Durham Research Online

Deposited in DRO:

19 January 2016

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Bourne, D.P. and Roper, S.M. (2015) 'Centroidal power diagrams, Lloyd's algorithm and applications to optimal location problems.', SIAM journal on numerical analysis., 53 (6). pp. 2545-2569.

Further information on publisher's website:

<http://dx.doi.org/10.1137/141000993>

Publisher's copyright statement:

© 2015, Society for Industrial and Applied Mathematics

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

CENTROIDAL POWER DIAGRAMS, LLOYD'S ALGORITHM, AND APPLICATIONS TO OPTIMAL LOCATION PROBLEMS*

D. P. BOURNE[†] AND S. M. ROPER[‡]

Abstract. In this paper we develop a numerical method for solving a class of optimization problems known as *optimal location* or *quantization* problems. The target energy can be written either in terms of atomic measures and the Wasserstein distance or in terms of weighted points and power diagrams (generalized Voronoi diagrams). The latter formulation is more suitable for computation. We show that critical points of the energy are *centroidal power diagrams*, which are generalizations of centroidal Voronoi tessellations, and that they can be approximated by a generalization of Lloyd's algorithm (Lloyd's algorithm is a common method for finding centroidal Voronoi tessellations). We prove that the algorithm is energy decreasing and prove a convergence theorem. Numerical experiments suggest that the algorithm converges linearly. We illustrate the algorithm in two and three dimensions using simple models of optimal location and crystallization (see online supplementary material).

Key words. Voronoi diagram, power diagram, Lloyd's method, optimal location problems, quantization

AMS subject classifications. 65K10, 65D99, 49M05

DOI. 10.1137/141000993

1. Introduction. In this paper we derive and analyze a numerical method for minimizing a class of energies that arises in economics (optimal location problems), electrical engineering (quantization), and materials science (crystallization and pattern formation). Applications are discussed further in section 1.5. These energies can be formulated either in terms of atomic measures and the Wasserstein distance (see (1.1)) or in terms of generalized Voronoi diagrams (see (1.9)). These formulations are equivalent, but (1.1) is more common in the applied analysis literature (e.g., [4], [8]) and (1.9) is more common in the computational geometry and quantization literature (e.g., [11], [13]). Importantly for us, formulation (1.9) is much more convenient for numerical work. We work with formulation (1.9) throughout the paper after first deriving it from (1.1) in sections 1.1 and 1.2. We start from (1.1) rather than directly from (1.9) in order to highlight the connection between the different communities.

1.1. Wasserstein formulation of the energy. Let Ω be a bounded subset of $\mathbb{R}^d$, $d \geq 2$, and $\rho : \Omega \rightarrow [0, \infty)$ be a given density on Ω . Let $f : [0, \infty) \rightarrow \mathbb{R}$. We consider the following class of discrete energies, which are defined on sets of weighted points $\{\mathbf{x}_i, m_i\}_{i=1}^N \in (\Omega \times (0, \infty))^N$, $\mathbf{x}_i \neq \mathbf{x}_j$ if $i \neq j$:

$$(1.1) \quad F(\{\mathbf{x}_i, m_i\}) = \sum_{i=1}^N f(m_i) + d^2 \left(\rho, \sum_{i=1}^N m_i \delta_{\mathbf{x}_i} \right).$$

The second term is the square of the Wasserstein distance between the density ρ and the atomic measure $\sum_{i=1}^N m_i \delta_{\mathbf{x}_i}$. It is defined below in (1.2). This energy models, e.g.,

*Received by the editors December 22, 2014; accepted for publication (in revised form) August 24, 2015; published electronically November 5, 2015.

<http://www.siam.org/journals/sinum/53-6/100099.html>

[†]Department of Mathematical Sciences, Durham University, Science Laboratories, Durham, DH1 3LE, UK (David.Bourne@durham.ac.uk).

[‡]School of Mathematics and Statistics, University of Glasgow, Glasgow, G12 8QW, UK (Steven.Roper@glasgow.ac.uk).

the problem of optimally locating resources (such as recycling points, polling stations, or distribution centers) in a city or country Ω with population density ρ . The points $\mathbf{x}_i$ are the locations of the resources, and the weights m_i represent their size. The first term of the energy penalizes the cost of building or running the resources. The second term penalizes the total distance between the population and the resources. In our case the Wasserstein distance $d(\cdot, \cdot)$ can be defined by

$$(1.2) \quad d^2 \left(\rho, \sum_{i=1}^N m_i \delta_{\mathbf{x}_i} \right) = \min_{T: \Omega \rightarrow \{\mathbf{x}_i\}_{i=1}^N} \left\{ \sum_{i=1}^N \int_{T^{-1}(\mathbf{x}_i)} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x} : \int_{T^{-1}(\mathbf{x}_i)} \rho d\mathbf{x} = m_i \forall i \right\}.$$

See, e.g., [30]. In two dimensions the minimization problem (1.2) can be interpreted as the following optimal partitioning problem: The map T partitions a city (for example), occupying Ω , with population density ρ into N regions, $\{T^{-1}(\mathbf{x}_i)\}_{i=1}^N$. Region $T^{-1}(\mathbf{x}_i)$ is assigned to the resource (e.g., polling station) located at point $\mathbf{x}_i$ of size m_i . The optimal map T does this in such a way as to minimize the total distance squared between the population and the resources subject to the constraint that each resource can meet the demand of the population assigned to it.

The Wasserstein distance is well defined provided that the weights m_i are positive and satisfy the mass constraint

$$(1.3) \quad \sum_i m_i = \int_{\Omega} \rho(\mathbf{x}) d\mathbf{x}.$$

It can be shown that $d(\cdot, \cdot)$ is a metric on measures and that it metrizes weak convergence of measures, meaning that if ρ_n converges to ρ , then $d(\rho, \rho_n) \rightarrow 0$. See, e.g., [30, Ch. 7]. It is not necessary to be familiar with measure theory or the Wasserstein distance since we will soon reformulate the minimization problem $\min F$ as a more elementary computational geometry problem involving generalized Voronoi diagrams (power diagrams).

The given data for the problem are Ω , f , ρ . We assume that

$$(1.4) \quad \Omega \text{ is compact and convex, } \rho \in C^0(\Omega), \quad \rho \geq 0,$$

$$(1.5) \quad f \in C^1([0, \infty)), \quad f \text{ is concave, } f(0) \geq 0, \quad \lim_{m \rightarrow 0} \frac{f(m)}{m} = +\infty.$$

Assumptions (1.5)₂, (1.5)₃ imply that

$$(1.6) \quad f \text{ is subadditive: } f(m_1) + f(m_2) \geq f(m_1 + m_2) \quad \forall m_1, m_2 \geq 0.$$

The number N of weighted points is *not* prescribed and is an unknown of the problem: The goal is to minimize F over the set of at most countably many weighted points $\{\mathbf{x}_i, m_i\}$ subject to the constraint (1.3).

If $f(0) > 0$, then it is easy to see that F has a minimizer and that minimizers consist of a finite number of weighted points (since minimizing F reduces to minimizing a continuous function on a compact subset of $\mathbb{R}^M$ for some sufficiently large M). If $f(0) = 0$, then the existence of a minimizer of F is more tricky and follows from (1.5)₄, (1.6), and a characterization of lower semicontinuous functionals of measures.

See [8, Thm. 2.1]. Since f is locally Lipschitz continuous, minimizers consist of a finite number of weighted points [8, Thm. 4.1].

The optimal number N of weighted points is determined by the competition between the two terms of F . The first term is minimized when $N = 1$, due to the subadditivity of f . The infimum of the second term is zero, which is obtained in the limit $N \rightarrow \infty$ (this is because the measure $\rho d\mathbf{x}$ can be approximated arbitrarily well with Dirac masses, e.g., by using a convergent quadrature rule, and because the Wasserstein distance $d(\cdot, \cdot)$ metrizes weak convergence of measures).

Energies of the form of F and generalizations have received a great deal of attention in the applied analysis literature; e.g., [4] and [8] study the existence and properties of minimizers for broad classes of optimal location energies. There is far less work, however, on numerical methods for such problems. Exceptions include the case of (1.1) with $f = 0$, which has been well studied numerically. This is discussed in section 1.4.

1.2. Power diagram formulation of the energy. Minimizing F numerically is challenging due to the presence of the Wasserstein term, which is defined implicitly in terms of the solution to the optimal transportation problem (1.2). This is an infinite-dimensional linear programming problem in which every point in Ω has to be assigned to one of the N weighted points $(\mathbf{x}_i, m_i)$. Therefore, even evaluating the energy F is expensive. One option is to discretize ρ so that (1.2) becomes a finite-dimensional linear programming problem. This is still costly, however, and it turns out that by exploiting a deep connection between optimal transportation theory and computational geometry we can reformulate the minimization problem $\min F$ in such a way that we can avoid solving (1.2) altogether.

First we need to introduce some terminology from computational geometry. The *power diagram* associated to a set of weighted points $\{\mathbf{x}_i, w_i\}_{i=1}^N$, where $\mathbf{x}_i \in \Omega$ are distinct, $w_i \in \mathbb{R}$, is the collection of subsets $P_i \subseteq \Omega$ defined by

$$(1.7) \quad P_i = \{\mathbf{x} \in \Omega : |\mathbf{x} - \mathbf{x}_i|^2 - w_i \leq |\mathbf{x} - \mathbf{x}_k|^2 - w_k \ \forall k\}.$$

The individual sets P_i are called *power cells* (or cells) of the power diagram. The power diagram is sometimes called the Laguerre diagram, or the radical Voronoi diagram. If all the weights w_i are equal, we obtain the standard Voronoi diagram; see Figure 1. From (1.7) we see that the power cells P_i are obtained by intersecting half planes and are therefore convex polytopes (or the intersection of convex polytopes with Ω in the case of cells that touch $\partial\Omega$): in dimension $d = 3$ the cells are convex polyhedra, and in dimension $d = 2$ the cells are convex polygons. Note that some of the cells may be empty. Comprehensive treatments of Voronoi diagrams include [3], [27].

Given weighted points $\{\mathbf{x}_i, m_i\}_{i=1}^N \in (\Omega \times (0, \infty))^N$, let T_* be the minimizer in (1.2). The optimal transport regions $\{T_*^{-1}(\mathbf{x}_i)\}_{i=1}^N$ form a power diagram. There exists $\{w_i\}_{i=1}^N \in \mathbb{R}^N$ such that the power diagram $\{P_i\}_{i=1}^N$ generated by $\{\mathbf{x}_i, w_i\}_{i=1}^N$ satisfies $P_i = T_*^{-1}(\mathbf{x}_i)$ for all i (up to sets of $\rho d\mathbf{x}$ -measure zero). Conversely, if $\{P_i\}_{i=1}^N$ is any power diagram with generators $\{\mathbf{x}_i, w_i\}_{i=1}^N$, then

$$(1.8) \quad d^2\left(\rho, \sum_{i=1}^N m_i \delta_{\mathbf{x}_i}\right) = \sum_{i=1}^N \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho d\mathbf{x}, \quad \text{where} \quad m_i = \int_{P_i} \rho d\mathbf{x}.$$

These results can be shown using Brenier's Theorem [30, Thm. 2.12] or the Kantorovich Duality Theorem [30, Thm. 1.3]. See [25, Thms. 1 and 2] or [5, Prop. 4.4].

FIG. 1. A comparison of a standard Voronoi diagram (left) with a power diagram (right). The location of the generators in both cases is the same, but the power diagram carries additional structure via the weights associated with each generator. The size of the weights in the power diagram is indicated by the radii of the dashed circles. Notice that in the power diagram it is possible for the generator to lie outside the cell or for the cell associated with a generator to be empty (the Voronoi diagram has 20 cells, and the power diagram has 19 cells). The geometrical construction of the power diagram in terms of the generator locations and the circles is simple; for each point $\mathbf{x}$ construct a tangent line from $\mathbf{x}$ to the circles centered at $\mathbf{x}_i$ with radii r_i ; the length of the tangent line is called the power of the point $\mathbf{x}$, and the point $\mathbf{x}$ belongs to the power cell that has minimum power. The weights of the generators in this case are $w_i = r_i^2$.

As far as we are aware, these results first appeared in [2], although they are not stated in the language of Wasserstein distances.

Equation (1.8) gives an explicit formula for the Wasserstein distance, without the need to solve a linear programming problem, provided that the weights m_i can be written as $\int_{P_i} \rho(\mathbf{x}) d\mathbf{x}$ for some power diagram $\{P_i\}$ (with generating points $\mathbf{x}_i$). In practice actually finding this power diagram involves solving another linear programming problem (the generating weights w_i come from the solution to the dual linear programming problem to (1.2); see [5, Prop. 4.4]), but in our case this can be avoided since we are interested in minimizing F rather than evaluating it at any given point.

