

A SIMPLE GEOMETRIC METHOD FOR NAVIGATING THE ENERGY LANDSCAPE OF CENTROIDAL VORONOI TESSELLATIONS*

IVAN GONZALEZ[†], RUSTUM CHOKSI[†], AND JEAN-CHRISTOPHE NAVE[†]

Abstract. Finding optimal (or low energy) centroidal Voronoi tessellations (CVTs) on a two-dimensional domain is a challenging problem. One must navigate an energy landscape whose desirable critical points have sufficiently small basins of attractions that are inaccessible with Monte Carlo initialized gradient descent methods. We present a simple deterministic method for efficiently navigating the energy landscape in order to *access* these low energy CVTs. The method has two parameters and is based upon each generator moving away from the closest neighbor by a certain distance. We give a statistical analysis of the performance of this hybrid method comparing with the results of a large number of runs for both Lloyd’s method and state-of-the-art quasi-Newton methods. Stochastic alternatives are also considered.

Key words. centroidal Voronoi tessellation, global optimization, energy ground state, optimal vector quantizer, regularity measures

AMS subject classifications. 49Q10, 65D18, 68U05, 82B80

DOI. 10.1137/20M1335534

1. Introduction. A fundamental problem in information theory and discrete geometry is known, respectively, as optimal quantization and optimal centroidal Voronoi tessellations (CVT); see [14, 15, 20, 40] for an overview of the many concrete applications. Let us present the problem in its simplest form where the underlying density is assumed to be uniform. Consider a bounded $\Omega \subset \mathbb{R}^d$ with a collection of N distinct points $\mathbf{X} := \{x_i\}_{i=1}^N \subset \Omega$ referred to as *generators*. The collection of points \mathbf{X} gives rise to a Voronoi tessellation $\mathcal{V}(\mathbf{X}) = \{V_i\}_{i=1}^N$ of Ω , where

$$(1.1) \quad V_i(\mathbf{X}) = \{y \in \Omega \mid \|x_i - y\| < \|x_j - y\| \ \forall j \neq i\}, \quad i = 1, \dots, N.$$

In other words, Voronoi cell V_i contains the points of Ω closer to x_i than to any other generator. We define for any \mathbf{X} the nonlocal energy

$$(1.2) \quad F(\mathbf{X}) = \int_{\Omega} \text{dist}^2(y, \mathbf{X}) \, dy = \sum_{i=1}^N \int_{V_i} \|y - x_i\|^2 \, dy =: \sum_{i=1}^N F_i,$$

and minimize F over the possible positions \mathbf{X} of the generators. As seen in section 2, criticality of this energy gives rise to a CVT; that is, a placement of the generators $\{x_i\}_{i=1}^N$ such that they are exactly the centroids of their associated Voronoi cell V_i , i.e.,

$$x_i = c_i := \frac{1}{|V_i|} \int_{V_i} y \, dy, \quad i = 1, \dots, N.$$

In the context of information theory and particularly in vector quantization [20], the set \mathbf{X} is viewed as a *quantizer* that discretely models data; here, uniformly distributed

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section May 4, 2020; accepted for publication (in revised form) January 5, 2021; published electronically April 29, 2021.

<https://doi.org/10.1137/20M1335534>

[†]Department of Mathematics and Statistics, McGill University, Montreal, QC, H3A 0B9 Canada (ivan.gonzalez@mail.mcgill.ca, rustum.choksi@mcgill.ca, jean-christophe.nave@mcgill.ca).

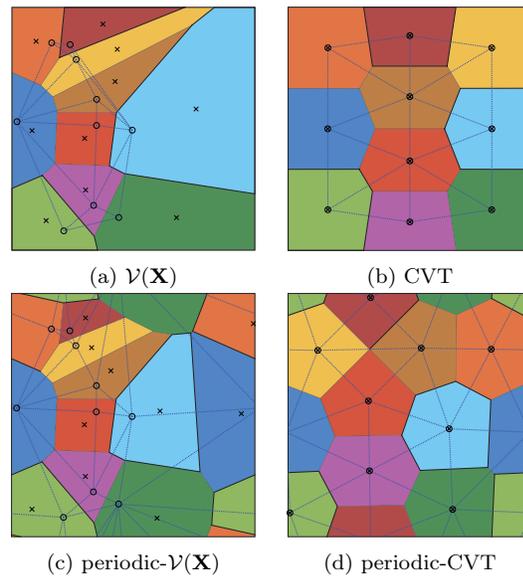


FIG. 1. Two examples of the two-dimensional framework for Voronoi tessellations with $N = 10$. The top row considers a bounded square Ω while the bottom row shows Ω as a square flat torus. Sets of generators $\mathbf{X} := \{x_i\}_{i=1}^N$ are marked as “o,” and centroids c_i of the corresponding cells V_i are marked as “x.” (a) Generic collection \mathbf{X} and associated Voronoi tessellation. (b) A centroidal Voronoi tessellation, i.e., generators x_i and respective centroids c_i coincide for all i . (c) The same sampling of generators found in (a) to emphasize the changes in centroids and in the connectivity of the tessellation when on the torus. (d) A periodic centroidal Voronoi tessellation (PCVT). The dual graphs of the boundary sets $\cup_{i \leq N} \partial V_i$ (i.e., the Delaunay triangulations) are shown in dotted lines. Notice the superior regularity of the tessellation and its corresponding triangulation in both centroidal cases.

over Ω . The *quantization error* is given by $F(\mathbf{X})$ with the *optimal quantizer* being the minimizer \mathbf{X}^* with the least error (alternatively, *the CVT with lowest energy*). Figure 1 contains some visual examples to illustrate these notions.

The energy (1.2) has a wealth of critical points (CVTs) [33, 46], and *low energy* CVTs have tiny basins of attractions making them difficult or impossible to find via gradient based descent with random initializations. Finding optimal or at least low energy CVTs with desirable geometric properties is of fundamental importance in many concrete applications; cf. [14, 15, 20], for example: (1) spatial optimization, (2) mesh generation and numerical analysis, and (3) vector quantization. To this end, one does have a benchmark for the optimal geometry in the limit $N \rightarrow \infty$ wherein we dispense with shape and boundary effects. Indeed, Gershgorin’s conjecture [19] addresses the periodic structure of the optimal quantizer as $N \rightarrow \infty$. The conjecture is completely solved in two dimensions wherein the optimal Voronoi cell is the regular hexagon, corresponding to generators on a triangular lattice. However, to date it remains open in three dimensions (3D) wherein the belief is that the optimal Voronoi *period-cell* is the truncated octahedron, corresponding to generators on a body centered cubic (BCC) lattice; see [2, 9, 16].

The purpose of this paper is to present and assess in two dimensions (2D) a simple deterministic method for efficiently navigating the energy landscape in order to *access* low energy CVTs which are otherwise inaccessible with Monte Carlo initializations coupled to gradient based descent methods (i.e., taking optimal results among

hundreds of thousands to millions of gradient based descents on randomly sampled initial configurations). The proposed method has two parameters, namely the *preconditioning number* K and the *probing number* Q , and is based upon a very simple scheme: each generator *moving away from the closest neighbor (MACN)* by a certain distance (displacement); i.e., the basic *MACN* scheme moves each generator in the opposite direction from its closest neighbor in \mathbf{X} . For this scheme we consider two choices of displacement:

- The individual displacement $\|x_i - c_i\|$ of x_i to the centroid of its Voronoi cell. Our scheme using this displacement is labelled as the *MACN-c* step.
- A fixed displacement δ for all generators where

$$(1.3) \quad \delta := \frac{1}{4} \sqrt{\frac{|\Omega|}{N}}$$

is set in terms of the average distance between generators in the approximate regular hexagonal lattice. Our scheme with this distance is labelled as the *MACN- δ* step.

Our method is then a *coupling* procedure repeating three steps. Starting from a random initialization (placement of generators in Ω) we have the following:

- Step 1. Iterate the *MACN-c* step K times as a *preconditioning*.
- Step 2. Implement Lloyd's, or any other deterministic method surveyed in section 3, minimizing (1.2) to a CVT.
- Step 3. A single *MACN- δ* step working as a geometric based *annealing* to disrupt the CVT from Step 2.
- Step 4. Repeat Steps 1–3 a total number of Q times, ending at Step 2.

Step 1 results in a configuration which is *close* to a CVT. However, this is not the point: it results in a configuration which lies in the basin of an *energetically desirable* CVT. By *energetically desirable* we mean two things: (i) it has low energy in the sense that its energy is comparable with the optimal result of any standard gradient based descent algorithm assessed over a “large” number of runs (for us with $N \sim 1000$, a large number of runs is on the order of 100000). (ii) The same holds true for the measures of regularity described in section 2. Step 2 achieves this *energetically desirable* CVT. Step 3 breaks away from this basin to another basin which can contain a more optimal CVT. In the end, one run of our method will have probed a number Q of potentially different basins of attraction and their corresponding CVTs.

The choice of δ in (1.3) is subtle: it is large enough in order to change basins but sufficiently small in order to not lose the desired regularity achieved thus far. In two dimensions, the optimal configuration is partial to N regular hexagonal Voronoi cells. In this case, δ can be thought of as half the distance from the generator to its Voronoi cell boundary plane. Indeed, in the case of regular hexagonal Voronoi cells, the six neighboring generators are equally distant. In this scenario, Step 3 chooses one of the six closest generators according to some pre-established tie breaking rule.

The probing number Q encapsulates a degree of freedom common to all coupling methods (see section 3). The preconditioning number K is chosen to remain constant over the Q stages. This simple choice of constant K is shown to produce excellent results on 2D domains. Moreover, preliminary implementation with a constant K has also produced similar results on the sphere. The adaptation of the *MACN* method

to the sphere and possibly other manifolds is, however, outside the scope of the present paper and will be addressed in subsequent work. On the other hand, one could devise several strategies—with higher degrees of freedom—for a generalized sequence $\{K_q\}_{q=0}^{Q-1}$; for example, taking an initial value K_0 combined with a suitable decaying profile adapted to both present and past history of the energy along the dynamics. Such variants could potentially improve convergence and energy efficiency of the method. It is, however, difficult to find consistent strategies that will systematically render desirable results across different problem sizes. Our crude tuning using a constant preconditioning number does render desirable results across different domains and problem sizes (cf. section 5).

We implement and assess our method on two choices of 2D domains, both having periodic boundary conditions in order to dispense with boundary effects. To start, we work on a primary domain for which the ground state is known: the periodic regular hexagon, which can be tessellated into N regular hexagons provided N is suitably chosen [10]. We then work on a primary domain that does not permit a perfect regular tiling: the flat square torus. Here there is always frustration due to the size effects, and it is surprisingly unclear as to the true nature of the lowest energy state. With the number of generators N taken from the range 1000 to 4000, we show that our hybrid method with $Q \leq 10$, implemented with less than two dozens initializations, readily finds states with far lower energy (and other metrics of optimal regularity) than the ones accessible with Lloyd's method, or any state-of-the-art gradient based descent method, assessed over hundreds of thousands of random initializations; see Figure 2. In sections 5 and 6 we present the full details of this comparison, emphasizing the role of statistics for assessing ours and other numerical methods.

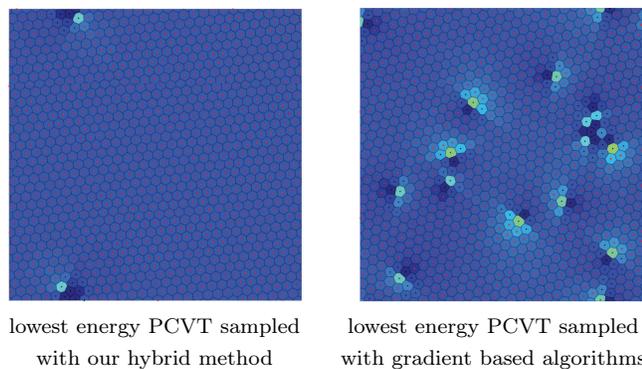


FIG. 2. Case $N = 1000$ on the square torus; comparative performance of our hybrid method with $Q = 10$ (left) versus the lowest energy configuration obtained during our statistical sampling of the landscape with selected gradient based methods (right). The PCVT on the left is energetically 548% closer to the nonachievable regular hexagonal lattice than the one on the right. Full statistical details on this comparison is found in Figure 9 and Table 3. The color of each V_i scales with F_i to indicate local energy contribution. The reader is referred to section 2 and subsequent figures for further details on the color map.

The motivation and scope of our hybrid method is twofold. First, there is the direct goal of generating low energy CVTs which has an impact in many applications such as the ones already presented above. In contrast to methods which are based upon initial sampling or building on regular hexagons, our method is based only on basic structures of the energy (centroids and distance functions). Moreover, on

domains like the square torus, optimality is subtle with regard to the presence of nonregular hexagons and defects (nonhexagonal cells): indeed, there is probably not even one perfectly regular hexagon in the Voronoi diagram of the optimal energy configuration.

