



HAL
open science

On the sequential convergence of Lloyd's algorithms

Léo Portales, Elsa Cazelles, Edouard Pauwels

► **To cite this version:**

Léo Portales, Elsa Cazelles, Edouard Pauwels. On the sequential convergence of Lloyd's algorithms. 2024. hal-04593982v2

HAL Id: hal-04593982

<https://hal.science/hal-04593982v2>

Preprint submitted on 2 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

On the sequential convergence of Lloyd’s algorithms

Léo Portales
IRIT, TSE, INP Toulouse
leo.portales@irit.fr

Elsa Cazelles
CNRS, IRIT, Université de Toulouse
elsa.cazelles@irit.fr

Edouard Pauwels
TSE, Université Toulouse Capitole
edouard2.pauwels@ut-capitole.fr

July 2, 2024

Abstract

Lloyd’s algorithm is an iterative method that solves the quantization problem, i.e. the approximation of a target probability measure by a discrete one, and is particularly used in digital applications. This algorithm can be interpreted as a gradient method on a certain quantization functional which is given by optimal transport. We study the sequential convergence (to a single accumulation point) for two variants of Lloyd’s method: (i) *optimal quantization* with an arbitrary discrete measure and (ii) *uniform quantization* with a uniform discrete measure. For both cases, we prove sequential convergence of the iterates under an analyticity assumption on the density of the target measure. This includes for example analytic densities truncated to a compact semi-algebraic set. The argument leverages the log analytic nature of globally subanalytic integrals, the interpretation of Lloyd’s method as a gradient method and the convergence analysis of gradient algorithms under Kurdyka-Lojasiewicz assumptions. As a by-product, we also obtain definability results for more general semi-discrete optimal transport losses such as transport distances with general costs, the max-sliced Wasserstein distance and the entropy regularized optimal transport loss.

1 Introduction

Quantization aims to approximate a target probability measure μ on \mathbb{R}^d by a discrete, finitely supported measure. We consider two quantization settings induced by an optimization problem in Wasserstein distance. First, optimal quantization consists in approximating the target measure μ with a discrete probability measure supported on N points $y_1, \dots, y_N \in \mathbb{R}^d$ with weights $\pi \in \Delta_N$, the $N - 1$ dimensional unit simplex, by solving the following optimization problem:

$$\min_{y_1, \dots, y_N \in \mathbb{R}^d} \min_{\pi \in \Delta_N} W_2^2 \left(\mu, \sum_{i=1}^N \pi_i \delta_{y_i} \right). \quad (1.1)$$

It is a central problem in signal processing to convert a continuous-time signal into a digital one [24, 34]. Second, uniform quantization consists in approximating μ with a discrete measure, corresponding to constraints on the optimal quantization weights:

$$\min_{y_1, \dots, y_N \in \mathbb{R}^d} W_2^2 \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right). \quad (1.2)$$




Figure 1: (Left) Target gaussian mixture μ with two components truncated on a disk. (Middle) Optimal quantization of μ with 20 points (blue) and their corresponding Voronoi cells (in red) after 250 iterations of Lloyd’s algorithm. (Right) Uniform quantization of μ with 20 points (blue) and the corresponding power cells (in red) after 5 iterations of Lloyd’s algorithm adjusted for uniform quantization. The diameter of a blue dot is proportional to its weight. The algorithms are randomly initialized and run using the PyMongeAmpere library¹.

In both problem, W_2^2 denotes the squared 2-Wasserstein distance, the optimal transport distance for the squared Euclidean cost between probability measures (see e.g. [42, 41, 38]). The Wasserstein distance is well suited in this context since it allows to handle the semi-discrete nature of the quantization task. The main question addressed in this work is that of the asymptotic convergence of the well known Lloyd iterative algorithmic method for problems (1.1) and (1.2). Note that both problems are non-convex and iterative solvers typically aim to find a stationary point rather than solve the global problem.

Lloyd-type algorithms. In his early paper [30], Lloyd proposes an algorithm to solve the least squares quantization problem in the univariate setting (without mentioning optimal transport at that time). Lloyd defines a quantization scheme using intervals and their centroids and introduces an iterative algorithm consisting in alternating optimization steps on the intervals and on the centroids. The generalization of this method to the multivariate setting involves *centroidal Voronoi cells* [19] which can be interpreted as optimal transport maps in the optimal quantization formulation in (1.1). In this setting, Lloyd’s method alternates between evaluation of Voronoi cells (given centroids), and update of centroids (given Voronoi cells).

Lloyd’s method can be extended to the case of uniform quantization in (1.2), for which the Voronoi cells encompass a volume constraint [4]. We refer to them as *power cells*. These regions are also associated to optimal transport maps in the uniform quantization problem in (1.2) [14], and algorithms have been developed for computing these power cells [26]. Both methods, for optimal or uniform quantization, share the same overall “alternating” structure and we will refer to both of them as Lloyd’s method. They are illustrated in Figure 1, generated using the PyMongeAmpere library¹. A precise description of Lloyd’s method for optimal and uniform quantization is given in Section 2 with Algorithm 1 and 2.

We address the question of sequential convergence of the iterates produced by Lloyd’s method for solving (1.1) and (1.2). More precisely, we consider the asymptotic stabilization of the support points y_1, \dots, y_N of the Dirac masses to a critical point of the quantization functional. This question was already considered in previous literature. For optimal quantization, the first analyses were proposed for the univariate case [43, 25] with special cases such as the strict log-concavity of the target measure density [21, 19]. In the multivariate setting, it is known that accumulation points of Lloyd’s iteration are critical points of the loss

¹<https://github.com/mrgt/PyMongeAmpere>

in (1.1) [18, 20, 35]. In particular if such critical points are isolated, an assumption very difficult to check in practice, then the algorithm converges. For uniform quantization, similar results exist. It was shown for example that accumulation points of Lloyd’s method are critical points of the functional in (1.2), leaving the convergence of the whole sequence open [33].

In this work, we prove convergence of the sequence generated by both Lloyd’s methods under an explicit definability assumption on the density of the target measure μ . Note that, we do not consider convergence toward the global minimizer of the quantization functional and limit ourselves to critical points.

Convergence of gradient sequences and rigidity of semi-discrete optimal transport losses. Both results rely on the interpretation of Lloyd’s methods as gradient schemes for problems (1.1) and (1.2) [19, 33]. In a non-convex setting, the iterates of gradient sequences converge under Lojasiewicz inequality, for example if the underlying loss function is analytic [1]. Kurdyka proposed a generalization of Lojasiewicz’s property that holds true for all functions that are definable in an o-minimal structure [28], which we refer to as the Kurdyka-Lojasiewicz (KL) inequality. Definable functions represent a broad and versatile class of functions and covers many applications. For example this class is closed under composition, partial minimization / maximization, etc. Convergence of first order methods under KL assumptions have now become standard in non-convex optimization, see for example [3, 8] and references therein. This is the main path that we follow.

Applying Kurdyka’s result to (1.1) and (1.2) requires to justify that these functions are definable in some o-minimal structure. This is far from direct as the description of these functionals involve integrals and o-minimal structures are not stable under integration in general. This is one of the great challenges of this field, several partial answers are known, one of them allowing to treat the definability of globally subanalytic integrals [29, 11, 10] which have a log-analytic nature. This is the crucial step behind our convergence results, an observation already made in [7] in a stochastic optimization context. We work under the assumption that the compactly supported target measure has a globally subanalytic density (see Assumption 1). This allows to justify the resulting log analytic nature of the losses in (1.1) and (1.2). Convergence then follows using the well known connection between Lloyd’s method and the gradient algorithm and its convergence analysis under KL assumptions.

The proposed analysis opens the question of definability of semi-discrete optimal transport losses beyond the 2-Wasserstein distance in (1.1) and (1.2). We justify the definable nature of several such losses involving general optimal transport costs, sliced Wasserstein distance [37, 9, 5], the max-sliced Wasserstein distance [27, 15], and entropy regularized optimal transport [13]. These results are of independent interest, and may be relevant to study numerical applications of optimal transport.

Organization of the paper and notations. We start with a description of Lloyd’s methods and state our main results in Section 2. The central convergence arguments based on KL inequality are described in Section 3, and the definability of the underlying losses is described in Section 4. Extension to broader optimal transport losses is given in Section 5.

Regarding technical elements discussed in the introduction, we recall the definition of KL inequality and the main result of Kurdyka in Section 3.1, while the necessary background on o-minimal structures and globally subanalytic sets is given in Section 4. Furthermore, the precise definition of the W_2 distance in the semi-discrete setting of (1.1) and (1.2) is postponed to Section 5.1.