We use this connection between the Wasserstein distance and power diagrams to rewrite the energy F in new variables, changing variables from $\{\mathbf{x}_i, m_i\}_{i=1}^N \in (\Omega \times (0, \infty))^N$ to $\{\mathbf{x}_i, w_i\}_{i=1}^N \in (\Omega \times \mathbb{R})^N$, with $\mathbf{x}_i$ distinct and w_i such that no cells are empty. By the results above, minimizing F is equivalent to minimizing

$$(1.9) \quad E(\{\mathbf{x}_i, w_i\}) = \sum_{i=1}^N \left\{ f(m_i) + \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x} \right\}$$

where $\{P_i\}$ is the power diagram generated by $\{\mathbf{x}_i, w_i\}$ and $m_i := \int_{P_i} \rho d\mathbf{x}$. The equivalence of E and F is in the following sense: Given $\{\mathbf{x}_i, w_i\}_{i=1}^N \in (\Omega \times \mathbb{R})^N$ and the corresponding power diagram $\{P_i\}_{i=1}^N$, (1.8) implies that

$$E(\{\mathbf{x}_i, w_i\}) = F(\{\mathbf{x}_i, m_i\}) \quad \text{for} \quad m_i = \int_{P_i} \rho(\mathbf{x}) d\mathbf{x}.$$

Conversely, it can be shown (e.g., [5, Prop. 4.4]) that given any $\{\mathbf{x}_i, m_i\}_{i=1}^N \in (\Omega \times (0, \infty))^N$, there exists $\{w_i\}_{i=1}^N \in \mathbb{R}^N$ such that the power diagram $\{P_i\}_{i=1}^N$ generated by $\{\mathbf{x}_i, w_i\}_{i=1}^N$ satisfies $\int_{P_i} \rho(\mathbf{x}) d\mathbf{x} = m_i$ for all i . Then it follows from (1.8) that $F(\{\mathbf{x}_i, m_i\}) = E(\{\mathbf{x}_i, w_i\})$. The weights $\{w_i\}_{i=1}^N \in \mathbb{R}^N$ are unique up to the addition of a constant; it is easy to see from (1.7) that $\{w_i + c\}_{i=1}^N$ and $\{w_i\}_{i=1}^N$ generate the same power diagram.

While the energies E and F are equivalent, from a numerical point of view it is far more practical to work with E since it can be easily evaluated, unlike F , since computing power diagrams is easy while solving the linear programming problem (1.2) is not. In the rest of the paper we focus on finding local minimizers of E .

1.3. Centroidal power diagrams and a generalized Lloyd algorithm.

From now on we will write $(\mathbf{X}, \mathbf{w}) = ((\mathbf{x}_1, \dots, \mathbf{x}_N), (w_1, \dots, w_N)) \in \Omega^N \times \mathbb{R}^N$ to denote the generators of a power diagram. In this section we introduce an algorithm for finding critical points of $E = E(\mathbf{X}, \mathbf{w})$.

Let $\mathcal{G}^N \subset \Omega^N \times \mathbb{R}^N$ be the smaller class of generators such that no two generators coincide and there are no empty cells:

$$(1.10) \quad \mathcal{G}^N = \{(\mathbf{X}, \mathbf{w}) \in \Omega^N \times \mathbb{R}^N : \mathbf{x}_i \neq \mathbf{x}_j \text{ if } i \neq j, P_i \neq \emptyset \forall i\}.$$

Define $\boldsymbol{\xi} : \mathcal{G}^N \rightarrow \Omega^N$ and $\boldsymbol{\omega} : \mathcal{G}^N \rightarrow \mathbb{R}^N$ by

$$(1.11) \quad \boldsymbol{\xi}(\mathbf{X}, \mathbf{w}) := (\xi_1(\mathbf{X}, \mathbf{w}), \dots, \xi_N(\mathbf{X}, \mathbf{w})), \quad \boldsymbol{\omega}(\mathbf{X}, \mathbf{w}) := (\omega_1(\mathbf{X}, \mathbf{w}), \dots, \omega_N(\mathbf{X}, \mathbf{w})),$$

where

$$(1.12) \quad \xi_i(\mathbf{X}, \mathbf{w}) := \frac{1}{m_i(\mathbf{X}, \mathbf{w})} \int_{P_i(\mathbf{X}, \mathbf{w})} \mathbf{x} \rho(\mathbf{x}) d\mathbf{x}, \quad \omega_i(\mathbf{X}, \mathbf{w}) := -f'(m_i(\mathbf{X}, \mathbf{w})).$$

Here $P_i(\mathbf{X}, \mathbf{w})$ is the i th power cell in the power diagram generated by $(\mathbf{X}, \mathbf{w})$, and $m_i(\mathbf{X}, \mathbf{w})$ is its mass:

$$m_i(\mathbf{X}, \mathbf{w}) = \int_{P_i(\mathbf{X}, \mathbf{w})} \rho(\mathbf{x}) d\mathbf{x}.$$

Note that $\xi_i(\mathbf{X}, \mathbf{w})$ is the centroid (or center of mass) of the i th power cell. We will sometimes denote this by $\bar{\mathbf{x}}_i$. In section 2 we show that critical points of E are fixed points of the Lloyd maps:

$$\nabla E(\mathbf{X}, \mathbf{w}) = \mathbf{0} \quad \Longleftrightarrow \quad (\boldsymbol{\xi}(\mathbf{X}, \mathbf{w}), \boldsymbol{\omega}(\mathbf{X}, \mathbf{w})) = (\mathbf{X}, \mathbf{w})$$

(up to the addition of a constant vector to $\mathbf{w}$; see Proposition 2.6 for a precise statement). The condition $\boldsymbol{\xi}(\mathbf{X}, \mathbf{w}) = \mathbf{X}$ means that the power diagram generated by $(\mathbf{X}, \mathbf{w})$ has the property that $\mathbf{x}_i$ is the centroid of its power cell P_i for all i . These special types of power diagrams are called *centroidal power diagrams* (CPDs). This is in analogy with centroidal Voronoi tessellations (CVTs), which are special types of Voronoi diagrams with the property that the generators of the Voronoi diagram are the centroids of the Voronoi cells. See [11] for a nice survey of CVTs. Note also that CVTs can be viewed as a special type of CPD where all the weights are equal; $w_i = c$ for all i , $c \in \mathbb{R}$, since power diagrams with equal weights are just Voronoi diagrams. Other examples of CPDs appear in section 2.4.

The following algorithm is an iterative method for finding fixed points of $(\boldsymbol{\xi}, \boldsymbol{\omega})$ and therefore critical points of E .

Algorithm 1. The generalized Lloyd algorithm for finding critical points of E .

Initialization: Choose $N_0 \in \mathbb{N}$ and $(\mathbf{X}^0, \mathbf{w}^0) \in \mathcal{G}^{N_0}$.

At each iteration:

- (1) Update the generators: Given $(\mathbf{X}^k, \mathbf{w}^k) \in \mathcal{G}^{N_k}$, compute the corresponding power diagram and define $(\mathbf{X}^{k+1}, \mathbf{w}^{k+1}) \in \Omega^{N_k} \times \mathbb{R}^{N_k}$ by

$$\mathbf{X}^{k+1} = \boldsymbol{\xi}(\mathbf{X}^k, \mathbf{w}^k), \quad \mathbf{w}^{k+1} = \boldsymbol{\omega}(\mathbf{X}^k, \mathbf{w}^k).$$

- (2) Remove empty cells: Compute the power diagram $\{P_i^{k+1}\}_{i=1}^{N_k}$ generated by $(\mathbf{X}^{k+1}, \mathbf{w}^{k+1})$ and let

$$J = \{j \in \{1, \dots, N_k\} : P_j^{k+1} = \emptyset\}.$$

For all $j \in J$, remove $(\mathbf{x}_j^{k+1}, w_j^{k+1})$ from the list of generators. Then replace N_k with $N_{k+1} = N_k - |J|$.

In particular, this algorithm computes CPDs, and it is a generalization of Lloyd's algorithm [22], which is a popular method for computing CVTs. See [11]. The classical Lloyd algorithm is recovered from our generalized Lloyd algorithm by simply taking the weights to be constant at each iteration, e.g., $\mathbf{w}^k = \mathbf{0}$ for all k . Due to this relation, we refer to $\boldsymbol{\xi}$ and $\boldsymbol{\omega}$ as generalized Lloyd maps.

Somewhat imprecisely, the role of the update $\mathbf{x}_i^{k+1} = \boldsymbol{\xi}_i(\mathbf{X}^k, \mathbf{w}^k)$ in step (1) of the algorithm can be thought of as being to decrease the second term of the energy (while possibly increasing the first term). To be precise,

$$\int_{P_i^k} |\mathbf{x} - \mathbf{x}_i^{k+1}|^2 \rho \, d\mathbf{x} \leq \int_{P_i^k} |\mathbf{x} - \mathbf{x}_i^k|^2 \rho \, d\mathbf{x}.$$

Note that both integrals are over P_i^k . The role of the update $w_i^{k+1} = -f'(m_i(\mathbf{X}^k, \mathbf{w}^k))$ can be thought of as being to decrease the first term of the energy (while possibly increasing the second term). Since f is concave, then f' is nonincreasing and $w_i^{k+1} \geq w_j^{k+1}$ if $m_i^k \geq m_j^k$. This suggests that the weight update transfers mass to larger cells from smaller neighboring cells, which would decrease the first term of the energy since f is concave. This reasoning is not completely correct since both updates take place simultaneously and the size of cells depends in a complicated way on both the weights and locations of all the generators. Nevertheless, the algorithm is in fact energy decreasing; see Theorem 3.1.

Step (2) of the algorithm means that, given $N_0 \in \mathbb{N}$ and $(\mathbf{X}^0, \mathbf{w}^0) \in \mathcal{G}^{N_0}$, the algorithm can converge to a fixed point $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$ with $N < N_0$. This means that the algorithm can partly correct for an incorrect initial guess N_0 (recall that we are minimizing $E(\mathbf{X}, \mathbf{w})$ over $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$ and over N). It is still possible, however, that the algorithm converges to a local minimizer of E , possibly with a nonoptimal value of N . Note also that the algorithm can eliminate generators, but it cannot create them. Therefore, it is impossible for the algorithm to find a global minimizer of E if the initial value of N_0 is less than the optimal value.

Algorithm 1 was introduced for the special case of $d = 2$, $\rho = 1$, $f(m) = \sqrt{m}$ in [5, sec. 4]. In the current paper we extend it to the broader class of energies (1.9), analyze it (prove that it is energy decreasing and that it converges; see Theorems 3.1 and 3.3), and implement it in both two and three dimensions. In addition, the derivation here, unlike in [5], is accessible to those not familiar with measure theory

and optimal transport theory since we work with formulation (1.9) rather than (1.1).

1.4. The case $f = 0$ and N fixed: CVTs and Lloyd's algorithm. Setting $f = 0$ in (1.1) and fixing N gives the energy

$$F_N(\{\mathbf{x}_i, m_i\}) = d^2 \left(\rho, \sum_{i=1}^N m_i \delta_{\mathbf{x}_i} \right).$$

It is necessary to fix N since otherwise this has no minimizer; the infimum is zero, which is obtained in the limit $N \rightarrow \infty$ by approximating ρ with Dirac masses. It can be shown that minimizing F_N is equivalent to minimizing

$$E_N(\{\mathbf{x}_i\}) = \sum_{i=1}^N \int_{V_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x},$$

where $\{V_i\}_{i=1}^N$ is the Voronoi diagram generated by $\{\mathbf{x}_i\}_{i=1}^N$:

$$V_i = \{\mathbf{x} \in \Omega : |\mathbf{x} - \mathbf{x}_i| \leq |\mathbf{x} - \mathbf{x}_k| \forall k\}.$$

See [5, sec. 4.1]. Numerical minimization of E_N has been well studied. A necessary condition for minimality is that $\{\mathbf{x}_i\}_{i=1}^N$ generates a CVT. CVTs can be easily computed using the classical Lloyd algorithm. See, e.g., [11]. Convergence of the algorithm is studied in [10], [11], and [28], among others, and there is a large literature on CVTs and Lloyd's algorithm. However, we are not aware of any work (other than [5]) on numerical minimization of E for $f \neq 0$.

