UNIVERSITÉ PARIS-CITÉ

Habilitation à diriger des recherches - Spécialité Mathématiques

THOMAS LEBLÉ

Les log-gas et leurs amis

Rapporteurs :

Christophe GARBAN	Université Lyon 1
Mathieu LEWIN	Université Paris-Dauphine
Mikhail SODIN	Université de Tel-Aviv

Soutenue le 10 juin 2026, devant le jury composé de :

Laure DUMAZ	Examinatrice
Christophe GARBAN	Rapporteur
Nathaël GOZLAN	Président du jury
Mathieu LEWIN	Rapporteur
Sylvia SERFATY	Examinatrice

Remerciements

Dankst du dem Mann von nebenan ?
Der ist Soldat, wenn du Goethe liest

Franz-Joseph Degenhardt, *Der Mann von nebenan.*

Un grand merci à Christophe Garban, Mathieu Lewin et Misha Sodin d'être rapporteurs de ce manuscrit. Vous êtes pour moi des modèles, et je suis profondément touché de votre considération envers mes travaux.

Merci à Laure Dumaz, ma grande sœur académique depuis quinze ans, de me faire l'honneur de participer au jury, et merci à Nathael Gozlan de le présider – nous sommes aussi voisins scientifiques ! Enfin, merci à Sylvia Serfaty de m'avoir initié au monde de la recherche et d'y être toujours pour moi un repère.

Merci à mes co-auteur.ice.s de m'apprendre des maths, et des façons d'en faire. C'est mieux à plusieurs. Un mot particulier à Ofer Zeitouni et à David Dereudre pour leur inépuisable enthousiasme. *Ein besonderer Dank gilt Martin Huesmann und dem ganzen Team in Münster für die vielen schönen Aufenthalte.*

Merci à Yuri Bakhtin, Alexei Borodin, Paul Bourgade, Djalil Chafaï, Alice Guionnet, Mylène Maïda, Nicolas Rougerie, Ofer Zeitouni pour leurs encouragements, leurs remarques, leurs conseils. Quelques mots ont parfois une grande importance.

Mes amitiés confraternelles à Jeanne Boursier, Raphaël Butez, Loren Coquille, Benoît Dagallier, Ujué Etayo, Will Feldman, Lisa Hartung, Ronan Herry, Jonas Jalowy, Gaultier Lambert, Flavien Léger, Alon Nishry, Elliot Paquette, Hugo Vanneuville pour les partages mathématiques au fil des années.

Merci à Alexis, Charles, Sylvain pour nos années New Yorkaises, et à Fréderic, Yannick, Antoine et Antoine pour les discussions scientifiques et bien davantage.

Depuis ma thèse, j'ai eu la chance de travailler dans des conditions excellentes pour lesquelles je remercie le Courant Institute, l'Institute for Advanced Study, et bien sûr le CNRS, qui offre le privilège incroyable de mener sa vie de chercheur avec une liberté totale.

Merci aux membres du MAP5 pour l'ambiance conviviale qui y règne, et notamment à Anne et Antoine pour leur rôle à la direction du département. Merci à Marie-Hélène, Gladys et Martine pour leur patience et leur soutien logistique toujours efficace.

Salut aux nombreuses professions dévalorisées dont le travail profite au monde académique.

To those for whom I wanted to seem smart.

Table des matières

Overview

In this manuscript, I present some questions that I have explored since the end of my PhD, defended in February 2016 under the supervision of Sylvia Serfaty. They are grouped in three themes:

1. *The one-dimensional log-gas and its infinite-volume limit* Sine_β. The Sine_β process was introduced by Valkó-Virág as the infinite-volume limit for the microscopic behavior, in the bulk of the spectrum, of certain important random matrix models. During my post-doc, with the help of several co-authors, we gave two "physical" characterizations of Sine_β:
   a) By proving that it is the unique minimizer of a certain free energy functional that I had introduced during my PhD (joint work with M. Erbar and M. Huesmann [EHL21], which appeared in CPAM).
   b) By proving DLR equations for Sine_β, which is the natural framework for infinite-volume Gibbs states (joint work with D. Dereudre, A. Hardy and M. Maïda [DHLM21], which appeared in CPAM).

   In this first chapter, I briefly mention a CLT for the fluctuations of smooth linear statistics of the Sine_β process that I proved in [Leb19] (published in IMRN), building on the DLR equations. It is the "infinite-volume version" of [BLS18] written with F. Bekerman and S. Serfaty and published in EJP.
2. *The two-dimensional log-gas, or one-component plasma.* This important model of statistical mechanics has been a central object of my research since the beginning of my PhD. In the past years, I made two main contributions as a single author:
   a) A proof of the long-standing conjecture concerning its hyperuniformity at all temperatures [Leb26].
   b) A preliminary study of its infinite-volume limit, through the proof of DLR equations [Leb24].
3. *The notion of hyperuniformity for general point processes*, and its connections with other "rigidity" properties, with a focus on the two-dimensional setting. With M. Huesmann, we recently clarified two aspects of hyperuniformity:
   a) In [DFHL24], with D. Dereudre et D. Flimmel, we thoroughly studied hyperuniformity of point processes that can be described as "lattices plus stationary perturbations", giving a rigorous proof of some statements found in the physics literature, and producing new (counter)-examples. The paper is under revision at Annals of Applied Probability.
   b) In [HL25], we showed a chain of implications relating hyperuniformity, Coulomb energy, and Wasserstein distance.

1. A statistical physics picture of Sine_β

1.1. The one-dimensional log-gas and its infinite-volume limit

For $N \geq 2$ and $\beta > 0$, we consider a probability measure $\mathbb{P}_N^\beta$ on $\mathbb{R}^N$ with the following density:

$$\mathtt{d}\mathbb{P}_N^\beta(x_1, \ldots, x_N) := \frac{1}{\mathrm{Z}_N^\beta} \prod_{i<j} |x_i - x_j|^\beta \exp\left(-\beta \sum_{i=1}^N N \cdot \frac{1}{2} x_i^2\right) \mathtt{d}x_1 \ldots \mathtt{d}x_N, \tag{1.1.1}$$

where Z_N^β is a normalization constant, and $\mathtt{d}x$ the Lebesgue measure. $\mathbb{P}_N^\beta$ is called a *Gaussian β-ensemble* or a *finite-volume one-dimensional log-gas*. The reason for these names is that the model has two interpretations: one coming from Random Matrix Theory (RMT), and another from statistical physics.

1. When β is equal to 1, 2 or 4, the measure $\mathbb{P}_N^\beta$ coincides with the joint law of the N random eigenvalues in certain classical models of *random matrices*, called the Gaussian Orthogonal, Unitary, or Symplectic Ensemble. Those are $N \times N$ matrices filled with Gaussian coefficients that are identically distributed (with a proper choice of the variance) and independent up to imposing a suitable symmetry to the matrix (respectively: symmetric, Hermitian, self-dual), so that the eigenvalues are all real a.s.
   Those three "Gaussian ensembles" are among the oldest and most studied models in RMT.
2. Physically, $\mathbb{P}_N^\beta$ corresponds to the canonical Gibbs measure of a system made of N particles interacting via a pairwise repulsive logarithmic potential, and confined by a quadratic external field $x \mapsto N \cdot \frac{1}{2}x^2$ whose strength scales with N. Indeed, (1.1.1) can certainly be re-written as:

$$\mathtt{d}\mathbb{P}_N^\beta(x_1, \ldots, x_N) := \frac{1}{\mathrm{Z}_N^\beta} \exp\left(-\beta \left(\sum_{1 \leq i < j \leq N} -\log|x_i - x_j| + \sum_{1 \leq i \leq N} N \cdot \frac{1}{2} x_i^2\right)\right) \mathtt{d}x_1 \ldots \mathtt{d}x_N. \tag{1.1.2}$$

   In that context, Z_N^β corresponds to the *partition function* and β is called the *inverse temperature*, while the exponential term is a *Boltzmann's factor*, of the form $\exp(-\beta \times \mathrm{Energy}(x_1, \ldots, x_N))$, the energy being given here by:

$$\mathrm{Energy}(x_1, \ldots, x_N) := \overbrace{\sum_{1 \leq i < j \leq N} -\log|x_i - x_j|}^{\text{logarithmic pairwise interaction}} + \overbrace{\sum_{1 \leq i \leq N} N \cdot \frac{1}{2} x_i^2}^{\text{confining external potential}}$$

   From a statistical physics point of view, one-dimensional log-gases are interesting toy models due to the fact that the underlying interaction is singular and, most importantly, *long-range* - one way to state this last property is to observe that the force field $\nabla \log$ is not integrable at infinity on $\mathbb{R}$.

The connection between random eigenvalues and particles of a one-dimensional system with logarithmic interactions for $\beta = 1, 2, 4$ has first been pointed out and investigated by Dyson in a famous series of papers [Dys62] and brought up again in [DMPS95], see also Forrester [For10] for a comprehensive account of this topic. However, it took several decades for people to realize that this connection actually extends to all values of β - indeed, for every $\beta > 0$, there exists a model of random matrices with independent entries, known as the *tridiagonal model for the Gaussian β-Ensemble*, discovered in [DE02], whose eigenvalues have a joint law given by $\mathbb{P}_N^\beta$. For all $\beta > 0$, and *in finite volume*, we thus have a correspondence:

$$\boxed{\text{Random eigenvalues} \longleftrightarrow \text{Particles (of a one-dimensional log-gas)}} \tag{1.1.3}$$

The goal of this chapter is to explore the correspondence (1.1.3) in the infinite-volume ($N \to \infty$) limit.

The circular ensemble and the periodic log-gas

When looking for a matrix model with "one-dimensional" eigenvalues, one could also think of unitary matrices, whose eigenvalues lie on the unit circle. The set of unitary matrices possesses a natural probability measure, which historically speaking seems to be the first "random matrix" model ever considered, namely *the Haar measure on the unitary group* of dimension N. Dyson observes that the distribution of eigenvalues *can again be computed explicitely* and is given by:

$$\mathtt{d}\mathrm{P}_N^{circ}(z_1,\dots,z_N) := \frac{1}{\mathrm{Z}'_N}\prod_{i<j}|z_i-z_j|^2\,\mathtt{d}z_1\dots\mathtt{d}z_N, \tag{1.1.4}$$

where the reference measure $\mathtt{d}z_i$ for each eigenvalue is the Lebesgue measure on the unit circle. This is of course reminiscent of (1.1.1), without the confining potential, but with the hard spatial constraint of living on the unit circle. This model is called the Circular Unitary Ensemble (CUE).

In the same series of papers, Dyson introduces two other "circular" matrix models, with eigenvalues on the unit circle, which are related to the orthogonal group and symplectic group although they do not correspond to the Haar measure on those groups. Dyson computes the eigenvalue distribution for both models, which look again like (1.1.4), but with the power 2 replaced by 1 (Circular Orthogonal Ensemble) or 4 (Circular Symplectic Ensemble). This naturally leads to introducing the "general β" version:

$$\mathtt{d}\mathrm{P}_{N,\beta}^{circ}(z_1,\dots,z_N) := \frac{1}{\mathrm{Z}'_{N,\beta}}\prod_{i<j}|z_i-z_j|^\beta\,\mathtt{d}z_1\dots\mathtt{d}z_N, \tag{1.1.5}$$

known as the *Circular β-Ensemble*. From the point of view of statistical physics, $\mathrm{P}_{N,\beta}^{circ}$ can be written as:

$$\mathtt{d}\mathrm{P}_{N,\beta}^{circ}(z_1,\dots,z_N) := \frac{1}{\mathrm{Z}'_{N,\beta}}\exp\left(-\beta\sum_{i<j}-\log|z_i-z_j|\right)\mathtt{d}z_1\dots\mathtt{d}z_N,$$

which corresponds to the Gibbs measure of a "periodic log-gas" i.e. a finite log-gas with periodic boundary conditions.