The Euclidean norm and the dot product in \mathbb{R}^d are denoted $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$. We define respectively the open and closed d -dimensional Euclidean ball centered in x with radius r as $B_d(x, r)$ and $\overline{B}_d(x, r)$.

2 Main results

We use the notation $Y \in (\mathbb{R}^d)^N$ to denote N points y_1, \dots, y_N in \mathbb{R}^d which are candidate support points for quantization. This section contains our main results regarding the sequential convergence of Lloyd's algorithms for optimal and uniform quantization.

2.1 Assumptions on the target measure

We will consider the following assumptions on the target measure μ and its density probability function f .

Assumption 1. μ is a probability measure on \mathbb{R}^d , with compact support, absolutely continuous with respect to Lebesgue measure, with globally subanalytic density f .

Absolute continuity and compacity ensure that Lloyd's iterations are well defined, see for example [20, Assumption 3.1]. The global subanalyticity assumption is a rigidity assumption which will allow to invoke KL inequality for our convergence analysis. The precise definition of global subanalyticity is given in Section 4.1. This is a very versatile notion for which we provide a few examples below.

Example 2.1. *The following are globally subanalytic density functions:*

- *The uniform density on a compact basic semi-algebraic set, of the form $S = \{x, P_i(x) \geq 0, i = 1, \dots, m\} \subset \mathbb{R}^d$, for some $m \in \mathbb{N}$, where $P_i: \mathbb{R}^d \rightarrow \mathbb{R}$ are polynomials.*
- *Densities of the form $x \mapsto f(x)\mathbb{1}_S(x)$ where S is a compact semi-algebraic set (as above) and $f: U \rightarrow \mathbb{R}$ is an analytic function on an open domain $U \supset S$ (Lemma 4.6). These include truncated Gaussians for example.*
- *Semi-algebraic densities including integrable rational functions and their restrictions to compact sets.*
- *Any finite mixture of the above density examples. This allows to consider nonsmooth and possibly discontinuous densities.*
- *For instance, any normalized image considered as a piecewise constant probability density on the plane is globally subanalytic.*

Remark 2.2. *Under Assumption 1, the density f is strictly positive on a full measure open dense set in $\text{Supp } \mu$, the support of the measure μ . Indeed the set $\{x, f(x) > 0\}$ is globally subanalytic and can be partitioned into a finite number of embedded submanifolds in \mathbb{R}^d [17, 4.8]. It is enough to keep those of dimension d only since those of dimension at most $d - 1$ have zero Lebesgue measure. We deduce that $S = \{x, f(x) > 0\}$ is non-empty, globally subanalytic and open, we have $\text{Supp } \mu = \text{cl}_{\mathbb{R}^d} S$ and $\dim(\text{Supp } \mu \setminus S) < d$, using [17, 4.7], which is the claimed statement.*

2.2 Optimal Quantization

In addition to the main hypothesis above, the target measure is assumed to have convex support throughout this section.

Assumption 2. μ is a probability measure as in Assumption 1, with convex support.

Algorithm 1: Lloyd's algorithm for optimal quantization

Input: μ absolutely continuous probability measure on \mathbb{R}^d with convex compact support, $Y_0 \in \text{Supp } \mu \setminus D_N$ as in (2.3).

for $n \in \mathbb{N}$ **do**

$Y_{n+1} = T_N(Y_n)$ where T_N is the Voronoi barycentric mapping in (2.5).

Partial minimization of the objective function in (1.1) with respect to the weights π leads to the *optimal quantization functional* $G_N: (\mathbb{R}^d)^N \rightarrow \mathbb{R}$, see e.g. [24],

$$G_N(Y) = \min_{\pi \in \Delta_N} \frac{1}{2} W_2^2 \left(\mu, \sum_{i=1}^N \pi_i \delta_{y_i} \right) = \frac{1}{2} \int \min_{i=1, \dots, N} \|x - y_i\|^2 f(x) dx, \quad (2.1)$$

where we recall that f stands for the density of μ . Problem (1.1) therefore amounts to minimize the objective G_N . This functional is tightly connected to the so-called Voronoi cells $(V_i(Y))_{i=1, \dots, N}$ associated to the support points $Y = (y_i)_{i=1}^N \in (\mathbb{R}^d)^N$ [19, Proposition 3.1], and given for $i = 1, \dots, N$ by

$$V_i(Y) = \left\{ x \in \mathbb{R}^d \mid \forall j = 1, \dots, N : \|x - y_i\|^2 < \|x - y_j\|^2 \right\}. \quad (2.2)$$

The Voronoi cells in (2.2) are open and disjoint, they are all non-empty as long as Y does not belong to the generalized diagonal defined by

$$D_N = \{Y \in (\mathbb{R}^d)^N \mid \exists i \neq j : y_i = y_j\}. \quad (2.3)$$

Therefore, for any $Y \notin D_N$, the optimal quantization functional in (2.1) takes the form

$$G_N(Y) = \frac{1}{2} \sum_{i=1}^N \int_{V_i(Y)} \|x - y_i\|^2 f(x) dx. \quad (2.4)$$

Note that a local minimizer to (1.1) on $\text{Supp } \mu$ does not lie in D_N [19, Proposition 3.5]. Lloyd's algorithm for computing optimal quantizers is then defined by fixed point iterations of the Voronoi barycentric mapping $T_N: \text{Supp } \mu \setminus D_N \rightarrow \text{Supp } \mu \setminus D_N$ defined by

$$T_N(Y) = \left(\frac{\int_{V_i(Y)} x d\mu(x)}{\mu(V_i(Y))} \right)_{i=1, \dots, N}, \quad (2.5)$$

Then the map T_N is indeed well defined on $\text{Supp } \mu \setminus D_N$ since, for distinct centroids in $\text{Supp } \mu$, the Voronoi cells in (2.2) are non-empty and have non-zero μ measure. Furthermore T_N takes values in $\text{Supp } \mu \setminus D_N$ as the centroids remain distinct since the Voronoi cells are disjoint, and they belong to the support of μ since it is convex by Assumption 2. This is summarized in Algorithm 1 and we prove in Section 3.2 the following theorem regarding convergence of the iterates.

Theorem 2.3. *Let μ be as in Assumption 2, then the iterates of Lloyd's algorithm for optimal quantization in Algorithm 1 converge to a critical point of G_N in (2.1).*

2.3 Uniform Quantization

The objective in (1.2) is called the *uniform quantization functional* $F_N: (\mathbb{R}^d)^N \rightarrow \mathbb{R}$, and is defined as follows

$$F_N(Y) = \frac{1}{2} W_2^2 \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right). \quad (2.6)$$

The uniform quantization problem consists in finding minimizers of F_N . Similar to optimal quantization, the functional can be explicitly described by power cells (also called Laguerre cells [33]). These generalize Voronoi cells and characterize the optimal transport plan in W_2 between the target measure and the discrete measure (see also (5.2)). If Y is not in the generalized diagonal D_N , then F_N in (2.6) can be written with the Kantorovich dual formulation introduced as follows for $w \in \mathbb{R}^N$ (see for example [41, Theorem 5.9] and [33, Section 2] for details on the uniform quantization setting)

$$\begin{aligned} \Phi(Y, w) &:= \frac{1}{2} \left\{ \int_{\mathbb{R}^d} \left(\min_{i=1, \dots, N} \|x - y_i\|^2 - w_i \right) f(x) dx + \frac{1}{N} \sum_{i=1}^N w_i \right\} \\ &= \frac{1}{2} \left\{ \sum_{i=1}^N \int_{\text{Pow}_i(Y, w)} (\|x - y_i\|^2 - w_i) f(x) dx + \frac{1}{N} \sum_{i=1}^N w_i \right\} \end{aligned} \quad (2.7)$$

where the power cells $\text{Pow}_i(Y, w)$ are defined for $i = 1, \dots, N$ as

$$\text{Pow}_i(Y, w) = \{x \in \mathbb{R}^d \mid \forall j = 1, \dots, N : \|x - y_i\|^2 - w_i \leq \|x - y_j\|^2 - w_j\}. \quad (2.8)$$

Kantorovich's duality states that for any $Y \notin D_N$, the uniform quantization functional in (2.6) is given by

$$F_N(Y) = \max_{w \in \mathbb{R}^N} \Phi(Y, w). \quad (2.9)$$

It can be shown (see Remark 3.8) that the maximum in (2.9) is attained and unique up to the addition of a constant. Therefore, it defines a unique set of power cells (2.8), and additionally the μ measure of each power cell is $\frac{1}{N}$.