1.5. Applications. Energies of the form (1.9), or equivalently (1.1), arise in many applications.

1.5.1. Simple model of pattern formation: Block copolymers. The authors first came in contact with energies of the form (1.1) in a pattern formation problem in materials science [5]. The following energy is a simplified model of phase separation for two-phase materials called block copolymers, for the case where one phase has a much smaller volume fraction than the other:

$$(1.13) \quad E(\{\mathbf{x}_i, w_i\}) = \sum_{i=1}^N \left\{ \lambda m_i^{\frac{d-1}{d}} + \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 d\mathbf{x} \right\},$$

where $m_i = \int_{P_i} 1 d\mathbf{x} = |P_i|$ and $d = 2$ or 3 . The measure $\nu = \sum_i m_i \delta_{\mathbf{x}_i}$ represents the minority phase. In three dimensions, $d = 3$, this represents N small spheres of the minority phase centered at $\{\mathbf{x}_i\}_{i=1}^N$. The weights m_i give the relative size of the spheres. These spheres are surrounded by a "sea" of the majority phase. In two dimensions, $d = 2$, the measure ν represents N parallel cylinders of the minority phases and Ω is a cross-section perpendicular to the axes of the cylinders. The first term of E penalizes the surface area between the two phases and so prefers phase separation ($N = 1$), and the second term prefers phase mixing ($N = \infty$). The parameter λ represents the repulsion strength between the two phases. Equation (1.13) is the special case of (1.9) with $\rho = 1$ and $f(m) = \lambda m^{\frac{d-1}{d}}$.

This energy can be viewed as a toy model of the popular Ohta-Kawasaki model of block copolymers (see, e.g., [9]). Like the Ohta-Kawasaki energy, it is nonconvex and nonlocal (in the sense that evaluating E involves solving an auxiliary infinite-dimensional problem). Unlike the Ohta-Kawasaki energy, however, it is discrete,

which makes it much more amenable to numerics and analysis. In general it can be viewed as a simplified model of nonconvex, nonlocal, energy-driven pattern formation, and it has applications in materials science outside block copolymers, e.g., to crystallization. It is also connected to the Ginzburg–Landau model of superconductivity [6, p. 123–124].

In [5] it was demonstrated numerically that for $d = 2$ minimizers of E tend to a hexagonal tiling as $\lambda \rightarrow 0$ (in the sense that the power diagram generated by $\{\mathbf{x}_i, w_i\}$ tends to a hexagonal tiling). This was proved in [6], and it agrees with block copolymer experiments, where in some parameter regime the minority phase forms hexagonally packed cylinders. It was conjectured in [5] that for the case $d = 3$, minimizers of E tend to a body-centered cubic (BCC) lattice as $\lambda \rightarrow 0$ (meaning that $\{\mathbf{x}_i\}$ tend to a BCC lattice and $w_i \rightarrow 0$). A brief examination of this conjecture can be found in the online supplementary material. In particular, numerical minimization of E in three dimensions suggests that the BCC lattice is at least a local minimizer of E when Ω is a periodic box. Again, this agrees with block copolymer experiments, where in some parameter regime the minority phase forms a BCC lattice.

1.5.2. Quantization. Energies of the form (1.9) can be used for data compression using a technique called *vector quantization*. By taking $f = 0$ in (1.9) and evaluating the resulting energy at $w_i = 0$ for all i , so that the power diagram $\{P_i\}_{i=1}^N$ generated by $\{\mathbf{x}_i, 0\}_{i=1}^N$ is just the Voronoi diagram $\{V_i\}_{i=1}^N$ generated by $\{\mathbf{x}_i\}_{i=1}^N$, we obtain the energy

$$(1.14) \quad D(\{\mathbf{x}_i\}) = \sum_{i=1}^N \int_{V_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x} \equiv \int_{\Omega} \min_i |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x}.$$

This is known in the quantization literature as the *distortion*. See [15, sec. 33] for a mathematical introduction to vector quantization and [13] and [14] for comprehensive treatments. Roughly speaking, the points $\mathbf{x}$ of Ω represent signals (e.g., parts of an image or speech) and $\mathbf{x}_i$ represent codewords in the codebook $\{\mathbf{x}_i\}_{i=1}^N$. The function ρ is a probability density on the set of signals Ω . If a signal $\mathbf{x}$ belongs to the Voronoi cell V_i , then the encoder assigns to it the codeword $\mathbf{x}_i$, which is then stored or transmitted. D measures the quality of the encoder, the average distortion of signals. The minimum value of D is called the *minimum distortion*.

In practice distortion is minimized subject to a constraint on the number of bits in the codebook. The codewords $\mathbf{x}_i$ are mapped to binary vectors before storage or transmission. In *fixed-rate* quantization all these vectors have the same length. In *variable-rate quantization* the length depends on the probability density ρ : Let $m_i = \int_{V_i} \rho d\mathbf{x}$ be the probability that a signal lies in Voronoi cell V_i . If m_i is large, then $\mathbf{x}_i$ should be mapped to a short binary vector since it occurs often. For cells with lower probabilities, longer binary vectors can be used. The *rate* of an encoder has the form

$$R = \sum_{i=1}^N l_i m_i,$$

where l_i is the length of the binary vector representing $\mathbf{x}_i$. Note that R is the expected value of the length. Distortion D is decreased by choosing more codewords. On the other hand, this means that the rate R , and hence the storage/transmission cost, is increased. Optimal encoders can be designed by trading off distortion against rate

by minimizing energies of the form $\lambda R + D$, where λ is a parameter determining the tradeoff. See [14, p. 2342]. Our energy (1.9) generalizes this: Take $l_i = l(1/m_i)$ for some concave function l so that $m \mapsto l(1/m)m$ is concave. In addition, l should be increasing so that the code length decreases as the probability m increases. We replace the Voronoi cells in (1.14) with power cells, which means that signals in power cell P_i are mapped to codeword $\mathbf{x}_i$. Then the energy $\lambda R + D$ has the form of (1.9):

$$E(\{\mathbf{x}_i, w_i\}) = \sum_{i=1}^N \left\{ f(m_i) + \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x} \right\}, \quad \text{where} \quad f(m) = \lambda l\left(\frac{1}{m}\right) m.$$

1.5.3. Optimal location of resources. As discussed in section 1.1, energies of the form (1.1) and (1.9) can be used to model the optimal location of resources $\{\mathbf{x}_i\}$ in a city or country Ω with population density ρ . The resources have size m_i , serve region P_i , and cost $f(m_i)$ to build or run. The assumption that f is concave, introduced for mathematical convenience to prove Theorem 3.1, is also natural from the modeling point of view; along with the assumption $f(0) \geq 0$ it implies that f is subadditive, which corresponds to an economy of scale. The energy trades off building/running costs against distance between the population and the resources.

1.5.4. Other applications and connections. Energies of the form (1.9), usually with $f = 0$, also arise in data clustering and pattern recognition (k -means clustering) [16], [24], image compression (this is a special case of vector quantization) [11, sec. 2.1], numerical integration [11, sec. 2.2], [15, p. 497–499], and convex geometry (packing and covering problems; approximation of convex bodies by convex polytopes) [15, sec. 33]. Taking $f \neq 0$ in (1.9) gives the algorithm more freedom, e.g., to automatically select the number of data clusters in addition to their location, based on a cost per cluster. Energies involving the Wasserstein distance also arise from the time-discrete gradient flow formulation of PDEs [18].

Recently there has been a resurgence of interest in centroidal Voronoi and power diagrams and Lloyd's algorithm [7], [20], [29]. Voronoi diagrams have also gained a lot of interest in the materials science community, e.g., to model solid foams [1] and grains in metals [19], although this is usually done in a manner more heuristic than energy minimization. Global minimizers of E can be difficult to find if they have a large value of N , and the generalized Lloyd algorithm tends to converge to local minimizers. These often resemble grains in metals (see Figure 5), which suggests that energy minimization might be a good method to produce Representative Volume Elements for the finite element simulation of materials with microstructure.

1.6. Structure of the paper. The generalized Lloyd algorithm, Algorithm 1, is derived in section 2. In section 3 we prove that it is energy decreasing, prove a convergence theorem, and study its structure. Implementation issues, such as how to compute power diagrams, are discussed in the online supplementary material. Numerical illustrations in two and three dimensions are given in section 4, with further illustrations in the online supplementary material.

2. Derivation of the algorithm. In this section we derive the generalized Lloyd algorithm, Algorithm 1, which is a fixed point method for the calculation of stationary points of the energy E , defined in (1.9). Calculating the gradient of E requires care since this involves differentiating the integrals appearing in the definition of E with respect to their domains. We perform this calculation in sections 2.2 and 2.3 after introducing some notation in section 2.1.

2.1. Notation for power diagrams. Throughout this paper we take Ω to be a compact, convex subset of $\mathbb{R}^d$, $d \geq 2$. We will take $d = 2$ or 3 for purposes of illustration, but the theory developed applies for all $d \geq 2$.

Given weighted points $(\mathbf{X}, \mathbf{w}) = ((\mathbf{x}_1, \dots, \mathbf{x}_N), (w_1, \dots, w_N)) \in \mathcal{G}^N$ and the associated power diagram $\{P_i\}_{i=1}^N$, we introduce the following notation:

$$(2.1) \quad d_{ij} = |\mathbf{x}_j - \mathbf{x}_i|, \quad \mathbf{n}_{ij} = \frac{\mathbf{x}_j - \mathbf{x}_i}{d_{ij}}, \quad F_{ij} = P_i \cap P_j,$$

$$(2.2) \quad m_i = \int_{P_i} \rho(\mathbf{x}) d\mathbf{x}, \quad m_{ij} = \int_{F_{ij}} \rho(\mathbf{x}) d\mathbf{x},$$

$$(2.3) \quad \bar{\mathbf{x}}_i = \frac{1}{m_i} \int_{P_i} \mathbf{x} \rho(\mathbf{x}) d\mathbf{x}, \quad \bar{\mathbf{x}}_{ij} = \frac{1}{m_{ij}} \int_{F_{ij}} \mathbf{x} \rho(\mathbf{x}) d\mathbf{x},$$

$$(2.4) \quad J_i = \{j \neq i : P_i \cap P_j \neq \emptyset\}.$$

Here d_{ij} is the distance between points $\mathbf{x}_i$ and $\mathbf{x}_j$; $\mathbf{n}_{ij}$ is the unit vector pointing from $\mathbf{x}_i$ to $\mathbf{x}_j$; the set F_{ij} is the *face* common to both cells P_i and P_j ; m_i is the mass of cell P_i ; m_{ij} is the mass of face F_{ij} ; $\bar{\mathbf{x}}_i$ is the *center of mass* of the cell P_i ; and $\bar{\mathbf{x}}_{ij}$ is the center of mass of face F_{ij} . The set of indices of the neighbors of cell P_i is given by the index set J_i . In the case $d = 2$ the power cells are convex polygons, and, rather than referring to the intersections of neighboring cells as faces, we refer to them as edges.

Recall that we sometimes write $P_i(\mathbf{X}, \mathbf{w})$ for the power cells generated by $(\mathbf{X}, \mathbf{w})$, instead of simply P_i , to emphasize that the power diagram is generated by $(\mathbf{X}, \mathbf{w})$. Similarly, we will sometimes write $m_i(\mathbf{X}, \mathbf{w})$ for the mass of the i th power cell. From (1.7) it is easy to see that adding a constant $c \in \mathbb{R}$ to all the weights generates the same power diagram: $P_i(\mathbf{X}, \mathbf{w} + \mathbf{c}) = P_i(\mathbf{X}, \mathbf{w})$ for all i , where $\mathbf{c} = (c, \dots, c) \in \mathbb{R}^N$. Let $\mathbb{R}_+ = [0, \infty)$, and let $\mathbf{m} : \Omega^N \times \mathbb{R}^N \rightarrow \mathbb{R}_+^N$ be the function defined by

$$(2.5) \quad \mathbf{m}(\mathbf{X}, \mathbf{w}) = (m_1(\mathbf{X}, \mathbf{w}), \dots, m_N(\mathbf{X}, \mathbf{w})),$$

which gives the mass of all of the cells generated by $(\mathbf{X}, \mathbf{w})$. Note that some of the cells may be empty (at most $N - 1$ of them), in which case the corresponding components of $\mathbf{m}$ take the value zero. Given a density $\rho : \Omega \rightarrow [0, \infty)$, let the space of admissible masses be

$$(2.6) \quad \mathcal{M}^N = \left\{ \mathbf{M} \in \mathbb{R}_+^N : \sum_{i=1}^N M_i = \int_{\Omega} \rho(\mathbf{x}) d\mathbf{x} \right\}.$$

Throughout this paper $\mathbf{I}_m$ denotes the m -by- m identity matrix.

