
Gradient Based Clustering

Aleksandar Armacki¹ Dragana Bajovic² Dusan Jakovetic³ Soumya Kar¹

Abstract

We propose a general approach for distance based clustering, using the gradient of the cost function that measures clustering quality with respect to cluster assignments and cluster center positions. The approach is an iterative two step procedure (alternating between cluster assignment and cluster center updates) and is applicable to a wide range of functions, satisfying some mild assumptions. The main advantage of the proposed approach is a simple and computationally cheap update rule. Unlike previous methods that specialize to a specific formulation of the clustering problem, our approach is applicable to a wide range of costs, including non-Bregman clustering methods based on the Huber loss. We analyze the convergence of the proposed algorithm, and show that it converges to the set of appropriately defined fixed points, under arbitrary center initialization. In the special case of Bregman cost functions, the algorithm converges to the set of centroidal Voronoi partitions, which is consistent with prior works. Numerical experiments on real data demonstrate the effectiveness of the proposed method.

1. Introduction

Clustering is a fundamental problem in unsupervised learning and is ubiquitous in various applications and domains, (Chandola et al., 2009), (Pediredla & Seelamantula, 2011), (Jain, 2010), (Dhillon et al., 2003). K -means (Lloyd, 1982) is a classical and widely adopted method for clustering. For a given target number K of clusters, K -means proceeds iteratively by alternating between two steps: 1) cluster assignment, i.e., assign each data point to its closest (in terms of the Euclidean distance) cluster; and 2) finding cluster

centers, i.e., position each cluster’s center at the average of the data points currently assigned to the cluster. Besides K -means, popular clustering methods include its improved version K -means++ (Arthur & Vassilvitskii, 2007), as well as K -modes (Huang, 1997), K -medians (Arya et al., 2004), (Arora et al., 1998), etc.

It is well-known, e.g., (Selim & Ismail, 1984), that K -means can be formulated as a joint minimization of a loss function with respect to two groups of variables: 1) binary variables that encode cluster assignments; and 2) continuous variables that designate cluster centers, where the corresponding loss function is a squared Euclidean norm. This K -means representation has motivated a class of new clustering methods called Bregman clustering (Banerjee et al., 2005), where the squared Euclidean norm is replaced with arbitrary Bregman divergence (Bregman, 1967), such as Kullback-Leibler, Mahalanobis, etc. An appealing feature of Bregman clustering is that the introduction of a different loss (other than squared Euclidean) does not harm computational efficiency, as, despite a more involved loss function, the cluster center finding step is still akin to K -means, i.e., it corresponds to computing an average vector.

Several relevant clustering methods have been proposed that also generalize the squared Euclidean norm of K -means and that *do not correspond to a Bregman divergence*. For example, clustering methods based on the Huber loss (Huber, 1964) have been shown to exhibit good clustering performance and exhibit a high degree of robustness to noisy data, (Pediredla & Seelamantula, 2011), (Liu et al., 2019). However, several challenges emerge when generalizing clustering beyond Bregman divergences. First, the cluster center finding step—that corresponds to minimizing the loss with respect to cluster center variables—is no longer an average-finding operation and may be computationally expensive. Second, convergence and stability results for clustering beyond Bregman divergences are limited. For example, reference (Pediredla & Seelamantula, 2011) shows a local convergence to a stationary point, assuming that the algorithm starts from an accurate cluster assignment.

In this paper, we propose a novel generalized clustering algorithm for a broad class of loss functions, and we provide a comprehensive convergence (stability) analysis for the algorithm. The assumed class of losses includes sym-

¹Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA ²Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia ³Faculty of Sciences, University of Novi Sad, Novi Sad, Serbia. Correspondence to: Aleksandar Armacki <aarmacki@andrew.cmu.edu>.

metric Bregman divergences (squared Euclidean norm, Mahalanobis, Jensen–Shannon, etc.), but more importantly, includes non-Bregman losses such as the Huber loss. The main novelty of the algorithm is that, at the cluster center finding step, the exact minimization of the loss function is replaced with a single gradient step with respect to the loss, hence significantly reducing computational cost in general. We prove that the algorithm converges to the appropriately defined stationary points associated with the joint loss with respect to the cluster assignment and cluster center variables, with arbitrary initialization. Numerical experiments on real data demonstrate that involving the cheap cluster center update incurs no or negligible loss both in clustering performance (appropriately measured accuracy) and in iteration-wise convergence speed, hence opening room for significant computational savings. We also show by simulation that the proposed method with the Huber loss exhibits a high degree of robustness to noisy data. While this is in line with prior findings on Huber-based clustering (Liu et al., 2019), (Pediredla & Seelamantula, 2011), the proposed Huber-based method exhibits stronger theoretical convergence guarantees than those offered by the previous work.

We now briefly review the literature to help us contrast the paper with existing work. Gradient based clustering has been explored in the context of the K -means cost in (MacQueen, 1967), (Bottou & Bengio, 1995). (MacQueen, 1967) analyzes a gradient based update rule for K -means, while (Bottou & Bengio, 1995) demonstrate that the standard centroid based solution of the K -means problem is equivalent to performing a Newton’s method in each step. However, their analysis only concerns the squared Euclidean cost. Our work is considerably more general and can be applied to costs such as the Huber loss, or a class of Bregman divergences. (Monath et al., 2017) propose a gradient-based approach for the problem of hierarchical clustering. (Paul et al., 2021) use adaptive gradient methods to design a unified framework for robust center-based clustering, applicable to a large class of Bregman divergences.

A similar approach is used in the robotics community, in the context of coverage control problems, e.g. (Cortes et al., 2004), (Schwager, 2009). However, the focus of their work is on continuous time gradient flow, designed for robot motion in an environment that is typically an infinite set. Additionally, the authors in (Cortes et al., 2004) propose a family of discrete time algorithms, that converge to sets of centroidal Voronoi partitions, if the cost is squared Euclidean distance. On the other hand, our work focuses on a discrete time gradient algorithm, designed for clustering a finite set of points. We explicitly characterize the conditions under which the method converges, and extend the notion of distance to other metrics, beyond the Euclidean distance.

Paper organization. The remainder of the paper is organized as follows. Section 2 formally defines the clustering problem. Section 3 describes the proposed method. Section 4 presents the main results. Section 5 presents an analysis of the fixed points the algorithm converges to. Section 6 presents numerical experiments, and Section 7 concludes the paper. Appendix A contains proofs of the main lemma’s. Appendix B contains proofs of some technical results used throughout the paper. Appendix C contains some additional numerical experiments. The notation used throughout the paper is introduced in the next paragraph.

Notation. \mathbb{R} denotes the set of real numbers, while \mathbb{R}^d denotes the corresponding d -dimensional vector space. More generally, for a vector space V , we denote by V^K its K -dimensional extension. \mathbb{R}_+ denotes the set of non-negative real numbers. We denote by \mathbb{N} the set of non-negative integers. $\|\cdot\| : \mathbb{R}^d \mapsto \mathbb{R}_+$ represents the standard Euclidean norm, while $\langle \cdot, \cdot \rangle : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ denotes the inner product. ∇ denotes the gradient operator, i.e., $\nabla_x f(x, y)$ denotes the gradient of the cost f with respect to variable x . $[N]$ denotes the set of integers up to and including N , i.e., $[N] = \{1, \dots, N\}$. In the algorithm description and throughout the analysis we use subscript to denote the iteration counter, while the value in the parenthesis corresponds to the particular center/cluster. In other words, $x_t(i)$ stands for the i -th cluster center at iteration t . Same holds for clusters, i.e., $C_t(i)$ denotes the i -th cluster at iteration t , corresponding to the subset of the data points assigned to cluster i , at iteration t .

2. Problem formulation

In this section we formalize the clustering problem, and propose a general cost, that subsumes many of the previous clustering formulations.

Let (\mathbb{R}^d, g) represent the standard d -dimensional real vector space, and a corresponding distance function. Let $\mathcal{D} \subset \mathbb{R}^d$ be a finite set, with an associated probability measure $\mu_{\mathcal{D}}$. For some $K > 1$, the problem of clustering the points in \mathcal{D} into K clusters can be cast as

$$\min_{x \in \mathbb{R}^{Kd}} \sum_{y \in \mathcal{D}} p_y \min_{i \in [K]} g(x(i), y)^2, \quad (1)$$

where $x = [x(1)^T, \dots, x(K)^T] \in \mathbb{R}^{Kd}$ represent the candidate cluster centers and $p_y \in (0, 1)^1$, given by $p_y := \mu_{\mathcal{D}}(y)$, represent problem independent weights, that measure the importance of data points $y \in \mathcal{D}$. In the case when g is the standard Euclidean distance, (1) is known in the literature as the K -means problem (Awasthi & Balcan, 2014).

¹Note that, while a standard probability measure can take values in $[0, 1]$, we implicitly assume two things: the support of $\mu_{\mathcal{D}}$ is the whole set \mathcal{D} , and \mathcal{D} contains at least two distinct points.

Another problem similar in nature to (1) is given by

$$\min_{x \in \mathbb{R}^{Kd}} \sum_{y \in \mathcal{D}} p_y \min_{i \in [K]} g(x(i), y), \quad (2)$$

and for g being the Euclidean distance, is known in the literature as K -medians (Arora et al., 1998). Both problems have been well studied, and are known to be NP-hard (Vattani, 2009), (Awasthi et al., 2015), (Megiddo & Supowit, 1984). Many algorithms for solving (1) and (2) exist, guaranteeing convergence to locally optimal solutions, e.g. (Lloyd, 1982), (MacQueen, 1967), (Banerjee et al., 2005), (Telgarsky & Vattani, 2010), (Arora et al., 1998), (Arya et al., 2004). However, all of the algorithms are specialized for solving either the K -means or the K -medians problem, and hence are not generally applicable.