Moreover, for all $\beta>0$, there is a *matrix model* extending the ones known for $\beta=1,2,4$ and whose eigenvalue distribution matches $\mathrm{P}_{N,\beta}^{circ}$, thus establishing again the connection (1.1.3) between "random eigenvalues" and "one-dimensional log-gases" in finite volume.

The infinite-volume limit

A basic result in the theory of log-gases is that when N is very large, the eigenvalues of a β-ensemble (or equivalently, the particles of a $1d$ log-gas) tend to arrange themselves on, and close to, the interval $[-2,2]$, with a typical macroscopic density approaching the famous *semicircle law* $x\mapsto\frac{1}{2\pi}\sqrt{4-x^2}$ of Wigner. The nearest-neighbor distance between them is thus of order N^{-1}. Hence, in order to study the microscopic behavior, one should "zoom in" by a factor N. For any fixed $\mathtt{x}\in\mathbb{R}$, we define the random point configuration $\mathbf{X}_{N,\mathtt{x}}$ by:

$$\mathbf{X}_{N,\mathtt{x}} := \sum_{i=1}^{N}\delta_{N(x_i-\mathtt{x})}. \tag{1.1.6}$$

This quantity encodes the local point configuration of the eigenvalues / particles as seen around $\mathtt{x}$. We let Config be the set of point configurations, finite or infinite, endowed with a local topology for which two configurations are close if they "look the same on a big interval". In particular, this allows for a sequence of finite configurations to converge to an infinite one. The set $\mathcal{P}(\mathrm{Config})$ of probability measures on Config is then endowed with the topology of weak convergence.

Proving the existence of a limit to the law of $\mathbf{X}_{N,\mathtt{x}}$ as $N\to\infty$ is a difficult question. Indeed:

— In any given interval, the configuration $\mathbf{X}_{N,\mathtt{x}}$ could a priori have anywhere between 0 and N points. Controlling this requires to understand the system *down to the microscopic scale*.

— Proving that any fixed interval contains in fact $\mathcal{O}(1)$ points *on average* would give some compactness, but would still not ensure the existence of an actual limit - and not just of limit points.

Existence of Sine_β

For the "classical" cases corresponding to β equal to 1, 2 or 4, the measure $\mathbb{P}_N^\beta$ enjoys additional properties of algebraic nature, making its analysis amenable to explicit computations. Most importantly, when $\beta = 2$, the finite volume point process is found to be *determinantal*, which means that all the k-point correlation functions of $\mathbb{P}_N^\beta$ can be expressed as a $k \times k$ determinant involving a certain explicit and fairly simple *kernel*, namely some function $\mathbb{R} \to \mathbb{R}$ which depends on N. In that case, proving existence of a limit for the point process $\mathbf{X}_{N,\mathsf{x}}$ as $N \to \infty$ can be reduced to finding a limit for the *kernel* itself, "zoomed" around the point x, which turns out to be a fairly simple analytic task.

If $\mathsf{x} \in (-2, 2)$, the limiting kernel is proportional to $x \mapsto \frac{\sin(x)}{x}$, and the limiting process was thus named [1] the *Sine-process*, which we denote by Sine_2 to recall that it comes from $\beta = 2$. In the cases $\beta = 1, 4$, the "determinantal" feature is replaced by a "Pfaffian" one, which is slightly less convenient. Nonetheless, proving existence of a limit process still boils down to a reasonably explicit computation. By analogy with the case $\beta = 2$, the two corresponding processes are called Sine_1 and Sine_4.

However, for $\beta \notin \{1, 2, 4\}$ (which we will refer to as the "general β" case), the "determinantal" or "Pfaffian" techniques fall apart completely and we are left with the two difficulties mentioned above. The first situation that has been understood for general β is the so-called *edge behavior*, obtained when zooming around the edge of the system by taking $\mathsf{x} = 2$. The limiting process [RRV11] is called [2] Airy_β.

Concerning the *bulk*, when $\mathsf{x} \in (-2, 2)$, Valkó-Virág proved, in a major breakthrough [VV09], that:

Theorem (Existence of Sine_β)**.** The local point configuration $\mathbf{X}_{N,\mathsf{x}}$ does converge in distribution as $N \to \infty$ to the law of an infinite, translation-invariant, non-empty point process called Sine_β.

This limiting point process is not only shown to *exist*, but is also characterized in an explicit enough way that one can say something meaningful about it, for instance:

— A central limit theorem for the number of points in large intervals [KVV12].

— Probability estimates for "overcrowding" of finite intervals by "way too many points" [HV17].

— How "close" Sine_β and $\mathrm{Sine}_{\beta'}$ are when $|\beta - \beta'| \ll 1$ [VV20].

The description of Sine_β introduced in [VV09] starts by viewing the finite-volume measure $\mathbb{P}_N^\beta$ as the distribution of eigenvalues for the random tridiagonal matrix discoved in [DE02]. They then go through an involved argument that eventually leads them to a one-parameter family of coupled SDE's indexed by $\lambda \in \mathbb{R}$, of the form: $\mathtt{d}\alpha_\lambda = \lambda\beta e^{-\beta t}\mathtt{d}t + \mathrm{Re}\left(\left(e^{-i\alpha_\lambda} - 1\right)\mathtt{d}Z_t\right)$, where $(Z_t)_t$ is a complex Brownian motion. One then gets a point process by thinking of $\frac{1}{2\pi}\alpha_\lambda(\infty)$ as giving the number of points in $[0, \lambda]$.

For the circular models (1.1.5), one would like to "zoom in" around any point on the unit circle and study the microscopic process there. In this case, the infinite-volume limit was derived in [KS09]. It is interesting, and not obvious, to know that [Nak14] *the infinite-volume limit of the Circular β-Ensemble coincides with the infinite-volume (bulk) limit of the Gaussian β-Ensemble.* One significant upside of working with the Circular ensemble rather than the Gaussian ensemble as a finite-volume model giving rise to Sine_β is that the law of $\mathrm{P}_{N,\beta}^{circ}$ has a built-in rotational invariance. As $N \to \infty$, this becomes a translation-invariance for the limiting process. Note that translation-invariance of Sine_β is a consequence of the work of Valkó-Virág, but is in no way obvious when thinking of it as the microscopic limit of a Gaussian β-ensemble, because the measure $\mathbb{P}_N^\beta$ (1.1.1) is *not* translation-invariant.

Other descriptions of Sine_β have been given later - as the spectrum of a certain random operator [VV17], or as the zeroes of a random analytic function [VV22]. None of those characterizations use the "physical" interpretation of $\mathbb{P}_N^\beta$ as the Gibbs measure of a one-dimensional system with logarithmic interactions.

In the next two sections, we retrieve a "physical" interpretation for the Sine_β process:

— By characterizing Sine_β as the unique minimizer of a certain free energy functional.

— By showing that Sine_β satisfies Dobrushin-Lanford-Ruelle (DLR) equations.

1. One could argue that a more appropriate choice would have been to name it the *Sinc*-process.
2. Because for $\beta = 2$ one gets a determinantal point process whose kernel is an Airy function

1.2. Uniqueness of the free energy minimizer

Start in finite volume and consider an arbitrary "energy" functional $H_N : \mathbb{R}^N \to \mathbb{R}$, which we assume for simplicity to be smooth and compactly supported, so that the integral defining the partition function

$$\widetilde{Z}_{N,\beta} := \int_{\mathbb{R}^N} \exp\left(-\beta H_N(x_1, \dots, x_N)\right) \mathtt{d}x_1 \dots \mathtt{d}x_N$$

is convergent, and that the Gibbs density associated to the energy H_N at inverse temperature $\beta > 0$,

$$\mathtt{d}\mathrm{Gibbs}^{H_N}_{N,\beta}(x_1, \dots, x_N) := \frac{1}{\widetilde{Z}_{N,\beta}} \exp\left(-\beta H_N(x_1, \dots, x_N)\right) \mathtt{d}x_1 \dots \mathtt{d}x_N,$$

is well-defined. Recall that the entropy of a probability measure μ on $\mathbb{R}^N$ is defined as:

$$\mathsf{Ent}[\mu] := \int_{\mathbb{R}^N} \log \frac{\mathtt{d}\mu}{\mathtt{d}x_1 \dots \mathtt{d}x_N} \mathtt{d}\mu$$

when the Radon-Nykodym derivative exists, and is set to $\mathsf{Ent}[\mu] = +\infty$ otherwise. Finally, let us introduce the "average energy" of such a measure μ as being simply given by:

$$\mathsf{H}_N(\mu) := \int_{\mathbb{R}^N} H_N(x_1, \dots, x_N) \mathtt{d}\mu. \tag{1.2.1}$$

Then the following holds:

Fact 1.2.1. *The Gibbs measure* $\mathrm{Gibbs}^{H_N}_{N,\beta}$ *is the* unique *minimizer of the "free energy" functional*

$$\mu \mapsto \beta \mathsf{H}_N(\mu) + \mathsf{Ent}[\mu].$$

The goal of this section is to find an analogous statement in the infinite-volume limit, identifying the "infinite-volume Gibbs measure" Sine_β as the unique minimizer of an "infinite-volume free energy".

Note that $\mu \mapsto \mathsf{H}_N(\mu)$ is linear in μ and $\mu \mapsto \mathsf{Ent}[\mu]$ is strictly convex, hence the free energy itself is strictly convex. The central question will be to find a suitable infinite-volume convexity argument.

The infinite-volume free energy functional

In [LS17], written with S. Serfaty during my PhD under her supervision, we studied large deviations for the microscopic behavior of one-dimensional log-gases. In a nutshell, our approach was the following:

1. Instead of looking at the point configuration $\mathbf{X}_{N,\mathsf{x}}$ seen at microscopic scale near a given point x (1.1.6), study the microscopic behavior *averaged over* x and, to this aim, consider an observable known in the mathematical statistical mechanics literature as the *empirical field* $\mathbf{Emp}_N$, or rather its "tagged" version $\overline{\mathbf{Emp}}_N$. One way to define those is to first choose x uniformly at random in the region of space where we know that the particles concentrate (here the support of the semicircle law, namely the interval $[-2, 2]$) and *then* look at $\mathbf{X}_{N,\mathsf{x}}$. Formally, this means studying the following random variable, with values in the space of probability measures on $[-2, 2] \times \mathrm{Config}$:

$$\overline{\mathbf{Emp}}_N := \fint_{[-2,2]} \delta_{(\mathsf{x}, \mathbf{X}_{N,\mathsf{x}})} \mathtt{d}\mathsf{x}.$$

   $\overline{\mathbf{Emp}}_N$ is the *tagged empirical field*, the "tag" referring to the first coordinate, which keeps track of the point around which we zoomed in. Despite the addition of a tag, the observable $\overline{\mathbf{Emp}}_N$ remains an average quantity, which is not quite as precise as the understanding of each $\mathbf{X}_{N,\mathsf{x}}$ separately.
2. Prove a large deviation principle for $\overline{\mathbf{Emp}}_N$. This means that we introduced a *rate function*, denoted by $\overline{\mathcal{F}}_\beta$, and shown that the probability of observing $\overline{\mathbf{Emp}}_N$ near a certain probability measure $\overline{\mathsf{P}}$ was, roughly speaking, given by $\exp\left(-N\overline{\mathcal{F}}_\beta(\overline{\mathsf{P}})\right)$. The processes $\overline{\mathsf{P}}$ for which $\overline{\mathcal{F}}_\beta$ is big represent very unlikely behaviors for $\overline{\mathbf{Emp}}_N$, and, in contrast, *the typical behavior of* $\overline{\mathbf{Emp}}_N$ *as* $N \to \infty$ *is to be close to minimizers of* $\overline{\mathcal{F}}_\beta$.