We introduce $B_N: (\mathbb{R}^d)^N \setminus D_N \rightarrow (\mathbb{R}^d)^N \setminus D_N$ the Laguerre barycentric application defined by

$$B_N(Y) = \left(N \int_{\text{Pow}_i(Y, \phi(Y))} x d\mu(x) dx \right)_{i=1, \dots, N} \quad (2.10)$$

where $\phi(Y)$ denotes an element of $\text{argmax}_{w \in \mathbb{R}^N} \Phi(Y, w)$ in (2.7). Note that the Laguerre cells are disjoint and non-empty outside of the generalized diagonal D_N so that B_N indeed has values in $(\mathbb{R}^d)^N \setminus D_N$. Lloyd's algorithm consists in fixed point iterations of the map B_N as summarized in Algorithm 2. Our main result regarding this algorithm is the following, whose proof is postponed to Section 3.3.

Theorem 2.4. *Let μ be as in Assumption 1, then the iterates of Lloyd's algorithm for uniform quantization in Algorithm 2 converge to a critical point of F_N in (2.6).*

Remark 2.5. *Contrary to optimal quantization, the convexity of the support is not required here for the algorithm to be well defined, by construction of the power cells at optimality in (2.9). Additionally, the identity in (2.9) only holds if $Y \notin D_N$. It is nonetheless possible to extend it to equality cases, that is when there exists $i, j \in \{1, \dots, N\}, i \neq j$, such that $y_i = y_j$, and this is done in Theorem 5.2. This remark does not have any influence on our algorithmic results since the iterates remain away from the diagonal, but it will be essential to justify that the function F_N in (2.6) is definable on the whole space.*

Algorithm 2: Lloyd’s algorithm for uniform quantization

Input: μ absolutely continuous probability measure on \mathbb{R}^d with compact support,
 $Y_0 \in (\mathbb{R}^d)^N \setminus D_N$ as in (2.3).
for $n \in \mathbb{N}$ **do**
 $Y_{n+1} = B_N(Y_n)$ where B_N is the Laguerre barycentric mapping in (2.10).

2.4 Rigidity of quantization functionals and semi-discrete optimal transport losses

The key in obtaining our convergence results is to interpret Lloyd’s algorithms as gradient descent and rely on convergence analysis under Kurdyka-Lojasiewicz assumptions [1, 7]. A sufficient condition for applying this analysis is that the associated loss functions are definable in an o-minimal structure [28]. While definability is stable under many operations, such as composition, inverse, and projection, it is not a general rule that infinite summations or integrals preserve definability. For the specific case of globally subanalytic integrands, the resulting partial integrals have a log-analytic nature [29, 11, 10], so that such parameterized integrals are definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$ [16]. This was also observed in the context of stochastic optimization in [7]. We prove in Section 5 that several loss functions based on semi-discrete optimal transport divergences are definable in this o-minimal structure under Assumption 1, including the optimal quantization functional in (2.1) and the uniform quantization functional in (2.6). This definability result holds for optimal transport with general globally subanalytic cost functions (Section 6 in [41]), the max-sliced Wasserstein distance [27] and the entropy regularized Wasserstein divergence [23]. These constitute technical extensions of our main sequential convergence analysis which are of independent interest.

3 Convergence of the iterates of Lloyd’s algorithms

This section provides the proof arguments for Theorems 2.3 and 2.4. The main device is the interpretation of Lloyd’s algorithms as gradient methods combined with Kurdyka-Lojasiewicz (KL) inequality. This allows to invoke sequential convergence results for gradient methods on functions verifying KL inequality [1, 3]. The KL inequality is verified in our cases by Kurdyka’s sufficient condition, which is the definability of the quantization loss functions in (2.1) and (2.9). The proof of these definability results is independent of the convergence analysis and is given in Section 4.2.

3.1 KL inequality

Lojasiewicz’s gradient inequality is of the form $\|f(x) - f(y)\|^\theta \leq C\|\nabla f(x)\|$ with $\theta \in]0, 1[$ for a fixed y , and was first identified for real analytic functions [32]. Later on, Kurdyka [28] proposed a generalization to functions definable in an o-minimal structure now known as the Kurdyka-Lojasiewicz (KL) inequality. The KL inequality usually takes the following form.

Definition 3.1 (Definition 7 in [2]). *Let V be an open set of \mathbb{R}^d and \bar{Y} be a point in V . We say that a \mathcal{C}^1 function $F : V \rightarrow \mathbb{R}$ has the KL property at \bar{Y} if there exist $\eta \in]0, +\infty[$, a neighborhood U of \bar{Y} in V and a continuous positive concave function $\Psi : [0, \eta[\rightarrow \mathbb{R}_+$ such that:*

1. $\Psi(0) = 0$.
2. $\Psi \in \mathcal{C}^1(]0, \eta[)$.

3. $\forall t \in]0, \eta[, \Psi'(t) > 0$.

4. $\forall Y \in U \cap \{Y \in \mathbb{R}^d \mid F(\bar{Y}) < F(Y) < F(\bar{Y}) + \eta\}$ the KL inequality holds:

$$\Psi'(F(Y) - F(\bar{Y})) \|\nabla F(Y)\| \geq 1.$$

If F has the KL property at all $\bar{Y} \in V$, we say that F has the KL property on V .

Theorem 3.2 (Theorem 1 in [28]). *Let V be an open set on \mathbb{R}^d and $f: V \rightarrow \mathbb{R}$ be C^1 and definable in an o-minimal structure. Then f satisfies the KL property, as described in Definition 3.1.*

We postpone a precise definition and detailed discussion of definability of semi-discrete optimal transport losses to Section 4. In particular, we will show that the objective functions G_N and F_N defined in (2.1) and (2.9) are definable in some o-minimal structure, under Assumption 1.

As a consequence of Theorem 3.2, in the following we may use the KL inequality for G_N and F_N to prove the convergence of Lloyd's iterates. Lloyd's algorithms can indeed be interpreted as gradient methods (see (3.1) and (3.5)) for which convergence analysis under KL assumptions has become standard [1, 2, 3, 8].

3.2 Lloyd's sequence for optimal quantization

Lloyd's method for optimal quantization, given in Algorithm 1, can be understood as a gradient descent algorithm on the optimal quantization objective function G_N in (2.1). Throughout this section, we interpret $Y \in (\mathbb{R}^d)^N$ as a matrix in $\mathbb{R}^{N \times d}$. The function G_N is differentiable outside of the generalized diagonal, and its gradient is explicit, as stated in the following proposition.

Proposition 3.3 (Proposition 6.2 in [19]). *Let μ be as in Assumption 2 and $Y \in (\mathbb{R}^d)^N \setminus D_N$. Then G_N is differentiable in Y with gradient*

$$\nabla G_N(Y) = M(V(Y)) (Y - T_N(Y)), \quad (3.1)$$

where $M(V(Y)) \in \mathbb{R}^{N \times N}$ is the diagonal matrix with entries $(\mu(V_i(Y)))_{i=1, \dots, N}$, with $(V_i(Y))_{i=1, \dots, N}$ the Voronoi cells in (2.2), and T_N is the Voronoi barycentric mapping defined in (2.5).

We denote $(Y_n)_{n \geq 0}$ the Lloyd sequence in Algorithm 1. Since $Y_n \in \text{Supp } \mu \setminus D_N$ for all n , the corresponding Voronoi cells are non empty and the diagonal matrix $M(V(Y_n))$ in Proposition 3.3 is invertible for all $n \in \mathbb{N}$. Therefore, Y_n can be rewritten for all $n \geq 1$ in a pre-conditioned gradient descent formulation as

$$Y_{n+1} = Y_n - M(V(Y_n))^{-1} \nabla G_N(Y_n). \quad (3.2)$$

It turns out that the diagonal matrix $M(V(Y_n))$ remains bounded away from 0 for all $n \in \mathbb{N}$ as a consequence of [20, Theorem 3.6] since Assumption 1 ensures that [20, Assumption 3.1] holds; see Remark 2.2.

Proposition 3.4 (Corollary 3.7 in [20]). *Let μ be as in Assumption 2 and $(Y_n)_{n \geq 0}$ be the Lloyd sequence defined in Algorithm 1, then there exists $\ell > 0$ such that for all $i = 1, \dots, N$*

$$\inf_{n \geq 0} \mu(V_i(Y_n)) \geq \ell. \quad (3.3)$$

As a consequence, we get the following strong descent condition.

