Riesz energy functionals and their applications

By

Oleksandr Vlasiuk

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Mathematics

August 31, 2018

Nashville, Tennessee

Approved:

Edward B. Saff, Ph.D.

Douglas Hardin, Ph.D.

George Hornberger, Ph.D.

Marcelo Disconzi, Ph.D.

Wöden Kusner, Ph.D.

To my father

ACKNOWLEDGMENTS

I am deeply grateful to Professor Edward Saff for advising me, for his unyielding support and thoughtful feedback throughout the program. His guidance, both mathematical and human, was immensely important to me, and has been of great help over those years.

To Professor Douglas Hardin, who has essentially become my informal co-adviser, I am indebted for his invariably energetic and generous attitude in both mathematics and cycling, and for the numerous insights into the field he shared during our discussions. I have greatly enjoyed working alongside the Minimal energy group, which at different times included Wöden Kusner, Timothy Michaels, Alexander Reznikov, Brian Simanek, Yujian Su, and Ken Yang.

I am grateful to all the members of the Mathematics Department at Vanderbilt, both faculty and staff, for creating a welcoming and stimulating environment. Special thanks are due to my fellow graduate students: Arman Darbinyan, Zach Gaslowitz, Colin Klaus, Kelly O'Connell, Armenak Petrosyan, Spencer Sasarita, and Ryan Solava, without whose company this journey would have never been this enjoyable.

The Spring 2018 ICERM program has had a big impact on my understanding of the subject of this dissertation and introduced me to many important results. The credit for this should be fully given to all the wonderful mathematicians who participated in it, and especially Dmitriy Bilyk, Juan Criado del Rey, Ujué Etayo, Alexey Glazyrin, and Mircea Petrache.

I thank my family for loving and supporting an occasionally frustrated mathematician-to-be. They raised my spirits and encouraged me for longer and more efficiently than anyone else, and I owe them much more than could possibly be expressed here. I also thank Ira for checking out the quenelle lyonnaise with me and, most importantly, for being the idea-woman.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1

Preliminaries

1.1 Introduction

The goal of the present text is to establish some fundamental asymptotic properties of an important class of positive definite kernels of energy functionals, and to provide an outline of their applications to manifold discretization and computations with radial basis functions. By an energy functional on discrete configurations in Ω we understand a mapping of the form

$$E(\cdot\,; g, q) : (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N) \in \Omega^N \mapsto \sum_{i \neq j} g(\boldsymbol{x}_i, \boldsymbol{x}_j) + \tau(N) \sum_i q(\boldsymbol{x}_i), \qquad (\star)$$

defined on discrete subsets of a compact set $\Omega \subset \mathbb{R}^p$ of Hausdorff dimension $d := \dim_H \Omega$. Functions g and q are lower semicontinuous and $g(\boldsymbol{x}, \boldsymbol{y})$ is radial, i.e. it depends only on the distance between its arguments $\|\boldsymbol{x} - \boldsymbol{y}\|$ (notation $\|\cdot\|$ stands for the Euclidean distance in $\mathbb{R}^p$). The purpose of the factor $\tau(N)$ is to ensure that the two sums in the RHS have the same order asymptotics as $N \to \infty$, so it will in general depend on both s and d; we shall return to it in a moment. Here and in the following we identify configurations $\{\boldsymbol{x}_i\}$ with their counting probability measures,

$$(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N) \longleftrightarrow \frac{1}{N} \sum_{i=1}^{N} \delta_{\boldsymbol{x}_i},$$

and so the expression $(\star)$ can be seen as a multiple of the continuous energy functional

$$I(\cdot\,; g, q) : \mu \in \mathcal{P}(\Omega) \mapsto \iint\limits_{\Omega \times \Omega} g(\boldsymbol{x}, \boldsymbol{y}) d\mu(\boldsymbol{x}) d\mu(\boldsymbol{y}) + \int\limits_{\Omega} q(\boldsymbol{x}) d\mu(\boldsymbol{x}) \qquad (\star\star)$$

defined on $\mathcal{P}(\Omega)$, the space of probability measures supported on Ω, in the special case of a discrete μ with the diagonal $\boldsymbol{x} = \boldsymbol{y}$ removed. Indeed, denote the subset of $\mathcal{P}(\Omega)$ of discrete measures supported on N vectors by

$$\mathcal{P}_N(\Omega) := \left\{ \mu \in \mathcal{P}(\Omega) \,:\, \mu = \frac{1}{N} \sum_{i=1}^{N} \delta_{\boldsymbol{x}_i},\, \boldsymbol{x}_i \neq \boldsymbol{x}_j \text{ for } i \neq j \right\}.$$

Then for a kernel g continuous on $\Omega \times \Omega$ and any $\mu_N \in \mathcal{P}_N(\Omega)$, setting $\tau(N) = N$ gives

$$I(\mu_N; g, q) = \frac{1}{N^2} E(\mu_N; g, q) + \frac{1}{N^2} \sum_i g(\boldsymbol{x}_i, \boldsymbol{x}_i). \qquad (1.1)$$

As the second sum in the RHS goes to zero as $N \to \infty$, equation (1.1) shows, unsurprisingly, that the asymptotic behavior of the minima of $E(\cdot\,; g, q)$ as $N \to \infty$ is determined by the minima of $I(\cdot\,; g, q)$ in the case of a continuous kernel g; as is shown in the classical potential theory, the same applies to the case of a $g(\boldsymbol{x}, \boldsymbol{y})$ with an integrable singularity on the diagonal $\boldsymbol{x} = \boldsymbol{y}$. As the asymptotic order remains N^2 for any integrable kernel, $\tau(N) = \tau_{s,d}(N) = N$, $s < d$.

The importance of minima and minimizers of the functionals ($\star$)–($\star\star$) lies in their interpretation as ground states of a certain physical system. Expression ($\star$) corresponds to the Hamiltonian of the collection $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_N)$ of particles with pairwise interaction energy $g(\cdot\,, \cdot)$ subject to an external potential q. Optimization of the Hamiltonian over the phase space leads to the ground state of the discrete system or, when $N \to \infty$, to the mean-field limit.

The present text targets the Riesz kernel $g(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$, especially the *hypersingular* case of $s \geq d$, although most of the results in Section 1.4.1 concerning the case of integrable g apply to positive definite lower semicontinuous kernels that are monotone decreasing as a function of radius. A discussion of properties of the Riesz kernel that make it quite remarkable is contained in Section 1.4. We have remarked that when g has an integrable singularity, using $\tau(N) = N$ in ($\star$) ensures that minimizers of $E(\cdot\,; g, q)$ over $\mathcal{P}_N(\Omega)$ converge to those of $I(\cdot\,; g, q)$ over $\mathcal{P}(\Omega)$; this result therefore applies to the Riesz kernel when $s < d$.

In the hypersingular case, $I(\mu; g, q)$ is not defined for measures μ such that $\mu \ll \mathcal{H}_d$, the Hausdorff measure on Ω, whereas the functional $E(\cdot\,; g, q)$ is well-defined on $\mathcal{P}_N(N)$ for all N, which raises the question of the continuous functional describing its asymptotic behavior as $N \to \infty$. This behavior was first described in the Poppy-seed bagel theorem (PSB) of the seminal paper by Hardin and Saff [62] in the case $q \equiv 0$ (see Theorem 1.3 below). It was shown that the minima of $E(\cdot\,; g, 0)$ grow as $N^{1+s/d}$ when $N \to \infty$ and $s > d$; this implies in particular that one has to set $\tau(N) = \tau_{s,d}(N) = N^{s/d}$ for $s > d$. Furthermore, identifying a collection $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_N) \in \Omega^N$ with its normalized counting measure, and minimizing the functional $E(\cdot\,; g, 0)$ over such collections, one obtains discrete measures weak*-converging to $\mathcal{H}_d$ on Ω and possessing asymptotically optimal local properties of covering and separation (the formal definitions can be found below). It therefore became clear that the hypersingular kernel might be particularly well-suited as a means to produce discrete configurations in a controlled fashion; this objective can be loosely stated in the following form.

Problem. *Given a compact rectifiable set $\Omega \subset \mathbb{R}^p$ with a probability measure μ on it, absolutely continuous with respect to the appropriately restricted Hausdorff measure $\mathcal{H}_d$, approximate μ on Ω with discrete measures.*

Here (d-)rectifiability$d$-rectifiable set means it is the image of a bounded subset of $\mathbb{R}^d$ under a Lipschitz mapping. The local behavior of discrete approximants is an important factor, influencing the quality of meshes [98] and meshless solvers [29]; although not the primary motivation for the Riesz energy-based discretization, it turns out that in practice the Voronoi tessellations

of the minimizers have a highly regular structure. The method of [62] was later successfully applied to nonuniform distributions by Borodachov, Hardin, and Saff [15, 18]. The present text is an extension of the above work; the contents Chapter 2 is the next natural step in this direction, introducing an external field instead of a multiplicative weight. Thereby we can confine optimal configurations to specified sets. Furthermore, the notion of Γ-convergence introduced in Section 1.5 below allows us to give a unified treatment of the continuous functionals describing the asymptotics of the minimizers of both integrable and non-integrable g_s, as well as of the weighted hypersingular g_s with an external field.

1.2 Overview of results

In this section we shall give the approximate formulations of the main results of the dissertation. The technicalities are largely omitted here; full details can be found in the corresponding statement(s) referenced in parentheses. We write $M_{s,d}$ for a certain numerical constant, which can be computed with arbitrary precision, and depends only on the s and d, see Section 1.4.2. We write ω_N for an N-point subset of Ω, so $\omega_N := \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_N\} \subset \Omega$; furthermore, whenever we deal with a Riesz energy functional $E(\cdot\,; g, q)$, its minimizer in $\mathcal{P}_N(\Omega)$ is denoted by $\hat{\omega}_N$. Recall also that we identify ω_N with elements of $\mathcal{P}_N(\Omega)$. The Dirac delta-function supported at $\boldsymbol{x}$ is denoted by $\delta_{\boldsymbol{x}}$; we denote $(t)_+ := \max(0, t)$, positive part of the number. The closed Euclidean ball of radius r in $\mathbb{R}^p$, centered at $\boldsymbol{x}$, is written as $B(\boldsymbol{x}, r)$. In this overview $s > d = \dim_H \Omega$; note that in this range for s, the value of $\tau(N)$ in $(\star)$ has to be set to $\tau(N) = \tau_{s,d}(N) = N^{s/d}$.

First-order term of the hypersingular Riesz energy with an external field

Theorem (Theorem 2.1)**.** *Assume that the compact set Ω is d-rectifiable and $\mathcal{H}_d(\Omega) > 0$. Let further $L_1 = L_1(q, \Omega)$ be the (unique) constant defined so that*

$$d\mu^q := M_{s,d}\,(L_1 - q)_+^{d/s}\; d\mathcal{H}_d$$

is a probability measure on Ω. Then every sequence $\hat{\omega}_N$, $N \geq 2$, of minimizers of the functional $E(\cdot\,; g_s, q)$ on Ω^N, $N \geq 2$, satisfies

$$\lim_{N\to\infty} \frac{E(\hat{\omega}_N; g_s, q)}{N^{1+s/d}} = \int \frac{L_1 d + q(\boldsymbol{x}) s}{d+s}\; d\mu_q(\boldsymbol{x}) =: S(q, \Omega).$$

Furthermore, if ω_N, $N \geq 2$, is a sequence of N-point configurations on Ω satisfying

$$\lim_{N\to\infty} \frac{E(\omega_N; g_s, q)}{N^{1+s/d}} = S(q, \Omega),$$

then

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\omega_N}\delta_{\boldsymbol{x}} \overset{*}{\rightharpoonup} \mu^q, \quad N\to\infty.$$

The support of the limiting distribution is thus characterized as a certain sublevel set of the external field $q(\boldsymbol{x})$. A sequence of configurations attaining the asymptotics $S(q,\Omega)$ is called *asymptotically optimal*. Note that the uniqueness of the weak* limit of an asymptotically optimal sequence does not depend on the convexity of q, and is entirely a consequence of the strong repulsive properties of g_s.

An important feature of the minimizers of $E(\cdot\,;g_s,q)$ is that, like those of the unweighted $E(\cdot\,;g_s,0)$, they have the optimal order minimal separation distances and (on the support of μ_q) covering distance, see theorem below. For this result we shall require that Ω be *Ahlfors regular with dimension* d, that is, for some positive constants $c_0,\ C_0$, inequalities

$$c_0R^d \le \mathcal{H}_d(B(\boldsymbol{x},R)\cap\Omega) \le C_0R^d \tag{1.2}$$

must hold for all $\boldsymbol{x}\in\Omega$ and $0<R\le \operatorname{diam}(\Omega)$, diameter of Ω. In the following statement, $\Omega(L)=\{\boldsymbol{x}\in\Omega : q(\boldsymbol{x})\le L\}$ denotes an L-sublevel set of function q, and $\operatorname{dist}(\boldsymbol{x},\hat\omega_N) = \min\{\|\boldsymbol{x}-\boldsymbol{y}\| : y\in\hat\omega_N\}$ is the distance from point $\boldsymbol{x}$ to the set $\hat\omega_N$, induced by the Euclidean metric on Ω.

Theorem (Theorems 2.7 and 2.10)**.** *If $\Omega\subset\mathbb{R}^p$ is compact and Ahlfors regular with dimension d, then for every sequence of minimizers $\hat\omega_N$, $N\ge 2$, there exists a constant C_1 such that*

$$\min\left\{\|\boldsymbol{x}-\boldsymbol{y}\| : \boldsymbol{x}\ne\boldsymbol{y};\ \boldsymbol{x},\boldsymbol{y}\in\hat\omega_N\right\} \ge C_1N^{-1/d} \qquad N\ge 2.$$

If in addition Ω is d-rectifiable, then for each $h>0$ there is a constant C_2, such that for every $\boldsymbol{x}\in\Omega(L_1-h)$ there holds

$$\operatorname{dist}(\boldsymbol{x},\hat\omega_N)\le C_2N^{-1/d}, \qquad N\ge 2.$$

Riesz energy on fractal sets

Results discussed in the previous section required that Ω be d-rectifiable; on the other hand, observe that minimizers of $E(\cdot\,;g,q)$ are well-defined provided Ω is compact in $\mathbb{R}^p$. It is therefore natural to consider behavior of the Riesz energy functionals on compact non-rectifiable sets, in particular on the class of self-similar fractals. The latter are a natural choice given the scale-invariance of the Riesz kernel; that is, minimizing the unweighted Riesz energy on similar sets leads to similar minimizers.

Let us recall some basic definitions first. A similitude $\psi:\mathbb{R}^p\to\mathbb{R}^p$ is a mapping that can be

written as

$$\psi(\boldsymbol{x}) = rA(\boldsymbol{x}) + \boldsymbol{z}$$

for some orthogonal matrix $A \in O(p)$, a vector $\boldsymbol{z} \in \mathbb{R}^p$, and a contraction ratio $0 < r < 1$. A self-similar fractal, as defined by Hutchinson [69], is a compact $\Omega \subset \mathbb{R}^p$ comprising the fixed points of a collection of similitudes $\{\psi_m\}_{m=1}^M$ with contraction ratios r_m, $1 \le m \le M$, that is, satisfying

$$\Omega = \bigcup_{m=1}^{M} \psi_m(\Omega),$$

where the union is disjoint. Such mappings $\{\psi_m\}$ are said to satisfy the open set condition if there exists a bounded open set $V \subset \mathbb{R}^p$ such that

$$\bigcup_{m=1}^{M} \psi_m(V) \subset V,$$

where the sets in the union are disjoint.

In [13] Borodachov, Hardin, and Saff considered, in particular, properties of $E(\cdot\,; g_s, 0)$ on fractal sets; they have shown that for the minimum values of $E(\cdot\,; g_s, 0)$ on a self-similar fractal Ω there does not exist an asymptotic limit, even assuming that all the contraction ratios of Ω are equal: $r_1 = \ldots = r_M$. Using their results, Calef [32] found a sequence of minimizers on a union of a rectifiable set with a self-similar fractal for which the weak* limit does not exist; consult Section 3.1 for further details.

Chapter 3 contains, among others, the following new results. Note that we use Fraktur for sequences of positive integers, as in $\mathfrak{N} \subset \mathbb{N}$; see also Section 1.3 and Glossary in Appendix A for notation.

Theorem 1.1 (Theorem 3.4)**.** *Let $\Omega \subset \mathbb{R}^p$ be a compact self-similar fractal satisfying the open set condition, and $\dim_H \Omega = d < s$. If $\{\underline{\omega}_N : N \in \underline{\mathfrak{N}}\}$, is a sequence of configurations for which*

$$\lim_{\underline{\mathfrak{N}} \ni N \to \infty} \frac{E(\underline{\omega}_N; g_s, 0)}{N^{1+s/d}} = \liminf_{N \to \infty} \frac{E(\hat{\omega}_N; g_s, 0)}{N^{1+s/d}},$$

then the corresponding sequence of empirical measures converges weak*:

$$\underline{\nu}_N \xrightarrow{\ *\ } \frac{\mathcal{H}_d}{\mathcal{H}_d(\Omega)}, \qquad \underline{\mathfrak{N}} \ni N \to \infty.$$

As usual, here $\mathcal{H}_d$ is the d-dimensional Hausdorff measure on Ω, and $\hat{\omega}_N$ stands for the minimizer of the $E(\cdot\,; g_s, 0)$ Riesz energy in $\mathcal{P}_N(\Omega)$. For a sequence $\mathfrak{N}$, we write

$$\{\mathfrak{N}\} := \lim_{\mathfrak{N} \ni N \to \infty} \{\log_M N\},$$

where $\{\cdot\}$ in the RHS stands for the fractional part, and

$$E_s(\mathfrak{N}) := \lim_{\mathfrak{N} \ni N \to \infty} \frac{E(\hat{\omega}_N; g_s, 0)}{N^{1+s/d}},$$

if either limit exists.

Theorem 1.2 (Theorem 3.6)**.** *If Ω is a self-similar fractal with equal contraction ratios, and two sequences $\mathfrak{N}_1$, $\mathfrak{N}_2 \subset \mathbb{N}$ are such that*

$$\{\mathfrak{N}_1\} = \{\mathfrak{N}_2\},$$

then

$$E_s(\mathfrak{N}_1) = E_s(\mathfrak{N}_2).$$

In particular, the latter limits exist. Moreover, the function $g_{s,d} : \{\mathfrak{N}\} \mapsto E_s(\mathfrak{N})$ is continuous on $[0, 1]$.

Applications to meshless methods

In this section (and in the eponymous chapter) we will refer to the elements of ω_N as *nodes*. The "node" terminology is common in the *radial-basis functions* (RBF) literature [51], and reflects that the discrete measures as above ordinarily serve as the nodes of interpolatory/FD stencils. Furthermore, for simplicity we assume that Ω has the full dimension, so $d = p$.

In order to use the minimizers of $(\star)$-type functionals in constructing discrete sets efficiently, one needs a suitably fast optimization method. On the other hand, geo-modelling applications that were of special interest to us often demand enormous node sets on rather complicated domains, usually with non-differentiable boundary; for example, Section 4.4.1 considers the Earth-shaped Ω_{etopo}. Our objective was the following.

Problem. *Guarantee that for a given function $\rho(\cdot)$, for every node $\boldsymbol{x} \in \omega_N$ the distance to the nearest neighbor $\Delta(\boldsymbol{x})$ satisfies $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$, that is, the two quantities differ only up to a constant factor.*

(In Chapter 4 we refer to the function ρ as *radial density*.) High dimensionality of the optimization problem arising from $(\star)$, combined with nontrivial domain constraints creates a formidable obstacle to applying any of the widely-used optimization solvers that involve the gradient (for example, the popular L-BFGS-B [30]). The difficulty lies in the large number local of minima of $(\star)$, and the complexity of domain Ω.

To tackle the above issues, we have developed the procedure for node distribution in Chapter 4; its general outline is as follows. Assume that the relation $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$ corresponds to the measure μ on Ω; the latter will necessarily be continuous with respect to $\mathcal{H}_d$. First, discretize

the target distribution μ using stratified quasi-Monte Carlo sequence [31], built from periodic (i.e., using a periodized version of the Euclidean metric) minimizers of $E(\cdot\,; g, q)$; then, apply the modified gradient flow of the truncated version of the Riesz energy with multiplicative weight, chosen so that the limiting distribution of the flow is exactly μ. More precisely, our construction performs T iterations, moving the i-th node position on the t-th iteration, $1 \le t \le T$, denoted by $\boldsymbol{x}_i^{(t)}$, in the direction of vector

$$\boldsymbol{g}_i^{(t)} = s\,\rho\left(\boldsymbol{x}_i^{(t)}\right)^s \sum_{k=1}^{K} \frac{\boldsymbol{x}_i^{(t)} - \boldsymbol{x}_{j(i,k)}^{(t)}}{\|\boldsymbol{x}_i^{(t)} - \boldsymbol{x}_{j(i,k)}^{(t)}\|^{s+2}}, \quad 1 \le i \le N. \tag{1.3}$$

Here $\boldsymbol{x}_{j(i,k)}^{(t)}$, $1 \le k \le K$, are the K nearest neighbors of $\boldsymbol{x}_i^{(t)}$. The use of flow (1.3) reduces the necessity of minimizing the whole energy functional to an implementation of line search that determines the length of the step in direction $\boldsymbol{g}_i^{(t)}$; in fact, even fixed size stepping with some elementary adjustments has proven to work well. Chapter 4 describes using a point inclusion check and a fraction of the distance to the nearest neighbor $\Delta(\boldsymbol{x}_i)$:

$$\boldsymbol{x}_i^{(t+1)} = \begin{cases} \boldsymbol{x}_i^{(t)} + \dfrac{\Delta\left(\boldsymbol{x}_i^{(t)}\right)}{t+C} \dfrac{\boldsymbol{g}_i^{(t)}}{\|\boldsymbol{g}_i^{(t)}\|} & \text{if this sum is inside } \Omega; \\ \boldsymbol{x}_i^{(t)}, & \text{otherwise,} \end{cases} \qquad 1 \le i \le N,\ 1 \le t \le T.$$

The multiplicative approach is employed to reduce the computational complexity of the algorithm, and the theoretic groundwork is provided by [19]. Additional simplification is attained due to using the truncated gradient, as shown in (1.3); see discussion on p.82. An even more general treatment is possible, comprising both kernel truncation and external field. For the precise formulation of the algorithm see Section 4.3.1. The procedure outlined here is scalable (our implementation [102] handles up to about 3 million nodes on a standard laptop) due to a complexity that is essentially linear in the number of nodes N, provides locally regular node sets, handles complex densities well, and does not create artifacts on non-smooth boundaries.

Γ-convergence

To reveal the common features of the asymptotic behavior of minimizers of Hamiltonians ($\star$) for all values of s, in Section 1.5 we shall introduce the notion of Γ-convergence. It will be shown that Γ-convergence of energy functionals leads to weak* convergence of their minimizers. This behavior is the same for both integrable and non-integrable kernels g; the analytic tools necessary to establish it however turn out to be somewhat different, and Chapter 2 is dedicated to proving the properties underlying the Γ-convergence in the hypersingular case with external field. The corresponding argument for the integrable case is classical, and we shall sketch out its essential

ideas in Section 1.5.1. The unified view of Γ-convergence allows one to obtain the expression for the minimizer of the hypersingular Riesz energy functional including both a multiplicative weight and an external field:

$$E(\boldsymbol{x}_1, \dots, \boldsymbol{x}_N; g_s, \kappa, q) := \sum_{i \neq j} \kappa(\boldsymbol{x}_i, \boldsymbol{x}_j) g_s(\boldsymbol{x}_i, \boldsymbol{x}_j) + \tau(N) \sum_i q(\boldsymbol{x}_i), \tag{1.4}$$

where $s \geq d = \dim_H \Omega$, as well as gives as simple proof of the uniqueness of the limit of an asymptotically optimal sequence in $\mathcal{P}(\Omega)$, see Sections 2.5 and 1.5.2 respectively.

Observe that the positive definiteness of the Riesz kernels implies also that the functional in $(\star\star)$ is convex on $\mathcal{P}(\Omega)$, and thus has a unique solution. An additional motivation for considering the continuous form $I(\cdot\,; g, q)$ is also that, assuming that the kernel $g(\cdot\,, \cdot)$ is positive definite and integrable, the first integral in the RHS of $(\star\star)$ defines scalar product on $\mathcal{P}(\Omega)$, and therefore a Hilbert space, in which the weak convergence is equivalent to the weak* topology in (Ω) (see also $\mathbf{2}^{\ddagger}$ in Section 1.4.1). Optimization of $(\star\star)$ for q identically equal to zero is then equivalent to the problem of approximation of the minimizing distribution with discrete measures in the aforementioned Hilbert space.

1.3 Notation and layout

The bold typeface is reserved for vectors, $\boldsymbol{x} \in \mathbb{R}^p$; Fraktur is used for sequences of integers, $\mathfrak{N} \subset \mathbb{N}$. All the distributions we discuss are supported on some compact subset of $\mathbb{R}^p$, denoted by Ω; its Hausdorff dimension $\dim_H \Omega = d$ is fixed within each chapter. In Chapter 4, $p = d$, and thus Ω has full dimension (in particular, the dimension of the ambient space is not fixed throughout the text). We write $\mathcal{H}_d$ for the d-dimensional Hausdorff measure normalized to coincide with the respective Lebesgue measure. The scaling factor $\tau(N)$ in $(\star)$ is defined as

$$\tau(N) = \tau_{s,d}(N) := \begin{cases} N, & s < d, \\ N \log N, & s = d, \\ N^{s/d}, & s > d. \end{cases} \tag{1.5}$$

We shall suppress the reference to s and d in $\tau_{s,d}(N)$ whenever it cannot cause confusion. The same applies to the asymptotic order $\mathcal{T}(N)(N) = \mathcal{T}_{s,d}(N) := N\tau_{s,d}(N)$. For further reference, consult the list of common symbolic notation and important terms in the Glossary.

The layout of the present text is following. We collect all the necessary information from classical potential theory and calculus of variations in the Chapter 1, and develop a description of the proofs contained in Chapters 2–3 in terms of Γ-convergence. Chapter 2 is dedicated to minimizing the hypersingular Riesz kernel with an external field; Chapter 3 considers the case of

Ω being a fractal. Chapter 4 discusses applications of the Riesz energy functionals, primarily for generation of node configurations used to form RBF-FD stencils. Appendices include the Glossary and excerpts from the listings of relevant routines. Additional details about the contents of each chapter can be found in the respective introductory section.

1.4 Riesz kernel

1.4.1 Potential-theoretic case $0 < s < d$

The classical potential theory apparently goes back at least to Gauss; other notable contributors (in roughly chronological order) include Poincaré, Vallée-Poussin, the Riesz brothers, Leja, Frostman, Szegő, Keldysh, Lavrentiev, Brelot, Ohtsuka, Choquet, Fuglede, Rakhmanov, Totik, Saff, and Dragnev. The field has a number of important connections to other areas of mathematics, for example, theory of PDEs, polynomial approximation, martingales, theory of holomorphic functions, measure theory, etc. The present section highlights some essential properties of integrable Riesz kernels as well as the properties of respective energies and potentials.

Recall that we are dealing with the problem of minimization of the integral operator $I(\cdot\,; g, q)$ defined in $(\star\star)$ over the set of probability measures supported on Ω. The specific assumptions of smoothness on Ω are postponed until the later formulations; in general we shall need at least that Ω be rectifiable. We shall write $I(\mu)$ whenever this cannot cause confusion, so

$$I(\mu) = \iint_{\Omega\times\Omega} g(\boldsymbol{x}, \boldsymbol{y}) d\mu(\boldsymbol{x}) d\mu(\boldsymbol{y}) + \int_{\Omega} q(\boldsymbol{x}) d\mu(\boldsymbol{x}),$$

for a lower semicontinuous kernel $g(\cdot\,,\cdot)$; in some cases we shall restrict g to positive definite radial kernels. We shall further use $g_s(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$ to denote the Riesz s-kernel. One more important notation to be introduced here is

$$U^{\mu}(\boldsymbol{x}; g) := \int_{\Omega} g(\boldsymbol{x}, \boldsymbol{y}) d\mu(\boldsymbol{y}),$$

the potential of measure μ with the kernel g; as before, whenever it cannot create confusion, we shall write simply $U^{\mu}(\boldsymbol{x})$. Potential of a measure with a Riesz s-kernel g_s is abbreviated to $U^{\mu}_s(\boldsymbol{x})$, and the respective energy to $I_s(\mu)$.

We are ready for a discussion of the Riesz energy in the integrable case. The three lists that follow contain properties of integrable kernels, the potentials they generate, and the corresponding energies obtained by integrating the potentials. Although we focus on the Riesz s-kernel, some properties are stated for a general lower semicontinuous or positive definite kernel. It should be observed that by positive definiteness we understand simply the positivity of the Fourier transform, or as the complete monotonicity of the radial function $\widetilde{g}(r)$ defined by $g(\boldsymbol{x}, \boldsymbol{y}) =$

$g(\|\boldsymbol{x}-\boldsymbol{y}\|^2) =: \widetilde{g}(r^2)$. Finally, suppose for simplicity that Ω is a sufficiently smooth d-dimensional manifold in $\mathbb{R}^p$.

We shall not use most of the properties of integrable kernels below in any of the following chapters; the cursory overview here is intended only as an illustration of the depth of classical potential theory. Certainly no claim for generality or completeness is made, instead we refer the reader to the several excellent sources [83, 74, 100, 93, 96].

1. A Riesz kernel $g_s(\boldsymbol{x},\cdot)$ is lower semicontinuous as a function of one variable $\boldsymbol{x}$.

2. The kernel $g_s(\boldsymbol{x},\cdot)$ is i) superharmonic for $s \in (0, \max\{0, d-2\}]$, ii) subharmonic for $s \in [\max\{0, d-2\}, d)$.

3. Riesz kernel $g_s(\boldsymbol{x},\boldsymbol{0}) = \|\boldsymbol{x}\|^{-s}$ has positive Fourier transform[1]:

$$\left(\|\boldsymbol{x}\|^{-s}\right)^{\wedge} = \gamma_s \|\boldsymbol{\lambda}\|^{s-p},$$

where $\gamma_s = \pi^{s-p/2}\frac{\Gamma((p-s)/2)}{\Gamma(s/p)}$. Consequently, the Riesz kernel is positive definite in the sense of Bochner.

4. It follows from the previous item that Riesz kernels g_s have semigroup properties in $\mathbb{R}^p$ for $s < p$; it is convenient to denote $f_t(\boldsymbol{x}) := \gamma_{d-t}^{-1}\|\boldsymbol{x}\|^{t-p} = \gamma_{d-t}^{-1} g_{p-t}(\boldsymbol{x},\boldsymbol{0})$, then

$$f_r(\boldsymbol{x}) * f_t(\boldsymbol{x}) = f_{r+t}(\boldsymbol{x}),$$

where $*$ denotes convolution in $\mathbb{R}^p$. This equality can be analytically extended to pairs r, t satisfying $\Re(r+t) < p$ and $r, t \neq p+2n,\ n = 0, 1, 2, \ldots$, [74, (1.1.11)]. Note that the relation between exponents s in g_s and t in f_t is $s + t = p$.

5. g_{d-2} is the fundamental solution of the Poisson equation:

$$-\Delta g_{d-2} = c_p \delta_{\boldsymbol{0}}$$

with $c_2 = 2\pi$ and $c_p = (p-2)\mathcal{H}_{p-1}(\mathbb{S}^{p-1}) = 4\pi^2\gamma_{p-2}$.

The above properties of the kernel g_s imply corresponding properties of $U^\mu(\boldsymbol{x}; g_s)$. Perhaps the most important is that in the harmonic case $s = d-2$, $U^\mu(\boldsymbol{x} g_{d-2})$ gives the solution to the Poisson problem with the data given by the generating measure μ.

1†. If g is a lower semicontinuous kernel, $U^\mu(\boldsymbol{x}, g)$ is l.s.c. as a function of $\boldsymbol{x}$. Moreover, it is l.s.c. in the weak* topology on Ω as a function of μ (this result is sometimes called *principle of descent*).

[1] where $\widehat{f}(\boldsymbol{\lambda}) := \int e^{-2\pi i \boldsymbol{x}\cdot\boldsymbol{\lambda}} d\boldsymbol{x}$.

2†. $U^\mu(\boldsymbol{x}, g)$ of a (sub-, super-)harmonic kernel g is itself (sub-, super-)harmonic outside $\operatorname{supp}\mu$. Moreover, if g is superharmonic, $U^\mu(\boldsymbol{x}, g)$ is also superharmonic on $\mathbb{R}^p$.

3†. By **2†**, Riesz potentials satisfy a maximum principle that depends on the exponent s. They also satisfy the following weak maximum principle for all values of s, $0 < s < p$. Suppose

$$U_s^\mu(\boldsymbol{x}) \le M, \qquad \boldsymbol{x} \in \operatorname{supp}\mu.$$

Then

$$U_s^\mu(\boldsymbol{x}) \le 2^s M, \qquad \boldsymbol{x} \in \mathbb{R}^p.$$

4†. For every function $F(\boldsymbol{x})$, supported on Ω and $p+2$ times continuously differentiable, and any $0 < s < d$, there exists a signed absolutely continuous measure on Ω such that, in the notation of **4**,

$$F(\boldsymbol{x}) = U_s^\mu(\boldsymbol{x}) = (f_{p-s} * \mu)(\boldsymbol{x}).$$

Moreover, the density of μ is given by

$$\frac{d\mu}{d\mathcal{H}_d} = (f_{s-p} * F)(\boldsymbol{x}).$$

Note that in the harmonic case $s = p-2$ it is sufficient to require that F have 3 continuous derivatives; it follows from **5** that the corresponding density of μ is then given by $-\frac{1}{4\pi^2}\Delta F$.

5†. If $U_s^\mu(\boldsymbol{x})$ is continuous on $\operatorname{supp}\mu$, it is continuous on $\mathbb{R}^p$.

6†. If $d = p$, any superharmonic function $f(\boldsymbol{x})$ can be (uniquely) represented on Ω as

$$f(\boldsymbol{x}) = U_{p-2}^\mu(\boldsymbol{x}) + h(\boldsymbol{x}), \qquad \boldsymbol{x} \in \Omega,$$

for an appropriate choice of a measure μ supported on Ω, and a harmonic function h; both μ and h are uniquely determined by Ω. This result is called the *Riesz representation theorem for superharmonic functions.*

7†. The lower semicontinuity property of $U^\mu(\boldsymbol{x})$ in μ can be made more precise as follows. If a sequence of probability measures $\{\mu_n : n \ge 1\}$ is such that

$$\mu_n \xrightarrow{*} \mu, \qquad n \to \infty,$$

then

$$\liminf_{n\to\infty} U^{\mu_n}(\boldsymbol{x}) = U^\mu(\boldsymbol{x}) \qquad \mathcal{H}_d\text{-a.e.}$$

The "$\mathcal{H}_d$-a.e." above can be replaced with "except for sets of capacity zero". Notation $\xrightarrow{*}$

here stands for the convergence in the weak* topology, see (1.6) below. This statement is known as the *lower envelope theorem* [26, 91, 93, Theorem 6.9].

Lastly, we formulate the properties of energies $I(\mu; g)$ and especially of $I_s(\mu)$.

1‡. $I(\mu)$ of a lower semicontinuous kernel g is itself lower semicontinuous in the weak* topology, as a functional on the space of probability measures on Ω.

2‡. As can be seen from **3** for example, a positive definite kernel such as $I_s(\mu)$ corresponds to a scalar product on the space of probability measures on Ω. Indeed, fix $0 < s < d$ and define the *mutual energy* of two probability measures μ, ν on Ω by

$$I[\mu, \nu] := \iint\limits_{\Omega\times\Omega} g_s(\boldsymbol{x}, \boldsymbol{y}) d\mu(\boldsymbol{x}) d\nu(\boldsymbol{y}).$$

It can be verified that the self-energy $I(\mu) = I[\mu, \mu]$ is nonnegative for any signed measure, and is equal to zero only if $\mu \equiv 0$. Clearly, the quadratic form $I[\mu, \nu]$ can then serve as a scalar product to introduce Hilbert space topology on the vector space of all signed measures on Ω with finite energies.

It is important to note that the weak topology in the resulting Hilbert space is equivalent to the usual weak* topology when restricted to the set of probability measures supported on Ω [74, pp. 88–89].

3‡. In the harmonic case $s = p - 2$, the energy integral can be represented as

$$I_{p-2}(\mu) = \frac{1}{c_p} \int_{\mathbb{R}^p} \|\nabla U^{\mu}_{p-2}(\boldsymbol{x})\|^2 d\boldsymbol{x}.$$

The proof of this equality consists in applying **5** together with integration by parts.

1.4.2 Hypersingular case $s \geq d$

We call the Riesz kernels $g_s(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$ *hypersingular* when $s \geq d = \dim_H \Omega$. The reason to consider such kernels will become evident from the following Theorem 1.3, which shows that the minimizers of (2.1) with $q \equiv 0$ are well-distributed (in a sense to be made precise later) on the set Ω, which need not be the case when $s < d$.

For the purposes of studying the asymptotic behavior of N-point configurations ω_N on a compact Ω we consider their normalized counting measures. Recall that such measures $N^{-1} \sum_{\boldsymbol{x}\in\omega_N} \delta_x$ are said to *weak* converge* to the measure μ if

$$\forall f \in C(\Omega), \qquad \frac{1}{N} \sum_{\boldsymbol{x}\in\omega_N} f(\boldsymbol{x}) \longrightarrow \int f \, d\mu \quad \text{as } N \to \infty, \tag{1.6}$$

where, as usual, $C(\Omega)$ denotes the family of all continuous functions defined on Ω. Weak* convergence is denoted by $\overset{*}{\rightharpoonup}$.

In this chapter we further need to impose some regularity conditions on the underlying compact set. A set $\Omega \subset \mathbb{R}^p$ is said to be *d-rectifiable* if it is the image of a bounded subset of $\mathbb{R}^d$ under a Lipschitz mapping. Note that any subset of a d-rectifiable set is also d-rectifiable. Here and below we write $\mathbb{B}^d$ for the d-dimensional unit ball. We use $\mathcal{H}_d$ to denote the d-dimensional Hausdorff measure on $\mathbb{R}^p$ normalized so that $[0,1]^d \subset \mathbb{R}^d \subset \mathbb{R}^p$ has unit volume, and by $\mathcal{H}_d$ its restriction to Ω. In particular, for a d-rectifiable Ω, $\mathcal{H}_d(\Omega) < \infty$. We shall also employ the following compact notation (see ($\star$), p.1):

$$\begin{aligned} E_s(\omega_N) &:= E(\omega_N; g_s, 0) \\ \mathcal{E}_s(\Omega, N) &:= \inf_{\omega_N \subset \Omega} E(\omega_N; g_s, 0). \end{aligned} \tag{1.7}$$

This infimum is attained for compact sets Ω because the Riesz s-kernel $\|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$ is lower semicontinuous on $\Omega \times \Omega$.

Theorem 1.3 ([63, 16])**.** *(Poppy-seed bagel theorem) Suppose $s \geq d$ and $\Omega \subset \mathbb{R}^p$ is d-rectifiable and compact. If $s = d$, it is further assumed that Ω is a subset of a d-dimensional C^1 manifold. Then for $s = d$*

$$\lim_{N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^2 \log N} = \frac{\mathcal{H}_d(\mathbb{B}^d)}{\mathcal{H}_d(\Omega)}, \tag{1.8}$$

while for $s > d$, the following limit exists:

$$\lim_{N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} = \frac{C_{s,d}}{\mathcal{H}_d(\Omega)^{s/d}}, \tag{1.9}$$

where $C_{s,d}$ is a finite positive constant independent of Ω and p, and $1/0 = +\infty$. Furthermore, if $\mathcal{H}_d(\Omega) > 0$ and $\{\omega_N\}_{N\geq 2}$ is any sequence of N-point configurations on Ω satisfying

$$\lim_{N\to\infty} \frac{E_s(\omega_N)}{\mathcal{E}_s(\Omega, N)} = 1, \tag{1.10}$$

then

$$\frac{1}{N} \sum_{\boldsymbol{x}\in\omega_N} \delta_{\boldsymbol{x}} \overset{*}{\rightharpoonup} \frac{d\mathcal{H}_d}{\mathcal{H}_d(\Omega)}, \quad N \to \infty. \tag{1.11}$$

This theorem is sometimes described as the Poppy-seed bagel theorem (PSB), a name that alludes to discrete equilibrium configurations on the torus. It first appeared in [63, Theorem 2.1], and in the present generality in [16, Theorems 1–3].

In particular, the theorem holds for any compact $\Omega \subset \mathbb{R}^d$ as well as any compact subset of a smooth d-dimensional manifold. To be consistent with the notation of (1.9), we define $C_{d,d}$

according to (1.8):

$$C_{d,d} := \mathcal{H}_d(\mathbb{B}^d) = \frac{\pi^{d/2}}{\Gamma\left(d/2+1\right)}, \quad d \geq 1, \tag{1.12}$$

where Γ is the standard gamma function. It is known for $d = 1,\ s > 1$ that

$$C_{s,1} = 2\zeta(s), \quad s > 1, \tag{1.13}$$

where ζ is the Riemann zeta function, see e.g. [79], and the claimed result of universal optimality of E_8 and the Leech lattice, implies in particular

$$C_{s,d} = |\Lambda_d|^{s/d}\zeta_{\Lambda_d}(s), \qquad s > d, \tag{1.14}$$

for $d = 8, 24$. This claim has been recently announced by Cohn, Kumar, Miller, Radchenko, and Viazovska, who are using the methods of [36] by the same authors. Here Λ_d denotes the lattice achieving optimal packing in these two dimensions: E_8 and the Leech lattice, respectively; $|\Lambda_d|$ stands for the volume of the corresponding fundamental cell and ζ_{Λ_d} is the corresponding Epstein zeta-function. However, the exact value of $C_{s,d}$ is unknown for all $d \geq 2$, $d \neq 8, 24$. In the cases $d = 2,\ 4$, the conjectured value is also described by (1.14) with Λ_d, respectively, the hexagonal and D_4 lattices; see [22, Conjecture 2]. It was shown [22, Proposition 1] that the conjectured values (1.14) serve as upper bounds for their respective $C_{s,d}$. For more details, consult [64] and [22].

One way of generalizing Theorem 1.3 so that it yields non-uniform limiting distributions was studied in [16], where the Riesz potential is multiplied by a weight satisfying semicontinuity conditions. More precisely, one minimizes the energy

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, \kappa, 0) = \sum_{\boldsymbol{x} \neq \boldsymbol{y} \in \omega_N} \frac{\kappa(\boldsymbol{x}, \boldsymbol{y})}{\|\boldsymbol{x} - \boldsymbol{y}\|^s},$$

for a non-negative weight function κ on $\Omega \times \Omega$. Our present goal is to develop an alternate approach by introducing an external field equipped with a suitable scaling factor that depends on the number of points N.

With regard to practical implementation, it is worth mentioning that by using a localized weight $\kappa(\cdot\,,\cdot) := \kappa_N(\cdot\,,\cdot)$ that depends on the number of points, one can lower the computational complexity of $E_s^w(\omega_N)$. This approach is investigated in [20], and we shall return to it in Chapter 4. On the other hand, a number of papers are dedicated to producing well-distributed discrete configurations by drawing them from a suitable random process with, perhaps, further local optimization, see for example [1], [5], [78].