The problems (1), (2), can be equivalently defined as follows. For any $K > 1$, we call $C = (C(1), \dots, C(K))$ a partition of \mathcal{D} , if $\mathcal{D} = \cup_{i \in [K]} C(i)$ and $C(i) \cap C(j) = \emptyset$, for $i \neq j$. Denote by $\mathcal{C}_{K, \mathcal{D}}$ the set of all K -partitions of \mathcal{D} . The clustering problem (1) is then equivalent to

$$\min_{x \in \mathbb{R}^{Kd}, C \in \mathcal{C}_{K, \mathcal{D}}} J(x, C) = \sum_{i \in [K]} \sum_{y \in C(i)} p_y g(x(i), y)^2. \quad (3)$$

The problem (2) can be defined in the same way.

We propose to unify and generalize (1) and (2) as follows. Let $f : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}_+$, be a loss function that satisfies the following assumption.

Assumption 2.1. The loss function f is increasing with respect to the function g , i.e., for all $x, y, z \in \mathbb{R}^d$

$$g(x, y) \leq g(z, y) \text{ implies } f(x, y) \leq f(z, y).$$

We can then define the following general problem

$$\min_{x \in \mathbb{R}^{Kd}, C \in \mathcal{C}_{K, \mathcal{D}}} J(x, C) = \sum_{i=1}^K \sum_{y \in C(i)} p_y f(x(i), y). \quad (4)$$

Remark 2.2. Introducing the function f along with g allows us to naturally decouple the concepts of *cluster shape* and *location of cluster center*. In particular, the function g dictates the cluster shape, while the choice of function f determines the exact location of the cluster centers. We elaborate further on this in Section 5.

Remark 2.3. Compared to (3), the formulation (4) is more general, in the sense that, while the dependence of f on g is maintained, via Assumption 2.1, the function f provides more flexibility, as is illustrated by the following examples.

Example. For the choice $g(x, y) = \|x - y\|$, and $f(x, y) = g(x, y)$, the K -medians formulation is recovered. For the choice $g(x, y) = \|x - y\|$, and $f(x, y) = g(x, y)^2$, the K -means formulation is recovered. For the choice

$g(x, y) = f(x, y)$, being a Bregman distance, the Bregman divergence clustering formulation from (Banerjee et al., 2005) is recovered. For the choice $g(x, y) = \|x - y\|$, and $f(x, y) = \phi_\delta(g(x, y))$, where $\phi_\delta(x)$ is the Huber loss, the formulation from (Pediredla & Seelamantula, 2011) is recovered. We recall that the Huber loss is defined by

$$\phi_\delta(x) = \begin{cases} \frac{x^2}{2}, & |x| \leq \delta \\ \delta|x| - \frac{\delta^2}{2}, & |x| > \delta \end{cases}. \quad (5)$$

3. The proposed method

In this section we outline the proposed method for solving instances of (4) that satisfy some mild assumptions (see ahead Assumptions 3.1-3.5).

To solve (4), an iterative approach is proposed. Starting from an arbitrary initialization x_0 , at every iteration t , it maintains and updates the pair (x_t, C_t) , where $x_t := [x_t(1)^T, x_t(2)^T, \dots, x_t(K)^T]^T \in \mathbb{R}^{Kd}$ and $C_t := (C_t(1), \dots, C_t(K))$ represent stacks of centers and clusters at time $t \in \mathbb{N}$. The iterative approach consists of two steps:

1. *Cluster reassignment:* for each $y \in \mathcal{D}$, we find the index $i \in [K]$, such that

$$g(x_t(i), y) \leq g(x_t(j), y), \forall j \neq i, \quad (6)$$

and assign the point y to cluster $C_{t+1}(i)$.

2. *Center update:* for each $i \in [K]$, we perform the following update

$$x_{t+1}(i) = x_t(i) - \alpha \sum_{y \in C_{t+1}(i)} p_y \nabla_x f(x_t(i), y), \quad (7)$$

where $\alpha > 0$ is a fixed step-size.

Note that (7) can be written compactly as

$$x_{t+1} = x_t - \alpha \nabla_x J(x_t, C_{t+1}), \quad (8)$$

where $\nabla_x J(x_t, C_{t+1}) \in \mathbb{R}^{Kd}$ is the gradient of J with respect to x , whose i -th block of size d is given by

$$[\nabla_x J(x_t, C_{t+1})]_i = \sum_{y \in C_{t+1}(i)} p_y \nabla_x f(x_t(i), y). \quad (9)$$

In addition to Assumption 2.1, for our method to be applicable, we make the following assumptions on functions f , g and J .

Assumption 3.1. The distance function g is a metric, i.e., it satisfies the following properties: 1) $g(x, y) \geq 0$, and $g(x, y) = 0 \iff x = y$; 2) $g(x, y) = g(y, x)$; 3) $g(x, y) \leq g(x, z) + g(z, y)$.

Remark 3.2. Assumption 3.1 requires the distance function g , that dictates cluster assignment, to be a distance metric. Note that, with respect to (Banerjee et al., 2005), Bregman divergences are not necessarily symmetric, nor do they obey the triangle inequality. However, (Acharyya et al., 2013), (Chen et al., 2008) show that a large class of Bregman divergences, such as Mahalanobis distances, as well as Jensen-Shannon divergence, represent squares of metrics. Hence, for the choice $f(x, y)$ a Bregman divergence representing the square of a metric and $g(x, y) = \sqrt{f(x, y)}$, Assumption 3.1 is satisfied.

Assumption 3.3. The cost function f is coercive with respect to the first argument, i.e. $\lim_{\|x\| \rightarrow +\infty} f(x, y) = +\infty, \forall y \neq x$.

Remark 3.4. Assumption 3.3 ensures that the sequence of centers, $\{x_t\}$, generated by (8) remains bounded. It does so, by not allowing for x to grow infinitely without affecting the loss function f .

Assumption 3.5. The function J has co-coercive gradients in the first argument, i.e., for all $x, z \in \mathbb{R}^{Kd}$

$$\begin{aligned} & \langle \nabla_x J(x, C) - \nabla_z J(z, C), x - z \rangle \\ & \geq \frac{1}{L} \|\nabla_x J(x, C) - \nabla_z J(z, C)\|^2. \end{aligned}$$

Remark 3.6. Assumption 3.5 ensures that the sequence of centers, $\{x_t\}$, generated by (8) not only decreases the cost J , but also decreases the distance of the generated sequence $\{x_t\}$ to a stationary point x_* (or the set of stationary points in general), at every iteration.

Remark 3.7. Assumption 3.5 implies Lipschitz continuous gradients with respect to the first argument of the function J , as a result of the Cauchy-Schwartz inequality. As we show in Appendix B, Assumption 3.5 is satisfied for any function that is convex and has Lipschitz continuous gradients.

Remark 3.8. Note that Assumption 3.5 rules out non-smooth costs, such as K -medians, (2). However, when a desirable feature of the cost is robustness, smooth costs like the Huber loss can be used.

4. Convergence analysis

In this section the goal is to show that the method (6)-(7) converges to a fixed point.

To begin with, the notions of a fixed point and a set of optimal clusterings are defined.

Definition 4.1. The pair (x_*, C_*) is a fixed point of the clustering procedure (6)-(7), if the following holds:

1. Optimal clustering with respect to centers: for each $i \in [K]$, and each $y \in C_*(i)$, we have

$$g(x_*(i), y) \leq g(x_*(j), y), \forall j \neq i. \quad (10)$$

2. Optimal centers with respect to clustering:

$$\nabla_x J(x_*, C_*) = 0.$$

Definition 4.2. Let $x \in \mathbb{R}^{Kd}$ represent cluster centers. We say U_x is the set of optimal clusterings with respect to x , if for all clusterings $C \in U_x$, (6) is satisfied.

Definition 4.3. Let $x \in \mathbb{R}^{Kd}$ represent cluster centers. We define the set \bar{U}_x as the set of clusterings with respect to x such that: 1) $\bar{U}_x \subset U_x$; 2) $\forall C \in \bar{U}_x : \nabla_x J(x, C) = 0$.

Remark 4.4. As we show in Section 4, for a Bregman cost (of which the K -means problem is a special case) any fixed point, per Definition 4.1, represents a centroidal partition of the data, i.e., the centers $x_*(i)$ correspond to the means of clusters $C_*(i)$. This is consistent with results in (Banerjee et al., 2005), and shows that Definition 4.1 is a natural one.

Remark 4.5. In a slight abuse of terminology, we will refer to a point x as fixed point, if there exists a clustering C such that (x, C) satisfies Definition 4.1.

Remark 4.6. Note that, by Definition 4.3, a pair (x, C) is a fixed point if $C \in \bar{U}_x$.

The main result of the paper is stated in Theorem 4.7, which shows the convergence of the sequence of cluster centers to a fixed point.

Theorem 4.7. *Let Assumptions 2.1, 3.1, 3.3, 3.5 hold. For the step-size choice $\alpha < \frac{2}{L}$ and any $x_0 \in \mathbb{R}^{Kd}$, the sequence of centers $\{x_t\}$ generated by the algorithm (6)-(7), converges to a fixed point $x_* \in \mathbb{R}^{Kd}$, i.e., a point such that $\bar{U}_{x_*} \neq \emptyset$.*

The result of Theorem 4.7 is strong - for a fixed step-size, under arbitrary initialization, the proposed algorithm converges to a fixed point. In the context of K -means clustering, e.g. (Lloyd, 1982), (Banerjee et al., 2005), we achieve the same guarantees. In the context of different costs, e.g. Huber loss, compared to (Pediredla & Seelamantula, 2011), where the authors show convergence of the sequence of centers, under the assumptions that the clusters have already converged, and the initialization x_0 is sufficiently close to a fixed point, our results are much stronger - we guarantee that the full sequence $\{x_t\}$ converges to a fixed point, under arbitrary initialization. We also show that the clusters converge.

To prove Theorem 4.7, a series of intermediate lemmas is introduced. The proof outline follows a similar idea as the one developed in (Kar & Swenson, 2019).