Lemma 3.5. Let μ be as in Assumption 2 and $(Y_n)_{n \geq 0}$ be the Lloyd sequence defined Algorithm 1. Then, given ℓ as in (3.3), for all $n \in \mathbb{N}$,

$$G_N(Y_n) - G_N(Y_{n+1}) \geq \frac{\ell}{2} \|\nabla G_N(Y_n)\| \|Y_{n+1} - Y_n\|. \quad (\text{SDC})$$

Proof. Recall that $Y_n \notin D_N$ for all $n \in \mathbb{N}$. Let us define for all $Y, Z \in (\mathbb{R}^d)^N \setminus D_N$ and for $i = 1, \dots, N$ the function $\mathcal{H}_i(Y, Z) = \int_{V_i(Y)} \|x - z_i\|^2 f(x) dx$. Then for all $i = 1, \dots, N$, we have

$$\begin{aligned} \mathcal{H}_i(Y_n, Y_n) &= \int_{V_i(Y_n)} \|x - y_i^n\|^2 f(x) dx \\ &= \int_{V_i(Y_n)} \|x - y_i^{n+1}\|^2 f(x) dx + \int_{V_i(Y_n)} \|y_i^{n+1} - y_i^n\|^2 f(x) dx \\ &= \mathcal{H}_i(Y_n, Y_{n+1}) + \mu(V_i(Y_n)) \|y_i^{n+1} - y_i^n\|^2. \end{aligned}$$

where the first equality is verified since $\int_{V_i(Y_n)} \langle x - y_i^{n+1}, y_i^{n+1} - y_i^n \rangle d\mu(x) = 0$ by definition of the sequence $(Y_n)_{n \geq 0}$ in (3.2). Additionally, denoting $\mathcal{H}(Y, Z) = \sum_{i=1}^N \mathcal{H}_i(Y, Z)$, we get by Lemma 2.1 of [18] and using (2.4), that for all $n \in \mathbb{N}$,

$$\min_{Z \in (\mathbb{R}^d)^N} \mathcal{H}(Z, Y_n) = \mathcal{H}(Y_n, Y_n) = 2G_N(Y_n).$$

We deduce that

$$\begin{aligned} G_N(Y_n) &= \frac{1}{2} \sum_{i=1}^N \mathcal{H}_i(Y_n, Y_n) \geq G_N(Y_{n+1}) + \frac{1}{2} \sum_{i=1}^N \mu(V_i(Y_n)) \|y_i^{n+1} - y_i^n\|^2 \\ &\geq G_N(Y_{n+1}) + \frac{1}{2} \sum_{i=1}^N \mu(V_i(Y_n))^2 \|y_i^{n+1} - y_i^n\|^2 \\ &= G_N(Y_{n+1}) + \frac{1}{2} \|\nabla G_N(Y_n)\| \left(\sum_{i=1}^N \mu(V_i(Y_n))^2 \|y_i^{n+1} - y_i^n\|^2 \right)^{1/2} \\ &\geq G_N(Y_{n+1}) + \frac{1}{2} \|\nabla G_N(Y_n)\| \min_{i=1, \dots, N} \mu(V_i(Y_n)) \|Y_{n+1} - Y_n\|. \end{aligned} \quad (3.4)$$

From this we get condition (SDC) using Proposition 3.4. \square

This strong descent condition allows to prove sequential convergence provided that G_N is a KL function [1]. This is stated in the following lemma whose proof rely on definability results postponed to Section 4.

Lemma 3.6. Let μ be as in Assumption 2, then the optimal quantization functional G_N defined in (2.1) is a KL function on $(\mathbb{R}^d)^N \setminus D_N$.

Proof. This result is a consequence of Lemma 4.11 and Theorem 3.2. \square

The standard analysis of gradient method under KL assumption allows to conclude about convergence of Lloyd's iterates for optimal quantization, the proof of which is given below.

Proof of Theorem 2.3. From Lemma 3.6, we have that G_N is a KL function. We can then apply [1, Theorem 3.4], which requires the strong descent condition (SDC) and the property that $G_N(Y_{n+1}) = G_N(Y_n)$ implies $Y_{n+1} = Y_n$, a property satisfied in our case due to the descent condition and the expression in (3.2). The convergence of the iterates then follows. Note that [1, Theorem 3.4] is stated without a domain. However the main proof mechanism is a trap argument which is purely local and thus can easily be extended to a closed set containing all the iterates in an open domain. In our setting we work on $(\mathbb{R}^d)^N \setminus D_N$, and Proposition 3.4 ensures that the sequence $(Y_n)_{n \geq 0}$ in Algorithm 1 remains away from the generalized diagonal D_N so that the trap argument applies. See also [3] for a general account using KL functions. \square

3.3 Lloyd's sequence for uniform quantization

The objective function F_N in (2.9) is differentiable outside the generalized diagonal, and the expression of its gradient is given in the following proposition.

Proposition 3.7 (Proposition 1 in [33]). *Let μ be as in Assumption 1, then for all $Y \in (\mathbb{R}^d)^N \setminus D_N$, we have*

$$\nabla F_N(Y) = \frac{1}{N} (Y - B_N(Y)), \quad (3.5)$$

where B_N is the Laguerre barycentric mapping described in (2.10).

Remark 3.8. *The vector $\phi(Y) \in \operatorname{argmax}_{w \in \mathbb{R}^N} \Phi(Y, w)$ appearing in F_N in the definition of (2.9) is unique up to the addition of a constant (see e.g. Theorem 1.17 in [38]). In order to have uniqueness, it is sufficient to add a constraint on the set of maximizers, such as setting the mean of the components of w to be 0 as done in [26]. Furthermore, all elements in the argmax describe the same power cells.*

As in Section 3.2, Proposition 3.7 allows to interpret the Lloyd sequence for uniform quantization in Algorithm 2 as gradient iterations on the uniform quantization functional F_N in (2.9). Similarly it is possible to show that F_N restricted to $(\mathbb{R}^d)^N \setminus D_N$ is a KL function, based on its definability, whose proof is postponed to Section 4.

Lemma 3.9. *Let μ be as in Assumption 1, then the objective function F_N defined in (2.9) is a KL function on $(\mathbb{R}^d)^N \setminus D_N$.*

Proof. It is a direct consequence of Lemma 4.12 and Theorem 3.2. \square

We follow the same line as in the optimal quantization setting of Section 3.2.

Proof of Theorem 2.4. The iterates remain in $(\mathbb{R}^d)^N \setminus D_N$ so that by Proposition 3.7 we have $Y_{n+1} = Y_n - N \nabla F_N(Y_n)$ for all $n \in \mathbb{N}$. From the proof of Proposition 2 in [33], we have the following inequality

$$F_N(Y_n) - F_N(Y_{n+1}) \geq \frac{1}{2N} \|Y_{n+1} - Y_n\|^2,$$

for all $n \in \mathbb{N}$, which is a strong descent condition and also entails that $F_N(Y_{n+1}) = F_N(Y_n)$ implies $Y_{n+1} = Y_n$. Finally Lemma 3.9 ensures that F_N is a KL function. The convergence follows from [1, Theorem 3.4], see also [3] for general KL functions. Note that similarly as in the proof of Theorem 2.3, the trap argument can be extended to a compact subset containing all the iterates in $(\mathbb{R}^d)^N \setminus D_N$ as described in the proof of Proposition 2 in [33]. \square

4 Rigidity of semi-discrete optimal transport quantization losses

In this section we show the definability in an o-minimal structure of the optimal and uniform quantization functionals G_N and F_N defined respectively in (2.1) and (2.9). This finishes the proof arguments for our main convergence results for Lloyd's algorithms in Theorem 2.3 and 2.4 since by Kurdyka's result recalled in Theorem 3.2, they are therefore KL functions. We start with the introduction of o-minimal structures and recall the main results that will be needed further in this section.

4.1 O-minimal structures

O-minimal structures describe families of subsets of Euclidean spaces which preserve the favorable rigidity properties of semi-algebraic sets. The description is axiomatic. We refer the reader to the lecture notes from Coste [12] for an introduction, to [17] for an extensive presentation of consequences of o-minimality and to [31] for additional bibliographic pointers. We limit ourselves to structures expanding the field of real numbers. For statements involving definable objects, if not precised otherwise, the word “*definable*” implicitly means that all objects are definable in the same o-minimal structure.

Definition 4.1. Let $\mathcal{S} = (\mathcal{S}_n)_{n \in \mathbb{N}}$ be such that for each $n \in \mathbb{N}$, \mathcal{S}_n is a family of subsets of \mathbb{R}^n . It is called an o-minimal structure if it satisfies the following:

- for each $n \in \mathbb{N}$, \mathcal{S}_n contains all algebraic subsets (defined by polynomial equalities).
- for each $n \in \mathbb{N}$, \mathcal{S}_n is a boolean sub-algebra of \mathbb{R}^n : stable under intersection, union and complement.
- \mathcal{S} is stable under Cartesian product and projection on lower dimensional subspace.
- \mathcal{S}_1 is exactly the set of finite unions of points and intervals.

A set belonging to an o-minimal structure is said to be *definable* (in this structure). A function whose graph or epigraph is definable is also called *definable* (in this structure).