There is a second scope to our work. Probing nonconvex and nonlocal energy landscapes is a fundamental problem in physics and applied mathematics. Deterministic strategies for navigating such landscapes are far and few between. The CVT energy (1.2) is perhaps the simplest nontrivial example of such a landscape; while it is finite dimensional with a simple geometric characterization of criticality (namely a CVT), it is challenging (even in two dimensions) to navigate the landscape of CVTs. As such, it perfects an ideal problem to address deterministic navigation. Our results demonstrate that while the CVT energy landscape on the square torus with $N \sim 1000$ is indeed complex, our hybrid algorithm is able to efficiently navigate it with only 9–10 deterministic “annealing” steps.

We emphasize that our *MACN* algorithm is designed for navigating the CVT energy landscape. We make no claim that our algorithm has any direct applications for navigating general nonconvex energy landscapes. In fact, even for the closest energetic models of crystallographic particle interactions [3], it is unclear whether or not our work has any application.

In light of our two scopes we would like to point out that there is another method able to successfully probe the landscape, namely the global Monte Carlo method from [33] (cf. section 3). However, the novelty and change of paradigm here is that our method, motivated entirely by geometry, probes the energy in a completely deterministic way. Moreover, our method shows a probing improvement compared to the global Monte Carlo. Precisely, the number of computed CVTs needed to reach low energy configurations is lower with the *MACN* algorithm than with this alternative method. We elaborate on this comparison in section 6.

Finally let us remark on the empirical nature of our work. While we give a heuristic rationale for our *MACN* steps, the precise nature of the distance chosen is based in part on empirical tries. For the *MACN*- δ step, we experimented with other choices of distance; for example, the intrinsic length-scale of the Voronoi cell (cf. section 6). Our choice is one with sensibly the most effective performance. Overall, while we only have heuristics to explain certain aspects of our method, we feel the remarkable results justify its presentation and discussion here.

This paper is organized as follows: in section 2 we give a brief description of periodic centroidal Voronoi tessellations (PCVTs). We also discuss, besides the normalized energy, certain natural *local measures of regularity*, one of which is novel. Then, in section 3 we survey methods to generate and improve CVTs. We present our hybrid algorithm in section 4 and then an analysis for its predictions in section 5. Later, in section 6 we discuss the advantages and disadvantages of our method compared to stochastic alternatives and then finish with closing remarks and future directions in section 7.

2. Periodic CVTs and regularity measures. For the remainder of this paper we will consider 2D torii spaces, namely polygons Ω with opposite and pairwise identifiable sides that can be periodically extended to the plane. With the enforcement of these boundary conditions we have the definitions of a periodic-VT (PVT) and periodic-CVT (PCVT) by using the metric $\|\cdot\|_{\mathbb{T}}$ that is inherited from the Euclidian metric in the above definitions (1.1) and (1.2), respectively. See Figure 1 (c) and (d) for a concrete example. In addition, the dual graph of the set $\cup_{i \leq N} \partial V_i$ known as the

periodic-Delaunay triangulation (PDT) will be of importance. The PDT has an edge e_{ij} linking x_i and x_j iff they are neighbors in the PVT, in other words the index sets $\mathcal{N}_i := \{j \neq i \mid \partial V_j \cap \partial V_i \neq \emptyset\}$, $i = 1, \dots, N$, are intrinsic to the triangulation. We refer the reader to [40] for an introductory treatment on Delaunay triangulations and other variants of the Voronoi tessellation and to [46, 45, 6] for a detailed definition of the periodic framework.

For further details on a popularized use of the Voronoi/Delaunay duality, see, for example, [1, 8, 16] where *optimal triangulations*, used in mesh generation and numerical analysis, are defined as the duals of optimal CVTs as opposed to the result of directly optimizing triangulations over prescribed criteria. See also [7] for additional discussion on the topic, this time concerning the nonachievable Delaunay triangulation made of equilateral triangles in three dimensions (that, therefore, is not the dual of an achievable ground state CVT).

Centroidal characterizations. In addition to the geometrical property $x_i = c_i \forall i$, a PCVT admits a variational characterization via $F : \mathcal{D} \rightarrow \mathbb{R}_+$ given by (1.2) here $\mathcal{D} := \{\mathbf{X} \in \mathbb{R}^{2N} \mid x_i \neq x_j \forall j \neq i; x_i \in \Omega \forall i\}$ represents the set of nondegenerate configurations of generators in our periodic space. Depending on the application, the functional $F = \sum_{i=1}^N F_i$ is often referred to as the distortion, cost function, or potential. Alternatively, one can interpret each F_i from a physical point of view as the trace of the 2×2 inertia tensor of V_i with respect to x_i , i.e., the resistance of V_i to its rotation around an axis passing through x_i that is orthogonal to Ω .

It is a well-established fact in rigid mechanics (parallel axis theorem) that the trace of this tensor will be minimized whenever the orthogonal axis of rotation passes through the centroid. This locally translates to a Voronoi region being generated by its centroid and, therefore, globally translates to a PCVT. Indeed, one can formally prove this: by the definition of $\mathcal{V}(\mathbf{X})$ with the Euclidian metric we have according to [26, 12] that $\frac{\partial F}{\partial \mathbf{V}}(\mathbf{X}, \mathcal{V}(\mathbf{X})) \equiv 0$ and thus the Reynolds transport theorem gives the gradient components

$$D_i F(\mathbf{X}) := \frac{\partial F}{\partial x_i}(\mathbf{X}) = 2|V_i|(x_i - c_i), \quad i = 1, \dots, N,$$

which show that the gradient vanishing configurations of F are exactly PCVTs. Moreover, a nontrivial second application of Reynolds theorem yields the entries of the Hessian $D^2 F(\mathbf{X})$

$$(2.1) \quad \begin{cases} \frac{\partial^2 F}{\partial x_i^{(m)} \partial x_i^{(m)}} = 2|V_i| - \sum_{j \in \mathcal{N}_i} \frac{2}{\|x_i - x_j\|} \int_{\partial V_i \cap \partial V_j} (x_i^{(m)} - y^{(m)})^2 dy, \\ \frac{\partial^2 F}{\partial x_i^{(m)} \partial x_i^{(l)}} = - \sum_{j \in \mathcal{N}_i} \frac{2}{\|x_i - x_j\|} \int_{\partial V_i \cap \partial V_j} (x_i^{(m)} - y^{(m)})(x_i^{(l)} - y^{(l)}) dy, & m \neq l, \\ \frac{\partial^2 F}{\partial x_i^{(m)} \partial x_j^{(l)}} = \frac{2}{\|x_i - x_j\|} \int_{\partial V_i \cap \partial V_j} (x_i^{(m)} - y^{(m)})(x_j^{(l)} - y^{(l)}) dy, & j \in \mathcal{N}_i, \\ \frac{\partial^2 F}{\partial x_i^{(m)} \partial x_j^{(l)}} \equiv 0, & j \neq i, j \notin \mathcal{N}_i. \end{cases}$$

Furthermore, F has been shown to be $C^2(\mathcal{D})$ [31, 46], and although the Hessian is sparse, it is, in general, not definite. As noted in [33], this is a consequence of the energy functional being highly nonconvex with the presence of saddle points; see also [46].

Optimal PCVT, ground state energy, and regularity measures. Based upon the *hexagon theorem* [21, 37, 43], and in a similar spirit to [11, 44], we scale the

energy functional with respect to $N F_{hex}$. Here F_{hex} denotes the second moment of a regular hexagon with area $|\Omega|/N$, and a simple calculation shows that

$$F_{hex} = \frac{5}{18\sqrt{3}} \frac{|\Omega|^2}{N^2}.$$

The scaled energy that is thus independent of the size of the problem is

$$E(\mathbf{X}) := \frac{F(\mathbf{X})}{N F_{hex}} = \frac{1}{N} \sum_{i=1}^N \frac{F_i}{F_{hex}} =: \frac{1}{N} \sum_{i=1}^N E_i,$$

with the same scaling carrying out trivially to $DE(\mathbf{X})$ and $D^2E(\mathbf{X})$.

As an additional way to help us quantify the quality of a PCVT, and as it has been customary in the literature, we use the fraction of hexagonal cells:

$$H(\mathbf{X}) := \frac{\#\{V_i, i = 1, \dots, N \mid \#\mathcal{N}_i = 6\}}{N}.$$

However, regularity measures such as H or the Voronoi entropy used in [5] (and rediscussed in [4]) only give us information about the connectivity of the PDT and make no direct connection with the hexagon theorem requiring hexagons to be regular to achieve the energy $E = 1$. For this reason we create a refinement on H that takes into account the regularity of the hexagonal cells through their isoperimetric ratio. We recall that the isoperimetric ratio r of a polygon is the dimensionless ratio of its perimeter squared and area. Using $r(V_i)$ of each Voronoi cell together with the isoperimetric ratio of the regular hexagon $r_{hex} := 8\sqrt{3}$, and for given $\epsilon > 0$ small we introduce the following regularity measure of a tessellation $\mathcal{V}(\mathbf{X})$:

$$R_\epsilon(\mathbf{X}) := \frac{\#\{V_i, i = 1, \dots, N \mid \#\mathcal{N}_i = 6; |1 - \frac{r(V_i)}{r_{hex}}| \leq \epsilon\}}{N}.$$

That is, R_ϵ is the fraction of cells that are regular hexagons within an isoperimetric tolerance of ϵ . Generators whose cell is taken into account by R_ϵ will be depicted in red in the PVTs of Figure 3 and onwards.

In sections 4–6 we will rely on $E - 1$, H , and R_ϵ to measure the performance of our method and others in reaching high quality PCVTs. The data will also make the compelling case that R_ϵ is a more sensitive measure of low values of E than H (for appropriate ϵ). Finally, the statistics on these three quantities will be at the core of the conjecture that our method acts at a semiglobal scale.

3. Generating and improving CVTs. We divide this brief literature survey into methods available to compute CVTs and ways to enhance their quality.

Computing CVTs. These methods either rely on the characterization $x_i = c_i \forall i$ and seek to solve these nonlinear equations or are of variational character on $E(\mathbf{X})$:

1. Lloyd’s method, introduced in the seminal work [32], is unquestionably the most widespread method due to its simplicity. It iteratively applies the map $\mathbf{T} : \mathcal{D} \rightarrow \mathcal{D}$ defined by $\mathbf{T}(\mathbf{X}) = \mathbf{C}$, where $\mathbf{C} := \{c_i\}_{i=1}^N$ is the collection of centroids of $\mathcal{V}(\mathbf{X})$. We refer the reader to [14, 13, 18] for properties and analysis of the map.

Due to the importance of the method in the remainder of this paper we provide explicit pseudo-code as a subroutine in Algorithm 4.1.

2. McQueen's and probabilistic Lloyd's methods [28] generate small sets of random sampling points in Ω and use these to approximate the centroid of Voronoi regions.

3. Lloyd–Newton [12] uses Newton's root finding method on $\mathbf{S}(\mathbf{X}) := \mathbf{X} - \mathbf{T}(\mathbf{X})$ after some Lloyd steps to reach the vicinity of a root in \mathcal{D} . The authors also propose multilevel based extensions: one using an algebraic multigrid preconditioner combined with Gauss–Seidel iterations to solve the linear system involved in each Newton step and another using a nonlinear multigrid approach in solving $DE = 0$. See also [17] for further development. However, as discussed in [31], unstable CVTs may be produced (i.e., saddle points of E).

4. Newton's classical technique minimizes Hessian based quadratic models of $E \in C^2(\mathcal{D})$. With (2.1) available, the method converges at least quadratically when coupled with a line-search ensuring the strong Wolfe conditions [39]. Nonetheless, this method suffers from two downsides: (i) D^2E is often indefinite and needs to be altered, for example, by adding a “small” matrix to render it SPD prior to executing an incomplete Cholesky factorization; see [34, 30] for theory and low memory algorithms. (ii) The Hessian is expensive to populate due to the boundary integrals.

5. Quasi-Newton BFGS collection. These methods only use E and DE to give an iterative approximation of the inverse Hessian. They remain to this day the favored methods in the literature for fast CVT/PCVT computation due to their expected superlinear convergence whenever they are coupled with a line-search method that ensures the strong Wolfe conditions. The two families suited for medium/large scale problems that will be used in this paper are the following:

- (a) the low memory L-BFGS(M) in which the inverse Hessian approximation remains sparse and is computed recursively from the M previous approximations;
- (b) the Preconditioned-L-BFGS(M, T) (P-L-BFGS(M, T)) uses, every modulo T iterations, an SPD preconditioner matrix \tilde{A} (that does not necessarily need to approximate the Hessian) with the goal of redirecting the algorithm to a more suitable descent direction.

See [39] for a thorough description and analysis of the classical BFGS and L-BFGS(M) and [27, 31] for the use of P-L-BFGS(M, T). Finally, an explicit routine of the preconditioned algorithm is provided in Appendix A.

6. Non-Linear Conjugate Gradient (NLCG) methods generalize the classical CG used in quadratic programming. Several updates for the conjugate directions are available [22, 39]. One can also use a relevant preconditioner SPD matrix \tilde{A} to improve the descent-conjugate directions.