The following lemma shows that the proposed algorithm decreases the objective function J in each iteration.

Lemma 4.8. *For the sequence $\{(x_t, C_t)\}$, generated by (6)-(7), with $\alpha < \frac{2}{L}$, the resulting sequence of costs $\{J(x_t, C_t)\}$ is non-increasing.*

Proof of Lemma 4.8. To begin with, note that (6) together with Assumption 2.1 implies that the clustering reassignment step decreases the cost, i.e.

$$\begin{aligned} J(x_t, C_{t+1}) &= \sum_{i=1}^K \sum_{y \in C_{t+1}(i)} p_y f(x_t(i), y) \\ &\leq \sum_{i=1}^K \sum_{y \in C_t(i)} p_y f(x_t(i), y) = J(x_t, C_t). \end{aligned} \quad (11)$$

Next, using Lipschitz continuity of gradients of J (recall Remark 3.7), we have

$$\begin{aligned} J(x_{t+1}, C_{t+1}) &\leq J(x_t, C_{t+1}) + \frac{L}{2} \|x_{t+1} - x_t\|^2 \\ &\quad + \left\langle \nabla_x J(x_t, C_{t+1}), x_{t+1} - x_t \right\rangle. \end{aligned}$$

Using (8), we get

$$J(x_{t+1}, C_{t+1}) \leq J(x_t, C_{t+1}) - c(\alpha) \|\nabla_x J(x_t, C_{t+1})\|^2,$$

where $c(\alpha) = \alpha \left(1 - \frac{\alpha L}{2}\right)$. Choosing $\alpha < \frac{2}{L}$ ensures that $c(\alpha) > 0$, and combining with (11), we get

$$\begin{aligned} J(x_{t+1}, C_{t+1}) &\leq J(x_t, C_{t+1}) - c(\alpha) \|\nabla_x J(x_t, C_{t+1})\|^2 \\ &\leq J(x_t, C_t) - c(\alpha) \|\nabla_x J(x_t, C_{t+1})\|^2 \\ &\leq J(x_t, C_t), \end{aligned} \quad (12)$$

which completes the proof. \square

The remaining proofs can be found in Appendix A. The following lemma shows that, if two cluster centers are sufficiently close, the optimal clustering sets match.

Lemma 4.9. *Let $x \in \mathbb{R}^{Kd}$ represent cluster centers. Then, $\exists \epsilon > 0$, such that, for any center $x' \in \mathbb{R}^{Kd}$, satisfying $\max_{i \in [K]} g(x(i), x'(i)) < \epsilon$, we have $U_{x'} \subset U_x$.*

The next lemma shows that, if a limit point of the sequence of centers exists, it must be a fixed point.

Lemma 4.10. *Any convergent subsequence of the sequence $\{x_t\}$, generated by (6)-(7), converges to a fixed point.*

The next lemma proves a stronger result, namely, that the clusters converge in finite time.

Lemma 4.11. *For any convergent subsequence of the sequence of centers, $\exists s_0 > 0$, such that $\forall s \geq s_0 : C_{t_s+1} \in \bar{U}_{x_*}$, where x_* is the limit of the sequence $\{x_{t_s}\}$.*

The following lemma shows that the generated sequence of cluster centers stays bounded.

Lemma 4.12. *The sequence of cluster centers $\{x_t\}$, generated by (6)-(7), is bounded.*

The next lemma shows that, if a point in the sequence of centers is sufficiently close to a fixed point, then all the subsequent points remain in the neighborhood of the fixed point.

Lemma 4.13. *Let $\{x_t\}$ be the sequence of cluster centers generated by (6)-(7), with the step-size satisfying $\alpha < \frac{2}{L}$. Let x_* be a fixed point, in the sense that $\bar{U}_{x_*} \neq \emptyset$. Then, $\exists \epsilon_{x_*} > 0$, such that, $\forall \epsilon \in (0, \epsilon_{x_*})$, $\exists t_\epsilon > 0$, such that, if $\|x_{t_0} - x_*\| \leq \epsilon$, for some $t_0 > t_\epsilon$, then $\|x_t - x_*\| \leq \epsilon$, for all $t \geq t_0$.*

We are now ready to prove Theorem 4.7.

Proof of Theorem 4.7. By Lemma 4.8 and the fact that the corresponding sequence of costs $\{J(x_t, C_t)\}$ is nonnegative, we know this sequence converges to some $J_* \in \mathbb{R}_+$, by the monotone convergence theorem. On the other hand, by Bolzano-Weierstrass theorem and Lemma 4.12, the sequence $\{x_t\}$ has a convergent subsequence, $\{x_{t_s}\}$, with some $x_* \in \mathbb{R}^{Kd}$ as its limit. From the continuity of J and convergence of $\{x_{t_s}\}$, we can then conclude that $J_* = \lim_{s \rightarrow +\infty} J(x_{t_s}, C_{t_s}) = J(x_*, C_*)$, for some $C_* \in U_{x_*}$. Lemma 4.10 then implies that x_* is a fixed point. Finally, Lemmas 4.11 and 4.13 imply the convergence of the entire sequence $\{x_t\}$ to x_* . \square

Remark 4.14. We note that the convergence guarantees of our method are independent of the initialization. Therefore, our method is amenable to seeding procedures, such as K -means++.

5. Fixed point analysis

In this section we analyse the fixed points and their properties. To begin with, we formally define the notion of Voronoi partitions, e.g., (Okabe et al., 2000).

Definition 5.1. Let (V, d) be a metric space. For a set $X \subset V$, and $z = (z(1), \dots, z(K)) \in V^K$, we say that $P = (P(1), \dots, P(K))$ is a Voronoi partition of the set X , generated by z , with respect to the metric d , if P is a partition of X and additionally, for every $i \in [K]$

$$P(i) = \{x \in X : d(z(i), x) \leq d(z(j), x), \forall j \neq i\}.$$

From Definitions 4.1 and 5.1, it is clear that, for a fixed point (x_*, C_*) , the clustering C_* represents a Voronoi partition of \mathcal{D} , with respect to g , generated by x_* . Moreover, from Definition 4.2, it is clear that, for any point x , the set U_x represents the set of all possible Voronoi partitions of \mathcal{D} , generated by x .

From the cluster reassignment step (6), we can see that in our approach, the clusters represent Voronoi partitions with respect to g . It is known that different distance metrics induce different Voronoi partitions, e.g., (Okabe et al., 2000), and the choice of metrics affects the shape of the resulting partitions. For example, choosing $g_1(x, y) = \|x - y\|$, the standard Euclidean distance and $g_2(x, y) = \|x - y\|_A$, a Mahalanobis distance (see (14) ahead), would potentially result in different Voronoi partitions of the dataset. In that sense, the distance function g determines the *cluster shape*.

Using (9), the fixed point condition from Definition 4.1 is equivalent to

$$\begin{aligned} \forall i \in [K] : \nabla_x J_i(x_*(i), C_*(i)) = 0 &\iff \\ \forall i \in [K] : \sum_{y \in C(i)} p_y \nabla_x f(x_*(i), y) = 0. \end{aligned} \quad (13)$$

From (13), we can see that the exact location of a cluster center is determined by f . In that sense, the cost function f determines the *location of cluster centers*. For example, for the choice $g(x, y) = \|x - y\|$, $f_1(x, y) = \frac{1}{2}\|x - y\|^2$ and $f_2(x, y) = \phi_\delta(g(x, y))$, where ϕ_δ is the Huber loss defined in (5), we can see that in both cases the cluster shapes will be determined by the Euclidean distance metric. However, applying (13) to f_1 and f_2 , it can be shown that

$$\begin{aligned} x_1(i) &= \frac{1}{\mu_{\mathcal{D}}(C_1(i))} \sum_{y \in C_1(i)} p_y y, \\ x_2(i) &= \frac{\sum_{y \in C_2(i)} p_y y + \sum_{y \in \overline{C_2}(i)} \frac{\delta}{\|x_2(i) - y\|} p_y y}{\sum_{y \in C_2(i)} p_y + \sum_{y \in \overline{C_2}(i)} \frac{\delta}{\|x_2(i) - y\|} p_y}, \end{aligned}$$

where $\overline{C_2}(i) = \{y \in C_2(i) : \|x_2(i) - y\| > \delta\}$, $C_2(i) = \{y \in C_2(i) : \|x_2(i) - y\| \leq \delta\}$, $x_1(i)$ and $x_2(i)$ satisfy (13) for f_1 and f_2 respectively, and $\mu_{\mathcal{D}}(C(i)) = \sum_{y \in C(i)} p_y$, represents the measure of the i -th cluster. Hence, we see that the function f dictates the exact location of the cluster center within the cluster.

Remark 5.2. Note that, while a fixed point of Huber loss takes the form of $x_2(i)$, as defined above, it is not actually a trivially computable closed form solution, as both sides of the equality contain $x_2(i)$. Therefore, to obtain such a form in practice, an iterative solver is required.

5.1. Case study: Centroidal Voronoi Partitions

A Voronoi partition C of the set \mathcal{D} generated by x is called centroidal, if the generator of each partition corresponds to its center, i.e.

$$x(i) = \frac{1}{\mu_{\mathcal{D}}(C(i))} \sum_{y \in C(i)} p_y y, \quad \forall i \in [K].$$

The authors in (Banerjee et al., 2005) show that, if the cost function f is a Bregman divergence, the Lloyd-type

algorithm (Lloyd, 1982) is optimal, i.e., using centroidal Voronoi partitions results in the minimal loss in Bregman information. In what follows, we show that, for a Bregman divergence-type cost function, our algorithm converges to the set of centroidal Voronoi partitions. To this end, we first define the notion of Bregman divergence.

Definition 5.3. Let $\phi : \mathbb{R}^d \mapsto \mathbb{R}$ be a strictly convex, differentiable function. The Bregman divergence defined by ϕ is given by $d_\phi(p, q) = \phi(p) - \phi(q) - \langle \nabla \phi(q), p - q \rangle$.