4.1.1 Semi-algebraic subsets

Semi-algebraic sets represent the smallest o-minimal structure. In other words, any o-minimal structure contains all the semi-algebraic sets.

Definition 4.2. A set S is a basic semi-algebraic subset of \mathbb{R}^n if there exists a finite number of real polynomial functions $(f_i)_{i=1, \dots, N}$ and $(g_j)_{j=1, \dots, M}$ such that:

$$S = \{x \in \mathbb{R}^n \mid f_1(x) > 0, \dots, f_N(x) > 0, g_1(x) = 0, \dots, g_M(x) = 0\}.$$

A subset which is a finite union of basic semi-algebraic subsets is semi-algebraic.

The following fact is a consequence of Tarski-Seidenberg theorem, see for example [12].

Proposition 4.3. The collection of all semi-algebraic subsets, denoted \mathcal{SA} , forms an o-minimal structure.

4.1.2 Globally subanalytic subsets

The o-minimal structure of globally subanalytic sets is the one appearing in Assumption 1. It is the smallest structure which contains all restricted analytic functions, a formal definition is given below.

Definition 4.4. *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a restricted analytic function if there exists a function F analytic on an open set containing $[-1, 1]^n$, such that $f = F$ on $[-1, 1]^n$ and $f = 0$ otherwise.*

Proposition 4.5 ([40]). *There exists an o-minimal structure which contains the graph of all restricted analytic functions. There is a smallest such structure denoted by \mathbb{R}_{an} .*

A subset of \mathbb{R}_{an} is called globally subanalytic and definable functions in \mathbb{R}_{an} are also called globally subanalytic. Note that any o-minimal structure contains the semi-algebraic subsets, thereby they are also globally subanalytic. The following Lemma illustrates the versatility of globally subanalytic sets and allows to consider the probability densities in Example 2.1. This is a very well known fact for which we provide a proof for completeness in Appendix A.

Lemma 4.6. *Let U be an open subset of \mathbb{R}^n , $f : U \rightarrow \mathbb{R}$ be an analytic function, and $K \subset U$ be a semi-algebraic compact subset of U . Then $f|_K : K \rightarrow \mathbb{R}$ is globally subanalytic.*

4.1.3 Inclusion of the exponential function

The structure \mathbb{R}_{an} contains the graph of the exponential function restricted to any bounded interval (because it is restricted analytic), but not on the whole real line. There is a bigger o-minimal structure which contains the whole graph of the exponential function, as the following shows.

Proposition 4.7 ([16]). *There exists an o-minimal structure which contains all functions definable in \mathbb{R}_{an} and the graph of the exponential function. There is a smallest such structure denoted by $\mathbb{R}_{an,exp}$.*

Note that $\mathbb{R}_{an,exp}$ also contains the graph of the logarithm function which is identical to that of the exponential function up to a symmetry.

4.2 Definability of optimal transport quantization functionals

We start with the description of a technical result on the integration of globally subanalytic functions [10] and then describe how it applies to our semi-discrete quantization losses under Assumption 1.

4.2.1 Integration of globally subanalytic functions and partial minimization

O-minimal structures are not stable under integration in general. This question constitutes one of the great challenges of the field. It is therefore not possible to directly use definability properties of the objective functions G_N in (2.1) and F_N in (2.9). Several partial answers are known, one of them allowing to treat globally subanalytic integrands [29, 11, 10]. The following lemma is a direct consequence of [10, Theorem 1.3].

Lemma 4.8. *Let $f : \mathbb{R}^d \times U \rightarrow \mathbb{R}$ be a globally subanalytic function with $U \subset \mathbb{R}^n$ open (globally subanalytic). Define $y \mapsto F(y) := \int_{\mathbb{R}^d} f(x, y) dx$ and suppose it is well defined for all y in U . Then F is definable in $\mathbb{R}_{an,exp}$.*

Proof. By [10, Theorem 1.3], if there exist $(f_i)_{i=1,\dots,k}, (g_{i,j})_{\substack{i=1,\dots,k \\ j=1,\dots,l_i}}$ two families of globally subanalytic functions such that f can be written

$$(x, y) \mapsto f(x, y) = \sum_{i=1}^k f_i(x, y) \prod_{j=1}^{l_i} \log g_{i,j}(x, y), \quad (4.1)$$

then the parameterized integral $y \mapsto F(y) = \int_{\mathbb{R}^d} f(x, y) dx$ can also be written in the same form. By Lemma A.1, a function expressed as a sum and product of globally subanalytic functions and of the logarithm of globally subanalytic functions is definable in $\mathbb{R}_{\text{an,exp}}$. \square

We conclude this section with the following technical result, which is well known and whose proof is postponed to Appendix A.

Lemma 4.9.

- (i) Let $(f_j)_{j \in J}$ be a finite collection of definable functions. Then $\min_{j \in J} f_j$ and $\max_{j \in J} f_j$ are definable functions.
- (ii) Let $f : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a definable function and E be a definable subset of \mathbb{R}^n . Then, provided they are attained, $\min_{w \in E} f(\cdot, w)$ and $\max_{w \in E} f(\cdot, w)$ are definable functions.

Remark 4.10. Lemma 4.9 remains true for the infimum and the supremum.

4.2.2 Definability of the optimal quantization functional

Let us prove that the objective function G_N in (2.1) is definable.

Lemma 4.11. Under Assumption 1, the optimal quantization function G_N defined in (2.1) is definable in $\mathbb{R}_{\text{an,exp}}$.

Proof. Let f be the globally subanalytic function corresponding to the density of the target measure μ . By stability of definability for the minimum function (Lemma 4.9), we have that the function $(x, Y) \mapsto \min_{i=1,\dots,N} \|x - y_i\|^2 f(x)$ is globally subanalytic. Since this function is integrable we can apply Lemma 4.8 and obtain that G_N in (2.1) is definable in $\mathbb{R}_{\text{an,exp}}$. \square

4.2.3 Definability of the uniform quantization functional

We prove that the objective function F_N in (2.6) is definable. This is based on the representation in (2.7) and (2.9) for $Y \notin D_N$, see Remark 2.5.

Lemma 4.12. Under Assumption 1, the uniform quantization function F_N defined in (2.9) is definable in $\mathbb{R}_{\text{an,exp}}$.

Proof. Using Lemma 4.8, the function Φ defined in (2.7) is definable in $\mathbb{R}_{\text{an,exp}}$. From Lemma 4.9, the function F_N is also definable. Strictly speaking this argument works outside of the generalized diagonal, but the equality cases can be treated similarly as described in Remark 2.5 and Theorem 5.2. \square

5 Beyond semi-discrete W_2 losses

We extend the definability results of the previous section to more general *semi-discrete loss functions* $F : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$ defined by

$$F(Y) = D \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right), \quad (5.1)$$

where D denotes an optimal transport divergence, and μ is a probability measure. In particular we consider optimal transports with general cost (Section 5.1), the max-sliced Wasserstein distance (Section 5.2) and the entropy regularized optimal transport problem (Section 5.3). For all these cases, we prove definability in $\mathbb{R}_{\text{an,exp}}$ (introduced in Proposition 4.7) of the resulting semi-discrete loss in (5.1). These results illustrate the relevance of the log analytic nature of integral of globally subanalytic objects in a semi-discrete optimal transport context.

Throughout this section, the measure μ is assumed to satisfy Assumption 1 and we introduce a general cost function c as follows.

Assumption 3. *The cost function $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ is lower semicontinuous and globally subanalytic.*

Remark 5.1. *The compactness of the support of μ is not strictly required and could be replaced by integrability conditions. Yet, for the sake of simplicity, we consider μ as in Assumption 1, especially to preserve an homogeneous set of assumptions throughout the text. Additionally, the results of this section directly hold replacing the uniform weights in (5.1) by any fixed weights $(\pi_1, \dots, \pi_N) \in \Delta_N$.*

5.1 Kantorovich duality and definability for semi-discrete optimal transport losses with general costs

The following is a reformulation of Kantorovich duality for optimal transport, specified for semi-discrete losses; see e.g. [41, Theorem 5.9]. We pay special attention to ties in the Dirac mass support points. Let μ and c be as in Assumptions 1 and 3 and let $Y = (y_i)_{i=1}^N \in (\mathbb{R}^d)^N$ be a set of N points in \mathbb{R}^d . We set

$$\mathcal{T}_c \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) = \inf_{\gamma \in \Pi} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma(x, y), \quad (5.2)$$

where the infimum is taken over couplings (or plans) $\gamma \in \Pi$ whose first marginal is μ and whose second marginal is the discrete measure $\frac{1}{N} \sum_{i=1}^N \delta_{y_i}$. The loss functions in (1.1) and (1.2) correspond to the 2-Wasserstein distance squared for which the cost function c is the squared Euclidean distance, i.e. $c(x, y) = \|x - y\|^2$ for $x, y \in \mathbb{R}^d$.