2.2. The helper function H . Motivated by [10], where convergence of the classical Lloyd algorithm is studied, we introduce a helper function H defined by

$$(2.7) \quad H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M}) \\ := \sum_{i=1}^N \left\{ M_i w_i^1 + f(M_i) + \int_{P_i(\mathbf{X}^2, \mathbf{w}^2)} (|\mathbf{x} - \mathbf{x}_i^1|^2 - w_i^1) \rho(\mathbf{x}) d\mathbf{x} \right\},$$

where $(\mathbf{X}^k, \mathbf{w}^k) = ((\mathbf{x}_1^k, \dots, \mathbf{x}_N^k), (w_1^k, \dots, w_N^k))$ for $k \in \{1, 2\}$, $\mathbf{M} = (M_1, \dots, M_N)$, and the domain of H is $(\Omega^N \times \mathbb{R}^N) \times (\Omega^N \times \mathbb{R}^N) \times \mathbb{R}_+^N$. The energy E is recovered by choosing the arguments of H appropriately:

$$(2.8) \quad E(\mathbf{X}, \mathbf{w}) = H((\mathbf{X}, \mathbf{w}), (\mathbf{X}, \mathbf{w}), \mathbf{m}(\mathbf{X}, \mathbf{w})).$$

Note that H is invariant under addition of a constant to all the weights:

$$(2.9) \quad H((\mathbf{X}^1, \mathbf{w}^1 + \mathbf{c}_1), (\mathbf{X}^2, \mathbf{w}^2 + \mathbf{c}_2), \mathbf{M}) = H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M})$$

for all $\mathbf{c}_i = c_i(1, \dots, 1) \in \mathbb{R}^N$, $i \in \{1, 2\}$, provided that $\mathbf{M} \in \mathcal{M}^N$ and the N points $\mathbf{x}_1^2, \dots, \mathbf{x}_N^2$ are distinct. The following lemma will be used to prove that the generalized Lloyd algorithm is energy decreasing.

LEMMA 2.1 (properties of H). *Let $\boldsymbol{\xi}$, $\boldsymbol{\omega}$ be the Lloyd maps defined in (1.11). Then*

- (i) $\min_{\mathbf{X}^1 \in \Omega^N} H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M}) = H((\boldsymbol{\xi}(\mathbf{X}^2, \mathbf{w}^2), \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M})$,
- (ii) $H((\mathbf{X}, \mathbf{w}^1), (\mathbf{X}, \mathbf{w}), \mathbf{m}(\mathbf{X}, \mathbf{w})) = E(\mathbf{X}, \mathbf{w})$, *i.e., is independent of $\mathbf{w}^1$,*
- (iii) $H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M}) \geq H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^1, \mathbf{w}^1), \mathbf{M})$,
with equality if and only if $P_i(\mathbf{X}^1, \mathbf{w}^1) = P_i(\mathbf{X}^2, \mathbf{w}^2)$ for all i .
- (iv) *If f is concave, then*

$$\begin{aligned} \max_{\mathbf{M} \in \mathbb{R}_+^N} H((\mathbf{X}^1, \boldsymbol{\omega}(\mathbf{X}^2, \mathbf{w}^2)), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M}) \\ = H((\mathbf{X}^1, \boldsymbol{\omega}(\mathbf{X}^2, \mathbf{w}^2)), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{m}(\mathbf{X}^2, \mathbf{w}^2)). \end{aligned}$$

Proof. Property (i): For fixed $\mathbf{X}^2 \in \Omega^N$, $\mathbf{w}^1, \mathbf{w}^2 \in \mathbb{R}^N$, and $\mathbf{M} \in \mathcal{M}^N$, define the function $h: \Omega^N \rightarrow \mathbb{R}$ by $h(\mathbf{X}^1) := H((\mathbf{X}^1, \mathbf{w}^1), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M})$. Then

$$\frac{\partial h}{\partial \mathbf{x}_i^1}(\mathbf{X}^1) = 2 \int_{P_i(\mathbf{X}^2, \mathbf{w}^2)} (\mathbf{x}_i^1 - \mathbf{x}) \rho(\mathbf{x}) d\mathbf{x} = 2m_i(\mathbf{X}^2, \mathbf{w}^2)(\mathbf{x}_i^1 - \boldsymbol{\xi}_i(\mathbf{X}^2, \mathbf{w}^2))$$

by the definition (1.11) of $\boldsymbol{\xi}_i$. Therefore, $\boldsymbol{\xi}(\mathbf{X}^2, \mathbf{w}^2)$ is a critical point of h . Moreover, it is a global minimum point since h is convex:

$$\frac{\partial^2 h}{\partial \mathbf{x}_i^1 \partial \mathbf{x}_j^1} = \begin{cases} 2m_i(\mathbf{X}^2, \mathbf{w}^2) \mathbf{I}_d & \text{if } i = j, \\ \mathbf{0} & \text{if } i \neq j, \end{cases}$$

where $\mathbf{I}_d$ and $\mathbf{0}$ are the d -by- d identity and zero matrices. (Note that h is not necessarily strictly convex since $m_i(\mathbf{X}^2, \mathbf{w}^2)$ may be zero for some i , which is the case when the power cell $P_i(\mathbf{X}^2, \mathbf{w}^2)$ is empty.)

Property (ii) is immediate from the definitions of H and E .

Property (iii): This follows from the fact that for any partition $\{S_i\}_{i=1}^N$ of Ω we have

$$\sum_i \int_{S_i} (|\mathbf{x} - \mathbf{x}_i^1|^2 - w_i^1) \rho(\mathbf{x}) d\mathbf{x} \geq \sum_i \int_{P_i(\mathbf{X}^1, \mathbf{w}^1)} (|\mathbf{x} - \mathbf{x}_i^1|^2 - w_i^1) \rho(\mathbf{x}) d\mathbf{x}$$

with equality if and only if $\{S_i\}_{i=1}^N$ is the power diagram generated by $(\mathbf{X}^1, \mathbf{w}^1)$ (up to sets of $\rho d\mathbf{x}$ -measure zero). This follows since

$$\sum_i \int_{P_i(\mathbf{X}^1, \mathbf{w}^1)} (|\mathbf{x} - \mathbf{x}_i^1|^2 - w_i^1) \rho(\mathbf{x}) d\mathbf{x} = \int_{\Omega} \min_i \{|\mathbf{x} - \mathbf{x}_i^1|^2 - w_i^1\} \rho(\mathbf{x}) d\mathbf{x}.$$

Property (iv): Define $g(\mathbf{M}) = H((\mathbf{X}^1, \boldsymbol{\omega}(\mathbf{X}^2, \mathbf{w}^2)), (\mathbf{X}^2, \mathbf{w}^2), \mathbf{M})$. Then

$$\begin{aligned} g(\mathbf{m}(\mathbf{X}^2, \mathbf{w}^2)) - g(\mathbf{M}) \\ = \sum_i \{f(m_i(\mathbf{X}^2, \mathbf{w}^2)) + f'(m_i(\mathbf{X}^2, \mathbf{w}^2))(M_i - m_i(\mathbf{X}^2, \mathbf{w}^2)) - f(M_i)\} \geq 0 \end{aligned}$$

since f is concave, as required. $\square$

Remark 2.2 (relation between F and H). By results from [3, pp. 98–99], the energy F introduced in (1.1) is related to H by

$$F(\{\mathbf{x}_i, m_i\}) = \max_{\{w_i\}} H(\{\mathbf{x}_i, w_i\}, \{\mathbf{x}_i, w_i\}, \{m_i\}).$$

Therefore, the problem of minimizing F is equivalent to solving a saddle point problem for the function $\{\mathbf{x}_i, m_i, w_i\} \mapsto H(\{\mathbf{x}_i, w_i\}, \{\mathbf{x}_i, w_i\}, \{m_i\})$. If $\{\mathbf{x}_i, w_i\}$ is a fixed point of the Lloyd map ω , then Lemma 2.1(iv) implies that

$$E(\{\mathbf{x}_i, w_i\}) = \max_{\{m_i\}} H(\{\mathbf{x}_i, w_i\}, \{\mathbf{x}_i, w_i\}, \{m_i\}).$$

2.3. Critical points of E . In this section we show that critical points of E are fixed points of the Lloyd maps ξ, ω .

LEMMA 2.3 (partial derivatives of E). *The partial derivatives of E are*

$$(2.10) \quad \frac{\partial E}{\partial \mathbf{x}_i}(\mathbf{X}, \mathbf{w}) = 2m_i(\mathbf{x}_i - \xi_i(\mathbf{X}, \mathbf{w})) + \sum_{j=1}^N \frac{\partial m_j}{\partial \mathbf{x}_i}(w_j - \omega_j(\mathbf{X}, \mathbf{w})),$$

$$(2.11) \quad \frac{\partial E}{\partial w_i}(\mathbf{X}, \mathbf{w}) = \sum_{j=1}^N \frac{\partial m_j}{\partial w_i}(w_j - \omega_j(\mathbf{X}, \mathbf{w}))$$

for $i \in \{1, \dots, N\}$. In matrix notation,

$$(2.12) \quad \begin{pmatrix} \nabla_{\mathbf{X}} E \\ \nabla_{\mathbf{w}} E \end{pmatrix} = \begin{pmatrix} 2\hat{M} & \nabla_{\mathbf{X}} \mathbf{m} \\ \mathbf{0} & \nabla_{\mathbf{w}} \mathbf{m} \end{pmatrix} \begin{pmatrix} \mathbf{X} - \xi(\mathbf{X}, \mathbf{w}) \\ \mathbf{w} - \omega(\mathbf{X}, \mathbf{w}) \end{pmatrix},$$

where

$$(2.13) \quad \hat{M} := \text{diag}(m_1, \dots, m_N) \otimes \mathbf{I}_d = \text{diag}(m_1 \mathbf{I}_d, \dots, m_N \mathbf{I}_d).$$

Proof. From (2.8),

$$(2.14) \quad \frac{\partial E}{\partial \mathbf{x}_i}(\mathbf{X}, \mathbf{w}) = \frac{\partial H}{\partial \mathbf{x}_i^1} + \frac{\partial H}{\partial \mathbf{x}_i^2} + \sum_j \frac{\partial H}{\partial M_j} \frac{\partial m_j}{\partial \mathbf{x}_i},$$

where the derivatives of H are evaluated at $((\mathbf{X}, \mathbf{w}), (\mathbf{X}, \mathbf{w}), \mathbf{m}(\mathbf{X}, \mathbf{w}))$. The second term on the right-hand side is zero by Lemma 2.1(iii). Direct computation (as in the proof of Lemma 2.1(i),(iv)) gives

$$(2.15) \quad \frac{\partial H}{\partial \mathbf{x}_i^1} = 2m_i(\mathbf{x}_i - \xi_i), \quad \frac{\partial H}{\partial M_j} = w_j + f'(m_j(\mathbf{X}, \mathbf{w})).$$

Combining (2.14), (2.15), and the definition of ω_j yields (2.10).

Differentiating (2.8) with respect to w_i gives

$$(2.16) \quad \frac{\partial E}{\partial w_i}(\mathbf{X}, \mathbf{w}) = \frac{\partial H}{\partial w_i^1} + \frac{\partial H}{\partial w_i^2} + \sum_j \frac{\partial H}{\partial M_j} \frac{\partial m_j}{\partial w_i},$$

where the derivatives of H are evaluated at $((\mathbf{X}, \mathbf{w}), (\mathbf{X}, \mathbf{w}), \mathbf{m}(\mathbf{X}, \mathbf{w}))$. The first two terms on the right-hand side are zero by Lemma 2.1(ii),(iii). Therefore, combining (2.16) and (2.15)₂ yields (2.11). $\square$

Weighted graph Laplacian matrices. Given a power diagram $\{P_i(\mathbf{X}, \mathbf{w})\}$, define a graph G that has vertices $\mathbf{X}$ and edges given by the neighbor relations of the power diagram: $\mathbf{x}_i$ is connected by an edge to $\mathbf{x}_j$ if and only if $i \in J_j$ (and equivalently $j \in J_i$). If we associate a weight $u_{ij} = u_{ji}$ to each edge of this graph, then we can define the weighted graph Laplacian matrix $L = L(G, u)$ by

$$(2.17) \quad L_{ij} = \begin{cases} \sum_{k \in J_j} u_{jk} & \text{if } i = j, \\ -u_{ij} & \text{if } i \in J_j, \\ 0 & \text{otherwise.} \end{cases}$$

The symmetric matrix L is the difference between the weighted degree matrix and the weighted adjacency matrix of G . It is well known that the dimension of the null space of L equals the number of connected components of G . See [26, Thm. 3.1, p. 117]. In our case G is connected, and so, for any edge-weighting u , the null space of $L(G, u)$ is one-dimensional and is spanned by $(1, 1, \dots, 1)$. In an analogous way, one can define (block) weighted graph Laplacian matrices for vector-valued weights $\mathbf{u}_{ij}$.