There are several preconditioner matrices \tilde{A} , for example: the Hessian D^2E itself (along with the often necessary modification to make it SPD) or a Graph Laplacian G introduced in [25] whose purpose is to approximate $DE(\mathbf{X}) = 0$ to first order by the matrix equation $G\mathbf{X} = 0$.

Originally presented for the Ω -bounded case, below is our adaptation of G for the periodic case; denoting by $\tau_{i,j}$ the pyramid with base $\partial V_j \cap \partial V_i$ and apex x_i , the $N \times N$ matrix is then given by

$$(3.1) \quad G := \begin{cases} g_{ij} = -\int_{\tau_{i,j} \cup \tau_{j,i}} \rho(y) dy & \text{if } j \in \mathcal{N}_i, \\ g_{ii} = \sum_{j \in \mathcal{N}_i} |g_{ij}|, & \\ 0 & \text{otherwise.} \end{cases}$$

However, contrary to the original construction for Ω -bounded, our adaptation is sym-

metric and positive semidefinite (G given by (3.1) is not strictly diagonal dominant and can thus be singular). As a consequence we will need to use a modified Cholesky factorization of G .

Energetic improvements. Next, we survey initializations and other methods used jointly with the above algorithms in order to improve the quality of CVTs.

Initializations.

(a) A Greedy Edge-Collapsing initialization [36] that meshes Ω using more than N vertices; then it repeatedly uses an edge-collapsing scheme, and finally the decimated vertices are employed as the initial generators.

(b) Quasi-random samplings of Ω reduce the discrepancy of the initial cloud of generators, i.e., the sampling of each site depends on the position of the others. These QR samplings use low discrepancy sequences such as those of Halton, Hammersley, Niederreiter, and Sobol [38, 24, 23]; see also [41] for CVT results.

Couplings

(a) A hierarchical method [44] that refines a CVT by cleverly inserting new generators over the Delaunay triangulation. In this way the “regular” portions of the CVT “grow” when alternating with an energy descent method.

(b) An Atomic Operation method [35]. Here the authors establish three operations on “defects” that merge or split nonhexagonal cells prior to minimizing the energy; the process is then repeated.

(c) A global Monte Carlo method [33] which applies ideas from Simulated Annealing. This method starts at a CVT and after a specific random perturbation of generators, which is dependent on the size of each V_i , a new CVT is obtained by a nonlinear minimization method. If the new energy is lower than the previous one, then the algorithm automatically accepts that new configuration; otherwise it accepts it only according to a transitional probability that is dependent both on a cooling temperature and on the energy gap between the two CVTs.

While the method has been proven to find the ground state in infinite time [29], a crucial disadvantage in practice is the number of parameters that need to be adjusted to obtain good performances, namely, the initial temperature, the temperature decay to zero, the perturbation amplitude, and the number of iterations repeating the procedure.

Indeed, the above-mentioned literature shows that these initializations and couplings yield significant energetic improvements, even more so when combined.

4. Our hybrid algorithm. This section describes in more detail our three stage method and subsequently provides an individual and deeper insight on the $MACN$ - c and $MACN$ - δ dynamics, respectively.

Given a set of generators $\mathbf{X} := \{x_i\}_{i=1}^N$ and its associated $\mathcal{V}(\mathbf{X})$, a closest neighbor to x_i is denoted $x_{j_i^*}$; i.e., $x_{j_i^*}$ solves $\min_{j \neq i} \|x_i - x_j\| = \min_{j \in \mathcal{N}_i} \|x_i - x_j\|$, where \mathcal{N}_i was defined earlier to be the index set of Delaunay edge-connected neighbors to the site x_i . Notice that $x_{j_i^*}$ may not be unique. However, it suffices for our implementation to simply pick one solution via a pre-established tie breaking rule (e.g., selection of the closest candidate up to machine precision).

Our $MACN$ scheme is a displacement of x_i by a distance d_i in the opposite direction to $x_{j_i^*}$, i.e., a $MACN$ iteration with distances $\{d_i\}_{i=1}^N$ consists of the update

scheme

$$(4.1) \quad x_i \leftarrow x_i + d_i \frac{x_i - x_{j_i^*}}{\|x_i - x_{j_i^*}\|}, \quad i = 1, \dots, N.$$

Here d_i is a dummy quantity representing either the distance to the centroid $d_i := \|x_i - c_i\|$ for a *MACN-c* step or $d_i := \delta = \frac{1}{4} \sqrt{\frac{|\Omega|}{N}}$ for a *MACN- δ* step—as in (1.3).

Explicit pseudo-code of the overall method comprising the four steps listed in section 1 is provided in Algorithm 4.1 for completeness.

Algorithm 4.1 *MACN* algorithm over Q stages with Lloyd subroutine.

Input: (1) initial generators $\mathbf{X} = \{x_i\}_{i=1}^N$; (2) probing number Q ; (3) preconditioning number K ; (4) tol for convergence to a PCVT

for $q = 0 : Q - 1$ **do**

I. *MACN-c*:

for $k = 0 : K - 1$ **do**

1. compute $\text{PDT}(\mathbf{X})$, associated $\text{PVT}(\mathbf{X})$ and centroids $\{c_i\}_{i=1}^N$
2. extract $\{\mathcal{N}_i\}_{i=1}^N$ and find $\{j_i^*\}_{i=1}^N$ solving $\min_{j \in \mathcal{N}_i} \|x_i - x_j\|$ for each i
3. set $d_i = \|x_i - c_i\|$ and update \mathbf{X} : $x_i \leftarrow x_i \forall i$ with (4.1)

end for

II. *Reaching criticality*:

$[\mathbf{X}_q^*, \text{PDT}(\mathbf{X}_q^*), \text{PCVT}(\mathbf{X}_q^*)] = \text{Lloyd}(\mathbf{X}, tol)$

III. *MACN- δ* :

if $q < Q - 1$ **then**

1. extract $\{\mathcal{N}_i\}_{i=1}^N$ from $\text{PDT}(\mathbf{X}_q^*)$ and find $\{j_i^*\}_{i=1}^N$ solving $\min_{j \in \mathcal{N}_i} \|x_i - x_j\|$ for each i
2. set $d_i = \delta$ and get new \mathbf{X} by δ -perturbing \mathbf{X}_q^* : $x_i \leftarrow x_i \forall i$ with (4.1)

end if

end for

Output: $\{\mathbf{X}_q^*\}_{q=0}^{Q-1}$, a collection of stable local minimizers of E and their corresponding PCVTs and PDTs

subroutine $[\mathbf{X}, \text{PDT}(\mathbf{X}), \mathcal{V}(\mathbf{X})] = \text{Lloyd}(\mathbf{X}, tol)$

set $diff = \text{Inf}$

while $diff > tol$ **do**

1. compute $\mathcal{V}(\mathbf{X})$ (i.e., $\text{PVT}(\mathbf{X})$), centroids $\{c_i\}_{i=1}^N$, areas $\{|V_i|\}_{i=1}^N$ and set $diff = \|DE\|/N$
2. update \mathbf{X} : $x_i \leftarrow c_i \forall i$

end while

end subroutine

As briefly mentioned in the introduction, the K initial *MACN-c* iterations yield more evenly distributed sets of generators. This preconditioning is common in outcome to other initializations listed in section 3. Thus the success of *MACN-c*+Lloyd in reaching “low” energy PCVTs is not surprising. Then by disrupting the centroidal configuration with *MACN- δ* and introducing the $Q - 1$ supplementary repetitions we create a coupling technique involving a symbiosis of relaxation and contraction. By this we mean that the combination of our preconditioning and perturbation helps Lloyd’s algorithm in reaching lower energy states while, conversely, Lloyd’s method gets the system to PCVTs that our *MACN* stages use as a stepping-stone to successfully probe the landscape.

We emphasize next that the sequence $\{E(\mathbf{X}_q^*)\}_{q=0}^{Q-1}$ obtained with our method may not be strictly decreasing. It is, in fact, quite likely that our algorithm moves to a higher energy basin of attractions from one stage to another, thus resembling Simulated Annealing in that sense (while remaining completely deterministic). The advantage, however, is, again, the low probing number needed. With $Q \leq 10$ we are able to sample low energy states that are impossible or scarcely achievable by other deterministic methods (even when using a stochastic sampling of the landscape done with a great number of initial configurations); cf. section 5.

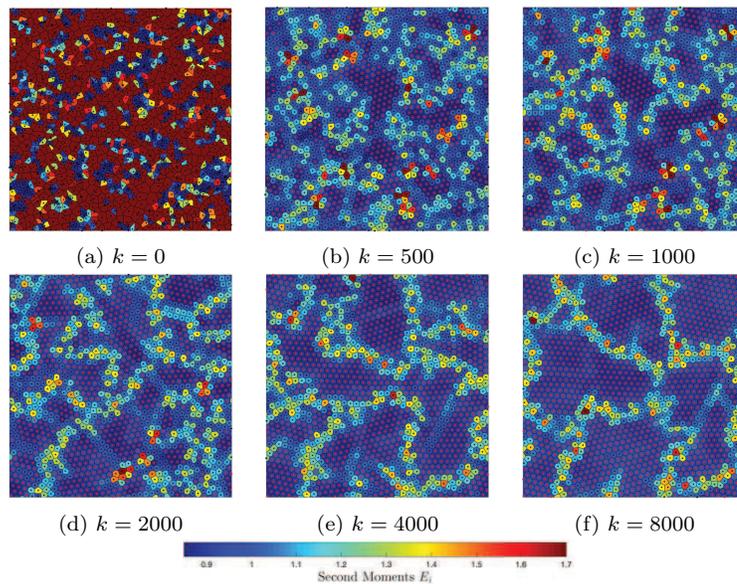
Concerning the complexity of our method, most of the software used to construct VT/PVT with Euclidean distance rely on an early construction of the Delaunay triangulation, e.g., the CGAL (Computational Geometry Algorithms Library)—<https://www.cgal.org> [42]—for C++ or the built-in MATLAB function `voronoin`. Thus one can extract $\mathcal{N}_i \forall i$ before the computation of the tessellation while keeping the same complexity. Moreover, the remaining difference between one *MACN-c* step and one Lloyd step is the computation of the indices $j_i^* \forall i$ which simply adds a lower order term $O(N)$ to the optimal overall complexity $O(N \ln(N))$ in two dimensions [40]. Hence, our overall scheme (4.1) benefits from the same low complexity of Lloyd’s algorithm as well as a comparable simple implementation. This is the reason we chose Lloyd’s method for our coupling rather than any other gradient based method.

MACN-c. We now focus on *MACN-c* as a stand-alone initialization in order to gain first insight on the functionality of this iterative scheme. With the concise example presented in Figure 3 we show the general long-term behavior of these dynamics over iterations $k = 0, \dots, K - 1$ with $K = 7 \times 10^4$.

The first immediate observation is the reduction of the discrepancy of the original cloud of points \mathbf{X} . Indeed, when applying *MACN-c* we obtain a significant decrease in energy, often close to two orders of magnitude when compared to a general random sampling. It is worth noting, however, that this scheme alone is not contractive due to the abrupt topological changes in the PDT from one iteration to another and thus will most likely fail to provide a PCVT by itself, even in the limit $K \rightarrow \infty$.

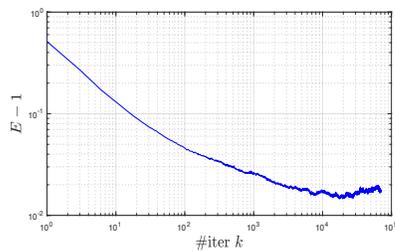
The second observation is that the system eventually reaches a “low” energy plateau and oscillates around it (see Figure 3 (h)). This suggests that in order to maximize energetic performance, the subsequent coupling with Lloyd’s algorithm should be performed when the *MACN-c* dynamics attain this regime. However, it is difficult to estimate a priori when the system will reach such a mesostability and even more so the required K might be too large. Thus, in practice K (more precisely $Q \times K$) should be chosen by the user to retain tractability of the overall method we presented in Algorithm 4.1. For this reason the presentation of our numerical results in section 5 will start by a parameter sweep over values of K that yield not necessarily the best energetic results but a suitable trade-off between computation time and energy.