We let $\lambda : (\mathbb{R}^d)^N \rightarrow \mathbb{R}^N$ be such that for all $i = 1, \dots, N$ and $Y \in (\mathbb{R}^d)^N$,

$$\lambda_i(Y) = \begin{cases} \frac{|\{j, j \geq i, y_j = y_i\}|}{N} & \text{if } y_k \neq y_i \forall k < i, \\ 0 & \text{otherwise.} \end{cases} \quad (5.3)$$

For $Y \notin D_N$, the generalized diagonal defined in (2.3), we have $\lambda_i(Y) = \frac{1}{N}$, $i = 1, \dots, N$. When there are ties with equal points, λ assign the mass of the whole group to the first index and zero to the remaining indices. Note that $\lambda_i \geq 0$ and $\sum_{i=1}^N \lambda_i = 1$ and the function λ is semi-algebraic (it is constant on finitely many disjoint pieces defined by linear equalities and their complements). We set $g_c : (\mathbb{R}^d)^N \times \mathbb{R}^N \rightarrow \mathbb{R}$ as follows

$$g_c(Y, w) = \int_{\mathbb{R}^d} \min_{i: \lambda_i(Y) > 0} \{c(x, y_i) - w_i\} d\mu(x) + \sum_{i=1}^N \lambda_i(Y) w_i. \quad (5.4)$$

Theorem 5.2. *Let μ be a probability measure on \mathbb{R}^d , $Y = (y_i)_{i=1}^N \in (\mathbb{R}^d)^N$ be a set of N points in \mathbb{R}^d and $c: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ be lower semicontinuous. Then*

$$\mathcal{T}_c \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) = \max_{w \in \mathbb{R}^n} g_c(Y, w),$$

where \mathcal{T}_c is given in (5.2) and g_c is given in (5.4).

Proof. This is a direct implication of [41, Theorem 5.9]. The design of the function λ in (5.4) ensures that

$$\nu := \frac{1}{N} \sum_{i=1}^N \delta_{y_i} = \sum_{i=1}^N \lambda_i(Y) \delta_{y_i},$$

where, furthermore, the set of weights $\lambda_i(Y) > 0$ corresponds to distinct points y_i , all ties being merged. A direct inspection of the function g_c in (5.4) shows that it does not depend on variables w_i corresponding to zero weights $\lambda_i(Y) = 0$. Therefore we may ignore them and consider $l \leq N$ non-zero weights $\lambda_i(Y)$. Assuming that they correspond to the first l weights (up to a permutation that keeps g_c and \mathcal{T}_c invariant), we obtain

$$\begin{aligned} \nu &= \sum_{i=1}^l \lambda_i(Y) \delta_{y_i}, \\ g_c(Y, w) &= \int_{\mathbb{R}^d} \min_{i=1, \dots, l} \{c(x, y_i) - w_i\} d\mu(x) + \sum_{i=1}^l \lambda_i(Y) w_i, \end{aligned}$$

where the points y_1, \dots, y_l are pairwise distinct and form the support of the measure ν . In this setting, we may interpret the dual variables w_1, \dots, w_l as a function $\phi \in L^1(\nu)$ by setting $\phi(y_i) = w_i$ for $i = 1, \dots, l$, since support points are pairwise distinct. The claimed result then follows from Kantorovich duality [41, Theorem 5.9] by noticing that with this interpretation one has $\int \phi d\nu = \sum_{i=1}^l \lambda_i(Y) w_i$ and $\min_{i=1, \dots, l} \{c(x, y_i) - w_i\} = -\max_{y \in \{y_1, \dots, y_l\}} \{\phi(y) - c(x, y)\} = -\phi^c(x)$, the c -transform of ϕ . \square

The following is our main definability result for this section.

Lemma 5.3. *Let μ be as in Assumption 1 and c be as in Assumption 3, then the loss function*

$$F_c : Y \in (\mathbb{R}^d)^N \mapsto \mathcal{T}_c \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) \quad (5.5)$$

is definable in $\mathbb{R}_{an,exp}$ where \mathcal{T}_c is defined in (5.2).

Proof. We aim at showing that the function g_c appearing in the Kantorovich dual formulation in Theorem 5.2, described in (5.4), is definable and the result will follow from Lemma 4.9.

First the function λ is semi-algebraic as it is constant on finitely many subsets defined by finitely many linear equalities and their complements. Under Assumption 1, we have

$$g_c(Y, w) = \int_{\mathbb{R}^d} \min_{i: \lambda_i(Y) > 0} \{c(x, y_i) - w_i\} f(x) dx + \sum_{i=1}^N \lambda_i(Y) w_i,$$

where the density f of μ and the cost c are globally subanalytic. This property is preserved under finite minimization and composition (Lemma 4.9 and Lemma A.1), and therefore the integrand in g_c is a globally subanalytic function of (x, Y, w) . Lemma 4.8 ensures that g_c is definable in $\mathbb{R}_{\text{an}, \text{exp}}$. This concludes the proof. \square

5.2 The sliced and max-sliced Wasserstein distances

Sliced Wasserstein. For any $\theta \in \mathbb{S}^{d-1}$ the unit sphere of \mathbb{R}^d , we set $P_\theta : x \in \mathbb{R}^d \mapsto \langle \theta, x \rangle$, the projection on the line directed by θ , and we let σ be the uniform probability measure over \mathbb{S}^{d-1} . Given μ and ν two probability measures on \mathbb{R}^d , the sliced Wasserstein distance is defined as follows:

$$\text{SW}_2^2(\mu, \nu) = \int_{\mathbb{S}^{d-1}} W_2^2(P_{\theta\#}\mu, P_{\theta\#}\nu) d\sigma(\theta). \quad (5.6)$$

Here W_2^2 still denotes the 2-Wasserstein distance squared and corresponds to the loss in (5.2) with $c(x, y) = (x - y)^2$ (the distance is taken between univariate measures). We recall that the pushforward operator of a measure μ in \mathbb{R}^d by a measurable map $T : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is defined as the measure $T\#\mu$ such that for all Borelian $B \subset \mathbb{R}^n$, $T\#\mu(B) = \mu(T^{-1}(B))$. We present definability of the sliced Wasserstein distance between two discrete probability measures. Note that even in the fully discrete setting, the divergence in (5.6) involves an integral over the sphere and therefore does not falls directly within the scope of definable functions.

Lemma 5.4. *Let $N, M \in \mathbb{N}^*$, then the following function is definable in $\mathbb{R}_{\text{an}, \text{exp}}$*

$$F_{\text{SW-discr}} : (\mathbb{R}^d)^N \times (\mathbb{R}^d)^M \times \Delta_N \times \Delta_M \rightarrow \mathbb{R}$$

$$(X, Y, a, b) \mapsto \text{SW}_2^2 \left(\sum_{i=1}^N a_i \delta_{x_i}, \sum_{j=1}^M b_j \delta_{y_j} \right).$$

Proof. By definition of the pushforward of a discrete measure by P_θ , the integrand in the definition of SW_2^2 in (5.6) corresponds to

$$(\theta, X, Y, a, b) \mapsto W_2^2 \left(\sum_{i=1}^N a_i \delta_{\langle \theta, x_i \rangle}, \sum_{j=1}^M b_j \delta_{\langle \theta, y_j \rangle} \right).$$

This function is semi-algebraic, for example it can be written explicitly as a finite dimensional linear program with semi-algebraic data or using quantile functions for finitely many support points, which are semi-algebraic. The result follows from Lemma 4.8. \square

The functional $F_{\text{SW-discr}}$ is studied in [39] in the context of machine learning applications, where convergence of a stochastic gradient descent algorithm is described using the weakly convex nature of the resulting loss function. Unlike the case of optimal transport with general costs, the definability of semi-discrete sliced Wasserstein loss does not follow directly from [10, Theorem 1.3]. Indeed, the main result states that a certain class of log-analytic functions is stable under integration. However, log-analytic functions are not stable under maxima or absolute value, which would be required to conclude regarding the semi-discrete sliced Wasserstein loss. Still, it is reasonable to believe that the semi-discrete loss in (5.1) is definable if D is the sliced Wasserstein distance, but this result is out of the scope of the present work.

Max-sliced Wasserstein. Replacing the integral in (5.6) by a maximum over the unit sphere leads to the max-sliced Wasserstein problem [27, 15] defined as

$$\text{MSW}_2^2(\mu, \nu) = \max_{\theta \in \mathbb{S}^{d-1}} W_2^2(P_{\theta\#}\mu, P_{\theta\#}\nu). \quad (5.7)$$

In practice, it requires an estimator for the maximum (see the remark following the Claim 3 in [15]). The following shows a definability result for the semi-discrete max-sliced Wasserstein distance.