It turns out that, by a fairly simple scaling argument, the tag does not play an important role, and that it is enough to study a certain function $\mathcal{F}_\beta$ on the space $\mathcal{P}(\mathrm{Config})$, whose minimizers encode the typical average microscopic behavior of a $1d$ log-gas (up to scaling), and which is of the form:

$$\mathcal{F}_\beta : \mathsf{P} \mapsto \beta \mathcal{W}(\mathsf{P}) + \mathcal{E}(\mathsf{P}), \tag{1.2.2}$$

where $\mathcal{W}$ is an "energy" term and $\mathcal{E}$ is an "entropy" term, making $\mathcal{F}_\beta$ a "free energy" functional in the language of statistical physics. More precisely:

— $\mathcal{E}(\mathsf{P})$ is the specific relative entropy, which is well-known in the context of infinite-volume statistical mechanics. Its definition is simple: take $R \gg 1$, consider the restrictions P_R, Π_R of P and of a Poisson point process Π to the interval $[-R, R]$, and define:

$$\mathcal{E}(\mathsf{P}) := \lim_{R\to\infty} \frac{1}{2R} \mathsf{Ent}\left[\mathsf{P}_R | \Pi_R\right].$$

If P is translation-invariant, the limit is shown to exist by a fairly simple super-additivity argument.

— $\mathcal{W}(\mathsf{P})$ is obtained by computing an expectation $\mathbb{E}_\mathsf{P}\left[\mathbb{W}(\mathbf{X})\right]$ under P, cf. (1.2.1). The quantity $\mathbb{W}(\mathbf{X})$ is the "renormalized energy" of a point configuration $\mathbf{X}$ as defined by Sandier-Serfaty in [SS12, SS15], see Appendix A. In short, this is the "right" way to define a logarithmic interaction energy for infinitely many particles on the real line.

Connection between Sine_β and $\mathcal{F}_\beta$

1. One the one hand, thanks to [VV09], we know that the point process $\mathbf{X}_{N,\mathsf{x}}$ converges to Sine_β for every $\mathsf{x} \in (-2, 2)$ up to a scaling depending on x.
2. On the other hand, thanks to the large deviation principle of [LS17], we know that the microscopic behavior *averaged over* $\mathsf{x} \in (-2, 2)$ concentrates (after rescaling) on minimizers of $\mathcal{F}_\beta$.

It is then not difficult to combine those statements and deduce that:

Corollary 1.2.2 ([LS17])**.** *The point process* Sine_β *is a minimizer of the free energy functional* $\mathcal{F}_\beta$*.*

In view of Fact 1.2.1, it is tempting to ask whether we can go further:

Is Sine_β *the* unique *minimizer of the free energy* $\mathcal{F}_\beta$*?*

Note that in infinite-volume one should not expect free energy minimizers to always be unique - at certain values of the inverse temperature, corresponding to when the system undergoes a *phase transition*, it might happen that there are two, more, or even infinitely many distinct minimizers. However, it is a rule of thumb that one-dimensional systems "rarely" undergo phase transitions.

Uniqueness of the minimizer would follow easily once we know that $\mathcal{F}_\beta$ is *strictly* convex. Unfortunately, both terms involved in the definition (1.2.2) of $\mathcal{F}_\beta$ are convex... but not strictly so! In fact, they are both linear in P.

For quite some time, this seemed like a serious obstacle, and I am very grateful to Alice Guionnet for her suggestion to have a look at the notion of "displacement convexity".

Displacement convexity as a tool to show uniqueness

Displacement convexity was introduced by McCann in [McC97]. The basic idea consists in changing the very notion of "interpolation" between two probability measures: for example, instead of considering that the middle point between two Dirac masses δ_0 and δ_1 is their mixture $\frac{1}{2}(\delta_0 + \delta_1)$, why not think of it as being a Dirac masss placed "in between", namely $\delta_{\frac{1}{2}}$? It means that we are switching from the usual linear structure on the space of measures to an interpolation induced by *transportation* maps between those probability measures.

We say that a map $T : \mathbb{R}^N \to \mathbb{R}^N$ *transports* μ_0 onto μ_1 when the pushforward $T\#\mu_0$ of μ_0 by T is μ_1. For such a T, and for $s \in [0, 1]$ we let $\mu_s := ((1 - s)\,\mathrm{Id} + sT)\#\mu_0$, this defines a one-parameter family of probability measures which clearly interpolates between μ_0 and μ_1. McCann then takes the natural next step in the form of a definition:

Definition 1.2.3 (Displacement convex functionals)**.** A functional F on the space of probability measures is said to be *displacement convex* when, for each μ_0, μ_1, the map $s \mapsto F(\mu_s)$ is convex on $[0,1]$.

Note that, technically, this definition depends on the choice of the transportation map T. The following table compares the convexity of some "typical" functionals for both notions of interpolation:

Functional	Usual interpolation	Displacement interpolation
$\mu \mapsto \int U(x)\mathtt{d}\mu(x)$	Linear	As convex as U
$\mu \mapsto \iint W(x-y)\mathtt{d}\mu(x)\mathtt{d}\mu(y)$, W convex	Not convex in general	Convex
Relative entropy	Strictly convex	Convex *if T is optimal*

Displacement convexity in the first two entries follows from a simple computation, which works for any choice of the transport T. The last entry is significantly more subtle, and requires that T be *optimal* for the quadratic cost, i.e. minimizes $\int |T(x)-x|^2 \mathtt{d}\mu_0(x)$ among possible transports from μ_0 to μ_1.

As an example of how things work, given μ_0 and μ_1, choosing a transport T from μ_0 to μ_1, and letting $\mu_{\frac{1}{2}}$ be the push-forward of μ_0 by $x \mapsto \frac{x+T(x)}{2}$, we can compare the interaction energy $\iint W(x-y)\mathtt{d}\mu(x)\mathtt{d}\mu(y)$ evaluated for $\mu = \mu_{\frac{1}{2}}$ to its value at μ_0 and μ_1 as follows:

$$\begin{aligned}
\iint W(x-y)\mathtt{d}\mu_{\frac{1}{2}}(x)\mathtt{d}\mu_{\frac{1}{2}}(y) &= \iint W\left(\frac{x+T(x)}{2} - \frac{y+T(y)}{2}\right)\mathtt{d}\mu_0(x)\mathtt{d}\mu_0(y) \\
&= \iint W\left(\frac{1}{2}(x-y) + \frac{1}{2}(T(x)-T(y))\right)\mathtt{d}\mu_0(x)\mathtt{d}\mu_0(y) \\
\leq \iint \frac{1}{2}\left(W(x-y) + W(T(x)-T(y))\right) &- \frac{1}{8}\left((x-y)-(T(x)-T(y))\right)^2 \inf_{[x-y,T(x)-T(y)]} W''\mathtt{d}\mu_0(x)\mathtt{d}\mu_0(y) \\
&\leq \frac{1}{2}\iint W(x-y)\,\mathtt{d}\mu_0(x)\mathtt{d}\mu_0(y) + \frac{1}{2}\iint W(x-y)\,\mathtt{d}\mu_1(x)\mathtt{d}\mu_1(y) \\
&\quad - \frac{1}{8}\iint \left((x-y)-(T(x)-T(y))\right)^2 \inf_{[x-y,T(x)-T(y)]} W''\mathtt{d}\mu_0(x)\mathtt{d}\mu_0(y). \quad (1.2.3)
\end{aligned}$$

The expression of the convexity "gain" on the very last line already illustrates two important technical points. First, if T is a translation, then there is no gain whatsoever. Second, if W'' is positive but tends to 0 at infinity, then in the last double integral we need to focus on a neighborhood of the diagonal, where $x-y$ and $T(x)-T(y)$ might not be too big, so that $\inf_{[x-y,T(x)-T(y)]} W''$ gets a good positive lower bound.

The lesson here for us is that by changing the point of view on what "interpolating between two measures" means, we can gain convexity for "interaction" functionals (second line) while preserving the convexity of "entropy" functionals (last line). Our goal thus become:

Goal: prove that $\mathcal{F}_\beta$ is strictly "displacement convex".

There are several obstacles:

— One cannot use a ready-made theory of optimal transportation between point processes. Those are probability measures on infinite-dimensional spaces, and the usual framework simply does not apply.

— One needs to ensure that any strict convexity gain that might appear for the energy functional thanks to this new point of view will survive in the infinite-volume limit, where we always *divide by the volume*!

In short, the solution is:

1. To start with infinite-volume objects.
2. To reduce them to a large but finite box, on which usual transportation theory applies.
3. To show some strict convexity there and to control its dependency with respect to the size of the box.
4. To finally return to the infinite-volume setting.

With Matthias Erbar and Martin Huesmann, we managed to implement this program, and showed:

Theorem 1 ([EHL21])**.** $\mathcal{F}_\beta$ has a unique minimizer among stationary point processes, which is Sine_β.

Sketch of the proof.

The proof goes by contradiction. We start with two stationary point processes $\mathsf{P}^0, \mathsf{P}^1$ such that $\mathsf{P}^0 \neq \mathsf{P}^1$, and assume that both are minimisers of $\mathcal{F}_\beta$. We cannot argue via displacement convexity directly on the level of $\mathsf{P}^0, \mathsf{P}^1$ since optimal transport theory for random stationary measures as initiated in [HS13, Hue16, EH15] was at the time not yet developed well enough to be directly applicable. Since then, there has been important progress [EHJM25], but that still would not allow for a direct infinite-volume argument.

Instead, we use transportation theory between *finite* measures in "large boxes", together with a careful approximation argument relying on a technical tool called "screening of electric fields". After all, the free energy functional itself is, by definition, a "large box limit", indeed:

$$\mathcal{F}_\beta(\mathsf{P}) = \beta\mathcal{W}(\mathsf{P}) + \mathcal{E}(\mathsf{P}) = \lim_{R\to\infty} \frac{1}{2R} \left(\beta\mathcal{W}_R(\mathsf{P}) + \mathcal{E}_R(\mathsf{P})\right),$$

where $\mathcal{W}_R, \mathcal{E}_R$ are quantities depending on the restriction of P to the line segment $[-R, R]$.