Computing the derivatives of m_j that appear in (2.10) and (2.11) is delicate since this involves differentiating the integrals $m_j = \int_{P_j(\mathbf{X}, \mathbf{w})} \rho d\mathbf{x}$ with respect to $\mathbf{x}_i$ and w_i . It turns out that these derivatives are weighted graph Laplacian matrices.

LEMMA 2.4 (weighted graph Laplacian structure of $\nabla_{\mathbf{X}} \mathbf{m}$ and $\nabla_{\mathbf{w}} \mathbf{m}$). *Let $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$ be the generators of a power diagram with the generic property that adjacent cells have a common face (a common edge in two dimensions). The partial derivatives of $\mathbf{m}(\mathbf{X}, \mathbf{w})$ are*

$$\frac{\partial m_j}{\partial \mathbf{x}_i} = \begin{cases} \sum_{k \in J_j} \frac{m_{jk}}{d_{jk}} (\bar{\mathbf{x}}_{jk} - \mathbf{x}_j) & \text{if } i = j, \\ -\frac{m_{ij}}{d_{ij}} (\bar{\mathbf{x}}_{ij} - \mathbf{x}_i) & \text{if } i \in J_j, \\ \mathbf{0} & \text{otherwise,} \end{cases} \quad \frac{\partial m_j}{\partial w_i} = \begin{cases} \sum_{k \in J_j} \frac{m_{jk}}{2d_{jk}} & \text{if } i = j, \\ -\frac{m_{ij}}{2d_{ij}} & \text{if } i \in J_j, \\ 0 & \text{otherwise} \end{cases}$$

for $i, j \in \{1, \dots, N\}$. In particular, the N -by- N matrix $\nabla_{\mathbf{w}} \mathbf{m}$, which has components $[\nabla_{\mathbf{w}} \mathbf{m}]_{ij} = \partial m_j / \partial w_i$, is the weighted graph Laplacian matrix of $G(\mathbf{X}, \mathbf{w})$ with respect to the weights $\frac{m_{ij}}{2d_{ij}}$. Therefore, the null space of $\nabla_{\mathbf{w}} \mathbf{m}$ is one-dimensional and is spanned by $(1, 1, \dots, 1) \in \mathbb{R}^N$. Note that $(1, 1, \dots, 1)$ also belongs to the null space of the (Nd) -by- N matrix $\nabla_{\mathbf{X}} \mathbf{m}$, which has d -by-1 blocks $[\nabla_{\mathbf{X}} \mathbf{m}]_{ij} = \partial m_j / \partial \mathbf{x}_i$.

Proof. Given the power diagram $\{P_j\}_{j=1}^N$ generated by $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$, let $\{P_j^t\}_{j=1}^N$ be the power diagram generated by $(\mathbf{X}^t, \mathbf{w}^t) := (\mathbf{X} + t\tilde{\mathbf{X}}, \mathbf{w} + t\tilde{\mathbf{w}})$ for some $\tilde{\mathbf{X}} \in (\mathbb{R}^d)^N$, $\tilde{\mathbf{w}} \in \mathbb{R}^N$. For t in a small enough neighborhood of zero, this family of power diagrams has the same number of cells, and each cell has the same number of faces, as the power diagram generated by $(\mathbf{X}, \mathbf{w})$ (this follows from the assumption that adjacent cells have a common face). Let $\varphi^t : \Omega \rightarrow \Omega$ be any flow map with the properties that φ^0 is the identity map, $\varphi^t(\mathbf{X}) = \mathbf{X}^t$, $\varphi^t(P_j) = P_j^t$ for all j , and φ^t maps the faces of P_j to the faces of P_j^t for all j . Fix j and consider

$$(2.18) \quad m_j(\mathbf{X}^t, \mathbf{w}^t) = \int_{P_j^t} \rho d\mathbf{x} = \int_{\varphi^t(P_j)} \rho d\mathbf{x}.$$

Define $V(\mathbf{x}) = \frac{d}{dt} \varphi^t(\mathbf{x})|_{t=0}$. By the Reynolds Transport Theorem, differentiating

(2.18) with respect to t and evaluating at $t = 0$ gives

$$(2.19) \quad \sum_{i=1}^N \frac{\partial m_j}{\partial \mathbf{x}_i} \cdot \tilde{\mathbf{x}}_i + \frac{\partial m_j}{\partial w_i} \tilde{w}_i = \int_{\partial P_j} \rho V \cdot \mathbf{n} dS = \sum_{k \in J_j} \int_{F_{jk}} \rho V \cdot \mathbf{n}_{jk} dS.$$

Now we compute $V \cdot \mathbf{n}_{jk}$. Choose a face $F_{jk} = P_j \cap P_k$ and some point $\mathbf{x} \in F_{jk}$. Then $\mathbf{x}^t := \varphi^t(\mathbf{x}) \in F_{jk}^t = P_j^t \cap P_k^t$, and so it satisfies

$$|\mathbf{x}^t - \mathbf{x}_j^t|^2 - w_j^t = |\mathbf{x}^t - \mathbf{x}_k^t|^2 - w_k^t.$$

Differentiating with respect to t and setting $t = 0$ gives

$$(2.20) \quad 2(\mathbf{x} - \mathbf{x}_j) \cdot (V(\mathbf{x}) - \tilde{\mathbf{x}}_j) - \tilde{w}_j = 2(\mathbf{x} - \mathbf{x}_k) \cdot (V(\mathbf{x}) - \tilde{\mathbf{x}}_k) - \tilde{w}_k.$$

Recall that $\mathbf{n}_{jk} = (\mathbf{x}_k - \mathbf{x}_j)/d_{jk}$. Therefore, rearranging (2.20) and dividing by d_{jk} yields

$$(2.21) \quad V(\mathbf{x}) \cdot \mathbf{n}_{jk} = \frac{(\mathbf{x} - \mathbf{x}_j) \cdot \tilde{\mathbf{x}}_j - (\mathbf{x} - \mathbf{x}_k) \cdot \tilde{\mathbf{x}}_k}{d_{jk}} + \frac{\tilde{w}_j - \tilde{w}_k}{2d_{jk}}.$$

Substituting this into (2.19) and using (2.2)₂ and (2.3)₂ gives

$$\sum_{i=1}^N \frac{\partial m_j}{\partial \mathbf{x}_i} \cdot \tilde{\mathbf{x}}_i + \frac{\partial m_j}{\partial w_i} \tilde{w}_i = \sum_{k \in J_j} \frac{m_{jk}}{d_{jk}} [(\bar{\mathbf{x}}_{jk} - \mathbf{x}_j) \cdot \tilde{\mathbf{x}}_j - (\bar{\mathbf{x}}_{jk} - \mathbf{x}_k) \cdot \tilde{\mathbf{x}}_k] + \frac{m_{jk}}{2d_{jk}} (\tilde{w}_j - \tilde{w}_k).$$

The derivatives in Lemma 2.4 can be read off from this equation by making suitable choices of $(\tilde{\mathbf{X}}, \tilde{\mathbf{w}})$. $\square$

Remark 2.5. The fact that $(1, 1, \dots, 1) \in \mathbb{R}^N$ belongs to the null space of the matrix $\nabla_{\mathbf{w}} \mathbf{m}$ corresponds to the fact that the power diagram has fixed total mass and that it is invariant under the addition of a constant to all its weights:

$$(2.22) \quad \sum_j m_j = \int_{\Omega} \rho(\mathbf{x}) d\mathbf{x}, \quad m_j(\mathbf{X}, \mathbf{w} + (c, c, \dots, c)) = m_j(\mathbf{X}, \mathbf{w}).$$

Differentiating the first equation with respect to w_i gives $\sum_j \partial m_j / \partial w_i = 0$ for all i , and so $(1, 1, \dots, 1)$ belongs to the null space of $\nabla_{\mathbf{w}} \mathbf{m}$. Differentiating the second equation with respect to c and then setting $c = 0$ gives $\sum_i \partial m_j / \partial w_i = 0$ for all j , and so $(1, 1, \dots, 1)$ belongs to the null space of $(\nabla_{\mathbf{w}} \mathbf{m})^T$ (which equals $\nabla_{\mathbf{w}} \mathbf{m}$ since $\nabla_{\mathbf{w}} \mathbf{m}$ is symmetric).

The main result of this section is the following proposition.

PROPOSITION 2.6 (critical points of E are fixed points of the Lloyd maps). *Let $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$ be a critical point of E . Assume that the power diagram generated by $(\mathbf{X}, \mathbf{w})$ has the generic property that adjacent cells have a common face (a common edge in two dimensions). Then, up to the addition of a constant to the weights, $(\mathbf{X}, \mathbf{w})$ is a fixed point of the Lloyd maps ξ and ω :*

$$(2.23) \quad \xi(\mathbf{X}, \mathbf{w}) = \mathbf{X}, \quad \omega(\mathbf{X}, \mathbf{w}) = \mathbf{w} + \mathbf{c},$$

where $\mathbf{c} = c(1, 1, \dots, 1) \in \mathbb{R}^N$. In particular, critical points of E are CPDs.

Proof. Equation (2.11) yields

$$\mathbf{0} = \nabla_{\mathbf{w}} E = \nabla_{\mathbf{w}} \mathbf{m}(\mathbf{w} - \omega(\mathbf{X}, \mathbf{w})).$$

By Lemma 2.4, $\omega(\mathbf{X}, \mathbf{w}) = \mathbf{w} + \mathbf{c}$ for some $\mathbf{c} = c(1, 1, \dots, 1) \in \mathbb{R}^N$. Since $\mathbf{c}$ belongs to the null space of $\nabla_{\mathbf{X}} \mathbf{m}$, (2.10) implies that

$$(2.24) \quad \mathbf{0} = \frac{\partial E}{\partial \mathbf{x}_i}(\mathbf{X}, \mathbf{w}) = 2m_i(\mathbf{x}_i - \boldsymbol{\xi}_i(\mathbf{X}, \mathbf{w})).$$

By assumption the power diagram generated by $(\mathbf{X}, \mathbf{w})$ has no empty cells. Therefore, $m_i \neq 0$ for any i and (2.24) gives $\mathbf{X} - \boldsymbol{\xi}(\mathbf{X}, \mathbf{w}) = \mathbf{0}$, as required. $\square$

Remark 2.7 (examples of critical points of E). Any CVT of Ω with the property that all cells have the same mass is a critical point of E . If $\rho = \text{constant}$ and Ω is a domain with nice symmetry, e.g., a square or a disc, then it is easy to write down several, in fact infinitely many, CVTs with this property and hence find infinitely many critical points of E (although not all will be local minima). The highly nonconvex nature of the energy landscape makes it difficult to find global minima. See section 4.1.

2.4. Centroidal power diagrams. In this section we give more examples of centroidal power diagrams (CPDs) and critical points of E .

Example 2.8 (CPDs in one dimension). All partitions of intervals are CPDs. If $a = y_0 < y_1 < \dots < y_N = b$ is any partition of $[a, b]$, then the following choice of generators gives rise to the centroidal power cells $P_i = [y_{i-1}, y_i]$ with respect to the density $\rho = 1$:

$$x_i = \frac{y_{i-1} + y_i}{2}, \quad i = 1, \dots, N,$$

$$w_1 = 0, \quad w_i = w_{i-1} + \frac{1}{4}(y_i - y_{i-2})(y_i + y_{i-2} - 2y_{i-1}), \quad i = 2, \dots, N.$$

Example 2.9 (CPDs in two dimensions). Figure 2 gives two examples of a CPD in two dimensions for $\rho = 1$. These are nontrivial examples in the sense that they are not centroidal *Voronoi tessellations* (CVTs). The diagrams were generated using a modification of the classical Lloyd algorithm, in which the weights of the generators are fixed and only the locations are updated at each iteration (to the centroids of the power cells). For the examples shown, the generators were initially arranged in a square lattice with a checkerboard pattern of weights, the generator locations were perturbed very slightly, and the modified Lloyd algorithm was applied. This procedure produced the patterns shown in Figure 2.

Having obtained examples of CPDs, we address the question of which CPDs can arise as critical points of E .

