (in particular, the  $k$  largest ones) must contain the core of one of the clusters. Further, since this component accounts for most points in  $C_i$ , the angular separation ensures that the empirical mean of the connected component is in  $S_i^{(\text{good})}$ . Before we move on to the proof, we state some properties of the connected components.

**Lemma C.8.**

1. Any connected component only contains points from a single cluster.

2. For all  $i, j$ ,  $S_i^{(\text{core})} \supseteq S_{i,j}^{(\text{nice})}$ . In particular,  $S_i^{(\text{core})} \cap X$  is nonempty.
3. Fix  $C_i, C_j$ . Let  $A_{i,j} = \{x \in C_i \mid \langle x - \mu_i, \frac{\mu_i - \mu_j}{\|\mu_i - \mu_j\|} \rangle \leq \beta\Delta\}$ . Then,  $|A_{i,j}| \geq \frac{\beta}{1+\beta}|C_i|$ .
4. For all  $i, j$ ,  $A_{i,j}$  is connected in  $G$ .
5. For all  $i$ , the largest component  $K_i$  in cluster  $C_i$  contains  $A_{i,j}$  for all  $j \neq i$ .

*Proof.*

1. Let  $x \in C_i$  and  $y \in C_j$ . Then  $\|x - y\| \geq |\langle x - y, u \rangle| > \rho$ , thus no edge connecting points in different clusters is added to  $G$ .
2. For  $x \in S_{i,j}^{(\text{nice})}$ ,  $\|(x - \mu_i)_{(V)}\| < \frac{1}{\varepsilon}(\Delta - \|(x - \mu_i)_{(u)}\|)$ , hence  $\|x - \mu_i\| < \Delta/\varepsilon$ . An argument similar to that in Lemma B.6 shows that  $S_{i,j}^{(\text{nice})} \cap X$  is nonempty and hence  $S_{i,j}^{(\text{core})} \cap X$  is nonempty.
3.  $\mu_i$  is the mean of the points in cluster  $C_i$ . By an averaging argument,  $|A_{i,j}|\Delta - (|C_i| - |A_{i,j}|)\beta\Delta \geq 0$ . Rearranging,  $|A_{i,j}| \geq \frac{\beta}{1+\beta}|C_i|$ .
4. We show that all of  $A_{i,j}$  is connected. Recall  $S_{i,j}^{(\text{nice})}$  is nonempty; let  $x \in S_{i,j}^{(\text{nice})}$ . For  $y \in A_{i,j}$ ,  $\|x - y\|^2 = \|(x - y)_{(u)}\|^2 + \|(x - y)_{(V)}\|^2 < ((\beta + 1)\Delta)^2 + ((\beta + 1)\Delta/\varepsilon)^2$ . Thus for  $\rho = \Omega(\beta\Delta/\varepsilon)$ , all of  $A_{i,j}$  is connected.
5. By part 2,  $S_i^{(\text{core})} \cap I$  is nonempty; let  $a_i \in S_i^{(\text{core})} \cap I$  and let  $K_i$  be the connected component containing  $a_i$ . As  $S_i^{(\text{core})} \supseteq S_{i,j}^{(\text{nice})}$ , and  $S_{i,j}^{(\text{nice})} \cap I$  is nonempty,  $K_i$  contains  $A_{i,j}$  by the above argument. The size of this connected component is at least  $\beta/(1 + \beta)|C_i|$ . As  $\beta \geq 1$ ,  $K_i$  is the largest component in the cluster.

□

Lemma C.9 states that each of the  $k$  largest component must lie in a different cluster which proves that each of the  $k$  means  $\{a_1, \dots, a_k\}$  come from a different cluster. While Lemma C.10 states that each  $a_i$  lie inside a good region. Together, they imply Lemma C.5, i.e. each  $a_i$  comes from a different good area.