As a consequence of strict convexity of ϕ , we have $d_\phi \geq 0$, and $d_\phi(p, q) = 0 \iff p = q$. However, in general, d_ϕ is not a metric. Therefore, in our framework, Bregman divergences are used as $f(x, y) = d_\phi(y, x)$. To define an appropriate metric g , we rely on the works (Acharyya et al., 2013), (Chen et al., 2008), that show a rich class of Bregman divergences that represent squares of metrics. Examples include Mahalanobis distance based Bregman divergences, as well as the Jensen-Shannon entropy. We show in Appendix B that, on a properly defined support, the Jensen-Shannon entropy satisfies Assumptions 2.1-3.5. Here, we define the Mahalanobis distance based Bregman divergences and show how they fit our framework. Let $A \in \mathbb{R}^{d \times d}$ be a symmetric positive definite matrix. The corresponding Bregman divergence is then given by

$$d_\phi(x, y) = \frac{1}{2}(x - y)^T A(x - y). \quad (14)$$

This class of Bregman divergences is covered by our formulation, for the choice

$$\begin{aligned} f(x, y) &= \frac{1}{2}(x - y)^T A(x - y), \\ g(x, y) &= \|x - y\|_A, \end{aligned}$$

where $\|x\|_A := \sqrt{\langle Ax, x \rangle}$.

Lemma 5.4. *Let f be a Bregman divergence, satisfying Assumptions 2.1-3.5. Then, the gradient clustering algorithm converges to the set of centroidal Voronoi partitions.*

The proof of Lemma 5.4 can be found in Appendix A.

5.2. Case study: Beyond Centroidal Voronoi Partitions

Note that, in the case the cost used is a Bregman distance, the fixed point has a closed-form solution (39). Therefore, in each iteration of the algorithm, it is possible to compute the optimal cluster center, which is exactly what the Lloyd algorithm does. The Lloyd algorithm (Lloyd, 1982), and its generalization (Banerjee et al., 2005), perform the following two steps:

1. *Cluster reassignment:* for each $y \in \mathcal{D}$, find the cluster center $i \in [K]$, such that

$$g(x_t(i), y) \leq g(x_t(j), y), \quad \forall j \neq i,$$

and assign the point y to cluster $C_{t+1}(i)$.

2. *Center update*: for each $i \in [K]$, perform the following update

$$x_{t+1}(i) = \frac{1}{\mu_{\mathcal{D}}(C_{t+1}(i))} \sum_{y \in C_{t+1}(i)} p_y y. \quad (15)$$

The authors in (Bottou & Bengio, 1995) analyze the update rule (15) and show that it corresponds to performing a Newton step in each iteration. The authors in (Banerjee et al., 2005) show an even stronger result - in the case f is a Bregman divergence, the update (15) corresponds to the optimal update, in terms of minimizing the Bregman information.

From that perspective, naively extending the Lloyd’s algorithm to a general cost f would correspond to

$$x_{t+1}(i) = \arg \min_{x(i)} \sum_{y \in C_{t+1}(i)} \frac{p_y}{\mu_{\mathcal{D}}(C_{t+1}(i))} f(x(i), y). \quad (16)$$

Performing the update (16) would require solving an optimization problem in each iteration. This computation might be prohibitively expensive. In this case, the update (7) is preferred, as computing the gradient is a feasible, and in many cases cheap operation.

An example of such a function is the Huber loss, defined in (5). Huber loss provides robustness, e.g., (Ke & Kanade, 2005), (Liu et al., 2019), as it behaves like the squared loss for points whose modulus is smaller than a given threshold, while it grows only linearly for points whose modulus is beyond the threshold. Therefore, Huber loss implicitly gives more weight to points with smaller modulus.

In our framework, Huber loss is used as

$$f(x, y) = \phi_{\delta}(\|x - y\|) = \begin{cases} \frac{1}{2} \|x - y\|^2, & \|x - y\| \leq \delta \\ \delta \|x - y\| - \frac{\delta^2}{2}, & \|x - y\| > \delta \end{cases}. \quad (17)$$

A closed form expression satisfying (16), for the cost (17) does not exist. Therefore, to perform the update (16) in practice, requires solving an optimization problem in every iteration. On the other hand, from (5) and (17), we have

$$\nabla_x f(x, y) = \begin{cases} (x - y), & \|x - y\| \leq \delta \\ \delta \frac{x - y}{\|x - y\|}, & \|x - y\| > \delta \end{cases},$$

hence the gradient update is straightforward to compute. Note that computing the gradient update of the Huber loss corresponds to performing gradient clipping, effectively dampening the contribution of points that are far away from the current center estimate. We show in Appendix B that Huber loss satisfies Assumptions 2.1-3.5.

6. Numerical experiments

In this section we demonstrate the effectiveness of the proposed method. The experiments presented in this section were performed on the MNIST dataset (LeCun et al.). In Appendix C we present additional numerical experiments, performed on the Iris dataset (Fisher, 1936). Throughout the experiments, we assume a uniform distribution over the data, i.e., $\mu_{\mathcal{D}}(y_i) = \frac{1}{N}$, $\forall i = 1, \dots, N$, with $\mathcal{D} = \{y_1, \dots, y_N\}$.

The MNIST training dataset consists of handwritten digits, along with the corresponding labels. The data is initially normalized (divided by the highest value in the dataset), so that each pixel belongs to the $[0, 1]$ interval. Next, we select the first 500 samples of the digits 1 through 7. In total, our dataset consists of $N = 3500$ points, each being in $[0, 1]^{768}$ (as there are 28×28 pixels), with the number of underlying clusters $K = 7$.

For the first experiment, we utilised the gradient based clustering using the standard squared Euclidean cost. In our setup, that corresponds to: $f(x, y) = \frac{1}{2} \|x - y\|^2$, $g(x, y) = \|x - y\|$. We refer to the resulting method as gradient K -means and compare it with the standard K -means (Lloyd, 1982), (Banerjee et al., 2005). In line with our theory, we set the step-size equal to $\alpha = \frac{1}{L} = \frac{1}{N} = \frac{1}{3500}$. For a fair comparison, we set the initial centers of both methods to be the same. In particular, we take a random point from each class and set them as the initial centroids.

We run the clustering experiments for 20 times and present the mean performance (solid line), as well as the standard deviation (shaded region). The measure of performance used is the fraction of correctly clustered samples. Note that both methods are unsupervised, i.e., do not use labels when learning. However, we used the labels as ground truth, when comparing the clustering results. In order to account for a possible label mismatch, we checked all the possible label permutations when computing the clustering accuracy and chose the highest score as the true score. The results are presented in Figure 1.

Figure 1 shows that accuracy-wise, the gradient based K -means slightly outperforms the standard K -means. Speed-wise, the standard K -means update converges faster, which is to be expected, as the K -means update corresponds to performing the exact arg min step in each iteration.

For the second experiment, we added zero mean Gaussian noise to a fraction of points from all classes, thus introducing noise. In order to combat the noise, we use a Huber loss function for our gradient based clustering method. In our framework, the Huber loss is used as in (17). We compare the performance of the gradient based Huber loss clustering and the Huber based method from (Pediredla & Seelamantula, 2011). The authors in (Pediredla & Seelamantula, 2011) consider a method that is based on a fixed-point itera-




Figure 1. Accuracy of the Lloyd based K -means vs the gradient based K -means algorithm. Presents the accuracy of clustering digits 1 through 7 from the MNIST dataset.

tion, given by the recursion

$$x_{t+1}(i) = \frac{\sum_{y \in \underline{C}_t(i)} p_y y + \sum_{y \in \overline{C}_t(i)} \frac{\delta}{\|x_t(i) - y\|} p_y y}{\sum_{y \in \underline{C}_t(i)} p_y + \sum_{y \in \overline{C}_t(i)} \frac{\delta}{\|x_t(i) - y\|} p_y},$$

where $\overline{C}_t(i) = \{y \in C_t(i) : \|x_t(i) - y\| > \delta\}$, $\underline{C}_t(i) = \{y \in C_t(i) : \|x_t(i) - y\| \leq \delta\}$. The authors also suggest initializing the method by doing one round of Lloyd’s algorithm from a random starting point. For fairness of comparison, we initialize both the gradient Huber and the method from (Pediredla & Seelamantula, 2011) (which we refer to as “Huber” in the figures) in this way.

As in the previous experiment, we report the average results over 20 runs, along with the standard deviation. We consider the effects of changing the percentage of noisy samples and changing the variance of the noise. In all the experiments, we fix the Huber loss parameter to $\delta = 10$ and use the same step-size as in the standard K -means case, i.e., $\alpha = \frac{1}{3500}$. The results are presented in the Figure 2 below.

Figure 2 shows the performance of the Huber loss gradient method vs the method from (Pediredla & Seelamantula, 2011), when the percentage of noisy samples and variance of noise vary. Comparing the rows, i.e., different percentage of noisy samples, we can see that both methods perform better when the percentage of noise is lower, as expected. Comparing the columns, i.e., different variance levels, we can see that our method is comparable to (Pediredla & Seelamantula, 2011) for variance 1, but slightly outperforms the competing method for variance 2. Therefore our method exhibits a similar or better performance, with a small loss in speed. However, our method provides much better convergence guarantees, as it provably converges for arbitrary initialization, while the method (Pediredla & Seelamantula,




Figure 2. Performance of Huber loss gradient vs the method from (Pediredla & Seelamantula, 2011) on MNIST data. The rows correspond to percentage of noisy samples being 10% and 20%, with columns corresponding to variance of noise being 1 and 2, respectively (e.g., the upper left image corresponds to 10% of noisy samples, with variance 1, etc.).

2011) provides only local convergence guarantees, when already in a neighborhood of the stationary point.