MACN- δ . We again emphasize the empirical nature of our choice for the distance δ (1.3) that was found among scalings of $\sqrt{|\Omega|/N}$ (the linear length-scale of the geometry in question). In particular, the a priori mysterious factor $1/4$ was chosen



k	0	500	1000	2000	4000	8000
$E - 1$	0.98007	0.02933	0.02598	0.02132	0.01875	0.01608
H (%)	29.33	77.86	81.60	84.80	87.46	89.60
R_ϵ (%)	0.00	25.60	34.13	42.13	52.26	62.86

(g) measures specific to configurations (a)–(f)



(h) energy profile showing mesostability

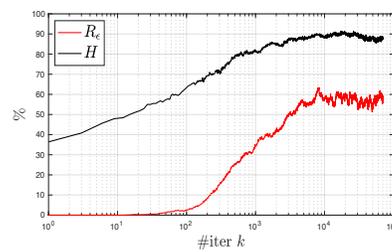
(i) R_ϵ and H profiles

FIG. 3. Example of MACN- c dynamics with $N = 1500$ in the square torus Ω . The k th iterations displayed start from the uncorrelated uniform random sampling of the domain shown in (a); the table contains the measures of the respective tessellations. The coloring represents E_i , and generators in red are those taken into account by R_ϵ ($\epsilon = 0.5\%$). Finally, we plot the three regularity measures against the iterate number up to $k = 7 \times 10^4 = K$ to show the different regimes; the system's mesostability after $k \approx 10^4$ is apparent.

to be in the small threshold of such scalings that simultaneously allows a change of basin of attraction while remaining small enough to preserve a “certain regularity” in the structure and topology of the tessellation. Perhaps the strongest ascertainment is

that δ equals the intrinsic length-scale of the regular hexagonal tessellation, namely $|V|/|\partial V|$ where V here represents the regular hexagon with inscribed circle of diameter $\sqrt{|\Omega|/N}$. More details on this latter interpretation of δ and on the global behavior of other variants of $MACN$ for PCVT dislocations are provided in section 6.

We finish the section with the example of Figure 4 showing a prototypical PCVT dislocation using $MACN-\delta$.

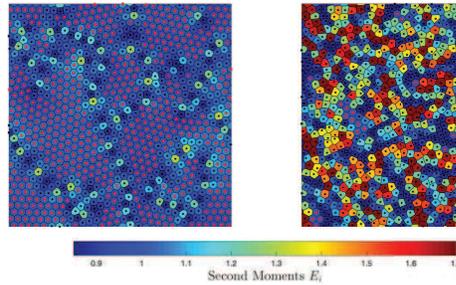


FIG. 4. Example of a $MACN-\delta$ perturbation with $N = 1000$ in a square torus Ω . On the left is a generic PCVT, and on the right is its resulting perturbation with $\delta = \frac{1}{4}\sqrt{|\Omega|/N}$. The colormap represents energy, and hexagonal cells with red generators are regular within the isoperimetric tolerance $\epsilon = 0.5\%$.

5. Numerical results. A total of six examples with different N 's are presented below: the first two are on the regular hexagonal torus while the remaining four are on the square torus. In the former the perfect honeycomb lattice is attainable iff $N = a^2 + ab + b^2 \forall a, b \in \mathbb{N}$ [10] while in the latter the ground state is unknown due to size frustration. For conciseness of the exposition we will from now on fix the probing number to $Q = 10$; this value will be large enough to demonstrate the power of our method vis-à-vis the alternatives. Our study of the six cases thus begins by finding a preconditioning number K offering a suitable trade-off between time and energy performance. Once K is chosen we deepen the analysis of the performance of our method and then finish the section with a comparison of R_ϵ and H as measures faithful to E .

Throughout this section we use uniformly uncorrelated initial configurations over Ω and our hybrid method is compared with L-BFGS(7) and P-L-BFGS(20,20). For the latter our implementation of the pseudo-code of Appendix A uses the periodic adaptation of the preconditioner matrix $\tilde{A} = G$ as described in (3.1). Additionally, we include results obtained with Lloyd's algorithm in our periodic set up to have a solid point of reference for past and future work since this is the only algorithm with a complete lack of tuning parameters.

Following collectively the results of [31, 44, 25] as well as our own implementation of some of the deterministic methods recalled in section 3, we believe that the two quasi-Newton choices of comparison paint a good overview of the current deterministic state-of-the-art methods: in particular, $\tilde{A} = D^2E$ as well as other (M, T) values were tested for the torus but did not achieve noticeable systematic improvements in the regularity measures.

We further emphasize that, while the objective and main contribution of this paper is to establish a dynamical and fully deterministic way of sampling energy basins with as few parameters as possible, we tested the global Monte Carlo method from

[33] on the torus; we report results later in section 6 where the energy outcomes are comparable to ours and we discuss the comparative advantages of both methods. Furthermore, although not adapted to the periodic boundaries, quasi-random samplings of Ω were also tested prior to quasi-Newton minimization; the energies achieved out of 1000 runs were comparable with those of the basins sampled by Lloyd's method out of the 10000 runs initialized with uncorrelated distributions that are discussed below.

Let us next introduce some notation; once K is fixed, our hybrid method and its lowest sampled energy E^{min} will primarily be compared with respect to the reference energy E^{ref} given by the minimal energy PCVT obtained from a large batch of initial configurations using the three gradient based methods. More precisely, E^{min} will be recorded with our method over 100 or 1000 initial configurations (depending on the example) while E^{ref} will be the lowest energy among 100000 runs for each of L-BFGS(7) and P-L-BFGS(20,20) as well as 10000 runs for Lloyd's algorithm.

We then define the following performance ratio for the sampled minimal energies:

$$(5.1) \quad \tau := \frac{E^{min} - 1}{E^{ref} - 1}.$$

We will also employ the empirical cumulative distribution functions (ECDFs) $f_{E-1}, f_{R_\epsilon}, f_H$ of our respective regularity measures as well as the values $f_{E-1}^*, f_{R_\epsilon}^*, f_H^*$ obtained when evaluating the ECDFs of our hybrid method at $E^{ref}, R_\epsilon^{ref}, H^{ref}$, respectively; these quantities will establish the frequency of PCVTs for which our hybrid algorithm outperforms the best comparative method. Other basic statistics provided on regularity measures include averages and standard deviations taken over the designated number of runs; we denote them by $\langle \cdot \rangle$ and $\sigma(\cdot)$, respectively. Finally, we fix the isoperimetric tolerance to be $\epsilon \equiv 0, 5\%$. Further details on this choice will be discussed at the end of this section.

A final note on our energy measurements: no quadrature was involved (i.e., exact calculations were performed), and a tolerance was used on $\|DE\|/N$ guaranteeing that the values of energy listed in all the tables and figures are accurate at least up to the significant digits provided.

Choosing K . The two scenarios on the hexagonal torus Ω (allowing the honeycomb tiling) are with $N = 973 = 17^2 + 17 \times 19 + 19^2$ and $N = 2029 = 25^2 + 25 \times 27 + 27^2$. On the square torus we'll work with the values $N = \{n \times 1000\}_{n=1}^4$. For these set ups we run our hybrid method on a reduced set of 15 initial configurations using a selected list of K 's; the energy results shown in Figure 5 will allow us to choose trade-off values between time and energy. We remark on the general decrease tendency over the Q hybrid stages but note that the decrease is not monotone in K .

Note as well that because the first two values of N on the square torus are close to those on the hexagonal torus there is no need to run sweeps for $N = 1000$ and 2000; we'll just retain the same parameter values.

The graphs of $\langle E - 1 \rangle$ in Figure 5 (a) and (b) suggest that we pick $K = 6000$ for Examples 1 and 3 and $K = 8000$ for Examples 2 and 4. For Examples 5 and 6, however, energy averages do not provide clear insight. We turn then to minimums from which Figure 5 (c) and (d) suggest we pick $K = 8000$ and $K = 12000$, respectively.

Hexagonal torus Ω .

Example 1. With $N = 973$ and $K = 6000$ chosen, we run our method on a larger batch of 100 initial configurations; it reaches the ground state with $E^{min} - 1 \approx 1e-14$ (highest precision allowed by our implementation) while the optimal PCVT from the

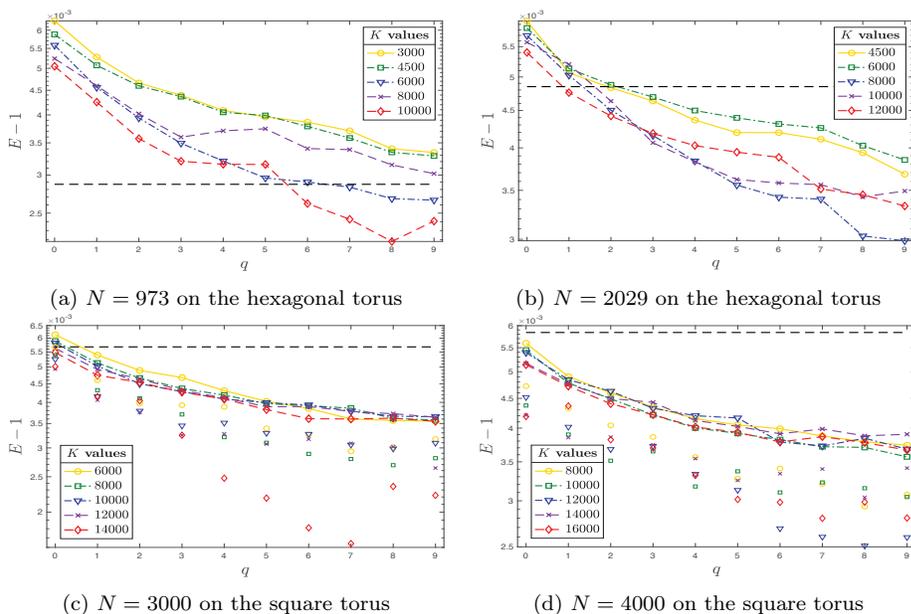


FIG. 5. *K*-sweeps: the joint markers represent $\langle E - 1 \rangle$ while the isolated ones represent minimums over 15 initial configurations. In black dotted lines are the values $E^{ref} - 1$ appearing in Tables 1, 2, 5, and 6 for a first comparison.

comparative methods is $E^{ref} - 1 = 0.00287$ (achieved by L-BFGS(7)); see Figures 6 and 7 as well as the statistics summary of Table 1. In particular, Figure 6 (c) shows how the symbiotic blocks of $MACN-\delta+MACN-c$ act on probing non-PCVT configurations with energy close to (if not below) E^{ref} .

Notice at last how the values f_{E-1}^* , $f_{R_\epsilon}^*$, f_H^* indicate that our method outperforms the others for a significant fraction of the runs. In particular, not only do we get the ratio τ (5.1) to be sensibly zero but on average one needs to run our hybrid algorithm on approximately three uniformly sampled initial configurations to obtain an energy lower than E^{ref} . In other words, we only need ≈ 30 PCVTs so that our way of probing the energy landscape achieves comparable results to the sampling of the basins done by L-BFGS(7) out of 100000 runs.

Example 2. For $N = 2029$ we first remark from the reduced sets of runs from Figure 5 (b) that energy averages and their deviation from E^{ref} immediately compare favorably to those in Example 1. This is the first indicator that our hybrid method is less affected than the comparative methods by the increasing nonconvexity of the energy with N . The same conclusion can be drawn from the data from a hundred runs in Tables 1 and 2; more precisely we have that although our hybrid method did not achieve the honeycomb structure, our optimal PCVT with $E^{min} - 1 = 0.00150$ is still far more regular and gets $\tau^{-1} \approx 3.2$ times closer to the ground state than the compared methods. Furthermore, the impressive value $f_{E-1}^* = 1.00$ achieved suggests that despite the increased nonconvexity with N , the required number of PCVTs needed to navigate the energy landscape, in a similar fashion as the random sampling made with gradient based methods, has decreased to nine (i.e., less than one full run).

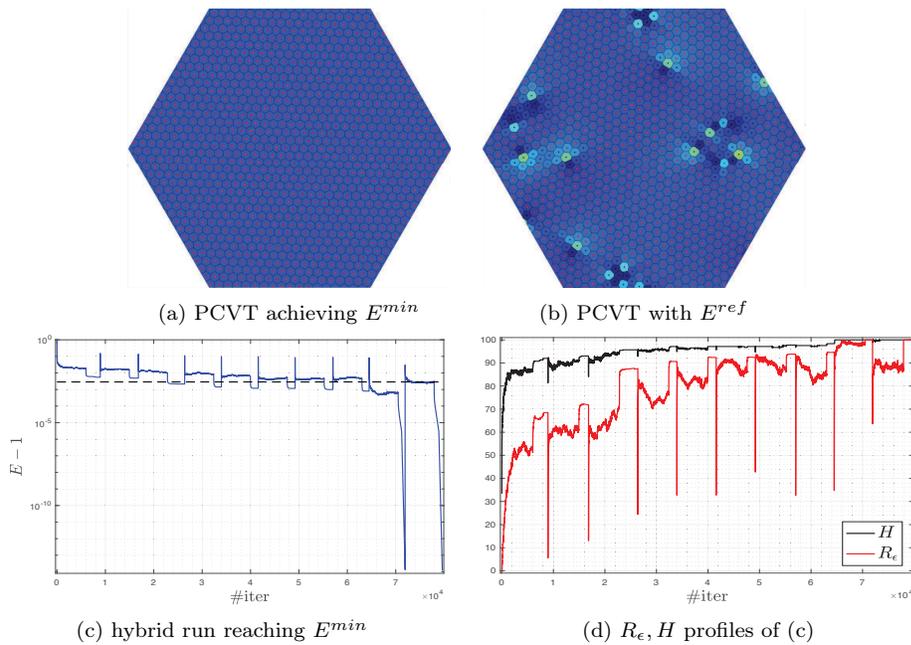


FIG. 6. Optimal PCVTs from Example 1: (a) Ground state with $E^{min} - 1 \approx 1e - 14$ reached by our hybrid method. (b) Configuration carrying $E^{ref} - 1 = 0.00287$ achieved by L-BFGS(7). (c) and (d) The measures of a hybrid run that reached E^{min} among the larger batch of 100 runs with $K = 6000$ (the dashed black line designates $E^{ref} - 1$ again).