1.5 Γ-convergence

Let X be a metric space and denote by $\mathcal{U}$ the topology on X corresponding to its metric. Suppose that the functionals F, $F_n : X \to \mathbb{R}$, $n \geq 1$, satisfy

$\mathbf{1}^\Gamma$. for every sequence $\{x_n\} \subset X$ such that $x_n \to x$, $n \to \infty$, there holds $\liminf_{n\to\infty} F_n(x_n) \geq F(x)$;

$\mathbf{2}^\Gamma$. for every $x \in X$ there exists a sequence $\{x_n\} \subset X$ converging to it and such that $\lim_{n\to\infty} F_n(x_n) = F(x)$.

We shall then say that the sequence $\{F_n\}$ is Γ-*converging* to the functional F on $(X,\mathcal{U})$ (on X *equipped* with topology $\mathcal{U}$); in symbols, $F_n \xrightarrow{\Gamma} F$. The sequence in $\mathbf{2}^\Gamma$ is called the *recovery sequence*. The notion of Γ-convergence originated in the theory of elliptic operators, and was developed by De Giorgi in the 1970s [21]. Its primary value for our purposes consists in that in conjunction with certain compactness property, it guarantees that minimizers of F_n converge to those of F. Namely, suppose that there exists a compact subset K of the metric space X, such that $\inf_X F_n = \inf_K F_n$ for all $n \geq 1$. Then the following result holds.

Theorem 1.4. *If a sequence of functionals $\{F_n\}$ on a metric space X satisfies the above compactness property and Γ-converges to F, then*

1. *F attains its minimum and $\min F = \lim_{n\to\infty} \inf F_n$*
2. *if $\{x_n\}$ is a sequence of (global) minimizers of F_n, converging to an $x \in X$, then x is a (global) minimizer for F.*

The proof is straightforward; see for example [21, Theorem 2.1]. In the context of energy minimization, the metric space $X = \mathcal{P}(\Omega)$ is chosen to be the space of all probability measures on Ω with the weak* topology; this topology is metrizable, and X is itself compact by the Banach-Alaoglu theorem.

To define the sequence of functionals on $\mathcal{P}(\Omega)$, recall that the set $\mathcal{P}_N(\Omega) \subset \mathcal{P}(\Omega)$ consists of measures of the form $\frac{1}{N}\sum_{i=1}^N \delta_{\boldsymbol{x}_i}$, and that the Riesz energy functional $(\star)$

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, q) = \sum_{i\neq j} g_s(\boldsymbol{x}_i, \boldsymbol{x}_j) + \tau(N) \sum_i q(\boldsymbol{x}_i),$$

is defined (only) on $\mathcal{P}_N(\Omega)$. An extension to the whole $\mathcal{P}(\Omega)$ can now be written as

$$\frac{1}{\mathcal{T}(N)} E_N(\mu; g_s, q) := \begin{cases} \frac{1}{\mathcal{T}(N)} E(\{\boldsymbol{x} \,:\, \boldsymbol{x} \in \operatorname{supp}\mu\}; g_s, q) & \text{if } \mu \in \mathcal{P}_N(\Omega), \\ +\infty, & \text{otherwise.} \end{cases} \tag{1.15}$$

Here as before, the scaling factor $\tau(N) = \tau_{s,d}(N)$ is defined by (1.5); $\mathcal{T}(N) = N\tau(N)$. The same construction applies to an any kernel with a singularity on the diagonal, provided that its asymptotics is known. Thus, the remaining part of this section, as well as the results of Chapter 2 can be generalized to similar (approximately) scale-invariant kernels, as for example the weighted kernels $E(\cdot\,; g_s, \kappa, q)$ of [15], see (1.4). It is known in the case of an integrable kernel g that the sequence of minimizers of the extended functionals $E_N(\mu; g_s, q)$ converges to the minimizer of the continuous energy functional ($\star\star$), and also that the functionals themselves Γ-converge to ($\star\star$). Similarly, the results of Chapter 2 yield in particular a proof that in the hypersingular case $s \geq d$, (1.15) also Γ-converges to a functional on absolutely continuous measures in $\mathcal{P}(\Omega)$ (the absolute continuity here is w.r.t. $\mathcal{H}_d$ on Ω), defined by

$$S(\mu; g_s, q) := \begin{cases} C_{s,d} \int_\Omega \varphi(\boldsymbol{x})^{1+s/d}\, d\mathcal{H}_d(\boldsymbol{x}) + \int_\Omega q(\boldsymbol{x})\varphi(\boldsymbol{x})\, d\mathcal{H}_d(\boldsymbol{x}), & \mu = \varphi\, d\mathcal{H}_d, \\ +\infty, & \text{otherwise.} \end{cases} \tag{1.16}$$

Here φ is the Radon-Nikodym derivative; $C_{s,d}$ stands for the constant from the PSB Theorem 1.3. Recall that a d-rectifiable set is a Lipschitz image of a closed d-dimensional set; a d-regular set satisfies inequalities (1.2). Recall also the classical integral functional $I(\cdot\,; g_s, q)$, defined in ($\star\star$). Gamma-convergence properties of the first-order asymptotics of the Riesz kernel can be summarized in the following propositions.

Proposition 1.5. *Suppose the set Ω is d-rectifiable and d-regular. If q is continuous on Ω, and the kernel g_s is the integrable Riesz kernel $g_s(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$ with $s < d = \dim_H \Omega$, then*

$$\frac{1}{N^2} E_N(\cdot\,; g_s, q) \xrightarrow{\Gamma} I(\cdot\,; g_s, q), \qquad N \to \infty,$$

on $\mathcal{P}(\Omega)$ equipped with the weak topology.*

Proposition 1.6. *Suppose Ω is a d-rectifiable set. If q is continuous on Ω and kernel g_s is the hypersingular Riesz kernel, $g_s(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$ with $s > d = \dim_H \Omega$, then*

$$\frac{1}{N^{1+s/d}} E_N(\cdot\,; g_s, q) \xrightarrow{\Gamma} S(\cdot\,; g_s, q), \qquad N \to \infty,$$

on $\mathcal{P}(\Omega)$ equipped with the weak topology.*

The remaining part of this section is dedicated to the proofs of these two statements. Proposition 1.5 follows from a well-known [96] argument, that applies to a general integrable kernel, whereas Proposition 1.6 can be seen as a reformulation of the results of [65]. Their comparison reveals the difference in the asymptotic structures of energies for $s < d$ and $s \geq d$: constructing a recovery sequence for the integrable case requires a delicate argument that we present in Section 1.5.1, whereas a hypersingular recovery sequence is produced by simply

taking the minimizers of the corresponding functional; it should however be noted that the latter approach is possible only provided the asymptotic analysis in Chapter 2 is in place. This distinction illustrates certain rigidity that the hypersingular minimizers possess, which makes them especially appealing for various applications, see Chapter 4.

1.5.1 Γ-convergence for the integrable case

In the present section we shall outline a proof of Γ-convergence in order to contrast it with the corresponding proofs in the hypersingular case, see Sections 1.5.2 and 2.5. The exposition here follows [96]; we shall omit some technical details, since all the following chapters deal exclusively with the hypersingular case, and refer the interested reader to [94, 96]. The reader should compare the proof of Γ-convergence for an integrable Riesz kernel given here, with the proof of Proposition 1.6 in the following section, covering the hypersingular case. Although the d-dimensional Hausdorff measure plays an important part in both proofs, its roles are quite different; it serves as an estimate for the Γ-limit in the former, and appears as the unique minimizer of a convex variational functional in the latter. Not surprisingly, it is the reduction to the variational argument that takes the most effort in the hypersingular case; the subsequent treatment is easier due to convexity of the resulting Γ-limit.

Proof of Proposition 1.5. Suppose $\{\mu_N\} \subset \mathcal{P}$ converges weak* to $\mu \in \mathcal{P}(\Omega)$; we shall first verify the property $\mathbf{1}^\Gamma$. Note that if $\{\mu_N\}$ does not contain a subsequence of the form $\{\mu_n \in \mathcal{P}_n(\Omega)\}$, the inequality in $\mathbf{1}^\Gamma$ is trivial. Thus, without loss of generality, $\mu_N \in \mathcal{P}_N(\Omega)$.

Observe that, since the term $\int_\Omega q d\mu$ is linear in μ, the functional $I(\mu; g_s, q)$ is lower semicontinuous in μ by the semicontinuity of g_s and q, see $\mathbf{1}^\ddagger$. Property $\mathbf{1}^\Gamma$ then follows simply from the lower semicontinuity of $I(\mu; g_s, q)$ in μ, and the representation (see $(\star)$ and $(\star\star)$ for the difference in scalings)

$$\frac{1}{N^2} E_N(\mu_N\,; g_s, q) = I(\mu_N\,; g_s, q). \tag{1.17}$$

To verify property $\mathbf{2}^\Gamma$, fix a probability measure $\mu \in \mathcal{P}(\Omega)$; we shall construct a sequence μ_N converging weak* to μ. Suppose the density $d\mu/\mathcal{H}_d$ is continuous on Ω; this is not a restrictive requirement, since μ can be approximated with a mollifying sequence, and then a diagonal argument applies to the latter. We shall further assume for simplicity that the parametrization of Ω, denoted by $\psi : \mathcal{C}^d \to \Omega$, is bi-Lipschitz and defined on the unit cube $\mathcal{C}^d \subset \mathbb{R}^d$.

Following [96], we shall choose a sequence $s_N, N \geq 1$, such that $N^{-1/d} = o(s_N)$ and $o(s_N) = o(1)$, $N \to \infty$, and split the cube $\mathcal{C}^d$ into equal cubes $\mathcal{C}_m$, $1 \leq m \leq M$ with side length in the interval $[s_N, 2s_N)$. By [94, Lemma 7.4, p.143], for every N there exist numbers n_m such that $\sum_m n_m = N$ and

$$|n_m - N\mu[\psi(\mathcal{C}_m)]| < 1, \qquad 1 \leq m \leq M, \tag{1.18}$$

where we used that the mapping ψ is a bijection. It is further possible to place n_m points $\{\boldsymbol{x}_{j,m}\}_{j=1}^{n_m}$ in each cube $\mathcal{C}_m$, so that for any pair of distinct points $\boldsymbol{x}_{i,l}$, $\boldsymbol{x}_{j,m} \in \mathcal{C}^d$ there holds

$$\|\boldsymbol{x}_{i,m} - \boldsymbol{x}_{j,m}\| \geq \frac{3c}{N^{1/d}}, \tag{1.19}$$

for a constant c independent of N. (An explicit construction of such points is given by the hypersingular Riesz minimizers, see Theorem 2.7.) Note that by (1.18), the counting measures of the discrete collections constructed in this way converge weak* to μ; denote such measures by μ_N, $N \geq 2$, for the rest of this proof. We shall verify that they satisfy the equality in $\mathbf{2}^\Gamma$. Note that by the continuity of q, convergence of the external field term is readily implied by the weak* convergence; it therefore suffices to deal with the case of $q \equiv 0$.

By (1.17) and the above reasoning, we need to show that

$$\lim_{N\to\infty} I(\mu_N; g_s, 0) = I(\mu_N; g_s, 0)$$

or, in the integral notation,

$$\lim_{N\to\infty} \iint_{\Omega\times\Omega} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}) = \iint_{\Omega} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}).$$

Let us fix $\delta > 0$ and write D_δ for the δ-neighborhood of diag $(\Omega \times \Omega)$ in l_∞ product metric. Splitting the integral in the LHS gives

$$\iint_{\Omega} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}) = \iint_{D_\delta} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}) + \iint_{D_\delta^c} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}),$$

where M^c denotes the complement of the set M. Since $g_s(\boldsymbol{x},\boldsymbol{y})$ is continuous outside D_δ, $\mu_N \overset{*}{\rightharpoonup} \mu$ implies

$$\lim_{N\to\infty} \iint_{D_\delta^c} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu_N(\boldsymbol{x}) d\mu_N(\boldsymbol{y}) = \iint_{D_\delta^c} g_s(\boldsymbol{x},\boldsymbol{y}) d\mu(\boldsymbol{x}) d\mu(\boldsymbol{y}).$$

To finish the proof, we will show that the integral $I(\mu_N\,; g_s, 0)$ over D_δ vanishes as $\delta \to 0$ uniformly in N. This is clearly the case for $I(\mathcal{H}_d; g_s, 0)$, where $\mathcal{H}_d$ is restricted to Ω, so it suffices to estimate the integral with respect to μ_N by that with respect to $\mathcal{H}_d$. By (1.19), there holds $B(\boldsymbol{x}_i, cN^{-1/d}) \cap B(\boldsymbol{x}_j, cN^{-1/d})$, $i \neq j$, and for any pair of points $\boldsymbol{x}$, $\boldsymbol{y}$ from two such distinct balls, $\|\boldsymbol{x} - \boldsymbol{y}\| \leq 2\|\boldsymbol{x}_i - \boldsymbol{x}_j\|$. From the lower bound of the d-regularity of Ω,

$$\mathcal{H}_d[\Omega \cap B(\boldsymbol{x}_i, cN^{-1/d})] \geq c_0 c^d N^{-1}, \qquad 1 \leq i \leq N.$$

This allows to obtain the aforementioned estimate in terms of $I(\mathcal{H}_d; g_s, 0)$ as follows. For large

enough N, $2cN^{-1/d} < \delta$, thus, writing $B_i := B(\boldsymbol{x}_i, cN^{-1/d})$, we have

$$\begin{aligned}\iint_{D_\delta} g_s(\boldsymbol{x},\boldsymbol{y})d\mu_N(\boldsymbol{x})d\mu_N(\boldsymbol{y}) &= \frac{1}{N^2}\sum_{\substack{\|\boldsymbol{x}_i-\boldsymbol{x}_j\|<\delta\\ i\neq j}} g_s(\boldsymbol{x}_i,\boldsymbol{x}_j)\\ &\leq \frac{2^s}{c_0^2c^d}\sum_{\substack{\|\boldsymbol{x}_i-\boldsymbol{x}_j\|<\delta\\ i\neq j}}\iint_{B_i\times B_j} g_s(\boldsymbol{x},\boldsymbol{y})d\mathcal{H}_d(\boldsymbol{x})d\mathcal{H}_d(\boldsymbol{y})\\ &\leq \frac{2^s}{c_0^2c^d}\iint_{D_{2\delta}} g_s(\boldsymbol{x},\boldsymbol{y})d\mathcal{H}_d(\boldsymbol{x})d\mathcal{H}_d(\boldsymbol{y}).\end{aligned}$$

As explained above, the RHS tends to zero when $\delta \to 0$. □

Remark 1.7. In the above argument, it was essential that the kernel is continuous away from diag $(\Omega \times \Omega)$ and monotone decreasing as a function of distance; a similar argument therefore applies to any integrable kernel satisfying these assumptions.

1.5.2 Γ-convergence for the hypersingular case

Before proceeding to prove Proposition 1.6, we shall rewrite (2.5) to show the provenance of (1.16), and also derive from variational principles that $S(\cdot\,; g_s, q)$ is indeed minimized by the measure μ^q in (2.5). As follows from (2.5)–(2.6), the expression for the density $d\mu^q/d\mathcal{H}_d(\boldsymbol{x})$,

$$\varphi^q(\boldsymbol{x}) = \frac{d\mu^q}{d\mathcal{H}_d}(\boldsymbol{x}) = \left(\frac{L_1 - q(\boldsymbol{x})}{C_{s,d}(1+s/d)}\right)_+^{d/s}, \tag{1.20}$$

where the constant L_1 is chosen so that $\varphi^q\, d\mathcal{H}_d$ is a probability measure, allows to rewrite the energy asymptotics (2.6) as

$$\begin{aligned}\lim_{N\to\infty}\frac{E(\hat{\boldsymbol{x}}_1,\ldots,\hat{\boldsymbol{x}}_N; g_s, q)}{\mathcal{T}_{s,d}(N)} &= \int_\Omega \frac{L_1 + sq(\boldsymbol{x})/d}{1+s/d}\, d\mu^q(\boldsymbol{x})\\ &= C_{s,d}\int_\Omega \varphi^q(\boldsymbol{x})^{1+s/d}\, d\mathcal{H}_d(\boldsymbol{x}) + \int_\Omega q(\boldsymbol{x})\varphi^q(\boldsymbol{x})\, d\mathcal{H}_d(\boldsymbol{x})\\ &= S(\mu^q; g_s, q),\end{aligned}$$

where $\{\hat{\boldsymbol{x}}_i\}$ is a collection minimizing $E(\cdot\,; g_s, q)$ on $\mathcal{P}_N(\Omega)$. This implies Γ-convergence to (1.16) on sequences of minimizers. The limit in the LHS is denoted by $S(q, \Omega)$ in (2.6), Chapter 2.

Let us now apply a variational argument to show that the density φ^q minimizes $S(\varphi d\mathcal{H}_d; g_s, q)$; for brevity, we shall write simply $S(\varphi; g_s, q)$. First, observe that $S(\varphi; g_s, q)$ is convex in φ, and must therefore attain a minimum on the set $\int_\Omega \varphi d\mathcal{H}_d = 1$. Suppose $\hat{\varphi}$ is a minimizer, and let $\delta\varphi$

satisfy $\int_\Omega \delta\varphi d\mathcal{H}_d = 0$. Taking Gâteaux derivative at $\hat{\varphi}$ in the direction $\delta\varphi$ gives

$$\begin{aligned} d[S(\hat{\varphi}; g_s, q); \delta\varphi] &= \lim_{t\to 0} \frac{1}{t} \left[S\left((\hat{\varphi} + \delta\varphi); g_s, q\right) - S\left(\hat{\varphi}; g_s, q\right)\right] \\ &= \lim_{t\to 0} \frac{1}{t} \left[\int_\Omega \delta\varphi(\boldsymbol{x}) t \left(C_{s,d}(1 + s/d)\hat{\varphi}(\boldsymbol{x})^{s/d} + q(\boldsymbol{x}) \right) d\mathcal{H}_d(\boldsymbol{x}) + o(t) \right] \\ &= \int_\Omega \delta\varphi(\boldsymbol{x}) \left(C_{s,d}(1 + s/d)\hat{\varphi}(\boldsymbol{x})^{s/d} + q(\boldsymbol{x}) \right) d\mathcal{H}_d(\boldsymbol{x}). \end{aligned}$$

Since $\hat{\varphi}$ is a minimizer, equality $d[S(\hat{\varphi}; g_s, q); \delta\varphi] = 0$ must hold for every $\delta\varphi$ such that $\int_\Omega \delta\varphi \, d\mathcal{H}_d = 0$. This implies that the factor

$$f(\boldsymbol{x}) := C_{s,d}(1 + s/d)\hat{\varphi}(\boldsymbol{x})^{s/d} + q(\boldsymbol{x})$$

is an $\mathcal{H}_d$-a.e. constant. Indeed, let $\delta\varphi = f - \int_\Omega f d\mathcal{H}_d / \mathcal{H}_d(\Omega)$; there holds $\int_\Omega \delta\varphi \, d\mathcal{H}_d = 0$, therefore

$$\begin{aligned} 0 &= \int_\Omega \delta\varphi \, f d\mathcal{H}_d \\ &= \int_\Omega \delta\varphi \, \hat{\varphi} \, d\mathcal{H}_d - \int_\Omega \frac{\delta\varphi}{\mathcal{H}_d(\Omega)} \left(\int_\Omega f d\mathcal{H}_d \right) d\mathcal{H}_d \\ &= \int_\Omega \left(f - \frac{1}{\mathcal{H}_d(\Omega)} \int_\Omega f d\mathcal{H}_d \right)^2 d\mathcal{H}_d, \end{aligned}$$

as desired. Let $L_1 = \int_\Omega f d\mathcal{H}_d$ so that $f = L_1$ $\mathcal{H}_d$-a.e.; then the definition of f and (1.20) imply $\hat{\varphi} = \varphi^q$ $\mathcal{H}_d$-a.e. Note that we did not make any smoothness assumptions on $\hat{\varphi}$ except that all the above integrals exist.

Proof of Proposition 1.6. To verify the property $\mathbf{1}^\Gamma$ of the definition of Γ-convergence, suppose a sequence $\{\mu_N\} \subset \mathcal{P}(\Omega)$ weak* converges to $\mu \in \mathcal{P}(\Omega)$; observe that if

$$\liminf_{N\to\infty} \frac{1}{\mathcal{T}(N)} E_N(\mu_N; g_s, q) = +\infty \geq S(\mu; g_s, q),$$

so the inequality in $\mathbf{1}^\Gamma$ holds trivially, and its proof is complete. It will therefore suffice to assume that the limit in the last equation is finite. In particular, $\{\mu_N\}$ must contain a subsequence comprising only elements from $\mathcal{P}_N(\Omega)$, so passing to a subsequence if necessary we now suppose that the sequence of discrete measures μ_N, $N \geq 2$, is such that

$$\mu_N \xrightarrow{*} \mu, \qquad \mu_N \in \mathcal{P}_N(\Omega).$$

Furthermore, Lemma 2.25 implies that μ must be absolutely continuous with respect to $\mathcal{H}_d$ for the above limit to be finite. Note also that this explains setting $S(\cdot\,; g_s, q)$ to $+\infty$ on measures with nontrivial singular components in the Lebesgue decomposition with respect to $\mathcal{H}_d$, as $+\infty$

is the only value for which construction of a recovery sequence is possible.

By Theorem 2.3, provided continuity of the density $d\mu/d\mathcal{H}_d$, which we shall assume for simplicity, there exists a function q^* such that the minimizers of (1.15) also converge to μ; denote them by μ_N^*, $N \geq 2$. Furthermore, it follows from Theorem 2.1 that for the sequence $\{\mu_N^*\}$, there holds

$$\frac{1}{N^{1+s/d}} E_N(\mu_N^*; g_s, q^*) \to S(\mu\,; g_s, q^*), \qquad N \to \infty.$$

Note that both q and q^* are continuous; the latter follows from the continuity of the density of μ. Using this together with the weak* convergence of the two sequences μ_N, μ_N^* to μ, we can conclude

$$\begin{aligned}
\liminf_{N\to\infty} \frac{1}{N^{1+s/d}} E_N(\mu_N; g_s, q) &= \liminf_{N\to\infty} \frac{1}{N^{1+s/d}} E_N(\mu_N; g_s, q^*) + \int_\Omega (q - q^*)\, d\mu \\
&\geq \lim_{N\to\infty} \frac{1}{N^{1+s/d}} E_N(\mu_N^*; g_s, q^*) + \int_\Omega (q - q^*)\, d\mu \\
&= S(\mu\,; g_s, q^*) + \int_\Omega (q - q^*)\, d\mu = S(\mu\,; g_s, q).
\end{aligned}$$

This gives property $\mathbf{1}^\Gamma$ in the definition of Γ-convergence.

To establish property $\mathbf{2}^\Gamma$, first note that taking μ_N^* as the recovery sequence immediately implies $\mathbf{2}^\Gamma$ for any μ such that $S(\mu\,; g_s, q) < +\infty$. On the other hand, $S(\mu\,; g_s, q) = +\infty$ implies that μ contains a nonzero singular component with respect to $\mathcal{H}_d$, and by Lemma 2.25,

$$\lim_{N\to\infty} \frac{1}{N^{1+s/d}} E_N(\mu_N; g_s, q) = +\infty = S(\mu\,; g_s, q)$$

for every sequence $\{\mu_N\}$ of probability measures weak* converging to μ. □

Remark 1.8. Note that the assumption of continuity of density $d\mu/d\mathcal{H}_d$ is not very restrictive; the general case can be reduced to it by a mollification argument, as outlined in the proof of Proposition 1.5.

Remark 1.9. The above proof shows also that any sequence $\{\mu_N\} \subset \mathcal{P}(\Omega)$ with an optimal asymptotics

$$\lim_{N\to\infty} \frac{1}{\mathcal{T}(N)} E_N(\mu_N; g_s, q)$$

must converge weak* to the measure μ^q. This is a corollary of the convexity of the functional $S(\cdot\,; g_s, q)$, by which it has a unique minimizer, and property $\mathbf{1}^\Gamma$ of the Γ-convergence.

Chapter 2

First-order term of the hypersingular Riesz energy with an external field

In the language of Section 1.5, the present chapter contains the proof of Γ-convergence of the functional ($\star$) (also in display (2.1) below) for a hypersingular kernel g_s in the presence of an external field. We derive the first-order asymptotics of minimal values of the functional, as well as the weak* limit of minimizing configurations; local separation and covering properties are established as well.

The outline of the chapter is as follows. In the following section we specialize the notation from Chapter 1 to the case of hypersingular Riesz kernel in the presence of an external field. Section 2.2 contains an extension of the PSB Theorem 1.3 to this case; it also includes results on separation and covering of minimizing configurations. We discuss several numerical examples in Section 2.3. Finally, Section 2.4 contains proofs of the results stated in Section 2.2. This chapter reproduces with minor modifications the contents of a paper by Douglas Hardin, Edward Saff, and the author [66], published in the SIAM Journal on Mathematical Analysis.

2.1 Overview

We are concerned with the problem of minimizing the discrete Riesz s-energy of N particles constrained to a compact subset Ω of $\mathbb{R}^p$ of Hausdorff dimension d under the influence of an external field $q(\boldsymbol{x})$. More precisely, we minimize

$$E^q_{s,d}(\omega_N) := E(\boldsymbol{x}_1, \dots \boldsymbol{x}_N; g_s, q) = \sum_{\substack{\boldsymbol{x}\neq\boldsymbol{y}\\ \boldsymbol{x},\boldsymbol{y}\in\omega_N}} \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + \tau(N) \sum_{\boldsymbol{x}\in\omega_N} q(\boldsymbol{x}), \quad s \geq d, \tag{2.1}$$

where from (1.5),

$$\tau(N) = \begin{cases} N^{s/d}, & s > d, \\ N \log N, & s = d, \end{cases}$$

over N-element subsets $\omega_N = \{\boldsymbol{x}_1, \dots \boldsymbol{x}_N\} \subset \Omega$.

The factor $\tau(N)$ is chosen so that the two terms on the RHS of (2.1) have the same order of growth as $N \to \infty$, equal to $\mathcal{T}(N) = \mathcal{T}_{s,d}(N) = N\tau_{s,d}(N)$. Since we shall fix the values of $s\,, d$ in all arguments, we shall often write simply $\mathcal{T}(N)$ if there is no risk of confusion. Here we consider only the case when s is chosen greater than or equal to the dimension of the set Ω because for $s < d$ such external field problems come under the umbrella of classical potential theory and have been well studied, as described in Chapter 1.

One motivation to consider this energy expression is that (under mild conditions on the set Ω) for any probability measure μ on Ω that is absolutely continuous with respect to the d-dimensional Hausdorff measure restricted to Ω, there is an easily described external field $q(\boldsymbol{x})$ for which the normalized counting measures of the minimizers of (2.1) weak* converge to μ (formal definitions are given in the next two subsections).

For $s > d$ the minimizers of (2.1) are shown to have optimal orders of separation and covering as $N \to \infty$. Minimization of (2.1) therefore provides well-distributed nodes on compact sets, which can be used for a number of applications; for example, meshless methods [84], halftoning [106] and sensor deployment [70].

Recall that external fields arise in the *Gauss variational problem*, which involves minimizing the functional introduced in equation ($\star\star$),

$$I(\mu\,; g, q) = \iint\limits_{\Omega} g(\boldsymbol{x}, \boldsymbol{y}) d\mu(\boldsymbol{x}) d\mu(\boldsymbol{y}) + \int\limits_{\Omega} q(\boldsymbol{x}) d\mu(\boldsymbol{x}),$$

for a pair of fixed integrable lower semicontinuous functions $g : \Omega \times \Omega \to \mathbb{R} \cup \{+\infty\},\ q : \Omega \to \mathbb{R} \cup \{+\infty\}$ over the probability measures $\mu \in \mathcal{P}(\Omega)$ supported on Ω. The classical work of Ohtsuka [85] deals with this question when Ω is locally compact and g is integrable on Ω. The case $g(\boldsymbol{x}, \boldsymbol{y}) = \log \frac{1}{\|\boldsymbol{x}-\boldsymbol{y}\|}$ and a number of its applications to constructive analysis are extensively treated in the book [93] by Saff and Totik. More recently, the question of solvability of the Gauss variational problem was considered by Zorii in [108] and [109]; several references to general monographs on the subject can be found in Section 1.4.1.

As follows from (1.5), for an integrable kernel g, the problem of minimization of ($\star$) over N-element sets $\omega_N \subset \Omega$ deals with

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g, q) = \sum_{\substack{\boldsymbol{x} \neq \boldsymbol{y} \\ \boldsymbol{x}, \boldsymbol{y} \in \omega_N}} g(\boldsymbol{x}, \boldsymbol{y}) + N \sum_{\boldsymbol{x} \in \omega_N} q(\boldsymbol{x}), \qquad N = 2, 3, \ldots. \tag{2.2}$$

Such problems have been studied by Petrache, Rougerie and Serfaty in [87], [92] for the Riesz s-kernel

$$g_s(\boldsymbol{x}, \boldsymbol{y}) = \|\boldsymbol{x} - \boldsymbol{y}\|^{-s}, \tag{2.3}$$

with $s < d$, and for the logarithmic kernel in [93]. An earlier series of papers [41], [23] and [24] by Brauchart, Dragnev, and Saff explores minima of (2.2) when Ω is a d-dimensional sphere and $d - 2 \leq s < d$. The paper [8] by Bilogliadov discusses minimizing ($\star\star$) with the Riesz kernel for $s = 1$ and a rotationally symmetric q over probability measures supported on the 2-dimensional unit sphere.

We consider configurations of points restricted to a compact set $\Omega \subset \mathbb{R}^p$, such that $\mathcal{H}_d(\Omega) > 0$, $d \leq p$. The *external field* $q : A \to (-\infty, \infty]$ is assumed to be lower semicontinuous and finite on a

subset of Ω of positive $\mathcal{H}_d$-measure. We write $\mathring{M}$ for the interior of a set $M \subset \mathbb{R}^p$, and $\overline{M}$ for its closure. For a real number t, let $(t)_+ := \max(0,\, t)$. The closed ball in $\mathbb{R}^p$ of radius r centered at the point $\boldsymbol{x}$ is denoted by $B(\boldsymbol{x}, r)$. Notation $L^1(\Omega, \mu)$ stands for the class of functions integrable on the set Ω with respect to measure μ.

The *minimal* (s, d, q)*-energy* of the set Ω over all N-point subsets of Ω is given by (see (2.1))

$$\mathcal{E}^q_{s,d}(\Omega, N) := \inf\{E^q_{s,d}(\omega_N) : \omega_N \subset \Omega,\ \#\omega_N = N\}, \tag{2.4}$$

where $\#S$ denotes the cardinality of a set S. Since q is lower semicontinuous and Ω is compact, there exists a configuration of N charges $\hat{\omega}_N$ for which the infimum in (2.4) is attained; i.e.,

$$E^q_{s,d}(\hat{\omega}_N) = \mathcal{E}^q_{s,d}(\Omega, N).$$

Such a configuration $\hat{\omega}_N$ will be called an *N-point (s, d, q)-energy minimizer on Ω*.

2.2 Main results

2.2.1 A Poppy-seed bagel theorem for (s, d, q)-energy

The following two results extend Theorem 1.3 to (s, d, q)-energy.

Theorem 2.1. *Assume $0 < d \le p$, and $s \ge d$. Let $\Omega \subset \mathbb{R}^p$ be a d-rectifiable compact set, $\mathcal{H}_d(\Omega) > 0$, and in the case $s = d$ require additionally that Ω be a subset of a d-dimensional C^1-manifold. Further assume that q is a lower semicontinuous function on Ω and finite on a set of positive $\mathcal{H}_d$-measure. Define L_1 and μ^q to be the positive constant and probability measure determined, respectively, by*

$$\int_\Omega \left(\frac{L_1 - q(\boldsymbol{x})}{C_{s,d}(1 + s/d)}\right)^{d/s}_+ d\mathcal{H}_d(\boldsymbol{x}) = 1, \qquad d\mu^q := \left(\frac{L_1 - q(\cdot)}{C_{s,d}(1 + s/d)}\right)^{d/s}_+ d\mathcal{H}_d, \tag{2.5}$$

where $C_{s,d}$ for $s \ge d$ is the same as in Theorem 1.3. Then

$$\lim_{N\to\infty} \frac{\mathcal{E}^q_{s,d}(\Omega, N)}{\mathcal{T}_{s,d}(N)} = S(q, \Omega) := \int \frac{L_1 + sq(\boldsymbol{x})/d}{1 + s/d}\, d\mu^q(\boldsymbol{x}). \tag{2.6}$$

Furthermore, if $\{\omega_N\}_{N\ge 2}$ is any sequence of asymptotically (s, d, q)-energy minimizing configurations on Ω; that is,

$$\lim_{N\to\infty} \frac{E^q_{s,d}(\omega_N)}{\mathcal{T}_{s,d}(N)} = S(q, \Omega), \tag{2.7}$$

then

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\omega_N}\delta_{\boldsymbol{x}} \xrightarrow{*} d\mu^q \quad \text{as } N\to\infty. \tag{2.8}$$

Remark 2.2. As with Theorem 1.3, this result holds on the (possibly) larger class of sets Ω satisfying $\mathcal{H}_d(\Omega)=\mathcal{M}_d(\Omega)$, where $\mathcal{M}_d$ is the d-dimensional Minkowski content.

As an application of Theorem 2.1, we deduce a method for constructing a sequence of (s,d,q)-energy minimizing collections $\hat{\omega}_N$ such that their normalized counting measures weak* converge to a given distribution.

Theorem 2.3. *Let the assumptions of Theorem 2.1 on the set Ω and numbers s,d,p hold. Assume further $\rho: A\to[0,\infty)$ is an upper semicontinuous function, such that $\rho\, d\mathcal{H}_d$ is a probability measure. Then the lower semicontinuous function $q: A\to(-\infty,0]$ given by*

$$q(\boldsymbol{x}) := -M_{s,d}\rho(\boldsymbol{x})^{s/d}, \quad \text{where } M_{s,d} := C_{s,d}(1+s/d), \tag{2.9}$$

is such that any sequence $\{\hat{\omega}_N\}_{N\geq 2}$ of (s,d,q)-energy minimizers converges weak* to $\rho\, d\mathcal{H}_d$:

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\hat{\omega}_N}\delta_x \xrightarrow{*} \rho\, d\mathcal{H}_d, \quad N\to\infty. \tag{2.10}$$

In particular, for $s=d$ equation (2.10) *holds with* (2.9) *taking the form*

$$q(\boldsymbol{x}) := -\frac{2\pi^{d/2}\rho(\boldsymbol{x})}{\Gamma\left(d/2+1\right)}. \tag{2.11}$$

Remark 2.4. The reader will no doubt observe that except for the case $d=1$ which is covered in (1.13), the usefulness of the last theorem is limited by the lack of knowledge of the value of $C_{s,d}$ when $d\geq 2$. Fortunately, the limit distribution in equation (2.10) is stable under perturbations of the constant $M_{s,d}$: small error in the value of $C_{s,d}$ used in (2.9) only leads to small errors in the resulting weak* limit of minimizers. We quantify this statement in Proposition 2.5 below.

Another possible way of overcoming this difficulty is modifying the problem of minimizing (2.1) so that the charges are restricted to an unbounded set Ω. It will be addressed in a future work.

Proposition 2.5. *Assume that in Theorem 2.3 one uses an approximate value of $C_{s,d}$ satisfying*

$$C'_{s,d} = (1+\Delta)C_{s,d}$$

with a fixed Δ. Let also $\rho(\boldsymbol{x})\geq\delta>0$ for all $\boldsymbol{x}\in\Omega$ and write $q'(\boldsymbol{x})$ for the external field defined with $C'_{s,d}$ instead of $C_{s,d}$ in (2.9). *Then for $\Delta < M_{s,d}/\left(1+(\|\rho\|_\infty\delta^{-1})^{s/d}\right)$, the weak* limit of*

the (s,d,q')-energy minimizers has density $\rho' = d\mu^{q'}/d\mathcal{H}_d$ satisfying

$$|\rho'(\boldsymbol{x}) - \rho(\boldsymbol{x})| \le \Delta \frac{d(1 + \|\rho\|_\infty^{s/d}/\rho(\boldsymbol{x})^{s/d})}{sM_{s,d}} + o(\Delta), \qquad \Delta \to 0. \tag{2.12}$$

Example 2.6. Consider the problem of minimization of $(4,1,q)$-energy on the interval $[0,2]$, where

$$q = (\boldsymbol{x}-1)^2 + \frac{1}{2}.$$

Formula (1.13) gives the exact value of $C_{4,1}$, which enables us to plot the density of $\mu^q_{[0,2]}$ on $[0,2]$. For comparison, we also plot the densities of asymptotic distributions obtained for non-exact values of $C_{4,1}$ by taking $\Delta = \pm 0.3, \pm 0.25, \pm 0.2, \pm 0.1, \pm 0.05, \pm 0.03$ in Proposition 2.5.

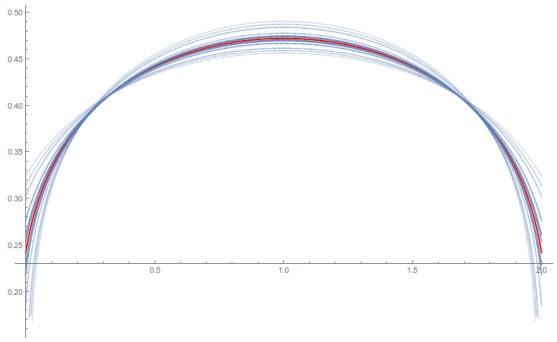

Figure 2.1: The exact graph of density $d\mu^q_{[0,2]}$ is pictured in red, graphs for perturbed values of $C_{4,1}$ are in blue.

2.2.2 Separation and covering properties of minimal configurations

Let

$$\delta(\omega_N) := \min_{\substack{\boldsymbol{x}\neq\boldsymbol{y},\\ \boldsymbol{x},\boldsymbol{y}\in\omega_N}} \|\boldsymbol{x}-\boldsymbol{y}\|,$$

be the *separation distance* of configuration ω_N. We write $\Omega(u) := \{\boldsymbol{x}\in\Omega : q(\boldsymbol{x}) \le u\}$ for a $u \in \mathbb{R}$.

Theorem 2.7. *Let $0 < d \le p$ and $s \ge d$. Let $\Omega \subset \mathbb{R}^p$ be compact with $\mathcal{H}_d(\Omega) > 0$, and let q be a nonnegative lower semicontinuous function on Ω. Then there exists a constant $C(\Omega,s,d,q)$ such that for each N-point (s,d,q)-energy minimizer $\hat{\omega}_N \subset \Omega$*

$$\delta(\hat{\omega}_N) \ge C(\Omega,s,d,q) \begin{cases} N^{-1/d} & s > d, \\ (N\log N)^{-1/d} & s = d, \end{cases} \qquad N \ge 2.$$

To prove Theorem 2.7, we will need the following lemma which is also of independent interest.

For a sequence of configurations $\{\omega_N\}_{N\geq 2}$ we consider the quantity

$$U(\boldsymbol{x},\omega_N) := \sum_{\substack{\boldsymbol{y}\in\omega_N:\\ \boldsymbol{y}\neq\boldsymbol{x}}} \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + q(\boldsymbol{x})\tau_{s,d}(N). \tag{2.13}$$

Lemma 2.8. *Let the assumptions of Theorem 2.7 be satisfied. Then there exists a constant $C(\Omega,s,d,q)$ such that for every (s,d,q)-energy minimizing configuration $\hat{\omega}_N$, $N\geq 2$, and each $\boldsymbol{x}\in\hat{\omega}_N$ there holds*

$$U(\boldsymbol{x},\hat{\omega}_N) \leq C(\Omega,s,d,q)\begin{cases} N^{s/d} & s>d;\\ N\log N & s=d,\end{cases} \qquad N\geq 2. \tag{2.14}$$

Corollary 2.9. *Let the assumptions of Theorem 2.7 hold. Then there exists a constant $C = C(\Omega,s,d,q)$ such that for all $N\geq 2$, the minimizers $\hat{\omega}_N$ are contained in the set $\Omega(C)$.*

Due to this corollary, the sets $\{\hat{\omega}_N\}_{N\geq 2}$ for the problem of minimizing the (s,d,q)-energy on the whole space $\mathbb{R}^p$ are restricted to a compact set, provided that for some compact Ω and a large enough cube $\mathcal{C}^p := [-R,R]^p$ with $\Omega\subset\mathcal{C}^p$, the value C in (2.14) is such that $q(\boldsymbol{x})>C$ for any $\boldsymbol{x}$ not in $\mathcal{C}^p$. Such a problem is then equivalent to energy minimization on $\mathcal{C}^p$ only.

To prove the covering property of (s,d,q)-energy minimizers, we will need the notion of Ahlfors regularity [39, Definition 1.13]. A set $\Omega\subset\mathbb{R}^p$ with $\mathcal{H}_d(\Omega)>0$ is called *d-regular with respect to μ* if there are positive constants c_0, C_0 and a positive locally finite Borel measure μ, such that

$$c_0R^d \leq \mu\big(B(\boldsymbol{x},R)\cap\Omega\big) \leq C_0R^d \tag{2.15}$$

for all $\boldsymbol{x}\in\Omega$ and $0<R\leq\operatorname{diam}(\Omega)$. In the case $\mu=\mathcal{H}_d$, the set Ω is called *Ahlfors regular with dimension d*.

For an $\boldsymbol{x}\in\Omega$ and an N-point collection ω_N define

$$\operatorname{dist}(\boldsymbol{x},\omega_N) := \min_{\boldsymbol{y}\in\omega_N}\|\boldsymbol{y}-\boldsymbol{x}\|,$$

the *covering radius at $\boldsymbol{x}$* with respect to ω_N.

Theorem 2.10. *Let $0<d\leq p$ and $s>d$. Assume $\Omega\subset\mathbb{R}^p$ is compact, d-rectifiable and Ahlfors regular with dimension d. Assume also $q\geq 0$ is a continuous function. Let $\boldsymbol{x}\in\Omega(L_1-h)$ for some $h>0$, where L_1 is defined in Theorem 2.1. Then for each sequence of (s,d,q)-energy minimizers $\{\hat{\omega}_N\}_{N\geq 2}$, there exists a constant $C(\Omega,h,s,d,q)$ such that*

$$\operatorname{dist}(\boldsymbol{x},\hat{\omega}_N) \leq C(\Omega,h,s,d,q)N^{-1/d}, \qquad N\geq 2.$$

A sequence of configurations $\{\omega_N\}_{N\geq 1}$ is said to be *quasi-uniform in* $M \subset \Omega$ if the ratio

$$\gamma(\boldsymbol{x}; \omega_N, \Omega) := \operatorname{dist}(\boldsymbol{x}, \omega_N)/\delta(\omega_N) \tag{2.16}$$

is bounded uniformly for all $\boldsymbol{x} \in M$ and $N \geq 1$. From Theorems 2.7 and 2.10 we have the following result.