DEFINITION 2.10 (monotone power diagrams). *Given a density $\rho \in L^1(\Omega; [0, \infty))$, a power diagram $\{\mathbf{x}_i, w_i\}$ in Ω is monotone with respect to ρ if it satisfies the following:*

If $w_i > w_j$, then $m_i \geq m_j$. If $m_i = m_j$, then $w_i = w_j$.

PROPOSITION 2.11 (all monotone CPDs are critical points of E for some f). *Let $\rho \in L^1(\Omega; [0, \infty))$. Let $\{\mathbf{x}_i, w_i\}$ generate a CPD in Ω . The following are equivalent:*

1. *The CPD is monotone with respect to ρ .*
2. *There exists a function f satisfying (1.5) such that the CPD is a critical point of E for this choice of f .*

The proof, which is constructive, appears in the online supplementary material. This proposition implies that critical points of E are not only CPDs but in fact monotone CPDs.

Example 2.12 (monotone CPDs). Both examples in Figure 2 are monotone CPDs with respect to $\rho = 1$. In each case we can construct an admissible f that is affine in

FIG. 2. CPDs of 36 points with density $\rho = 1$. In the first example (left) the 18 small cells have weights $w \approx -0.01389$, and the 18 large cells have weights $w = 0$. In the second example (center) the 18 small cells have weights $w \approx -0.01667$, and the 18 large cells have weights $w = 0$. Both diagrams were obtained by applying a modified form of Lloyd's algorithm (in which the locations are updated to the centroids but the weights are fixed) to an initial checkerboard pattern in which the generators lie on a square lattice with an alternating pattern of weights. Both examples are fixed points of Algorithm 1 with an appropriately constructed f .

a neighborhood of the points m_i and has slopes $f'(m_i) = -w_i$. The f for Figure 2 (left) is shown in Figure 3. By construction, the power diagrams $\{P_i\}_{i=1}^{36}$ shown in Figure 2 are local minimizers of E for their respective f . To see this, consider a small perturbation of $\{P_i\}_{i=1}^{36}$. Since f is affine in a neighborhood of $\{m_i\}_{i=1}^{36}$, each iteration of the generalized Lloyd algorithm produces a power diagram with exactly the same weights as $\{P_i\}_{i=1}^{36}$ provided that the perturbation is small enough. Moreover, since $\{P_i\}_{i=1}^{36}$ was constructed using the classical Lloyd algorithm (with the weights fixed), the generator locations $\{x_i^k\}_{i=1}^{36}$ produced by the generalized Lloyd algorithm converge to the generator locations of $\{P_i\}_{i=1}^{36}$ for small enough perturbations.

Remark 2.13 (optimal partitions with cells of different sizes). Typically optimal partitions exhibit cells of roughly the same size (see section 4). Example 2.12 provides evidence that there exist optimal partitions with cells of different sizes. (Note, however, that we have only shown that these are *local* minimizers.) In [4, sec. 3.4] a closely related optimal location energy is studied, and analytical evidence is given for the existence of optimal partitions with different cell sizes. Example 2.12 is a concrete example. This example is also loosely related to the problem of creating materials with designer microstructure. The function f could be thought of as a control to produce a desired microstructure, represented by a monotone CPD.

3. Properties of the algorithm. Our main result is the following theorem.

THEOREM 3.1. *Assume that f is concave. Then the generalized Lloyd algorithm is energy decreasing:*

$$E(\mathbf{X}^{n+1}, \mathbf{w}^{n+1}) \leq E(\mathbf{X}^n, \mathbf{w}^n),$$

where $\mathbf{X}^{n+1} = \boldsymbol{\xi}(\mathbf{X}^n, \mathbf{w}^n)$, $\mathbf{w}^{n+1} = \boldsymbol{\omega}(\mathbf{X}^n, \mathbf{w}^n)$, $(\mathbf{X}^n, \mathbf{w}^n) \in \mathcal{G}^N$. The inequality is strict unless $(\mathbf{X}^{n+1}, \mathbf{w}^{n+1}) = (\mathbf{X}^{n+2}, \mathbf{w}^{n+2})$, i.e., unless the algorithm has converged.

FIG. 3. The leftmost CPD in Figure 2 is a monotone CPD and is a critical point of the energy with appropriately constructed f . For this CPD, the construction of f following the proof of Proposition 2.11 gives the function $f(m)$ shown here (left). As a test, our generalized Lloyd algorithm using this f was applied to a random initial configuration of 36 cells. The algorithm converged to the CPD with two different cell sizes shown above (right).

Proof. The proof follows easily by stringing together the properties of H from Lemma 2.1:

$$\begin{aligned}
E(\mathbf{X}^n, \mathbf{w}^n) &= H((\mathbf{X}^n, \mathbf{w}^{n+1}), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^n, \mathbf{w}^n)) && \text{(by Lemma 2.1(ii))} \\
&= H((\mathbf{X}^n, \omega(\mathbf{X}^n, \mathbf{w}^n)), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^n, \mathbf{w}^n)) && \text{(by definition of } \mathbf{w}^{n+1}) \\
&\geq H((\mathbf{X}^n, \omega(\mathbf{X}^n, \mathbf{w}^n)), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})) && \text{(by Lemma 2.1(iv))} \\
&= H((\mathbf{X}^n, \mathbf{w}^{n+1}), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})) && \text{(by definition of } \mathbf{w}^{n+1}) \\
&\geq H((\xi(\mathbf{X}^n, \mathbf{w}^n), \mathbf{w}^{n+1}), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})) && \text{(by Lemma 2.1(i))} \\
&= H((\mathbf{X}^{n+1}, \mathbf{w}^{n+1}), (\mathbf{X}^n, \mathbf{w}^n), \mathbf{m}(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})) && \text{(by definition of } \mathbf{X}^{n+1}) \\
&\geq H((\mathbf{X}^{n+1}, \mathbf{w}^{n+1}), (\mathbf{X}^{n+1}, \mathbf{w}^{n+1}), \mathbf{m}(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})) && \text{(by Lemma 2.1(iii))} \\
&= E(\mathbf{X}^{n+1}, \mathbf{w}^{n+1}) && \text{(by (2.8)).}
\end{aligned}$$

By Lemma 2.1(iii) the last inequality is strict unless $P_i(\mathbf{X}^{n+1}, \mathbf{w}^{n+1}) = P_i(\mathbf{X}^n, \mathbf{w}^n)$ for all i , up to sets of $\rho d\mathbf{x}$ -measure zero, in which case x_i^{n+2} (which is the centroid of $P_i(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})$) equals x_i^{n+1} (which is the centroid of $P_i(\mathbf{X}^n, \mathbf{w}^n)$) and

$$w_i^{n+2} = -f'(|P_i(\mathbf{X}^{n+1}, \mathbf{w}^{n+1})|) = -f'(|P_i(\mathbf{X}^n, \mathbf{w}^n)|) = w_i^{n+1},$$

as required. $\square$

Remark 3.2 (elimination of generators is energy decreasing). The generalized Lloyd algorithm removes generators corresponding to empty cells; i.e., if $P_i^n = \emptyset$, then the generator pair (x_i^n, w_i^n) is removed in step (2) of Algorithm 1. The assumption that $f(0) \geq 0$ ensures that removing generators is energy decreasing.

Recall from (1.10) that $\mathcal{G}^N$ is the set of N generators such that no two generators coincide and that the corresponding power diagram has no empty cells. The energy decreasing property of the algorithm can be used to prove the following convergence

result, which is a generalization of the convergence theorem for the classical Lloyd algorithm [10, Thm. 2.6].

THEOREM 3.3 (convergence of the generalized Lloyd algorithm). *Assume that f is concave. Assume that E has only finitely many critical points with the same energy. Let $(\mathbf{X}^k, \mathbf{w}^k)$ be a sequence generated by Algorithm 1. Let K be large enough such that, for all $k \geq K$, $(\mathbf{X}^k, \mathbf{w}^k) \in \mathcal{G}^N$ for N fixed, i.e., there is no elimination of generators after iteration K . If the sequence $(\mathbf{X}^k, \mathbf{w}^k)_{k \geq K}$ is a compact subset of $\mathcal{G}^N$, then it converges to a critical point of E .*

Proof. This follows by combining a minor modification of the proof of the Global Convergence Theorem from [23, p. 206] with a convergence theorem for the classical Lloyd algorithm [10, Thm. 2.5]. Note that the Lloyd maps ξ_i, ω_i and the energy E are continuous on $\mathcal{G}^N$ by the continuity of the mass and first and second moments of mass of the power cells P_i , and the continuity of f .

Let $(\mathbf{X}^{k_j}, \mathbf{w}^{k_j})$ be a convergent subsequence converging to $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$. By the continuity of E on $\mathcal{G}^N$, $E(\mathbf{X}^{k_j}, \mathbf{w}^{k_j}) \rightarrow E(\mathbf{X}, \mathbf{w})$. Take J large enough so that $E(\mathbf{X}^{k_J}, \mathbf{w}^{k_J}) - E(\mathbf{X}, \mathbf{w}) < \varepsilon$. By Theorem 3.1 the whole sequence $E(\mathbf{X}^k, \mathbf{w}^k)$ converges to $E(\mathbf{X}, \mathbf{w})$ since for all $k > k_J$

$$0 \leq E(\mathbf{X}^k, \mathbf{w}^k) - E(\mathbf{X}, \mathbf{w}) \leq E(\mathbf{X}^k, \mathbf{w}^k) - E(\mathbf{X}^{k_J}, \mathbf{w}^{k_J}) + E(\mathbf{X}^{k_J}, \mathbf{w}^{k_J}) - E(\mathbf{X}, \mathbf{w}) < \varepsilon.$$

Next we check that $(\mathbf{X}, \mathbf{w})$ is a fixed point of the Lloyd maps and hence a critical point of E . Consider the sequence $(\mathbf{X}^{k_{j_l}-1}, \mathbf{w}^{k_{j_l}-1})$. By the compactness of $(\mathbf{X}^k, \mathbf{w}^k)$ there is a subsequence $(\mathbf{X}^{k_{j_l}-1}, \mathbf{w}^{k_{j_l}-1})$ converging to $(\mathbf{X}_-, \mathbf{w}_-) \in \mathcal{G}^N$. The continuity of the Lloyd maps on $\mathcal{G}^N$ implies that

$$(\xi(\mathbf{X}^{k_{j_l}-1}, \mathbf{w}^{k_{j_l}-1}), \omega(\mathbf{X}^{k_{j_l}-1}, \mathbf{w}^{k_{j_l}-1})) = (\mathbf{X}^{k_{j_l}}, \mathbf{w}^{k_{j_l}}) \rightarrow (\xi(\mathbf{X}_-, \mathbf{w}_-), \omega(\mathbf{X}_-, \mathbf{w}_-)).$$

But $(\mathbf{X}^{k_{j_l}}, \mathbf{w}^{k_{j_l}}) \rightarrow (\mathbf{X}, \mathbf{w})$. Therefore, $(\xi(\mathbf{X}_-, \mathbf{w}_-), \omega(\mathbf{X}_-, \mathbf{w}_-)) = (\mathbf{X}, \mathbf{w})$. Since $E(\mathbf{X}^k, \mathbf{w}^k) \rightarrow E(\mathbf{X}, \mathbf{w})$, we obtain that

$$E(\mathbf{X}_-, \mathbf{w}_-) = E(\mathbf{X}, \mathbf{w}) = E(\xi(\mathbf{X}_-, \mathbf{w}_-), \omega(\mathbf{X}_-, \mathbf{w}_-))$$

and thus, by Theorem 3.1, $(\xi(\mathbf{X}_-, \mathbf{w}_-), \omega(\mathbf{X}_-, \mathbf{w}_-)) = (\mathbf{X}, \mathbf{w})$ is a fixed point of the Lloyd maps.

We have shown that any accumulation point of $(\mathbf{X}^k, \mathbf{w}^k)$ is a fixed point of the Lloyd maps and, by the energy decreasing property of the algorithm, all accumulation points have the same energy. Therefore, by the first assumption of the theorem, it follows that $(\mathbf{X}^k, \mathbf{w}^k)$ has only finitely many accumulation points.