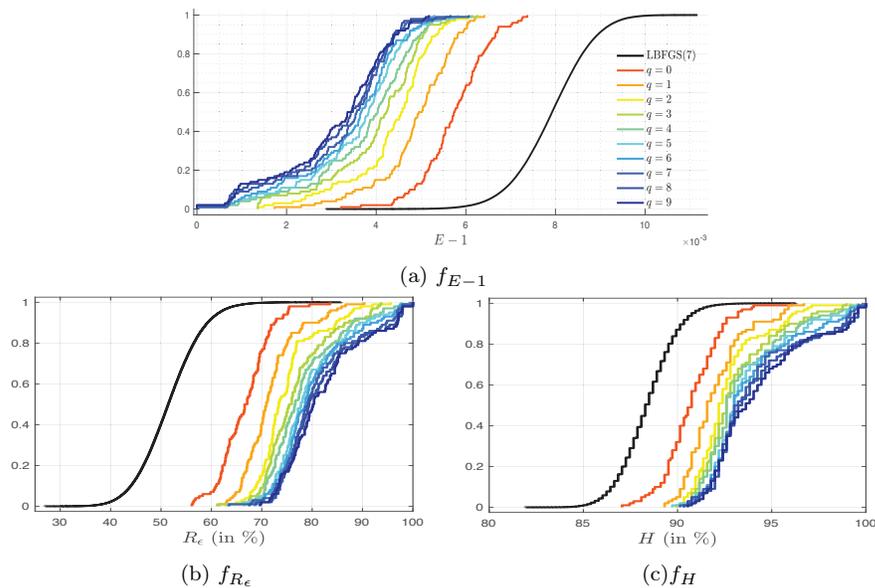


FIG. 7. Data from Example 1: ECDFs of the three regularity measures for 100 hybrid runs with $K = 6000$ along with 100000 runs of L-BFGS(7).

Square torus Ω . We now work with Ω being the square torus and $N = \{n \times 1000\}_{n=1}^4$ to investigate the behavior of our method when the shape effects are tightened.

Example 3. For $N = 1000$ with $K = 6000$ (as in Example 1), we implemented a large batch of 1000 runs of our algorithm so as to have more robust statistics that are presented in Figure 9 and Table 3. We further wish to emphasize, through the monotone skewness in q of the ECDFs and histograms, the increasing performance tendency from one stage to another in this scenario where the contraction and relaxation phases are more constrained by the domain's shape.

The behavior is illustrated in the mosaic of PVTs from Figure 8 where snapshots of the run achieving $E^{min} - 1 = 5.27e-4$ were taken. In particular, one appreciates how stage after stage we achieve the following:

- (i) The *MACN-c* dynamics make defective (nonhexagonally regular) regions of the PVTs “communicate” with each other by creating a flow between regions with high average of individual energy E_i and ones with sensibly lower value, then exhibiting a clear nonlocal behavior.
- (ii) Lloyd’s algorithm contracts the system while preserving the localization of the defects.
- (iii) *MACN- δ* preserves regions of hexagonal regularity better and better as the energy of the perturbed PCVT diminishes while, similarly to *MACN-c*, creating a nonlocal “communication” between defects.

At last, we point out that the progressive constriction of the defect “interfaces” we observe after the Lloyd block (the two middle columns on the mosaic) is recurrent across all set ups that were tested; this is simply a consequence of the remarkable navigation of the energy that our hybrid method performs. The video animating the iterations of this particular run can be found in the accompanying supplementary material (Animation_Emin_MACN_Stages_N1000.mp4 [local/web 22.7MB]). It will appear in more detail how the combination of these symbiotic blocks seems to have a similar effect to a grain boundary evolution algorithm in polycrystals when looking at the produced sequence of PCVTs.

Example 4. For $N = 2000$ with $K = 8000$ (as in Example 2) our results for 100 runs are summarized in Table 4. Here P-L-BFGS(20,20) achieved $E^{ref} - 1 = 0.00495$ which yields the ratio $\tau^{-1} \approx 2.6$. We note that despite the nonnegligible performance decrease of τ^{-1} when compared to Example 2, the values of f_{E-1}^* above 90% for $\mathbf{X}_{q \geq 4}^*$ indicate that, statistically, the overall comparative efficiency is remarkably similar to what we obtained in the hexagonal torus domain. This suggests that the aforementioned change in τ could be mainly attributed to the increased rigidity of the square torus.

On the other hand, when restricting our attention to the distributions of E achieved thus far from 100 runs or more, we make the point that there is no such discrepancy since our energy scaling shows remarkable robustness for each method (e.g., $\langle E - 1 \rangle$ measurements seem to remain comparable regardless of N and Ω ; this will be seen as well in the remaining cases).

Example 5. When running our algorithm on a total of 100 initial configurations with $K = 8000$ for $N = 3000$ and comparing against P-L-BFGS(20,20), we get $\tau^{-1} \approx 3.1$ as well as values of f_{E-1}^* above 90% for $\mathbf{X}_{q \geq 1}^*$; see Table 5. The increasing relative performance with N is once again apparent.

Example 6. At last we consider $N = 4000$ and 100 runs with $K = 12000$; see Table 6. The reader can appreciate how this set up comes to further corroborate the assertions made thus far about the nature of our hybrid algorithm, namely:

- (i) The monotone decrease of $\langle E - 1 \rangle$ in terms of q for batches of 100 runs or more. Additionally, we've seen the robustness of E in terms of N for the PCVTs produced by each method individually. This suggests, for example, that within $Q = 10$ stages our hybrid method is able to get on average twice as close to the nonachievable regular hexagonal configuration as the quasi-Newton methods.
- (ii) As the nonconvexity of the problem increases with N , it is of course harder to get closer to the ground state (the smallest energies recorded increase regardless of the method), yet we observe that, as the nonconvexity increases, we require fewer stages of our method so that f_{E-1}^* surpasses 90%. This shows that the manner in which our method probes the energy landscape manages to overcome this stiffness remarkably better than gradient based methods combined with random-like sampling.

Graphics providing visual insight on the results of Examples 2, 4, 5, and 6 can be found (SuppMaterial_NavigatingCVTLandscape.pdf [local/web 5.43MB]) in the accompanying supplementary material.

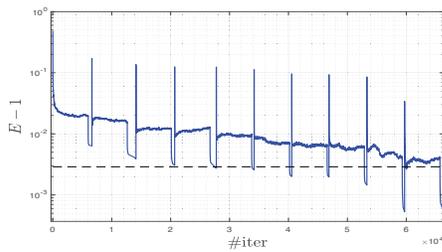
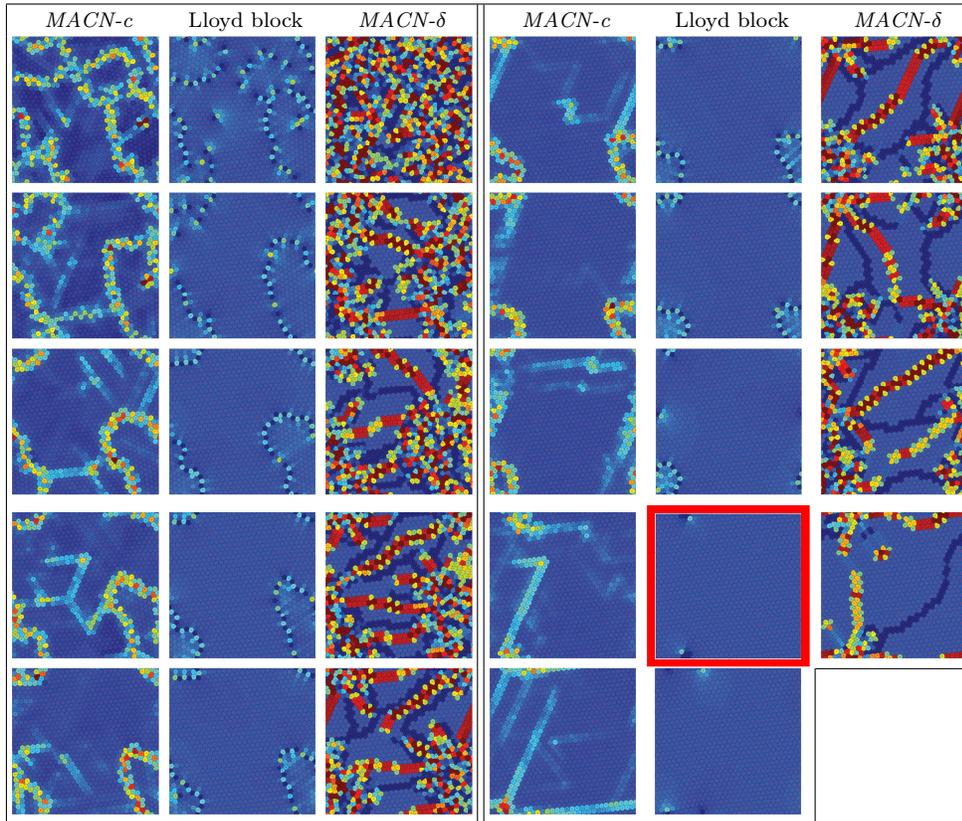
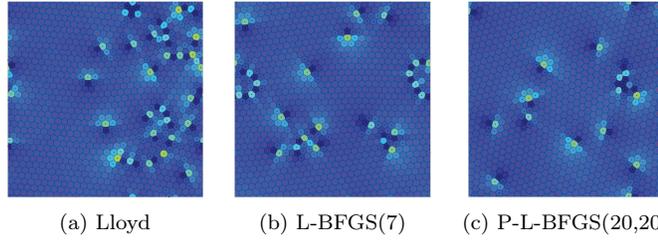
Scope on R_ϵ . We have taken ϵ to be fixed at 0.5% because this value is bounded above by the deviation from r_{hex} obtained from the δ -perturbation of a single generator in the honeycomb PCVT, yet it remains big enough so that the data clearly shows the following:

- A higher variation $|\Delta E|$ between consecutive iterations in our method results in a higher $|\Delta R_\epsilon|$ than $|\Delta H|$, regardless of the block *MACN-c/Lloyd/MACN- δ* .
- We have systematically that $|f_{E-1}^* - f_{R_\epsilon}^*| < |f_{E-1}^* - f_H^*|$.
- The ECDFs of H present larger discontinuity jumps than the ones of R_ϵ , meaning that for given \mathbf{X}^* the number of computed PCVTS states sharing the value $H(\mathbf{X}^*)$ is higher than the one sharing $R_\epsilon(\mathbf{X}^*)$.

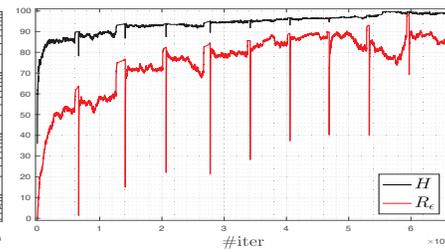
These combined observations point out that R_ϵ is indeed a measure more faithful to E and a better indicator of “well-distributed” PCVTs than H is. We provide further insight on this matter in Table 7 through the correlation ratio

$$\varrho := \frac{\sigma_{R_\epsilon} \text{COV}(E - 1, H)}{\sigma_H \text{COV}(E - 1, R_\epsilon)}$$

and in Figure 10 through scatter plots of the data from Example 3.



(d) hybrid run achieving E^{min}



(e) R_e, H profiles associated with (d)

FIG. 8. Configurations with $N = 1000$ from Example 3: (a) optimal PCVT obtained with Lloyd's method with $E - 1 = 0.00466$, (b) optimal PCVT obtained with L-BFGS(7) having value $E - 1 = 0.00349$, (c) minimal energy configuration obtained among the three comparative methods—achieved by P-L-BFGS(20,20)—with $E^{ref} - 1 = 0.00289$. The mosaic shows the last iteration of each of the three blocks of our hybrid method across the 10 stages of the run that achieved $E^{min} - 1 = 5.27e-4$; the PCVT carrying that minimal value is framed in a red box. These images are to be read starting on the left side of the vertical double black bar and then on the right; each row is for a stage $q = 0, \dots, 9$. Both left columns are the last iteration of the MACN-c blocks, the middle columns contain the PCVTs from the Lloyd blocks, and the right columns are their respective MACN- δ dislocations. Finally, (d) and (e) are the regularity measures' profiles of the run depicted in the mosaic, i.e., the one that achieved E^{min} among the 1000 runs using $K = 6000$.