Corollary 2.11. *Let $s > d$. Assume $\Omega \subset \mathbb{R}^p$ is compact, d-rectifiable and Ahlfors regular with dimension d. Suppose also that $q : \Omega \to \mathbb{R}$ is a continuous function. Then for any sequence of (s, d, q)-energy minimizers $\{\hat{\omega}_N\}_{N\geq 2}$ on Ω, sequence of subsets $\{\hat{\omega}_N \cap \Omega(L_1 - h)\}_{N\geq 2}$ is quasi-uniform in $\Omega(L_1 - h)$ for any $h > 0$. That is, for some constant $C = C(\Omega, h, s, d, q)$ there holds:*

$$\gamma(\boldsymbol{x}; \omega_N, \Omega(L_1 - h)) \leq C(\Omega, h, s, d, q), \qquad \boldsymbol{x} \in \Omega(L_1 - h),\ N \geq 2.$$

2.3 Examples and numerics

All the results of this section were obtained by using default Mathematica routines (*FindMinimum*) to minimize the energy functional, starting with a randomly generated collection of point charges. We will write $L_1(q, \Omega)$ to show explicitly the set on which we are solving the minimization problem and the external field acting on it.

In this section $\mathbf{e}_z := (0,\ 0,\ 1)^T$ is the basis vector.

Example 2.12. Consider the problem of minimizing (2.1) with $s = 2$ and an external field

$$q_{\mathrm{a}}(\boldsymbol{x}) = \cos(3 \arccos\langle \boldsymbol{x}, \mathbf{e}_z\rangle)^{16}$$

on the unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$. According to (1.12), $C_{2,2} = \pi$. Equation (2.5) for $L_1(q_{\mathrm{a}}, \mathbb{S}^2)$ in this case is

$$\int_{\mathbb{S}^2} \left(\frac{L - q_{\mathrm{a}}(\boldsymbol{x})}{2\pi} \right)_+ d\mathcal{H}_2(\boldsymbol{x}) = 1, \tag{2.17}$$

solving it for L gives $L_1(q_{\mathrm{a}}, \mathbb{S}^2) \approx 0.65448$. Figure 2.2 is the graph of q_{a} depending on $\langle \boldsymbol{x}, \mathbf{e}_z\rangle$.

Density of $\mu_{\mathbb{S}^2}^{q_{\mathrm{a}}}$ for this external field is

$$d\mu_{\mathbb{S}^2}^{q_{\mathrm{a}}}(\boldsymbol{x}) = \left(\frac{L_1(q_{\mathrm{a}}, \mathbb{S}^2) - \cos(3 \arccos \boldsymbol{x} \cdot \mathbf{e}_z)^{16}}{2\pi} \right)_+ d\mathcal{H}_d.$$

Using the numeric method described above, we obtain an approximate minimizer ω_{a} pictured in Figure 2.3.

Evaluating separation distance for ω_{a} gives $\delta(\omega_{\mathrm{a}}) \approx 0.0813$. Covering radius for the middle strip is $\eta_{\mathrm{mid}} \approx 0.0829$, and for the other two $\eta_{\mathrm{polar}} \approx 0.0727$, whence mesh ratio is $\gamma_{\mathrm{mid}} \approx 1.02$

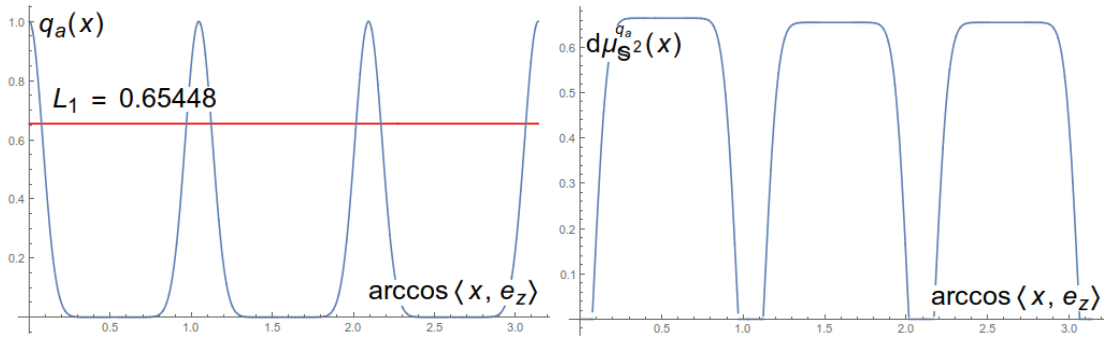

Figure 2.2: Left: graph of $q_{\mathrm{a}}(\boldsymbol{x})$, right: $d\mu_{\mathbb{S}^2}^{q_{\mathrm{a}}}(\boldsymbol{x})$ from Example 2.12. The horizontal axis is $\arccos(\boldsymbol{x} \cdot \mathbf{e}_z)$.

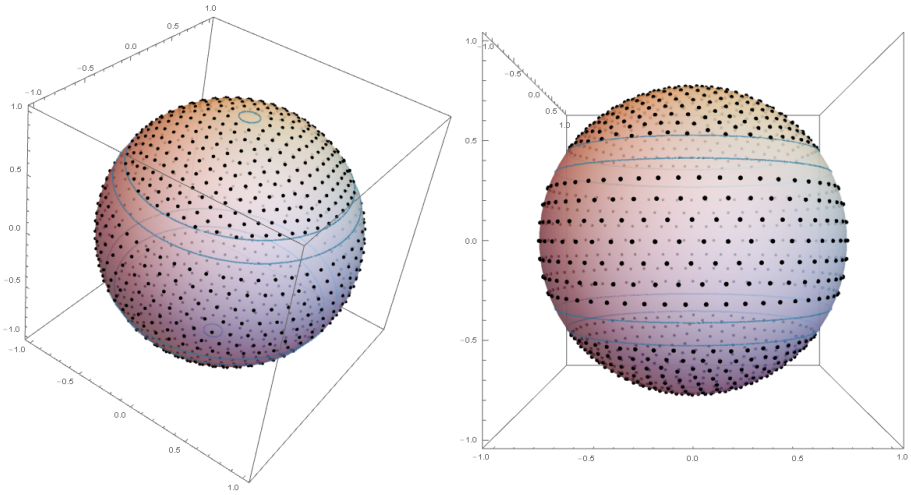

Figure 2.3: Two views of an approximate 1000-point $(2, 2, q_{\mathrm{a}})$-energy minimizer ω_{a} from Example 2.12. The latitudinal circles denote the boundaries of $\operatorname{supp} \mu_{\mathbb{S}^2}^{q_{\mathrm{a}}}$; i.e., $\{x : q_{\mathrm{a}}(\boldsymbol{x}) = L_1(q_{\mathrm{a}}, \mathbb{S}^2)\}$.

and $\gamma_{\mathrm{polar}} \approx 0.8942$ respectively.

Example 2.13. Again, let $\Omega = \mathbb{S}^2 \subset \mathbb{R}^3$, $s = d = 2$. Let us construct a sequence of discrete collections $\{\hat{\omega}_N\}_{N \geq 2}$ weak* converging to the probability distribution with density proportional to

$$\rho_{\mathrm{b}}(\boldsymbol{x}) = \begin{cases} 10\cos(4\theta) + 11, & 0 \leq \theta < \pi/4, \\ 1, & \pi/4 \leq \theta < 3\pi/4, \\ 10\cos(4\theta) + 11, & 3\pi/4 \leq \theta, \end{cases} \tag{2.18}$$

where $\theta = \arccos(\boldsymbol{x} \cdot \mathbf{e}_z)$. The external field with such a sequence of minimizers is provided by Theorem 2.3. Writing ρ for the normalization of (2.18), $\rho(\boldsymbol{x}) := \rho_{\mathrm{b}}(\boldsymbol{x}) / \int_{\mathbb{S}^2} |\rho_{\mathrm{b}}| \, d\mathcal{H}_2 \approx \rho_{\mathrm{b}}(\boldsymbol{x})/5.581722$, equation (2.9) gives the following external field:

$$q_{\mathrm{b}}(\boldsymbol{x}) := -2\pi\rho,$$

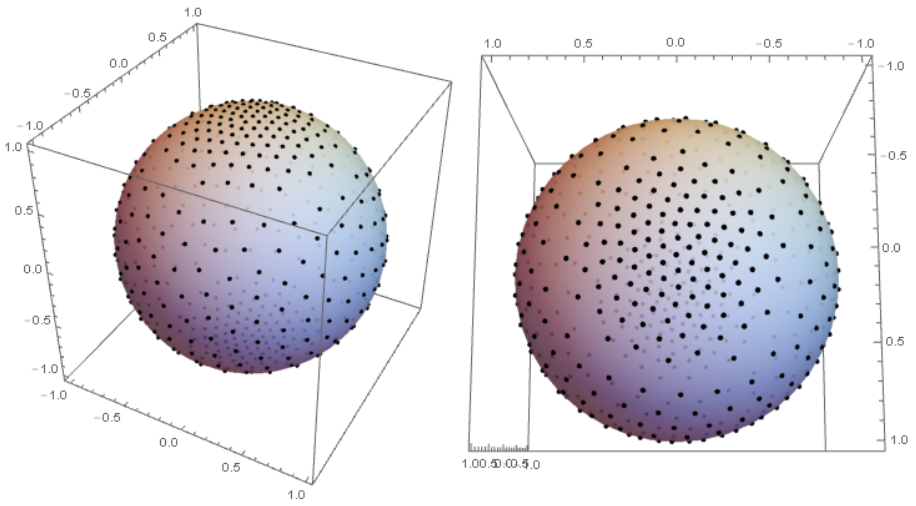

Figure 2.4: Two views of an approximation of 500-point $(2, 2, q_{\mathrm{b}})$-energy minimizer ω_{b} from Example 2.13.

where we used again that $C_{2,2} = \pi$.

An approximate discrete minimizer of this $(2, 2, q_{\mathrm{b}})$-energy is shown in the Figure 2.4. Note how higher density of $\mu_{\mathbb{S}^2}^{q_{\mathrm{b}}}$ (equivalently, larger values of ρ) causes charges to concentrate near the poles. Evaluating separation distance for the pictured configuration ω_{b} gives $\delta(\omega_{\mathrm{b}}) \approx 0.0777$, covering radius $\eta_{\mathrm{b}} \approx 0.1681$. The mesh ratio of ω_{b} is therefore: $\gamma(\omega_{\mathrm{b}}, \mathbb{S}^2) \approx 2.163$.

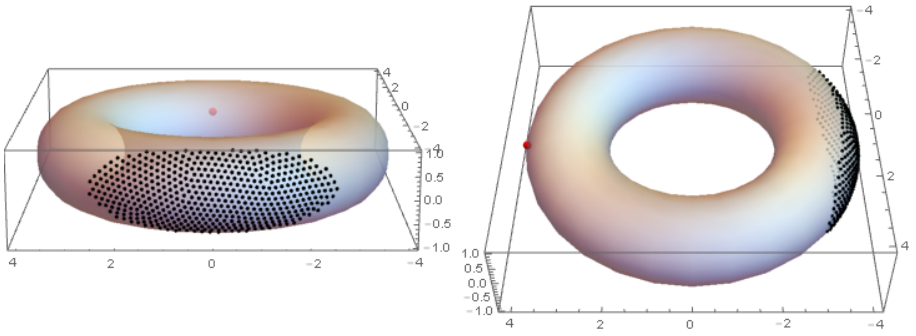

Figure 2.5: An approximation of 500-point $(8, 2, q_{\mathrm{c}})$-energy minimizer ω_c from Example 2.14. The red dot marks position $(4,\ 0,\ 0)^T$, where the repelling external field q_{c} is centered.

Example 2.14. In this example the underlying set Ω is a 2-dimensional torus with inner radius $r_i = 2$, outer radius $r_o = 4$, centered at the origin. In particular, the point $(4, 0, 0)$ lies on the outer side of its surface. Consider the problem of minimizing $(8, 2, q_{\mathrm{c}})$-energy with the external field

$$q_{\mathrm{c}}(\boldsymbol{x}) := \|\boldsymbol{x} - (4,\ 0,\ 0)^T\|^{-2}.$$

A resulting approximation of 500-point minimizer ω_{c} is shown in Figure 2.5. Separation distance

for this collection is $\delta(\omega_c) \approx 0.125339$.

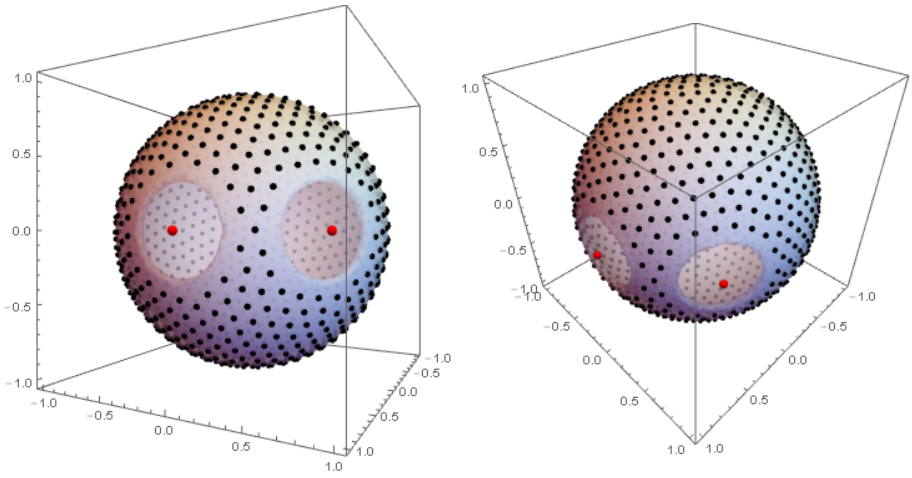

Figure 2.6: Two views of an approximation of 1000-point energy minimizer from Example 2.15. The support of $\mu_{\mathbb{S}^2}^{q_d}$ is highlighted as are the positions of the fixed "repelling charges" that create the external field q_d.

Example 2.15. Let us now consider an example of repelling field on the sphere $\mathbb{S}^2 \subset \mathbb{R}^3$. Namely, we will minimize the $(4, 2, q_d)$-energy, where

$$q_d := 10^{-3}\left(\left\|\boldsymbol{x} - (1,0,0)^T\right\|^{-4} + \left\|\boldsymbol{x} - (0.5691,\ 0.8223,\ 0)^T\right\|^{-4}\right).$$

The second repelling charge is a randomly selected point in the first quadrant of Oxy plane; factor 10^{-3} is used merely for convenience purposes.

An approximate 1000-point minimizer ω_d is shown in Figure 2.6. The shaded region marks the support of $\mu_{\mathbb{S}^2}^{q_d}$, obtained using formulas (2.5) with $C_{4,2} \approx 5.7834$ computed by the formula for its conjectured value (1.14). In other words, the shaded set is $\{x : q_d(\boldsymbol{x}) \le L_1(q_d, \mathbb{S}^2)\} \approx \{x : q_d(\boldsymbol{x}) \le 0.127\}$ (thus the complement of the support in the sphere consists of two circular-like regions). The separation distance of the pictured configuration is $\delta(\omega_d) \approx 0.1015$.

Example 2.16. Finally, consider a 1-dimensional example. We will minimize the $(4, 1, q_e)$-energy on the interval $[0, 2]$, where

$$q_e(\boldsymbol{x}) := (\boldsymbol{x} - 1.6)^4 + 40(\boldsymbol{x} - 0.2)^4(\boldsymbol{x} - 1.6)^2.$$

Substituting values of $s, d, C_{s,d}$ (the latter using formula (1.13)) into (2.5) gives that the weak* limit of minimizers is the measure with density

$$d\mu_{[0,2]}^{q_e} \approx \left(\frac{5.9574 - q_e(\cdot)}{10.8232}\right)_+^{1/4} d\mathcal{H}_1.$$

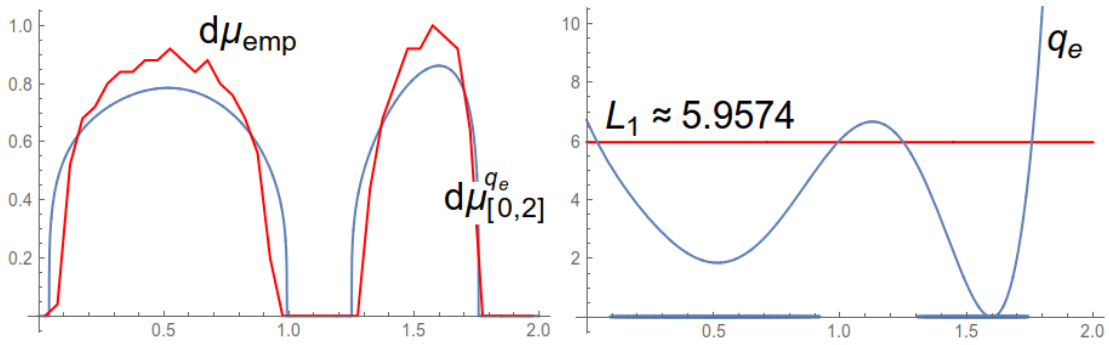

Figure 2.7: Left, empirical density $d\mu_{\rm emp}$ of $\omega_{\rm e}$, an approximate 500-point $(4,1,q_{\rm e}(\boldsymbol{x}))$-energy minimizer from Example 2.16 (red) overlaid with the graph of $d\mu^{q_{\rm e}}_{[0,2]}$ (blue); right, $\omega_{\rm e}$ and the graph of external field $q_{\rm e}$.

Figure 2.7 shows graphs of empirical density computed for $\omega_{\rm e}$ and $d\mu^{q_{\rm e}}_{[0,2]}$, as well as the graph of $q_{\rm e}$. Separation distance of the pictured configuration is $\delta(\omega_{\rm e}) \approx 0.0051$.

2.4 Proofs

For $s \geq d$, we denote by $\mathcal{R}^d_s$ the collection of all compact d-rectifiable sets $\Omega \subset \mathbb{R}^p$ with $\mathcal{H}_d(\Omega) > 0$ and, in the case $s = d$, additionally require Ω to be a subset of a d-dimensional C^1-manifold in $\mathbb{R}^p$. For $\Omega \in \mathcal{R}^d_s$ and a suitable external field q, the values of L_1 and $S(q,\Omega)$ are defined by (2.5) and (2.6), respectively. For real sequences $\{a_N\}_1^\infty$, $\{\beta\}_1^\infty$ we shall use the notation $\alpha \sim \beta,\ N \to \infty$ to mean $\alpha/\beta_N \to 1,\ N \to \infty$.

Observe that in the formula (2.1) the scaling factor $\mathcal{T}(N)$ depends on N, the number of elements in ω_N. We will occasionally need to evaluate the (s,d,q)-energy of a discrete $\omega \subset \Omega$ with $\#\omega \neq N$ and the scaling factor $\mathcal{T}(N)$, that is, the value of

$$E^q_{s,d}(\omega, N) := E^0_{s,d}(\omega) + \tau(N) \sum_{\boldsymbol{x}\in\omega} q(\boldsymbol{x}), \quad s \geq d. \tag{2.19}$$

Throughout this section $\mathbb{N}$ stands for the set of positive integers. For $s \geq d$ we also define

$$\underline{g}^q_{s,d}(\Omega) := \liminf_{N\to\infty} \frac{\mathcal{E}^q_{s,d}(\Omega,N)}{\mathcal{T}(N)}, \qquad \bar{g}^q_{s,d}(\Omega) := \limsup_{N\to\infty} \frac{\mathcal{E}^q_{s,d}(\Omega,N)}{\mathcal{T}(N)},$$

and if the limits coincide, the common value is denoted by

$$g^q_{s,d}(\Omega) = \underline{g}^q_{s,d}(\Omega) = \bar{g}^q_{s,d}(\Omega).$$

Remark 2.17. Note that for $s \geq d$ both the lower and upper asymptotic limits are finite

if q is finite-valued on a set of positive measure. Indeed, then there exists an L_* such that $\mathcal{H}_d(\{\boldsymbol{x} \in \Omega : q(\boldsymbol{x}) \le L_*\}) > 0$, so the fact that $\underline{g}^q_{s,d}(\Omega)$ and $\bar{g}^q_{s,d}(\Omega)$ are finite follows from Theorem 1.3 and a simple observation: for any two functions $q_1,\ q_2$ satisfying the hypotheses of Theorem 2.1, the inequality $q_1(\boldsymbol{x}) \le q_2(\boldsymbol{x}),\ \forall \boldsymbol{x} \in \Omega$, implies $\mathcal{E}^{q_1}_{s,d}(\Omega, N) \le \mathcal{E}^{q_2}_{s,d}(\Omega, N)$. It suffices to put $q_1 := q,\ q_2(\boldsymbol{x}) \equiv L_*$ and restrict the minimization problem to $\{\boldsymbol{x} \in \Omega : q(\boldsymbol{x}) \le L_*\}$.

Remark 2.18. If q is finite valued on a set of positive measure, then the constant L_1 in (2.5) is finite:

$$L_1 \le C_{s,d} \left(\mathcal{H}_d(\{\boldsymbol{x} \in \Omega : q(\boldsymbol{x}) < L_*\})\right)^{-s/d} (1 + s/d) + L_*,$$

where L_* is as in Remark 2.17.

Remark 2.19. We will use in many computations that if a sequence $\{a_N\}_{N \in \mathfrak{N}}$ with $a_N > 0$ satisfies $\lim_{\mathfrak{N} \ni N \to \infty} a_N/N = \alpha \ge 0$, then

$$\lim_{\substack{N \to \infty \\ N \in \mathcal{N}}} \frac{\mathcal{T}(a_N)}{\mathcal{T}(N)} = \alpha^{1+s/d}. \tag{2.20}$$

2.4.1 Proofs of the main theorems

We first establish a few lemmas that will be used in the proof of Theorem 2.1.

Lemma 2.20. *Let $u, v > 0$ and q_0, q_1 be real. Then the function*

$$F(t) := tq_0 + (1-t)q_1 + ut^{1+s/d} + v(1-t)^{1+s/d}, \quad s \ge d, \tag{2.21}$$

has a unique minimum on $[0,1]$. If there is some t_* in $(0,1)$ that satisfies

$$\frac{q_1 - q_0}{1 + s/d} = ut_*^{s/d} - v(1-t_*)^{s/d}, \tag{2.22}$$

then the minimum occurs at t_*. Otherwise, the minimum occurs at $t_* \in \{0,1\}$ such that $q_{1-t_*} = \min\{q_0, q_1\}$.

Proof. As F is strictly convex on $[0,1]$, it has a unique minimum in this interval. Differentiating yields $F'(t) = q_0 - q_1 + (1 + s/d)(ut^{s/d} - v(1-t)^{s/d})$. If there is a $t_* \in (0,1)$ satisfying (2.22), then $F'(t_*) = 0$ and the minimum of F occurs at the value t_*. Otherwise, F is strictly monotone in $[0,1]$, and the minimum must occur at an endpoint. In fact, the minimum is at $t_* \in \{0,1\}$ such that $q_{1-t_*} = \min\{q_0, q_1\}$. □

Lemma 2.21. *Let μ be a finite Radon measure on the set $\Omega \subset \mathbb{R}^p$ and $q : \Omega \to \mathbb{R}$ be measurable with respect to μ. Then for every $\varepsilon > 0$ and μ-a.e. point $\boldsymbol{x} \in \mathbb{R}^p$ there exists a positive number*

$R = R(\boldsymbol{x}, \varepsilon)$ *such that*

$$\frac{\mu\left[\{\boldsymbol{z} \in \Omega \cap B(\boldsymbol{x}, r) : |q(\boldsymbol{z}) - q(\boldsymbol{x})| < \varepsilon\}\right]}{\mu[B(\boldsymbol{x}, r)]} > 1 - \varepsilon \tag{2.23}$$

for all $r < R$.

Proof. Consider the following partition of set Ω:

$$\Omega = \bigcup_{m=1}^{\infty} \{(m-1)\varepsilon \le q(\boldsymbol{x}) < m\varepsilon\} =: \bigcup_{m=1}^{\infty} \Omega_m. \tag{2.24}$$

By the Lebesgue-Besicovitch differentiation theorem (cf. [42, 1.7.1] or [81, Corollary 2.14]) (or *Lebesgue's density theorem* in this case), for μ-a.e. point $\boldsymbol{x} \in A_m,\ m = 1, 2, \ldots,$

$$\lim_{r \downarrow 0} \frac{\mu[A_m \cap B(\boldsymbol{x}, r)]}{\mu[B(\boldsymbol{x}, r)]} = 1. \tag{2.25}$$

Therefore, (2.25) holds for μ-a.e. point $\boldsymbol{x} \in \Omega$. In particular, fix such a point $\boldsymbol{x} \in A_m$. Because $(A_m \cap B(\boldsymbol{x}, r)) \subset \{\boldsymbol{z} \in B(\boldsymbol{x}, r) : |q(\boldsymbol{z}) - q(\boldsymbol{x})| < \varepsilon\}$, equation (2.25) implies (2.23) for small enough R. □

Lemma 2.22. *Let* $\Omega \subset \mathbb{R}^p$ *satisfy* $\Omega = \cup_{m=1}^{M} A_m$, *where* A_m *are sets from the class* $\mathcal{R}_s^d$. *Let also the function* q *be defined and lower semicontinuous on* Ω. *Assume that a sequence of configurations* $\omega_n \subset \Omega$, $n \in \mathfrak{N}$, *is such that*

1. $\omega_n = \bigcup_{m=1}^{M} \omega_n^m$ *and* $\omega_n^m \subset A_m$;
2. $\omega_n^k \cap \omega_n^l = \emptyset$ *if* $k \neq l$;
3. $\lim_{\mathfrak{N} \ni n \to \infty} \#\omega_n^m / n = \alpha_m,\ 1 \le m \le M$.

Then

$$\liminf_{\mathfrak{N} \ni n \to \infty} \frac{E_{s,d}^q(\omega_n, n)}{\mathcal{T}(n)} \ge \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \sum_{m=1}^{M} \alpha_m \min_{\boldsymbol{x} \in A_m} q(\boldsymbol{x}). \tag{2.26}$$

Proof. Observe that the minima on the RHS of (2.26) are attained due to the lower semicontinuity

of q. For the LHS of (2.26) there holds

$$\begin{aligned}
\liminf_{\mathfrak{N}\ni n\to\infty} \frac{E^q_{s,d}(\omega_n,n)}{\mathcal{T}(n)} &= \liminf_{\mathfrak{N}\ni n\to\infty} \frac{1}{\mathcal{T}(n)} \left(\sum_{\substack{\boldsymbol{x}\neq\boldsymbol{y}\\ \boldsymbol{x},\boldsymbol{y}\in\omega_n}} \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + \frac{\mathcal{T}(n)}{n} \sum_{\boldsymbol{x}\in\omega_n} q(\boldsymbol{x}) \right) \\
&\geq \liminf_{\mathfrak{N}\ni n\to\infty} \frac{1}{\mathcal{T}(n)} \sum_{m=1}^{M} \sum_{\substack{\boldsymbol{x}\neq\boldsymbol{y}\\ \boldsymbol{x},\boldsymbol{y}\in\omega_n^m}} \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + \liminf_{n\to\infty} \frac{1}{n} \sum_{m=1}^{M} \sum_{\boldsymbol{x}\in\omega_n^m} q(\boldsymbol{x}) \\
&\geq \liminf_{\mathfrak{N}\ni n\to\infty} \sum_{m=1}^{M} \frac{\mathcal{T}(\#\omega_n^m)}{\mathcal{T}(n)} \frac{E^0_{s,d}(\omega_n^m)}{\mathcal{T}(\#\omega_n^m)} + \liminf_{n\to\infty} \sum_{m=1}^{M} \frac{\#\omega_n^m}{n} \min_{\boldsymbol{x}\in A_m} q(\boldsymbol{x}) \\
&\geq \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \sum_{m=1}^{M} \alpha_m \min_{\boldsymbol{x}\in A_m} q(\boldsymbol{x}),
\end{aligned}$$

where for the last inequality we used (2.20) and Theorem 1.3. □

Remark 2.23. Observe that the only assertion about $\#\omega_n$ we make is (3).

Corollary 2.24. *Let the assumptions of Lemma 2.22 be satisfied and suppose q_m, $1 \leq m \leq M$, are numbers such that the closure of $Z_m := \{\boldsymbol{x} \in A_m : q(\boldsymbol{x}) < q_m\}$ has $\mathcal{H}_d$-measure zero. Then*

$$\liminf_{\mathfrak{N}\ni n\to\infty} \frac{E^q_{s,d}(\omega_n,n)}{\mathcal{T}(n)} \geq \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \sum_{m=1}^{M} \alpha_m q_m. \tag{2.27}$$

Proof. Let $\mathfrak{N}' \subset \mathfrak{N}$ be such that

$$\liminf_{\mathfrak{N}\ni n\to\infty} \frac{E^q_{s,d}(\omega_n,n)}{\mathcal{T}(n)} = \lim_{\mathfrak{N}'\ni n\to\infty} \frac{E^q_{s,d}(\omega_n,n)}{\mathcal{T}(n)}.$$

Then

$$\begin{aligned}
\lim_{\mathfrak{N}'\ni n\to\infty} \frac{E^q_{s,d}(\omega_n,n)}{\mathcal{T}(n)} &\geq \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \lim_{\mathfrak{N}'\ni n\to\infty} \frac{1}{n} \sum_{m=1}^{M} \sum_{\boldsymbol{x}\in\omega_n^m} q(\boldsymbol{x}) \\
&\geq \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \lim_{\mathfrak{N}'\ni n\to\infty} \frac{1}{n} \left(\sum_{m=1}^{M} \frac{\#(\omega_n^m \cap (\Omega \setminus Z_m))}{n} \right) \\
&= \sum_{m=1}^{M} \alpha_m^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_m)^{s/d}} + \sum_{m=1}^{M} \alpha_m q_m.
\end{aligned}$$

□

Lemma 2.25. *Let the set Ω be such that the assumptions of Theorem 2.1 hold. Assume that a*

sequence of N-point configurations $\{\omega_N\}_{N\geq 2}$ in Ω satisfies

$$\limsup_{\substack{N\to\infty\\ N\in\mathfrak{N}}} \frac{E^q_{s,d}(\omega_N)}{\mathcal{T}(N)} < +\infty \tag{2.28}$$

and

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\omega_N} \delta_{\boldsymbol{x}} \xrightarrow{\ *\ } d\mu, \qquad \mathfrak{N}\ni N\to\infty, \tag{2.29}$$

for some Borel probability measure μ on Ω. Then μ is $\mathcal{H}_d$-absolutely continuous.

Proof. Indeed, otherwise let $E\subset\Omega$ be a Borel set such that $\mathcal{H}_d(E)=0$ and $\mu(E)>0$. Since μ is inner regular as a Borel measure on a Radon space, [59, 434K(b)], without loss of generality E is closed. For an $\varepsilon>0$ pick $r>0$ such that $E_r:=\{\boldsymbol{x}\in A:\operatorname{dist}(\boldsymbol{x},E)\leq r\}$ satisfies $\mathcal{H}_d(E_r)<\varepsilon$; observe that E_r is closed. By the definition of weak* convergence, $\liminf_{\mathfrak{N}\ni N\to\infty}\frac{1}{N}\#\{\boldsymbol{x}\in\omega_N : \boldsymbol{x}\in E_r\}\geq\mu(E)$. Then according to Theorem 1.3 and the limit (2.20),

$$\begin{aligned}\liminf_{\substack{N\to\infty\\ N\in\mathfrak{N}}} \frac{E^0_{s,d}(\omega_N\cap E_r,N)}{\mathcal{T}(N)} &= \liminf_{\substack{N\to\infty\\ N\in\mathfrak{N}}} \frac{E^0_{s,d}(\omega_N\cap E_r,N)}{\mathcal{T}(\#(\omega_N\cap E_r))}\frac{\mathcal{T}(\#(\omega_N\cap E_r))}{\mathcal{T}(N)}\\ &\geq \frac{C_{s,d}}{\mathcal{H}_d(E_r)^{s/d}}\mu(E)^{1+s/d}\geq\frac{C_{s,d}}{\varepsilon^{s/d}}\mu(E)^{1+s/d}.\end{aligned}$$

As ε was arbitrary, this contradicts (2.28). Thus μ must be $\mathcal{H}_d$-absolutely continuous. □

Lemma 2.26. *Let the assumptions of Theorem 2.1 be satisfied. Let also the sequence of N-point configurations $\{\omega_N\}_{N\in\mathfrak{N}}$ be such that*

$$\lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N)}{\mathcal{T}(N)} = \underline{g}^q_{s,d}(\Omega) \tag{2.30}$$

and

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\omega_N} \delta_{\boldsymbol{x}} \xrightarrow{\ *\ } d\mu, \qquad \mathfrak{N}\ni N\to\infty. \tag{2.31}$$

Assume that $\{B_m\}_{m=1}^M$, $M\geq 1$, is a collection of closed pairwise disjoint balls such that $\mathcal{H}_d(B_m)>0$, $\mathcal{H}_d(\partial B_m)=0$ and $\mathcal{H}_d\left(\{\boldsymbol{z}\in B_m : q(\boldsymbol{z})\leq q_m\}\right)\geq(1-\delta)\mathcal{H}_d(B_m)$, $m=1,\dots,M$, for some positive $\delta<1$.

Then

$$\lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N\cap(\cup_m B_m),N)}{\mathcal{T}(N)} \leq \min\left\{\sum_{m=1}^M \frac{C_{s,d}\,\alpha_m^{1+s/d}}{((1-\delta)\mathcal{H}_d(B_m))^{s/d}} + q_m\alpha_m\right\}, \tag{2.32}$$

where the minimum is taken over $\alpha_m\geq 0$ such that $\sum\alpha_m=\sum\mu(B_m)$.

In particular, there exists a sequence $\{\omega_N^0\}_{N\in\mathfrak{N}}$ *for which* (2.32) *is an equality with* ω_N^0 *in place of* ω_N.

Proof. Fix an $\varepsilon > 0$ satisfying $\varepsilon < 1-\delta$. Consider the set $\{\boldsymbol{z} \in B_m : q(\boldsymbol{z}) \le q_m\},\ m = 1, \dots, M$. By the inner regularity of measure $\mathcal{H}_d$, it has a closed subset B'_m contained in a ball concentric with B_i of smaller radius, for which $\mathcal{H}_d(B'_m) > (1-\delta-\varepsilon)\mathcal{H}_d(B_m)$. Let a sequence of N-point configurations $\{\omega_N^0\}_{N\in\mathfrak{N}}$ in Ω be such that $\omega_N^0 \cap (\Omega \setminus \cup_m B_m) = \omega_N \cap (\Omega \setminus \cup_m B_m)$, and such that for $1 \le m \le M$, the collection $\omega_N^0 \cap B_m$ is a minimizer of the $(s,d,0)$-energy in B'_m (in particular, is contained in it).

Equation (2.30) and Lemma 2.25 imply that $\mu(\partial B_m) = 0,\ 1 \le m \le M$. Hence the weak* convergence in (2.31) implies $\lim \#\{\boldsymbol{x} \in \omega_N : \boldsymbol{x} \in B_m\}/N \to \mu(B_m)$ when $\mathfrak{N} \ni N \to \infty$, [7, Theorem 2.1].

We will further assume that the following limits exist $\alpha_m := \lim_{\mathfrak{N}\ni N\to\infty} \#(\omega_N^0 \cap B_m)/N$, $1 \le m \le M$. The assumptions on $\#(\omega_N^0 \cap B_m)$ mean that $\sum_m \alpha_m = \sum_m \mu(B_m)$. Finally, we observe that by the construction of the sets B'_m, there exists a positive r such that $\operatorname{dist}(\cup_m B'_m, \Omega \setminus \cup_m B_m) \ge r$. Recall that

$$E^q_{s,d}(\omega_N^0, N) = E^q_{s,d}\left(\omega_N^0 \cap (\cup_m B_m), N\right) + E^q_{s,d}\left(\omega_N^0 \cap (\Omega \setminus \cup_m B_m), N\right) + \sum_{\substack{\boldsymbol{x},\boldsymbol{y}\in\omega_N^0,\\ \boldsymbol{x}\in\cup_m B_m,\\ \boldsymbol{y}\in\Omega\setminus\cup_m B_m}} \|\boldsymbol{x}-\boldsymbol{y}\|^{-s}. \tag{2.33}$$

Because $\omega_N^0 \cap \cup_m B_m = \omega_N^0 \cap \cup_m B'_m$ and because of the lower bound r for the distance between $\cup_m B'_m$ and $\Omega \setminus \cup_m B_m$, every term in the last sum is at most r^{-s}.

Using the previous equation and the definition of $\underline{g}^q_{s,d}(\Omega)$, we have:

$$\begin{aligned} 0 &\le \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N^0, N)}{\mathcal{T}(N)} - \underline{g}^q_{s,d}(\Omega) = \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \left(\frac{E^q_{s,d}(\omega_N^0)}{\mathcal{T}(N)} - \frac{E^q_{s,d}(\omega_N)}{\mathcal{T}(N)} \right) \\ &\le \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N^0 \cap (\cup_m B_m), N)}{\mathcal{T}(N)} - \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N \cap (\cup_m B_m), N)}{\mathcal{T}(N)} \\ &\quad + \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{N^2}{\mathcal{T}(N)} r^{-s}. \end{aligned} \tag{2.34}$$

Since $\lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} N^2 r^{-s}/\mathcal{T}(N) = 0$, there holds

$$\lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N \cap (\cup_m B_m), N)}{\mathcal{T}(N)} \le \lim_{\substack{N\to\infty\\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\omega_N^0 \cap (\cup_m B_m), N)}{\mathcal{T}(N)}. \tag{2.35}$$

From equation (2.30) and Lemma 2.25 follows that $\mu(\partial B_m) = 0,\ 1 \le m \le M$. Hence the weak*

convergence in (2.31) implies $\lim \#\{\boldsymbol{x} \in \omega_N : \boldsymbol{x} \in B_m\}/N \to \mu(B_m)$ when $\mathfrak{N} \ni N \to \infty$, [7, Theorem 2.1]. The construction of the sequence $\{\omega_N^0\}_{N\in\mathfrak{N}}$ and the limit (2.20) therefore imply

$$\lim_{\substack{N\to\infty \\ N\in\mathcal{N}}} \frac{E_{s,d}^q(\omega_N^0 \cap (\cup_m B_m)\,, N)}{\mathcal{T}(N)} \le \sum_{m=1}^{M} \left(\frac{C_{s,d}\,\alpha_m^{1+s/d}}{((1-\delta-\varepsilon)\mathcal{H}_d(B_m))^{s/d}} + q_m\alpha_m \right). \tag{2.36}$$

We have so far only imposed the conditions that $\alpha_1, \ldots, \alpha_M$ are nonnegative and sum to $\sum_m \mu(B_m)$. Taking $\varepsilon \to 0+$ and minimizing over such α_m in (2.36) gives (2.32). □

We first prove Theorem 2.1 for the case that q is a suitable simple function. The general case then follows by approximating an arbitrary lower semicontinuous q with such functions.

Lemma 2.27. *Let $\Omega \subset \mathbb{R}^p$ be a set from $\mathcal{R}_s^d$, and B_m, $1 \le m \le M$, be a collection of pairwise disjoint closed balls such that $\mathcal{H}_d(B_m) > 0$ and $\mathcal{H}_d(\Omega \cap \partial B_m) = 0$, $1 \le m \le M$. Assume also q is a lower semicontinuous function and for $D := \overline{\Omega \setminus \cup_m B_m}$,*

$$q(\boldsymbol{x}) = \begin{cases} q_0, & \mathcal{H}_d\text{-a.e. } \boldsymbol{x} \in D, \\ q_m, & \mathcal{H}_d\text{-a.e. } \boldsymbol{x} \in B_m, \quad 1 \le m \le M, \end{cases} \tag{2.37}$$

for positive q_m, $0 \le m \le M$.

Then equation (2.6) holds for the set Ω and function q.

Proof. For convenience, let $A_0 := D$, $A_m := B_m$, $1 \le m \le M$, in this proof. We will first verify that for some positive $\{\hat\alpha_m\}_{m=0}^M$ that add up to one,

$$\lim_{N\to\infty} \frac{\mathcal{E}_{s,d}^q(\Omega, N)}{\mathcal{T}_{s,d}(N)} = \sum_{m=0}^{M} \left(\frac{C_{s,d}\,\hat\alpha_m^{1+s/d}}{\mathcal{H}_d(A_m)^{s/d}} + q_m\hat\alpha_m \right), \tag{2.38}$$

where the values of $\hat\alpha_m$, $0 \le m \le M$ are such that

$$(\hat\alpha_0, \ldots, \hat\alpha_M) := \operatorname*{arg\,min}_{\substack{\alpha_m\ge 0, \\ \sum\alpha_m = 1}} \sum_{m=0}^{M} \left(\frac{C_{s,d}\,\alpha_m^{1+s/d}}{\mathcal{H}_d(A_m)^{s/d}} + q_m\alpha_m \right). \tag{2.39}$$

(*i*). Due to the weak* compactness of the set Ω, Corollary 2.24 implies

$$\underline{g}_{s,d}^q(\Omega) \ge \min_{\substack{\alpha_m\ge 0, \\ \sum\alpha_m = 1}} \sum_{m=0}^{M} \left(\frac{C_{s,d}\,\alpha_m^{1+s/d}}{\mathcal{H}_d(A_m)^{s/d}} + q_m\alpha_m \right). \tag{2.40}$$

Let a closed $D' \subset D$ satisfy $q(\boldsymbol{x}) \equiv q_0$, $\boldsymbol{x} \in D'$ and $\mathcal{H}_d(D') > (1-\varepsilon)\mathcal{H}_d(D)$. By the same argument as in the proof of Lemma 2.26, for a fixed $\varepsilon > 0$ we construct a sequence of N-element sets $\{\omega_N^0\}_{N\in\mathfrak{N}}$ such that the subsets $\omega_N^0 \cap D$ and $\omega_N^0 \cap B_m$, $1 \le m \le M$ are $(s,d,0)$-energy minimizing in D' and

B'_m respectively. Recall that B'_m are closed subsets of $\Omega \cap B_m$ satisfying $\mathcal{H}_d(B'_m) > (1-\varepsilon)\mathcal{H}_d(B_m)$ and dist $(\cup_m B'_m, D) > 0$. As in Lemma 2.26, we construct $\{\omega_N^0\}_{N\geq 1}$ so that the following limits exist and are finite $\alpha_0 := \lim_{\mathfrak{N}\ni N\to\infty} \#(\omega_N^0 \cap D')/N$ and $\alpha_m := \lim_{\mathfrak{N}\ni N\to\infty} \#(\omega_N^0 \cap B_m)/N$, $1 \leq m \leq M$. Since $\mathcal{E}^q_{s,d}(\Omega, N) \leq E^q_{s,d}(\omega_N^0, N)$, equation (2.33) implies

$$\bar{g}^q_{s,d}(\Omega) \leq \lim_{N\to\infty} \frac{E^q_{s,d}(\omega_N^0, N)}{\mathcal{T}(N)} = \sum_{m=0}^{M} \left(\frac{C_{s,d}\, \alpha_m^{1+s/d}}{((1-\varepsilon)\mathcal{H}_d(A_m))^{s/d}} + q_m \alpha_m \right). \tag{2.41}$$

This gives (2.38) after taking $\varepsilon \to 0+$.