1. Let $\mathsf{P}^0_R, \mathsf{P}^1_R$ be the restriction of $\mathsf{P}^0, \mathsf{P}^1$ to Λ_R. Assume that there are almost surely $2R$ points in $[-R, R]$. We may thus see $\mathsf{P}^0_R, \mathsf{P}^1_R$ as probability measures on $[-R, R]^{2R}$, apply classical results about optimal transportation, and find an optimal transport map T_R which pushes P^0_R onto P^1_R.

   In fact, it is not true that $\mathsf{P}^0, \mathsf{P}^1$ have almost surely $2R$ points in Λ_R, so we first perform a version of the "screening procedure" of Sandier-Serfaty. This highly technical procedure has for effect to modify the configurations near the extremities of $[-R, R]$ in order to enforce the correct number of points, while not changing too much the average energy, nor the entropy. Screening requires some conditions in order to be applied, we need to guarantee that they are verified with high probability under $\mathsf{P}^0, \mathsf{P}^1$.

2. We let $\mathsf{P}^{\mathsf{h}}_R$ be the half-interpolate of $\mathsf{P}^0_R, \mathsf{P}^1_R$ along the displacement T_R, i.e. the push-forward of P^0_R by $\frac{1}{2}(\mathrm{Id} + \mathsf{T}_R)$. As recalled above, [McC97] guarantees that the relative entropy is displacement convex:

   $$\mathcal{E}_R[\mathsf{P}^{\mathsf{h}}_R] \leq \frac{1}{2}\left(\mathcal{E}_R(\mathsf{P}^0) + \mathcal{E}_R(\mathsf{P}^1)\right).$$

   At this stage, the convexity is not strict - this would be equivalent to having a form of "positive Ricci curvature" on the space of point configurations, which is not to be expected.

   Nevertheless, the interaction potential $-\log$ is strictly convex on $\mathbb{R}^*_+$, so the interaction energy $\mathcal{W}_R$ (which encodes the logarithmic interaction of the particles) is also displacement convex. A computation similar to (1.2.3) yields:

   $$\mathcal{W}_R[\mathsf{P}^{\mathsf{h}}_R] \leq \frac{1}{2}\left(\mathcal{W}_R(\mathsf{P}^0) + \mathcal{W}_R(\mathsf{P}^1)\right) - \mathrm{Gain}_R,$$

   where $\mathrm{Gain}_R > 0$ is some quantitative positive gain due to the *strict* convexity of the interaction.

   A significant part of the work, using the fact that $\mathsf{P}^0, \mathsf{P}^1$ are stationary in a crucial way, consists in showing that the gain is at least proportional to R, namely $\mathrm{Gain}_R \geq \mathfrak{g}R$ for some $\mathfrak{g} > 0$.

   So far, we have thus constructed a point process $\mathsf{P}^{\mathsf{h}}_R$ within $[-R, R]$, such that:

   $$\beta\mathcal{W}_R(\mathsf{P}^{\mathsf{h}}_R) + \mathcal{E}_R(\mathsf{P}^{\mathsf{h}}_R) \leq \frac{1}{2}\left(\beta\mathcal{W}_R(\mathsf{P}^0_R) + \mathcal{E}_R(\mathsf{P}^0_R) + \beta\mathcal{W}_R(\mathsf{P}^1_R) + \mathcal{E}_R(\mathsf{P}^1_R)\right) - \mathfrak{g} \times R.$$

3. We turn $\mathsf{P}^{\mathsf{h}}_R$ into a process on the full line by pasting independent copies of itself on disjoint copies of $[-R, R]$, followed by an "averaging" step to make it stationary. The relative entropy is additive with respect to these independent copies, and the interaction between two independent copies is almost zero. Since $\mathfrak{g} > 0$, we obtain a global candidate P^{h} such that:

   $$\mathcal{F}_\beta(\mathsf{P}^{\mathsf{h}}) < \frac{1}{2}\left(\mathcal{F}_\beta(\mathsf{P}^0) + \mathcal{F}_\beta(\mathsf{P}^1)\right) = \min \mathcal{F}_\beta.$$

   which is absurd. Hence $\mathcal{F}_\beta$ must have a unique minimizer.

1.3. DLR equations and number-rigidity

As mentioned ealier, the density $\mathbb{P}_N^\beta$ (1.1.1) has a physical interpretation because, when written in the form (1.1.2), it presents itself as a Gibbs density, i.e. is of the form:

$$\mathsf{d}\mathbb{P}(\mathbf{X}) \propto \exp\left(-\beta \mathsf{Energy}(\mathbf{X})\right) \mathsf{dVol}(\mathbf{X}), \tag{1.3.1}$$

where $\mathbf{X}$ is the state of the system (here a N-tuple of points in $\mathbb{R}$), Energy is some "energy" functional (here the logarithmic interaction of the particles plus a confining potential), and Vol is a reference "volume" measure (here the Lebesgue measure on $\mathbb{R}^N$).

When proving the existence of an infinite-volume limit, Valkó & Virág did not rely on this point of view, and their description of Sine_β is therefore not "Gibbsian". They leave a naive yet natural question open:

Does Sine_β *admit a representation of the form* (1.3.1)*?*

It is a *natural* question to ask, all the more since we do know good infinite-volume candidates for the "energy" term, namely the renormalized energy of Sandier-Serfaty, as well as for the reference measure, whose role should probably be played by a Poisson point process since it represents "infinitely many independent points". It is, however, also a *naive* question, because of the following basic fact:

Fact 1.3.1. *The only stationary point process that has a density with respect to the Poisson point process is... the Poisson point process itself.*

We thus need to re-formulate our question. The "right" way to ask it turns out to be:

Does Sine_β *satisfy Dobrushin-Lanford-Ruelle equations?*

We tackled this question with D. Dereudre, A. Hardy and M. Maïda in [DHLM21].

Making sense of DLR equations: the "move" functions

The Dobrushin-Lanford-Ruelle (DLR) formalism has been known for a long time as the one really suited to describe infinite-volume processes[3] P which are "Gibbsian". It consists in finding a Gibbs measure representation not for the whole density of P (which is hopeless because of Fact 1.3.1) but for the local densities $\mathsf{P}\left(\mathbf{X}_\Lambda | \mathbf{X}_{\overline{\Lambda}}\right)$ in each compact set Λ *given the configuration outside* Λ. The goal is then to write:

$$\mathsf{dP}\left(\mathbf{X}'_\Lambda | \mathbf{X}_{\overline{\Lambda}}\right) \propto \exp\left(-\beta\left(\mathsf{Energy}(\mathbf{X}'_\Lambda) + \mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}})\right)\right) \mathsf{dVol}_\Lambda(\mathbf{X}'_\Lambda), \tag{1.3.2}$$

with a good choice of the "energy" and "volume" terms. In general, the self-interaction energy term $\mathsf{Energy}(\mathbf{X}'_\Lambda)$ of the configuration within Λ is not very mysterious - since Λ is compact, this should be similar to the finite-volume interaction energy occuring in the very definition of our finite-volume model. What can be more puzzling is the meaning that one should give to the interaction energy between $\mathbf{X}'_\Lambda$, which is a finite configuration within the compact set Λ, and the "outside" configuration $\mathbf{X}_{\overline{\Lambda}}$, *which is infinite*. When the interaction potential has a short-range (think of spin systems with "nearest-neighbor" interactions, like the Ising model) this is of course not of problem, because $\mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}})$ will only involve finitely many terms. Here, however, we are dealing with the logarithmic potential which is *long-range*.

In order to give a meaning to DLR equations for the one-dimensional log-gas, one thus needs to make sense of "the logarithmic interaction energy between $\mathbf{X}'_\Lambda$ and $\mathbf{X}_{\overline{\Lambda}}$", the main obstacle being the long-range nature of the pair potential. Technically, it boils down to giving a sense to integrals of the form:

$$\mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}}) \overset{??}{:=} \iint_{\Lambda \times \overline{\Lambda}} -\log|x-y| \mathtt{d}\mathbf{X}'_\Lambda(x) \mathtt{d}\left(\mathbf{X}_{\overline{\Lambda}} - \mathrm{Leb}\right)(y). \tag{1.3.3}$$

The Lebesgue measure appearing here is an heritage of the confining potential: indeed, one might set

$$V_\Lambda := x \mapsto \int_{\overline{\Lambda}} \log|x-y| \mathtt{d}\mathrm{Leb}(y),$$

3. Historically, a lot of this theory have been devoted to *lattice spin systems*, i.e. models in which the state of the system is given by a random "spin configuration", namely a map from the lattice (e.g. $\mathbb{Z}^d$) to a certain "spin space" (e.g. $\{\pm 1\}$). There is however a sizeable literature on DLR equations for point processes.

and view V_Λ as some kind of external potential felt by the particles within Λ. Admitting that (1.3.3) is indeed the right quantity to define, its convergence is still not obvious. The following two things will help:

— In the context of Gibbs measures, interaction energies only need to be defined "up to a constant", and morally speaking this constant *might be infinite*. In other words, what matters is not really to give a meaning to the interaction energy $\mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}})$ but to the difference

$$\mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}}) - \mathsf{Energy}(\mathbf{X}^{\mathrm{ref}}_\Lambda, \mathbf{X}_{\overline{\Lambda}}),$$

where $\mathbf{X}^{\mathrm{ref}}_\Lambda$ is a "reference" point configuration withing Λ, or alternatively to

$$\mathsf{M}_\Lambda(\mathbf{X}, \mathbf{X}') := \mathsf{Energy}(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}}) - \mathsf{Energy}(\mathbf{X}_\Lambda, \mathbf{X}_{\overline{\Lambda}}). \tag{1.3.4}$$

We call M_Λ a "move function" - it represents the energy cost of moving the points from one position to another within Λ, while feeling the effect of a fixed exterior configuration.

— We do not need to consider all possible exterior configurations $\mathbf{X}_{\overline{\Lambda}}$, but only the ones that are "typical" under Sine_β. We can then use our knowledge on how "regular" a typical sample of Sine_β is and how much it "resembles" the Lebesgue measure. At a technical level, it means we can import the quantitative understanding of the cancellations between $\mathbf{X}_{\overline{\Lambda}}$ and Leb in (1.3.3) as developed by Serfaty and co-authors, for example through so-called "discrepancy estimates", which control the difference between the (second moment of the) number of points in a fixed interval and the length of that interval.

Our first result is the existence of "Move functions" $\mathsf{M}_\Lambda(\mathbf{X}, \mathbf{X}')$ in the following sense:

Theorem 2 ([DHLM21])**.** For all compact sets $\Lambda \subset \mathbb{R}$, for Sine_β-almost every point configuration $\mathbf{X}$ and for all point configurations $\mathbf{X}'$ *such that* $\mathbf{X}'$ *and* $\mathbf{X}$ *have the same number of points in* Λ, the quantity $\mathsf{M}_\Lambda(\mathbf{X}, \mathbf{X}')$ is well-defined.