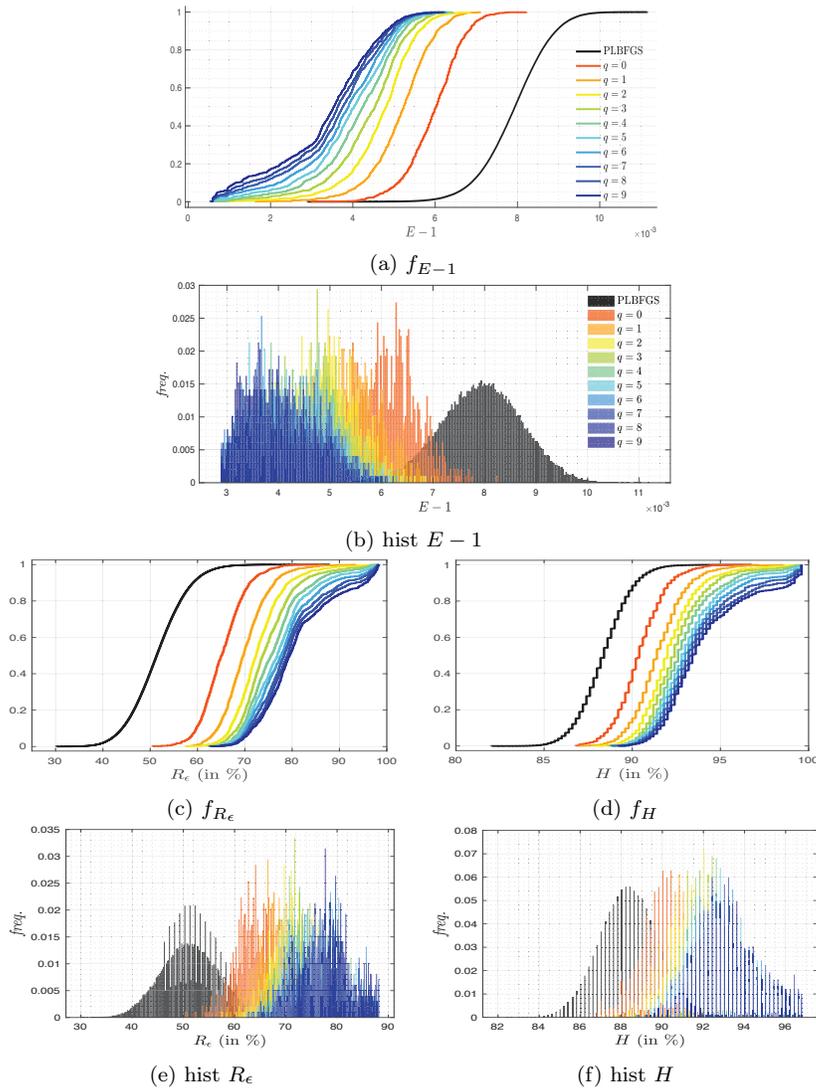


FIG. 9. Distributions from Example 3 for $N = 1000$ with $K = 6000$: ECDFs and histograms of regularity measures for the 1000 hybrid runs as well as for the 100000 runs of $P-L-BFGS(20,20)$. Further details on these measures are given in Table 3.

TABLE 1

Statistics of Example 1 for $N = 973$ with 100 runs of our hybrid method using $K = 6000$; the values $E^{min} - 1$, R_ϵ^{min} , H^{min} as well as $E^{ref} - 1$, R_ϵ^{ref} , H^{ref} are boldfaced.

	Hybrid									Lloyd	LBFGS	PLBFGS	
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*				
$(E-1)$	0.00574	0.00488	0.00435	0.00401	0.00377	0.00356	0.00342	0.00328	0.00316	0.00309	0.00807	0.00790	0.00789
σ_{E-1}	0.00075	0.00087	0.00095	0.00105	0.00113	0.00116	0.00116	0.00121	0.00123	0.00124	0.00082	0.00082	0.00082
max	0.00737	0.00641	0.00626	0.00625	0.00582	0.00613	0.00529	0.00606	0.00565	0.00517	0.01102	0.01116	0.01114
min	0.00321	0.00173	0.00137	0.00135	0.00062	$\approx 1e-14$	0.00383	0.00287	0.00349				
f_{E-1}^*	0.00	0.04	0.10	0.15	0.19	0.23	0.26	0.28	0.33	0.36	---	---	---
(R_ϵ)	66.71	71.76	76.94	76.88	78.22	79.48	80.22	81.11	81.81	82.33	50.42	51.56	51.63
σ_{R_ϵ}	5.01	5.57	6.06	6.69	7.05	7.33	7.34	7.83	7.90	7.96	5.99	5.98	5.98
max	83.55	90.33	95.58	93.73	97.73	100	100	100	100	100	79.44	85.81	82.01
min	56.11	61.56	62.79	61.15	66.49	63.20	68.65	63.51	67.52	69.68	31.55	26.92	28.57
$1 - f_{R_\epsilon}^*$	0.00	0.02	0.08	0.12	0.15	0.19	0.20	0.21	0.24	0.25	---	---	---
(H)	90.72	91.98	92.66	93.11	93.38	93.69	93.85	94.08	94.23	94.35	88.26	88.38	88.39
σ_H	1.51	1.60	1.76	2.00	2.15	2.23	2.34	2.52	2.59	2.60	1.46	1.47	1.47
max	96.09	96.71	99.17	98.97	99.58	100	100	100	100	100	94.65	96.30	95.68
min	87.05	89.31	89.92	89.92	90.33	89.72	90.33	90.13	90.33	90.54	82.52	81.91	81.91
$1 - f_H^*$	0.00	0.01	0.05	0.10	0.14	0.16	0.16	0.20	0.22	0.21	---	---	---

TABLE 2

Statistics of Example 2 for $N = 2029$ with 100 runs of our hybrid method using $K = 8000$.

	Hybrid									Lloyd	LBFGS	PLBFGS	
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*				
$(E-1)$	0.00559	0.00482	0.00445	0.00423	0.00397	0.00379	0.00370	0.00357	0.00342	0.00335	0.00802	0.00784	0.00783
σ_{E-1}	0.00047	0.00053	0.00053	0.00046	0.00047	0.00050	0.00054	0.00054	0.00059	0.00065	0.00057	0.00057	0.00057
max	0.00659	0.00598	0.00567	0.00548	0.00529	0.00515	0.00511	0.00493	0.00473	0.00518	0.00985	0.01013	0.01034
min	0.00401	0.00361	0.00301	0.00304	0.00245	0.00249	0.00202	0.00176	0.00176	0.00150	0.00572	0.00485	0.00523
f_{E-1}^*	0.08	0.53	0.77	0.93	0.99	0.98	0.98	0.98	1.00	0.99	---	---	---
(R_ϵ)	67.41	72.06	74.25	75.58	77.10	78.23	78.75	79.63	80.44	80.93	50.75	51.98	52.03
σ_{R_ϵ}	3.12	3.42	3.41	2.90	2.93	3.17	3.37	3.33	3.61	4.02	4.17	4.15	4.15
max	77.62	81.27	84.22	82.65	85.36	85.60	88.12	90.58	90.53	92.65	67.12	72.10	71.21
min	61.60	64.31	65.94	67.12	68.45	68.85	69.29	70.08	72.15	69.09	37.16	35.18	35.97
$1 - f_{R_\epsilon}^*$	0.07	0.50	0.73	0.90	0.97	0.97	0.98	0.99	1.00	0.99	---	---	---
(H)	90.99	92.04	92.61	92.90	93.32	93.57	93.77	93.95	94.19	94.33	88.38	88.50	88.50
σ_H	0.89	0.96	1.01	0.89	0.91	0.95	0.99	0.97	1.09	1.21	0.99	1.01	1.01
max	93.39	94.48	95.66	95.56	95.46	96.05	96.15	96.64	96.64	98.02	92.70	93.49	93.00
min	88.71	89.94	89.74	90.04	91.22	91.42	90.93	91.03	92.11	90.93	85.11	84.42	84.22
$1 - f_H^*$	0.00	0.06	0.21	0.29	0.42	0.61	0.64	0.68	0.78	0.80	---	---	---

TABLE 3

Statistics from Example 3 with $N = 1000$ and $K = 6000$: 1000 runs of our hybrid method versus the comparative algorithms. The values $E^{min} - 1$, R_ϵ^{min} , H^{min} as well as $E^{ref} - 1$, R_ϵ^{ref} , H^{ref} are boldfaced.

	Hybrid									Lloyd	LBFGS	PLBFGS	
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*				
$(E-1)$	0.00597	0.00516	0.00469	0.00437	0.00412	0.00392	0.00374	0.00358	0.00346	0.00333	0.00848	0.00790	0.00791
σ_{E-1}	0.00068	0.00075	0.00083	0.00092	0.00099	0.00106	0.00113	0.00117	0.00121	0.00126	0.00080	0.00081	0.00081
max	0.00822	0.00710	0.00686	0.00631	0.00643	0.00628	0.00592	0.00592	0.00610	0.00622	0.01106	0.01142	0.01116
min	0.00297	0.00161	7.40e-4	5.96e-4	6.26e-4	5.91e-4	5.82e-4	5.50e-4	5.27e-4	5.50e-4	0.00466	0.00349	0.00289
f_{E-1}^*	0.000	0.009	0.030	0.058	0.114	0.143	0.188	0.227	0.262	0.283	---	---	---
(R_ϵ)	65.32	70.05	72.93	74.76	76.26	77.43	78.47	78.48	80.24	81.03	47.77	51.57	51.51
σ_{R_ϵ}	4.67	4.96	5.42	5.93	6.33	6.78	7.19	7.50	7.75	8.06	5.79	5.90	5.87
max	86.20	93.60	98.20	98.20	98.20	98.20	98.20	98.20	98.40	98.30	74.59	82.50	88.09
min	50.50	57.60	57.50	62.30	62.30	63.50	62.70	64.40	63.30	62.60	27.89	28.30	30.10
$1 - f_{R_\epsilon}^*$	0.000	0.004	0.017	0.034	0.048	0.074	0.103	0.145	0.166	0.192	---	---	---
(H)	90.42	91.50	92.15	92.56	92.91	93.22	93.50	93.73	93.95	94.16	87.71	88.38	88.37
σ_H	1.38	1.50	1.64	1.74	1.89	2.04	2.19	2.32	2.42	2.53	1.42	1.44	1.44
max	97.20	98.60	99.60	99.60	94.00	95.79	96.79						
min	86.80	87.00	87.80	88.60	89.20	88.80	89.20	88.80	89.40	89.20	82.59	82.00	82.00
$1 - f_H^*$	0.001	0.003	0.019	0.029	0.042	0.061	0.079	0.112	0.138	0.161	---	---	---

TABLE 4

Statistics out of 100 hybrid method runs from Example 4 with $N = 2000$ and $K = 8000$.

	Hybrid									Lloyd	LBFGS	PLBFGS	
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*				
$(E-1)$	0.00472	0.00500	0.00462	0.00433	0.00418	0.00402	0.00395	0.00377	0.00368	0.00355	0.00842	0.00784	0.00786
σ_{E-1}	0.00043	0.00044	0.00053	0.00057	0.00056	0.00060	0.00059	0.00061	0.00060	0.00060	0.00057	0.00057	0.00057
max	0.00671	0.00582	0.00584	0.00586	0.00552	0.00581	0.00523	0.00533	0.00520	0.00551	0.01051	0.01013	0.01034
min	0.00471	0.00345	0.00256	0.00196	0.00226	0.00195	0.00197	0.00193	0.00191	0.00204	0.00585	0.00511	0.00495
f_{E-1}^*	0.04	0.42	0.75	0.88	0.93	0.97	0.99	0.99	0.99	0.99	---	---	---
(R_ϵ)	66.70	71.17	73.43	75.27	76.12	76.95	77.46	78.60	79.03	79.85	48.21	51.95	51.82
σ_{R_ϵ}	2.97	2.80	3.47	3.49	3.51	3.67	3.65	3.73	3.72	3.69	4.12	4.15	4.18
max	73.70	81.65	86.65	89.85	87.90	89.90	90.00	90.10	89.85	89.45	65.80	72.54	71.95
min	60.00	66.05	64.70	65.85	66.80	65.90	69.25	69.65	70.45	67.75	33.55	35.39	35.75
$1 - f_{R_\epsilon}^*$	0.02	0.25	0.63	0.83	0.91	0.90	0.94	0.96	0.97	0.99	---	---	---
(H)	90.75	91.77	92.40	92.85	93.15	93.37	93.46	93.74	93.88	94.07	87.84	88.48	88.45
σ_H	0.84	0.91	1.01	1.03	1.03	1.09	1.05	1.10	1.08	1.06	1.01	1.01	1.02
max	92.40	94.70	95.95	96.40	96.50	96.40	96.40	96.50	96.65	97.00	91.70	93.50	93.30
min	88.90	89.80	90.00	90.10	90.00	90.40	90.50	91.10	91.90	91.10	83.90	84.00	83.90
$1 - f_H^*$	0.00	0.04	0.10	0.23	0.33	0.38	0.44	0.58	0.56	0.72	---	---	---

TABLE 5
 Statistics out of 100 hybrid method runs from Example 5 with $N = 3000$ and $K = 8000$.