(ii). Fix an A_m with strictly positive $\hat{\alpha}_m$, say, A_0 and assume $q_0 \leq q_m$ for definiteness. Pick any of the remaining sets A_k, $1 \leq k \leq M$ and denote $\beta = \beta(k) := \alpha_0 + \alpha_k$. Consider the terms on the RHS of (2.39) that contain either α_0 or α_k:

$$\bar{F}(\alpha_1, \alpha_k) := \sum_{m=0,k} \left(\frac{C_{s,d}\, \alpha_m^{1+s/d}}{\mathcal{H}_d(A_m)^{s/d}} + q_m \alpha_m \right). \tag{2.42}$$

Now choose the coefficients of the function $F(t)$ in Lemma 2.20 so that

$$F(t) = \; t^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_0)^{s/d}} + t \frac{q_0}{\beta^{s/d}} + (1-t)^{1+s/d} \frac{C_{s,d}}{\mathcal{H}_d(A_k)^{s/d}} + (1-t) \frac{q_k}{\beta^{s/d}}, \tag{2.43}$$

then $\beta^{1+s/d} F(\alpha_0/\beta) = \bar{F}(\alpha_0, \alpha_k)$. Because of (2.39), it must be that $\hat{\alpha}_0/\beta$ is the value $\hat{t} \in (0, 1]$ for which the minimum of $F(t)$ is attained. According to Lemma 2.20, either $\hat{t} = 1$, or

$$\frac{q_k - q_0}{C_{s,d}(1+s/d)} = \left(\frac{\hat{\alpha}_0}{\mathcal{H}_d(A_0)} \right)^{s/d} - \left(\frac{\hat{\alpha}_k}{\mathcal{H}_d(A_k)} \right)^{s/d}. \tag{2.44}$$

Equation (2.44) thus applies to any pair of sets in $A_0, \ldots, A_M$ provided both the corresponding $\hat{\alpha}_m$'s is positive. Also, if $\hat{\alpha}_k > 0$, then $q_k < q_l$ for every l such that $\hat{\alpha}_l = 0$. To summarize, for some L_1 there holds

$$\left(\frac{\hat{\alpha}_m}{\mathcal{H}_d(A_m)} \right)^{s/d} = \left(\frac{L_1 - q_m}{C_{s,d}(1+s/d)} \right)_+, \qquad 0 \leq m \leq M. \tag{2.45}$$

It follows from $\sum_m \hat{\alpha}_m = 1$ that the first of equations (2.5) is satisfied for this L_1.

Finally, we can evaluate the RHS of (2.38):

$$\lim_{N\to\infty} \frac{\mathcal{E}^q_{s,d}(\Omega, N)}{\mathcal{T}_{s,d}(N)} = \sum_{m=0}^{M} \left(\hat{\alpha}_m \left(\frac{L_1 - q_m}{1+s/d} \right)_+ + q_m \hat{\alpha}_m \right) = \sum_{m=0}^{M} \hat{\alpha}_m \frac{L_1 + s q_m/d}{1+s/d},$$

where in the last equality we used $\hat{\alpha}_m = 0 \iff (L_1 - q_m)_+ = 0$. This implies (2.6) because from

(2.45), $\hat{\alpha}_0 = \mu^q(D)$ and $\hat{\alpha}_m = \mu^q(A_m),\ 1 \le m \le M$, for the μ^q defined in (2.5). $\square$

Proof of Theorem 2.1. Note that as q is lower semicontinuous on the compact set Ω, it is bounded below there, so we may assume without loss of generality q is positive.

(*i*). Let $\Omega \in \mathcal{R}_s^d$ and let $0 \le q < C$ for a positive constant C. We will further use that the restriction $\mathcal{H}_d$ is a Radon measure on $\mathbb{R}^p$. Namely, [81, Theorem 7.5] implies it is locally finite because Ω is the Lipschitz image of a compact set in $\mathbb{R}^d$. It is also Borel regular as a restriction of Hausdorff measure, see [81, Theorem 4.2]. Then by [81, Corollary 1.11], $\mathcal{H}_d$ is a Radon measure.

Fix from now on a number $0 < \varepsilon < 1/4$. Apply Lemma 2.21 to the measure $\mathcal{H}_d$ and function q, denote the set of $\boldsymbol{x} \in \Omega$ for which there exists an $R(\boldsymbol{x}, \varepsilon)$ as described in the Lemma by Ω', and consider covering of Ω' by the collection of closed balls $\{B(\boldsymbol{x}, r) : \boldsymbol{x} \in \Omega',\ 0 < r < R(\boldsymbol{x}, \varepsilon)\}$. Choose for each $\boldsymbol{x} \in \Omega'$ a sequence of radii $r_{\boldsymbol{x},k} \to 0,\ k \to \infty,\ k \in \mathbb{N}$, for which

$$\frac{\mathcal{H}_d\left[\{\boldsymbol{z} \in B(\boldsymbol{x}, r_{\boldsymbol{x},k}) : |q(\boldsymbol{z}) - q(\boldsymbol{x})| < \varepsilon\}\right]}{\mathcal{H}_d[B(\boldsymbol{x}, r_{\boldsymbol{x},k})]} > 1 - \varepsilon \tag{2.46}$$

and also $\boldsymbol{y} \in \Omega \cap B(\boldsymbol{x}, r_{\boldsymbol{x},k}) \implies q(\boldsymbol{y}) > q(\boldsymbol{x}) - \varepsilon$. The latter is possible due to the lower semicontinuity of q.

Let $\{B(\boldsymbol{x}, r_{\boldsymbol{x},k})\}$ be a Vitali cover of Ω', so one can apply the version of Vitali's covering theorem for Radon measures [81, Theorem 2.8] to produce a (at most) countable subcollection of pairwise disjoint $\{B_j := B(\boldsymbol{x}_j, r_j) : j \ge 1\}$ for which $\mathcal{H}_d\left(\Omega' \setminus \cup_{j \ge 1} B_j\right) = 0$. Using $\mathcal{H}_d(\Omega) < \infty$, $\{B_j\}_{j \ge 1}$ can be chosen so that $\mathcal{H}_d(\Omega \cap \partial B_j) = 0,\ j = 1, 2, \ldots$ (there are uncountable many options for the value of r_j, at most countably many of them positive). Since $\mathcal{H}_d(\Omega \setminus \Omega') = 0$, we can fix a $J \in \mathbb{N}$ such that $\mathcal{H}_d\left(\Omega \setminus \cup_{j=1}^J B_j\right) < \varepsilon$. Let $D := \overline{\Omega \setminus \cup_{j=1}^J B_j}$. As $\mathcal{H}_d(\partial B_j) = 0,\ 1 \le j \le J$, there holds $\mathcal{H}_d(D) < \varepsilon$.

Define the two simple functions $\overline{q}_\varepsilon,\ \underline{q}_\varepsilon$ to be constant on each $B_j,\ 1 \le j \le J$:

$$\begin{aligned} \overline{q}_\varepsilon(\boldsymbol{x}) &:= \begin{cases} q(\boldsymbol{x}_j) + \varepsilon, & \boldsymbol{x} \in B_j, \\ C, & \boldsymbol{x} \in D \setminus \bigcup_j B_j. \end{cases} \\ \underline{q}_\varepsilon(\boldsymbol{x}) &:= \begin{cases} \max(0, q(\boldsymbol{x}_j) - \varepsilon), & \boldsymbol{x} \in B_j \setminus D, \\ 0, & \boldsymbol{x} \in D. \end{cases} \end{aligned} \tag{2.47}$$

Such $\underline{q}_\varepsilon,\ \overline{q}_\varepsilon$ are lower semicontinuous on Ω. Lemma 2.27 gives equation (2.6) applied to $\underline{q}_\varepsilon$ and $\overline{q}_\varepsilon$ on Ω. Let $B_j' := \{\boldsymbol{z} \in \Omega \cap B_j : |q(\boldsymbol{z}) - q(\boldsymbol{x}_j)| < \varepsilon\}$. Then,

$$q(\boldsymbol{x}_j) - \varepsilon \le \underline{q}_\varepsilon(\boldsymbol{x}) \le q(\boldsymbol{x}) \le \overline{q}_\varepsilon(X) = q(\boldsymbol{x}_j)_\varepsilon, \quad \boldsymbol{x} \in B_j'. \tag{2.48}$$

In view of (2.46) for B_j' and $\mathcal{H}_d(D) < \varepsilon$, (2.48) implies that both $\underline{q}_\varepsilon$ and $\overline{q}_\varepsilon$ converge $\mathcal{H}_d$-a.e. to q as $\varepsilon \to 0+$. Since both are bounded by $C + 1$, the dominated convergence theorem is applicable,

and

$$\lim_{\varepsilon\to 0} S(\underline{q}_\varepsilon,\Omega)=\lim_{\varepsilon\to 0} S(\overline{q}_\varepsilon,\Omega)=S(q,\Omega). \tag{2.49}$$

We now estimate

$$\lim_{N\to\infty} \mathcal{E}^q_{s,d}(\Omega,N)/\mathcal{T}_{s,d}(N)$$

in terms of

$$\lim_{N\to\infty} \mathcal{E}^{\underline{q}_\varepsilon}_{s,d}(\Omega,N)/\mathcal{T}_{s,d}(N)$$

and

$$\lim_{N\to\infty} \mathcal{E}^{\overline{q}_\varepsilon}_{s,d}(\Omega,N)/\mathcal{T}_{s,d}(N).$$

Firstly, by construction $q(\boldsymbol{x})\ge \underline{q}_\varepsilon(\boldsymbol{x}),\ \boldsymbol{x}\in\Omega$, which gives

$$\mathcal{E}^q_{s,d}(\Omega,N)\ge \mathcal{E}^{\underline{q}_\varepsilon}_{s,d}(\Omega,N). \tag{2.50}$$

On the other hand,

$$\mathcal{E}^q_{s,d}(\Omega,N)\le \mathcal{E}^q_{s,d}(D\cup\bigcup_j B'_j,N)\le \mathcal{E}^{\bar q}_{s,d}(D\cup\bigcup_j B'_j,N)\le \frac{\mathcal{T}(N)}{(1-\varepsilon)^{s/d}} S(\overline{q}_\varepsilon,\Omega), \tag{2.51}$$

where the last inequality follows from (2.46) and (2.38). This proves (2.6).

(ii). It remains to prove equation (2.8) for a sequence $\{\omega_N\}_{N\ge 2}$ satisfying (2.7). Since the probability measures on Ω are weak* compact, one can pick a subsequence $\{\hat\omega_N\}_{N\in\mathfrak{N}}\subset\{\omega_N\}_{N\ge 2}$ for which the corresponding normalized counting measures have a weak* limit:

$$\frac{1}{N}\sum_{\boldsymbol{x}\in\hat\omega_N}\delta_x \overset{*}{\rightharpoonup} \mu, \qquad \mathfrak{N}\ni N\to\infty.$$

Then μ is $\mathcal{H}_d$-absolutely continuous by the Lemma 2.25. Set $\rho(\boldsymbol{x}):=\frac{d\mu}{d\mathcal{H}_d}(\boldsymbol{x})$.

Since the integral $\int_\Omega \rho\, d\mathcal{H}_d=1$ is finite, at $\mathcal{H}_d$-a.e. point $\boldsymbol{x}$ of Ω there holds

$$\lim_{r\to 0}\frac{1}{\mathcal{H}_d(B(\boldsymbol{x},r))}\int_{B(\boldsymbol{x},r)}\rho\, d\mathcal{H}_d=\rho(\boldsymbol{x}). \tag{2.52}$$

Fix two distinct points $\boldsymbol{x}_1,\ \boldsymbol{x}_2$ for which both (2.23) for measure $\mathcal{H}_d$ and (2.52) hold. Then for an arbitrary fixed $0<\varepsilon<\min(1/2,\rho(\boldsymbol{x}_1),\rho(\boldsymbol{x}_2),q(\boldsymbol{x}_1),q(\boldsymbol{x}_2))$ there exist closed disjoint balls $B_1:=B(\boldsymbol{x}_1,r_1),\ B_2:=B(\boldsymbol{x}_2,r_2)$ centered around $\boldsymbol{x}_1,\ \boldsymbol{x}_2$ such that equations

$$\frac{\mathcal{H}_d\left[\{\boldsymbol{z}\in\Omega\cap B_m:|q(\boldsymbol{z})-q(\boldsymbol{x}_m)|<\varepsilon\}\right]}{\mathcal{H}_d[B_m]}>1-\varepsilon, \qquad m=1,2, \tag{2.53}$$

$$\int_{B_m}|\rho(\boldsymbol{z})-\rho(\boldsymbol{x}_m)|\,d\mathcal{H}_d(\boldsymbol{z})<\varepsilon\mathcal{H}_d(B_m), \qquad m=1,2, \tag{2.54}$$

hold for all closed balls concentric with B_m of radius at most r_m. Without loss of generality we will also require $\mathcal{H}_d(\partial B_1) = \mathcal{H}_d(\partial B_2) = 0$ and that $q(\boldsymbol{x}) \geq q(\boldsymbol{x}_m) - \varepsilon$ for all $\boldsymbol{x} \in B_m,\ m = 1, 2$ (which can be assumed by lower semicontinuity). Let $q_m := q(\boldsymbol{x}_m),\ m = 1, 2$; let also $q_1 \leq q_2$.

Due to μ being absolutely continuous with respect to $\mathcal{H}_d$, the assumption $\mathcal{H}_d(\partial B_m) = 0,\ m = 1, 2$, and the limit (2.20), Lemma 2.22 implies

$$\lim_{\substack{N\to\infty \\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\hat{\omega}_N \cap (B_1 \cup B_2), N)}{\mathcal{T}(N)} \geq \sum_{m=1,2} \left(\frac{C_{s,d}\, \mu(B_m)^{1+s/d}}{\mathcal{H}_d(B_m)^{s/d}} + \mu(B_m)(q_m - \varepsilon) \right). \tag{2.55}$$

On the other hand, from Lemma 2.26:

$$\lim_{\substack{N\to\infty \\ N\in\mathcal{N}}} \frac{E^q_{s,d}(\hat{\omega}_N \cap (B_1 \cup B_2), N)}{\mathcal{T}(N)} \leq \min \left\{ \sum_{m=1,2} \frac{C_{s,d}\, \alpha_m^{1+s/d}}{((1-\varepsilon)\mathcal{H}_d(B_m))^{s/d}} + (q_m + \varepsilon)\alpha_m \right\} \tag{2.56}$$

with minimum taken over positive $\alpha_1,\ \alpha_2$ satisfying $\alpha_1 + \alpha_2 = \mu(B_1) + \mu(B_2)$. If we denote

$$\begin{aligned} (\hat{\alpha}_1, \hat{\alpha}_2) := \arg\min \Bigg\{ \sum_{m=1,2} \left(\frac{C_{s,d}\, \alpha_m^{1+s/d}}{((1-\varepsilon)\mathcal{H}_d(B_m))^{s/d}} + (q_m + \varepsilon)\alpha_m \right) \\ : \alpha_1 + \alpha_2 = \mu(B_1) + \mu(B_2) \Bigg\}, \end{aligned} \tag{2.57}$$

and argue as in the proof of Lemma 2.27, we obtain, similarly to (2.44), that $(\hat{\alpha}_1, \hat{\alpha}_2)$ satisfy

$$\frac{q_2 - q_1}{C_{s,d}(1 + s/d)} = \left(\frac{\hat{\alpha}_1}{(1-\varepsilon)\mathcal{H}_d(B_1)} \right)^{s/d} - \left(\frac{\hat{\alpha}_2}{(1-\varepsilon)\mathcal{H}_d(B_2)} \right)^{s/d}. \tag{2.58}$$

Inequalities (2.55)–(2.56) and the definition of $(\hat{\alpha}_1, \hat{\alpha}_2)$ give:

$$\begin{aligned} &\sum_{m=1,2} \left(\frac{C_{s,d}\, \mu(B_m)^{1+s/d}}{\mathcal{H}_d(B_m)^{s/d}} + \mu(B_m)(q_m - \varepsilon) \right) \\ &\leq \sum_{m=1,2} \left(\frac{C_{s,d}\, \hat{\alpha}_m^{1+s/d}}{((1-\varepsilon)\mathcal{H}_d(B_m))^{s/d}} + \hat{\alpha}_m(q_m + \varepsilon) \right) \\ &\leq \sum_{m=1,2} \left(\frac{C_{s,d}\, \mu(B_m)^{1+s/d}}{((1-\varepsilon)\mathcal{H}_d(B_m))^{s/d}} + \mu(B_m)(q_m + \varepsilon) \right). \end{aligned} \tag{2.59}$$

Observe that if in the above construction we fix the ball B_1 and allow $r_2 \to 0$, the first term on the RHS of (2.58) is bounded, so the ratio $\hat{\alpha}_2/\mathcal{H}_d(B_2)$ is bounded as well, say, $\hat{\alpha}_2/\mathcal{H}_d(B_2) \leq R_2$[1].

[1] due to the assumptions $q_1 \leq q_2$ and $\mathcal{H}_d(B_2)/\mathcal{H}_d(B_1) < \varepsilon$, the equality $\hat{\alpha}_1 + \hat{\alpha}_2 = \mu(B_1) + \mu(B_2)$, and equations

Let also r_2 be such that $\mathcal{H}_d(B_2)/\mathcal{H}_d(B_1) < \varepsilon$ and $\hat{\alpha}_2/\mathcal{H}_d(B_1) < \varepsilon$. Due to equation (2.54), there holds $|\mu(B_m)/\mathcal{H}_d(B_m) - \rho(\boldsymbol{x}_m)| < \varepsilon,\ m = 1, 2$. Dividing (2.59) through by $\mathcal{H}_d(B_1)$ for such a choice of r_2 gives:

$$\begin{aligned}
&C_{s,d}(\rho(\boldsymbol{x}_1) - \varepsilon)^{1+s/d} + (\rho(\boldsymbol{x}_1) - \varepsilon)(q_1 - \varepsilon) \\
\leq &\frac{C_{s,d}}{(1-\varepsilon)^{s/d}} \left(\frac{\hat{\alpha}_1}{\mathcal{H}_d(B_1)} \right)^{1+s/d} + \left(\frac{\hat{\alpha}_1}{\mathcal{H}_d(B_1)} \right) (q_1 + \varepsilon) \\
&\quad + \left(\frac{\hat{\alpha}_2}{\mathcal{H}_d(B_1)} \right) \frac{C_{s,d}}{(1-\varepsilon)^{s/d}} \left(\frac{\hat{\alpha}_2}{\mathcal{H}_d(B_2)} \right)^{s/d} + \left(\frac{\hat{\alpha}_2}{\mathcal{H}_d(B_1)} \right) (q_2 + \varepsilon) \qquad (2.60) \\
\leq &\frac{C_{s,d}}{(1-\varepsilon)^{s/d}} (\rho(\boldsymbol{x}_1) + \varepsilon)^{1+s/d} + (\rho(\boldsymbol{x}_1) + \varepsilon)(q_1 + \varepsilon) \\
&\quad + \varepsilon \left(\frac{C_{s,d}}{(1-\varepsilon)^{s/d}} (\rho(\boldsymbol{x}_2) + \varepsilon)^{1+s/d} + (\rho(\boldsymbol{x}_2) + \varepsilon)(q_2 + \varepsilon) \right),
\end{aligned}$$

Finally, because $\varepsilon > 0$ was arbitrary and the function $C_{s,d} t^{1+s/d} + q(\boldsymbol{x}_1)t,\ t \geq 0$ is monotone, inequalities (2.60) yield by the above discussion

$$\rho(\boldsymbol{x}_1) = \lim_{r_1 \to 0} \frac{\hat{\alpha}_1}{\mathcal{H}_d(B_1)}. \qquad (2.61)$$

We could similarly fix the ball B_2 first and ensure $\mathcal{H}_d(B_1)/\mathcal{H}_d(B_2) < \varepsilon$, taking $r_2 \to 0$ afterwards, thus also

$$\rho(\boldsymbol{x}_2) = \lim_{r_2 \to 0} \frac{\hat{\alpha}_2}{\mathcal{H}_d(B_2)}. \qquad (2.62)$$

In conjunction with (2.58) the last two equations give

$$\frac{q_2 - q_1}{C_{s,d}(1 + s/d)} = \rho(\boldsymbol{x}_1)^{s/d} - \rho(\boldsymbol{x}_2)^{s/d} \qquad (2.63)$$

for $\mathcal{H}_d \times \mathcal{H}_d$-a.e. pair $(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \Omega \times \Omega$. Due to the normalization property $\int_\Omega \rho(\boldsymbol{x})\, d\mathcal{H}_d = 1$ and the definition of L_1 in (2.5),

$$\rho(\boldsymbol{x}) = \left(\frac{L_1 - q(\boldsymbol{x})}{C_{s,d}(1 + s/d)} \right)_+^{d/s} \qquad \mathcal{H}_d\text{-a.e.} \qquad (2.64)$$

which coincides with the density in formula (2.8).

(iii). Finally, we turn to the case when the function q need not be bounded above. Consider

$$q_C(\boldsymbol{x}) := \begin{cases} q(\boldsymbol{x}), & q(\boldsymbol{x}) \leq C, \\ C, & \text{otherwise.} \end{cases}$$

(2.54) and (2.58), one can take $R_2 = \rho(\boldsymbol{x}_1) + \rho(\boldsymbol{x}_2) + 2$ as a rough estimate.

Recall that $\Omega(C) = \{\boldsymbol{x} \in \Omega : q(\boldsymbol{x}) \le C\}$ is a d-rectifiable set as a closed subset of Ω. The Theorem 2.1 is therefore applicable to each function q_C if seen as defined on $\Omega(C)$. By Remark 2.18, the value of L_1 is finite. For all $C \ge L_1$,

$$\operatorname{supp}(\mu^{q_C}) \cap \{x : q(\boldsymbol{x}) > C\} = \emptyset. \tag{2.65}$$

Inequality $q_C(\boldsymbol{x}) \le q(\boldsymbol{x})$ for all $\boldsymbol{x} \in \Omega$ implies

$$\mathcal{E}^{q_C}_{s,d}(\Omega, N) \le \mathcal{E}^{q}_{s,d}(\Omega, N), \quad N \ge 2,$$

so $S(q_C, \Omega) \le \underline{g}^q_{s,d}(\Omega)$. On the other hand, due to set inclusion,

$$\underline{g}^q_{s,d}(\Omega) \le \bar{g}^q_{s,d}(\Omega) \le \bar{g}^q_{s,d}(\Omega(C)) = S(q_C, \Omega(C)) = S(q_C, \Omega),$$

where the last two equalities follow from Theorem 2.1 applied to the function q_C and sets $\Omega(C)$ and Ω respectively, and equation (2.65). To summarize, $\underline{g}^q_{s,d}(\Omega) = \bar{g}^q_{s,d}(\Omega) = S(q, \Omega)$.

Let now $\{\omega_N\}_{N\ge 2}$ be a sequence satisfying (2.7). Fix a $C > L_1$. Because $q_C(\boldsymbol{x}) \le q(\boldsymbol{x})$ for all $\boldsymbol{x} \in \Omega$ and $S(q, \Omega) = S(q_C, \Omega)$, the sequence $\{\omega_N\}_{N\ge 2}$ is also asymptotically (s, d, q_C)-energy minimizing. Then by Theorem 2.1 this sequence converges weak* to $d\mu^{q_C}$, and it remains to observe that for $C > L_1$ it holds $d\mu^{q_C} = d\mu^q$, where the two measures are defined in equation (2.5). □

Proof of Theorem 2.3. The desired result is an immediate application of Theorem 2.1 since using equation (2.5) for the external field from (2.9) gives $L_1 = 0$, so the asymptotic distribution is indeed (2.10). □

Proof of Proposition 2.5. We have

$$q'(\boldsymbol{x}) = (1 + \Delta) q(\boldsymbol{x}).$$

According to (2.5), the equation

$$\int_\Omega \left(\frac{l - q'(\boldsymbol{x})}{M_{s,d}} \right)^{d/s}_+ d\mathcal{H}_d(\boldsymbol{x}) = 1 \tag{2.66}$$

for variable l has the unique solution $l = L_1'$. Using (2.9), it can be rewritten as

$$\int_\Omega \left(\rho^{s/d} + \frac{\Delta \rho^{s/d} + l}{M_{s,d}} \right)^{d/s}_+ d\mathcal{H}_d(\boldsymbol{x}) = 1, \tag{2.67}$$

which, in view of $\int_\Omega \rho \, d\mathcal{H}_d = 1$ and monotonicity of the function $(\cdot)_+$, shows that the solution of

(2.66) satisfies $|l| \le |\Delta| \, \|\rho\|_\infty^{s/d}$, that is,

$$|L_1'| \le |\Delta| \, \|\rho\|_\infty^{s/d}.$$

We will therefore write $L_1' = K\Delta$ with $|K| \le \|\rho\|_\infty^{s/d}$.

Let us now estimate the difference between densities $\rho' = d\mu^{q'} / d\mathcal{H}_d$ and $\rho = d\mu^q / d\mathcal{H}_d$ in terms of Δ. Factor out $\rho^{s/d}$ from the parentheses in (2.67) and observe that for $\Delta < M_{s,d} / \left(1 + (\|\rho\|_\infty \delta^{-1})^{s/d}\right)$ the expression inside is nonnegative, which allows to expand it up to $o(\Delta)$:

$$\rho' = \rho \left(1 + \frac{\Delta(1 + K/\rho^{s/d})}{M_{s,d}}\right)^{d/s} = \rho + \Delta \frac{d(1 + K/\rho^{s/d})}{s M_{s,d}} + o(\Delta), \qquad \Delta \to 0.$$

□

2.4.2 Proofs of separation and covering properties

To obtain point separation results we use techniques of [71, 12].

Proof of Lemma 2.8. Fix an $\boldsymbol{x} \in \hat{\omega}_N$. Because the minimal value of energy $E_{s,d}^q(\omega_N)$ is attained for $\hat{\omega}_N$, one must have

$$U(\boldsymbol{x}, \hat{\omega}_N) \le U(\boldsymbol{z}, \hat{\omega}_N), \quad \boldsymbol{z} \in \Omega, \tag{2.68}$$

where $U(\cdot\,, \hat{\omega}_N)$ is defined in (2.13). According to Frostman's lemma, [81, Theorem 8.8], for the set Ω there exists a positive Borel measure μ satisfying $\mu(\Omega) > 0$ and such that for all $\boldsymbol{x} \in \mathbb{R}^p$ and $R > 0$,

$$\mu\big(B(\boldsymbol{x}, R) \cap \Omega\big) \le R^d. \tag{2.69}$$

By continuity of measure μ from below there exists a positive constant H for which $\mu[\Omega(H)] \ge 2\mu(\Omega)/3$; this constant then depends on Ω and q, $H = H(\Omega, q)$. Observe that when q is bounded from above, H can be chosen equal to its upper bound. Let $r_0 := (\mu(\Omega)/3N)^{1/d}$. Consider the set

$$D_x := \Omega(H) \setminus \bigcup_{\substack{\boldsymbol{y} \in \hat{\omega}_N: \\ \boldsymbol{y} \ne \boldsymbol{x}}} B(\boldsymbol{y}, r_0).$$

From (2.69):

$$\mu(D_x) \ge 2\mu(\Omega)/3 - \sum_{\substack{\boldsymbol{y} \in \hat{\omega}_N: \\ \boldsymbol{y} \ne \boldsymbol{x}}} \mu\left(\Omega \cap B(\boldsymbol{y}, r_0)\right) \ge \mu(\Omega)/3.$$

Averaging $U(\boldsymbol{z}, \hat{\omega}_N)$ on D_x and taking into account (2.68) yields

$$\begin{aligned} U(\boldsymbol{x}, \hat{\omega}_N) \leq & \mu(D_x)^{-1} \int_{D_x} U(\boldsymbol{z}, \hat{\omega}_N) d\mu(\boldsymbol{z}) \\ \leq & \frac{3}{\mu(\Omega)} \left(\sum_{\substack{\boldsymbol{y} \in \hat{\omega}_N : \\ \boldsymbol{y} \neq \boldsymbol{x}}} \int_{\Omega \setminus B(\boldsymbol{y}, r_0)} \|\boldsymbol{z} - \boldsymbol{y}\|^{-s} \, d\mu(\boldsymbol{z}) + \frac{\mathcal{T}(N)}{N} \int_\Omega q(\boldsymbol{z}) \, d\mu(\boldsymbol{z}) \right). \end{aligned} \tag{2.70}$$

Denote $R_0 := \operatorname{diam}(\Omega)$. For the integrals in the sum (2.70), use (2.69) again:

$$\begin{aligned} I(\boldsymbol{y}, r) := & \int_{\Omega \setminus B(\boldsymbol{y}, r)} \|\boldsymbol{z} - \boldsymbol{y}\|^{-s} \, d\mu(\boldsymbol{z}) = \int_0^{r^{-s}} \mu \left\{ \boldsymbol{z} \in \Omega \setminus B(\boldsymbol{y}, r) : \|\boldsymbol{z} - \boldsymbol{y}\|^{-s} > t \right\} dt \\ \leq & \mu(\Omega) R_0^{-s} + \int_{R_0^{-s}}^{r^{-s}} \mu \left[\Omega \cap B(\boldsymbol{y}, t^{-1/s}) \right] dt \\ \leq & \begin{cases} r^{d-s} \, s/(s-d), & s > d, \\ (1 + d \log(R_0/r)), & s = d. \end{cases} \end{aligned} \tag{2.71}$$

This estimate is independent of $\boldsymbol{y}$. Using the definition of r_0, in the case $s > d$:

$$\begin{aligned} U(\boldsymbol{x}, \hat{\omega}_N) \leq & \frac{3}{\mu(\Omega)} \left(N I(\boldsymbol{y}, r_0) + N^{s/d} \int_\Omega q(\boldsymbol{z}) \, d\mu(\boldsymbol{z}) \right) \\ \leq & \frac{3}{\mu(\Omega)} \left(\frac{Ns}{r_0^{s-d}(s-d)} + \mu(\Omega) H N^{s/d} \right) \\ = & \left(\frac{s}{s-d} \left(\frac{3}{\mu(\Omega)} \right)^{s/d} + 3H \right) N^{s/d} = C(\Omega, s, d, q) N^{s/d}. \end{aligned}$$

Similarly, for $s = d$,

$$\begin{aligned} U(\boldsymbol{x}, \hat{\omega}_N) \leq & \frac{3}{\mu(\Omega)} N (1 + d \log(R_0/r_0)) + 3HN \log N \\ \leq & \, C(\Omega, d, q) N \log N = C(\Omega, d, q) N \log N. \end{aligned}$$

This proves the desired statement. □

Proof of Corollary 2.9. It is immediate from Lemma 2.8 and nonnegativity of q that each $\boldsymbol{x} \in \hat{\omega}_N$ satisfies $\boldsymbol{x} \in \Omega(C(\Omega, s, d, q))$ with the constant taken from (2.14). □

Proof of Theorem 2.7. Let $\|\boldsymbol{x}-\boldsymbol{y}\| = \delta(\hat{\omega}_N),\ \boldsymbol{x},\boldsymbol{y} \in \hat{\omega}_N$. From Lemma 2.8, for $N \geq 2$,

$$\delta(\hat{\omega}_N)^{-s} = \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} \leq U(\boldsymbol{x},\hat{\omega}_N) \leq U(\boldsymbol{x},\hat{\omega}_N) \leq C(\Omega,s,d,q) \begin{cases} N^{s/d} & s > d; \\ N\log N & s = d, \end{cases}$$

which implies the theorem. □

Similarly to the function (2.13), for an $r > 0$ and $\boldsymbol{y} \in \Omega$ let

$$U_r(\boldsymbol{x},\omega_N) := \sum_{\boldsymbol{y}\in\omega_N(\boldsymbol{x},r)} \|\boldsymbol{y}-\boldsymbol{x}\|^{-s}, \tag{2.72}$$

where $\omega_N(\boldsymbol{x},r) := \{\boldsymbol{y} \in \omega_N : \boldsymbol{y} \in B(\boldsymbol{x},r)\}$ for a fixed sequence of discrete configurations $\{\omega_N\}_{N\geq 2}$.

Lemma 2.28. *Let $s > d$. Assume that $\Omega \subset \mathbb{R}^p$ satisfies $\mathcal{H}_d(\Omega) > 0$, is compact and d-regular with respect to μ. Let $q \in L^1(\Omega,\mu)$ be a nonnegative lower semicontinuous function and $\{\hat{\omega}_N\}_{N\geq 2}$ be a sequence of (s,d,q)-energy minimizers. If for a point $\boldsymbol{x} \in \Omega$ and some $r > 0$,*

$$U_r(\boldsymbol{x},\hat{\omega}_N) \geq CN^{s/d}, \quad N \geq 2, \tag{2.73}$$

then

$$\operatorname{dist}(\boldsymbol{x},\hat{\omega}_N) \leq \widetilde{C}(C,\Omega,\mu,s,d)N^{-1/d}, \quad N \geq 2.$$

Proof. The proof follows the lines of [67]. By Theorem 2.7, there exists a $C_1 > 0$ such that

$$\delta(\hat{\omega}_N) \geq C_1/N^{1/d}, \quad N \geq 2.$$

Considering a subsequence if necessary, one may assume $\operatorname{dist}(\boldsymbol{x},\hat{\omega}_N) \geq C_1/2N^{1/d}$, since otherwise the statement of the Lemma follows immediately. Consider $r_0 := \varepsilon C_1/N^{1/d},\ 0 < \varepsilon < 1/2$ and put $B_y := \Omega \cap B(\boldsymbol{y},r_0),\ \boldsymbol{y} \in \hat{\omega}_N(\boldsymbol{x},r)$ for every $N \geq 2$. The collection $\{B_y\}$ defined in this way consists of disjoint sets. By construction, then for any $\boldsymbol{z} \in B_y$,

$$\|\boldsymbol{z}-\boldsymbol{x}\| \leq \|\boldsymbol{z}-\boldsymbol{y}\| + \|\boldsymbol{y}-\boldsymbol{x}\| \leq r_0 + \|\boldsymbol{y}-\boldsymbol{x}\| \leq (2\varepsilon+1)\|\boldsymbol{y}-\boldsymbol{x}\|, \quad \boldsymbol{y} \in \hat{\omega}(\boldsymbol{x},r)$$

where we used that $r_0 \leq 2\varepsilon \operatorname{dist}(\boldsymbol{x},\hat{\omega}_N) \leq 2\varepsilon\|\boldsymbol{y}-\boldsymbol{x}\|$. As Ω is d-regular with respect to μ, we obtain from the last equation

$$\|\boldsymbol{y}-\boldsymbol{x}\|^{-s} \leq \frac{(2\varepsilon+1)^s}{\mu(B_y)} \int_{B_y} \|\boldsymbol{z}-\boldsymbol{x}\|^{-s}\, d\mu(\boldsymbol{z}) \leq \frac{(2\varepsilon+1)^s}{c_0 r_0^d} \int_{B_y} \|\boldsymbol{z}-\boldsymbol{x}\|^{-s}\, d\mu(\boldsymbol{z}). \tag{2.74}$$

Also, for $\boldsymbol{z} \in B_y$:

$$\|\boldsymbol{z}-\boldsymbol{x}\| \geq \|\boldsymbol{y}-\boldsymbol{x}\| - \|\boldsymbol{z}-\boldsymbol{y}\| \geq (1-2\varepsilon)\|\boldsymbol{y}-\boldsymbol{x}\| \geq (1-2\varepsilon)\mathrm{dist}\,(\boldsymbol{x},\hat{\omega}_N) =: r_\varepsilon,$$

which implies

$$\bigcup_{\boldsymbol{y}\in\hat{\omega}(\boldsymbol{x},r)} B_y \subset \Omega \setminus B(\boldsymbol{x},r_\varepsilon).$$

We write $\widetilde{c_0} := (2\varepsilon+1)^s/c_0$. Summing equations (2.74) over $\boldsymbol{y} \in \hat{\omega}(\boldsymbol{x},r)$ and using (2.71),

$$\begin{aligned} U_r(\boldsymbol{x},\hat{\omega}_N) = \sum_{\boldsymbol{y}\in\hat{\omega}_N(\boldsymbol{x},r)} \|\boldsymbol{y}-\boldsymbol{x}\|^{-s} \leq \frac{\widetilde{c_0}}{r_0^d} \sum_{\boldsymbol{y}\in\hat{\omega}_N(\boldsymbol{x},r)} \int_{B_y} \|\boldsymbol{z}-\boldsymbol{x}\|^{-s}\, d\mu(\boldsymbol{z}) \leq \frac{\widetilde{c_0}}{r_0^d} \int\limits_{\Omega\setminus B(\boldsymbol{x},r_\varepsilon)} \|\boldsymbol{z}-\boldsymbol{x}\|^{-s}\, d\mu(\boldsymbol{z}) \\ = \frac{\widetilde{c_0}}{r_0^d} I(\boldsymbol{x},r_\varepsilon) \leq r_\varepsilon^{d-s} C_0 \frac{s\widetilde{c_0}}{(s-d)r_0^d} = N[\mathrm{dist}\,(\boldsymbol{x},\hat{\omega}_N)]^{d-s} \frac{C_0 s\,(1+2\varepsilon)^s(1-2\varepsilon)^{d-s}}{c_0 C_1^d (s-d)\,\varepsilon^d}. \end{aligned}$$

The RHS has the minimal value at $\varepsilon = \frac{d}{4s-2d} < 1/2$, if considered as function of ε. Summarizing, we have

$$U_r(\boldsymbol{x},\hat{\omega}_N) \leq \hat{C}(\Omega,\mu,s,d) N [\mathrm{dist}\,(\boldsymbol{x},\hat{\omega}_N)]^{d-s}.$$

Substitution of (2.73) gives

$$\mathrm{dist}\,(\boldsymbol{x},\hat{\omega}_N) \leq \left(\frac{\hat{C}(\Omega,\mu,s,d)N}{U_r(\boldsymbol{x},\hat{\omega}_N)}\right)^{1/(s-d)} \leq \left(\frac{\hat{C}(\Omega,\mu,s,d)N}{CN^{s/d}}\right)^{1/(s-d)},$$

which ends the proof. □

Recall that we write $\alpha_N \sim \beta_N,\ N\to\infty$ if $\alpha_N/\beta_N \to 1,\ N\to\infty$.

Lemma 2.29. *Let the assumptions of Theorem 2.10 hold. Then*

$$U_r(\boldsymbol{x},\hat{\omega}_N) \geq C(\Omega,h,s,d,q) N^{s/d}, \qquad N \geq 2. \tag{2.75}$$

Proof. For a fixed $\Delta > 0$, choose small enough $r > 0$, so that $\boldsymbol{z} \in B(\boldsymbol{x},r)$ implies $q(\boldsymbol{z}) \leq L_1 - h/2$ and $q(\boldsymbol{z}) \geq q(\boldsymbol{x}) - \Delta$ for all $\boldsymbol{x} \in \Omega(L_1-h)$. The choice of r thus depends on q,h,Ω,Δ. Note that by (2.5), $\boldsymbol{x} \in \mathrm{supp}\,(\mu^q)$. Suppose also that r satisfies $\mathcal{H}_d(\partial B(\boldsymbol{x},r)) = 0$ (such values of r are dense because $\mathcal{H}_d(\Omega) < \infty$). As above, $\hat{\omega}_N(\boldsymbol{x},r) = \{\boldsymbol{y} \in \hat{\omega}_N : \boldsymbol{y} \in B(\boldsymbol{x},r)\}$. Using equation (2.15) and Theorem 2.1, we have:

$$\begin{aligned} \mu^q[B(\boldsymbol{x},r)] = \int\limits_{B(\boldsymbol{x},r)} \left(\frac{L_1 - q(\boldsymbol{z})}{C_{s,d}(1+s/d)}\right)^{d/s} d\mathcal{H}_d(\boldsymbol{z}) \geq \\ \geq \left(\frac{h}{2C_{s,d}(1+s/d)}\right)^{d/s} c_0 r^d =: c(h,\Omega,s,d) r^d. \end{aligned} \tag{2.76}$$

For any $N \geq 2$, if $\boldsymbol{x} \in \hat{\omega}_N$ there is nothing to prove. Otherwise, as $\hat{\omega}_N$ is an optimal configuration, from (2.68) for every $\boldsymbol{y} \in \hat{\omega}_N(\boldsymbol{x}, r)$

$$\begin{aligned} U_r(\boldsymbol{x}, \hat{\omega}_N) + N^{s/d} q(\boldsymbol{x}) &\geq \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + \sum_{\substack{\boldsymbol{z}\in\hat{\omega}_N(\boldsymbol{x},r):\\ \boldsymbol{z}\neq\boldsymbol{y}}} \|\boldsymbol{z}-\boldsymbol{y}\|^{-s} + N^{s/d} q(\boldsymbol{y}) + \\ &\quad + \sum_{\substack{\boldsymbol{z}\in\hat{\omega}_N:\\ \boldsymbol{z}\notin\hat{\omega}_N(\boldsymbol{x},r)}} \left(\|\boldsymbol{z}-\boldsymbol{y}\|^{-s} - \|\boldsymbol{z}-\boldsymbol{x}\|^{-s}\right) \\ &\geq \|\boldsymbol{x}-\boldsymbol{y}\|^{-s} + \sum_{\substack{\boldsymbol{z}\in\hat{\omega}_N(\boldsymbol{x},r):\\ \boldsymbol{z}\neq\boldsymbol{y}}} \|\boldsymbol{z}-\boldsymbol{y}\|^{-s} + N^{s/d} q(\boldsymbol{y}) - N r^{-s}. \end{aligned}$$

Summing over all $\boldsymbol{y} \in \hat{\omega}_N(\boldsymbol{x}, r)$,

$$\begin{aligned} &(\#\hat{\omega}_N(\boldsymbol{x}, r) - 1) U_r(\boldsymbol{y}, \hat{\omega}_N) \geq \\ &\qquad \geq \sum_{\substack{\boldsymbol{y},\boldsymbol{z}\in\hat{\omega}_N(\boldsymbol{x},r):\\ \boldsymbol{z}\neq\boldsymbol{y}}} \|\boldsymbol{z}-\boldsymbol{y}\|^{-s} + N^{s/d} \sum_{\boldsymbol{y}\in\hat{\omega}_N(\boldsymbol{x},r)} (q(\boldsymbol{y}) - q(\boldsymbol{x})) - N^2 r^{-s} \\ &\qquad \geq \sum_{\substack{\boldsymbol{y},\boldsymbol{z}\in\hat{\omega}_N(\boldsymbol{x},r):\\ \boldsymbol{z}\neq\boldsymbol{y}}} \|\boldsymbol{z}-\boldsymbol{y}\|^{-s} - N^{s/d} \Delta \#\hat{\omega}_N(\boldsymbol{x}, r) - N^2 r^{-s}. \end{aligned} \tag{2.77}$$

Since $\mathcal{H}_d(\partial B(\boldsymbol{x}, r)) = 0$, there holds $\lim \#\hat{\omega}_N(\boldsymbol{x}, r)/N = \mu^q[B(\boldsymbol{x}, r)]$, $N \to \infty$. Using (2.76) and the Lemma 2.22 for the single set $B(\boldsymbol{x}, r)$ with $q(\cdot) \equiv 0$, we conclude

$$\liminf_{N\to\infty} N^{-1-s/d} \sum_{\substack{\boldsymbol{y},\boldsymbol{z}\in\hat{\omega}_N(\boldsymbol{x},r):\\ \boldsymbol{z}\neq\boldsymbol{y}}} \|\boldsymbol{z}-\boldsymbol{y}\|^{-s} \geq \frac{\mathcal{H}_d[B(\boldsymbol{x}, r)]}{C_{s,d}^{d/s}} \left(\frac{h}{2(1+s/d)}\right)^{1+d/s}.$$

Since $N^2 r^{-s} = o(N^{1+s/d})$, dividing (2.77) by $\#\hat{\omega}_N(\boldsymbol{x}, r) \sim N\mu^q[B(\boldsymbol{x}, r)]$ for $N \to \infty$ gives

$$\begin{aligned} &U_r(\boldsymbol{x}, \hat{\omega}_N) \geq \\ &\geq N^{s/d} \left(\frac{\mathcal{H}_d[B(\boldsymbol{x}, r)]}{\mu^q[B(\boldsymbol{x}, r)] C_{s,d}^{d/s}} \left(\frac{h}{2(1+s/d)}\right)^{1+d/s} - \Delta - N^{1-s/d} \frac{r^{-s}}{\mu^q[B(\boldsymbol{x}, r)]} \right) \\ &\geq N^{s/d} \left(\frac{(h/2)^{1+d/s}}{(1+s/d) L_1^{d/s}} - \Delta - N^{1-s/d} \frac{r^{-s-d}}{c(h, \Omega, s, d)} \right), \qquad N \to \infty, \end{aligned} \tag{2.78}$$

where we used

$$\mu^q[B(\boldsymbol{x}, r)] \leq \left(\frac{L_1}{C_{s,d}(1+s/d)}\right)^{d/s} \mathcal{H}_d[B(\boldsymbol{x}, r)].$$

If we put $\Delta = \frac{(h/2)^{1+d/s}}{2(1+s/d)L_1^{d/s}}$, the inequality (2.78) implies that there exists a constant $C =$

$C(\Omega, h, s, d, q)$ for which

$$U_r(\boldsymbol{x}, \hat{\omega}_N) \geq C(\Omega, h, s, d, q) N^{s/d}, \quad N \in \mathbb{N}. \tag{2.79}$$

□

Proof of Theorem 2.10. Follows from Lemma 2.28 and Lemma 2.29. □

2.5 Combined kernels

The *combined* Riesz energy functional was introduced in (1.4) for $s \geq d = \dim_H \Omega$ as

$$E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, \kappa, q) = \sum_{i \neq j} \kappa(\boldsymbol{x}_i, \boldsymbol{x}_j) g_s(\boldsymbol{x}_i, \boldsymbol{x}_j) + \tau(N) \sum_i q(\boldsymbol{x}_i). \tag{2.80}$$

Similarly to how it was done for $E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, q)$ in Section 1.5.2, one can consider a sequence of functionals on $\mathcal{P}(\Omega)$ using (2.80) and to compute its Γ-limit, defined on the measures absolutely continuous with respect to $\mathcal{H}_d$ on Ω. In this section we shall derive an expression for the limiting functional, which will be denoted by $S(\cdot\,; g_s, \kappa, q)$. First, by an analog to (1.15),

$$\frac{1}{\mathcal{T}(N)} E_N(\mu; g_s, \kappa, q) := \begin{cases} \frac{1}{\mathcal{T}(N)} E(\{\boldsymbol{x} \,:\, \boldsymbol{x} \in \operatorname{supp} \mu\}; g_s, q) & \text{if } \mu \in \mathcal{P}_N(\Omega), \\ +\infty, & \text{otherwise.} \end{cases} \tag{2.81}$$

It has been shown in Proposition 1.6 that

$$\frac{1}{\mathcal{T}(N)} E_N(\cdot\,; g_s, 1, q) \xrightarrow{\Gamma} S(\cdot\,; g_s, 1, q), \qquad N \to \infty,$$

where

$$S(\mu; g_s, \kappa, q) := \begin{cases} C_{s,d} \int_\Omega \kappa(\boldsymbol{x}, \boldsymbol{x}) \varphi(\boldsymbol{x})^{1+s/d} \, d\mathcal{H}_d(\boldsymbol{x}) + \int_\Omega q(\boldsymbol{x}) \varphi(\boldsymbol{x}) \, d\mathcal{H}_d(\boldsymbol{x}), & \mu = \varphi \, d\mathcal{H}_d, \\ +\infty, & \text{otherwise.} \end{cases}$$

Note that we reuse the notation of (1.16) by simply adding another argument. Further, suppose following [15] that κ is strictly positive and continuous on the diagonal $\operatorname{diag}(\Omega \times \Omega)$, that is, $(\boldsymbol{x}_n, \boldsymbol{y}_n) \to (\boldsymbol{x}_0, \boldsymbol{x}_0)$ in the product topology of $\Omega \times \Omega$ implies $\kappa(\boldsymbol{x}_n, \boldsymbol{y}_n) \to \kappa(\boldsymbol{x}_0, \boldsymbol{x}_0)$. The following statement is a straightforward generalization of Proposition 1.6.