Finally, the whole sequence $(\mathbf{X}^k, \mathbf{w}^k)$ converges to $(\mathbf{X}, \mathbf{w})$ by the following result, which is proved in [10, Thm. 2.5] for the classical Lloyd algorithm but holds for general fixed point methods of the form $z^{k+1} = T(z^k)$: If the sequence $\{z^k\}$ generated by $z^{k+1} = T(z^k)$ has finitely many accumulation points, T is continuous at them, and they are fixed points of T , then z^k converges. This completes the proof. $\square$

Remark 3.4 (assumptions of the convergence theorem). The assumption that E has only finitely many critical points with the same energy is true for generic domains Ω but not for all; e.g., if Ω is a ball and ρ is radially symmetric, then there could be infinitely many fixed points with the same energy by rotational symmetry. The assumption that $(\mathbf{X}^k, \mathbf{w}^k)_{k \geq K}$ is a compact subset of $\mathcal{G}^N$ is stronger. It means that in the limit there is no elimination of generators. We need this assumption since the Lloyd maps are not defined if there are empty cells; $P_i = \emptyset$ for some i . While numerical experiments suggest that cells do not disappear in the limit, it is difficult

to prove. For the classical Lloyd algorithm it was proved in one dimension in [10, Prop. 2.9] and in higher dimensions in [12]. For further convergence theorems for the classical Lloyd algorithm, see [11] and [28].

Remark 3.5 (interpretation of the Lloyd algorithm as a descent method). In the following proposition we study the structure of the generalized Lloyd algorithm. Recall that an iterative method is a descent method for an energy $\mathcal{E}$ if it can be written in the form

$$(3.1) \quad \mathbf{z}_{n+1} = \mathbf{z}_n - \alpha_n \mathbf{B}_n \nabla \mathcal{E},$$

where $\mathbf{B}_n$ is positive-definite, α_n is the step size, and $-\mathbf{B}_n \nabla \mathcal{E}$ is the step direction; e.g., $\mathbf{B}_n = \mathbf{I}$ is the steepest descent method, $\mathbf{B}_n = (D^2 \mathcal{E})^{-1}$, and $\alpha_n = 1$ is Newton's method. The following proposition asserts that the generalized Lloyd algorithm can be written in the form (3.1), but not that $\mathbf{B}_n$ is positive-definite, which we are unable to prove.

PROPOSITION 3.6. *The generalized Lloyd algorithm can be written in the form*

$$(3.2) \quad \begin{pmatrix} \mathbf{X}^{n+1} \\ \mathbf{w}^{n+1} \end{pmatrix} = \begin{pmatrix} \mathbf{X}^n \\ \mathbf{w}^n \end{pmatrix} - \mathbf{B}_n \begin{pmatrix} \nabla_{\mathbf{X}} E^n \\ \nabla_{\mathbf{w}} E^n \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \mathbf{c} \end{pmatrix},$$

where $\mathbf{B}_n$ is a square matrix of dimension $N(d+1)$ and $\mathbf{c} = c(1, 1, \dots, 1)^T$ for some $c \in \mathbb{R}$.

The proof is given in the online supplementary material.

Remark 3.7 (alternative algorithm). The following proposition gives explicit expressions for the derivatives of the Lloyd maps ξ and ω . These could be used to find critical points of E in an alternative way, e.g., by solving the nonlinear equations (2.23) using Newton's method. For recent work on quasi-Newton methods for the calculation of CVTs, see [17, 21].

PROPOSITION 3.8 (derivatives of the Lloyd maps). *Given a face F of a power diagram, define the matrix $\mathcal{S}(F)$ by*

$$\mathcal{S}(F) = \frac{1}{m(F)} \int_F \mathbf{x} \otimes \mathbf{x} \rho(\mathbf{x}) dS,$$

where $m(F) = \int_F \rho dS$ is the mass of the face. Let $f \in C^2([0, \infty))$. Let $(\mathbf{X}, \mathbf{w}) \in \mathcal{G}^N$ be the generators of a power diagram with the generic property that adjacent cells have a common face (a common edge in two dimensions). The derivatives of the Lloyd maps $\xi(\mathbf{X}, \mathbf{w})$ and $\omega(\mathbf{X}, \mathbf{w})$ are

$$\left(\frac{\partial \xi}{\partial \mathbf{X}} \right)_{ij} = \frac{\partial \xi_i}{\partial \mathbf{x}_j} = \begin{cases} \frac{1}{m_i} \sum_{k \in J_i} \frac{m_{ik}}{d_{ik}} (\mathcal{S}(F_{ik}) - \bar{\mathbf{x}}_{ik} \otimes \mathbf{x}_i + \bar{\mathbf{x}}_i \otimes (\mathbf{x}_i - \bar{\mathbf{x}}_{ik})) & \text{if } i = j, \\ -\frac{m_{ij}}{m_i d_{ij}} (\mathcal{S}(F_{ij}) - \bar{\mathbf{x}}_{ij} \otimes \mathbf{x}_j + \bar{\mathbf{x}}_i \otimes (\mathbf{x}_j - \bar{\mathbf{x}}_{ij})) & \text{if } j \in J_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

$$\left(\frac{\partial \xi}{\partial \mathbf{w}} \right)_{ij} = \frac{\partial \xi_i}{\partial w_j} = \begin{cases} \frac{1}{2m_i} \sum_{k \in J_i} \frac{m_{ik}}{d_{ik}} (\bar{\mathbf{x}}_{ik} - \bar{\mathbf{x}}_i) & \text{if } i = j, \\ -\frac{m_{ij}}{2m_i d_{ij}} (\bar{\mathbf{x}}_{ij} - \bar{\mathbf{x}}_i) & \text{if } j \in J_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

$$\left(\frac{\partial \boldsymbol{\omega}}{\partial \mathbf{X}}\right)_{ij} = \frac{\partial \omega_i}{\partial x_j} = \begin{cases} -f''(m_i) \sum_{k \in J_i} \frac{m_{ik}}{d_{ik}} (\bar{\mathbf{x}}_{ik} - \mathbf{x}_i) & \text{if } i = j, \\ f''(m_i) \frac{m_{ij}}{d_{ij}} (\bar{\mathbf{x}}_{ij} - \mathbf{x}_i) & \text{if } j \in J_i, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

$$\left(\frac{\partial \boldsymbol{\omega}}{\partial \mathbf{w}}\right)_{ij} = \frac{\partial \omega_i}{\partial w_j} = \begin{cases} -f''(m_i) \sum_{k \in J_i} \frac{m_{ik}}{2d_{ik}} & \text{if } i = j, \\ f''(m_i) \frac{m_{ij}}{2d_{ij}} & \text{if } j \in J_i, \\ 0 & \text{otherwise.} \end{cases}$$

We order the block matrices $\partial \boldsymbol{\xi} / \partial \mathbf{X}$, $\partial \boldsymbol{\xi} / \partial \mathbf{w}$, $\partial \boldsymbol{\omega} / \partial \mathbf{X}$, and $\partial \boldsymbol{\omega} / \partial \mathbf{w}$ so that they have dimensions (Nd) -by- (Nd) , (Nd) -by- N , N -by- (Nd) , and N -by- N .

The proof appears in the online supplementary material. Potentially these derivatives could also be used to prove convergence of the Lloyd algorithm by proving that the Lloyd map pair $(\boldsymbol{\xi}, \boldsymbol{\omega}) : \mathcal{G}^N \rightarrow \mathcal{G}^N$ is a contraction. These derivatives are also needed to evaluate the Hessian of E , which can be used to check the stability of fixed points.

PROPOSITION 3.9 (Hessian of E evaluated at fixed points). *Let $f \in C^2([0, \infty))$. If $(\mathbf{X}, \mathbf{w})$ is a fixed point of the Lloyd maps $\boldsymbol{\xi}$ and $\boldsymbol{\omega}$, i.e., if it satisfies (2.23), then the Hessian of E evaluated at $(\mathbf{X}, \mathbf{w})$ is*

$$\begin{pmatrix} E_{\mathbf{X}\mathbf{X}} & E_{\mathbf{X}\mathbf{w}} \\ E_{\mathbf{w}\mathbf{X}} & E_{\mathbf{w}\mathbf{w}} \end{pmatrix} = \begin{pmatrix} 2\hat{\mathbf{M}} & \nabla_{\mathbf{X}} \mathbf{m} \\ \mathbf{0} & \nabla_{\mathbf{w}} \mathbf{m} \end{pmatrix} \begin{pmatrix} \mathbf{I}_{Nd} - \frac{\partial \boldsymbol{\xi}}{\partial \mathbf{X}} & -\frac{\partial \boldsymbol{\xi}}{\partial \mathbf{w}} \\ -\frac{\partial \boldsymbol{\omega}}{\partial \mathbf{X}} & \mathbf{I}_N - \frac{\partial \boldsymbol{\omega}}{\partial \mathbf{w}} \end{pmatrix},$$

where $\mathbf{0}$ is the N -by- (Nd) zero matrix, $\hat{\mathbf{M}}$ is as defined in (2.13), $E_{\mathbf{X}\mathbf{X}}$ is the (Nd) -by- (Nd) block matrix with d -by- d blocks $\partial^2 E / \partial \mathbf{x}_i \partial \mathbf{x}_j$, $E_{\mathbf{w}\mathbf{w}}$ is the N -by- N matrix with entries $[E_{\mathbf{w}\mathbf{w}}]_{ij} = \partial^2 E / \partial w_i \partial w_j$, $E_{\mathbf{X}\mathbf{w}}$ is the (Nd) -by- N block matrix with d -by-1 blocks $\partial^2 E / \partial \mathbf{x}_i \partial w_j$, and $E_{\mathbf{w}\mathbf{X}}$ is the N -by- (Nd) block matrix with 1-by- d blocks $\partial^2 E / \partial w_i \partial \mathbf{x}_j$.

Proof. The proof follows immediately from (2.12). $\square$

4. Illustrations and applications. In this section we implement the algorithm in two dimensions. We use a crystallization problem to illustrate the typical flatness and nonconvexity of the energy landscape and the rate of convergence of the algorithm. In the online supplementary material we provide further examples of the algorithm with an optimal location problem with nonconstant ρ and a three-dimensional example where we test a conjecture about the optimality of the BCC lattice for a crystallization problem.

4.1. Nonconvexity and flatness of energy landscape. In this section we look for critical points of the two-dimensional block copolymer energy from section 1.5.1:

$$(4.1) \quad E(\{\mathbf{x}_i, w_i\}) = \sum_{i=1}^N \left\{ \lambda \sqrt{m_i} + \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 d\mathbf{x} \right\},$$

where $\mathbf{x}_i \in \Omega = [0, 1]^2$. This example first appeared in [5]. It is the special case of (1.9) with $\rho = 1$, $f(m) = \lambda \sqrt{m}$, where $\lambda > 0$ is a parameter representing the

FIG. 4. Flatness of the energy landscape: Some local minimizers of the energy (4.1) for $\lambda = 0.005$. The polygons are the power cells P_i , and the points are the generators $\mathbf{x}_i$. The weights w_i are not shown. The shading corresponds to the number of sides of the cells.

strength of the repulsion between the two phases of the block copolymer. The scaling of the energy suggests that the optimal value of N scales like $\lambda^{-\frac{2}{3}}$. Figure 4 shows local minimizers of E for $\lambda = 0.005$. We believe that the top left figure is a global minimizer. These were generated using 25000 random initial conditions to probe the nonconvex energy landscape. The energy has infinitely many critical points, e.g., every CVT of $[0, 1]^2$ with cells of equal area (such as the checkerboard configuration) is a critical point. The flatness of the energy landscape can be seen from the energy values in Figure 4.

As λ decreases it becomes harder to find global minimizers. Figure 5 shows two local minimizers for $\lambda = 10^{-5}$. The figure on the left was obtained by using the triangular lattice as an initial condition. It was proved in [6] that the triangular lattice is optimal in the limit $\lambda \rightarrow 0$. The figure on the right was obtained with a random initial condition. The “grains” of hexagonal tiling resemble grains in metals. This suggests that energies of the form (1.9) could be used to simulate material microstructure—for example, to produce Representative Volume Elements for finite element simulations [1].

4.2. Convergence rate. In this section we study the rate of convergence of the algorithm to critical points of the energy (4.1) with $\lambda = 0.005$. Figure 6 shows the logarithm of the approximate error of the energy plotted against the number of iterations n for three simulations with random initial conditions. The initial number of generators was $N = 6, 10, 25$, and there was no elimination of generators throughout

$$N = 1037, E = 4.775 \times 10^{-4}, \lambda = 10^{-5}$$

$$N = 1037, E = 4.787 \times 10^{-4}, \lambda = 10^{-5}$$

FIG. 5. Two local minimizers of the energy (4.1) for $\lambda = 10^{-5}$ with $N = 1037$ in both cases. In the first case the cell generators were initially arranged in a triangular lattice; in the second case they were distributed randomly.

the simulations. The approximate error was computed using the value of the energy at the final iteration. The graph shows that the energy converges linearly, meaning that the error at the n th iteration ε_n satisfies $\varepsilon_{n+1}/\varepsilon_n \rightarrow r$, where $r \in (0, 1)$ is the rate of convergence. We observe that the rate of convergence decreases as the number of generators increases and that $r \sim 1 - \frac{C}{N}$ for some constant C . In [10] it was found that for the classical Lloyd algorithm with $\rho = 1$ in one dimension the rate of convergence of the generators (rather than the energy) is approximately $1 - 1/(4\pi^2 N^2)$. This was found from the spectrum of the derivative of the Lloyd map. In principle the rate of convergence of the generalized Lloyd algorithm could be found using the derivatives given in Proposition 3.8. We believe that region $(\star)$ in the figure is the result of the Lloyd iterates passing close to a saddle point of the energy on the way to a local minimum.