7. Conclusion

We proposed an approach to clustering, based on the gradient of a generic loss function, that measures clustering quality with respect to cluster assignments and cluster center positions. The approach is based on a formulation of the clustering problem that unifies the previously proposed distance based clustering approaches. The main advantage of the algorithm, compared to the standard approaches is its applicability to a wide range of clustering problems, low computational cost, as well as the ease of implementation. We prove that the sequence of centers generated by the algorithm converges to an appropriately defined fixed point, under arbitrary center initialization. We further analyze the type of fixed points our algorithm converges to, and show consistency with prior works, in case the cost is a Bregman divergence. Most notably, the assumed generic formulation includes loss functions beyond Bregman divergences (such as the Huber loss), for which the K -means-type averaging cluster center update step is not appropriate, while the step that corresponds to exact minimization with respect to the loss is computationally expensive. To combat these challenges, the proposed method involves a single gradient step with respect to the loss to update cluster centers. Numerical experiments illustrate and corroborate the results.

Acknowledgements

We thank the reviewers for their useful comments and suggestions. The work of A. Armacki and S. Kar was partially supported by the National Science Foundation under grant CNS-1837607. The work of D. Bajovic and D. Jakovetic is supported by the European Union’s Horizon 2020 Research and Innovation program under grant agreements No 957337 and 871518. This paper reflects only the authors’ views and the European Commission cannot be held responsible for any use which may be made of the information contained therein.

References

- Acharyya, S., Banerjee, A., and Boley, D. Bregman divergences and triangle inequality. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, pp. 476–484. SIAM, 2013.
- Arora, S., Raghavan, P., and Rao, S. Approximation schemes for euclidean k -medians and related problems. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC ’98, pp. 106–113, Dallas, Texas, USA, 1998. Association for Computing Machinery. ISBN 0897919629. doi: 10.1145/276698.276718.
- Arthur, D. and Vassilvitskii, S. K-means++: The advantages of careful seeding. In *In Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1027–1035, New Orleans, Louisiana, 2007. SIAM. ISBN 9780898716245.
- Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., and Pandit, V. Local search heuristics for k -median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004. doi: 10.1137/S0097539702416402.
- Awasthi, P. and Balcan, M.-F. Center based clustering: A foundational perspective. 2014.
- Awasthi, P., Charikar, M., Krishnaswamy, R., and Sinop, A. K. The hardness of approximation of euclidean k -means. *arXiv preprint arXiv:1502.03316*, 2015. doi: 10.48550/ARXIV.1502.03316. URL <https://arxiv.org/abs/1502.03316>.
- Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. Clustering with bregman divergences. *Journal of Machine Learning Research*, 6(58):1705–1749, 2005. URL <http://jmlr.org/papers/v6/banerjee05b.html>.
- Bottou, L. and Bengio, Y. Convergence properties of the k -means algorithms. In Tesauro, G., Touretzky, D., and Leen, T. (eds.), *Advances in Neural Information Processing Systems*, volume 7. MIT Press, 1995. URL <https://proceedings.neurips.cc/paper/1994/file/a1140a3d0df1c81e24ae954d935e8926-Paper.pdf>.
- Bregman, L. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967. ISSN 0041-5553. doi: [https://doi.org/10.1016/0041-5553\(67\)90040-7](https://doi.org/10.1016/0041-5553(67)90040-7).
- Chandola, V., Banerjee, A., and Kumar, V. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3), jul 2009. ISSN 0360-0300. doi: 10.1145/1541880.1541882. URL <https://doi.org/10.1145/1541880.1541882>.
- Chen, P., Chen, Y., and Rao, M. Metrics defined by Bregman Divergences. *Communications in Mathematical Sciences*, 6(4):915 – 926, 2008. doi: [cms/1229619676](https://doi.org/10.1145/1541880.1541882).
- Cortes, J., Martinez, S., Karatas, T., and Bullo, F. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004. doi: 10.1109/TRA.2004.824698.
- Dhillon, I. S., Mallela, S., and Kumar, R. A divisive information-theoretic feature clustering algorithm for text classification. *Journal of Machine Learning Research (JMLR)*, 3:1265–1287, Mar 2003. ISSN 1532-4435.
- Fisher, R. A. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936. doi: <https://doi.org/10.1111/j.1469-1809.1936.tb02137.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-1809.1936.tb02137.x>.
- Huang, Z. Clustering large data sets with mixed numeric and categorical values. In *In The First Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 21–34, 1997.
- Huber, P. J. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73 – 101, 1964. doi: 10.1214/aoms/1177703732. URL <https://doi.org/10.1214/aoms/1177703732>.
- Jain, A. K. Data clustering: 50 years beyond k -means. *Pattern Recognition Letters*, 31(8):651–666, 2010. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2009.09.011>. URL <https://www.sciencedirect.com/science/article/pii/S0167865509002323>. Award winning papers from the 19th International Conference on Pattern Recognition (ICPR).

- Kar, S. and Swenson, B. Clustering with distributed data. *arXiv preprint arXiv:1901.00214*, 2019. doi: 10.48550/ARXIV.1901.00214. URL <https://arxiv.org/abs/1901.00214>.
- Ke, Q. and Kanade, T. Robust L1 norm factorization in the presence of outliers and missing data by alternative convex programming. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 1:739–746, 2005. doi: 10.1109/CVPR.2005.309.
- LeCun, Y., Cortes, C., and Burges, C. J. C. MNIST handwritten digit database. URL <http://yann.lecun.com/exdb/mnist/>.
- Liu, C., Sun, Q., and Tan, K. M. Robust convex clustering: How does fusion penalty enhance robustness? *arXiv preprint arXiv:1906.09581*, 2019. doi: 10.48550/ARXIV.1906.09581. URL <https://arxiv.org/abs/1906.09581>.
- Lloyd, S. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982. doi: 10.1109/TIT.1982.1056489.
- MacQueen, J. Some methods for classification and analysis of multivariate observations. In *In 5-th Berkeley Symposium on Mathematical Statistics and Probability*, number 14, pp. 281–297. University of California Press, 1967.
- Megiddo, N. and Supowit, K. J. On the complexity of some common geometric location problems. *SIAM Journal on Computing*, 13(1):182–196, 1984. doi: 10.1137/0213014. URL <https://doi.org/10.1137/0213014>.
- Monath, N., Kobren, A., Krishnamurthy, A., and McCallum, A. Gradient-based hierarchical clustering. In *Discrete Structures in Machine Learning Workshop, NIPS*, Long Beach, CA, USA, 2017.
- Nesterov, Y. *Lectures on Convex Optimization*. Springer Publishing Company, Incorporated, 2nd edition, 2018. ISBN 3319915770.
- Okabe, A., Boots, B., Sugihara, K., Chiu, S. N., and Kendall, D. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams, Second Edition*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Ltd., 2000. ISBN 9780470317013.
- Paul, D., Chakraborty, S., Das, S., and Xu, J. Uniform concentration bounds toward a unified framework for robust clustering. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 8307–8319. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/460b491b917d4185ed1f5be97229721a-Paper.pdf>.
- Pediredla, A. K. and Seelamantula, C. S. A Huber-loss-driven clustering technique and its application to robust cell detection in confocal microscopy images. In *2011 7th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pp. 501–506, 2011.
- Schwager, M. *A Gradient Optimization Approach to Adaptive multi-robot control*. PhD thesis, Massachusetts Institute of Technology, 2009.
- Selim, S. Z. and Ismail, M. A. K-means-type algorithms: A generalized convergence theorem and characterization of local optimality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(1):81–87, 1984. doi: 10.1109/TPAMI.1984.4767478.
- Telgarsky, M. and Vattani, A. Hartigan’s method: k-means clustering without voronoi. In Teh, Y. W. and Titterton, M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 820–827, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <https://proceedings.mlr.press/v9/telgarsky10a.html>.
- Vattani, A. The hardness of k-means clustering in the plane. 2009. URL https://cseweb.ucsd.edu/~avattani/papers/kmeans_hardness.pdf.

A. Missing proofs

In this section we provide the proofs omitted from the main body.

Proof of Lemma 4.9. For given cluster centers $x \in \mathbb{R}^{Kd}$ and each data point $y \in \mathcal{D}$, we denote by $\mathcal{K}_x^*(y)$ the set of cluster indices i whose centers $x(i)$ are closest to y :

$$\mathcal{K}_x^*(y) = \arg \min_{i \in [K]} g(x(i), y).$$

Define

$$\epsilon_0 := \min_{y \in \mathcal{D}} \min_{i \in [K] \setminus \mathcal{K}_x^*(y)} g(x(i), y) - g(x(c_y^*), y), \quad (18)$$

where c_y^* denotes an arbitrary cluster in $\mathcal{K}_x^*(y)$. By the construction of $\mathcal{K}_x^*(y)$ and finiteness of the set of data points \mathcal{D} , we have that $\epsilon_0 > 0$.

Let $\mathcal{X}_{x, \epsilon} := \{x' \in \mathbb{R}^{Kd} : g(x(i), x(i')) < \epsilon, \forall i \in [K]\}$, where $\epsilon > 0$. We show that, for each $x' \in \mathcal{X}_{x, \epsilon_0/2}$, for each $y \in \mathcal{D}$, there holds

$$\mathcal{K}_{x'}^*(y) \subseteq \mathcal{K}_x^*(y). \quad (19)$$

From (19), it is easy to see that any optimal cluster assignment with respect to x' , $C \in U_{x'}$, will also be optimal with respect to x , thus implying the claim of the lemma.