Proposition 2.30. *Suppose Ω is d-rectifiable . If κ is continuous at every $(\boldsymbol{x}, \boldsymbol{x}) \in \operatorname{diag}(\Omega \times \Omega)$,*

q is continuous on Ω, and g_s is the hypersingular Riesz kernel, then

$$\frac{1}{\mathcal{T}(N)} E_N(\cdot\,; g_s, \kappa, q) \xrightarrow{\Gamma} S(\mu; g_s, \kappa, q), \qquad N \to \infty, \tag{2.82}$$

on $\mathcal{P}(\Omega)$ equipped with the weak* topology.

Proof. Note that for any fixed $\delta > 0$, the contribution of pairs $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ for which $\|\boldsymbol{x}_i, \boldsymbol{x}_j\| \geq \delta$ to the LHS of (2.82) is zero, since there are at most $N^2 \ll \mathcal{T}(N)$ such pairs; this implies only the diagonal values $\kappa(\boldsymbol{x}, \boldsymbol{x})$ have impact on the Γ-limit. For additional details, see Lemma 2.26. Since the statement of the proposition is obvious when $\kappa \equiv C$ for some constant C, it suffices to construct a partition of Ω into a metrically separated collection of sets $\{\Omega_m\}_1^M$, with a remainder Ω_0 of small $\mathcal{H}_d$ measure, and to approximate κ with a constant on every element of the partition. The metric separation property will then allow to discard interactions of pairs $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ between different Ω_m, $1 \leq m \leq M$. Lastly, it will be convenient to assume that $q(\boldsymbol{x}) > 0$, $\boldsymbol{x} \in \Omega$; it is not a restrictive assumption since adding a constant to q translates into the same constant being added to both sides of (2.82).

To verify the property $\mathbf{1}^\Gamma$ of Γ-convergence, first fix a sequence $\{\mu_N\} \subset \mathcal{P}(\Omega)$ converging to $\mu \in \mathcal{P}(\Omega)$; just as in the proof of Proposition 1.6, it suffices to assume that $\mu_N \in \mathcal{P}_N(\Omega)$. By the same argument as in Proposition 1.6, μ must be absolutely continuous w.r.t. $\mathcal{H}_d$. Fix an $\varepsilon > 0$; we shall need a partition of Ω of the form

$$\Omega = \bigcup_{m=0}^{M} \Omega_m$$

with $\operatorname{dist}(\Omega_l, \Omega_m) \geq \sigma > 0$, $1 \leq l \neq m \leq M$; in addition, $\mathcal{H}_d(\partial\Omega_m) = 0$, $1 \leq m \leq M$ and all Ω_m are closed and such that $\bar{\kappa}_m - \underline{\kappa}_m < \varepsilon$ for $1 \leq m \leq M$. The latter is achieved by taking $\operatorname{diam}(\Omega_m)$ small enough, due to continuity of $\kappa(\boldsymbol{x}, \boldsymbol{x})$. By taking $\mathcal{H}_d(\Omega_0)$ small enough, we shall further guarantee that

$$\int_{\Omega_0} \left[C_{s,d} \kappa(\boldsymbol{x}, \boldsymbol{x}) \varphi(\boldsymbol{x})^{1+s/d} + q(\boldsymbol{x}) \varphi(\boldsymbol{x}) \right] d\mathcal{H}_d(\boldsymbol{x}) < \varepsilon,$$

where $\varphi := d\mu / d\mathcal{H}_d$. Such a partition is constructed in the proof of Theorem 2.1 using the Vitali covering lemma. As outlined above, we shall approximate $\kappa(\boldsymbol{x}, \boldsymbol{x})$ on Ω_m with $\underline{\kappa}_m := \min_{\Omega_m} \kappa(\boldsymbol{x}, \boldsymbol{x})$ and $\bar{\kappa}_m := \max_{\Omega_m} \kappa(\boldsymbol{x}, \boldsymbol{x})$; both are well-defined due to the continuity of κ on the diagonal.

As discussed above, the case of a constant $\kappa(\boldsymbol{x}, \boldsymbol{x})$ can be handled by the argument of Proposition 1.6, thereby giving

$$\frac{1}{\mathcal{T}(N)} E_N(\cdot\,; g_s, C, q) \xrightarrow{\Gamma} S(\cdot\,; g_s, C, q), \qquad N \to \infty, \quad 1 \leq m \leq M$$

for any constant C; observe that the continuity of q means that $S(\cdot\,; g_s, C, q)$ is a continuous perturbation of $S(\cdot\,; g_s, C, 0)$, [21, Proposition 2.3]. Denote the N-point supports of μ_N by ω_N, $N \geq 2$. We shall write

$$\mu^{(m)} := \frac{1}{\mu(\Omega_m)}\,\mu\big|_{\Omega_m}, \qquad \mu_N^{(m)} := \frac{1}{\mu_N(\Omega_m)}\,\mu_N\big|_{\Omega_m},$$

the for the renormalized restrictions of μ, μ_N to Ω_m, $1 \leq m \leq M$. In the case of $\mu(\Omega_m) = 0$ for some m, we shall join Ω_m with an Ω_l of positive measure and adjust the value of M below accordingly. By the property $\mathbf{1}^\Gamma$ of Γ-convergence, there holds

$$\lim_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \underline{\kappa}_m, q)}{\mathcal{T}(N\mu_N(\Omega_N))} \geq S(\mu^{(m)}; g_s, \underline{\kappa}_m, q), \quad 1 \leq m \leq M,$$

where we used the definition of the counting measures μ_N and, by the agreement made in Section 1.1, identified a counting measure with its support as an argument of an energy functional.

The second sum of $S(\cdot\,; g_s, \kappa, q)$, containing the external field $q(\boldsymbol{x}_i)$ terms, simply adds the $\int_\Omega q\varphi\, d\mathcal{H}_d$ term to the Γ-limit due to the stability under continuous perturbations [21, Proposition 2.3], so we shall focus on the weighted energy sum instead. Since $\mathcal{H}_d(\partial\Omega_m) = 0$, also $\mu(\partial\Omega_m) = 0$, and thus all the limits

$$\lim_{N\to\infty} \frac{\#\{\omega_N \cap \Omega_m\}}{N} = \lim_{N\to\infty} \mu_N(\Omega_m) = \mu(\Omega_m), \qquad 1 \leq m \leq M,$$

exist. This means, we can apply Lemma 3.7 (which is formulated for the unweighted Riesz energy, but applies here as well) to obtain

$$\begin{aligned}
\lim_{N\to\infty} \frac{1}{\mathcal{T}(N)} E_N(\mu_N; g_s, \kappa, 0) &\geq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} \liminf_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \kappa, 0)}{\mathcal{T}(N\mu_N(\Omega_N))} \\
&\geq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} \liminf_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \underline{\kappa}_m, 0)}{\mathcal{T}(N\mu_N(\Omega_N))} \\
&\geq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} S(\mu^{(m)}; g_s, \underline{\kappa}_m, 0).
\end{aligned}$$

To finish the proof of the property $\mathbf{1}^\Gamma$, we need to show that the sum in the RHS of the last inequality approximates the value of $S(\mu; g_s, \kappa, q)$. Indeed, first observe that

$$\begin{aligned}
S(\mu; g_s, \kappa, q) &= \sum_{m=0}^{M} \int_{\Omega_m} \Big[C_{s,d}\kappa(\boldsymbol{x}, \boldsymbol{x})\varphi(\boldsymbol{x})^{1+s/d} + q(\boldsymbol{x})\varphi(\boldsymbol{x}) \Big]\, d\mathcal{H}_d(\boldsymbol{x}) \\
&\leq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} S(\mu^{(m)}; g_s, \kappa, q) + \varepsilon,
\end{aligned}$$

since the functional $S(\cdot\,; g_s, \kappa, q)$ is defined on $\mathcal{P}(\Omega)$, and thus the respective densities $\varphi_m(\boldsymbol{x})$ have to be rescaled. Furthermore, for every $\boldsymbol{x} \in \Omega_m$, $0 \leq \kappa(\boldsymbol{x}, \boldsymbol{x}) - \underline{\kappa}_m \leq \bar{\kappa}_m - \underline{\kappa}_m < \varepsilon$; by taking ε small enough and applying the monotone convergence theorem, we complete the proof of the property $\mathbf{1}^\Gamma$ of Γ-convergence.

Fix a $\mu \in \mathcal{P}(\Omega)$. Constructing a recovery sequence for μ to verify the property $\mathbf{2}^\Gamma$ will involve a partition of Ω constructed to satisfy the same requirements as in the first part of the proof, but instead of approximating κ from below, we shall use $\bar{\kappa}_m$ to approximate it from above. First, for the remaining part of this proof let $\mu_N^{(m)}$ be the recovery sequences of $S(\cdot\,; g_s, \bar{\kappa}_m, q)$ for the restrictions $\mu^{(m)}$, $1 \leq m \leq M$, with the cardinalities $\#\omega_N^{(m)}$ adding up to N (thus assuming a notation not used elsewhere in this text; it will be rather helpful here); let also

$$\mu_N := \frac{1}{N} \sum_m \mu_N^{(m)} \#\omega_N^{(m)}.$$

The existence of such recovery sequences $\mu_N^{(m)}$ follows from Γ-convergence on individual Ω_m, $1 \leq m \leq M$. We thus have

$$\lim_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \bar{\kappa}_m, q)}{\mathcal{T}(N\mu_N(\Omega_N))} = S(\mu^{(m)}; g_s, \bar{\kappa}_m, q), \quad 1 \leq m \leq M,$$

from where

$$\begin{aligned}
\lim_{N\to\infty} \frac{1}{\mathcal{T}(N)} E_N(\mu_N; g_s, \kappa, 0) &= \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} \lim_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \kappa, 0)}{\mathcal{T}(N\mu_N(\Omega_N))} \\
&\leq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} \lim_{N\to\infty} \frac{E_N(\mu_N^{(m)}; g_s, \bar{\kappa}_m, 0)}{\mathcal{T}(N\mu_N(\Omega_N))} \\
&= \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} S(\mu^{(m)}; g_s, \bar{\kappa}_m, 0).
\end{aligned}$$

Since the positivity of κ and q implies

$$S(\mu; g_s, \kappa, q) \geq \sum_{m=1}^{M} \mu(\Omega_m)^{1+s/d} S(\mu^{(m)}; g_s, \kappa, q),$$

taking ε small enough and using the monotone convergence theorem completes the proof of (2.82). □

Corollary 2.31. *Any sequence $\{\omega_N\}$ that attains the lowest value of*

$$\lim_{N\to\infty} \frac{1}{\mathcal{T}(N)} E_N(\omega_N; g_s, \kappa, q)$$

converges weak to the probability measure $\mu^{\kappa,q}$ with density*

$$\varphi^{\kappa,q}(\boldsymbol{x}) = \frac{d\mu^{\kappa,q}}{d\mathcal{H}_d}(\boldsymbol{x}) := \left(\frac{L_1 - q(\boldsymbol{x})}{C_{s,d}(1+s/d)\kappa(\boldsymbol{x},\boldsymbol{x})}\right)_+^{d/s},$$

where L_1 is a normalizing constant.

Proof. The proof proceeds by a variational argument identical to the one given in Section 1.5.2. Uniqueness of the weak* limit follows from convexity of the functional $S(\cdot\,; g_s, \kappa, q)$ on probability measures $\mathcal{P}(\Omega)$ and convergence of minimizers, see Theorem 1.4. □

Chapter 3

Riesz energy on fractal sets

In the previous chapter we made a smoothness assumption on Ω: namely, that it be d-rectifiable. Observe that given an infinite set Ω, the *Riesz s-energy*

$$E_s(\omega_N) := E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, 0) = \sum_{i \neq j} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|^{-s}, \qquad N = 2, 3, 4, \ldots$$

on discrete subsets $\Omega \supset \omega_N = \{\boldsymbol{x}_i : 1 \leq i \leq N\}$ can always be defined, provided that $\Omega \subset \mathbb{R}^p$ is compact. Indeed, in this case it is a lower semicontinuous functional in the product topology on $(\mathbb{R}^p)^N$, which allows us to consider configurations minimizing it for a fixed N, as well as the asymptotics of $\mathcal{E}_s(\Omega, N) := \min\{E_s(\omega_N) : \omega_N \subset \Omega\}$ when $N \to \infty$.

As has been pointed out in Chapter 1, in the hypersingular case of $s > d := \dim_H \Omega$ for integer d, there exists extensive treatment of both the asymptotics of $\mathcal{E}_s(\Omega.N)$, and of the minimizing configurations in the weak* topology. Recall that we identify discrete sets $\omega_N = \{\boldsymbol{x}_i : 1 \leq i \leq N\} \subset \Omega$ with the corresponding (empirical) probability measures

$$\nu_N := \frac{1}{N} \sum_{i=1}^{N} \delta_{\boldsymbol{x}_i}.$$

Then, as summarized in Theorem 1.3 [61, 17], any sequence of minimizers of E_s on Ω^N weak* converges to the normalized d-dimensional Hausdorff measure $\mathcal{H}_d(\cdot \cap \Omega)/\mathcal{H}_d(\Omega)$ on Ω, provided that

$$\Omega = \Omega^{(0)} \cup \bigcup_{k=0}^{\infty} \Omega^{(k)}, \tag{3.1}$$

where each $\Omega^{(k)}$ is d-rectifiable and $\mathcal{H}_d(\Omega^{(0)}) = 0$, and that the d-dimensional Minkowski content of Ω coincides with $\mathcal{H}_d(\Omega)$. Moreover, under such assumptions the limit of $\mathcal{E}_s(\Omega, N)/N^{1+s/d}$, $N \to \infty$, exists; i.e., every sequence $\{\omega_N : N \geq 2\}$ of discrete subsets of Ω achieving the limit will converge weak* to $\mathcal{H}_d/\mathcal{H}_d(\Omega)$.

The requirement of Ω being representable as (3.1) is the weakest that is known to guarantee the Poppy-seed bagel theorem; on the other hand, it has been established [14, Proposition 2.6] that for a class of self-similar fractals F with $\dim_H F = d$, the limit of $\mathcal{E}_s(F, N)/N^{1+s/d}$ does not exist. Using this observation, [32] gives an example of a sequence of minimizers without a weak* limit.

In view of the above developments, it is natural to ask what can be said about weak* cluster points of $\{\nu_N : N \geq 2\}$ in the case when the underlying set Ω is not d-rectifiable; similarly, what is there to be said about cluster points of $\{\mathcal{E}_s(\Omega, N)/N^{1+s/d} : N \geq 2\}$. We will show

that when Ω belongs to the class of self-similar fractals satisfying the open set condition, any subsequence $\{\underline{\omega}_N : N \in \underline{\mathfrak{N}}\}$ achieving the limit inferior must converge weak* to $\mathcal{H}_d/\mathcal{H}_d(\Omega)$. Furthermore, in the case when contraction ratios of the fractal Ω are all equal, the energy limit $\mathcal{E}_s(\Omega, N)/N^{1+s/d}$, $N \in \mathfrak{N}$, along a sequence $\mathfrak{N} \subset \mathbb{N}$ can be characterized by the behavior of sequence itself.

The following section contains formal definitions and the prerequisites necessary to state our result; Section 3.2 formulates the main results, and Section 3.3 is dedicated to the proof of weak* convergence of sequences corresponding to the $\liminf$, as well as the results regarding a fractal with equal contraction ratios. This chapter is based on a joint work in progress by Alexander Reznikov and the author.

3.1 Self-similarity and open set condition

From now on in this chapter we shall be working with subsets of the Euclidean space $\mathbb{R}^p$. As above, bold typeface is used for its elements: $\boldsymbol{x} \in \mathbb{R}^p$. An open ball of radius r, centered at $\boldsymbol{x}$, is denoted by $B(\boldsymbol{x}, r)$. Recall also the standard definition of the weak* convergence: given a countable sequence $\{\mu_N : N \geq 1\}$ of probability measures supported on Ω and another probability measure μ,

$$\mu_N \xrightarrow{*} \mu,\ N \to \infty \quad \Longleftrightarrow \quad \int_A f(\boldsymbol{x}) d\mu_N(\boldsymbol{x}) \longrightarrow \int_A f(\boldsymbol{x}) d\mu(\boldsymbol{x}),\ N \to \infty,$$

for every $f \in C(\Omega)$. (Limits along nets are not necessary, as in this context weak* topology is metrizable.) A pair of sets $\Omega^{(1)}$, $\Omega^{(2)}$ will be called *metrically separated* if $\|\boldsymbol{x} - \boldsymbol{y}\| \geq \sigma > 0$ whenever $\boldsymbol{x} \in \Omega^{(1)}$ and $\boldsymbol{y} \in \Omega^{(2)}$. Recall that a *similitude* $\psi : \mathbb{R}^p \to \mathbb{R}^p$ can be written as

$$\psi(\boldsymbol{x}) = rM(\boldsymbol{x}) + \boldsymbol{z}$$

for some orthogonal matrix $M \in O(p)$, a vector $\boldsymbol{z} \in \mathbb{R}^p$, and a contraction ratio $0 < r < 1$. We shall assume the set $\Omega \subset \mathbb{R}^p$ to be a *self-similar fractal*, as defined by Hutchinson [69]. Namely, let $\Omega \subset \mathbb{R}^p$ be the compact set of fixed points of a collection of similitudes $\{\psi_m\}_{m=1}^M$ with contraction ratios r_m, $1 \leq m \leq M$, that is, satisfying

$$\Omega = \bigcup_{m=1}^{M} \psi_m(\Omega),$$

where the union is disjoint. Assume additionally that the maps $\{\psi_m\}$ satisfy the *open set condition*: there exists a bounded open $V \subset \mathbb{R}^p$ such that

$$\bigcup_{m=1}^{M} \psi_m(V) \subset V,$$

where the sets in the union are disjoint. Let d be such that

$$\sum_{m=1}^{M} r_m^d = 1. \tag{3.2}$$

Then d is equal to the Hausdorff dimension of Ω, $d = \dim_H \Omega$, [43, 80]. It will further be used that if Ω is a self-similar fractal satisfying the open set condition, then there holds $0 < \mathcal{H}_d(\Omega) < \infty$ and Ω is *d-regular* with respect to $\mathcal{H}_d$; that is, there exists a positive constant c, such that for every r, $0 < r \leq \mathrm{diam}(\Omega)$, and every $\boldsymbol{x} \in \Omega$,

$$c^{-1} r^d \leq \mathcal{H}_d(\Omega \cap B(\boldsymbol{x}, r)) \leq c r^d. \tag{3.3}$$

Consider the collection of all sequences $\{\omega_N : N \geq 2\}$ of discrete subsets of Ω, and let $\{\underline{\omega}_N : N \in \underline{\mathfrak{N}}\}$ denote a sequence such that

$$\lim_{\underline{\mathfrak{N}} \ni N \to \infty} \frac{E_s(\underline{\omega}_N)}{N^{1+s/d}} = \liminf_{N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} =: \underline{g}_{s,d}(\Omega), \tag{3.4}$$

and similarly, $\{\overline{\omega}_N : N \in \overline{\mathfrak{N}}\}$ a sequence for which

$$\lim_{\overline{\mathfrak{N}} \ni N \to \infty} \frac{E_s(\overline{\omega}_N)}{N^{1+s/d}} = \limsup_{N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} =: \bar{g}_{s,d}(\Omega). \tag{3.5}$$

Using the notation in (3.4)-(3.5), the precise statement of the non-existence of the limit of $\mathcal{E}_s(\Omega, N)/N^{1+s/d}$, $N \to \infty$ from [14] is as follows.

Proposition 3.1. *If Ω is a self-similar fractal as above, then there exists an $S_0 > 0$ such that for every $s > S_0$,*

$$0 < \underline{g}_{s,d}(\Omega) < \bar{g}_{s,d}(\Omega) < \infty.$$

It is further useful to recall that [80, Theorem 5.7] if a compact set is d-regular, it must have Hausdorff dimension d. Then the result of [32] can be fomulated in

Proposition 3.2. *Assume that the two d-regular compact sets $\Omega^{(1)}$, $\Omega^{(2)}$ are metrically separated and are such that $\Omega^{(1)}$ is a self-similar fractal as above and $\underline{g}_{s,d}(\Omega^{(2)}) = \bar{g}_{s,d}(\Omega^{(2)})$. Then for any sequence of minimizers of E_s, $\{\hat{\omega}_N : N \geq 2\}$, the corresponding sequence of measures $\{\hat{\nu}_N : N \geq 2\}$ does not have a weak* limit.*

In view of these two propositions, it is remarkable that the local properties of minimizers of E_s are fully preserved on self-similar fractals. Indeed, d-regularity of Ω can be readily used to obtain that any sequence of minimizers of E_s has the optimal orders of separation and covering. The following result was proved in [68]:

Proposition 3.3. *If $\Omega \subset \mathbb{R}^p$ is a compact d-regular set, $\{\hat{\omega}_N : N \geq 1\}$ a sequence of configurations minimizing E_s with $\hat{\omega}_N = \{\hat{\boldsymbol{x}}_i : 1 \leq i \leq N\}$, then there exist a constant $C_1 > 0$ such that for any $1 \leq i < j \leq N$,*

$$\|\hat{\boldsymbol{x}}_i - \hat{\boldsymbol{x}}_j\| \geq C_1 N^{-1/d}, \qquad N \geq 2,$$

and a constant $C_2 > 0$ such that for any $\boldsymbol{y} \in \Omega$,

$$\min_i \|\boldsymbol{y} - \hat{\boldsymbol{x}}_i\| \leq C_2 N^{-1/d}, \qquad N \geq 2.$$

As already noted, self-similar fractals have the d-regularity property; the same applies to finite unions and countable metrically separated unions of such sets.

3.2 Main results

As will become clear from Lemma 3.7, we require that the first order asymptotics of $\mathcal{E}_s(\Omega, N)$ grow faster than N^2 when $N \to \infty$; this is also the case when $s = d$, and the proofs in this section are fully applicable. We shall assume $s > d$ for simplicity; to obtain $s = d$ case, replace all instances of $N^{1+s/d}$ with $N^2 \log N$.

In accordance with the prior notation, we write $\underline{\omega}_N = \{\underline{\boldsymbol{x}}_i : 1 \leq i \leq N\}$ for the sequence of configurations with the lowest limit, and

$$\underline{\nu}_N = \frac{1}{N} \sum_{i=1}^{N} \delta_{\underline{\boldsymbol{x}}_i}, \quad N \in \underline{\mathfrak{N}}.$$

As described above, generally neither the weak* limit of minimizers of E_s, nor the limit of $\mathcal{E}_s(\Omega, N)/N^{1+s/d}$, $N \to \infty$, need to exist. We can still obtain the following statement about the sequence $\{\underline{\omega}_N : N \in \underline{\mathfrak{N}}\}$.

Theorem 3.4. *Let $\Omega \subset \mathbb{R}^p$ be a compact self-similar fractal satisfying the open set condition, and $\dim_H \Omega = d < s$. If $\{\underline{\omega}_N : N \in \underline{\mathfrak{N}}\}$, is a sequence of configurations for which*

$$\lim_{\underline{\mathfrak{N}} \ni N \to \infty} \frac{E_s(\underline{\omega}_N)}{N^{1+s/d}} = \underline{g}_{s,d}(\Omega),$$

then the corresponding sequence of empirical measures converges weak*:

$$\underline{\nu}_N \xrightarrow{*} \frac{\mathcal{H}_d}{\mathcal{H}_d(\Omega)}, \qquad \underline{\mathfrak{N}} \ni N \to \infty.$$

When the similitudes $\{\psi_m\}_{m=1}^M$ fixing Ω all have the same contraction ratio, it is natural to expect some additional symmetry of minimizers, associated with the M-fold scale symmetry of Ω.

Similarly, since the energy of interactions between particles in different $\Omega^{(m)}$ is at most of order N^2, see proof of Lemma 3.7 below, we expect that by acting with $\{\psi_m\}_{m=1}^M$ on a minimizer $\hat{\omega}_N$ with N large, we obtain a near-minimizer with MN elements. This heuristic is made rigorous in the following theorem.

Theorem 3.5. *Let $\Omega \subset \mathbb{R}^p$ be a self-similar fractal, fixed under M similitudes with the same contraction ratio, and $\mathfrak{M} = \{M^k n : k \geq 1\}$. Then the following limit exists*

$$\lim_{\mathfrak{M} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}}.$$

The previous theorem can be slightly extended to show that in the case of equal contraction ratios $r_1 = \ldots = r_m$, any sequence $\mathfrak{N}$ for which the corresponding sequence of minimizers has a limit must be of this form. Indeed, the following result establishes a bijection between cluster points of $\{\{\log_M N\} : N \in \mathfrak{N}\}$ and those of the sequence $\{\mathcal{E}_s(\Omega, N)/N^{1+s/d} : N \in \mathfrak{N}\}$. We shall need some notation first. For a sequence $\mathfrak{N}$, let

$$\{\mathfrak{N}\} := \lim_{\mathfrak{N} \ni N \to \infty} \{\log_M N\},$$

where $\{\cdot\}$ in the RHS denotes the fractional part, and

$$E_s(\mathfrak{N}) := \lim_{\mathfrak{N} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}},$$

if the corresponding limit exists.

Theorem 3.6. *If Ω is a self-similar fractal with equal contraction ratios, and two sequences $\mathfrak{N}_1, \mathfrak{N}_2 \subset \mathbb{N}$ are such that*

$$\{\mathfrak{N}_1\} = \{\mathfrak{N}_2\}, \tag{3.6}$$

then

$$E_s(\mathfrak{N}_1) = E_s(\mathfrak{N}_2). \tag{3.7}$$

In particular, the limits in (3.7) *exist. Moreover, the function $g_{s,d} : \{\mathfrak{N}\} \mapsto E_s(\mathfrak{N})$ is continuous on $[0, 1]$.*

3.3 Proofs

The key to proving Theorem 3.4 is that the hypersingular Riesz energy grows faster than N^2. We shall need this property in the following form.

Lemma 3.7. *Let a pair of compact sets $\Omega^{(1)}$, $\Omega^{(2)} \subset \mathbb{R}^p$ be metrically separated; let further*

$\{\omega_N \subset \Omega : N \in \mathfrak{N}\}$ *be a sequence for which the limits*

$$\lim_{\mathfrak{N}\ni N\to\infty} \frac{\#(\omega_N \cap \Omega^{(i)})}{N} = \beta^{(i)}, \quad i = 1, 2.$$

exist. Then

$$\liminf_{\mathfrak{N}\ni N\to\infty} \frac{E_s(\omega_N)}{N^{1+s/d}} \geq \left(\beta^{(1)}\right)^{1+s/d} \liminf_{\mathfrak{N}\ni N\to\infty} \frac{E_s(\omega_N \cap \Omega^{(1)})}{\#(\omega_N \cap \Omega^{(1)})^{1+s/d}} + \left(\beta^{(2)}\right)^{1+s/d} \liminf_{\mathfrak{N}\ni N\to\infty} \frac{E_s(\omega_N \cap \Omega^{(2)})}{\#(\omega_N \cap \Omega^{(2)})^{1+s/d}}.$$

Proof. We observe that with $\sigma = \operatorname{dist}(\Omega^{(1)}, \Omega^{(2)})$,

$$\left\| E_s(\omega_N) - \left(E_s(\omega_N \cap \Omega^{(1)}) + E_s(\omega_N \cap \Omega^{(2)})\right)\right\| = \sum_{\substack{\boldsymbol{x}_i\in\Omega^{(1)}\\ ,\boldsymbol{x}_j\in\Omega^{(2)}}} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|^{-s} \leq \sigma^{-s} N^2,$$

and use the definition of $\beta^{(i)}$, $i = 1, 2$, to obtain the desired equality. □

This is particularly useful due to the open set property. Consider a self-similar fractal satisfying it; since $\Omega \subset V$ for an open V and $\psi_m(V)$ are disjoint, there exists a $\sigma > 0$ such that $\operatorname{dist}(\psi_i(\Omega), \psi_j(\Omega)) \geq \sigma$ for $i \neq j$. Following [43], we will write

$$\Omega_{m_1\ldots m_l} := \psi_{m_1} \circ \ldots \circ \psi_{m_l}(\Omega), \qquad 1 \leq m_i \leq M, \quad l \geq 1.$$

Then $\operatorname{dist}(\Omega_{m_1\ldots m_l}, \Omega_{m'_1\ldots m'_l}) \geq r_{m_1} \ldots r_{m_k}\sigma$, where $k = \min\{i : m_i \neq m'_i\}$, so for a fixed M in the expression

$$\Omega = \bigcup_{m_1,\ldots,m_l=1}^{M} \Omega_{m_1\ldots m_l}$$

not only the union is disjoint, but also the sets $\Omega_{m_1\ldots m_l}$ are metrically separated.

Lemma 3.8. *If $\{\mu_N : N \in \mathfrak{N}\}$ is a sequence of probability measures on the set Ω, which for every $l \geq 1$ satisfies*

$$\lim_{\mathfrak{N}\ni N\to\infty} \mu_N(\Omega_{m_1\ldots m_l}) = \mu(\Omega_{m_1\ldots m_l}), \qquad 1 \leq m_1, \ldots, m_l \leq M,$$

for another probability measure μ on Ω, then

$$\mu_N \xrightarrow{*} \mu, \qquad \mathfrak{N} \ni N \to \infty.$$

Proof. Fix an $f \in C(\Omega)$; since Ω is compact, f is uniformly continuous on Ω. Let $\varepsilon > 0$ be also fixed. By the uniform continuity of f, there exists an $L_0 \in \mathbb{N}$ such that $|f(x) - f(y)| < \varepsilon$

whenever $x, y \in \Omega_{m_1,\ldots,m_l}$ for any $l \geq L_0$ and any set of indices $0 \leq m_1, \ldots, m_l \leq M$; this is possible due to

$$\operatorname{diam}(\Omega_{m_1,\ldots,m_l}) \leq r_{m_1} \ldots r_{m_l} \operatorname{diam}(\Omega) \leq \left(\max_{1 \leq m \leq M} r_m \right)^l \operatorname{diam}(\Omega).$$

Fix an $l \geq L_0$ until the end of this proof, then pick an $N_0 \in \mathfrak{N}$ so that for every $N \geq N_0$, there holds

$$|\mu_N(\Omega_{m_1 \ldots m_l}) - \mu(\Omega_{m_1 \ldots m_l})| < \varepsilon / M^l, \qquad 1 \leq m_1, \ldots, m_l \leq M.$$

Finally, let us write $f_{m_1 \ldots m_l} := \min_{\Omega_{m_1 \ldots m_l}} f(x)$ for brevity. Then for $N \geq N_0$,

$$\begin{aligned}
\left| \int_A f(x) d\mu_N(x) - \int_A f(x) d\mu(x) \right| & \\
&\leq \sum_{m_1,\ldots,m_l=1}^{M} \left| \int_{\Omega_m} (f(x) - f_{m_1 \ldots m_l}) d\mu_N(x) - \int_{\Omega_m} (f(x) - f_{m_1 \ldots m_l}) d\mu(x) \right| \\
&\quad + \sum_{m_1,\ldots,m_l=1}^{M} |(\mu_N(\Omega_m) - \mu(\Omega_m)) f_{m_1 \ldots m_l}| \\
&\leq 2\varepsilon + \varepsilon \|f\|_\infty,
\end{aligned}$$

where the estimate for the first sum used that both μ_N and μ are probability measures. This proves the desired statement. □

Note that the converse is also true: since the sets $\Omega_{m_1,\ldots,m_l}$ are metrically separated, convergence $\mu_N \overset{*}{\rightharpoonup} \mu$ of measures supported on Ω immediately implies (by Urysohn's lemma) $\mu_N(\Omega_{m_1 \ldots m_l}) \to \mu(\Omega_{m_1 \ldots m_l})$ for all $l \geq 1$ and all indices $1 \leq m_1, \ldots, m_l \leq M$.

Lemma 3.9. *If Ω is a compact d-regular set, then $0 < \underline{g}_{s,d}(\Omega) \leq \bar{g}_{s,d}(\Omega) < \infty$.*

Proof. We shall provide a concise proof for the sake of completeness; in fact, it suffices to follow the well-known approach, found for example in [72, 68]. First, for the lower bound observe that for any $\omega_N = \{\boldsymbol{x}_i : 1 \leq i \leq N\}$, applying Jensen's inequality to the function $t \mapsto t^{-s/d}$ gives

$$E_s(\omega_N) = \sum_{i \neq j} |\boldsymbol{x}_i - \boldsymbol{x}_j|^{-s} \geq \sum_{i=1}^{N} \left(\Delta_i^d \right)^{-s/d} \geq N \left(\frac{\sum_{i=1}^N \Delta_i^d}{N} \right)^{-s/d} = N^{1+s/d} \left(\sum_{i=1}^{N} \Delta_i^d \right)^{-s/d},$$

where we denote $\Delta_i := \min_{j:j \neq i} \|\boldsymbol{x}_j - \boldsymbol{x}_i\|$. Thus bounding the RHS of the previous equation from below will give the desired result. Let $\{B_i = B(\boldsymbol{x}_i, \Delta_i/3) : 1 \leq i \leq N\}$ be a collection of pairwise disjoint closed balls of radii $\{\Delta_i/3\}$, centered around $\{\boldsymbol{x}_i\}$. Then the first inequality in

d-regularity of Ω (3.3) implies

$$\mathcal{H}_d(\Omega) \geq \sum_{i=1}^{N} \mathcal{H}_d(\Omega \cap B_i) \geq c^{-1} \sum_{i=1}^{N} (\Delta_i/3)^d,$$

which yields

$$E_s(\omega_N) \geq \frac{N^{1+s/d}}{3^s \left(c\mathcal{H}_d(\Omega)\right)^{s/d}}. \tag{3.8}$$

To obtain an upper bound, consider a set of minimizers $\hat{\omega}_N = \{\hat{\boldsymbol{x}}_i : 1 \leq i \leq N\}$ of E_s; then for any $\boldsymbol{x} \in \Omega$ and $1 \leq i \leq N$, by optimality of the set $\hat{\omega}_N$ there holds

$$\min_{\boldsymbol{x}\in\Omega} \sum_{j:j\neq i} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_j\|^{-s} = \sum_{j:j\neq i} \|\hat{\boldsymbol{x}}_i - \hat{\boldsymbol{x}}_j\|^{-s} \leq \sum_{j:j\neq i} \|\boldsymbol{x} - \hat{\boldsymbol{x}}_j\|^{-s}.$$

Denoting $r = \left(\mathcal{H}_d(\Omega)/2cN\right)^{1/d}$ and integrating over $\boldsymbol{x} \in \Omega \setminus \bigcup_i B\left(\hat{\boldsymbol{x}}_i, r\right)$ gives by Fubini's theorem

$$\begin{aligned}\frac{\mathcal{H}_d(\Omega)}{2} \sum_{j:j\neq i} \|\hat{\boldsymbol{x}}_i - \hat{\boldsymbol{x}}_j\|^{-s} &\leq \sum_{j:j\neq i} \int_0^{r^{-s}} \mathcal{H}_d(\{\boldsymbol{x} \in \Omega \,:\, \|\boldsymbol{x} - \hat{\boldsymbol{x}}_j\| \leq t^{-1/s}\})\, dt \\ &\leq cN \int_0^{r^{-s}} t^{-d/s}\, dt = \frac{csN}{s-d} \left(\frac{\mathcal{H}_d(\Omega)}{2cN}\right)^{(d-s)/d}.\end{aligned}$$

After adding the above inequalities over $1 \leq i \leq N$:

$$\mathcal{E}_s(\Omega, N) \leq \frac{s(2c)^{s/d}}{s-d} \frac{N^{1+s/d}}{\mathcal{H}_d(\Omega)^{s/d}}.$$

□

The above argument gives a somewhat stronger statement, since both bounds hold for finite values of N; furthermore, each bound requires only one of the inequalities in (3.3). Equation (3.8) implies that for any sequence of configurations ω_N, $N \in \mathfrak{N}$, with

$$\lim_{\mathfrak{N}\ni N\to\infty} \frac{E_s(\omega_N)}{N^{1+s/d}} < \infty,$$

every weak* cluster point of ν_N, $N \in \mathfrak{N}$, must be absolutely continuous with respect to $\mathcal{H}_d$ on Ω. Lastly, observe that the proof of Lemma 3.9 shows also that for an arbitrary collection of distinct points $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{N-1} \subset \Omega$, the minimal value of the point energy is bounded:

$$\min_{\boldsymbol{x}\in\Omega} \sum_{j=1}^{N-1} \|\boldsymbol{x} - \boldsymbol{x}_j\|^{-s} \leq \frac{sc^{s/d}}{s-d} \left(\frac{\mathcal{H}_d(\Omega)}{2}\right)^{-s/d} N^{s/d}.$$

It will be subsequently used that the RHS in the last inequality does not depend on N. The above bound can therefore be summarized as follows.

Corollary 3.10. *Suppose Ω is a compact d-regular set, $\omega_N = \{\boldsymbol{x}_i : 1 \le i \le N\} \subset \Omega$, and $s > d$. Then the minimal point energy of ω_N is bounded by:*

$$\min_{\boldsymbol{x}\in\Omega} \sum_{j=1}^{N} \|\boldsymbol{x} - \boldsymbol{x}_j\|^{-s} \le CN^{s/d},$$

where C depends only on Ω, s, d.

Proof of Theorem 3.4. In view of the weak* compactness of probability measures in Ω, to establish existence of the weak* limit of $\underline{\nu}_N$, $N \in \underline{\mathfrak{N}}$, it suffices to show that any cluster point of $\underline{\nu}_N$, $N \in \underline{\mathfrak{N}}$, in the weak* topology is $\mathcal{H}_d/\mathcal{H}_d(\Omega)$ by a standard argument (see [37, Proposition A.2.7]). To that end, consider a subsequence of $\underline{\mathfrak{N}}$ for which the empirical measures $\underline{\nu}_N$ converge to a cluster point μ; for simplicity we shall use the same notation $\underline{\mathfrak{N}}$ for this subsequence.