Note that we require $\mathbf{X}$ and $\mathbf{X}'$ to have the same number of points, an assumption that is used to push a certain Taylor expansion until the point where log has received enough derivatives to be integrable. Once we know how to give a sense to Move functions, it is not too difficult to obtain:

Theorem 3 ([DHLM21])**.** Sine_β satisfies *canonical* DLR equations for Sine_β, namely something similar to (1.3.2), but with the additional question "how many points in Λ?":

$$\begin{aligned}&\mathsf{dP}\left(\mathbf{X}'_\Lambda | \mathbf{X}_{\overline{\Lambda}},\ \mathbf{X} \text{ and } \mathbf{X}' \text{ have the same number of points in } \Lambda.\right)\\ &\propto \exp\left(-\beta\left(\mathsf{Energy}(\mathbf{X}'_\Lambda) + \mathsf{M}_\Lambda(\mathbf{X}'_\Lambda, \mathbf{X}_{\overline{\Lambda}})\right)\right) \mathsf{d}\Pi_\Lambda\left(\mathbf{X}' |\ \mathbf{X} \text{ and } \mathbf{X}' \text{ have the same number of points in } \Lambda.\right).\end{aligned}$$

Here Π_Λ is a Poisson point process in Λ, which we condition to having a certain number of points.

This formulation is slightly unusual (most DLR equations in the literature use a grand-canonical formalism, with an arbitrary number of points inside Λ), but is in fact the natural one, as we now explain.

Number-rigidity

The notion of *number-rigidity* was introduced in [GP17].

Definition 1.3.2 (Number-rigidity)**.** A random point configuration $\mathbf{X}^\bullet$ is *number-rigid* when for every compact set Λ, the number of points of the configuration $\mathbf{X}^\bullet$ in Λ is almost surely equal to a certain deterministic function of the configuration $\mathbf{X}^\bullet$ *outside* Λ.

The Sine_2 process was known to be number-rigid [BS17], as well as Sine_1 and Sine_4 [BNQ19]. We prove:

Theorem 4 ([DHLM21])**.** For all $\beta > 0$, the point process Sine_β is number-rigid.

The proof uses our canonical DLR equations and an abstract stochastic geometry argument. A different proof, much closer to the spirit of Ghosh-Peres [GP17], appeared in [CN18].

The number-rigidity of Sine_β justifies a posteriori that conditioning on the number of points within Λ is actually *not an extra information to ask for*, once we have specified the exterior configuration $\mathbf{X}_{\overline{\Lambda}}$. Our canonical DLR equations are thus "the" DLR equations.

1.4. Application and perspectives

An application of DLR equations: fluctuations of linear statistics

In [BLS18], with F. Bekerman and S. Serfaty, we had used a "transportation" method to study the fluctuations of linear statistics $\sum_{i=1}^N \varphi(x_i)$ in β-ensembles. This is an old topic going back to Johansson [Joh98], the main point being that if φ is smooth enough (in practice, a few derivatives suffice), the variance of $\sum_{i=1}^N \varphi(x_i)$ remains bounded as $N \to \infty$, which is remarkable. Moreover, one usually gets an associated CLT, with some additional technical conditions regarding the test function φ and the support of the "equilibrium measure", which we will not discuss here. [BL18] had proven that the CLT remained true if φ was living at some mesoscopic scale. It was natural to ask for a similar result at the *microscopic* scale, namely for Sine_β. As the method of [BLS18] used the Gibbsian picture of β-ensembles in a crucial way, my adaptation to Sine_β relies on the DLR equations of [DHLM21].

Theorem 5 ([Leb19])**.** Let $\overline{\varphi}$ be a fixed test function of class C^4, compactly supported on $\mathbb{R}$, and for all $\ell > 0$, let $\varphi_\ell : x \mapsto \overline{\varphi}\left(\frac{x}{\ell}\right)$. Let $\mathbf{X}^\bullet$ be a random point configuration distributed according to Sine_β. The following convergence holds, in law, as $\ell \to \infty$,

$$\int_{\mathbb{R}} \varphi_\ell(x) \mathtt{d}\mathbf{X}^\bullet(x) - \int_{\mathbb{R}} \overline{\varphi}(x)\mathtt{d}x \implies \text{Gaussian r.v. of mean 0 and variance } \frac{2}{\beta}\|\overline{\varphi}\|^2_{H^{\frac{1}{2}}}.$$

Perspective: a fully physical approach to Sine_β?

All the results concerning Sine_β mentioned in this chapter rely at least... on its very existence, as proven by Valkó&Virág. It would be satisfying to have a self-standing "statistical physics" approach to the infinite-volume limit that would include a proof of its existence. Here is an outline of how this could go:

1. Study the finite-volume Gibbs measure and prove *local laws* that are valid *down to the microscopic scale* and imply a uniform-in-N bound on the expected number of points in any interval of size $\mathcal{O}(N^{-1})$.
   This was Luke Peilen's PhD thesis [Pei23]. The proof relies on a delicate bootstrap-in-scales argument.
2. Such local laws immediately ensure the existence of limit points for the law of $\mathbf{X}_{N,\mathsf{x}}$ (with x in the bulk) in some weak topology. Upgrade the convergence to a stronger topology and prove that any limit point for the law of $\mathbf{X}_{N,\mathsf{x}}$ satisfies DLR equations and is number-rigid.
   The proof should be close to the argument of [DHLM21], using the local laws to make up for the lack of stationarity. The two-dimensional version of this step is implemented in [Leb24].
3. Show that any point process which satisfies DLR equations is stationary.
   This is not trivial, but there is a basic strategy going back to [FP81, FP86]: construct good "approximate translations" and estimate their cost. An "approximate translation" would here be a map $\mathbb{R} \to \mathbb{R}$ which act as $x \mapsto x+1$ on $[-L, L]$ for some large L, and as $x \mapsto x$ far away. In short, one needs to control the effect of such a map on the interaction energy of an infinite point configuration, uniformly as $L \to \infty$. Again, this was done for the two-dimensional model in [Leb24].
4. Prove that *every stationary solution of the DLR equations minimizes the free energy functional* $\mathcal{F}_\beta$.
   Such a statement, together with the converse implication, is known in the literature as the DLR variational principle, and has been proven for a broad class of lattice spin systems, but the existing results do not apply to the case of one-dimensional log-gas. It requires a careful adaptation of the proof strategy to this context.
5. Recall that $\mathcal{F}_\beta$ has in fact a unique stationary minimizer, as proven in [EHL21]. Deduce that there is a unique stationary solution of DLR equations.
   Another option would be to prove directly the uniqueness of solutions to DLR equations, without going through the connection with the free energy functional, but there is no clear way to start.
6. Deduce that the the law of $\mathbf{X}_{N,\mathsf{x}}$ has a unique limit point, hence a limit... which we call Sine_β.

2. Progress on the 2DOCP

2.1. The two-dimensional one-component plasma

Let $N \geq 1$ be an integer, let $\Sigma_N \subset \mathbb{R}^2$ be the disk of center 0 and radius $\mathrm{R}_N := \sqrt{\frac{N}{\pi}}$. Denote by $\mathrm{X}_N := (x_1, \dots, x_N)$ a N-tuple of points in Σ_N and let $\mathbf{X}_N := \sum_{i=1}^N \delta_{x_i}$ be the associated atomic measure. Define the *logarithmic interaction energy* $\mathsf{F}_N(\mathbf{X}_N)$ as:

$$\mathsf{F}_N(\mathbf{X}_N) := \frac{1}{2} \iint\limits_{(x,y)\in\Sigma_N\times\Sigma_N, x\neq y} -\log|x-y| \mathtt{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(x) \mathtt{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(y), \tag{2.1.1}$$

where Leb denotes the Lebesgue measure on $\mathbb{R}^2$. We then define a probability density on the space of N-tuples of points in Σ_N by setting, for $\beta > 0$:

$$\mathtt{d}\mathbb{P}_N^\beta(\mathrm{X}_N) := \frac{1}{\mathrm{K}_N^\beta} \exp\left(-\beta \mathsf{F}_N(\mathbf{X}_N)\right) \mathtt{d}\mathrm{X}_N. \tag{2.1.2}$$

Here and below, $\mathtt{d}\mathrm{X}_N$ denotes the product Lebesgue measure $\mathtt{d}x_1 \dots \mathtt{d}x_N$ on $(\mathbb{R}^2)^N$, which serves as the reference measure. The quantity K_N^β is the *partition function*, namely the normalizing constant:

$$\mathrm{K}_N^\beta := \int_{\Sigma_N\times\cdots\times\Sigma_N} \exp\left(-\beta \mathsf{F}_N(\mathbf{X}_N)\right) \mathtt{d}\mathrm{X}_N.$$

The probability measure $\mathbb{P}_N^\beta$ is called the *canonical Gibbs measure of the two-dimensional one-component plasma* (2DOCP) at *inverse temperature* β. We denote the expectation under $\mathbb{P}_N^\beta$ by $\mathbb{E}_N^\beta$.

The 2DOCP is a well-studied model of statistical physics, also known as Coulomb or log-gas, or jellium. When defining the energy F_N in (2.1.1) and the canonical Gibbs measure $\mathbb{P}_N^\beta$ in (2.1.2), we think of $(x_1, \dots, x_N)$ as the positions of point particles in Σ_N which all carry the same electric charge $+1$, and which are immersed in a uniform neutralizing background of constant density on Σ_N. The logarithmic potential through which the "charges" interact can then be thought of as the electrostatic (Coulomb) interaction potential in dimension 2.

Discussion of the model

There are slightly different ways to define this model.

— Some physics papers work with an "infinitely extended equilibrium" Coulomb system, e.g. [MY80, Leb83, JLM93]. The mathematical existence of such infinite-volume limits is not yet clear for $\beta \neq 2$, see Section 2.3 for a discussion.

— In the statistical physics literature, when working in finite volume, it is common to place N particles in a "uniform neutralizing background of opposite charge" which occupies a certain domain Σ_N with constant density $\rho_N := \frac{-N}{|\Sigma_N|}$, as described above. There is then *perfect confinement* of the particles, in the sense that they are not allowed to live outside Σ_N. The domain Σ_N is not always explicitly chosen, though it often ends up being a disk, mostly by default or for the convenience of symmetry. There is also some interest for studying the 2DOCP on a sphere, which avoids having to deal with a boundary. Some authors state their results for different "reasonable shapes" as e.g. [SM76].

In practice, most results concerning the "bulk" of the system, namely the behavior of the particles living at a macroscopic distance from the boundary, do not depend on the choice of the boundary.

— In contrast, in the mathematical physics literature around the planar Coulomb gas (e.g. [ZW06, AHM11, SS15]) the particles/eigenvalues X_N are usually not confined *a priori* in a certain domain of the Euclidean space, but are rather subject to a certain external "confining" potential/field/weight, the usual choice being the quadratic potential. This model is sometimes called a *(two-dimensional) β-ensemble* by analogy with the one-dimensional models discussed in Chapter 1.

The macroscopic properties of the gas will depend on the choice of the potential, but to a large extent the microscopic behavior in the bulk of the system does not (there seems to be a form of universality of local statistics with respect to the potential, although no results have yet been proven) and coincides with what one gets using a "perfect confinement".

The Ginibre ensemble. When the inverse temperature parameter β is equal to 2, the two-dimensional Coulomb gas (with the "quadratic external potential" convention) admits an interpretation as the joint law of the complex eigenvalues of a $N \times N$ non-Hermitian random matrix, known as the (finite) Ginibre ensemble, after [Gin65]. The model belongs to the class of *determinantal processes* (see e.g. [HKPV06, Section 4.3.7]) and is amenable to exact computations. This is the only value for which a matrix model is known - in particular, the connection between "Random eigenvalues $\leftrightarrow$ Particles" of (1.1.3), valid for the log-gas in dimension 1, does not seem to hold in dimension 2 besides this particular case. This is also the only case for which we know:

— Existence of an infinite-volume limit, called the "infinite" Ginibre point process.