**Lemma C.9.** *The set of  $k$  largest components of  $G$  contains the largest component of each cluster.*

*Proof.* Let  $K_i$  be the largest component in  $C_i$  and let  $K'_j$  be a component in  $C_j$  that is not the largest. Then by the  $\beta$  parameter,  $|K_i| \geq \frac{\beta}{1+\beta}|C_i| > \frac{1}{1+\beta}|C_j| \geq |K'_j|$ . It follows that the  $k$  largest connected components are  $K_1, K_2, \dots, K_k$ . □

**Lemma C.10.** *The mean of points in  $K_i$  lies in  $S_i^{(\text{good})}$ .*

*Proof.* Fix  $i$  and consider  $K_i$ . For  $j \neq i$ ,  $K_i \subseteq A_{i,j}$ . Additionally,  $K_i \subseteq S_{i,j}^{(\text{cone})}$  a convex set. It follows that the mean of points in  $K_i$  lies in  $S_{i,j}^{(\text{cone})}$ . As  $K_i \supseteq S_i^{(\text{core})} \supseteq S_{i,j}^{(\text{nice})}$ , the points not contained in  $K_i$  have  $\langle x - \mu_i, u \rangle > 0$  and hence the mean of points in  $K_i$  lies in  $S_{i,j}^{(\text{nice})}$ . As this holds for each  $j \neq i$ , the mean of points in  $K_i$  lies in  $S_i^{(\text{good})}$ . □

## Appendix D Robust $k$ -means

We show in this section that a simple extension of algorithm 3.1 does well even in the presence of adversarial noise for instances with  $(\rho, \Delta, \varepsilon)$ -separation for large enough  $\rho$ . Specifically, we consider the following model.



Let  $X = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$  be a  $k$ -means clustering instance with optimal clustering  $C_1, \dots, C_k$ . We call  $X$  the set of *pure* points. An additional set of at most  $\eta n$ -many *impure* points  $Z \subset \mathbb{R}^d$  is added by an adversary. Our goal is to find a clustering of  $X \cup Z$  that agrees with  $C_1, \dots, C_k$  on the pure points.

Let  $w_{\max} = \max |C_i|/n$  and let  $w_{\min} = \min |C_i|/n$  be the maximum and minimum weight of clusters. We will assume that  $\eta < w_{\min}$ .

---

**Algorithm D.1.**

---

**Input:**  $X \cup Z, r, t$

- 1: **procedure** INITIALIZE
  - 2:   Create graph  $G$  on  $I \cup Z$  where points are connected by an edge iff they are at distance  $\leq r$
  - 3:   Remove vertices with vertex degree  $\leq t$
  - 4:   Let  $a_1, \dots, a_k \in \mathbb{R}^d$  where  $a_i$  is the mean of the  $i$ th largest connected component of  $G$
  - 5: **procedure** ASSIGN
  - 6:   Let  $C_1, \dots, C_k$  be the clusters obtained by assigning each point in  $I \cup Z$  to the closest  $a_i$
- 

**Theorem D.2.** *Given  $X \cup Z$  where  $X$  satisfies  $(\rho, \Delta, \varepsilon)$ -separation for*

$$\rho = \Omega\left(\frac{\Delta}{\varepsilon^2} \left(\frac{w_{\max} + \eta}{w_{\min} - \eta}\right)\right)$$

*and  $\eta < w_{\min}$ , there exists values of  $r, t$  such that Algorithm D.1 returns a clustering consistent with  $C_1, \dots, C_k$  on  $X$ . Algorithm D.1 runs in time  $O(n^2 d)$ .*

Just as in Theorem 3.2, this algorithm uses  $r$  and  $t$  as thresholds. Hence, it is possible to guess  $r$  from the  $\binom{n}{2}$  pairwise edge lengths and  $t$  from  $[n]$ .

Fix the following parameters.

$$\alpha = 2 \left( \frac{w_{\max} + \eta}{w_{\min} - \eta} \right), \quad r = (\alpha + 1)\Delta(1 + 2/\varepsilon), \quad t = w_{\min}\eta \frac{\alpha}{\alpha + 1}.$$

**Definition D.3.** Define the following regions.

- For  $i \neq j$ , let  $S_{i,j}^{(\text{e nice})} = \{x \in S_{i,j}^{(\text{cone})} \mid \langle x - \mu_i, u \rangle \leq \alpha\}$ .
- For  $i \neq j$ , let  $S_{i,j}^{(\text{r e nice})} = \{x \in \mathbb{R}^d \mid d(x, S_{i,j}^{(\text{e nice})}) \leq r\}$  be the robust nice region.
- Let  $S_i^{(\text{r good})} = \bigcap_{j \neq i} S_{i,j}^{(\text{r e nice})}$  be the robust good region.

It suffices to prove the following two lemmas.