As shown above, $\underline{\nu}_N(\Omega_{m_1\dots m_l}) \to \mu(\Omega_{m_1\dots m_l})$, $\underline{\mathfrak{N}} \ni N \to \infty$; this ensures that the quantities

$$\underline{\beta}_m := \mu(\Omega_m) = \lim_{\underline{\mathfrak{N}}\ni N\to\infty} \underline{\nu}_N(\Omega_m) = \lim_{\underline{\mathfrak{N}}\ni N\to\infty} \frac{\#(\underline{\omega}_N \cap \Omega_m)}{N}, \qquad m = 1, \dots, M,$$

are well-defined. From (3.4), separation of $\{\Omega_m\}$, and Lemma 3.7 follows

$$\begin{aligned} \underline{g}_{s,d}(\Omega) &= \sum_{m=1}^{M} \lim_{\underline{\mathfrak{N}}\ni N\to\infty} \frac{E_s(\underline{\omega}_N \cap \Omega_m)}{N^{1+s/d}} \ge \sum_{m=1}^{M} \underline{\beta}_m^{1+s/d} \liminf_{\underline{\mathfrak{N}}\ni N\to\infty} \frac{E_s(\underline{\omega}_N \cap \Omega_m)}{\#(\underline{\omega}_N \cap \Omega_m)^{1+s/d}} \\ &\ge \sum_{m=1}^{M} \underline{\beta}_m^{1+s/d} r_m^{-s}\, \underline{g}_{s,d}(\Omega). \end{aligned}$$

Consider the RHS in the last inequality. As a function of $\{\underline{\beta}_m\}$, it satisfies the constraint $\sum_m \underline{\beta}_m = 1$; note also that by the defining property (3.2) of d, there holds $\sum_m R_m = 1$ with $R_m := r_m^d$, $1 \le m \le M$. We have

$$\underline{g}_{s,d}(\Omega) \ge \inf\left\{ \sum_{m=1}^{M} \beta_m^{1+s/d} R_m^{-s/d} \;:\; \sum_{m=1}^{M} \beta_m = 1 \right\} \underline{g}_{s,d}(\Omega). \tag{3.9}$$

Level sets of the function $\sum_m \beta_m^{1+s/d} R_m^{-s/d}$ are convex, so the infimum is attained and unique; it is easy to check that the solution is at $\beta_m = R_m = r_m^d$, $1 \le m \le M$, and the minimal value is 1. Indeed, the corresponding Lagrangian is

$$L(\beta_1, \dots, \beta_M, \lambda) := \sum_{m=1}^{M} \beta_m^{1+s/d} R_m^{-s/d} - \lambda \sum_{m=1}^{M} \beta_m,$$

hence

$$\nabla L_{\beta_m} = (1+s/d)\left(\frac{\beta_m}{R_m}\right)^{s/d} - \lambda, \quad 1 \le m \le M,$$

and it remains to use $\beta_m \ge 0$, $1 \le m \le M$, and $\sum_m R_m = 1$, to conclude $\beta_m = R_m$, $1 \le m \le M$.

Since $0 < \underline{g}_{s,d}(\Omega) < \infty$ by Lemma 3.9, from (3.9) it follows

$$\underline{\beta}_m = r_m^d, \qquad m = 1, \dots, M.$$

Note that this argument shows also

$$\lim_{\underline{\mathfrak{N}} \ni N \to \infty} \frac{E_s(\underline{\omega}_N \cap \Omega_m)}{(\#(\underline{\omega}_N \cap \Omega_m))^{1+s/d}} = \underline{g}_{s,d}(\Omega),$$

so the above can be repeated recursively for sets $\Omega_{m_1 \dots m_l}$. Namely, for every $l \ge 1$ and $1 \le m, m_1, \dots, m_l \le M$,

$$\mu(\Omega_{m m_1 \dots m_l}) =: \underline{\beta}_{m m_1 \dots m_l} = r_m^d \, \underline{\beta}_{m_1 \dots m_l}.$$

Observe further that $\mathcal{H}_d$ satisfies

$$\mathcal{H}_d(\Omega_{m m_1 \dots m_l}) = r_m^d \mathcal{H}_d(\Omega_{m_1 \dots m_l})$$

by definition, so by Lemma 3.8 follows that every weak* cluster point of $\underline{\nu}_N$, $N \in \underline{\mathfrak{N}}$, is $\mathcal{H}_d/\mathcal{H}_d(\Omega)$, as desired. □

Proof of Theorem 3.5. Note that setting equal contraction ratios $r_1 = \dots = r_m = r$ in (3.2) gives $r^{-s} = M^{s/d}$. Consider the set function

$$\psi : \boldsymbol{x} \mapsto \bigcup_{m=1}^{M} \psi_m(\boldsymbol{x}), \quad \boldsymbol{x} \in \Omega,$$

and denote

$$\psi(\omega_\varepsilon) := \bigcup_{\boldsymbol{x} \in \omega_\varepsilon} \psi(x).$$

It follows from the open set condition that the union above is metrically separated; we shall denote the separation distance by σ. Observe that the definition of a similitude implies $\#(\psi(\omega_\varepsilon)) = M\#(\omega_\varepsilon)$. We then have for any configuration ω_N, $N \ge 2$,

$$\begin{aligned} \mathcal{E}_s(\Omega, MN) &\le E_s(\psi(\omega_N)) \le M r^{-s} E_s(\omega_N) + \sigma^{-s} N^2 M^2 \\ &= M^{1+s/d} E_s(\omega_N) + \sigma^{-s} N^2 M^2, \end{aligned}$$

and repeated application of the second inequality yields

$$\begin{aligned}\mathcal{E}_s(\Omega, M^k N) &\le E_s[\psi(\psi^{(k-1)}(\omega_N))] \le M^{1+s/d} E_s(\psi^{(k-1)}(\omega_N)) + \sigma^{-s}(M^{k-1}N)^2 M^2 \\ &\le (M^2)^{1+s/d} E_s(\psi^{(k-2)}(\omega_N)) + M^{1+s/d}\sigma^{-s}(M^{k-2}N)^2 M^2 + \sigma^{-s}(M^{k-1}N)^2 M^2 \\ &\le \dots \\ &\le (M^k)^{1+s/d} E_s(\omega_N) + \sigma^{-s} N^2 \sum_{l=1}^{k} (M^{l-1})^{1+s/d} (M^{k-l})^2 M^2.\end{aligned}$$

Estimating the geometric series in the last inequality, we obtain

$$\begin{aligned}\mathcal{E}_s(\Omega, M^k N) &\le (M^k)^{1+s/d} E_s(\omega_N) + \sigma^{-s} N^2 M^{2k+1-s/d} \sum_{l=1}^{k} M^{l(s/d-1)} \\ &\le (M^k)^{1+s/d} E_s(\omega_N) + \sigma^{-s} N^2 M^{2k-s/d} \left(2M^{(k+1)(s/d-1)}\right) \\ &= (M^k)^{1+s/d} E_s(\omega_N) + \frac{2}{\sigma^s M} N^{1-s/d} \left(M^k N\right)^{(1+s/d)}.\end{aligned} \tag{3.10}$$

Let now $\varepsilon > 0$ fixed; find ω_{N_0} such that $N_0 \in \mathfrak{M}$ and

$$\frac{E_s(\omega_{N_0})}{N_0^{1+s/d}} \le \liminf_{\mathfrak{M} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} + \varepsilon,$$

and in addition, $2N_0^{1-s/d} < \varepsilon \sigma^s M$. Then by (3.10) we have

$$\frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \le \frac{E_s(\omega_{N_0})}{N_0^{1+s/d}} + \varepsilon \le \liminf_{\mathfrak{M} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} + 2\varepsilon, \qquad \mathfrak{M} \ni N \ge N_0.$$

This proves the desired statement. □

In the following lemma we write $\mathfrak{N}(k), k \in \mathbb{N}$, to denote the k-th element of the sequence $\mathfrak{N} \subset \mathbb{N}$; we say that $\mathfrak{N}$ is *majorized* by a sequence $\mathfrak{M}$, if the inequality $\mathfrak{N}(k) < \mathfrak{M}(k)$ holds for every $k \ge 1$.

Lemma 3.11. *If $\mathfrak{M} \subset \mathbb{N}$ is a sequence such that the limit*

$$\lim_{\mathfrak{M} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}}$$

exists, then for any sequence of integers $\mathfrak{N} \subset \mathbb{Z}$ with $|N(k)|$ majorized by $\mathfrak{M}$ and satisfying $|\mathfrak{N}(k)| = o(\mathfrak{M}(k))$, $k \to \infty$, there holds

$$\lim_{(\mathfrak{M}+\mathfrak{N}) \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} = \lim_{\mathfrak{M} \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}}, \tag{3.11}$$

where the addition $\mathfrak{M}+\mathfrak{N}$ is performed elementwise.

Proof. First, observe that by passing to subsequences of $\mathfrak{M}$ and $\mathfrak{N}$, it suffices to assume $\mathfrak{N}(k)\geq 0$ and to show(3.11) for $\mathfrak{M}+\mathfrak{N}$ and $\mathfrak{M}-\mathfrak{N}$. Furthermore, inequality

$$\mathcal{E}_s[\Omega,(\mathfrak{M}+\mathfrak{N})(k)]\geq \mathcal{E}_s(\Omega,\mathfrak{M}(k))$$

holds by the definition of $\mathcal{E}_s$. Thus

$$\begin{aligned}\liminf_{(\mathfrak{M}+\mathfrak{N})\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}&\geq \lim_{k\to\infty}\frac{\mathcal{E}_s(\Omega,\mathfrak{M}(k))}{(\mathfrak{M}(k)+\mathfrak{N}(k))^{1+s/d}}\\ =\lim_{k\to\infty}\frac{\mathcal{E}_s(\Omega,\mathfrak{M}(k))}{(\mathfrak{M}(k))^{1+s/d}}\left(\frac{\mathfrak{M}(k)}{\mathfrak{M}(k)+\mathfrak{N}(k)}\right)^{1+s/d}&=\lim_{\mathfrak{M}\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}},\end{aligned}\tag{$*$}$$

in view of $\mathfrak{N}(k)=o(\mathfrak{M}(k))$. Similarly,

$$\limsup_{(\mathfrak{M}-\mathfrak{N})\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}\leq\lim_{\mathfrak{M}\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}.\tag{$*$}$$

For the converse estimates, use Corollary 3.10 to conclude that for every $N\in\mathbb{N}$ there holds

$$\mathcal{E}_s(\Omega,N+1)\leq\mathcal{E}_s(\Omega,N)+CN^{s/d}.$$

Applying this inequality $\mathfrak{N}(k)$ times to $\mathfrak{M}(k)$, we obtain

$$\mathcal{E}_s[\Omega,(\mathfrak{M}+\mathfrak{N})(k)]\leq\mathcal{E}_s(\Omega,\mathfrak{M}(k))+\mathfrak{N}(k)C[\mathfrak{M}(k)+\mathfrak{N}(k)]^{s/d},$$

which yields

$$\limsup_{(\mathfrak{M}+\mathfrak{N})\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}\leq\lim_{\mathfrak{M}\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}.\tag{$*$}$$

Finally, applying Corollary 3.10 $\mathfrak{N}(k)$ times to $\mathfrak{M}(k)-\mathfrak{N}(k)$ gives

$$\mathcal{E}_s[\Omega,\mathfrak{M}(k)]\leq\mathcal{E}_s[\Omega,(\mathfrak{M}-\mathfrak{N})(k)]+\mathfrak{N}(k)C\mathfrak{M}(k)^{s/d},$$

whence, using that $\mathfrak{N}(k)=o(\mathfrak{M}(k))$, $k\to\infty$,

$$\liminf_{(\mathfrak{M}-\mathfrak{N})\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}\geq\lim_{\mathfrak{M}\ni N\to\infty}\frac{\mathcal{E}_s(\Omega,N)}{N^{1+s/d}}.\tag{$*$}$$

Combining inequalities marked by $(*)$, we obtain the desired result. □

The proof of the previous lemma gives also the following.

Corollary 3.12. *If $\mathfrak{M}, \mathfrak{N} \subset \mathbb{N}$ are a pair of sequences such that*

$$\mathfrak{N}(k) \le \theta\, \mathfrak{M}(k), \qquad k \ge 1,$$

then

$$\liminf_{(\mathfrak{M}+\mathfrak{N})\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \ge \liminf_{\mathfrak{M}\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \cdot \left(\frac{1}{1+\theta}\right)^{1+s/d}$$

and

$$\limsup_{(\mathfrak{M}+\mathfrak{N})\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \le \limsup_{\mathfrak{M}\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} + \frac{C\theta}{1+\theta},$$

where C is the same as in Corollary 3.10.

Proof of Theorem 3.6. To show that $g_{s,d}(\cdot)$ is well-defined, it is necessary to verify that (i) existence of the limit $\{\mathfrak{N}\}$ implies that of the limit $E_s(\mathfrak{N})$, and (ii) the value of $E_s(\mathfrak{N})$ is uniquely defined by $\{\mathfrak{N}\}$. First assume that $\mathfrak{N}_1$, $\mathfrak{N}_2$ are multiples of (a subset of) the geometric series, that is, $\mathfrak{N}_i = \{M^k n_i \,:\, k \in \mathcal{K}_i\}$, $i = 1, 2$. Observe that (3.6) implies $\{\log_M n_1\} = \{\log_M n_2\}$ and let for definiteness $n_2 \ge n_1$; then $n_2 = M^{k_0} n_1$ for some $k_0 \ge 1$. It follows that $\mathfrak{N}_i \subset \mathfrak{N}_0$, $i = 1, 2$, with $\mathfrak{N}_0 = \{M^k n_0 \,:\, k \ge 1\}$. By Theorem 3.5, the limit

$$\lim_{\mathfrak{N}_0\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}}$$

exists, so it must be that the limits over subsequences of $\mathfrak{N}_0$

$$\lim_{\mathfrak{N}_i\ni N\to\infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}}, \qquad i = 1, 2,$$

also exist and are equal, so the function $g_{s,d}(\cdot)$ is well-defined on the subset of $[0, 1]$ of the form $\{\mathfrak{N}\}$ with $\mathfrak{N} = \{M^k n \,:\, k \in \mathcal{K}\}$.

Now let $\mathfrak{N}_1$, $\mathfrak{N}_2 \subset \mathbb{N}$ be arbitrary. Denote the common value of the limit $a := \{\mathfrak{N}_i\}$, $i = 1, 2$. We shall assume for definiteness that $a \in (0, 1)$; the case of $a \in \{0, 1\}$ can be handled similarly. In order to bound $\mathfrak{N}_i$ between two sequences of the type $\{M^k n_i \,:\, k \in \mathcal{K}_i\}$, discussed above, fix an $\varepsilon > 0$ such that $a - 2\varepsilon > 0$ and $a + 2\varepsilon < 1$, and find an $N_0 \in \mathbb{N}$, for which

$$|\{\log_M N_i\} - a| < \varepsilon, \qquad N_0 \le N_i \in \mathfrak{N}_i, \quad i = 1, 2. \tag{3.12}$$

By the choice of ε, the above equation gives $\lfloor\{\log_M N_1\}\rfloor = \lfloor\{\log_M N_2\}\rfloor$ when $N_0 \le N_i \in \mathfrak{N}_i$. Now let n_i, $i = 1, 2$ be such that

$$\begin{aligned} a - 2\varepsilon \le \{\log_M n_1\} \le a - \varepsilon \\ a + \varepsilon \le \{\log_M n_2\} \le a + 2\varepsilon. \end{aligned} \tag{3.13}$$

Replacing one of n_i, $i = 1, 2$, with its multiple, if necessary, we can guarantee that $0 < \log_M n_2 - \log_M n_1 < 4\varepsilon$. Consider a pair of sequences $\widetilde{\mathfrak{N}}_i = \{M^k n_i : k \geq \lceil \log_M N_0 \rceil\}$, $i = 1, 2$; observe that by the above argument, limits

$$E_s(\widetilde{\mathfrak{N}}_i) =: L_i, \qquad i = 1, 2,$$

along $\widetilde{\mathfrak{N}}_i$, $i = 1, 2$, both exist, and the inequality

$$\widetilde{\mathfrak{N}}_1(k) \leq N_i \leq \widetilde{\mathfrak{N}}_2(k), \qquad k = \lfloor \log_M N_i \rfloor, \quad N_0 \leq N_i \in \mathfrak{N}_i,\ i = 1, 2,$$

holds. By the definition of $\mathcal{E}_s$, and due to (3.12)–(3.13),

$$\limsup_{\mathfrak{N}_i \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \leq \lim_{k \to \infty} \frac{\mathcal{E}_s(\Omega, M^k n_2)}{(M^k n_1)^{1+s/d}} = \left(\frac{n_2}{n_1}\right)^{1+s/d} L_2, \qquad i = 1, 2,$$

and

$$\liminf_{\mathfrak{N}_i \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \geq \lim_{k \to \infty} \frac{\mathcal{E}_s(\Omega, M^k n_1)}{(M^k n_2)^{1+s/d}} = \left(\frac{n_1}{n_2}\right)^{1+s/d} L_1, \qquad i = 1, 2.$$

Combining the last two inequalities gives

$$\left(\frac{n_1}{n_2}\right)^{1+s/d} L_1 \leq \liminf_{\mathfrak{N}_i \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \leq \limsup_{\mathfrak{N}_i \ni N \to \infty} \frac{\mathcal{E}_s(\Omega, N)}{N^{1+s/d}} \leq \left(\frac{n_2}{n_1}\right)^{1+s/d} L_2,$$

so it suffices to show that L_2 can be made arbitrarily close to L_1 by taking $\varepsilon \to 0$. The latter follows from Corollary 3.12, and the choice of n_i, $i = 1, 2$:

$$0 \leq \frac{\widetilde{\mathfrak{N}}_2(k) - \widetilde{\mathfrak{N}}_1(k)}{\widetilde{\mathfrak{N}}_1(k)} = \frac{n_2}{n_1} - 1 \leq M^{4\varepsilon} - 1.$$

Taking $\varepsilon \to 0$ shows both that $E_s(\mathfrak{N}_1) = E_s(\mathfrak{N}_2)$, and that these two limits exist. The function $g_{s,d} : [0, 1] \to (0, \infty)$ is therefore well-defined. Note that repeating the above argument for $|\{\mathfrak{N}_1\} - \{\mathfrak{N}_2\}| < \varepsilon$ for a fixed positive ε gives a bound on $|E_s(\mathfrak{N}_1) - E_s(\mathfrak{N}_2)|$, which implies that $g_{s,d}$ is continuous. This completes the proof. □

Chapter 4

Applications to meshless methods

The present chapter discusses construction of stencils for the Radial Basis Function-generated Finite Differences (RBF-FD) method, by applying a Riesz energy gradient flow to an initial Quasi-Monte Carlo (Q-MC) configuration. This approach targets primarily meshless finite difference methods, see Section 4.2.1, but can be applied to a wide range of problems that require generation of a point cloud with controlled local properties.

The section structure is as follows: Section 4.2.1 outlines the RBF-FD method using Gaussian and Polyharmonic Spline kernels; Sections 4.2.2 and 4.2.3 introduce the two essential components of our approach, Riesz energy functionals and quasi-Monte Carlo methods. The main algorithm and its discussion are the subjects of Sections 4.3.1 and 4.3.2, respectively. Sections 4.4.1–4.4.3 offer applications of the algorithm; the corresponding run times are summarized in Section 4.4.4. Section 4.5.1 contains comparisons of the condition numbers of RBF-FD matrices with stencils on periodic Riesz minimizers, Halton nodes, and the Cartesian grid; Section 4.5.2 discusses the range of dimensions where the present method is applicable. Lastly, Section 4.6 is dedicated to numerical experiments with the mean and minimal separation distance of Riesz minimizers and irrational lattices. Throughout the chapter the underlying set has full dimension: $\dim_H \Omega = d = p$; we shall denote it by d. This chapter reproduces with minor modifications the joint paper with Natasha Flyer, Bengt Fornberg, and Timothy Michaels [104, 103], published electronically in Computers & Mathematics with Applications.

4.1 Formulation of the problem

In a number of important applications, usefulness of meshless methods in general, and of radial basis functions (RBFs) in particular, is well-known. They have found their way into high-dimensional interpolation, machine learning, spectral methods, vector-valued approximation and interpolation, just to name a few [105, 33, 49, 27, 97]. RBFs have multiple advantages, most importantly extreme flexibility in forming stencils (in the case of RBF-FD) and high local adaptivity; allowing spectral accuracy on irregular domains; the fact that the corresponding interpolation matrix (denoted by $\boldsymbol{A}$ below) is positive definite for several types of radial functions and does not suffer from instability phenomena characteristic of some alternative interpolation methods.

Applying RBF-FD stencils to building solvers requires an efficient way of distributing the *nodes* of basis elements in the domain, which can be either a solid or a surface. The tasks of modeling and simulation often call for massive numbers of nodes, so it is important to ensure

that the distribution process is easily scalable. One further has to be able to place the RBFs according to a certain density, as a method of local refinement, for example, at the boundary, or in regions of special interest. Yet another challenge arises when it is necessary to deal with complex or non-smooth domains and/or surfaces.

Recall [51] that an RBF is a linear combination of the form

$$\mathcal{S}(\boldsymbol{x}) = \sum_{k=1}^{K} c_k \varphi(\|\boldsymbol{x} - \boldsymbol{x}_k\|), \tag{4.1}$$

where $\varphi(\cdot)$ is a radial function, and $\boldsymbol{x}_k,\ k = 1, \ldots, K$, is a collection of pairwise distinct points in $\mathbb{R}^d$. A common choice of φ is the Gaussian $\varphi(r) = e^{-(\varepsilon r)^2}$, although one may also use $1/(1+(\varepsilon r)^2)$, $r^{2p}\log(r),\ p \in \mathbb{N}$, etc. In this discussion, we are not concerned with the distinctions between the different radial kernels, so the reader can assume that $\varphi(r) = e^{-(\varepsilon r)^2}$. In contrast to pseudospectral methods [49], RBF-FD approach means that to obtain a useful approximation of a function, or a differential operator, the nodes in expressions like (4.1) must be in the vicinity of the point $\boldsymbol{x}$, and therefore numerous stencils are constructed throughout the underlying set. It is well-known that the matrix

$$\boldsymbol{A} = \begin{bmatrix} \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_K\|) \\ \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_K\|) \\ \vdots & \vdots & & \vdots \\ \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_K\|) \end{bmatrix} \tag{4.2}$$

is positive definite if the nodes $\boldsymbol{x}_1 \ldots \boldsymbol{x}_K$ are all distinct [95], and so under this assumption there exist K-point RBF interpolants for any function data. A different question, however, is whether the matrix $\boldsymbol{A}$ will be well-conditioned: it is not the case, for example, when the nodes are placed on a lattice and $\varepsilon \to 0$, [53]. The other extreme, having low regularity, also does not provide a reliable source of nodes, as can be seen on the example of the Halton sequence [53]. Furthermore, node clumping can lead to instability of PDE solvers, [49]. To avoid this, one must guarantee that the nodes are separated. In effect, generally the quasi-uniform node sets generated by the present algorithm, or, for example, the one constructed by the third and fourth authors [52], perform better than either lattice or the Halton sequence.

In many applications, one has to ensure that the distance from a node $\boldsymbol{x}$ to its nearest neighbor behaves approximately as a function of the position of the node [52]. Prescribing this function, $\rho(\boldsymbol{x})(\boldsymbol{x})$, which we call the *radial density*, is a natural way to treat the cases when a local refinement is required in order to capture special features of the domain. In the present chapter we will describe a method of node placement for which the actual distance to the nearest neighbor, denoted by $\Delta(\boldsymbol{x}) = \min_{\boldsymbol{x}' \neq \boldsymbol{x}} \|\boldsymbol{x}' - \boldsymbol{x}\|$, satisfies the above description. To summarize, we are interested in a procedure for obtaining discrete configurations inside a compact set that

will:

- guarantee that $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$ (that is, differ only up to a constant factor) for a given function $\rho(\boldsymbol{x})$ with a reasonably wide choice of ρ;
- be suitable for mesh-free PDE discretizations using RBFs, i.e., produce well-separated configurations without significant node alignment;
- result in quasi-uniform node distributions also on the surface boundaries of the domain;
- be computationally efficient, easily scalable, and suitable for parallelization.

4.2 Choice of method

4.2.1 RBF-FD approximations

In this section we shall outline the common practices involving RBFs, in order to motivate the requirements that have to be imposed on the node distribution used in the respective computations. For a more in-depth discussion see one of [51, 105, 28, 44]. A significant portion of the RBF approach hinges on the theory of positive definite functions.

Suppose we need to approximate a linear operator $\mathfrak{L}$ acting on sufficiently smooth functions supported on Ω, given locally by their values at the nodes $\boldsymbol{x}_k$, $k = 1, \dots, K$. More specifically, we need to compute the value $\mathfrak{L}\xi(\boldsymbol{x}_0)$ for some fixed $\boldsymbol{x}_0 \in \Omega$ and a variable function ξ. A generalization of the standard [50] finite-difference (FD) approach consists in constructing weights w_k, $k = 1, \dots, K$, that recover the value of $\mathfrak{L}$ at $\boldsymbol{x}_0$ in the form

$$\mathfrak{L}\mathcal{S}(\boldsymbol{x}_0) = \sum_{k=1}^{K} w_k \mathcal{S}(\boldsymbol{x}_k), \tag{4.3}$$

for every $\mathcal{S}$ from a convenient functional space; the $\mathcal{S}$ is then chosen to interpolate ξ at the given nodes $\boldsymbol{x}_1, \dots, \boldsymbol{x}_K$. In our case, $\mathcal{S}$ is spanned by RBFs, so by analogy to the 1-dimensional case this method is called RBF-FD; there exists extensive literature covering different types of kernels and different applications [51, 48, 54, 11, 47, 46, 4, 45]. Note that the node localization is required due to that (i) local stencils lead to sparse matrices, and are thus much more suitable for computations, (ii) in most applications, $\mathfrak{L}$ is either an interpolation or a differential operator; both act locally, so it is natural to use local stencils.

For example, using the space of shifts of the Gaussian kernel $\varphi(r) = e^{-(\varepsilon r)^2}$, one arrives at an RBF interpolant

$$\mathcal{S}(\boldsymbol{x}) = \sum_{k=1}^{K} c_k \varphi(\|\boldsymbol{x} - \boldsymbol{x}_k\|).$$

In order to express $\mathfrak{L}\mathcal{S}(\boldsymbol{x}_0)$ as a functional of $\mathcal{S}(\boldsymbol{x}_k), k = 1, \ldots, K$, as in (4.3), it suffices to do so for the functions $\varphi_k(\boldsymbol{x}) = \varphi(\|\boldsymbol{x} - \boldsymbol{x}_k\|)$, $k = 1, \ldots, K$. The weights $\{w_k\}$ are then obtained as the solution to

$$\begin{bmatrix} \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_1 - \boldsymbol{x}_K\|) \\ \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_2 - \boldsymbol{x}_K\|) \\ \vdots & \vdots & & \vdots \\ \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_1\|) & \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_2\|) & \ldots & \varphi(\|\boldsymbol{x}_K - \boldsymbol{x}_K\|) \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_K \end{bmatrix} = \begin{bmatrix} \mathfrak{L}\varphi(\|\boldsymbol{x}_0 - \boldsymbol{x}_1\|) \\ \mathfrak{L}\varphi(\|\boldsymbol{x}_0 - \boldsymbol{x}_2\|) \\ \vdots \\ \mathfrak{L}\varphi(\|\boldsymbol{x}_0 - \boldsymbol{x}_K\|) \end{bmatrix}.$$

Observe that in order to find the interpolant $\mathcal{S}$ for ξ, $\{c_k\}$ are determined from the same system, with $\mathfrak{L}$ taken to be the identity map. The matrix on the left is denoted by $\boldsymbol{A}$ in (4.2); it is degenerate whenever any two of $\{\boldsymbol{x}_k\}$ coincide, and is ill-conditioned whenever any two are very close, due to the continuity of φ. This brings us to further considerations of how the stencil $\{\boldsymbol{x}_k\}$ can be chosen. By the above, it is necessary that the nodes be (i) distinct and well-separated, and (ii) localized inside the domain Ω. For a quasi-uniform node set, K nearest neighbors of a fixed node satisfy both conditions.

Observe that for any strictly positive definite kernel φ, provided $\{\boldsymbol{x}_k\}$ are all distinct, the interpolation matrix $\boldsymbol{A}$ is always invertible (the Gaussian is an example of such kernel). To summarize, the above expression for weights $\{w_k\}$ is the defining property of the RBF-FD methods with the Gaussian kernel.

Taking the limit of the shape parameter $\varepsilon \to 0$ can cause the interpolant s to diverge for other RBF kernels [56, 29] that contain ε; this phenomenon however does not occur for the Gaussian $\varphi(r) = e^{-(\varepsilon r)^2}$. The motivation for considering the "increasingly flat" limit $\varepsilon \to 0$ is that the resulting RBFs can be used to obtain highly accurate solutions of elliptic problems and approximants of smooth data [56, 75]. We now conclude the discussion of the Gaussian kernel and turn to its novel alternative.

In the recent years, there have been noteworthy advances in RBF-FD using Polyharmonic Spline (PHS) kernels, $\varphi(r) = r^{2p-1}$ or $\varphi(r) = r^{2p} \log r$, $p \in \mathbb{N}$; it has been shown [46, 4, 45] that using PHS-based RBF-FD leads to improved accuracy, stability, and eliminates the Runge phenomenon at the boundary of the domain [4], which is not the case in general [57]. Another benefit from using the PHS kernel is that it does not contain the shape parameter ε. The analytical property underlying existence of the weights $\{w_k\}$ for the PHS kernels is that the PHS are conditionally strictly positive definite [44, 82] and thus need a slightly different treatment, which we shall now outline.

To ensure unisolvency (uniqueness of the weights and interpolants), we need to augment the

$\mathcal{S}$ with polynomial terms: it is selected from the space defined by

$$\mathcal{S}(\boldsymbol{x}) = \sum_{k=1}^{K} c_k \varphi(\|\boldsymbol{x} - \boldsymbol{x}_k\|) + \sum_{i=1}^{\binom{l+d}{l}} b_i \pi_i(\boldsymbol{x}),$$

with $\{c_k\}$ satisfying the constraint

$$\sum_{k=1}^{K} c_k \pi_i(\boldsymbol{x}_k) = 0, \quad i = 1, 2, \ldots, \binom{l+d}{l},$$

where φ is now the PHS kernel, and $\binom{l+d}{l}$ is the dimension of the space of multivariate polynomials of degree up to l in $\mathbb{R}^d$; accordingly, π_i varies over the monomial basis for such polynomials.

The degree l has to satisfy $l = p - 1$ and $l = p$ for $\varphi(r) = r^{2p-1}$ and $\varphi(r) = r^{2p} \log r$, respectively [44, Chapter 8]. For example, when $\varphi(r) = r^2 \log r$ and $d = 2$, then $l = 1$ and the weights corresponding to an operator $\mathfrak{L}$ in $\mathbb{R}^2$ are determined by

$$\begin{bmatrix} & & & | & 1 & x_{11} & x_{12} \\ & \boldsymbol{A} & & | & \vdots & \vdots & \vdots \\ & & & | & 1 & x_{k1} & x_{k2} \\ - & - & - & + & - & - & - \\ 1 & 1 & 1 & | & & & \\ x_{11} & \ldots & x_{k1} & | & & 0 & \\ x_{12} & \ldots & x_{k2} & | & & & \end{bmatrix} \begin{bmatrix} w_1 \\ \vdots \\ w_k \\ - \\ w_{k+1} \\ w_{k+2} \\ w_{k+3} \end{bmatrix} = \begin{bmatrix} \mathfrak{L}\varphi(\|\boldsymbol{x}_0 - \boldsymbol{x}_1\|) \\ \vdots \\ \mathfrak{L}\varphi(\|\boldsymbol{x}_0 - \boldsymbol{x}_K\|) \\ - \\ \mathfrak{L}\,1 \\ \mathfrak{L}x_1(x_{01}) \\ \mathfrak{L}x_2(x_{02}) \end{bmatrix}, \tag{4.4}$$

where the matrix $\boldsymbol{A}$ is the same as in (4.2) with the PHS kernel; x_{kj} is the j-th coordinate of $\boldsymbol{x}_k$, $k = 1, \ldots, K$; $j = 1, 2$, and similarly for $\boldsymbol{x}_0 = (x_{01}, x_{02})^{\mathrm{tr}}$; $\mathfrak{L}1$, $\mathfrak{L}x_1$, $\mathfrak{L}x_2$ denote images of the constant and coordinate functions under $\mathfrak{L}$, respectively. Here, as before, the interpolation case is obtained by taking $\mathfrak{L}$ equal to the identity operator; compare the constraints on $\{c_k\}$ above. The generalization to larger values of l and higher dimensions follows along the same lines, with higher-degree monomials instead of linear terms [105, Chapters 8, 11.1].

For this and the other commonly used kernels, non-degeneracy follows from a strengthened form of Micchelli's theorem [88], see also [44, 105, 34, 82]; namely, the matrix in the LHS of the previous equation is non-degenerate for any unisolvent $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_K$. The remaining part of the discussion for the Gaussian kernel above is further applicable without any modifications. It should be noted that the optimal choice of the degree of PHS that needs to be used, does in general depend on the particular problem under consideration.

4.2.2 Riesz energy

To generate nodes both devoid of lattice alignment and having near-optimal local separation, we shall minimize a functional on the space of discrete subsets of Ω. Equivalently, one can think of the corresponding gradient flow moving the starting configuration to a suitable position. The desired properties of the minimizing configuration will then follow from the strong repulsion imposed by the functional.

Recall that for a fixed $s > d$ the *Riesz s-energy* is a functional $E_s : \Omega^N \mapsto (0, \infty)$ such that for a collection of vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N$ in Ω,

$$E_s(\boldsymbol{x}_1, \ldots \boldsymbol{x}_N) = \sum_{i \neq j} \frac{1}{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^s}. \tag{4.5}$$

As discussed in Chapter 1, there exists extensive literature dedicated to the collections minimizing this and derived functionals for $s \geq d$, their asymptotics and limiting measures, see for example [62, 15, 22]. Recall that in the case $s \geq d$, the distribution of minimizers of E_s coincides with the normalized Hausdorff measure on Ω; practically this means that the minimizers are uniform in the volumetric sense on Ω, that is, the number of nodes per unit volume is close to constant. In order to produce non-uniform nodes, we shall further add multiplicative weight to (4.5); this modification of (4.5) was first studied in [15]. The *weighted Riesz s-energy* with kernel $\kappa : \Omega \times \Omega \to (0, \infty)$ is the functional $E_s^\kappa : \Omega^N \mapsto (0, \infty)$, defined in the notation of (1.4) as

$$E_s^\kappa(\boldsymbol{x}_1, \ldots \boldsymbol{x}_N) := E(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N; g_s, \kappa, 0) = \sum_{i \neq j} \frac{\kappa(\boldsymbol{x}_i, \boldsymbol{x}_j)}{\|\boldsymbol{x}_i - \boldsymbol{x}_j\|^s}.$$

It has been shown [15], that the counting measures of minimizers of the weighted energy converge to the probability measure with volumetric density proportional to

$$\kappa(\boldsymbol{x}, \boldsymbol{x})^{-d/s} d\mathcal{H}_d(\boldsymbol{x}),$$

with $\mathcal{H}_d$ denoting the d-dimensional Hausdorff measure. More precisely, for any $\mathcal{B} \subset \Omega$ with boundary of zero measure there holds

$$\frac{1}{N} \sum_{i=1}^{N} \chi(\boldsymbol{x}_i; \mathcal{B}) \longrightarrow \frac{1}{Z(\kappa, \Omega)} \int_{\mathcal{B}} \kappa^{-d/s}(\boldsymbol{x}, \boldsymbol{x}) d\mathcal{H}_d(\boldsymbol{x}), \quad N \to \infty,$$

where $\chi(\cdot\,; \mathcal{B})$ is the indicator function of $\mathcal{B}$ and $Z(\kappa, \Omega)$ is the normalization constant. Of course, for most applications the set $\mathcal{B}$ will have zero-measure or even differentiable boundary.

It is worth noting that the previous equation shows that the distribution of nodes produced by minimizing E_s^κ depends only on the diagonal values of κ for large enough N. Indeed, this has

been explored in [18], where it is shown that omitting interactions of points at least r_N apart in E_s^κ and minimizing the resulting expression leads to the same distribution when $N \to \infty$; the sequence r_N, $N \geq 2$ here satisfies $r_N N^{1/d} \to \infty$, $N \to \infty$. Following [18], a weight $\kappa(\boldsymbol{x}, \boldsymbol{y})$ vanishing whenever $\|\boldsymbol{x} - \boldsymbol{y}\| > r_N$ is said to be *truncated*.

Configurations that minimize E_s^κ over Ω^N for a compact Ω are well-separated, that is, the quantity $\Delta(\{\boldsymbol{x}_1, \dots \boldsymbol{x}_N\}) = \min_i \Delta(\boldsymbol{x}_i)$ satisfies

$$\Delta(\{\boldsymbol{x}_1, \dots \boldsymbol{x}_N\}) \geq C N^{-1/d}, \quad N \geq 2. \tag{4.6}$$

We shall assume that the terms in E_s^κ for which $\boldsymbol{x}_j$ is not among the K nearest nodes to $\boldsymbol{x}_i$ are zero, a condition equivalent to truncating κ, provided the nodes are well-separated. Under this assumption, the expression for E_s^κ can be rewritten as

$$E_s^\kappa(\boldsymbol{x}_1, \dots \boldsymbol{x}_N) = \sum_{i=1}^{N} \sum_{k=1}^{K} \frac{\kappa\big(\boldsymbol{x}_i, \boldsymbol{x}_{j(i,k)}\big)}{\|\boldsymbol{x}_i - \boldsymbol{x}_{j(i,k)}\|^s}, \tag{4.7}$$

where nodes $\boldsymbol{x}_{j(i,k)}$, $k = 1, \dots, K$, are the K nearest neighbors of $\boldsymbol{x}$. That minimizers of (4.7) are well-separated can be shown by the standard argument from one of [62, 15, 18]. This further implies that they are quasi-uniform, which is the key property for us in view of the discussion in Section 4.2.1. As the form (4.7) makes clear, for the truncated kernel κ the E_s^κ can be computed in $O(NK)$ operations, unlike the $O(N^2)$ operations required to evaluate the functional for a non-vanishing κ. This, and the fact that (4.7) requires constant size memory for storage makes this form of E_s^κ the most useful for our purposes.

The value of the exponent s is chosen so that $s \geq d$ to ensure that the energy functional is sufficiently repulsive; it is known from the classical potential theory that for $s < d$ the minimal energy configurations are not necessarily uniform, and their local structure depends on the shape of the domain [74]. Property (4.6) holds for any $s > d$, when minimizing the energy over any fixed compact set $\Omega \subset \mathbb{R}^d$. While setting $\kappa(\boldsymbol{x}, \boldsymbol{y}) = f(\|\boldsymbol{x} - \boldsymbol{y}\|)$ for an arbitrary positive definite radial function $f(r)$ that grows fast enough when $r \to 0$ would produce similar results, we chose the Riesz energy because the properties of its minimizers are well understood.

Note that simply looking for minimizers of E_s^κ does not provide node sets satisfying $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$ for every $\boldsymbol{x}$; in fact, boundary nodes of such minimizers will often have smaller separations than desired. Since in such cases the boundary has a lower Hausdorff dimension, it does not influence the volumetric density, which agrees with the results above. With this motivation in mind, we are ready to introduce the second component of our method.

4.2.3 Quasi-Monte Carlo methods

To facilitate convergence of whichever optimization algorithm is used to find minimizers of (4.7), we can initialize it with a configuration that approximates the limiting measure. One has to rule out Monte Carlo methods due to the separation requirement: random points exhibit clustering [25], which makes deterministic post-processing, in particular by energy minimization, costly. Similarly, mitigating the clustering by purely probabilistic approaches, as for example *thinning* discussed in [77], or by sampling from a random process with repulsive properties [2, 6], does not generally yield satisfying results, since the separation can only be guaranteed on average. Instead, we turn to the quasi-Monte Carlo (Q-MC) approach. As has been pointed out at the end of Section 4.2.1, in order to ensure convergence of RBF-FD interpolants, the underlying node set must be (locally) unisolvent; for our purposes this just means that the nodes are in a generic position with respect to each other. The latter is clearly not the case for lattice-like Q-MC configurations, which explains why we resort to energy minimization. On the other hand, we choose not to use other popular Q-MC sequences, such as Halton nodes, as they do not necessarily lead to the best conditioned systems, see [55, Figure 5.1] and Figure 4.10 below, and are harder to handle when the distribution support Ω has complex geometry.

The key element of our construction lies in distributing the node set in a deterministic way so that to guarantee low discrepancy between the desired and the obtained radial densities. This is achieved by a Q-MC analog of the *stratification* of the Monte Carlo method [31]: nodes are distributed with piecewise constant (radial) density that approximates the desired one. We consider two different Q-MC sequences to draw from with (near-)constant radial density: irrational lattices and periodic Riesz minimizers. After dividing the set Ω into cube-shaped *voxels* $\mathcal{V}_m,\ m = 1, \dots, M^d$, each voxel is filled with nodes obtained in one of the two ways, appropriately scaled, then the weighted s-energy (4.7) of the whole node set is minimized. Although we discuss the radial density in the present chapter, an argument for the volumetric density can be produced along the same lines.

Yet another reason to make use of a Q-MC sequence is to avoid recursive data structures, which in some cases can be detrimental to the overall performance. Even though such structures have seen significant developments over the years, both dynamic update and parallelization for them remain challenging, [90, 107]. The present approach should therefore be understood as almost opposite to the well-known "quadtree" algorithm [60], that indeed has been used for meshless node generation [101]. Namely, as outlined above, our algorithm places nodes en masse inside the voxels to produce a rough approximation of the target distribution, and subsequently adjusts them by a gradient flow, which is straightforward to parallelize. Although this does involve the computation of the nearest neighbors in (4.7), which in practice will be done by constructing a k-d tree, by initializing the node configuration with a stratified Q-MC sequence we ensure the indices $j(i, k)$ in (4.7) will not undergo significant changes during the energy minimization stage,

so the k-d tree will not require intensive updates.

An *irrational lattice*(IL) is defined as a discrete subset of the d-dimensional unit cube $[0,1]^d$

$$\mathcal{L}_n = \left\{ \Big(\{C_1 + i/n\}, \{i\alpha_1\}, \{i\alpha_2\}, \ldots, \{i\alpha_{d-1}\} \Big) \ : \ i = 1, \ldots, n \right\}, \tag{4.8}$$

where $\{x\} = \mathrm{mod}(x,1) = x - \lfloor x \rfloor$ denotes the fractional part of x, $C_1 > 0$ is fixed, and $\alpha_1, \alpha_2, \ldots, \alpha_{d-1}$ are irrational numbers, linearly independent over the rationals. This terminology seems to be accepted in the low-discrepancy community [9], while closely related objects, when used for Q-MC purposes, are known as *Korobov/lattice point sets* [76].

The motivation for using an IL in this context is due to the existing results on the discrepancy of ILs. It is known for example, that the two-dimensional ILs have the optimal order of L^2 discrepancy, [10, 9]. Furthermore, in all dimensions ILs are uniformly distributed [73, Chapter 1.6], that is, the fraction of lattice points inside any rectangular box with faces parallel to the coordinate planes converges to its volume. The simple linear structure of ILs makes them especially attractive for SIMD-parallelization.

Another Q-MC sequence that has proven to suit our purposes consists of *periodic Riesz minimizers* on the unit flat torus, that is, n-point collections $\mathcal{M}_n,\ n \geq 1 = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ that minimize (4.5) on $([0,1]^d)^n$ with the Euclidean distance $\|\cdot\|$ replaced by the periodic metric

$$\|\boldsymbol{x} - \boldsymbol{y}\|_{\sim}^2 = \Pi(x_1 - y_1) + \Pi(x_2 - y_2) + \Pi(x_3 - y_3), \tag{4.9}$$

where $\Pi(x) = \min(x^2, (1-x)^2),\ 0 \leq x \leq 1$. It follows from [62] that such configurations have optimal separation and asymptotically uniform volumetric density. It follows from the numerics also, that in this case the nearest neighbor distances vary very little from node to node; this and that minimizing configurations do not suffer from the lattice-like alignment, makes their rescaled copies good candidates for the stratification.