To prove (19), fix an arbitrary data point y and fix an arbitrary $i \in \mathcal{K}_{x'}^*(y)$. We want to show that $i \in \mathcal{K}_x^*(y)$ as well, i.e., that cluster center $x(i)$ belongs to the set of cluster centers x closest to y . By the triangle inequality for g , we have

$$g(x(i), y) \leq g(x(i), x'(i)) + g(x'(i), y) < \frac{\epsilon_0}{2} + g(x'(j), y) \leq \frac{\epsilon_0}{2} + g(x(j), x'(j)) + g(x(j), y) < \epsilon_0 + g(x(j), y), \quad (20)$$

where in the second line we use the fact that $x' \in \mathcal{X}_{x, \epsilon_0/2}$ (for index i) and the fact that $i \in \mathcal{K}_{x'}^*(y)$, in the third line we apply the triangle inequality for g , and in the fourth line we use again the fact that x' is in the $\epsilon_0/2$ neighborhood of x (for index j). For the sake of contradiction, suppose now that $i \notin \mathcal{K}_x^*(y)$ and take $j \in \mathcal{K}_x^*(y)$ (note that (20) holds for all $j \in [K]$). Then, from (18) we have $g(x(i), y) \geq g(x(j), y) + \epsilon_0$, which clearly contradicts (20). This proves (19) and subsequently proves the lemma. \square

Proof of Lemma 4.10. Let $\{x_{t_s}\}_{s=0}^\infty$ be a convergent subsequence of $\{x_t\}$. Let x_* be its limit point and assume the contrary, that x_* is not a fixed point. By Definition 4.1, this means

$$\|\nabla_x J(x_*, C)\| > 0, \forall C \in U_{x_*}.$$

As the number of possible clusterings is finite, we can define

$$\min_{C \in U_{x_*}} \|\nabla_x J(x_*, C)\| = \epsilon_1 > 0. \quad (21)$$

From the assumption $x_{t_s} \rightarrow x_*$, we have that, for a fixed $\delta_* > 0$, there exists a sufficiently large $s_0 > 0$, such that

$$\forall i \in [K], \forall s \geq s_0 : \|x_{t_s}(i) - x_*(i)\| < \delta_*.$$

It then follows from the continuity of g that there exists a sufficiently large $s_0 > 0$, such that $g(x_{t_s}(i), x_*(i)) < \epsilon_*$. Per Lemma 4.9, we then have $C_{x_{t_s+1}} \in U_{x_{t_s}} \subset U_{x_*}, \forall s \geq s_0$. From (21), we have

$$\|\nabla_x J(x_*, C_{t_s+1})\| \geq \epsilon_1, \forall s \geq s_0. \quad (22)$$

Next, using the results established in Lemma 4.8, we have

$$\begin{aligned} J(x_{t+1}, C_{t+1}) &\leq J(x_t, C_t) - c(\alpha) \|\nabla_x J(x_t, C_{t+1})\|^2 \\ &\leq J(x_{t-1}, C_{t-1}) - c(\alpha) \|\nabla_x J(x_{t-1}, C_t)\|^2 - c(\alpha) \|\nabla_x J(x_t, C_{t+1})\|^2 \\ &\leq \dots \leq J(x_0, C_1) - c(\alpha) \sum_{r=0}^t \|\nabla_x J(x_r, C_{r+1})\|^2. \end{aligned}$$

Rearranging, we get

$$c(\alpha) \sum_{r=0}^t \|\nabla_x J(x_r, C_{r+1})\|^2 \leq J(x_0, C_1) - J(x_{t+1}, C_{t+1}) \leq J(x_0, C_1). \quad (23)$$

Additionally, we have

$$\sum_{j=0}^{s(t)} \|\nabla_x J(x_{t_j}, C_{t_{j+1}})\|^2 \leq \sum_{j=0}^t \|\nabla_x J(x_j, C_{j+1})\|^2, \quad (24)$$

where $s(t) = \sup\{j : t_j \leq t\}$. Combining (23) and (24), we get

$$c(\alpha) \sum_{j=0}^{s(t)} \|\nabla_x J(x_{t_j}, C_{t_{j+1}})\|^2 \leq J(x_0, C_1). \quad (25)$$

Noting that the term on the right hand side of (25) is finite and independent of t , and $s(t) \rightarrow +\infty$ as $t \rightarrow +\infty$, we can take the limit as $t \rightarrow +\infty$, to obtain

$$c(\alpha) \sum_{j=0}^{\infty} \|\nabla_x J(x_{t_j}, C_{t_{j+1}})\|^2 \leq J(x_0, C_1) < +\infty,$$

which implies

$$\lim_{s \rightarrow \infty} \|\nabla_x J(x_{t_s}, C_{t_{s+1}})\|^2 = 0.$$

Fix an $\epsilon > 0$. By the definition of limits, there exists a $s_1 > 0$, such that

$$\|\nabla_x J(x_{t_s}, C_{t_{s+1}})\| < \epsilon, \forall s \geq s_1.$$

On the other hand, from $x_{t_s} \rightarrow x_*$, there exists a $s_2 > 0$, such that

$$\|x_{t_s} - x_*\| < \epsilon, \forall s \geq s_2.$$

As $C_{x_{t_{s+1}}} \in U_{x_{t_s}} \subset U_{x_*}, \forall s \geq s_0$, for any $s \geq \max\{s_0, s_1, s_2\}$, we have

$$\|\nabla_x J(x_*, C_{t_{s+1}})\| \leq \|\nabla_x J(x_*, C_{t_{s+1}}) - \nabla_x J(x_{t_s}, C_{t_{s+1}})\| + \|\nabla_x J(x_{t_s}, C_{t_{s+1}})\| \leq L\|x_* - x_{t_s}\| + \epsilon < (L+1)\epsilon,$$

where we used the Lipschitz continuity of the gradients of J in the second inequality. As $\epsilon > 0$ was arbitrarily chosen, we can conclude

$$\|\nabla_x J(x_*, C_{t_{s+1}})\| \rightarrow 0, \quad (26)$$

which clearly contradicts (22). Hence, we can conclude that x_* is a fixed point, i.e.,

$$\exists C \in U_{x_*} : \|\nabla_x J(x_*, C)\| = 0.$$

□

Proof of Lemma 4.11. Let

$$\delta := \min_{C \in U_{x^*} \setminus \bar{U}_{x^*}} \|\nabla_x J(x^*, C)\|.$$

Note that, by construction of \bar{U}_{x^*} , it must be that $\|\nabla_x J(x^*, C)\| > 0$ for each $C \in U_{x^*} \setminus \bar{U}_{x^*}$, which together with the finiteness of $U_{x^*} \setminus \bar{U}_{x^*}$, implies $\delta > 0$.

For the sake of contradiction, suppose now that $C_{t_{s+1}} \in U_{x^*} \setminus \bar{U}_{x^*}$, infinitely often. Then, $\|\nabla_x J(x^*, C_{t_{s+1}})\| \geq \delta$ infinitely often, which clearly contradicts (26). □

Proof of Lemma 4.12. By Lemma 4.8, we have

$$\dots \leq J(x_t, C_{t+1}) \leq \dots \leq J(x_1, C_1) \leq J(x_0, C_1) < +\infty. \quad (27)$$

Recalling equation (4), for $x \in \mathbb{R}^{Kd}$ and a clustering C , we define

$$J_i(x(i), C(i)) = \sum_{y \in C(i)} p_y f(x(i), y),$$

so that

$$J(x, C) = \sum_{i=1}^K J_i(x(i), C(i)). \quad (28)$$

For the sake of contradiction, suppose that the sequence of centers $\{x_t\}$ is unbounded. This implies the existence of a cluster k and a subsequence t_s such that $\|x_{t_s}(i)\| \rightarrow +\infty$. For each t_s , let $\underline{t}_s = \max\{t \leq t_s : C_t(i) \neq \emptyset\}$, i.e., \underline{t}_s is the largest element in the sequence prior to t_s , such that the i -th cluster is non-empty.

Recalling the update rule (7), it is not hard to see that $x_{t_s}(i) = x_{\underline{t}_s}(i)$, for all s , implying $\|x_{\underline{t}_s}(i)\| \rightarrow +\infty$. By Assumption 3.3 and the fact that $C_{\underline{t}_s}(i)$ is nonempty for each s , we have

$$\lim_{\|x_{\underline{t}_s}(i)\| \rightarrow +\infty} J_k(x_{\underline{t}_s}(i), C_{\underline{t}_s}(i)) = +\infty. \quad (29)$$

Note that this is the case regardless of the clustering $C_{\underline{t}_s}$, as the dataset \mathcal{D} is finite, and therefore a bounded set. It is easy to see that unboundness of J_i implies unboundness of J , i.e., $\lim_{s \rightarrow +\infty} J(x_{\underline{t}_s}, C_{\underline{t}_s}) = +\infty$. But this contradicts (27), hence proving the claim of the lemma. \square

Proof of Lemma 4.13. Recall that, by Lemma 4.8, the sequence of costs $\{J(x_t, C_t)\}_{t \geq 0}$ is decreasing. Moreover, since $J(x, C) \geq 0$, we know that the limit of the sequence of costs exists and is finite. Let

$$J_* = \lim_{t \rightarrow \infty} J(x_t, C_t). \quad (30)$$

By assumption, $\bar{U}_{x_*} \neq \emptyset$. From the definition of \bar{U}_{x_*} , for all $C \in U_{x_*} \setminus \bar{U}_{x_*}$ we have

$$\|\nabla_x J(x_*, C)\| > 0. \quad (31)$$

As U_{x_*} is a finite set, we can define

$$\epsilon_1 = \min_{C \in U_{x_*} \setminus \bar{U}_{x_*}} \|\nabla_x J(x_*, C)\| > 0.$$

Let $\epsilon_* > 0$ be such that Lemma 4.9 holds. From the continuity of g , we have

$$\exists \delta_* > 0 \forall x : \|x - x_*\| < \delta_* \implies g(x, x_*) < \epsilon_*. \quad (32)$$

Define

$$\epsilon_{x_*} = \min \left\{ \delta_*, \frac{\epsilon_1}{L} \right\}. \quad (33)$$

For an arbitrary $\epsilon \in (0, \epsilon_{x_*})$, let $t_0 > 0$ be such that

$$J(x_t, C_t) \leq J_* + \frac{c(\alpha)}{2} (\epsilon_1 - L\epsilon)^2, \forall t \geq t_0, \quad (34)$$

with $c(\alpha)$ defined as in Lemma 4.8. Note that the choice of t_0 is possible, from (30) and the fact that $(\epsilon_1 - L\epsilon)^2 > 0$. Our goal now is to show that, for a fixed $\epsilon \in (0, \epsilon_{x_*})$, if for some $t : t \geq t_0$ and $\|x_t - x_*\| < \epsilon$, then $\|x_{t+1} - x_*\| < \epsilon$.