**Lemma D.4.** *Given  $X \cup Z$  where  $X$  satisfies the  $(\rho, \Delta, \varepsilon)$ -condition for  $\rho = \Omega(\alpha\Delta/\varepsilon^2)$  and  $\eta < w_{\min}$ , the ASSIGN subroutine recovers  $C_1, \dots, C_k$  correctly when initialized with  $k$  points  $\{a_1, \dots, a_k\}$  where  $a_i \in S_i^{(\text{r good})}$ .*

**Lemma D.5.** *Given  $X \cup Z$  where  $X$  satisfies the  $(\rho, \Delta, \varepsilon)$ -condition for  $\rho = \Omega(\alpha\Delta/\varepsilon^2)$  and  $\eta < w_{\min}$ , for the choices of  $r$  and  $t$  as above, the INITIALIZE subroutine provides  $k$  points  $\{a_1, \dots, a_k\}$  where  $a_i \in S_i^{(\text{r good})}$ .*

### D.1 Proof of Lemma D.4

We will show that for any  $a_i \in S_i^{(\text{r good})}$ ,  $a_j \in S_j^{(\text{r good})}$  and  $x \in C_i$ ,  $x$  is closer to  $a_i$  than  $a_j$ .

The following lemma describes the properties of the perpendicular bisector between  $a_i$  and  $a_j$ .

**Lemma D.6.** *There is some  $\rho \geq \Omega(\alpha\Delta/\varepsilon^2)$  such that for  $a_i \in S_{i,j}^{(\text{r e nice})}$  and  $a_j \in S_{j,i}^{(\text{r e nice})}$ ,*

1.  $\varepsilon \|(a_i - a_j)_{(u)}\| \geq \|(a_i - a_j)_{(v)}\|$ ,
2.  $\langle \frac{a_i + a_j}{2}, u \rangle < (\alpha + 1)\Delta/2 + r$ ,

$$3. \left\| \left( \frac{a_i + a_j}{2} \right)_{(V)} \right\| \leq (\alpha + 1)\Delta/\varepsilon + r.$$

*Proof.*

1. By triangle inequality,  $\|(a_i - a_j)_{(V)}\| \leq 2((\alpha + 1)\Delta/\varepsilon + r)$ . On the other hand,  $\|(a_i - a_j)_{(u)}\| \geq \rho - 2r$ . Thus the inequality holds for  $\rho \geq 2r + \frac{2}{\varepsilon}((\alpha + 1)\Delta/\varepsilon + r)$ .
2.  $\langle a_i + a_j, u \rangle = \langle a_i, u \rangle + \langle a_j, u \rangle < (D_{i,j}/2 + \alpha\Delta + r) + (-D_{i,j}/2 + \Delta + r) = (\alpha + 1)\Delta + 2r$ . Multiplying by  $1/2$  gives the desired inequality.
3.  $\|(a_i + a_j)_{(V)}\| \leq \|a_{i,(V)}\| + \|a_{j,(V)}\| < 2((\alpha + 1)\Delta/\varepsilon + r)$ . Multiplying by  $1/2$  gives the desired inequality.

□

We now prove Lemma D.4.

*Proof of Lemma D.4.* It suffices to show that for any  $a_i \in S_i^{(\text{r good})}$ ,  $a_j \in S_j^{(\text{r good})}$  and  $x \in C_i$ ,  $\|x - a_i\| < \|x - a_j\|$ . By definition  $S_i^{(\text{r good})} \subseteq S_{i,j}^{(\text{e nice})}$  for all  $j$ . Thus by Lemma D.6 above,

$$\begin{aligned} \left\langle x - \frac{a_i + a_j}{2}, a_i - a_j \right\rangle &= \langle x_{(u)}, (a_i - a_j)_{(u)} \rangle + \langle x_{(V)}, (a_i - a_j)_{(V)} \rangle \\ &\quad - \left\langle \left( \frac{a_i + a_j}{2} \right)_{(u)}, (a_i - a_j)_{(u)} \right\rangle - \left\langle \left( \frac{a_i + a_j}{2} \right)_{(V)}, (a_i - a_j)_{(V)} \right\rangle \\ &\geq \|x_{(u)}\| \|(a_i - a_j)_{(u)}\| - \frac{1}{\varepsilon} (\|x_{(u)}\| - \rho/2) \varepsilon \|(a_i - a_j)_{(u)}\| \\ &\quad - ((\alpha + 1)\Delta/2 + r) \|(a_i - a_j)_{(u)}\| - ((\alpha + 1)\Delta/\varepsilon + r) \varepsilon \|(a_i - a_j)_{(u)}\| \\ &= \rho/2 - \left( \frac{3}{2}(\alpha + 1)\Delta + (1 + \varepsilon)2r \right). \end{aligned}$$