— Two-point correlation functions, and an explicit expression for k-point correlation functions.

Most results about the 2DOCP at a general temperature have first been obtained for the Ginibre ensemble, often with (much) shorter proofs.

Mathematical treatment. As mathematical objects, Coulomb gases (under various forms: one-, two- or three-dimensional, with one or two components...) have attracted much interest. The recent lecture notes [Ser24] of S. Serfaty include an overview of motivations, ranging from constructive approximation to the study of the Quantum Hall Effect, via random matrix theory.

Concerning the 2DOCP itself, topics that have been investigated in the last years alone are very diverse and include: lower bounds on the minimal distance between points [Ame18], concentration inequalities for the empirical measure of the particles ([CHM18], see also [PG23] for a sharper bound), upper bounds for the local density of points [LRY19], generalizations to Riemannian manifolds ([Ber19], [GZ19]), Wegner's and clustering estimates [Tho24]... to quote only a few.

While the macroscopic behavior of 2DOCP's at all temperatures had been understood for some time, the seminal works of Sandier-Serfaty [SS12, SS15] opened the way to a statistical physics approach of the behavior at microscopic scale - here $N^{-\frac{1}{2}}$, which corresponds to the typical distance between nearest-neighbors.

Crystallization versus rigidity at all temperatures. A major open question concerning the 2DOCP is the physics conjecture, supported by various numerical simulations [AJ81, CLWH82, CC83], that it undergoes a liquid/solid phase transition around $\beta \approx 140$.

To pinpoint such a transition at a mathematical level, one would expect to see some property expressing "disorder" for $\beta < 140$ and "order" for $\beta > 140$. However, the results of this chapter all tend to indicate [1] that the 2DOCP *appears very "ordered" for all positive values.*

A first example of this phenomenon was found in [LS18, BBNY19, Ser23]. In short, if φ is a smooth enough test function, then

$$\mathrm{Var}\left[\sum_{i=1}^{N} \varphi(x_i)\right] = \mathcal{O}(1) \text{ as } N \to \infty, \text{ and is thus } \ll N$$

which expresses a strong form of rigidity compared to the i.i.d case, as far as fluctuations of smooth linear statistics are concerned. The next section deals with the case when φ is the indicator function.

1. It does not mean that there is no phase transition, but simply that we may need to keep looking for other observables!

2.2. Hyperuniformity of the 2DOCP

If φ is the indicator function of a disk $\mathfrak{D}$, then $\mathsf{Pts}(\mathbf{X}_N, \mathfrak{D}) := \sum_{i=1}^{N} \varphi(x_i)$ counts the random number of points falling into that disk. The variance of $\mathsf{Pts}(\mathbf{X}_N, \mathfrak{D})$ is called the *number variance* - for i.i.d. particles, this would be proportional to the area. We let $\mathsf{Dis}(\mathbf{X}_N, \mathfrak{D}) := \mathsf{Pts}(\mathbf{X}_N, \mathfrak{D}) - \mathrm{Leb}(\mathfrak{D})$ be the "discrepancy".

The term "hyperuniform(ity)" has been coined in the theoretical chemistry literature by S. Torquato (see [TS03, Tor18] for surveys), an alternative terminology due to J. Lebowitz is "superhomogeneous/ity". Hyperuniformity of a system is defined by the fact that:

> [Tor18, Section 1] *(...) the number variance of particles within a spherical observation window of radius R grows more slowly than the window volume in the large-R limit.*

The definition of hyperuniformity needs an adaptation for finite systems, as one should take both the size of the system and of the "spherical window" to infinity. In [Leb26], we prove the following:

Theorem 6 (Hyperuniformity of the 2DOCP at all temperatures.)**.** For each $N \geq 1$, let $\mathsf{x} = \mathsf{x}(N)$ be a point in the bulk of Σ_N and let $R = R(N)$ be such that $R(N) \to \infty$. Then the number variance in the disk of center x and radius R is $o(R^2)$ as $N \to \infty$.

This is valid for all $\beta > 0$ although some implicit constants depend on β. The precise statement reads:

Proposition 2.2.1 ([Leb26])**.** *Let $\delta > 0$ be fixed. For all N and R large enough (both depending on β and δ), for all x in Σ_N such that x is "in the bulk" in the following sense:*

$$\mathrm{dist}(\mathfrak{D}(\mathsf{x}, R), \partial\Sigma_N) \geq \delta\sqrt{N},$$

we have:

$$\mathbb{P}_N^{\beta}\left(\left\{|\mathsf{Dis}(\mathbf{X}_N, \mathfrak{D}(x, R))| \geq R(\log R)^{-0.3}\right\}\right) \leq \exp\left(-\log^{1.5} R\right).$$

In particular, we obtain that, as $N \to \infty$:

$$\mathrm{Var}[\mathsf{Pts}(\mathbf{X}_N, \mathfrak{D}(\mathsf{x}, R))] = \mathcal{O}\left(\frac{R^2}{\log^{0.6} R}\right) = o(R^2).$$

The power 0.6 is not optimized, but the method would not give anything better than $\log R$.

Context of the result. Hyperuniformity of the 2DOCP at all positive temperatures was a forty year old prediction in the statistical physics literature, see [2] [MY80, Leb83, Mar88, JLM93, LWL00] (which use a different terminology, as the term "hyperuniformity" had not been coined yet), or [Tor18] again: *"OCP fluid phases at [all] temperatures must always be hyperuniform"*.

The full physical prediction says not only that the number variance is negligible with respect to the area of the disk, but that it should even be *comparable to the perimeter* ($\mathcal{O}(R)$ and not only $o(R^2)$). In fact, the predictions go even beyond: in [JLM93], Jancovici, Lebowitz and Manificat made precise predictions concerning the probability of observing large charge fluctuations within the 2DOCP. Their statement is significantly more precise than hyperuniformity, as they argue that for all $\alpha > \frac{1}{2}$:

$$\mathbb{P}\left[\text{The discrepancy } \mathsf{Dis} \text{ in a disk of radius } R \text{ is larger than } R^{\alpha}\right] \sim \exp(-R^{\varphi(\alpha)})$$

where the rate $\varphi(\alpha) > 0$ is an explicit piecewise affine function of α. This was later checked for $\beta = 2$ through explicit computations, see [Shi06, OS08, FL21], while the general β case is open for small α, which is the most difficult range. Remarkably, although [JLM93] deals with Coulomb gases, it was verified in [NSV08] for a different model, namely the zeros of the Gaussian Entire Function (see below).

Our result [Leb26] is the first mathematical proof of hyperuniformity for $\beta \neq 2$, but remains far from the conjectured sharp bounds, and there is much left to be understood.

2. Those results are rigorous to the extent that authors make use of so-called "clustering assumptions", i.e. they assume properties of the two-point correlation function at large distances in order to derive certain identities (called "sum rules") which, among other things, imply hyperuniformity. As explained in [Mar88]: "*The results obtained in this way are exact (i.e. do not follow from approximations), but not all of them are rigorously proven, in so far as some reasonable properties (e.g. the type of decay of the correlations) are assumed to hold a priori.*". Unfortunately, obtaining mathematically rigorous statements about the large-distance properties of the 2DOCP's two-point correlation function (for $\beta \neq 2$) is extremely challenging.

Outline of the proof

Our main source of inspiration is [NSV08, Section 4], in which Nazarov-Sodin-Volberg study the probability of having large "charge" fluctuations in the disk $\mathfrak{D}(0,R)$ for the zeroes of the Gaussian Entire Function. This random function is given by:

$$z \mapsto \sum_{n\geq 0} \frac{1}{\sqrt{n!}} a_n z^n, \text{ where } (a_n)_{n\geq 0} \text{ are i.i.d. complex Gaussian variables,} \tag{2.2.1}$$

which almost surely defines an entire function on $\mathbb{C}$. The point process of its zeroes happens to bear some interesting resemblance with the Ginibre ensemble and, more generally, with the particles of a 2DOCP.

The argument of [NSV08] can be roughly summarized as follows: fix $\alpha \in \left(\frac{1}{2}, 1\right)$ and assume that there is a discrepancy of size R^α in the disk $\mathfrak{D}(0,R)$, i.e. that $\mathsf{Dis}(\mathbf{X}_N, \mathfrak{D}) := \mathsf{Pts}(\mathbf{X}_N, \mathfrak{D}) - \mathrm{Leb}(\mathfrak{D}) \geq R^\alpha$. Then:

1. Show that the discrepancy can be "captured" along the boundary $\partial\mathfrak{D}(0,R)$.
2. Split that boundary into $\approx R$ pieces of size ≈ 1. Take M large and apply a basic pigeonhole argument: there exists a family of $\approx \frac{R}{M}$ pieces that capture a discrepancy of size at least $\frac{R^\alpha}{M}$ and which are "well-separated" (distances between pieces are multiples of M).
3. Then comes the main probabilistic work:
   a) Show that these well-separated pieces are *approximately independent*.
   b) Show that the discrepancy on each piece (of size ≈ 1) is typically $\mathcal{O}(1)$.
   c) Show that the discrepancy on each piece is centered.
4. Apply Bernstein's concentration inequality: if $\{D_i\}_i$ is a family of $\approx \frac{R}{M}$ independent centered random variables of size $\mathcal{O}(1)$ then:
$$\mathbb{P}\left(\sum_i D_i \geq R^\alpha\right) \lesssim \exp\left(-\frac{R^{2\alpha-1}}{M}\right).$$
   Such a tail estimate yields sharp bounds on the number variance.

We follow a similar strategy, with much effort to translate it to the context of the 2DOCP.

Technical imports and new techniques

The technical core is the "statistical physics" approach to Coulomb gases developed by Serfaty et al. [SS15, RS16, PS17, LS18, AS21, Ser23], see [Ser24]. We develop three new technical tools.

1. The role of sub-systems. First, we put forward the role of so-called "sub-systems", which arise as restrictions of the full system to some region and which we view as two-dimensional Coulomb systems in their own right, possibly with a small global non-neutrality, and (most importantly) feeling the effect of some harmonic exterior potential. As such, this object is not new - it has been sometimes called a "conditional" or "local" measure in the literature, see e.g. [BBNY19, BEY14].

We provide a thorough study of its behavior, through a global law and local laws, which show that with high probability the sub-systems, although under the influence of this external potential, retain most of the "good" properties. A big part of the work consists in ensuring that the techniques developed for the "full" system can be extended to this more involved situation, which is sometimes very tedious.

The fact that the local laws extend to sub-systems allows one to treat Point (3) (b), which seems unsurprising, but requires to consider objects of size ~ 1, hence to control the system *down to the smallest scales*. Such local laws are also crucial for the next tool: an approximate translation-invariance result.

2. Approximate translation-invariance. Step (3) (c) is void in the context of [NSV08] because the underlying point process is infinite and translation-invariant, which implies that linear statistics are centered. However, it turns out to be a major roadblock when adapting the proof to a Coulomb gas context. Indeed, in sub-systems - or even in the full system if one does not impose artificial boundary conditions - there is absolutely no obvious translation-invariance.