The number of nodes in individual voxels is defined by the function ρ, so the resulting collection has piecewise constant density; refining the voxel partition leads to an improved piecewise approximation of the desired (e.g., smooth) density. In practice, the dependence between the number of nodes contained in the unit cube, and average/minimal nearest neighbor distance is tabulated in advance, and then inverted during the construction of the node set.

4.3 The algorithm

The interested reader will find a Matlab codebase implementing the algorithm described here, as well as the sources for all the figures contained in the present chapter, at [102].

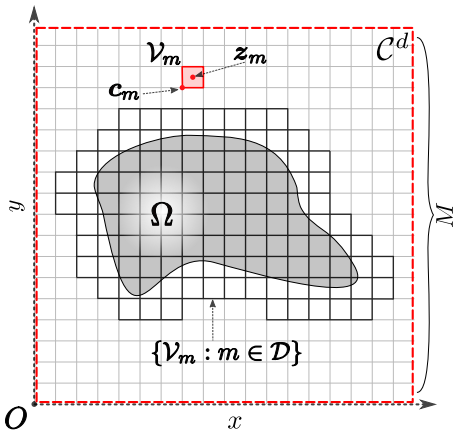

Figure 4.1: An illustration of some of the symbolic notation used in the algorithm below, in the case $d = 2$.

4.3.1 Formulation

If the nodes must be restricted to a certain compact set Ω, for example, support of a given indicator, we will refer to the set as *density support*, and to the indicator function as *point inclusion function*. We may further assume that Ω is contained in the d-dimensional unit cube $\mathcal{C}^d = [0,1]^d$ (see Figure 4.1 for some of the notation involved); the case of an arbitrary compact set then follows by choosing a suitable enclosing cube and applying scaling and translation. Suppose the radial density is prescribed by a Lipschitz-1 function; i.e., $|\rho(\boldsymbol{x})(\boldsymbol{x}) - \rho(\boldsymbol{y})| \leq \|\boldsymbol{x} - \boldsymbol{y}\|$. The reason for this assumption is the respective property of $\Delta(\boldsymbol{x})$, and is explained in further detail in the following section. Recall that we use an exponent $s > d$. We summarize the discussion in Section 4.2 into the following algorithm for generating nodes with radial density ρ:

Initial node layout.

`Step 0` Choose one of the two Q-MC sequences described in Section 4.2, $\{\mathcal{L}_n : n \geq 1\}$ or $\{\mathcal{M}_n : n \geq 1\}$, draw configurations with up to $n_{\max}$ nodes from it, and determine the mean nearest neighbor distance for its periodization by the integer lattice, denoted by $\bar{\Delta}_n$ for n nodes. Let $\lambda : (0,\infty) \to \{0,1,2,\ldots,n_{\max}\}$ be the interpolated inverse to $\langle \Delta_n \rangle : \{1,2,\ldots,n_{\max}\} \to (0,1]$.[1]

`Step 1` Partition $\mathcal{C}^d$ into M^d equal cube-shaped voxels $\mathcal{V}_m$, $m = 1,\ldots,M^d$ of side length $1/M$, with faces parallel to the coordinate planes. Let $\{\mathcal{V}_m : m \in \mathcal{D}\}$ be the subset for which at least one of the adjacent (i.e., sharing a face) voxels has a vertex inside Ω.

`Step 2` Let $\langle \rho_m \rangle$ be the average value of ρ at the 2^d vertices of a voxel $\mathcal{V}_m$, $m \in \mathcal{D}$. Place inside

[1] Note that both ILs and the minimizers can have the nearest neighbor distance of at most 1, due to periodicity. We therefore take $\lambda(x) = 0$ whenever $x > 1$.

$\mathcal{V}_m$ a scaled and translated version of the n_m-point IL (4.8), or of the n_m-point periodic Riesz minimizer, using n_m defined by

$$n_m = \lambda(\langle \rho_m \rangle M).$$

Repeat for each $m \in \mathcal{D}$.

Saturation and cleanup.

`Step 3` Let $\mathcal{E} \subset \mathcal{D}$ be the set of m for which voxels $\{\mathcal{V}_m : m \in \mathcal{E}\}$ contain no nodes and the centers $\{\boldsymbol{z}_m : m \in \mathcal{E}\}$ satisfy $\Delta(\boldsymbol{z}_m) > \rho(\boldsymbol{z}_m)$. Sort $\mathcal{E}$ by the increasing values of $\rho(\boldsymbol{z}_m)$. Repeat until $\mathcal{E}$ is empty: for every $m \in \mathcal{E}$ place a node in $\boldsymbol{z}_m$; recompute $\mathcal{E}$.

`Step 4` For all nonempty voxels, remove nodes outside Ω.

Repel-type iterations, boundary detection.

`Step 5` Perform T iterations of the partial gradient descent on the weighted s-energy functional (4.7) with $\kappa(\boldsymbol{x}, \boldsymbol{x}) = c\rho(\boldsymbol{x})^s$, using the K nearest neighbors of each node: Let the initial configuration be the 0-th iteration, $\boldsymbol{x}_i^{(0)} = \boldsymbol{x}_i$, $i = 1, \ldots, N$, with N denoting the total number of nodes distributed. On the t^{th} iteration, $1 \le t \le T$, given a node $\boldsymbol{x}_i^{(t)}$ with K nearest neighbors $\boldsymbol{x}_{j(i,k)}^{(t)}$, $k = 1, \ldots, K$, form the weighted vector sum

$$\boldsymbol{g}_i^{(t)} = s\,\rho\left(\boldsymbol{x}_i^{(t)}\right)^s \sum_{k=1}^{K} \frac{\boldsymbol{x}_i^{(t)} - \boldsymbol{x}_{j(i,k)}^{(t)}}{\|\boldsymbol{x}_i^{(t)} - \boldsymbol{x}_{j(i,k)}^{(t)}\|^{s+2}}, \quad 1 \le i \le N,$$

and the new node position can now be expressed as

$$\boldsymbol{x}_i^{(t+1)} = \begin{cases} \boldsymbol{x}_i^{(t)} + \dfrac{\Delta(\boldsymbol{x}_i^{(t)})}{t + C_2} \dfrac{\boldsymbol{g}_i^{(t)}}{\|\boldsymbol{g}_i^{(t)}\|} & \text{if this sum is inside } \Omega; \\ \boldsymbol{x}_i^{(t)}, & \text{otherwise,} \end{cases} \qquad 1 \le i \le N. \tag{4.10}$$

where C_2 is a fixed offset chosen to control the step size between $\boldsymbol{x}_i^{(t)}$ and $\boldsymbol{x}_i^{(t+1)}$. If a "pullback" function is provided from a neighborhood of Ω to its boundary, the condition of $\boldsymbol{x}_i^{(t+1)}$ being inside Ω is replaced with applying the pullback; furthermore, if the radial density has an easily computable gradient, or is changing rapidly, an additional term must be included in (4.10) (see discussion below).

Update the neighbor indices $j(i,k)$ after every few iterations.

`Step 6` If no boundary node set/pullback function is prescribed, define the boundary nodes as follows. Evaluate the point inclusion function for $\boldsymbol{x}_i \pm \Delta(\boldsymbol{x}_i)\boldsymbol{e}_l$, $l = 1, \ldots, d$, $i = 1, \ldots, N$, where $\boldsymbol{e}_l$ is the l-th basis vector. If at least one such point lies outside Ω, the $\boldsymbol{x}_i$ is considered to be a boundary node.

4.3.2 Discussion

Our assumption of ρ being Lipschitz-1 is natural, since $\Delta(\cdot)$ is always Lipschitz-1, if viewed as a function of position. To see this, consider any two nodes $\boldsymbol{x}, \boldsymbol{y}$, and let $\boldsymbol{x'}, \boldsymbol{y'}$ be their nearest neighbors, respectively, so that $\|\boldsymbol{x} - \boldsymbol{x'}\| = \Delta(\boldsymbol{x})$ and $\|\boldsymbol{y} - \boldsymbol{y'}\| = \Delta(\boldsymbol{y})$. It follows, $\Delta(\boldsymbol{x}) \leq \|\boldsymbol{x} - \boldsymbol{y'}\| \leq \|\boldsymbol{x} - \boldsymbol{y}\| + \Delta(\boldsymbol{y})$, which by symmetry implies Δ is Lipschitz-1.

Initial node layout.

In the parts of the density support with constant ρ, the nodes will locally look like a periodization of the initial Q-MC sequence; hence the average neighbor distance in `Step 0` is tabulated for the periodized version. Observe that there is some freedom in the notion of voxel adjacency used to define $\{\mathcal{V}_m : m \in \mathcal{D}\}$ in `Step 1`; for example, in the case of a highly non-convex domain Ω, it might be reasonable to denote the $3^d - 1$ voxels sharing a vertex with a given $\mathcal{V}_m$ as adjacent to it, rather than only the $2d$ voxels that it has a common face with. This would then ensure that no part of Ω will be omitted in the node allocation; imagine a long and thin peninsula in Ω containing no corners of $\mathcal{V}_m$, $m = 1, \ldots, M^d$. We have found however, that the subsequent repel iterations will guarantee that such a peninsula is adequately filled with nodes even when using only the face-adjacent voxels.

If the IL sequence is chosen in `Step 0`, the n_m-node set placed in voxel $\mathcal{V}_m$ at `Step 2` is an adjusted version of (4.8) as follows. Let for every $\mathcal{V}_m$ the corner with the smallest absolute value be $\boldsymbol{c}_m$; the points $\boldsymbol{c}_m$ are then vertices of a lattice. Before scaling and translating $\mathcal{L}_{n_m}$, apply a random permutation to the coordinates of each node in it, so that to remove long-range lattice structure from the distribution; we will denote such an operation by π_c. Then the IL in voxel $\mathcal{V}_m$ becomes

$$\mathcal{L}'_{n_m} = \boldsymbol{c}_m + \frac{f}{M}\pi_c(\mathcal{L}_{n_m}) + \frac{\boldsymbol{h}}{M}, \tag{4.11}$$

where

$$f = 1 - c_d\,(n_{\max})^{-1/d}, \qquad \boldsymbol{h} = \frac{1-f}{2} \cdot (1, 1, \ldots, 1)^{\mathrm{tr}},$$

with c_d depending only on the dimension. The quantities f and $\boldsymbol{h}$ ensure that the lattice points in $\mathcal{L}_m$ are inset into the voxel by about half the separation distance, avoiding poorly separated points along the voxel interfaces.

When the periodic minimizer sequence is selected in `Step 0`, the inset is defined similarly, but the permutation is just an identity, $\pi_c \equiv \mathrm{id}$, as the minimizers don't have the lattice structure. Likewise, the scaling factor and translation are

$$f_m = 1 - c_d\,(n_m)^{-1/d}, \qquad \boldsymbol{h}_m = \frac{1-f_m}{2} \cdot (1, 1, \ldots, 1)^{\mathrm{tr}}.$$

The analog of (4.11) thus takes the form

$$\mathcal{M}'_{n_m} = \boldsymbol{c}_m + \frac{f_m}{M}\mathcal{M}_{n_m} + \frac{\boldsymbol{h}_m}{M}. \tag{4.12}$$

As one would expect, the average separation for the sequence $\{\mathcal{M}_n\}$ is larger than that of $\{\mathcal{L}_n\}$ for the respective values of n. While the inset for the latter is necessary to account for the node proximity after periodization, for the former it serves to mitigate the effects of interfacing voxels containing different number of nodes. This is further discussed in the Appendix.

Saturation and cleanup.

Observe that after `Step 2`, voxels in $\{\mathcal{V}_m : m \in \mathcal{D}\}$ satisfying $\langle\rho_m\rangle M > 1$ do not contain any nodes. The goal of `Step 3` is therefore to remove any redundant sparsity that may be present whenever the radial density ρ is larger than $1/M$, as in this case the function λ in `Step 2` is set to zero. More careful geometric considerations would lead one to set $\lambda(x) > 0$ when $0 < x < \sqrt{d}/M$, the length of a voxel diagonal, and thus make λ dependent on the dimension; on the other hand, using the interval $0 < x < 1/M$ as we did appears to suffice due to correction of density in `Step 5`.

Note that in practice, when recomputing $\mathcal{E}$ in `Step 3`, to verify $\Delta(\boldsymbol{z}_{m_0}) > \rho(\boldsymbol{z}_{m_0})$ for a fixed $m_0 \in \mathcal{E}$ it is enough to check $\|\boldsymbol{z}_{m_0} - \boldsymbol{z}_m\| > \rho(\boldsymbol{z}_{m_0})$ for the previously selected $\boldsymbol{z}_m$ with $\rho(\boldsymbol{z}_m) < \rho(\boldsymbol{z}_{m_0})$. Indeed, let $\boldsymbol{z}_{m_0}$ be the center of $\mathcal{V}_{m_0}$. Then, by the definition of λ in `Step 2`, the radial density $\rho(\boldsymbol{z}_{m_0}) = (1 + D)/M$ for some $D > 0$, so the Lipschitz-1 property implies, for any $\boldsymbol{x}$ such that $\rho(\boldsymbol{x}) \le 1/M$ there holds $\|\boldsymbol{z}_{m_0} - \boldsymbol{x}\| \ge |\rho(\boldsymbol{z}_{m_0}) - \rho(\boldsymbol{x})| \ge D/M$. This ensures that distances from $\boldsymbol{z}_{m_0}$ to the nodes produced on `Step 2` satisfy

$$\|\boldsymbol{z}_{m_0} - \boldsymbol{x}\| \ge \frac{D}{1+D}\rho(\boldsymbol{z}_{m_0}).$$

This shows, when $D \gg 1$, not checking the inequality $\|\boldsymbol{z}_{m_0} - \boldsymbol{x}\| > \rho(\boldsymbol{z}_{m_0})$ leads to at most a bounded factor error. On the other hand, for $D \ll 1$ distances from $\boldsymbol{z}_{m_0}$ to the nodes from `Step 2` are also controlled: it follows from (4.11)–(4.12) that for n_m small, nodes in the voxel $\mathcal{V}_m$ have larger inset (depending on c_d). This analysis is certainly not rigorous; however, applying the partial gradient descent in `Step 5`, we are able to ensure that the ratio ρ/Δ is close to 1, as desired.

Observe that in `Step 2` the nodes are only placed in $\mathcal{V}_m$'s for which either of the adjacent voxels has corners inside the density support Ω, so removing nodes outside Ω in `Step 4` does not lead to much overhead. Furthermore, since the density is evaluated at the corners only and not at individual nodes, the total number of evaluations may be significantly reduced, which is especially useful when ρ is computationally expensive. It is essential here that due to the Lipschitz-1

property, ρ is well estimated by its values at the corners $\boldsymbol{c}_m$; specifically, $|\rho(\boldsymbol{x}) - \rho(\boldsymbol{c}_m)| \leq \sqrt{d}/2M$ with $\boldsymbol{c}_m$ the nearest voxel corner to $\boldsymbol{x}$.

Repel-type iterations, boundary detection.

The equality $\kappa(\boldsymbol{x}, \boldsymbol{x}) = c\rho(\boldsymbol{x})^s$ can be justified by observing that each node $\boldsymbol{x}$ of the target distribution must be contained in a ball of radius $\rho(\boldsymbol{x})$, not containing any of the other nodes, hence, the volumetric density must be inverse proportional to $\rho(\boldsymbol{x})^d$. On the other hand, minimizers of (4.7) converge to the distribution with volumetric density $\kappa(\boldsymbol{x}, \boldsymbol{x})^{-d/s}$; hence $\kappa(\boldsymbol{x}, \boldsymbol{x})^{d/s} = c\rho(\boldsymbol{x})^d$.

The vector sum in `Step 5` is the partial $\boldsymbol{x}$-gradient of the weighted Riesz s-energy (4.7) in the sense that a single summand of (4.7) is $e(\boldsymbol{x}, \boldsymbol{y}) = \kappa(\boldsymbol{x}, \boldsymbol{y})\|\boldsymbol{x} - \boldsymbol{y}\|^{-s}$, and thus its complete $\boldsymbol{x}$-gradient is equal to

$$\nabla_{\boldsymbol{x}}\, e(\boldsymbol{x}, \boldsymbol{y}) = -s\,\kappa(\boldsymbol{x}, \boldsymbol{y})(\boldsymbol{x} - \boldsymbol{y})\|\boldsymbol{x} - \boldsymbol{y}\|^{-s-2} + \nabla_{\boldsymbol{x}}\kappa(\boldsymbol{x}, \boldsymbol{y})\|\boldsymbol{x} - \boldsymbol{y}\|^{-s}.$$

For our purposes, the $\boldsymbol{y}$ here is one of the K nearest nodes to $\boldsymbol{x}$, and, since due to the Q-MC initialization there will be few isolated nodes, and since off-diagonal values of $\kappa(\boldsymbol{x}, \boldsymbol{y})$ do not influence the limiting distribution (for details see [18]), we assume $\kappa(\boldsymbol{x}, \boldsymbol{y}) \approx \kappa(\boldsymbol{x}, \boldsymbol{x})$ to rewrite the previous equation as

$$\nabla_{\boldsymbol{x}}\, e(\boldsymbol{x}, \boldsymbol{y}) = -s\,\rho\,(\boldsymbol{x})^s\,(\boldsymbol{x} - \boldsymbol{y})\|\boldsymbol{x} - \boldsymbol{y}\|^{-s-2} + s\nabla_{\boldsymbol{x}}\rho(\boldsymbol{x})\,\rho(\boldsymbol{x})^{s-1}\|\boldsymbol{x} - \boldsymbol{y}\|^{-s}. \tag{4.13}$$

As has been pointed out at the beginning of Section 4.3.2, in order to be meaningful as a radial density, the function ρ must be Lipschitz-1. Then by the Rademacher's theorem, $\nabla_{\boldsymbol{x}}\rho$ exists almost everywhere; this validates the use of it in (4.13) as well as the approximation $\kappa(\boldsymbol{x}, \boldsymbol{y}) \approx \kappa(\boldsymbol{x}, \boldsymbol{x})$. The ratio of the second term to the first one in (4.13) is bounded by $\nabla_{\boldsymbol{x}}\rho(\boldsymbol{x})\|\boldsymbol{x} - \boldsymbol{y}\|/\rho(\boldsymbol{x})$ and, provided that the distances from $\boldsymbol{x}$ to its nearest neighbors are close to the value of $\rho(\boldsymbol{x})$, is at most $c\nabla_{\boldsymbol{x}}\rho(\boldsymbol{x})$ for a constant c. This condition is satisfied because the chosen Q-MC sequences have very regular local structure. In practice, the node distance is small on the scale of the support and varies slowly, so the second term will have negligible impact on the direction of the gradient after normalization; besides, precise gradient computation may prove costly. For these reasons we omit the second term in equation (4.10). If it is necessary to deal with a fast-changing radial density, a trade-off between the computational costs and the resulting distribution properties must be sought.

It doesn't matter which minimization method is applied to the weighted s-energy, rather the gradient descent is chosen due to its simplicity. Note, the second case in (4.10), leading to shrinking of the line stepping distance, can be thought of as a simplistic backtracking line search; it turns out to be sufficient for our purposes. Furthermore, applying a more involved line search

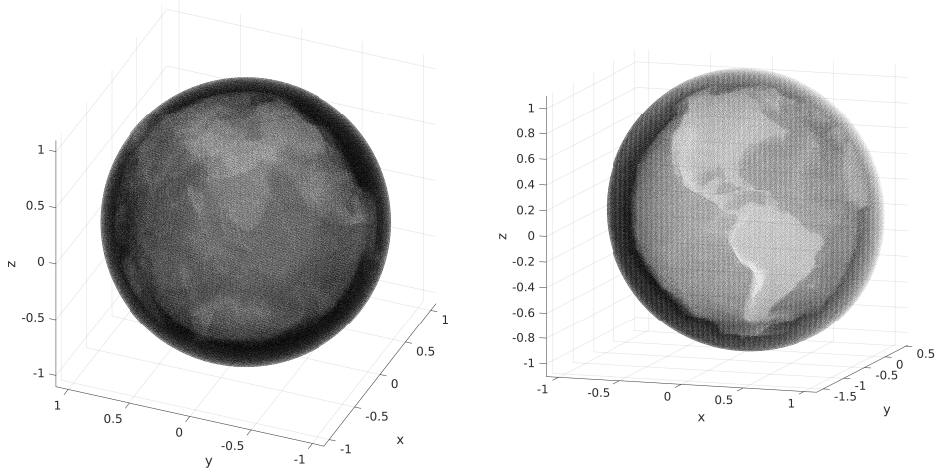

Figure 4.2: Left: a general view of a uniform node distribution in an atmospheric-like shell. Right: a separate view of the Western hemisphere.

may significantly degrade performance for complicated or nonsmooth domains.

The number of nearest neighbors K in (4.7), and the number of iterations T in `Step 5` can be adjusted to achieve a trade-off between execution speed/memory consumption/local separation. In our experiments,[2] even relatively small values of K and T produce good results: we used $K \approx T \approx 30$ for 1.36 million nodes with constant density in Section 4.4.1, and $K = 30$, $T = 200$ for 0.58 million and 0.36 million nodes with variable densities in Sections 4.4.2 and 4.4.3, respectively.

4.4 Sample applications

4.4.1 Atmospheric node distribution using surface data

We use the geodata [3] from the collection of global relief datasets produced by NOAA (National Oceanic and Atmospheric Administration), which contains a 1 arc-minute resolution model. We generate a sample configuration consisting of 1,356,566 nodes distributed uniformly inside an atmospheric-type shell Ω_{etopo}: the outer boundary of Ω_{etopo} is spherical, the inner one is an interpolation of the relief from ETOPO1 data, exaggerated by a factor of 100. The scale is chosen so that the average Earth radius, assumed to be 6,371,220 meters, has unit length; the radius of the outer boundary is set to 1.1, which corresponds to the height of 6,371 meters above the average radius, given the exaggeration factor.

The ETOPO1 dataset stores relief as a 21 600-by-10 800 array of elevations above the sea level; equivalently, of radial coordinates that correspond to the spherical angles defined by the array's

[2]The Matlab code we provide performs naive autotuning of K and T, using the total number of nodes to be placed. Although sufficient for demonstration purposes, there is room for improvement.

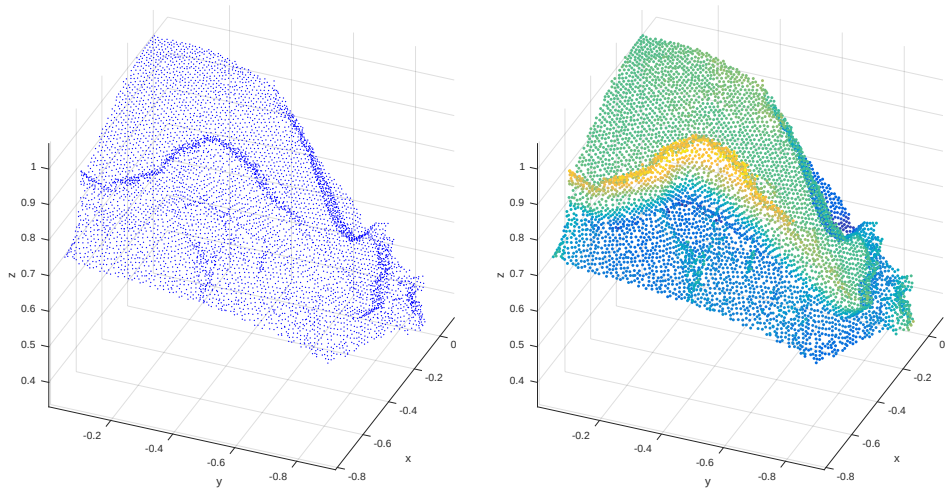

Figure 4.3: Surface subset: a fragment of the Western coast of South America. The nodes on the right are color-coded using heights.

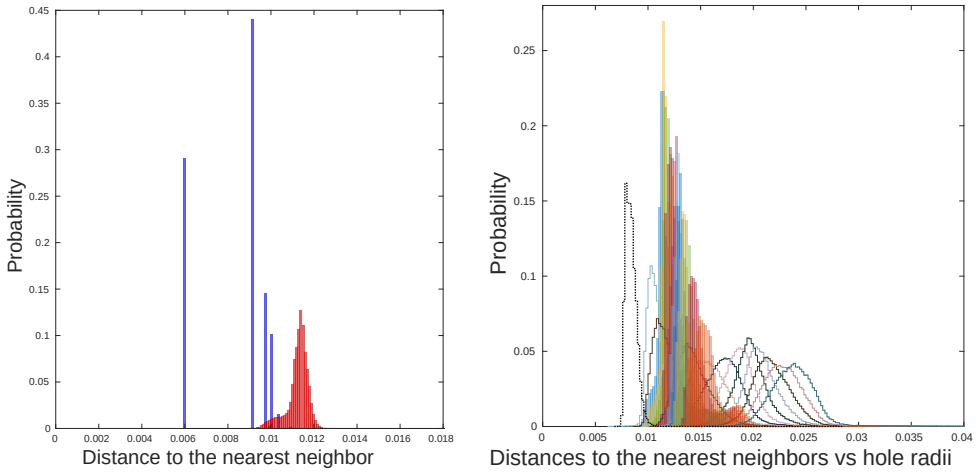

Figure 4.4: The effects of the repel procedure and hole radii. Left: probability distribution of the nearest-neighbor distances in the atmospheric node set, before (blue) and after (red) executing the repel subroutine. Right: distribution of distances to the 12 nearest neighbors for the whole configuration (color only), for the surface subset (contours), the hole radii (black dashed contour on the left).

indices. The data points are equispaced on lines of constant azimuth/inclination with angular distance $B = \pi/10\,800$ between them. To determine whether a given node $\boldsymbol{x} = (r_{\boldsymbol{x}}, a_{\boldsymbol{x}}, p_{\boldsymbol{x}}) = (r,\, a,\, p)$ belongs to Ω_{etopo}, its radial coordinate $r_{\boldsymbol{x}}$ was compared with a linear interpolation of the values of radii of three ETOPO1 points with the nearest spherical coordinates. For example, assume that such three points have the spherical coordinates (r_j, a_j, p_j), $j = 1, 2, 3$, where

$0 \le a \le 2\pi$ and $0 \le p \le \pi$ are the azimuth and polar angle, respectively, and

$$\begin{aligned} a_1 &= lB, \quad a_2 = (l+1)B, \qquad a_3 = lB, \qquad 0 \le l \le 21\,599; \\ p_1 &= mB, \qquad p_2 = mB, \qquad p_3 = (m+1)B, \quad 0 \le m \le 10\,799. \end{aligned}$$

Without loss of generality, the inequalities

$$lB \le a_{\boldsymbol{x}} < (l+1)B; \qquad mB \le p_{\boldsymbol{x}} < (m+1)B$$

hold true. The point inclusion function is defined in this case as

$$\chi(\boldsymbol{x}; \Omega_{\text{etopo}}) = \begin{cases} 1, & r_1 + \frac{a_{\boldsymbol{x}} - a_1}{B}(r_2 - r_1) + \frac{p_{\boldsymbol{x}} - p_1}{B}(r_3 - r_1) \;<\; r_{\boldsymbol{x}} \;<\; 1.1; \\ 0, & \text{otherwise,} \end{cases}$$

with 1.1 being the radius of the outer sphere in the chosen scale. In effect, the algorithm for evaluating the $\chi(\cdot\,; \Omega_{\text{etopo}})$ described here coincides with the star-shaped point location algorithm from [89, Section 2.2], applied to the interpolated Earth surface and the outer spherical boundary.

Our node set consists of 1,356,566 nodes with the nearest-neighbor separation close to the constant $\rho(\boldsymbol{x}) = 0.01124$, and our top priority was to ensure the low variance of the radial separation across the configuration, especially on the surface; the general view of the set is given in Figure 4.2. We used the piecewise IL with golden-ratio derived parameters $\alpha_1 = \sqrt{2}$, $\alpha_2 = (\sqrt{5} - 1)/\sqrt{2}$; regarding these α_1, α_2 see also the discussion in Appendix. Several statistics of the resulting set are presented in the following table; here again we used the common notation $\langle x \rangle$ for the averaged value of a quantity x. Notation Δ^k stands for the distance to the k-th nearest neighbor.

	Whole node set	Surface nodes
$\overline{\Delta^{12}(\boldsymbol{x})/\Delta^2(\boldsymbol{x})}$	1.3674	2.0353
$\overline{\Delta^4(\boldsymbol{x})/\Delta^1(\boldsymbol{x})}$	1.0859	1.34019
99th percentile of $\{\Delta(\boldsymbol{x}_i)\}$	0.012143	0.014444
$\overline{\Delta(\boldsymbol{x})}$	0.011243	0.010879
1st percentile of $\{\Delta(\boldsymbol{x}_i)\}$	0.009652	0.009340

Figure 4.3 illustrates the distribution of nodes close to the surface of Ω_{etopo}. No pullback function has been used, just the inclusion check performed as in (4.10). Observe that the near-surface nodes display no artifacts, and the spacing does not significantly depend on the local surface shape. The left subplot in the Figure 4.4 illustrates the effect of `Step 5` on the distribution of distances to the nearest neighbor. In the right subplot, we have collected distances to the 12 nearest neighbors for the whole configuration, and separately for the surface subset. The histogram also contains the distribution of *hole radii*, that is, distances from the Voronoi centers

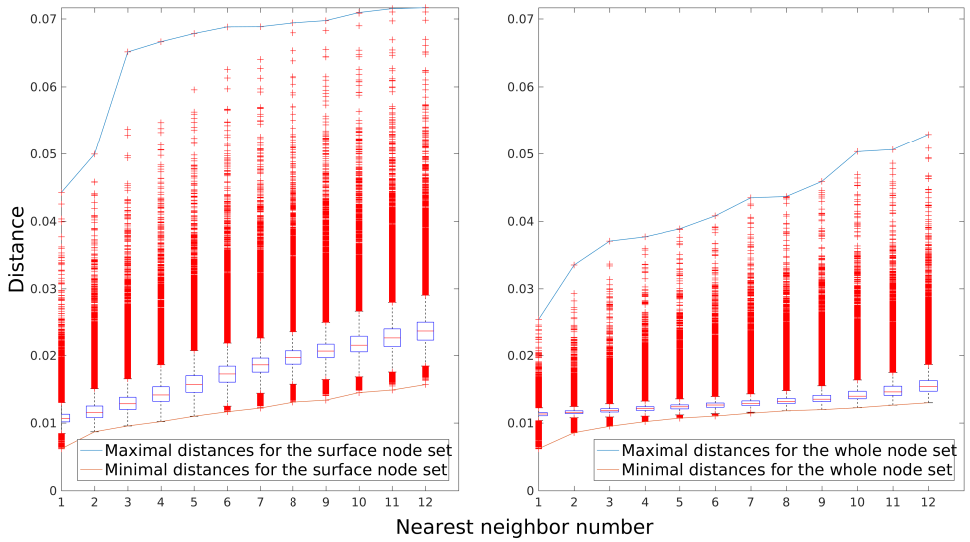

Figure 4.5: Distribution of distances to the 12 nearest neighbors for the atmospheric node configuration; medians and the 25th and 75th percentiles are shown. Left: the surface subset. Right: the whole set. Scales are the same in both subplots.

of the entire node configuration to their respective nearest nodes. It is a well-known fact that the Voronoi centers are local maxima of the distance from the node set [38], considered as a function on the whole space $\mathbb{R}^3$. Note that all the histograms on the right are normalized by probability, not by the node count.

The pair of plots in Figure 4.5 shows in detail the distribution of distances to the nearest neighbors in the sample node set. It has been produced using the standard Matlab routine *boxplot*. For each of the blue boxes corresponding to a specific nearest neighbor, the central mark is the median, the edges of the box denote the 25th and 75th percentiles. The red crosses mark outliers.

4.4.2 Point cloud

To demonstrate a nonuniform node distribution using our algorithm, we fix a collection of 100 points, $\mathcal{X}_{100}$, inside the cube $[-1,1]^3$, and consider the following radial density function:

$$\rho(\boldsymbol{x}) = \left(\Delta(\boldsymbol{x};\mathcal{X}_{100}) + \Delta^2(\boldsymbol{x};\mathcal{X}_{100})\right)/20,$$

where, as above, Δ^k is the distance to the k-th nearest neighbor to the nodes in $\mathcal{X}_{100}$. A possible interpretation of this density is a distribution that concentrates about a set of points $\mathcal{X}_{100}$, which are of particular interest for a certain model.

We proceed as in the `Step 5` of the algorithm, not using the full gradient expression described in (4.13). In fact, it is instructive to note that computing the second term in (4.13) would be quite cumbersome here in view of ρ being a piecewise-defined function. One could thus consider

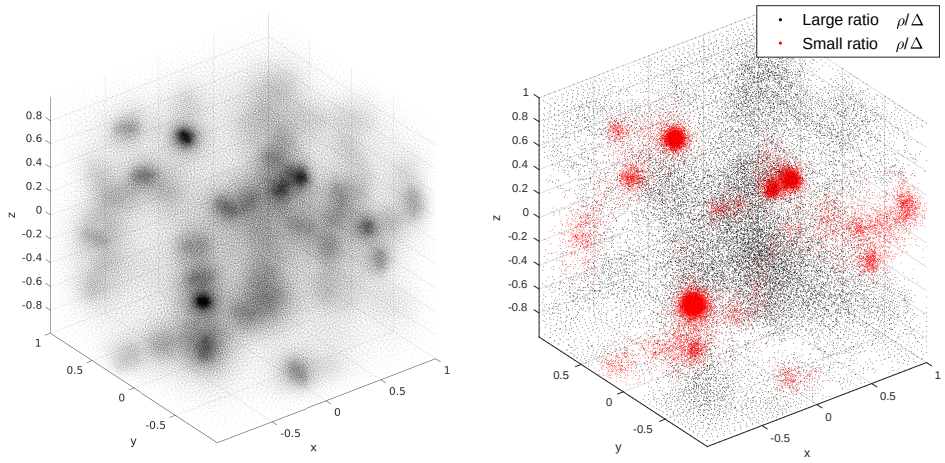

Figure 4.6: Left: the node set from Section 4.4.2. Right: node locations that contribute to the distribution of the ratio $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$ beyond the 5- and 95-percentiles.

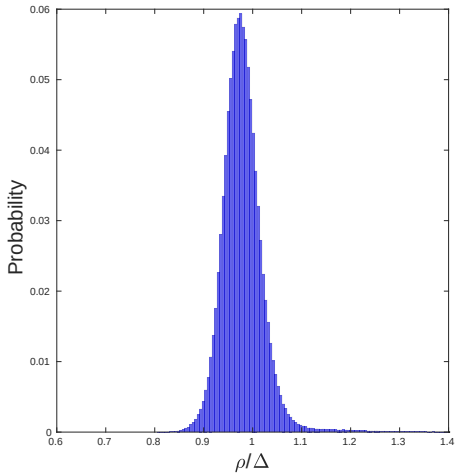

Figure 4.7: Distribution of the ratios $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$ for the node set in Section 4.4.2.

the density recovery for this distribution, Figure 4.7, as a validation of the gradient truncation approach in `Step 5`; cf. Section 4.3.2. The Q-MC voxels were drawn from the sequence $\{\mathcal{L}_n\}$ with the same lattice parameters as in Section 4.4.1, $\alpha_1 = \sqrt{2}$, $\alpha_2 = (\sqrt{5}-1)/\sqrt{2}$.

Figure 4.7 contains the distribution of the ratio $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$. The minimal and maximal values of the ratio are about 0.8099 and 1.8231 respectively; its mean value is 0.9797, and the variance is 0.0019. The 5- and 95-percentiles are 0.9208 and 1.0441, respectively; the right plot in Figure 4.6 highlights the outliers in the ratio distribution.

4.4.3 Spherical shell

The motivation for this example comes from atmospheric modeling. Representing the Earth surface by a sphere, we consider first a thin 3-dimensional shell Ω_{shell} of inner radius R_{inner} and outer radius $R_{\text{inner}} + H_{\text{atm}}$ with constant target separation h between points in the radial (vertical) direction, and the tangential (horizontal) separation to be $\tau(r) = C \cdot r$ at radius r, for some constant C. With typical choices of parameters, τ will be much larger than h, reflecting the much higher resolution needed in the vertical direction due to $H_{\text{atm}} \ll R_{\text{inner}}$. We make a radial change of variables, which can be written in spherical coordinates as $(r\,, a\,, p) \mapsto (\hat{r}(r),\, a,\, p)$, so that any configuration in Ω_{shell} having the 2-directional resolutions $\tau(r)$ and h will have isotropic resolution after the transformation. It is much easier to construct RBF bases in the isotropic case, hence our deliberation.

Following this change of variables, the radial/tangential node separations become, respectively,

$$\begin{aligned} \hat{\nu}(r) &= h \cdot \hat{r}'(r) \\ \hat{\tau}(r) &= C \cdot \hat{r}(r). \end{aligned} \tag{4.14}$$

Setting these two quantities to be equal, we obtain the ODE

$$\hat{r}'(r) = \frac{C}{h} \cdot \hat{r}(r)$$

with initial condition $\hat{r}(R_{\text{inner}}) = 1$, and its solution becomes

$$\hat{r}(r) = \exp\left(C \cdot \frac{r - R_{\text{inner}}}{h}\right).$$

From the second equation in (4.14) follows that our goal is to generate a node set in the $(\hat{r},\, a,\, p)$-space, whose separation is proportional to $\hat{r}$ and is equal in all directions: $\rho(\boldsymbol{x}) = C \cdot \|\boldsymbol{x}\|$. The outer radius of the image of Ω_{shell} in the $(\hat{r},\, a,\, p)$-space is a function of R_{inner} and H_{atm}; our model implies $R_{\text{inner}} = 6{,}371{,}220$, the mean radius of the Earth in meters, and $H_{\text{atm}} = 12{,}000$, the thickness of the atmospheric layer we are interested in. The constant C is determined by the desired tangential separation at the $r = R_{\text{inner}}$ level (see Figure 4.8).

Say, we intend to generate nodes corresponding to the 2 degree resolution on the spherical "Earth surface" and $h = 400$ meter vertical resolution. Due to the peculiarities of atmospheric modeling, we would like to fix two much denser sets of nodes on the inner and outer boundary of Ω_{shell}; specifically, we are using 12,100 approximate Riesz energy minimizers on a sphere, appropriately rescaled. The interior nodes are generated using our algorithm, and then `Step 5` is modified so that to leave the boundary subset intact. This, however, causes a difficulty: the separation distances between the interior and the surface nodes must remain large; on the other hand, our generic formulation of `Step 5` does not account for the much higher concentration of

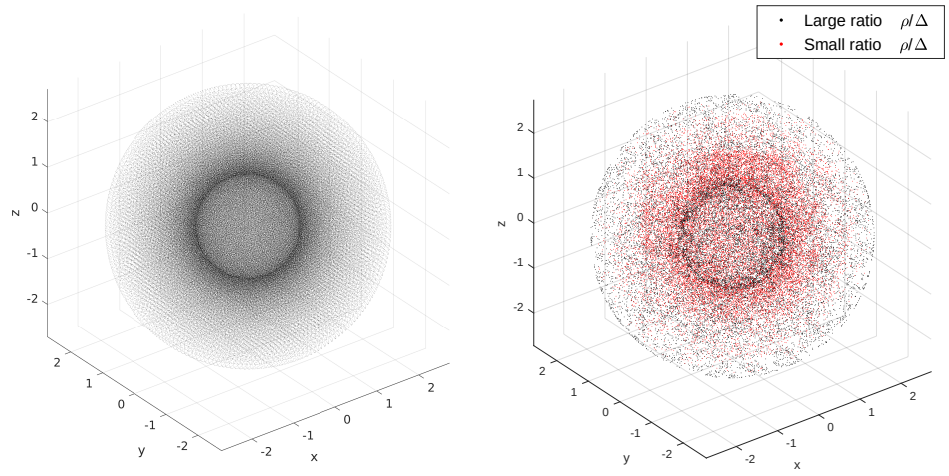

Figure 4.8: Left: the node set from Section 4.4.3. Right: node locations that contribute to the distribution of the ratio $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$ beyond the 5- and 95-percentiles.

nodes on the surface, which causes excessive repelling force, seen in the oscillations of the radial distribution, Figure 4.9. Mitigating this effect requires artificially weakening the repulsive force caused by the boundary nodes, a straightforward task using our codebase. Instead, we show in Figure 4.9 the performance of the generic algorithm, to illustrate complications that may arise when applying it to specialized problems.

The set Ω_{shell} can be challenging for the basic form of our algorithm, as described in Section 4.3.1: obtaining satisfying convergence requires using the full version of gradient descent (4.13). The reasons for it being more difficult to tackle than, say, Ω_{etopo} in Section 4.4.1, are that due to convexity of the outer boundary, the weighted s-energy minimizers on it are denser than on the sphere with radius $\hat{r}(R_{\text{inner}} + H_{\text{atm}}) - 10^{-3}$, for example; see also discussion at the end of Section 4.2.2. Getting rid of the artifacts at the endpoints of the radial distribution is done by using the full gradient, weakening the repulsion of the fixed boundary nodes, and *not* striving for the full convergence of a minimization method applied to the Riesz energy.

In this example, we used the $\{\mathcal{M}_n\}$ sequence to fill individual voxels. The left subplot in Figure 4.9 contains the distribution of the ratio $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$. The minimal and maximal values of the ratio are about 0.9165 and 1.8989 respectively; its mean value is 1.0226, and the variance is 0.0024. The 5- and 95-percentiles are 0.9782 and 1.0717, respectively.

4.4.4 Run times

The execution times (in seconds) for the above examples are summarized in Table 4.1, where, as before, K and T stand for the number of nearest neighbors and the number of iterations used in the repel procedure in `Step 5`, respectively. The fifth column contains times required to fill the voxels selected at `Step 1` with configurations from either $\{\mathcal{L}_n\}$ or $\{\mathcal{M}_n\}$ and to remove any redundant

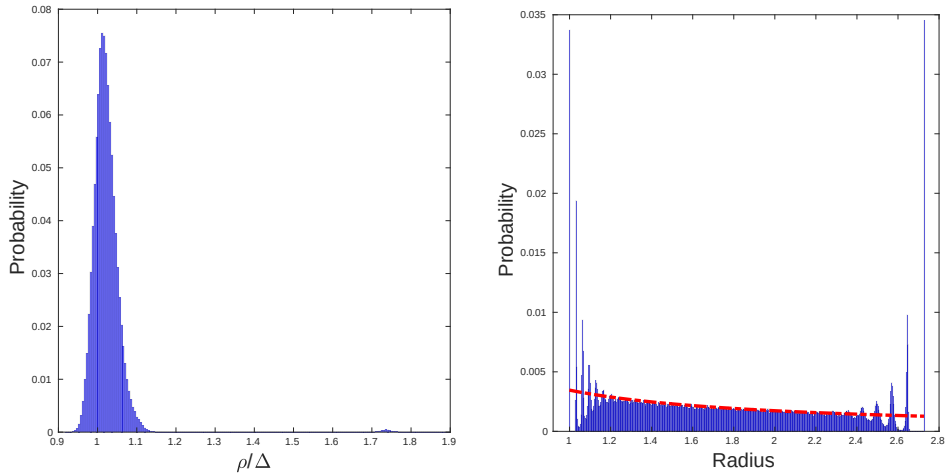

Figure 4.9: Left: Distribution of the ratios $\rho(\boldsymbol{x})/\Delta(\boldsymbol{x})$ for the node set in Section 4.4.3. Right: Radial node distribution, actual (blue) and the theoretic (red) continuous component; i.e., without the δ-function spikes at the endpoints.