First note that, if $t \geq t_0$ and $\|x_t - x_*\| < \epsilon$, it holds that $C_{t+1} \in \bar{U}_{x_*}$. To see this, assume the contrary, $\|x_t - x_*\| < \epsilon$ and $C_{t+1} \notin \bar{U}_{x_*}$. It follows from (33) that

$$\|x_t - x_*\| < \delta_*.$$

From (32) and Lemma 4.9, we then have $U_{x_t} \subset U_{x_*}$, and hence, $C_{t+1} \in U_{x_*}$. Using Lipschitz continuity of gradients of J , we get

$$\|\nabla_x J(x_t, C_{t+1}) - \nabla_x J(x_*, C_{t+1})\| \leq L\|x_t - x_*\| \leq L\epsilon. \quad (35)$$

As $C_{t+1} \notin \bar{U}_{x_*}$, from (31), we have

$$\|\nabla_x J(x_*, C_{t+1})\| \geq \epsilon_1. \quad (36)$$

Applying the triangle inequality, (35) and (36), we get

$$\|\nabla_x J(x_t, C_{t+1})\| \geq \epsilon_1 - L\epsilon. \quad (37)$$

Note that, by (33), the right-hand side of (37) is positive. Combining (12), (34) and (37), we have

$$\begin{aligned} J(x_{t+1}, C_{t+1}) &\leq J(x_t, C_t) - c(\alpha)\|\nabla_x J(x_t, C_{t+1})\|^2 \leq J_* + \frac{c(\alpha)}{2}(\epsilon_1 - L\epsilon)^2 - c(\alpha)\|\nabla_x J(x_t, C_{t+1})\|^2 \\ &\leq J_* + \frac{c(\alpha)}{2}(\epsilon_1 - L\epsilon)^2 - c(\alpha)(\epsilon_1 - L\epsilon)^2 < J_*, \end{aligned}$$

which is a contradiction. Hence, $C_{t+1} \in \bar{U}_{x_*}$.

Using Assumption 3.5, the update rule (8), and the fact that $C_{t+1} \in \bar{U}_{x_*}$, we have

$$\begin{aligned} \|x_{t+1} - x_*\|^2 &= \|x_t - \alpha\nabla_x J(x_t, C_{t+1}) - x_*\|^2 = \|x_t - x_*\|^2 + \alpha^2\|\nabla_x J(x_t, C_{t+1})\|^2 - 2\alpha\langle \nabla_x J(x_t, C_{t+1}), x_t - x_* \rangle \\ &\leq \|x_t - x_*\|^2 - \alpha\left(\frac{2}{L} - \alpha\right)\|\nabla_x J(x_t, C_{t+1})\|^2 \leq \|x_t - x_*\|^2 < \epsilon^2, \end{aligned} \quad (38)$$

where the second inequality follows from the step-size choice $\alpha < \frac{2}{L}$. Therefore, we have shown that

$$\|x_t - x_*\| < \epsilon \implies \|x_{t+1} - x_*\| < \epsilon.$$

The same result holds for all $s > t$ inductively, which proves the claim. \square

Proof of Lemma 5.4. To this end, we want to show that, for an arbitrary fixed point (x_*, C_*) of the algorithm, the pair produces a centroidal Voronoi partition.

From Definition 4.1, it is clear that C_* is a Voronoi partition of the dataset, generated by x_* . Now, let $f(x, y)$ be a Bregman divergence, for some strictly convex ϕ . From the definition of Bregman divergence, we then have

$$\nabla_x f(x, y) = -\nabla\phi(x) + \nabla\phi(x) - \nabla^2\phi(x)(y - x) = \nabla^2\phi(x)(x - y).$$

Combining with (13), we get, for all $i \in [K]$

$$0 = \sum_{y \in C_*(i)} p_y \nabla_x f(x_*(i), y) = \nabla^2\phi(x_*(i)) \left(\sum_{y \in C_*(i)} p_y (x_*(i) - y) \right).$$

From the strict convexity of ϕ , we have

$$\sum_{y \in C_*(i)} p_y \nabla_x f(x_*(i), y) = 0 \iff \sum_{y \in C_*(i)} p_y (x_*(i) - y) = 0 \iff x_*(i) = \frac{1}{\mu_{\mathcal{D}}(C_*(i))} \sum_{y \in C_*(i)} p_y y. \quad (39)$$

We have shown that the generators of Voronoi partitions correspond to their respective centers, which completes the proof. \square

B. Some technical results

In this section we show some technical results used in the paper. The next lemma is taken from (Nesterov, 2018). For the sake of completeness, we provide the proof here.

Lemma B.1. *Let $f : \mathbb{R}^d \mapsto \mathbb{R}$ be convex and have Lipschitz continuous gradients. Then, f has co-coercive gradients.*

Proof. Define the function:

$$\phi_x(z) = f(z) - \langle \nabla f(x), z \rangle.$$

It is straightforward to see that ϕ_x maintains convexity, for any $x \in \mathbb{R}^d$. It then follows that the point x is a minimizer of ϕ_x . Next, we use the following lower-bound for functions with Lipschitz continuous gradients (the proof can be found in (Nesterov, 2018)):

$$\frac{1}{2L} \|\nabla f(x)\|^2 \leq f(x) - f(x^*), \quad (40)$$

where x^* is a minimizer of f . Substituting ϕ_x in equation (40), we get

$$\phi_x(y) - \phi_x(x) = f(y) - \langle \nabla f(x), y \rangle - f(x) + \langle \nabla f(x), x \rangle \geq \frac{1}{2L} \|\nabla \phi_x(y)\|^2 = \frac{1}{2L} \|\nabla f(y) - \nabla f(x)\|^2.$$

Applying the same steps to ϕ_y , and summing the resulting inequalities, gives the desired result. \square

The following lemma shows that Huber loss satisfies Assumptions 2.1-3.5.

Lemma B.2. *Huber loss-based cost satisfies Assumptions 2.1-3.5.*

Proof. Note that Huber loss is an increasing function on the domain of interest, $[0, +\infty)$. By definition,

$$\begin{aligned} g(x, y) &= \|x - y\|, \\ f(x, y) &= \phi_\delta(g(x, y)), \end{aligned}$$

hence Assumptions 2.1 and 3.1 are satisfied. By the same argument, for a fixed y , we have

$$\lim_{\|x\| \rightarrow +\infty} f(x, y) = +\infty,$$

satisfying Assumption 3.3.

Next, note that f is a convex function, as a composition of convex functions. By Lemma B.1, it suffices to show that $f(x, y)$ has Lipschitz continuous gradients. The gradient of f is given by

$$\nabla_x f(x, y) = \begin{cases} (x - y), & \|x - y\| \leq \delta \\ \delta \frac{x - y}{\|x - y\|}, & \|x - y\| > \delta \end{cases}.$$

We differentiate between the following cases:

1. $\|x - y\|, \|z - y\| \leq \delta$. We then have

$$\|\nabla f(x, y) - \nabla f(z, y)\| = \|(x - y) - (z - y)\| = \|x - z\|.$$

2. $\|x - y\| \leq \delta, \|z - y\| > \delta$ (the case when $\|x - y\| > \delta, \|z - y\| \leq \delta$ is analogous). We then have

$$\begin{aligned} \|\nabla f(x, y) - \nabla f(z, y)\| &= \left\| (x - y) - \frac{\delta}{\|z - y\|} (z - y) \right\| = \left\| (x - z) + \left(1 - \frac{\delta}{\|z - y\|}\right) (z - y) \right\| \\ &\leq \|x - z\| + \left(1 - \frac{\delta}{\|z - y\|}\right) \|z - y\| = \|x - z\| + \|z - y\| - \delta. \end{aligned}$$

Next, using the triangle inequality and $\|x - y\| \leq \delta$, we get

$$\|z - y\| \leq \|x - z\| + \|x - y\| \leq \|x - z\| + \delta.$$

Rearranging and substituting in the equation above, we get

$$\|\nabla f(x, y) - \nabla f(z, y)\| \leq 2\|x - z\|.$$

3. $\|x - y\|, \|z - y\| > \delta$. Without loss of generality, assume $\|x - y\| \leq \|z - y\|$. We then have

$$\begin{aligned} \|\nabla f(x, y) - \nabla f(z, y)\| &= \delta \left\| \frac{x - y}{\|x - y\|} - \frac{z - y}{\|z - y\|} \right\| = \delta \left\| \frac{x - y}{\|x - y\|} - \frac{z - y}{\|z - y\|} \pm \frac{x - y}{\|z - y\|} \right\| \\ &\leq \delta \left(\frac{1}{\|x - y\|} - \frac{1}{\|z - y\|} \right) \|x - y\| + \delta \frac{\|x - z\|}{\|z - y\|} \leq \delta \frac{\|z - x\|}{\|z - y\|} + \delta \frac{\|z - x\|}{\|z - y\|} \\ &\leq 2\|z - x\|, \end{aligned}$$

where we use the triangle inequality and $\|x - y\| \leq \|z - y\|$ in the first inequality, while the last inequality stems from $\|z - y\| > \delta$.

Hence, we have shown that, $\forall x, y, z \in \mathbb{R}^d$

$$\|\nabla f(x, y) - \nabla f(z, y)\| \leq 2\|x - z\|.$$

By Lemma B.1, we see that Assumption 3.5 is satisfied, thus proving the claim. \square

The following lemma shows that Jensen-Shannon divergence satisfies Assumptions 2.1-3.5, on a properly defined support.