	Hybrid										Lloyd	LBFGS	PLBFGS
	X_1^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*	X_9^*			
$(E-1)$	0.00577	0.00499	0.00462	0.00443	0.00427	0.00412	0.00396	0.00380	0.00368	0.00361	0.00842	0.00782	0.00786
σ_{E-1}	0.00044	0.00044	0.00047	0.00049	0.00056	0.00057	0.00057	0.00380	0.00053	0.00060	0.00057	0.00047	0.00047
max	0.00678	0.00587	0.00566	0.00547	0.00528	0.00530	0.00529	0.00519	0.00533	0.00540	0.01051	0.00985	0.00986
min	0.00425	0.00318	0.00314	0.00245	0.00199	0.00179	0.00208	0.00231	0.00218	0.00219	0.00585	0.00583	0.00568
f_{E-1}^*	0.44	0.98	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---	---	---
(R_e)	66.41	71.15	73.42	74.52	75.518	76.44	77.38	78.32	79.03	79.53	48.21	52.12	51.89
σ_{R_e}	2.97	2.83	2.97	3.20	3.60	3.57	3.54	3.31	3.34	3.72	4.12	3.42	3.42
max	76.93	82.63	82.06	85.96	89.30	90.36	89.23	87.66	88.53	87.53	65.80	66.13	67.20
min	60.40	65.10	66.40	67.83	69.00	69.40	69.53	70.33	69.23	67.80	33.55	38.36	38.46
$1 - f_{R_e}^*$	0.40	0.95	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---	---	---
(H)	90.67	91.79	92.35	92.70	92.95	93.18	93.44	93.72	93.97	94.11	87.84	88.53	88.46
σ_H	0.94	0.86	0.80	0.98	1.09	1.12	1.10	1.03	1.05	1.14	1.01	0.83	0.83
max	93.76	95.40	95.00	96.53	97.20	97.53	97.26	96.73	96.93	96.73	91.70	92.06	92.73
min	88.66	90.00	90.66	90.66	91.00	91.00	90.86	91.46	91.00	90.86	83.90	85.20	85.13
$1 - f_H^*$	0.03	0.17	0.34	0.45	0.60	0.64	0.76	0.81	0.89	0.89	---	---	---

TABLE 6
 Statistics out of 100 hybrid method runs from Example 6 with $N = 4000$ and $K = 12000$.

	Hybrid										Lloyd	LBFGS	PLBFGS
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*	X_9^*			
$(E-1)$	0.00535	0.00474	0.00441	0.00424	0.00404	0.00396	0.00380	0.00376	0.00375	0.00368	0.00833	0.00781	0.00786
σ_{E-1}	0.00033	0.00036	0.00039	0.00038	0.00040	0.00042	0.00042	0.00046	0.00047	0.00048	0.00041	0.00040	0.00040
max	0.00626	0.00545	0.00517	0.00507	0.00502	0.00486	0.00486	0.00472	0.00479	0.00499	0.00989	0.00956	0.00946
min	0.00448	0.00389	0.00338	0.00321	0.00314	0.00300	0.00268	0.00255	0.00240	0.00228	0.00645	0.00584	0.00601
f_{E-1}^*	0.96	1.00	---	---	---								
(R_e)	69.11	72.61	74.64	75.75	76.93	77.44	78.43	78.72	78.66	79.15	48.84	52.21	51.86
σ_{R_e}	2.24	2.36	2.51	2.49	2.50	2.65	2.66	2.93	2.83	2.94	2.96	2.96	2.95
max	75.25	78.15	81.15	82.50	82.55	84.15	84.60	86.87	87.27	88.22	61.65	66.02	65.40
min	61.75	67.85	69.75	70.12	71.17	71.80	71.52	73.17	72.40	71.57	38.25	40.12	40.25
$1 - f_{R_e}^*$	0.94	1.00	---	---	---								
(H)	91.28	92.19	92.74	92.98	93.34	93.49	93.79	93.85	93.87	93.96	88.00	88.57	88.45
σ_H	0.64	0.73	0.77	0.78	0.75	0.77	0.76	0.79	0.77	0.83	0.72	0.72	0.72
max	92.90	93.90	94.60	95.30	95.10	95.52	95.40	95.95	96.30	96.60	91.30	92.00	91.60
min	89.35	90.65	91.35	91.15	91.45	91.75	91.80	92.20	91.90	91.65	85.50	85.59	85.60
$1 - f_H^*$	0.16	0.53	0.80	0.91	0.98	0.99	0.99	1.00	0.99	0.98	---	---	---

TABLE 7
 Correlation ratios ρ for each method from the PCVT data presented in Examples 1–6.

Ex	Hybrid										Lloyd	LBFGS	PLBFGS
	X_0^*	X_1^*	X_2^*	X_3^*	X_4^*	X_5^*	X_6^*	X_7^*	X_8^*	X_9^*			
1	1.0621	1.0924	1.0838	1.0675	1.0550	1.0687	1.0674	1.0581	1.0523	1.0578	1.0750	1.0663	1.0666
2	1.0691	1.0594	1.0732	1.1228	1.1250	1.1451	1.1594	1.1924	1.1515	1.1035	1.0817	1.0688	1.0683
3	1.0636	1.0717	1.0746	1.0744	1.0667	1.0605	1.0561	1.0555	1.0476	1.0457	1.0822	1.0671	1.0678
4	1.0995	1.0690	1.0646	1.0964	1.0759	1.0649	1.0743	1.0806	1.0960	1.0849	1.0869	1.0692	1.0673
5	1.0644	1.0608	1.0752	1.0490	1.0382	1.0230	1.0389	1.0449	1.0455	1.0430	1.0820	1.0689	1.0679
6	1.0566	1.0586	1.0742	1.0803	1.0724	1.0736	1.0602	1.0447	1.0508	1.0449	1.0842	1.0688	1.0685

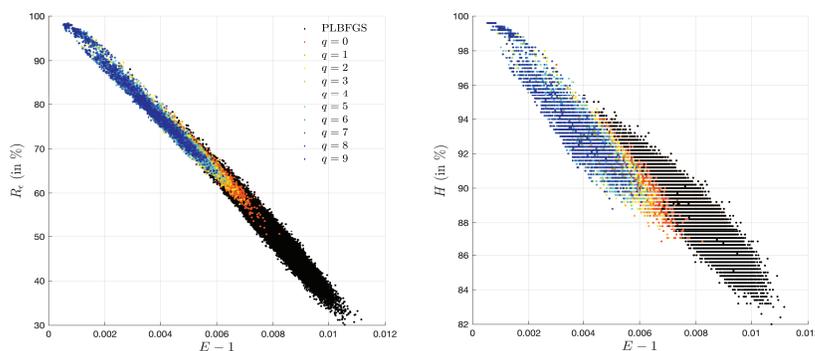


FIG. 10. Scatter plots of E versus H and E versus R_e displaying the same data as Figures 9 (a)–(f) from Example 3.

6. Deterministic versus stochastic. In this section we present evidence of both positive and negative aspects of the fully deterministic nature of our method vis-à-vis stochastic alternatives. For this we start by comparing our *MACN* algorithm with the global Monte Carlo Method (*MCM*) presented in [33] and we finish the section by considering variants of centroidal dislocations to ratify our use of *MACN*- δ .

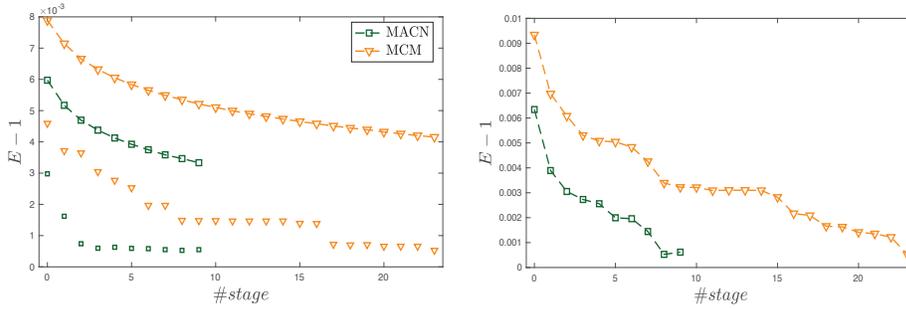


FIG. 11. Comparison of $Q = 10$ stages of *MACN* ($K = 6000$) and 24 stages of *MCM* for $N = 1000$. Left: Averages $\langle E - 1 \rangle$ over 1000 runs are shown with joint markers while minimums achieved are isolated ones. Right: The energy of the PCVTs obtained for the respective runs that first achieved the lowest recorded value $E^{min} - 1 = 5.27e-4$.

MACN versus MCM. We implemented the *MCM* method for $N = 1000$ with the parameter values used for constant densities in [33]; the results along with the *MACN* data from Example 3 of section 5 are summarized in Figure 11. We first depict the energy averages $\langle E - 1 \rangle$ and minimums obtained over 1000 runs, and we also show the energy of each PCVT obtained during two runs, respectively achieving $E^{min} - 1 = 5.27e-4$ at the earliest iteration.

It becomes clear from these statistics that our method has a lower probing number than the comparative stochastic approach; e.g., Figure 11 (right) shows we need to compute 9 versus 24 PCVTs to achieve the same low energy basins of attraction as *MCM*.

Moreover, as already noted in section 3, a crucial advantage of our method is the simple tuple $\{Q, K\}$ of parameters to adjust compared to the more complex set $\{K, T_0, T_k, h\}$ used by *MCM* (note that K serves a different purpose in each method). On the other hand, *MCM* benefits from having a substantially lower computation time. This is due to the fact that: (1) a quasi-Newton method is used instead of Lloyd’s method and (2) the lack of preconditioning spares a total of $K \times Q$ iterations of complexity $O(N \log(N))$ compared to *MACN*; thus making *MCM* still desirable for applications where fast computation time is critical.

Random direction versus MACN motions. As defined in (1.3), the value of δ is crucial since it has the peculiarity of making our perturbation stage preserve a certain regularity in the structure of the tessellation. However, the direction of the perturbation seems to be of primary importance compared to the step size when the latter is fixed. To gain insight on this matter, we define three variants of our

perturbation step:

1. Inspired by the relation between δ and the intrinsic length-scale of the regular hexagonal lattice discussed in section 4, the first variant consists of moving away from the closest neighbor by the length-scale proper to each cell in the PCVT. Precisely, the perturbation follows

$$(6.1) \quad x_i \leftarrow x_i + \frac{|V_i|}{|\partial V_i|} \frac{x_i - x_{j_i^*}}{\|x_i - x_{j_i^*}\|}, \quad i = 1, \dots, N.$$

2. The next variant contemplates δ as in (1.3) but chooses a random neighbor $x_j, j \in \mathcal{N}_i$, to move away from, thus not necessarily being the closest one.
3. Finally, we consider each generator moving by the distance δ and at a random angle $\theta_i \in [0, 2\pi)$ taken from the uniform distribution.

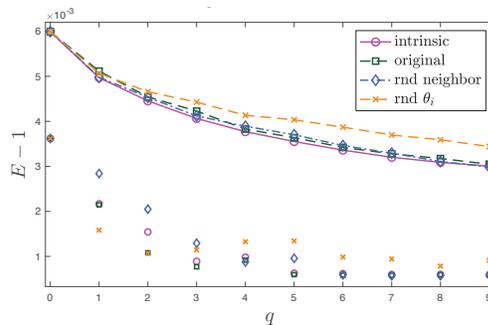


FIG. 12. Comparison in the energy performance of the original $MACN-\delta$ step and its three variants across stages of our hybrid method with $N = 1000$ and $K = 6000$. The joint markers represent averages while the isolated ones represent minima over 100 runs.

Figure 12 illustrates how close the performance of our variants of the $MACN-\delta$ step are from one another but with particular distinction of the random angle θ_i perturbation, being then of some reassurance that our original neighbor-guided dislocation of PCVTs is better suited than some random search in a δ -vicinity.

7. Closing remarks and future directions. We have introduced and assessed a simple deterministic method for navigating the energy landscape of CVTs in two dimensions. This deterministic coupling algorithm (i) only has two degrees of freedom, (ii) shows remarkable robustness with respect to the increasing nonconvexity of the energy as N grows larger, (iii) statistically allows us to systematically obtain configurations closer to the ground state compared to the current state-of-the-art deterministic methods, and (iv) finds energetically comparable results to the leading

stochastic method while needing fewer CVT computations (lower probing number). We also introduced the isoperimetric ratio via R_ϵ as an indicator of low energy CVTs.

We point out that while we prioritized simplicity of the method in this paper, the algorithm's performance could be further improved by the following: (i) using initial quasi-random distributions, (ii) replacing Lloyd's block with other gradient based descent methods that satisfy Wolfe conditions, and most importantly (iii) by introducing a suitable decay in the sequence $\{K_q\}_{q=0}^{Q-1}$ (possibly adapted to $\{E(\mathbf{X}_q^*)\}_{q=0}^{Q-1}$). The point made is that even with the crude tunings made on $K_q \equiv K$ in section 5, the resulting regularity measures are remarkable.