Thus for some  $\rho \geq \Omega(\alpha\Delta/\varepsilon^2)$ , this quantity is greater than zero, i.e., for all  $a_i \in S_i^{(\text{r good})}$ ,  $a_j \in S_j^{(\text{r good})}$ , and  $x \in C_i$ ,  $x$  is closer to  $a_i$  than to  $a_j$ . □

## D.2 Proof of Lemma D.5

We prove Lemma D.5 using the below proof outline.

*Proof outline.*

1. For any  $i \neq j$ , the set of vertices  $S_{i,j}^{(\text{e nice})} \cap X$  forms a clique and the size of this clique is greater than  $t$ . In particular, no vertex in  $S_{i,j}^{(\text{e nice})}$  is deleted.
2. Fix  $i$ . For all  $j \neq i$ , the vertices  $S_{i,j}^{(\text{e nice})} \cap X$  belong to a single connected component. Let  $K_i$  be this connected component.
3. Before vertex deletion (and after), no vertex is adjacent to pure points from different clusters.
4. After vertex deletion, every remaining point lies in  $S_i^{(\text{r good})}$  for some  $i$ . Hence by part 2, every connected component contains pure points from at most a single cluster. In particular,  $K_1, \dots, K_k$  are distinct.
5. Consider  $K_i$  as above and an arbitrary connected component  $K$  not in  $K_1, \dots, K_k$ . Then  $|K_i| > |K|$ . In particular, the  $k$  largest components of  $G$  are  $K_1, \dots, K_k$ .
6. The mean of  $K_i$  lies in  $S_i^{(\text{r good})}$ .

□

*Proof.*

1. The diameter of  $S_{i,j}^{(\text{e nice})}$  is  $\text{diam}(S_{i,j}^{(\text{e nice})}) \leq (\alpha + 1)\Delta(1 + 2/\varepsilon) = r$ . Thus every pair of points in this region is connected. Recall that  $\mu_i$  is the mean of the pure points in cluster  $C_i$ . By an averaging argument,  $|S_{i,j}^{(\text{e nice})} \cap X| \Delta - (|C_i| - |S_{i,j}^{(\text{e nice})} \cap X|)\alpha\Delta \geq 0$ . Rearranging,  $|S_{i,j}^{(\text{e nice})} \cap I| \geq \frac{\alpha}{\alpha+1}|C_i| \geq \frac{\alpha}{\alpha+1}nw_{\min} = t$ .
2. Fix  $j$ . Let  $j \neq i$ . Recall  $S_{i,j}^{(\text{nice})} \cap X$  is nonempty; let  $x \in S_{i,j}^{(\text{nice})} \cap X$ . Then  $\|x - \mu_i\| \leq \Delta/\varepsilon$ . We show that for any  $j' \neq i$ , the connected component containing  $x$  contains  $S_{i,j'}^{(\text{e nice})} \cap X$ . Let  $y \in S_{i,j'}^{(\text{e nice})} \cap X$ . Then  $\|y - x\| \leq \|y - \mu_i\| + \|x - \mu_i\| \leq (\alpha + 1)\Delta/\varepsilon + \alpha\Delta + \Delta/\varepsilon \leq \Delta(\alpha + 1)(1 + 2/\varepsilon) = r$ .
3. Pure points in different clusters are at distance at least  $\rho$  whereas two vertices sharing a neighbor must be at distance at most  $2r$ . Thus the inequality holds for  $\rho \geq \Omega(\alpha\Delta/\varepsilon)$ .
4. Suppose a point  $p$  not in  $\bigcup_i S_i^{(\text{r good})}$  remains after vertex deletion. By part 3 above,  $p$  can only be connected to pure points in a single cluster. Suppose it is connected to pure points in cluster  $C_i$ . By assumption, there exists a  $j$  such that  $p \notin S_{i,j}^{(\text{r e nice})}$ . We bound the degree of  $p$  above by the number of points in  $X \setminus S_{i,j}^{(\text{e nice})}$  and the  $\eta n$ -many impure points, i.e.,  $\deg(p) \leq \eta n + \frac{|C_i|}{\alpha+1} \leq n(\eta + \frac{w_{\max}}{\alpha+1})$ . By our choice of  $t$ , we have that  $\deg(p) < t$ . Thus  $p$  is deleted and all remaining points lie in  $\bigcup_i S_i^{(\text{r good})}$ .  
For any  $i, j$ , the minimum distance between  $S_i^{(\text{r good})}$  and  $S_j^{(\text{r good})}$  is at least  $\rho - 2r$ . For some  $\rho \geq \Omega(\alpha\Delta/\varepsilon)$  then, the distance between these regions is greater than  $\rho - 2r > r$  and no connected component contains pure points from multiple clusters.
5. As in part 2, the size of  $K_i$  is bounded below by the averaging argument  $|K_i| \geq \frac{\alpha}{\alpha+1}|C_i|$ . Let  $K$  belong to cluster  $C_j$ . Then by part 5, the size of the connected component  $K$  is bounded above by the number of remaining points after  $K_j$  is removed and the  $\eta n$ -many impure points, i.e.,  $|C_j| \leq \frac{1}{\alpha+1}|C_j| + \eta n$ . Then by our choice of  $\alpha$ ,  $|K| < |K_i|$ .
6. By part 4,  $K_i \subseteq S_i^{(\text{r good})}$ . As  $S_i^{(\text{r good})}$  is convex, the mean of  $K_i$  also lies in  $S_i^{(\text{r good})}$ .