We state an "approximate translation-invariance" result, valid for both the full system and sub-systems, which is crucial for Step (3) (c). The point is that that translating a function (or more precisely, averaging

over translations) acts as a mollification and enables us to compare the expectation of a non-smooth linear statistics, like the discrepancy, to a smoother one.

There is a series of results in mathematical statistical mechanics *à la* H.-O. Georgii devoted to prove translation invariance of infinite volume Gibbs measures in contexts where stationarity is not built-in, see e.g. [Geo99, Ric07], following earlier works by Fröhlich-Pfister [FP81, FP86]. The basic idea is to construct suitable "localized translations" in the form of diffeomorphisms acting as a given translation in a large box while leaving the majority of the system unchanged, and to control the effect on the energy of such changes of variables. However, this last aspect is very challenging when working with Coulomb interactions, and ends up requiring a very careful choice of the "localized translation" as well as a painstaking analysis of its effect on the interaction energy.

Following the cryptic but inspiring wisdom found in a remark of [Sim14, Chap. 3, Sec. III.7], we construct a localized translation that varies very slowly (in terms of its H^1 norm). We then proceed to carefully revisit a difficult computation done in [Ser23]. This, together with Serfaty's "smallness of anisotropy" trick (which is also highly non-trivial), allows us to conclude that our localized translation can be chosen to have an arbitrarily small effect on the energy. An easy-to-state consequence is the following:

Corollary 2.2.2 ([Leb26])**.** *Let* G *be a bounded function on the space* Conf *of point configurations, which is local in the sense that* $\mathsf{G}(\mathbf{X})$ *depends only on the restriction of* $\mathbf{X}$ *to some bounded set of* $\mathbb{R}^2$*, and let* $\vec{v} \in \mathbb{R}^2$*. We have:*

$$\mathbb{E}\left[\mathsf{G}(\mathbf{X}_N)\right] = \frac{1}{2}\left(\mathbb{E}\left[\mathsf{G}(\mathbf{X}_N + \vec{v})\right] + \mathbb{E}\left[\mathsf{G}(\mathbf{X}_N - \vec{v})\right]\right) + o_N(1).$$

If $\mathbf{X}_N$ converges in distribution (up to extraction) to a random infinite point configuration $\mathbf{X}$, we get:

$$\mathbb{E}\left[\mathsf{G}(\mathbf{X})\right] = \frac{1}{2}\left(\mathbb{E}\left[\mathsf{G}(\mathbf{X} + \vec{v})\right] + \mathbb{E}\left[\mathsf{G}(\mathbf{X} - \vec{v})\right]\right),$$

which implies that $\mathbf{X}$ and $\mathbf{X} + \vec{v}$ have the same distribution.

3. Approximate independence of sub-systems. Finally, one needs to find some kind of independence between the sub-systems to treat Step (3) (a). We introduce a new "approximate independence" argument for sub-systems that are well-separated (far enough from each other), conditionally on the *number of points* in each of them. The simple idea is that two domains Λ_i, Λ_j with $\mathsf{n}_i, \mathsf{n}_j$ points have a contribution to the interaction energy F_N (2.1.1) given, to first order, by:

$$-\left(\mathsf{n}_i - |\Lambda_i|\right)\left(\mathsf{n}_j - |\Lambda_j|\right)\log \operatorname{dist}(\Lambda_i, \Lambda_j),$$

while the precise arrangements of the points inside each domain should matter only to a lower order. The statements look as follows: for $i = 1, \dots, \mathcal{N}$ we fix a region Λ_i within the system and a function G_i which is Λ_i-local, and we write:

$$\mathbb{E}_N^\beta\left[\prod_{i=1}^{\mathcal{N}} \mathsf{G}_i(\mathbf{X}_N)\right] \leq \exp\left(2\beta \sup_{\mathbf{X}_N} \mathrm{ErrorCI}(\mathbf{X}_N)\right) \times \sup_{\mathbf{X}_{\mathrm{Ext}}, \{\mathsf{n}_i\}} \prod_{i=1}^{\mathcal{N}} \mathbb{E}^\beta_{\mathsf{n}_i, \Lambda_i, \mathrm{V}^{\mathrm{ext}}}\left[\mathsf{G}_i(\mathbf{X})\right].$$

On the left-hand side, the expectation is taken under the true Gibbs measure (2.1.2) while on the right-hand side, we work with certain "sub-systems", one for each Λ_i, all feeling the effect of an external potential $\mathrm{V}^{\mathrm{ext}}$ created by the configuration $\mathbf{X}_{\mathrm{Ext}}$ outside $\bigcup_i \Lambda_i$. Those local Gibbs measures depend on the number of points in each Λ_i, denoted by n_i. They appear to be "independent" in the sense that we get a product of expectations in the right-hand side. A similar lower bound holds with $-$ErrorCI instead of ErrorCI.

Despite our phrasing, one should not think of distant subsystems as being truly independent: the size of ErrorCI is not $o(1)$. But it might be small *in comparison* to other things. In particular, events that depend on several distant subsystems and that are *very unlikely for independent subsystems* are also very unlikely for the true Gibbs measure.

2.3. DLR equations, number-rigidity, translation-invariance

For x in Σ_N, and if X_N is sampled from $\mathbb{P}_N^\beta$, define $\mathbf{X}_{N,\mathsf{x}}$ (the local point configuration *seen from* x) as:

$$\mathbf{X}_{N,\mathsf{x}} := \sum_{i=1}^{N} \delta_{x_i - \mathsf{x}}.$$

In [AS21, Corollary 1.1], it is proven that for all (possibly N-dependent) choices of x *in the bulk* of Σ_N, the law of $\mathbf{X}_{N,\mathsf{x}}$ converges as $N \to \infty$ *up to extraction of a subsequence* to the law P_∞^β of some random infinite point configuration.

The argument of [AS21] is an abstract existence result based on compactness, in particular it does not say anything about the point processes. In [Leb24], I gave a "physical" description of those limit points in terms of DLR equations, as well as some properties valid for all $\beta > 0$.

DLR equations

As for the one-dimensional log-gas, a first issue to overcome is how to make sense of DLR equations. In fact the logarithmic potential is even more "long-range" in dimension 2, in the sense that:

$$\int_{|x|\geq 1} \log|x|\mathtt{d}x = +\infty, \quad \int_{|x|\geq 1} \|\nabla \log|x|\|\mathtt{d}x = +\infty, \text{ and even } \int_{|x|\geq 1} \|\nabla^2 \log|x|\|\mathtt{d}x = +\infty.$$

Let $\Lambda \subset \mathbb{R}^2$ a compact subset. Giving sense to DLR equations implies defining the interaction between a configuration of points in Λ and one in $\overline{\Lambda}$. A naive way would be to set (cf. (1.3.3) in the $1d$ case):

$$\mathsf{M}_\Lambda^?(\mathbf{X}', \mathbf{X}) := \iint_{\Lambda\times\overline{\Lambda}} -\log|x-y|\mathtt{d}\mathbf{X}'(x)\mathtt{d}(\mathbf{X}-\mathrm{Leb})(y),$$

This however might not make sense, even when invoking the fine estimates of [LS18, BBNY19, Ser23] on the typical size of the linear statistics. In short, this is due to the fact that $\nabla \log|x-\cdot|$ (barely) fails to be in L^2 near infinity.

Similarly to what was done in the $1d$ case, see Section 1.3, we consider instead a *difference of energies* and introduce some "abstract" *reference point configuration* $\mathbf{X}^{\mathrm{abs}}$ *in* Λ *that has the same number of points as* $\mathbf{X}'$. The actual choice does not matter, we can simply take

$$\mathbf{X}^{\mathrm{abs}} := \mathsf{Pts}(\mathbf{X}', \Lambda) \times \delta_0,$$

i.e. we place the correct number of points, all of them located at the origin, and we then define:

$$\mathsf{M}_\Lambda(\mathbf{X}', \mathbf{X}) = \widetilde{\mathsf{M}}_\Lambda(\mathbf{X}', \mathbf{X}) - \widetilde{\mathsf{M}}_\Lambda(\mathbf{X}, \mathbf{X}) = \iint_{\Lambda\times\overline{\Lambda}} -\log|x-y|\mathtt{d}\left(\mathbf{X}'-\mathbf{X}\right)(x)\mathtt{d}\left(\mathbf{X}-\mathrm{Leb}\right)(y).$$

which is the cost of "moving points from $\mathbf{X}$ to $\mathbf{X}'$ within Λ", cf. (1.3.3), (1.3.4). We are able to prove that this makes sense for typical configurations generated by any possible limit point P_∞^β.

Theorem 7 ([Leb24])**.** For all compact sets $\Lambda \subset \mathbb{R}$, for P_∞^β-almost every point configuration $\mathbf{X}$ and for all point configurations $\mathbf{X}'$ *such that* $\mathbf{X}'$ *and* $\mathbf{X}$ *have the same number of points in* Λ, the quantity $\mathsf{M}_\Lambda(\mathbf{X}', \mathbf{X})$ is well-defined.

This result is significantly more subtle than its one-dimensional counterpart, although the statements look the same. The basic controls on $\mathbf{X} - \mathrm{Leb}$ (e.g. discrepancy estimates) analogous to those used in $1d$ now fail to ensure convergence. The solution was to:

— Decompose the space using a smooth partition of unity on dyadic scales, instead of a "sharp" one.

— Applying fine bounds for the fluctuations of smooth linear statistics, instead of "rough", energy-based, ones. Those are known in finite-volume and need to first be transferred to the limit point P_∞^β.

As in Section 1.3, once the "Move" functions are defined, one can consider *canonical* DLR equations.

Properties of the limit points

Theorem 8 ([Leb24])**.** For all $\beta > 0$, and all choices of x *in the bulk* of Σ_N, the limit points P_∞^β are:

1. Number-rigid (see Definition 1.3.2).
2. Solutions to the canonical DLR equations.
3. Translation-invariant.

All of this is new for $\beta \neq 2$. In fact, the mere definition of DLR equations for the 2DOCP was an open problem, see e.g. [GL17, Sec. 12]. Some comments:

— The number-rigidity result might not come as a surprise, but it provides a natural one-parameter family of two-dimensional number-rigid point processes which (conjecturally) interpolate between a Poisson point process ($\beta \to 0$) and a lattice ($\beta \to \infty$).

— The fact that translation-invariance holds for β arbitrarily large is physically significant. Indeed, it is expected (see e.g. [Lew22, Sec. VI] for a survey) that around $\beta \approx 140$, the system becomes a "solid", and that as $\beta \to \infty$ it becomes a lattice. Our result shows that even in its hypothetical "solid" phase, and despite the long-range nature of the interaction, the 2DOCP does not break translation-invariance.

Of course, a very important question is left unsolved for $\beta \neq 2$:

Does the law of the local point configuration $\mathbf{X}_{N,\mathsf{x}}$ have an actual limit*, and not just limit points?*

A positive answer for some $\beta > 0$ would yield the existence of a "Ginibre-β" point process, whereas a negative one would hint at a phase transition.