Example	K	T	N	Q-MC distribution times, s	Repel times, s
Atmospheric nodes	33	29	$1,356,566$	5	89
Point cloud	30	200	$577,321$	4	840
Spherical shell	30	200	$358,915$	1	144

Table 4.1: Timings of the examples in Sections 4.4.1–4.4.3.

nodes as in `Step 4`. All the computations were performed on a dedicated machine with 40 GB RAM and an 8-core *Intel Xeon* CPU. Note that the basic Q-MC node sets for both sequences were precomputed, and the pre-computation times are not included in the table. Computation of configurations in $\{\mathcal{L}_n\}$ for $1 \le n \le 200$ took less than 1 second. An implementation of the $\{\mathcal{M}_n\}$ sequence for $1 \le n \le 200$ took 4311 seconds to generate; coordinates of the resulting minimizers as well as the corresponding average separation distances are distributed with the associated codebase [102].

4.5 Final observations and comparisons

4.5.1 Comparisons

Of the two Q-MC sequences we considered, the periodic Riesz minimizers appear more promising, being devoid of the lattice structure and having high space utilization. On the other hand, we have successfully applied ILs as an elementary uniform configuration. One could use different

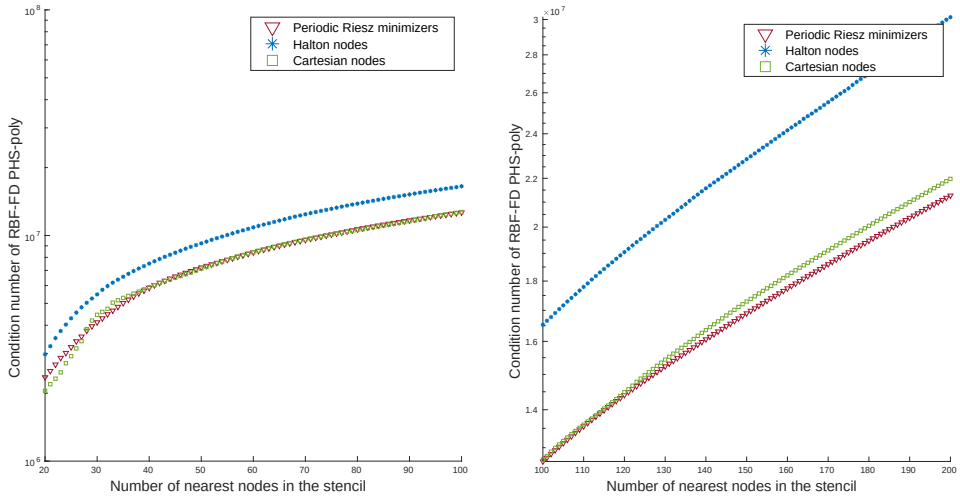

Figure 4.10: Average condition numbers of the joint RBF-FD PHS-based matrices for order 1- and 2-differential operators. Left: Between 20 and 100 nodes in the stencil. Right: Between 100 and 200 nodes in the stencil.

sets of irrational parameters $\alpha_1, \ldots, \alpha_{d-1}$ for different numbers of nodes in a voxel. Although this might be useful in mitigating the non-isotropic behavior of ILs, it makes hard to control node separation at voxel interfaces.

Another quasi-uniform node set commonly used in Q-MC methods can be constructed from the Halton sequence [76, Chapter 5.4], an example of a low-discrepancy sequence. To see how the Halton nodes compare to Riesz minimizers for RBF-FD methods, we have computed condition numbers of PHS-based RBF-FD matrix in the LHS of (4.4), using operators $\partial/\partial x_i$, $i = 1, 2, 3$, and $\partial^2/\partial x_i \partial x_j$, $i \leq j = 1, 2, 3$ as $\mathfrak{L}$; the resulting LHS then constitutes the joint system for the weights corresponding to these nine differential operators. We used the RBF kernel $\varphi(r) = r^5$, and the polynomials in the interpolation space were of degree at most 2, see Section 4.2.1. The computations were performed for the Riesz and Halton nodes, and the uniform Cartesian grid. The stencils consisted of K nearest neighbors of a random vector with Gaussian distribution, centered around $(0.5, 0.5, 0.5)^{\mathrm{tr}}$; the evaluation point $\boldsymbol{x}_0$ was taken equal to the random vector itself. The nearest neighbors were drawn from 1000 nodes of the respective sequence, uniformly distributed over the unit cube. The Riesz nodes were produced by minimizing periodic energy (4.5) with the distance (4.9).

Figure 4.10 contains a comparison of the condition numbers of RBF-FD matrices for the three sequences. The values shown are averages of the condition numbers for 500 random stencil centers $\boldsymbol{x}_0$; the averaging was introduced to eliminate the rather unpredictable dependence on $\boldsymbol{x}_0$, and to display the underlying trend. We omitted values of K below 20 from the plots, as all the three node sets resulted in relatively ill-conditioned systems; this was to be expected, as the recommended stencil size is roughly twice the number of linearly independent polynomials in the

interpolation space (there are 10 monomials of degree at most 2 in $\mathbb{R}^3$) [46].

4.5.2 Range of applications

Our method has proven very efficient for slowly varying radial densities that are small (recall that small radial density means a large number of nodes per unit volume) compared to the entire node set scale, and is capable of handling very complex underlying sets. The range of dimensions where the algorithm can be used efficiently is determined by the applicability of Q-MC initialization and the nearest neighbor searches: the repelling iterations for Riesz energy in `Step 5` are largely (with a proper value of s) dimension-agnostic. A shortcoming that is common to all quasi-Monte Carlo methods (but of little practical relevance) is a much worse performance (measured by L^2 discrepancy), compared to Monte-Carlo distribution, in dimensions starting at about 15 [31]. Furthermore, using the uniform grid to detect the support Ω, as is done in `Step 1`–`Step 2`, becomes unfeasible already for $d = 10$; instead, one needs an efficient way to determine which corners of the grid are in some sense close to Ω. This is certainly not a feature of our approach, but a manifestation of the curse of dimensionality: treatment of a complicated high-dimensional set is a computationally intensive task. Regarding finding the nearest neighbors it should be noted that common implementations of k-d trees are efficient only up to about $d = 20$; additionally, the k-d tree approach is faster than the full brute force search only if $N \gg 2^d$ [58]. On the other hand, as has already been noted, our repelling procedure does not require frequent updates of the search tree, as the local adjacency largely remains intact.

The suggested algorithm is very local, and it therefore must be straightforward to add multi-resolution and adaptive refinement, as is widely done for grids [35, 40], yet as of this writing, our proof-of-concept implementation does not include these features. Still, we would like to observe that refining the voxel structure is indeed easier than refining a mesh, since no geometry is taken into account. This partially addresses the previous remark on detection of Ω in high dimensions.

The closest set of goals to what we have presented here, that we're aware of, is posed in the pioneering paper [99]; our method is crafted for full-dimensional domains, and apparently performs faster in this case. The bubble packing algorithm in [99] is conceptually similar to the greedy filling of centers in `Step 3`, while physical relaxation is an alternative to the energy minimization we employ; of course, the idea of relaxation can also be found in a number of related references, and is a well-known approach in this context, see for example [86]. Our method requires computing the gradient of the desired radial density in the cases when the outer boundary of the underlying set is uniformly convex, and/or when the radial density changes quickly. Alternatively, fine partition of the set is necessary. Either solution, however, may be computationally expensive.

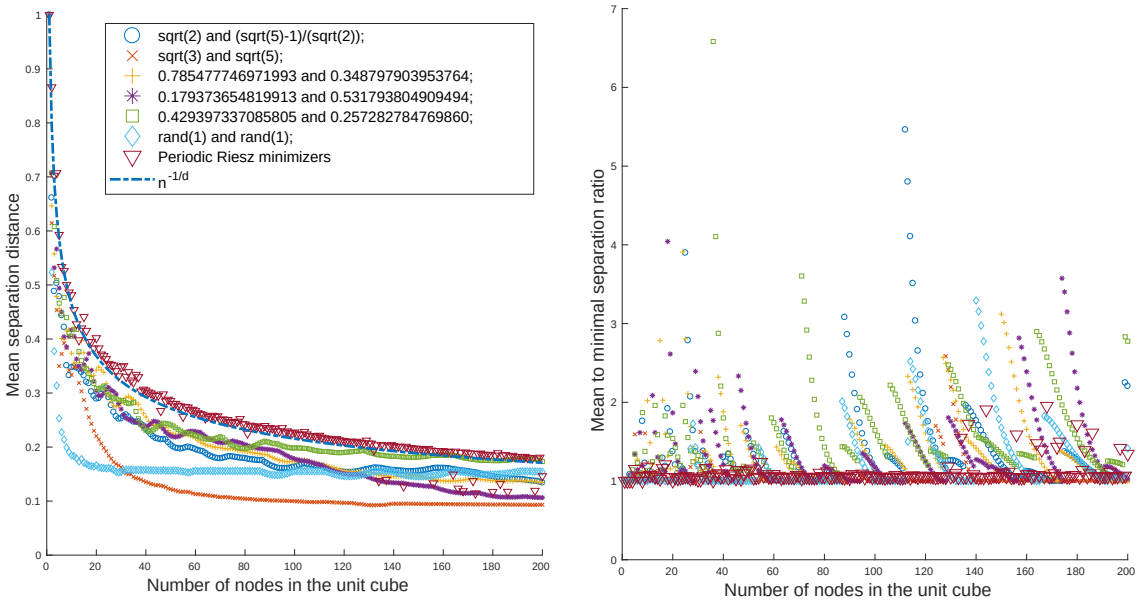

Figure 4.11: Left: dependence of the mean separation distances on the number of nodes in the unit cube for different values of parameters $\alpha_1,\ \alpha_2$; the $n^{-1/d}$ decay rate shown as a dash-dot line. Right: ratios of the mean separation distances to the minimal ones for the same configurations.

4.6 Separation properties of sequences $\{\mathcal{L}_n\}$ and $\{\mathcal{M}_n\}$

This section deals with the results of our numerical experiments, set in the 3-dimensional space. The function $\lambda(r)$ used in `Step 2` is the number of nodes in the unit cube $[0,1]^3$, placed according to (4.8), or obtained by minimizing the Riesz s-energy (4.5) with periodic metric, such that the mean separation distance of these nodes is the closest to r. To compute $\lambda(r)$ for the periodization of $\{\mathcal{L}_n\}$, we tabulate mean separations $\langle\Delta_n\rangle$ in a sample configuration comprising $\mathcal{L}_n$ and its $26 = 3^3 - 1$ copies, obtained translating $\mathcal{L}_n$ by the vectors $\big\{(i,j,k)^{\rm tr} : i,j,k \in \{0,\pm 1\}$ and $|i|+|j|+|k|>0\big\}$. The tabulated dependence of separation on n is then inverted and interpolated using a piecewise cubic Hermite interpolating polynomial. The reason to consider separation distance between configurations in 3^d cubes in dimension d (and not a single cube with a single instance of $\mathcal{L}_n$) is to account for the boundary effects. Likewise, to compute $\lambda(r)$ for the Riesz minimizers, the mean separation of $\mathcal{M}_n$ is tabulated for $1 \le n \le n_{\max}$, then the inverse dependence is interpolated. No copies of $\mathcal{M}_n$ are considered alongside the original configuration, since periodicity condition is already included in the metric (4.9).

In general, putting too many nodes in individual voxels is justified only if the radial density function ρ varies slowly. For our applications, $n_{\max} \le 100$ was sufficient. The left plot in Figure 4.11 illustrates the delicate dependence of the separation distances of ILs on the lattice parameters. While any set of irrational quantities $\alpha_1, \dots, \alpha_{d-1}$ in (4.8) that are linearly independent over rationals will give a uniformly distributed IL as n grows, certain values may perform better than

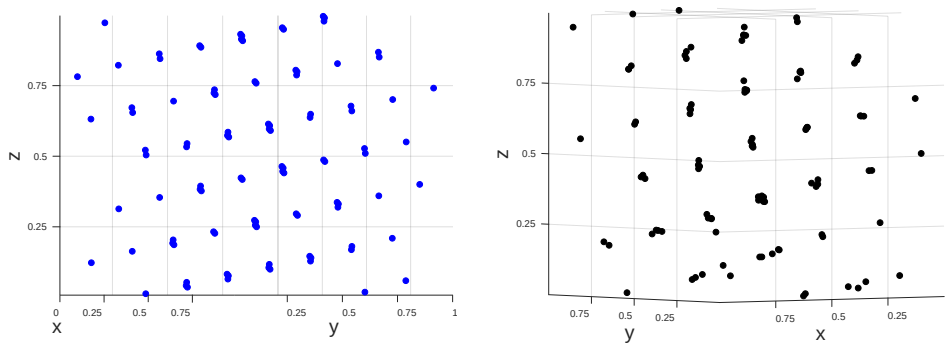

Figure 4.12: Left: A cross-section of the IL $\mathcal{L}_{100}$ with parameters 0.179373654819913 and 0.531793804909494. Right: A (different) cross-section of the periodic Riesz minimizer $\mathcal{M}_{100}$.

the others. In particular, adjustments can be made to improve the distribution for small values of n. For example, it is known from [9] that a 2-dimensional IL generated by the golden ratio has optimal L^2 discrepancy. Numerical experiments have shown that its 3-dimensional analog with parameters $\alpha_1 = \sqrt{2}$, $\alpha_2 = (\sqrt{5}-1)/\sqrt{2}$ does perform well for large numbers of nodes; yet by carrying out a Monte Carlo search for the parameters maximizing separation distance in (4.8), we found several (necessarily rational) pairs that performed at least just as well for up to $n = 200$, see Figure 4.11.

Curiously enough, a pair of random numbers drawn uniformly from $[0, 1]$ (shown in the legend as rand(1)), consistently performed better than the pair $\sqrt{3}$ and $\sqrt{5}$, starting at $n \approx 40$. We were able to reproduce this behavior in a number of runs; in fact, we haven't seen a random pair that wouldn't always outperform $\sqrt{3}$ and $\sqrt{5}$ after a fairly small n.

The second graph in Figure 4.11 shows the ratios of the mean to minimal separation distances $\langle\Delta_n\rangle/\Delta_n$ for the same range of n. In both subfigures, Riesz periodic minimizers clearly stand out, by having the largest mean separation (left), and by smallest ratios (right). This means, the nearest neighbor distances $\Delta(\boldsymbol{x})$ vary little from node to node in the $\{\mathcal{M}_n\}$ sequence. We conclude this section by presenting in Figure 4.12 a pair of cross-sections of the IL $\mathcal{L}_{100}$ and the configuration $\mathcal{M}_{100}$ that look remarkably similar. In fact, we found the vague resemblance between the low-energy periodic configurations and lattice structures, similar to ILs, quite interesting, given the connection between packing and Riesz energy minimization [62], and that the highest packing density in the 3-dimensional space is achieved, in particular, by the hcp lattice [38].

Appendix A

Glossary

Term	**Description**	**Pages**
Ahlfors regular set	see d-regular set	4, 27
$\mathbb{B}^d$	the unit ball in $\mathbb{R}^d$	13
$B(\boldsymbol{x}, r)$	a ball of radius of radius r, centered at $\boldsymbol{x} \in \mathbb{R}^p$	3, 4, 24
$\mathcal{C}^d$	unit cube $[0,1]^d$	17, 78
$\overline{M}$	closure of a set M	24
M^c	complement of a set M	18
$\boldsymbol{c}_m$	the closest to the origin corner of $\mathcal{V}_m$	78, 82
$\Delta(\boldsymbol{x})$	distance from the node $\boldsymbol{x}$ to its nearest neighbor in ω_N	71
$\Delta(\{\boldsymbol{x}_1, \dots \boldsymbol{x}_N\})$	separation of the configuration	75
$\Delta^k(\boldsymbol{x}),\ k = 1, \dots, K$	distance from $\boldsymbol{x}$ to the k-th nearest neighbor	85, 86
$\bar{\Delta}_n$	mean separation of the periodized $\mathcal{L}_n$, $\mathcal{M}_n$	78
$\operatorname{diam}(\Omega)$	diameter of the set Ω	4, 27, 61
d-rectifiable set	the image of a bounded subset of $\mathbb{R}^d$ under a Lipschitz mapping	2, 13, 16, 50
d-regular set	satisfying $c_0 R^d \le \mathcal{H}_d(B(\boldsymbol{x}, R) \cap \Omega) \le C_0 R^d$	4, 27
$E(\boldsymbol{x}_1, \dots, \boldsymbol{x}_N; g_s, \kappa, q)$	combined Riesz kernel, showing all the arguments explicitly	8, 50
$\frac{1}{\mathcal{T}(N)} E_N(\mu; g_s, q)$	extension of the discrete Riesz energy to $\mathcal{P}(\Omega)$	15
$\frac{1}{\mathcal{T}(N)} E_N(\mu; g_s, \kappa, q)$	extension of the weighted discrete Riesz energy to $\mathcal{P}(\Omega)$	50
$E^q_{s,d}$	short notation for the unweighted Riesz functional $E(\cdot\,; g_s, 1, q)$ with external field	22
E^κ_s	short notation for the weighted Riesz functional $E(\boldsymbol{x}_1, \dots, \boldsymbol{x}_N; g_s, \kappa, 0)$ without external field	74
E_s	short notation for the unweighted Riesz functional $E(\cdot\,; g_s, 1, 0)$ without external field	13, 55
$\{t\}$	fractional part of the real number t	6, 59, 77
Γ-convergence	convergence of functionals on metric spaces, preserving minimizers	iv, 3, 15, 17, 22, 51

Term	Description	Pages
$\mathcal{H}_d$	d-dimensional Hausdorff measure	13
$\chi(\cdot\,;\Omega)$	indicator function of the set Ω	74, 85
$\mathring{M}$	interior of a set M	24
κ	multiplicative weight for the Riesz energy functional	8, 14, 50, 74, 75, 95
λ	interpolated inverse of $\bar{\Delta}_n$	78
$\mathcal{L}_n,\ n \geq 1$	n-point irrational lattice	77
$\langle t \rangle$	mean value of the quantity t	78, 85
$\mathcal{M}_n,\ n \geq 1$	n-point periodic Riesz minimizer	77
n_m	number of nodes in $\mathcal{V}_m$	79, 80
$\|\cdot\|$	Euclidean distance in $\mathbb{R}^p$	1
Ω	target distribution support	1–3, 56, 78
ω_N	a subset of Ω containing N elements	3, 6, 12
$(t)_+$	$\max(0,\, t)$, positive part of the number	24
$\mathcal{P}_N(\Omega)$	measures in $\mathcal{P}(\Omega)$, supported on N points	1, 3
$\mathcal{P}(\Omega)$	the space of probability measures supported on Ω	1
PSB	Poppy-seed bagel theorem	iv, 2, 16, 22, 24
Q-MC	Quasi-Monte Carlo	69
$(r,\, a,\, p)$	spherical coordinates	84
RBF-FD	Radial Basis Function-generated Finite Differences	69
$\rho(\boldsymbol{x})$	objective radial density	70, 78
$S(\mu; g_s, q)$	Γ-limit of the unweighted discrete energies with external field	16
$S(\mu; g_s, q)$	Γ-limit of the weighted discrete energies with external field	50
σ	separation distance between metrically separated sets	56, 60
$\operatorname{supp}\mu$	support of the measure μ	11, 29
$\mathcal{T}(N)$	order of the asymptotics of the Riesz energy	8
$\tau(N)$	asymptotic scaling factor for the external field term in Riesz energy	8, 22
$\mathcal{V}_m,\ m = 1, \ldots, M^d$	cube-shaped voxel in $\mathbb{R}^d$	76, 78
$\xrightarrow{*}$	convergence in the weak* topology	13, 74
$\boldsymbol{x};\ \boldsymbol{x}_i,\ i = 1, \ldots, N$	points in $\mathbb{R}^d$; nodes of the configuration	1, 8
$\boldsymbol{x}_{j(i,k)},\ k = 1, \ldots, K$	the k-th nearest neighbor of $\boldsymbol{x}_i$	7, 75, 79
$\boldsymbol{z}_m$	center of $\mathcal{V}_m$	78

Appendix B

Repel-type code listings

For the complete source code of the routines belows, see repository at [102].

B.1 Full dimension case

Below is an excerpt from the Matlab routine performing repelling iterations discussed in `Step 5-Step 6` of the main algorithm of Chapter 4.

```
%% Main loop
for iter=1:repel_steps
    if mod(iter,10) == 1
        [IDX, ~] = knnsearch(cnf', cnf(:,1:N_moving)', 'k', k_value+1);
        IDX = IDX(:,2:end)';
    end
%% Vectors from nearest neighbors
    cnf_repeated = reshape(repmat(cnf(:,1:N_moving),k_value,1),dim,[]);
    cnf_repeated_concentric = cnf_repeated./sqrt(sum(cnf_repeated.*cnf_repeated
        ,1));
    knn_cnf = cnf(:,IDX);
    knn_differences = cnf_repeated - knn_cnf;
    knn_norms_squared = sum(knn_differences.*knn_differences,1);
%% Weights using radial density
    riesz_weights = compute_riesz(knn_norms_squared);
    if isa(densityF,'function_handle')
        knn_density =  densityF(knn_cnf);
        density_weights = compute_weights(knn_density);
        weights = s*density_weights .* riesz_weights ./ ...
                                                    knn_norms_squared;
    else
        weights = s*riesz_weights./knn_norms_squared;
    end
%% Sum up over the nearest neighbors
    gradient = bsxfun(@times,weights,knn_differences);
    gradient = reshape(gradient, dim, k_value, []);
    gradient = reshape(sum(gradient,2), dim, []);
%% Add noise and renormalize
    if isa(noise,'function_handle')
        gradient = gradient + noise() * mean(sqrt(sum(gradient.*gradient,1)));
    end
    directions = gradient./sqrt(sum(gradient.*gradient,1));
    step = sqrt(min(reshape(knn_norms_squared,k_value,[]),[],1));
```

```
    cnf_tentative = cnf(:,1:N_moving) +...
                            directions(:,1:N_moving).*step/(offset+iter-1);
%% Detect the domain
    if exist('in_domainF', 'var') && isa(in_domainF,'function_handle')
        domain_check = in_domainF( cnf_tentative(1,:), cnf_tentative(2,:),
            cnf_tentative(3,:));
    else
        domain_check = ~any((cnf_tentative<-A/2.0) | (cnf_tentative>A/2.0),1);
    end
    if isa(pullbackF,'function_handle')
        cnf(:,~domain_check) = pullbackF(cnf_tentative(:,~domain_check));
    end
    cnf(:,domain_check) = cnf_tentative(:,domain_check);
end
toc
```

B.2 Repelling on implicit surfaces

The following listing contains an excerpt from a Matlab routine distributing nodes on an implicit surface. Note that the pullback in this case is done by solving a mini-Newton's method problem.

```
%% Main loop
for iter=1:repel_steps
    if mod(iter,10) == 1
        [IDX, ~] = knnsearch(cnf', cnf(:,1:N_moving)', 'k', k_value+1);
        IDX = IDX(:,2:end)';
    end
%% Vectors from nearest neighbors
    cnf_repeated = reshape(repmat(cnf(:,1:N_moving),k_value,1),dim,[]);
    knn_cnf = cnf(:,IDX);
    knn_differences = cnf_repeated - knn_cnf;
    knn_norms_squared = sum(knn_differences.*knn_differences,1);
%% Weights using radial density
    riesz_weights = compute_riesz(knn_norms_squared);
    if isa(densityF,'function_handle')
        knn_density =  abs(densityF(knn_cnf)) + .1;
        weights = s* riesz_weights ./ knn_norms_squared ./ knn_density;
    else
        weights = s*riesz_weights./knn_norms_squared;
    end
%% Sum up over the nearest neighbors
    gradient = bsxfun(@times,weights,knn_differences);
    gradient = reshape(gradient, dim, k_value, []);
    gradient = reshape(sum(gradient,2), dim, []);
%
```

```
    surfnormals = ngrad./sqrt(sum(ngrad.*ngrad,1));
    tangentgrad = gradient - surfnormals .* sum(gradient.*surfnormals, 1);
    if mod(iter,50) == 1
        tangentgradnorm = sqrt(max(sum(tangentgrad .* tangentgrad, 1)))
    end
%
    directions = tangentgrad./sqrt(sum(tangentgrad.*tangentgrad,1));
    step = sqrt(min(reshape(knn_norms_squared,k_value,[]),[],1));
    cnf_tentative = cnf(:,1:N_moving) +...
                            directions(:,1:N_moving).*step/(offset+iter-1);
%% Pullback to surface
    h =   surfF(cnf_tentative) ;
    ngrad = gradF(cnf_tentative);
    while max(abs( h )) > 1e-4
         cnf_tentative = cnf_tentative - ngrad .* h ./ sum(ngrad .* ngrad, 1);
         h =   surfF(cnf_tentative)  ;
         ngrad = gradF(cnf_tentative);
    end
    cnf = cnf_tentative;
end
```

BIBLIOGRAPHY

[1] Alishahi, K., and Zamani, M. The spherical ensemble and uniform distribution of points on the sphere. *Electron. J. Probab. 20* (2015).

[2] Alishahi, K., and Zamani, M. The spherical ensemble and uniform distribution of points on the sphere. *Electron. J. Probab. 20* (2015).

[3] Amante, C., and Eakins, B. W. ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. `https://doi.org/10.7289/V5C8276M`, 2009.

[4] Bayona, V., Flyer, N., Fornberg, B., and Barnett, G. A. On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs. *J. Comput. Phys. 332* (2017), 257–273.

[5] Beltrán, C., Marzo, J., and Ortega-Cerdà, J. Energy and discrepancy of rotationally invariant determinantal point processes in high dimensional spheres. `Arxiv:1511.02535` (2015).

[6] Beltrán, C., Marzo, J., and Ortega-Cerdà, J. Energy and discrepancy of rotationally invariant determinantal point processes in high dimensional spheres. *J. Complex. 37* (2016), 76–109.

[7] Billingsley, P. *Convergence of probability measures*. John Wiley & Sons, 1999.

[8] Bilogliadov, M. Weighted energy problem on the unit sphere. `Arxiv:1510.06420` (2015).

[9] Bilyk, D. The L^2 discrepancy of irrational lattices. In *Monte Carlo and Quasi-Monte Carlo Methods 2012*. Springer, 2013, pp. 289–296.

[10] Bilyk, D., Temlyakov, V. N., and Yu, R. The L^2 Discrepancy of Two-Dimensional Lattices. In *Springer Proc. Math. Stat.*, vol. 25. 2012, pp. 63–77.

[11] Bollig, E. F., Flyer, N., and Erlebacher, G. Solution to PDEs using radial basis function finite-differences (RBF-FD) on multiple GPUs. *J. Comput. Phys. 231*, 21 (2012), 7133–7151.

[12] Borodachov, S. V., Hardin, D. P., and Saff, E. B. *Minimal discrete energy on rectifiable sets*. In preparation.

[13] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics of best-packing on rectifiable sets. *Proc. Am. Math. Soc. 135*, 08 (2007), 2369–2381.

[14] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics of best-packing on rectifiable sets. *Proc. Am. Math. Soc. 135*, 08 (2007), 2369–2381.

[15] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics for discrete weighted minimal Riesz energy problems on rectifiable sets. *Trans. Amer. Math. Soc. 360*, 3 (2008), 1559–1580.

[16] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics for discrete weighted minimal Riesz energy problems on rectifiable sets. *Trans. Amer. Math. Soc. 360*, 3 (2008), 1559–1580.

[17] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics for discrete weighted minimal Riesz energy problems on rectifiable sets. *Trans. Am. Math. Soc. 360*, 03 (2008), 1559–1581.

[18] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Low Complexity Methods For Discretizing Manifolds Via Riesz Energy Minimization. *Found. Comput. Math. 14*, 6 (2014), 1173–1208.

[19] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Low complexity methods for discretizing manifolds via Riesz energy minimization. *Found. Comput. Math. 14*, 6 (2014), 1173–1208.

[20] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Low complexity methods for discretizing manifolds via Riesz energy minimization. *Found. Comput. Math. 14*, 6 (2014), 1173–1208.

[21] Braides, A. *Local Minimization, Variational Evolution and* Γ*-Convergence*. Lecture Notes in Mathematics. Springer International Publishing, 2014.

[22] Brauchart, J., Hardin, D., and Saff, E. The next-order term for optimal Riesz and logarithmic energy asymptotics on the sphere. In *Contemporary Mathematics*, J. Arvesú and G. Lagomasino, Eds., vol. 578. American Mathematical Society, Providence, Rhode Island, 2012, pp. 31–61.

[23] Brauchart, J. S., Dragnev, P. D., and Saff, E. B. Riesz extremal measures on the sphere for axis-supported external fields. *J. Math. Anal. Appl. 356*, 2 (2009), 769–792.

[24] Brauchart, J. S., Dragnev, P. D., and Saff, E. B. Riesz external field problems on the hypersphere and optimal point separation. *Potential Anal. 41*, 3 (2014), 647–678.

[25] Brauchart, J. S., Reznikov, A. B., Saff, E. B., Sloan, I. H., Wang, Y. G., and Womersley, R. S. Random point sets on the sphere—hole radii, covering, and separation. *Experimental Mathematics* (2016), 1–20.

[26] Brelot, M. Lectures on Potential Theory. 158.

[27] Broomhead, D. S. and Lowe, D. Multivariable Functional Interpolation and Adaptive Networks. *Complex Syst. 2* (1988), 321– 355.

[28] Buhmann, M. D. *Radial Basis Functions*. Cambridge University Press, Cambridge, 2003.

[29] Buhmann, M. D., Dinew, S., and Larsson, E. A note on radial basis function interpolant limits. *IMA J. Numer. Anal. 30*, 2 (2010), 543–554.

[30] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. Y. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput. 16*, 5 (1995), 1190–1208.

[31] Caflisch, R. E. Monte Carlo and quasi-Monte Carlo methods. *Acta Numerica 7* (1998), 1.

[32] Calef, M. T. A sequence of discrete minimal energy configurations that does not converge in the weak-star topology. *J. Math. Anal. Appl. 395*, 2 (2012), 550–558.

[33] Chang, C., and Lin, C. LIBSVM. *ACM Trans. Intell. Syst. Technol. 2*, 3 (2011), 1–27.

[34] Cheney, W., and Light, W. *A course in approximation theory*, vol. 101 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2009. Reprint of the 2000 original.

[35] Clawpack Development Team. Clawpack software, 2017. Version 5.4.0.

[36] Cohn, H., Kumar, A., Miller, S., Radchenko, D., and Viazovska, M. The sphere packing problem in dimension 24. *Ann. Math. 185*, 3 (2017), 1017–1033.

[37] Conway, J. B. *A course in functional analysis*, 2^{nd} ed., vol. 96 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1990.

[38] Conway, J. H., and Sloane, N. J. A. *Sphere packings, lattices and groups*, third ed., vol. 290 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, New York, 1999. With additional contributions by E. Bannai, R. E. Borcherds, J. Leech, S. P. Norton, A. M. Odlyzko, R. A. Parker, L. Queen and B. B. Venkov.

[39] David, G., and Semmes, S. *Analysis of and on uniformly rectifiable sets*, vol. 38. American Mathematical Soc., 1993.

[40] Debreu, L., Vouland, C., and Blayo, E. AGRIF: Adaptive grid refinement in fortran. *Computers & Geosciences 34*, 1 (2008), 8–13.

[41] Dragnev, P. D., and Saff, E. B. Riesz spherical potentials with external fields and minimal energy points separation. *Potential Anal. 26*, 2 (2007), 139–162.

[42] Evans, L. C., and Gariepy, R. F. *Measure theory and fine properties of functions*, vol. 5. CRC press, 1991.

[43] Falconer, K. J. *The geometry of fractal sets*, vol. 85 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1986.

[44] Fasshauer, G. E. *Meshfree approximation methods with MATLAB*, vol. 6 of *Interdisciplinary Mathematical Sciences*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2007. With 1 CD-ROM (Windows, Macintosh and UNIX).

[45] Flyer, N., Barnett, G. A., and Wicker, L. J. Enhancing finite differences with radial basis functions: Experiments on the Navier-Stokes equations. *J. Comput. Phys. 316* (2016), 39–62.

[46] Flyer, N., Fornberg, B., Bayona, V., and Barnett, G. A. On the role of polynomials in RB-FD approximations: I. Interpolation and accuracy. *J. Comput. Phys. 321* (2016), 21–38.

[47] Flyer, N., Lehto, E., Blaise, S., Wright, G. B., and St-Cyr, A. A guide to RBF-generated finite differences for nonlinear transport: Shallow water simulations on a sphere. *J. Comput. Phys. 231*, 11 (2012), 4078–4095.

[48] Flyer, N., Wright, G. B., and Fornberg, B. Radial Basis Function-Generated Finite Differences: A Mesh-Free Method for Computational Geosciences. *Handb. Geomathematics* (2013), 1–30.

[49] Fornberg, B. *A practical guide to pseudospectral methods*. Cambridge University Press, 1996.

[50] Fornberg, B. Calculation of Weights in Finite Difference Formulas. *SIAM Rev. 40*, 3 (1998), 685–691.

[51] Fornberg, B., and Flyer, N. *A Primer on Radial Basis Functions with Applications to the Geosciences*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2015.

[52] Fornberg, B., and Flyer, N. Fast generation of 2-D node distributions for mesh-free PDE discretizations. *Comput. Math. with Appl. 69*, 7 (2015), 531–544.

[53] Fornberg, B., and Flyer, N. Solving PDEs with radial basis functions. *Acta Numerica 24* (2015), 215–258.

[54] Fornberg, B., Larsson, E., and Flyer, N. Stable Computations with Gaussian Radial Basis Functions. *SIAM J. Sci. Comput. 33*, 2 (2011), 869–892.

[55] Fornberg, B., Lehto, E., and Powell, C. Stable calculation of Gaussian-based RBF-FD stencils. *Comput. Math. with Appl. 65*, 4 (2013), 627–637.

[56] Fornberg, B., Wright, G., and Larsson, E. Some Observations Regarding Interpolants in the Limit of Flat Radial Basis Functions. *Comput. Math. with Appl. 47*, 1 (2004), 37–55.

[57] Fornberg, B., and Zuev, J. The Runge phenomenon and spatially variable shape parameters in RBF interpolation. *Comput. Math. with Appl. 54*, 3 (2007), 379–398.

[58] Freidman, J. H., Bentley, J. L., and Finkel, R. A. An Algorithm for Finding Best Matches in Logarithmic Expected Time. *ACM Trans. Math. Softw. 3*, 3 (1977), 209–226.

[59] Fremlin, D. H. *Measure theory*, vol. 4. Torres Fremlin, 2000.

[60] Frey, P. J., and George, P.-L. *Quadtree-octree Based Methods*. ISTE, 2010, pp. 163–199.

[61] Hardin, D., and Saff, E. Discretizing Manifolds via Minimum Energy Points. *Not. AMS*, x (2004), 1186–1194.

[62] Hardin, D. P., and Saff, E. B. Minimal Riesz Energy Point Configurations for Rectifiable d-Dimensional Manifolds. *Adv. Math. 193*, 1 (2003), 174–204.

[63] Hardin, D. P., and Saff, E. B. Minimal Riesz energy point configurations for rectifiable d-dimensional manifolds. *Adv. Math. 193* (2005), 174–204.

[64] Hardin, D. P., Saff, E. B., and Vlasiuk, O. V. Generating point configurations via hypersingular Riesz energy with an external field. *SIAM J. Math. Anal. 49*, 1 (2017), 646–673.

[65] Hardin, D. P., Saff, E. B., and Vlasiuk, O. V. Generating Point Configurations via Hypersingular Riesz Energy with an External Field. *SIAM J. Math. Anal. 49*, 1 (2017), 646–673.

[66] Hardin, D. P., Saff, E. B., and Vlasiuk, O. V. Generating Point Configurations via Hypersingular Riesz Energy with an External Field. *SIAM J. Math. Anal. 49*, 1 (Jan. 2017), 646–673.

[67] Hardin, D. P., Saff, E. B., and Whitehouse, J. T. Quasi-uniformity of minimal weighted energy points on compact metric spaces. *J. Complexity 28*, 2 (2012), 177–191.

[68] Hardin, D. P., Saff, E. B., and Whitehouse, J. T. Quasi-uniformity of minimal weighted energy points on compact metric spaces. *J. Complex. 28*, 2 (2012), 177–191.

[69] Hutchinson, J. Fractals and self-similarity. *Indiana Univ. Math. J. 30* (1981), 713–747.

[70] Johnson, M. P., Sariöz, D., Bar-Noy, A., Brown, T., Verma, D., and Wu, C. W. More is more: the benefits of denser sensor deployment. *ACM Trans. Sens. Netw. 8*, 3 (2012), 22.

[71] Kuijlaars, A. B. J., and Saff, E. B. Asymptotics for minimal discrete energy on the sphere. *Trans. Amer. Math. Soc. 350*, 2 (1998), 523–538.

[72] Kuijlaars, A. B. J., and Saff, E. B. Asymptotics for minimal discrete energy on the sphere. *Trans. Amer. Math. Soc. 350*, 2 (1998), 523–538.

[73] Kuipers, L., and Niederreiter, H. *Uniform distribution of sequences*. Dover Publications, Mineola, N.Y, 2006.

[74] Landkof, N. S. *Foundations of Modern Potential Theory*. No. 180 in Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen. Springer, Berlin, 1972. OCLC: 579100.

[75] Larsson, E., Lehto, E., Heryudono, A., and Fornberg, B. Stable Computation of Differentiation Matrices and Scattered Node Stencils Based on Gaussian Radial Basis Functions. *SIAM J. Sci. Comput. 35*, 4 (2013), A2096–A2119.

[76] Lemieux, C. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer Series in Statistics. Springer New York, New York, NY, 2009.

[77] Link, W. A., and Eaton, M. J. On thinning of chains in MCMC. *Methods Ecol. Evol. 3*, 1 (2012), 112–115.

[78] Mak, S., and Roshan Joseph, V. Support points. `Arxiv:1511.02535` (2016).

[79] Martínez-Finkelshtein, A., Maymeskul, V., Rakhmanov, E. A., and Saff, E. B. Asymptotics for minimal discrete Riesz energy on curves in $\mathbb{R}^d$. *Canad. J. Math. 56*, 3 (2004), 529–552.

[80] Mattila, P. *Geometry of sets and measures in Euclidean spaces*, vol. 44 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995. Fractals and rectifiability.

[81] Mattila, P. *Geometry of sets and measures in Euclidean spaces: fractals and rectifiability*. No. 44. Cambridge University Press, 1995.

[82] Micchelli, C. A. Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constr. Approx. 2*, 1 (1986), 11–22.

[83] Mizuta, Y. *Potential Theory in Euclidean Spaces*. No. 6 in GAKUTO international series Mathematical sciences and applications. Gakkōtosho, Tokyo, 1996. OCLC: 37153855.

[84] Nguyen, V. P., Rabczuk, T., Bordas, S., and Duflot, M. Meshless methods: a review and computer implementation aspects. *Math. Comput. Simulation 79*, 3 (2008), 763–813.

[85] Ohtsuka, M. On potentials in locally compact spaces. *J. Sci. Hiroshima Univ. Ser. A-I Math. 25*, 2 (1961), 135–352.

[86] Persson, P.-O., and Strang, G. A Simple Mesh Generator in MATLAB. *SIAM Rev. 46*, 2 (2004), 329–345.

[87] Petrache, M., and Serfaty, S. Next order asymptotics and renormalized energy for Riesz interactions. *J. Inst. Math. Jussieu* (2015), 1–69.

[88] Powell, M. J. D. Five Lectures on Radial Basis Functions. *IMM Lect.*, December 2004 (2004), 27.

[89] Preparata, F. P., and Shamos, M. I. *Computational Geometry*. Springer New York, 1985.

[90] Procopiuc, O., Agarwal, P. K., Arge, L., and Vitter, J. S. Bkd-tree: A dynamic scalable kd-tree. In *Advances in Spatial and Temporal Databases*. Springer Berlin Heidelberg, 2003, pp. 46–65.

[91] Reznikov, A., Saff, E. B., and Vlasiuk, O. V. A Minimum Principle for Potentials with Application to Chebyshev Constants. *Potential Anal. 47*, 2 (2017), 235–244.

[92] Rougerie, N., and Serfaty, S. Higher-dimensional Coulomb gases and renormalized energy functionals. *Commun. Pure Appl. Math. 69*, 3 (2016), 519–605.

[93] Saff, E., and Totik, V. *Logarithmic potentials with external fields*, vol. 316. Springer Science & Business Media, 2013.

[94] Sandier, E., and Serfaty, S. *Vortices in the Magnetic Ginzburg-Landau Model*. Birkhäuser, Boston, 2007. OCLC: 458258305.

[95] Schoenberg, I. J. Metric Spaces and Completely Monotone Functions. *Ann. Math. 39*, 4 (1938), 811.

[96] Serfaty, S. *Coulomb Gases and Ginzburg–Landau Vortices*. 2015.

[97] Shankar, V., Wright, G. B., Kirby, R. M., and Fogelson, A. L. A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction–diffusion equations on surfaces. *Journal of Scientific Computing 63*, 3 (2014), 745–768.

[98] Shimada, K. *Physically-based mesh generation: Automated triangulation of surfaces and volumes via bubble packing*. ProQuest LLC, Ann Arbor, MI, 1993. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[99] Shimada, K., and Gossard, D. C. Bubble mesh. In *Proc. third ACM Symp. Solid Model. Appl. - SMA '95* (New York, New York, USA, 1995), ACM Press, pp. 409–419.

[100] Tsuji, M. *Potential Theory in Modern Function Theory*, 2d ed. ed. Chelsea Pub. Co, New York, 1975.

[101] Varma, U. M., Rao, S. V. R., and Deshpande, S. M. Point distribution generation using hierarchical data structures. In *Proc. ECCOMAS 2004* (Jyväskylä, 2004).

[102] Vlasiuk, O., and Michaels, T. Boxed lattices and Riesz minimizers for RBF computations. `https://github.com/OVlasiuk/3dRBFnodes.git`, 2017. [Online; accessed 12-October-2017].

[103] Vlasiuk, O., Michaels, T., Flyer, N., and Fornberg, B. Fast high-dimensional node generation with variable density. *ArXiv171005011 Math* (Oct. 2017).

[104] Vlasiuk, O., Michaels, T., Flyer, N., and Fornberg, B. Fast high-dimensional node generation with variable density. *Comput. Math. Appl.* (Aug. 2018).

[105] Wendland, H. *Scattered Data Approximation*. Cambridge University Press, 2004.

[106] Wu, C. W., Trager, B., Chandu, K., and Stanich, M. A Riesz energy based approach to generating dispersed dot patterns for halftoning applications. In *IS&T/SPIE Electronic Imaging* (2014), International Society for Optics and Photonics, pp. 90150Q–90150Q.

[107] Zhou, K., Hou, Q., Wang, R., and Guo, B. Real-time KD-tree construction on graphics hardware. *ACM Trans. Graph. 27*, 5 (2008), 1.

[108] Zorii, N. Equilibrium potentials with external fields. *Ukrainian Math. J. 55*, 9 (2003), 1423–1444.

[109] Zorii, N. Necessary and sufficient conditions for the solvability of the Gauss variational problem. *Ukrainian Math. J. 57*, 1 (2005), 70–99.