Lemma B.3. *Let $P_\epsilon \subset \mathbb{R}^d$, for some $\epsilon > 0$, define the restricted probability simplex, i.e.*

$$P_\epsilon = \left\{ p \in \mathbb{R}^d : \sum_{i=1}^d p_i = 1, \epsilon \leq p_i < 1 \right\}. \quad (41)$$

Then, the Jensen-Shannon divergence based cost satisfies Assumptions 2.1-3.5 on P_ϵ .

Proof. By the definition of Jensen-Shannon divergence, we have

$$D_{JS}(y \parallel x) = \frac{1}{2} D_{KL}(y \parallel m) + \frac{1}{2} D_{KL}(x \parallel m),$$

where $m = \frac{x+y}{2}$, and $D_{KL}(\cdot \parallel \cdot)$ is the Kullback-Leibler divergence, defined by

$$D_{KL}(x \parallel y) = \sum_{i=1}^d x_i \log \frac{x_i}{y_i}.$$

It is shown in (Acharyya et al., 2013) that the Jensen-Shannon divergence represents the square of a metric. Therefore, for

$$\begin{aligned} g(x, y) &= \sqrt{D_{JS}(y \parallel x)}, \\ f(x, y) &= D_{JS}(y \parallel x), \end{aligned}$$

Assumptions 2.1 and 3.1 are satisfied. Since the domain of interest, given by (41) is bounded, Assumption 3.3 is not of interest.

We next show that D_{JS} is convex and has Lipschitz continuous gradients on P_ϵ . A basic computation yields that the partial derivative of D_{JS} , with respect to x_i , is given by

$$\frac{\partial}{\partial x_i} D_{JS}(y \parallel x) = \frac{1}{2} \log \frac{2x_i}{x_i + y_i}. \quad (42)$$

It is then straightforward to see that the Hessian of D_{JS} is a diagonal matrix, whose i -th diagonal element is given by

$$\frac{\partial^2}{\partial x_i^2} D_{JS}(y \parallel x) = \frac{y_i}{2(x_i + y_i)^2}. \quad (43)$$

Since $x, y \in P_\epsilon$, the expression in (43) is positive, hence D_{JS} is convex on P_ϵ . Next, from (42), for any $x, y, z \in P_\epsilon$, we have

$$\left| \frac{\partial}{\partial x_i} D_{JS}(y \parallel x) - \frac{\partial}{\partial z_i} D_{JS}(y \parallel z) \right| = \frac{1}{2} \left| \log \frac{x_i}{z_i} + \log \frac{z_i + y_i}{x_i + y_i} \right| \leq \max \left\{ \left| \log \frac{x_i}{z_i} \right|, \left| \log \frac{z_i + y_i}{x_i + y_i} \right| \right\}.$$

Without loss of generality, assume $x_i \geq z_i$. We then have

$$\left| \log \frac{x_i}{z_i} \right| = \log \frac{x_i}{z_i} \leq \frac{x_i}{z_i} - 1 = \frac{x_i - z_i}{z_i} \leq \frac{1}{\epsilon} (x_i - z_i),$$

and

$$\begin{aligned} \left| \log \frac{z_i + y_i}{x_i + y_i} \right| &= \log \frac{x_i + y_i}{z_i + y_i} \leq \frac{x_i + y_i}{z_i + y_i} - 1 = \frac{x_i - z_i}{z_i + y_i} \\ &\leq \frac{1}{\epsilon} (x_i - z_i), \end{aligned}$$

where we used $\log x \leq x - 1$ in the above inequalities. Hence, we have shown that

$$\left| \frac{\partial}{\partial x_i} D_{JS}(y \| x) - \frac{\partial}{\partial z_i} D_{JS}(y \| z) \right| \leq \frac{1}{\epsilon} |x_i - z_i|.$$

By definitions of the gradient and norm, it then follows that

$$\left\| \nabla_x D_{JS}(y \| x) - \nabla_z D_{JS}(y \| z) \right\| \leq \frac{1}{\epsilon} \|x - z\|,$$

which shows Lipschitz continuity of the gradients of D_{JS} on P_ϵ . Hence, by Lemma B.1, D_{JS} satisfies Assumption 3.5 on P_ϵ . \square

Remark B.4. Note that in general, Jensen-Shannon divergence does not satisfy Assumptions 3.3 and 3.5. However, in certain problems, where the restricted probability simplex of the form (41) is a natural domain of choice, the Jensen-Shannon divergence can be applied in our framework. One such example is soft clustering under uncertainty - where no class can be ruled out with certainty, nor can a point belonging to any class be taken with certainty. Hence, for an appropriately selected ϵ , the restricted probability simplex (41) represents a natural domain.

C. Additional experiments

In this section we present some additional numerical experiments. The experiments were performed on the Iris dataset (Fisher, 1936). We begin by describing the Iris dataset.

The Iris dataset consists of three species of the *Iris* flower, *Iris setosa*, *Iris virginica* and *Iris versicolor*, along with the corresponding labels. Each of the species has 50 samples, so that the total number of samples is 150. Each sample consists of 4 features, being the length and the width of the sepals and petals of the flowers. In total, the dataset consists of $N = 150$ points, with the number of underlying clusters $K = 3$.

For the first experiment, we utilised the gradient based clustering using the standard squared Euclidean cost. We compare the performance of the gradient K -means with the standard K -means. The step-size is set to $\alpha = \frac{2}{L} = \frac{2}{2N} = \frac{1}{150}$.

We run the clustering experiments for 20 times and present the mean performance (solid line), as well as the standard deviation (shaded region). The measure of performance used is the fraction of correctly clustered samples. Note that both methods are unsupervised, i.e., do not use labels when learning. However, we used the labels as ground truth, when comparing the clustering results. In order to account for a possible label mismatch, we checked all the possible label permutations when computing the clustering accuracy and chose the highest score as the true score. The results are presented in Figure 3. Figure 3 shows that accuracy-wise, the gradient based K -means performs identically to the standard K -means, at a negligible speed loss.

For the second experiment, we added zero mean Gaussian noise to a fraction of points from all classes, thus introducing noise. We again compare the performance of the gradient based Huber loss clustering and the Huber based method from (Pediredla & Seelamantula, 2011). Again, we follow the suggestions in (Pediredla & Seelamantula, 2011) and initialize both the methods by doing one round of Lloyd's algorithm from a random starting point.

As in the previous experiment, we report the average results over 20 runs, along with the standard deviation. We consider the effects of changing the percentage of noisy samples and changing the variance of the noise. In all the experiments, we fix the Huber loss parameter to $\delta = 5$ and use the same step-size as in the standard K -means version, i.e., $\alpha = \frac{1}{150}$. The results are presented in the Figure 4 below.




Figure 3. Accuracy of the Lloyd based K -means vs the gradient based K -means algorithm. Presents the accuracy of clustering flowers from the Iris dataset.

Figure 4 shows the performance of the Huber loss gradient method vs the method from (Pediredla & Seelamantula, 2011), when the percentage of noisy samples and variance of noise vary. The step-size was the same as in the standard gradient K -means case. Comparing the rows, i.e., different percentage of noisy samples, we can see that both methods perform identically both accuracy and speed-wise, when the percentage of noisy samples is lower. However, the gradient based Huber method outperforms (Pediredla & Seelamantula, 2011) when the percentage of noisy samples is higher, more significantly when the variance is higher as well (bottom right image). Comparing the columns, i.e., different variance levels, we can see that both methods perform better when the variance of noise is lower.

Gradient Based Clustering

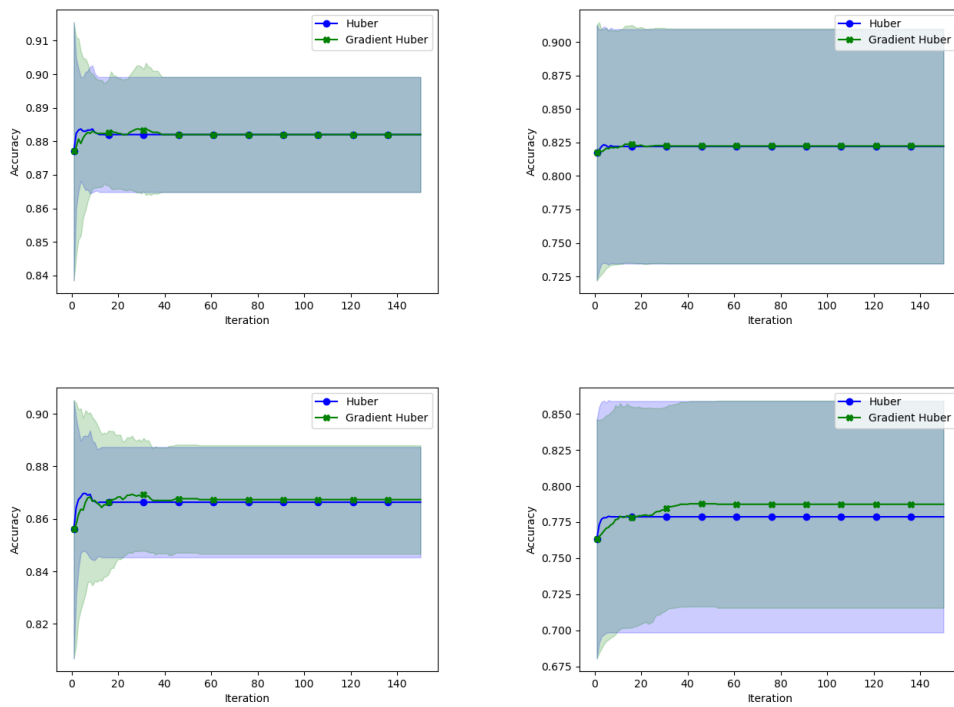


Figure 4. Performance of Huber loss gradient vs the method from (Pediredla & Seelamantula, 2011) on Iris data. The rows correspond to percentage of noisy samples being 10% and 20%, with columns corresponding to variance of noise being 1 and 2, respectively (e.g., the upper left image corresponds to 10% of noisy samples, with variance 1, etc.).