Lemma 5.5. *Let μ be as in Assumption 1 and c be as in Assumption 3, then the objective function*

$$F_{\text{MSW}} : Y \in (\mathbb{R}^d)^N \mapsto \text{MSW}_2^2 \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) \quad (5.8)$$

is definable in $\mathbb{R}_{\text{an,exp}}$, where MSW_2^2 is defined in (5.7).

Proof. For any $\theta \in \mathbb{S}^{d-1}$, we have that the pushforward of the discrete measure by P_θ is given by $\frac{1}{N} \sum_{i=1}^N \delta_{\langle y_i, \theta \rangle}$. Using Kantorovich duality (Theorem 5.9 in [41]), following the same construction as in Section 5.1 we have

$$\begin{aligned} \text{MSW}_2^2 \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) &= \max_{\theta \in \mathbb{S}^{d-1}, w \in \mathbb{R}^N} \int_{\mathbb{R}} \min_{\lambda_i(Y) > 0} \{ (t - \langle y_i, \theta \rangle)^2 - w_i \} dP_{\theta\#}\mu(t) + \sum_{i=1}^N \lambda_i(Y) w_i \\ &= \max_{\theta \in \mathbb{S}^{d-1}, w \in \mathbb{R}^N} \int_{\mathbb{R}^d} \min_{\lambda_i(Y) > 0} \{ \langle x - y_i, \theta \rangle^2 - w_i \} f(x) dx + \sum_{i=1}^N \lambda_i(Y) w_i, \end{aligned}$$

where the first equality follows from Theorem 5.2 and λ is given as in (5.3) for $d = 1$ applied to the projections $\langle \theta, y_i \rangle$, $i = 1, \dots, N$. The second equality corresponds the change of variable for image measures (for example [6, Theorem 3.6.1]). The integrand is globally subanalytic so by Lemma 4.8, the integral is definable in $\mathbb{R}_{\text{an,exp}}$, and so is the maximum by Lemma 4.9. □

5.3 Entropic regularization

Entropic regularization is widely used in optimal transport and can for example induce desirable computational and algorithmic features [13] (see also section 4.1 of [36]). This regularization was considered in a semi-discrete setting in [23, Section 2]. Given $\epsilon > 0$ and a general cost function $c: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$, the entropy regularized optimal transport loss W_ϵ between two probability measures μ and ν is defined as follows

$$W_\epsilon(\mu, \nu) = \min_{\gamma \in \Pi} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma(x, y) + \epsilon \mathbf{H}(\gamma | \mu \otimes \nu) \quad (5.9)$$

where Π denotes the set of measures on $\mathbb{R}^d \times \mathbb{R}^d$ whose first and second marginals are μ and ν , and $\mu \otimes \nu$ is the product measure. The function H denotes the Kullback-Leibler divergence, which is defined for probability measures γ, ζ on $\mathbb{R}^d \times \mathbb{R}^d$ by:

$$H(\gamma|\zeta) = \int_{\mathbb{R}^d \times \mathbb{R}^d} \left[\log \left(\frac{d\gamma}{d\zeta}(x, y) \right) - 1 \right] d\gamma(x, y)$$

and $\frac{d\gamma}{d\zeta}$ denotes the relative density of γ with respect to ζ (the divergence is infinite if γ is not absolutely continuous with respect to ζ).

Lemma 5.6. *Let μ be as in Assumption 1 and c be as in Assumption 3. Suppose that $R > 0$ is such that $\text{Supp } \mu \subset B_d(0, R)$, the unit ball of radius R , and c is in addition analytic on an open set containing $B_d(0, R) \times B_d(0, R)$. Then given a fixed $\epsilon > 0$, the entropy regularized transport functional (restricted to the R -radius ball)*

$$F_\epsilon : Y \in B_d(0, R)^N \longmapsto W_\epsilon \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) \quad (5.10)$$

is definable in $\mathbb{R}_{an,exp}$.

Proof. Using [23, Proposition 2.1], we have that W_ϵ can be written in dual form

$$\begin{aligned} W_\epsilon \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) &= \max_{w \in \mathcal{C}(B_d(0, R))} \int_{\mathbb{R}^d} -\epsilon \log \left(\frac{1}{N} \sum_{i=1}^N \exp \left(\frac{w(y_i) - c(x, y_i)}{\epsilon} \right) \right) d\mu(x) \\ &\quad + \frac{1}{N} \sum_{i=1}^N w(y_i) - \epsilon \end{aligned} \quad (S_\epsilon)$$

where the soft-max approximation of the c -transform of w appears and the maximum is over continuous functions. We first remark that the addition of a constant to w in (S_ϵ) does not change the value of the objective, so we may assume for example that $w(y_1) = 0$, without modifying the value of the maximum. Furthermore, it is known that since c is \mathcal{C}^∞ by assumption, the optimal potential w in (S_ϵ) is Lipschitz continuous with the same Lipschitz constant as the cost c , say on $B_d(0, R)$ [22, Proposition 1]. More precisely, it can be deduced from the smooth c -transform formulation of the dual variable and a first order condition. Hence the potential w in (S_ϵ) may be assumed to be uniformly bounded by a constant K which only depends on c and R . Now we may use the same device as in Theorem 5.2 and obtain

$$\begin{aligned} W_\epsilon \left(\mu, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) &= \max_{\|w\|_\infty \leq K, w_1=0} \int_{\mathbb{R}^d} -\epsilon \log \left(\sum_{i=1}^N \lambda_i(Y) \exp \left(\frac{w_i - c(x, y_i)}{\epsilon} \right) \right) d\mu(x) \\ &\quad + \sum_{i=1}^N \lambda_i(Y) w_i - \epsilon \end{aligned} \quad (5.11)$$

where λ is the semi-algebraic function in (5.3) which allows to break ties. The representation in (5.11) involves the logarithm restricted to a compact interval of the form $[a, b]$ where $a > 0$, and the exponential function also restricted to a compact interval. Both are restricted analytic and hence globally subanalytic as in Proposition 4.5. Definable functions are stable by composition as described in Lemma A.1, therefore the expression in (5.11) involves the integral of a globally subanalytic function, and by Lemma A.1 it is definable in $\mathbb{R}_{an,exp}$. We conclude using the stability of definable functions under maxima, Lemma 4.9. \square

Remark 5.7. *The result in Lemma 5.6 also holds if ϵ is considered as a variable, restricted to a compact interval of \mathbb{R}_+^* .*

6 Conclusion

The rigidity of o-minimal structures, and the log-analytic nature of globally subanalytic integrals allowed to prove sequential convergence of Lloyd-type algorithms under definability assumptions of the density of the target measure μ . A by-product of the analysis is the introduction of the powerful tools of o-minimal geometry in a semi-discrete optimal transport context. The main assumption on the target measure μ is that it stems from a globally subanalytic density. While the class of such densities is extremely broad, it is not clear how restrictive this assumption is and how it relates to practice. For example, it seems difficult to practically distinguish between absolutely continuous measures for which the global subanalyticity assumption on the density holds or not. Finally, exploring the nature of critical points resulting from Lloyd’s iterations is a relevant next step (global/local minima, saddle point) which would require a second order analysis.

Acknowledgments

The idea of using KL inequality for integral functionals based on definability was proposed by Jérôme Bolte along numerous discussions on the geometrical approach to optimization. This work benefited from financial support from the French government managed by the National Agency for Research under the France 2030 program, with the reference "ANR-23-PEIA-0004". Edouard Pauwels acknowledges the support of Institut Universitaire de France (IUF), the AI Interdisciplinary Institute ANITI funding, through the French "Investments for the Future – PIA3" program under the grant agreement ANR-19-PI3A0004, Air Force Office of Scientific Research, Air Force Material Command, USAF, under grant numbers FA8655-22-1-7012, ANR Chess (ANR-17-EURE-0010), ANR Regulia and ANR Bonsai.

A Additional proofs and technical results

The following lemma will be used throughout the text, see for example [12, Exercise 1.11].