It would be natural to explore our global method in two different settings. First, explore our algorithm on the 3D cubic torus wherein we expect similar comparative results with the appearance of the BCC lattice and truncated octahedron Voronoi cells. Second, explore our algorithm on the 2-sphere. As we previously mentioned in the introduction, work in progress shows that our algorithm works well on the 2-sphere, systematically obtaining configurations closer to the ground state. Of particular interest on the sphere is the nature of the ground state. For certain values of N , for example $N = 32$, the folklore suggests that optimality is tied to *the soccer ball configuration*: hexagons except for exactly 12 regular pentagonal defects. Our results on the 2-sphere as well as structural results of the ground state are forthcoming. We note, however, that in the extension of our method to other manifolds the increased complexity of the underlying computation of the Voronoi tessellation and centroids needs to be compensated by reducing the product $K \times Q$. In particular, time tractability will be at the cost of energy efficiency.

A different related question is the inclusion of an underlying inhomogeneous probability density ρ over Ω wherein the energy (1.2) takes the form $F(\mathbf{X}) = \sum_{i=1}^N \int_{V_i} \|y - x_i\|^2 \rho(y) dy$. However, here it is unclear how to choose the distance δ in the *MACN* annealing step. In order to gain better insight on how to tackle this generalized problem, we strive to cast both our *MACN* iterations as a low order approximation of some yet-to-be-determined gradient flow involving the Voronoi energy.

Appendix A. Explicit pseudo-code for the P-L-BFGS method.

Algorithm A.1 P-L-BFGS(M, T).

Prior definitions: to ease notation we define at each iteration k

$$s_k := \mathbf{X}^{(k+1)} - \mathbf{X}^{(k)}, \quad y_k := DE(\mathbf{X}^{(k+1)}) - DE(\mathbf{X}^{(k)}),$$

$$\rho_k := \frac{1}{y_k^\top s_k}, \quad \text{and } H_0^{(k)} := \frac{s_{k-1}^\top y_{k-1}}{y_{k-1}^\top y_{k-1}} I$$

Input: (i) initial iterate $\mathbf{X}^{(0)}$; (ii) integer parameters M and T ; (iii) tolerance tol for convergence

set $k = 0$

set $diff = \text{Inf}$

while $diff > tol$ **do**

 set $q = DE(\mathbf{X}^{(k)})$

 1st *L-BFGS update*

for $i = k - 1 : -1 : k - M$ **do**

$a_i = \rho_i s_i^\top q$

$q \leftarrow q - a_i y_i$

end for

Redirect search direction

if $k \bmod T = 0$ **then**

 construct preconditioner matrix \tilde{A}_k and solve the system $\tilde{A}_k r = q$

else

 construct $H_0^{(k)}$ and set $r = H_0^{(k)} q$

end if

 2nd *L-BFGS update*

for $i = k - M : k - 1$ **do**

$r \leftarrow r + s_i (a_i - \rho_i y_i^\top r)$

end for

 set descent direction $p^{(k)} = -r$

 update iterate $\mathbf{X}^{(k+1)} = \mathbf{X}^{(k)} + \alpha^{(k)} p^{(k)}$, where $\alpha^{(k)}$ is a step length satisfying the strong Wolfe conditions

if $k > M$ **then**

 erase the tuple $\{s_{k-M}; y_{k-M}\}$

 compute and store $\{s_k; y_k\}$

end if

$diff = \|DE\|/N$

$k \leftarrow k + 1$

end while

Output: \mathbf{X}^* , a stable local minimizer of E and its corresponding PCVT and PDT

REFERENCES

- [1] P. ALLIEZ, D. COHEN-STEINER, M. YVINEC, AND M. DESBRUN, *Variational tetrahedral meshing*, ACM Trans. Graph., 24 (2005), pp. 617–625.
- [2] E. S. BARNES AND N. J. A. SLOANE, *The optimal lattice quantizer in three dimensions*, SIAM J. Algebraic Discrete Methods, 4 (1983), pp. 30–41, <https://doi.org/10.1137/0604005>.
- [3] X. BLANC AND M. LEWIN, *The crystallization conjecture: A review*, EMS Surv. Math. Sci., 2 (2015), pp. 225–306.
- [4] E. BORMASHENKO, M. FRENKEL, AND I. LEGCHENKOVA, *Is the Voronoi entropy a true entropy? Comments on “Entropy, Shannon’s measure of information and Boltzmann’s H-theorem,”* Entropy 2017, 19, 48, Entropy, 21 (2019), 251.
- [5] E. BORMASHENKO, M. FRENKEL, A. VILK, I. LEGCHENKOVA, A. A. FEDORETS, N. E. AKTAEV, L. A. DOMBROVSKY, AND M. NOSONOVSKY, *Characterization of self-assembled 2D patterns with Voronoi entropy*, Entropy, 20 (2018), 956.
- [6] M. CAROLI AND M. TEILLAUD, *Computing 3D periodic triangulations*, in Algorithms–ESA 2009, Lecture Notes in Comput. Sci. 5757, Springer, Berlin, 2009, pp. 59–70.
- [7] J. CHEN, D. WANG, AND Q. DU, *Linear finite element superconvergence on simplicial meshes*, Math. Comp., 83 (2014), pp. 2161–2185.
- [8] L. CHEN AND J.-C. XU, *Optimal Delaunay triangulations*, J. Comput. Math., 22 (2004), pp. 299–308.
- [9] R. CHOKSI AND X. Y. LU, *Bounds on the geometric complexity of optimal centroidal Voronoi tessellations in 3D*, Comm. Math. Phys., 377 (2020), pp. 2429–2450.
- [10] P. N. CHUNG, M. A. FERNANDEZ, N. SHAH, L. S. VIEIRA, AND E. WIKNER, *Perimeter-minimizing pentagonal tilings*, Involve, 7 (2014), pp. 453–478.
- [11] J. CONWAY AND N. SLOANE, *Voronoi regions of lattices, second moments of polytopes, and quantization*, IEEE Trans. Inform. Theory, 28 (1982), pp. 211–226.
- [12] Q. DU AND M. EMELIANENKO, *Acceleration schemes for computing centroidal Voronoi tessellations*, Numer. Linear Algebra Appl., 13 (2006), pp. 173–192.
- [13] 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, <https://doi.org/10.1137/040617364>.
- [14] Q. DU, V. FABER, AND M. GUNZBURGER, *Centroidal Voronoi tessellations: Applications and algorithms*, SIAM Rev., 41 (1999), pp. 637–676, <https://doi.org/10.1137/S0036144599352836>.
- [15] Q. DU, M. GUNZBURGER, AND L. JU, *Advances in studies and applications of centroidal Voronoi tessellations*, Numer. Math. Theory Methods Appl., 3 (2010), pp. 119–142.
- [16] Q. DU AND D. WANG, *The optimal centroidal Voronoi tessellations and the Gershó’s conjecture in the three-dimensional space*, Comput. Math. Appl., 49 (2005), pp. 1355–1373.
- [17] M. EMELIANENKO, *Fast multilevel CVT-based adaptive data visualization algorithm*, Numer. Math. Theory Methods Appl., 3 (2010), pp. 195–211.
- [18] 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, <https://doi.org/10.1137/070691334>.
- [19] A. GERSHO, *Asymptotically optimal block quantization*, IEEE Trans. Inform. Theory, 25 (1979), pp. 373–380.
- [20] S. GRAF AND H. LUSCHGY, *Foundations of Quantization for Probability Distributions*, Lecture Notes in Math. 1730, Springer-Verlag, Berlin, 2000.
- [21] P. M. GRUBER, *A short analytic proof of Fejes Tóth’s theorem on sums of moments*, Aequationes Math., 58 (1999), pp. 291–295.
- [22] W. W. HAGER AND H. ZHANG, *A survey of nonlinear conjugate gradient methods*, Pac. J. Optim., 2 (2006), pp. 35–58.
- [23] J. H. HALTON, *On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals*, Numer. Math., 2 (1960), pp. 84–90.
- [24] J. M. HAMMERSLEY, *Monte Carlo methods for solving multivariable problems*, Ann. New York Acad. Sci., 86 (1960), pp. 844–874.
- [25] J. C. HATELEY, H. WEI, AND L. CHEN, *Fast methods for computing centroidal Voronoi tessellations*, J. Sci. Comput., 63 (2015), pp. 185–212.
- [26] M. IRI, K. MUROTA, AND T. OHYA, *A fast Voronoi-diagram algorithm with applications to geographical optimization problems*, in System Modelling and Optimization, Lect. Notes Control Inf. Sci. 59, Springer, Berlin, 1984, pp. 273–288.
- [27] L. JIANG, R. H. BYRD, E. ESKOW, AND R. B. SCHNABEL, *A Preconditioned L-BFGS Algorithm with Application to Molecular Energy Minimization*, Tech. report, Department of

- Computer Science, University of Colorado at Boulder, Boulder, CO, 2004.
- [28] L. JU, Q. DU, AND M. GUNZBURGER, *Probabilistic methods for centroidal Voronoi tessellations and their parallel implementations*, *Parallel Comput.*, 28 (2002), pp. 1477–1500.
 - [29] Z. LI AND H. A. SCHERAGA, *Monte Carlo-minimization approach to the multiple-minima problem in protein folding*, *Proc. Nat. Acad. Sci. U.S.A.*, 84 (1987), pp. 6611–6615.
 - [30] C.-J. LIN AND J. J. MORÉ, *Incomplete Cholesky factorizations with limited memory*, *SIAM J. Sci. Comput.*, 21 (1999), pp. 24–45, <https://doi.org/10.1137/S1064827597327334>.
 - [31] Y. LIU, W. WANG, B. LÉVY, F. SUN, D.-M. YAN, L. LU, AND C. YANG, *On centroidal Voronoi tessellation-energy smoothness and fast computation*, *ACM Trans. Graph.*, 28 (2009), 101.
 - [32] S. LLOYD, *Least squares quantization in PCM*, *IEEE Trans. Inform. Theory*, 28 (1982), pp. 129–137.
 - [33] L. LU, F. SUN, H. PAN, AND W. WANG, *Global optimization of centroidal Voronoi tessellation with Monte Carlo approach*, *IEEE Trans. Vis. Comput. Graph.*, 18 (2012), pp. 1880–1890.
 - [34] T. MCSWEENEY, *Modified Cholesky Decomposition and Applications*, Master's thesis, University of Manchester, Manchester, UK, 2017.
 - [35] Y. MEN, Z. SHEN, D. KHAN, AND D.-M. YAN, *Improving regularity of the centroidal Voronoi tessellation*, in *SIGGRAPH '18: ACM SIGGRAPH 2018 Posters*, ACM, 2018, 66.
 - [36] M. MORIGUCHI AND K. SUGIHARA, *A new initialization method for constructing centroidal Voronoi tessellations on surface meshes*, in *Proceedings of the 3rd International Symposium on Voronoi Diagrams in Science and Engineering*, IEEE, 2006, pp. 159–165.
 - [37] D. NEWMAN, *The hexagon theorem*, *IEEE Trans. Inform. Theory*, 28 (1982), pp. 137–139.
 - [38] H. NIEDERREITER, *Random Number Generation and Quasi-Monte Carlo Methods*, CBMS-NSF Regional Conf. Ser. in Appl. Math. 63, SIAM, 1992, <https://doi.org/10.1137/1.9781611970081>.
 - [39] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, 2nd ed., Springer, New York, 2006.
 - [40] A. OKABE, B. BOOTS, K. SUGIHARA, AND S. N. CHIU, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd ed., Wiley Ser. Probab. Stat. 501, John Wiley & Sons, Chichester, UK, 2009.
 - [41] J. QUINN, F. SUN, F. C. LANGBEIN, Y.-K. LAI, W. WANG, AND R. R. MARTIN, *Improved initialization for centroidal Voronoi tessellation and optimal Delaunay triangulation*, *Comput.-Aided Des.*, 44 (2012), pp. 1062–1071.
 - [42] THE CGAL PROJECT, *CGAL User and Reference Manual*, CGAL Editorial Board, 5.0 ed., 2019, <https://doc.cgal.org/5.0/Manual/packages.html>.
 - [43] L. F. TÓTH, *Lagerungen in der Ebene auf der Kugel und im Raum*, Springer-Verlag, Berlin, Göttingen, Heidelberg, 1953.
 - [44] L. WANG, F. HÉTRUY-WHEELER, AND E. BOYER, *A hierarchical approach for regular centroidal Voronoi tessellations*, *Comput. Graph. Forum*, 35 (2016), pp. 152–165.
 - [45] D.-M. YAN, K. WANG, B. LÉVY, AND L. ALONSO, *Computing 2D periodic centroidal Voronoi tessellation*, in *Proceedings of the Eighth International Symposium on Voronoi Diagrams in Science and Engineering*, IEEE, 2011, pp. 177–184.
 - [46] J. ZHANG, M. EMELIANENKO, AND Q. DU, *Periodic centroidal Voronoi tessellations*, *Int. J. Numer. Anal. Model.*, 9 (2012), pp. 950–969.