□

## Appendix E Experimental Results

Table 3: Values of  $\varepsilon$  satisfying Lemma 2.4

Dataset	Minimum $\varepsilon$	Average $\varepsilon$	Maximum $\varepsilon$
Wine	0.0115	0.0731	0.191
Wine (normalized)	0.000119	0.0394	0.107
Iris	0.00638	0.103	0.256
Iris (normalized)	0.00563	0.126	0.237
Banknote Auth.	0.00127	0.00127	0.00127
Banknote (norm.)	0.00175	0.00175	0.00175
Letter Recognition	3.22e-05	0.0593	0.239
Letter Rec. (norm.)	8.49e-06	0.0564	0.247

**Ground truth recovery.** The clustering returned by our algorithm recovers well ( $\approx 97\%$ ) the solution returned by Lloyd’s algorithm initialized with the ground truth for Wine, Iris, and Banknote Authentication across normalized and unnormalized datasets.

Table 4: Values of  $(\rho, \varepsilon, \Delta)$  satisfied by  $(1 - \eta)$ -fraction of points

Dataset	$\eta$	$\varepsilon$	minimum $\rho/\Delta$	average $\rho/\Delta$	maximum $\rho/\Delta$
Wine	0.05	0.1	0.355	0.992	2.19
		0.01	0.374	1	2.2
	0.1	0.1	0.566	1.5	3.05
		0.01	0.609	1.53	3.07
Wine (normalized)	0.05	0.1			
		0.01	0.399	1.06	2.29
	0.1	0.1	0.451	1.3	2.66
		0.01	0.735	1.96	3.62
Iris	0.05	0.1	0.156	2.47	5.37
		0.01	0.263	2.88	6.43
	0.1	0.1	0.398	4.35	7.7
		0.01	0.496	5.04	9.06
Iris (normalized)	0.05	0.1	0.0918	1.89	3.08
		0.01	0.213	2.21	3.4
	0.1	0.1	0.223	3.74	7.12
		0.01	0.391	4.42	8.3
Banknote Auth.	0.05	0.1	0.0731	0.0731	0.0731
		0.01	0.198	0.198	0.198
	0.1	0.1	0.264	0.264	0.264
		0.01	0.398	0.398	0.398
Banknote (norm.)	0.05	0.1			
		0.01	0.197	0.197	0.197
	0.1	0.1	0.246	0.246	0.246
		0.01	0.474	0.474	0.474
Letter Recognition	0.05	0.1			
		0.01	0.168	2.06	6.96
	0.1	0.1	0.018	2.19	7.11
		0.01	0.378	3.07	11.4
Letter Rec. (norm.)	0.05	0.1			
		0.01	0.157	1.97	7.14
	0.1	0.1			
		0.01	0.378	2.92	11.2