Lemma A.1. *The composition of two functions definable in the same o-minimal structure is definable.*

Proof of Lemma 4.6. Let $\overline{B_{1,\infty}(\bar{x}_1, r)}, \dots, \overline{B_{l,\infty}(\bar{x}_l, r)}$ be a finite covering of K by l closed balls for the infinite norm with r sufficiently small so that for all $i = 1, \dots, l$, we have $\overline{B_{i,\infty}(\bar{x}_i, r)} \subset U$. Fix $i \in \{1, \dots, l\}$ and consider the translation and scaling transformation

$$A_i : x \in \mathbb{R}^n \mapsto \frac{1}{r}(x - \bar{x}_i) \in \mathbb{R}^n$$

so that $A_i \left(\overline{B_{i,\infty}(\bar{x}_i, r)} \right) = [-1, 1]^n$. These are affine bijections, so that $A_i(U)$ is an open set containing $[-1, 1]^n$. By composition of analytic functions, we get that $f \circ A_i^{-1} : A_i(U) \rightarrow \mathbb{R}$ is analytic. In particular, since $[-1, 1]^n \subset A_i(U)$ we have that the restriction $f \circ A_i^{-1}|_{[-1, 1]^n}$ is restricted analytic, hence globally subanalytic.

Finally, by denoting $B_i := \overline{B_{i,\infty}(\bar{x}_i, r)}$, we have that $f|_{B_i} = f \circ A_i^{-1}|_{[-1, 1]^n} \circ A_i$ where A_i is affine, hence definable (its graph is a subspace). Therefore $f|_{B_i}$ is also globally subanalytic by Lemma A.1. The

graph of $f|_K$ can then be written:

$$\Gamma(f|_K) = \Gamma\left(f|_{\bigcup_{i=1}^l B_i \cap K}\right) = \bigcup_{i=1}^l \Gamma(f|_{B_i \cap K}).$$

For all $i \in \{1, \dots, N\}$, we therefore have $f|_{B_i \cap K} = f|_{B_i} \mathbb{1}_K$, which is the product of two globally subanalytic functions as $\mathbb{1}_K$ is the indicator function of a semi algebraic set. Therefore $f|_{B_i \cap K}$ is globally subanalytic, for example using Lemma A.1 for stability under product, and we conclude using the stability of globally subanalytic sets with finite unions. □

Proof of Lemma 4.9. We prove the results (i) and (ii) for the minimum. The maximum case follows a symmetric proof using the hypograph.

First, we recall that by definition of the epigraph, $\text{epi}\left(\min_{j \in J} f_j\right) = \bigcup_{j \in J} \text{epi}(f_j)$. We thus conclude for (i) since definable subsets are stable by finite unions.

For the second property (ii), we denote by $p : \mathbb{R}^{d+n+1} \rightarrow \mathbb{R}^{d+1}$ the canonical projection on the first d coordinates and the last coordinate. We aim at showing the following equality

$$\text{epi}\left(\min_{w \in E} f(\cdot, w)\right) = p(\text{epi}(f)).$$

Since the minimum is attained by hypothesis, we have

$$\begin{aligned} (x, y) \in \text{epi}\left(\min_{w \in E} f(\cdot, w)\right) &\iff \min_{w \in E} f(x, w) \leq y \\ &\iff \exists w' \in E \text{ such that } f(x, w') \leq y \\ &\iff (x, w', y) \in \text{epi}(f) \\ &\iff (x, y) \in p(\text{epi}(f)) \end{aligned}$$

which concludes the proof. □

References

- [1] P-A. Absil, R. Mahony, and B. Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM J. Optim.*, 6(2):531–547, 2005.
- [2] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Łojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.
- [3] H. Attouch, J. Bolte, and B.F. Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. *Mathematical Programming*, 137(1):91–129, 2013.

- [4] M. Balzer, T. Schlömer, and O. Deussen. Capacity-constrained point distributions: A variant of Lloyd’s method. *ACM Transactions on Graphics (TOG)*, 28(3):1–8, 2009.
- [5] S. Bobkov and M. Ledoux. One-dimensional empirical measures, order statistics, and Kantorovich transport distances. *American Mathematical Society*, 261(1259), 2019.
- [6] V. Bogachev. *Measure Theory*, volume 1. Springer Berlin, Heidelberg, 2007.
- [7] J. Bolte, T. Le, and E. Pauwels. Subgradient sampling for nonsmooth nonconvex minimization. *SIAM Journal on Optimization*, 33(4):2542–2569, 2023.
- [8] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.
- [9] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 1(51):22–45, 2015.
- [10] R. Cluckers and D.J. Miller. Stability under integration of sums of products of real globally subanalytic functions and their logarithms. *Duke Mathematical Journal*, 156(2):311 – 348, 2011.
- [11] G. Comte, J-M Lion, and J-P Rolin. Nature log-analytique du volume des sous-analytiques. *Illinois Journal of Mathematics*, 44(4):884–888, 2000.
- [12] M. Coste. *An introduction to o-minimal geometry*. Istituti editoriali e poligrafici internazionali Pisa, 2000.
- [13] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [14] F. De Goes, K. Breeden, V. Ostromoukhov, and M. Desbrun. Blue noise through optimal transport. *ACM Transactions on Graphics (TOG)*, 31(6):1–11, 2012.
- [15] I. Deshpande, Y.T. Hu, R. Sun, A. Pyrros, N. Siddiqui, S. Koyejo, Z. Zhao, D. Forsyth, and A. Schwing. Max-Sliced Wasserstein distance and its use for GANs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10648–10656, 2019.
- [16] L. Dries and C. Miller. On the real exponential field with restricted analytic functions. *Israel Journal of Mathematics*, 92:427, 1995.
- [17] L. Dries and C. Miller. Geometric categories and o-minimal structures. *Duke Mathematical Journal*, 84(2):497 – 540, 1996.
- [18] Q. Du, M. Emelianenko, and L. Ju. Convergence of the Lloyd algorithm for computing centroidal Voronoï tessellations. *SIAM Journal on Numerical Analysis*, 44(1):102–119, 2006.
- [19] Q. Du, V. Faber, and M. Gunzburger. Centroidal Voronoï tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
- [20] M. Emelianenko, L. Ju, and A. Rand. Nondegeneracy and weak global convergence of the Lloyd algorithm in \mathbb{R}^d . *SIAM Journal on Numerical Analysis*, 46(3):1423–1441, 2008.

- [21] P.E. Fleischer. Sufficient conditions for achieving minimum distortion in a quantizer. *IEEE Int. Conv. Rec.*, 1, 1964.
- [22] A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of Sinkhorn divergences. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1574–1583. PMLR, 2019.
- [23] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [24] S. Graf and H. Luschgy. *Foundations of Quantization for Probability Distributions*, volume 1730. Springer Berlin, Heidelberg, 2000.
- [25] J. Kieffer. Exponential rate of convergence for Lloyd’s method I. *IEEE Transactions on Information Theory*, 28(2):205–210, 1982.
- [26] Q. Kitagawa, J. Méridot and B. Thibert. Convergence of a Newton algorithm for semi-discrete optimal transport. *Journal of the European Mathematical Society*, 21, 2019.
- [27] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. Rohde. Generalized Sliced Wasserstein distances. *Advances in neural information processing systems*, 32, 2019.
- [28] K. Kurdyka. On gradients of functions definable in o-minimal structures. *Annales de l’Institut Fourier*, 48(3):769–783, 1998.
- [29] J.-M. Lion and J.-P. Rolin. Intégration des fonctions sous-analytiques et volumes des sous-ensembles sous-analytiques. *Annales de l’Institut Fourier*, 48(3):755–767, 1998.
- [30] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [31] T.L. Loi. Lecture 1: O-minimal structures. In *The Japanese-Australian Workshop on Real and Complex Singularities: JARCS III*, volume 43, pages 19–31. Australian National University, Mathematical Sciences Institute, 2010.
- [32] S. Lojasiewicz. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.
- [33] Q. Méridot, F. Santambrogio, and C. Sarrazin. Non-asymptotic convergence bounds for Wasserstein approximation using point clouds. In *Advances in Neural Information Processing Systems*, volume 34, pages 12810–12821. Curran Associates, Inc., 2021.
- [34] G. Pagès. Introduction to vector quantization and its applications for numerics. *ESAIM: proceedings and surveys*, 48:29–79, 2015.
- [35] G. Pagès and J. Yu. Pointwise convergence of the Lloyd algorithm in higher dimension. *SIAM Journal on Control and Optimization*, 54(5):2354–2382, 2016.
- [36] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

- [37] J. Rabin, G. Peyré, J. Delon, and M. Bernot. Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision: Third International Conference, SSVM*, pages 435–446. Springer, 2012.
- [38] F. Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- [39] E. Tanguy. Convergence of SGD for training neural networks with Sliced Wasserstein losses. *Transactions on Machine Learning Research*, 2023.
- [40] L. Van den Dries. A generalization of the tarski-seidenberg theorem, and some nondefinability results. *Bulletin of the American Mathematical Society*, 15(2):189–193, 1986.
- [41] C. Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg, 2008.
- [42] C. Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.
- [43] X. Wu. On convergence of Lloyd’s method I. *IEEE Transactions on Information Theory*, 38(1):171–174, 1992.