5. Concluding remarks.

5.1. Limitations of the algorithm. First we discuss the assumptions on the data given in (1.4), (1.5).

The assumption that Ω is convex ensures that the centroid of each power cell lies in Ω . Without this assumption the algorithm could produce an unfeasible solution with $\mathbf{x}_i \notin \Omega$ for some i . For example, if Ω is the annulus $A(r_1, r_2)$ centered at the origin, $\rho = 1$, and f is chosen suitably, then E is minimized when $N = 1$ by $(\mathbf{x}_1, w_1)$ in which $|\mathbf{x}_1| = r_1$ (the generator lies on the interior boundary of the annulus) and w_1 is irrelevant (in the case where there is only one cell the weight is not determined). The generalized Lloyd algorithm, however, initialized with $N_0 = 1$, would return $\mathbf{x} = \mathbf{0} \notin \Omega$. This strong limitation on the shape of Ω means that the algorithm cannot be used to solve optimal location problems in highly nonconvex countries like Scotland. We plan to address this issue in a future paper.

The concavity assumption on f was used to prove Theorem 3.1, which asserts

FIG. 6. Rate of convergence of the generalized Lloyd algorithm to critical points of the energy (4.1) with $\lambda = 0.005$. The approximate deviation of the energy from its minimum against the number of iterations is plotted on semilog axes for three simulations, each using random initial conditions. The initial number of generators in the three cases was $N = 6, 10, 25$, and there was no elimination of generators throughout the simulations. We see that the algorithm converges linearly with rate r . The rate was computed by fitting straight lines to the data.

that step (1) of the algorithm decreases the energy at every iteration. The assumption $f(0) \geq 0$ ensures that iteration step (2) is energy decreasing as well. These are also reasonable modeling assumptions for many applications, as discussed in section 1.5, and the energy decreasing property is used to prove the convergence theorem. The concavity of f , however, is not necessary for the existence of a minimizer of E , which merely requires that f be subadditive (along with the growth condition $(1.5)_4$ and sufficient regularity). For example, E has a minimizer if we take f to be the convex, subadditive function $f(m) = e^{-m}$. Therefore, there is a gap between the assumptions needed for existence of a minimizer and those needed for the performance of the algorithm.

As discussed in section 1.3, another limitation of the algorithm is that, while it can annihilate generators (step (2)), it cannot create them. Therefore, the initial guess N_0 for the optimal number of generators should be an overestimate. This limitation could be addressed by using a simulated annealing method to randomly introduce new generators at certain iterations. This could also be used to prevent the algorithm from getting stuck at a local minimizer.

5.2. Generalizations. While we have focused on energy (1.9), our general methodology could be easily applied to broader classes of optimal location energies where the first term is more general, e.g., to

$$E(\{\mathbf{x}_i, w_i\}) = g(\{\mathbf{x}_i, m_i\}) + \sum_{i=1}^N \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^2 \rho(\mathbf{x}) d\mathbf{x},$$

where $m_i = \int_{P_i} \rho d\mathbf{x}$. Our algorithm can also be modified to minimize the following energy, which is obtained from (1.1) by replacing the square of the 2-Wasserstein

distance with the p th power of the p -Wasserstein distance, $p \in [1, \infty)$:

$$F_p(\{\mathbf{x}_i, m_i\}) = \sum_{i=1}^N f(m_i) + d_p^p\left(\rho, \sum_{i=1}^N m_i \delta_{\mathbf{x}_i}\right).$$

See [30, Chap. 7] for the definition of $d_p(\cdot, \cdot)$. In this case the energy can be rewritten in terms of what we call p -power diagrams. These are a generalization of power diagrams where the cells generated by $\{\mathbf{x}_i, w_i\}$ are defined by

$$P_i = \{\mathbf{x} \in \Omega : |\mathbf{x} - \mathbf{x}_i|^p - w_i \leq |\mathbf{x} - \mathbf{x}_k|^p - w_k \ \forall k\}.$$

For $p = 2$ this is just the power diagram. For $p = 1$ this is known as the Apollonius diagram (or the additively weighted Voronoi diagram, or the Voronoi diagram of disks). For general p there does not seem to be a standard name, although they fall into the class of generalized Dirichlet tessellations, or generalized additively weighted Voronoi diagrams. It can be shown that minimizing F_p is equivalent to minimizing

$$(5.1) \quad E_p(\{\mathbf{x}_i, w_i\}) = \sum_{i=1}^N \left\{ f(m_i) + \int_{P_i} |\mathbf{x} - \mathbf{x}_i|^p \rho(\mathbf{x}) d\mathbf{x} \right\},$$

where $\{P_i\}$ is the p -power diagram generated by $\{\mathbf{x}_i, w_i\}$ and $m_i := \int_{P_i} \rho d\mathbf{x}$. See [5, sec. 4.2]. Critical points of E_p can be found using a modification of the generalized Lloyd algorithm where for each i the map ξ_i returns the p -centroid of the p -power cell P_i , i.e., $\xi_i(\mathbf{X}, \mathbf{w})$ satisfies the equation

$$(5.2) \quad \int_{P_i} (\xi_i - \mathbf{x}) |\xi_i - \mathbf{x}|^{p-2} d\mathbf{x} = \mathbf{0}.$$

See [5, Thm. 4.16]. For the case $p = 2$ this equation just says that ξ_i is the centroid of P_i . Therefore, in principle the algorithm can be extended to all $p \in [1, \infty)$. In practice it is much harder to implement. Except for the cases $p = 1, 2$, we are not aware of any efficient algorithms for computing p -power diagrams. This is due to the fact that for $p \neq 2$ the boundaries between cells are curved (unless all the weights are equal). In addition, evaluating the Lloyd map $\xi(\mathbf{X}, \mathbf{w})$ involves solving the nonlinear equation (5.2). We plan to say more about these aspects in a future paper.

Acknowledgments. The generalized Lloyd algorithm was derived for a special case in collaboration with Mark Peletier [5]. The authors would like to thank Alexander Rand for insightful comments that strengthened the paper. All plots were prepared using Gnuplot.

REFERENCES

- [1] J. ALSAYEDNOOR, P. HARRISON, AND Z. GUO, *Large strain compressive response of 2-d periodic representative volume element for random foam microstructures*, Mech. Mater., 66 (2013), pp. 7–20.
- [2] F. AURENHAMMER, F. HOFFMANN, AND B. ARONOV, *Minkowski-type theorems and least-squares clustering*, Algorithmica, 20 (1998), pp. 61–76.
- [3] F. AURENHAMMER, R. KLEIN, AND D.-T. LEE, *Voronoi Diagrams and Delaunay Triangulations*, World Scientific, Singapore, 2013.
- [4] G. BOUCHITTÉ, C. JIMENEZ, AND R. MAHADEVAN, *Asymptotic analysis of a class of optimal location problems*, J. Math. Pures Appl., 95 (2011), pp. 382–419.

- [5] D.P. BOURNE, M.A. PELETIER, AND S.M. ROPER, *Hexagonal patterns in a simplified model for block copolymers*, SIAM J. Appl. Math., 74 (2014), pp. 1315–1337.
- [6] D.P. BOURNE, M.A. PELETIER, AND F. THEIL, *Optimality of the triangular lattice for a particle system with Wasserstein interaction*, Comm. Math. Phys., 329 (2014), pp. 117–140.
- [7] A. BRIEDEN AND P. GRITZMANN, *On optimal weighted balanced clusterings: Gravity bodies and power diagrams*, SIAM J. Discrete Math., 26 (2012), pp. 415–434.
- [8] G. BUTTAZZO AND F. SANTAMBROGIO, *A mass transportation model for the optimal planning of an urban region*, SIAM Rev., 51 (2009), pp. 593–610.
- [9] R. CHOKSI, M.A. PELETIER, AND J.F. WILLIAMS, *On the phase diagram for microphase separation of diblock copolymers: An approach via a nonlocal Cahn–Hilliard functional*, SIAM J. Appl. Math., 69 (2009), pp. 1712–1738.
- [10] Q. DU, M. EMELIANENKO, AND L. JU, *Convergence of the Lloyd algorithm for computing centroidal Voronoi tessellations*, SIAM J. Numer. Anal., 44 (2006), pp. 102–119.
- [11] Q. DU, V. FABER, AND M. GUNZBURGER, *Centroidal Voronoi tessellations: Applications and algorithms*, SIAM Rev., 41 (1999), pp. 637–676.
- [12] M. EMELIANENKO, L. JU, AND A. RAND, *Nondegeneracy and weak global convergence of the Lloyd algorithm in $\mathbb{R}^d$* , SIAM J. Numer. Anal., 46 (2008), pp. 1423–1441.
- [13] A. GERSHO AND R.M. GRAY, *Vector Quantization and Signal Compression*, Springer, New York, 1992.
- [14] R.M. GRAY AND D.L. NEUHOFF, *Quantization*, IEEE Trans. Inform. Theory, 44 (1998), pp. 2325–2382.
- [15] P.M. GRUBER, *Convex and Discrete Geometry*, Springer, New York, 2007.
- [16] J.A. HARTIGAN, *Clustering Algorithms*, Wiley, New York, 1975.
- [17] J.C. HATELEY, H. WEI, AND L. CHEN, *Fast methods for computing centroidal Voronoi tessellations*, J. Sci. Comput., 63 (2015), pp. 185–212.
- [18] R. JORDAN, D. KINDERLEHRER, AND F. OTTO, *The variational formulation of the Fokker–Planck equation*, SIAM J. Math. Anal., 29 (1998), pp. 1–17.
- [19] P.J.J. KOK AND F.N.M. KORVER, *Modelling of complex microstructures in multi phase steels: Geometric considerations for building an RVE*, in Proceedings of the 10th International Conference on Computational Plasticity, Barcelona, Spain, 2009.
- [20] L.J. LARSSON, R. CHOKSI, AND J.-C. NAVE, *An iterative algorithm for computing measures of generalized Voronoi regions*, SIAM J. Sci. Comput., 36 (2014), pp. A792–A827.
- [21] Y. LIU, W. WANG, B. LEVY, F. SUN, D.-M. YAN, L. LU, AND C. YANG, *On centroidal Voronoi tessellation—energy smoothness and fast computation*, ACM Trans. Graphics, 28 (2009), 101.
- [22] S.P. LLOYD, *Least squares quantization in PCM*, IEEE Trans. Inform. Theory, 28 (1982), pp. 129–137.
- [23] D.G. LUENBERGER AND Y. YE, *Linear and Nonlinear Programming*, 3rd ed., Springer, New York, 2008.
- [24] J.B. MACQUEEN, *Some methods for classification and analysis of multivariate observations*, in Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability, 1, University of California Press, Berkeley, CA, 1967, pp. 281–297.
- [25] Q. MÉRIGOT, *A multiscale approach to optimal transport*, Comput. Graph. Forum, 30 (2011), pp. 1583–1592.
- [26] B. MOHAR, *Graph Laplacians*, in Topics in Algebraic Graph Theory, Cambridge, UK, 2004, pp. 113–136.
- [27] A. OKABE, B. BOOTS, K. SUGIHARA, AND S.N. CHIU, *Spatial Tessellations. Concepts and Applications of Voronoi Diagrams*, 2nd ed., Wiley, New York, 2000.
- [28] J. SABIN AND R. GRAY, *Global convergence and empirical consistency of the generalized Lloyd algorithm*, IEEE Trans. Inform. Theory, 32 (1986), pp. 148–155.
- [29] M. THORPE, F. THEIL, A.M. JOHANSEN, AND N. CADE, *Convergence of the k-means minimization problem using Γ -convergence*, preprint, <http://arxiv.org/abs/1501.01320>, 2015.
- [30] C. VILLANI, *Topics in Optimal Transportation*, AMS, Providence, RI, 2003.