2.4. Perspective: finding a liquid phase for the 2DOCP

From numerical simulations, the phase portrait of the 2DOCP is expected to look like this:

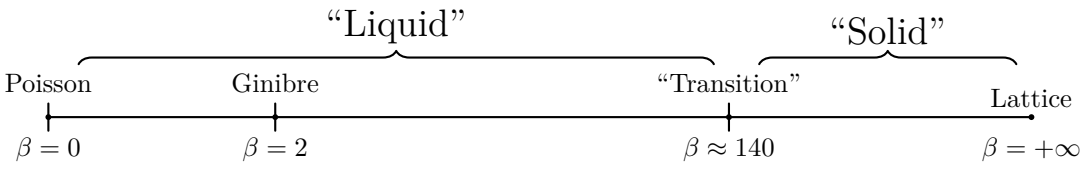

However, so far, the only part that is rigorously proven is that as $\beta \to 0$, minimizers of the free energy functional $\mathcal{F}_\beta$ converge to the Poisson point process - this is a result from my PhD thesis [Leb16]. The other extreme point, namely the $\beta = +\infty$ behavior, is a major open problem - and any significant progress for the range $\beta \gg 1$ seems unlikely to happen as long as the ground state itself is not rigorously understood. The most realistic hope for progress in a genuine "statistical physics" understanding of the 2DOCP seems to reside in the "high-temperature", or $\beta \ll 1$, region. Indeed, "high-temperature" behaviors are usually well-behaved and easier to understand, because the interaction becomes very weak. The literature contains some general results, among which *Dobrushin's uniqueness theorem* as well as , but they do not apply to point processes with such long range interactions as ours. We are left with the following objective:

Goal: prove that for β small enough, the 2DOCP is in a "liquid phase" in the following sense:

— *The DLR equations have a unique solution, which is also the unique minimizer of the free energy.*

— *Two-point correlations decay exponentially fast (this is known for $\beta = 2$, in fact the decay is Gaussian).*

3. Hyperuniformity & friends in dimension 2

3.1. Hyperuniformity as "order within disorder"

In the statistical physics literature of the 1980's (see e.g. [MY80, Leb83]) it has been observed that some classical systems of particles with long-range interactions exhibit *cancellation of charge fluctuations*, namely that the variance of the number of charges in a domain increases proportionally to the surface area of the domain's boundary, and not to the volume. The notion of *hyperuniformity*, or equivalently of *super-homogeneous* distributions of points was introduced in [GJL02, TS03], with applications in theoretical chemistry and physics. We refer to the survey by S. Torquato [Tor18] for an overview of the considerable literature that has been devoted to this topic - and continues to grow.

One important motivation to study hyperuniform processes is that they can exhibit a form of "order within disorder". Indeed:

— The fact that their rescaled number variance goes to 0 can be interpreted as a form of collective order.

— Their truncated correlation functions may decay fast at infinity - hence there is no "long-range order" in the usual statistical physics sense.

Beyond hyperuniformity, one can think of other notions expressing order, e.g. number-rigidity as defined above (Definition 1.3.2), or the two properties that we will introduce below: finiteness of Coulomb energy, and finiteness of Wasserstein distance to Lebesgue. A mind-opening observation is that it is possible to present all those features of "order" and yet to be "disordered" in the classical sense. An excellent example of this phenomenon is the infinite Ginibre ensemble:

— It is "as hyperuniform as it gets", number-rigid, has finite Coulomb energy and finite Wasserstein distance to Lebesgue.

— Its two-point correlation functions decay faster than any exponential (like a Gaussian).

Terminology around hyperuniformity. Let $\mathbf{X}^\bullet$ be a *random* point configuration. We say that $\mathbf{X}^\bullet$ is *hyperuniform* when:

$$\sigma(r) := \frac{\mathrm{Var}\left[|\mathbf{X}^\bullet \cap \mathrm{B}_r|\right]}{|\mathrm{B}_r|} \longrightarrow 0 \text{ as } r \to \infty, \tag{3.1.1}$$

where $|\mathrm{B}_r|$ denotes the volume of a ball of radius r. The quantity $\mathrm{Var}[|\mathbf{X}^\bullet \cap \mathrm{B}_r|]$ is called the *number variance*, while σ is the rescaled number variance.

This definition should be understood by comparison to the fundamental example of a random point configuration given by the Poisson point process. Recall that for a Poisson point process $\mathbf{X}^\bullet$ with intensity 1, the number of points $|\mathbf{X}^\bullet \cap \mathrm{B}_r|$ follows a Poisson distribution of parameter $|\mathrm{B}_r|$ for all $r > 0$, and as a straightforward consequence the number variance is precisely equal to the volume $|\mathrm{B}_r|$. For a random point configuration $\mathbf{X}^\bullet$ to be hyperuniform thus means that its number variance (in large balls) is asymptotically "much smaller than for a Poisson point process".

Of course, the volume $|\mathrm{B}_r|$ is proportional to r^{d} so when (3.1.1) holds one can wonder how slow the number variance growth actually is compared to r^{d}. In the terminology of [Tor18], $\mathbf{X}^\bullet$ is said to be *hyperuniform of class I* when in fact:

$$\mathrm{Var}\left[|\mathbf{X}^\bullet \cap \mathrm{B}_r|\right] = \mathcal{O}(r^{\mathsf{d}-1}) \text{ as } r \to \infty. \tag{3.1.2}$$

Amusingly, the original definition for hyperuniformity was more stringent than (3.1.1) and corresponds to what is now called "class-I hyperuniformity" (3.1.2). In the meantime, it has been realized that many interesting processes satisfy (3.1.1) but not (3.1.2).

3.2. About hyperuniformity of perturbed lattices

The basic example of a class-I hyperuniform point process is the "stationary lattice" defined by:

$$\mathbf{L} := \{x + U,\ x \in \Lambda\},$$

where Λ is a lattice in $\mathbb{R}^{\mathsf{d}}$, and U is uniformly distributed in a fundamental domain of Λ (e.g. $[0,1]^d$ when $\Lambda = \mathbb{Z}^d$). Since (3.1.2) is the slowest possible growth (cf. [Cos21, Thm. 4.2], [Bec87, Thm. 2A]), stationary lattices are "as hyperuniform as it gets". A natural extension is the class of *perturbed lattices*, of the form:

$$\mathbf{PL} := \{x + U + \mathrm{p}_x,\ x \in \Lambda\},$$

where U is a uniform random shift as before, while $(\mathrm{p}_x)_{x\in\Lambda}$ is a family of random variables, independent of U, thought of as "perturbations", with values in $\mathbb{R}^{\mathsf{d}}$.

Taking an arbitrary family would be too general, so we need to make the following assumption: the perturbations $(\mathrm{p}_x)_{x\in\Lambda}$ are identically distributed, and moreover [1] their joint distribution is Λ-invariant in the sense that for all $x_\star$ in Λ, $(\mathrm{p}_x)_{x\in\Lambda}$ and $(\mathrm{p}_{x+x_\star})_{x\in\Lambda}$ have the same distribution.

Do such perturbations affect hyperuniformity of lattices? The following was known in all dimensions:

Theorem ([GS75])**.** Assume that the perturbations are *independent and identically distributed*. If p_0 has a *finite first moment*, then **PL** remains *hyperuniform of class I*.

In fact, even if p_0 has only a finite moment of order $0 < s < 1$, it is not hard to check that **PL** remains hyperuniform (but not necessarily of class I) when using i.i.d perturbations.

For a *dependent* perturbation field $(\mathrm{p}_x)_{x\in\Lambda}$, hyperuniformity was discussed in the physics literature [Gab04, KT18], with the central claim that, in dimension 2, existence of a second moment for p_0 would be enough to preserve hyperuniformity. In [DFHL24], we settled the question as follows:

Theorem 9 ((Non)-hyperuniformity of perturbed lattices, [DFHL24])**.**

— For $\mathsf{d} = 1$ or 2, if $\mathbb{E}[|\mathrm{p}_0|^{\mathsf{d}}]$ is finite, then **PL** is hyperuniform.

— For all $\mathsf{d} \geq 1$, if X is a positive random variable with $\mathbb{E}[X^{\mathsf{d}}] = +\infty$, then there exist perturbations with $|\mathrm{p}_0| \leq X$ almost surely, such that **PL** has *infinite number variance in finite balls*.

— For all $\mathsf{d} \geq 3$, taking $\Lambda = \mathbb{Z}^{\mathsf{d}}$, for all $\varepsilon > 0$ there exist perturbations $(\mathrm{p}_x)_{x\in\mathbb{Z}^{\mathsf{d}}}$ with $|\mathrm{p}_0| \leq \varepsilon$ almost surely, such that the resulting perturbed lattice satisfies $\lim_{r\to+\infty} \frac{\mathrm{Var}[|\mathbf{PL}\cap \mathrm{B}_r|]}{|\mathrm{B}_r|} = +\infty$.

The first item shows that the physicists' claim was indeed correct. The second item shows that this sufficient condition for hyperuniformity is particularly sharp, because anything else can give rise to processes which not only are *not hyperuniform*, but have in fact *infinite number variance in finite balls*. The last item shows that, in dimensions $\mathsf{d} \geq 3$, even arbitrarily small perturbations can break hyperuniformity.

The proof of the first item follows the steps of the argument found in the physics literature [Gab04].

Exotic behavior of the number variance. To prove the sharpness of our conditions (second and third items of the statement), we construct some families of point processes by hand, as fine-tuned mixtures of some elementary "blocks", for which we can control the asymptotic number variance explicitely. As a byproduct of our methods, we obtain the following result, which amusingly was absent from the literature, and expresses the fact that the rescaled number variance can tend to 0 at an arbitrary slow speed.

Theorem 10 (Arbitrarily slow decay of the number variance, [DFHL24])**.**
Let $\tilde{\sigma} : (0,+\infty) \to (0,+\infty)$ be a non-increasing function with [2] $\lim_{r\to+\infty} \tilde{\sigma}(r) = 0$. For all $\mathsf{d} \geq 1$, there exists a perturbed lattice **PL** such that:

1. **PL** is hyperuniform.
2. The rescaled number variance σ of **PL** satisfies, for some $c > 0$, that $\sigma(r) \geq c\tilde{\sigma}(r)$ for $r \geq 1$.

1. It implies that the perturbed lattice **PL** itself is stationary.
2. We emphasize that $\tilde{\sigma}$ can really be chosen tending to 0 as slow as one wants.

3.3. Hyperuniformity, transportation, and Coulomb

Denote by $\mathtt{HU}$ the property of being hyperuniform (see (3.1.1)). We consider here two other notions.

1. Finite Coulomb energy ($\mathtt{Coul}$). First, we say that a stationary point process $\mathbf{X}^\bullet$ has *finite regularized Coulomb energy* when there exists a stationary *electric field* $\mathrm{E}^\bullet$ compatible with $\mathbf{X}^\bullet$ (a random vector field which could be the gradient of the Coulomb potential generated by the random points of $\mathbf{X}^\bullet$) that satisfies: $\mathbb{E}\left[|\mathrm{E}^\bullet_\eta(0)|^2\right] < +\infty$, where $\eta \in (0,1)$ is any short-distance truncation lengthscale.
See Appendix A for a precise definition.

2. Finiteness of Wasserstein distance to Lebesgue ($\mathtt{Wass}_2$). Let $\mathbf{m}_0, \mathbf{m}_1$ be two measures on $\mathbb{R}^\mathsf{d}$ with the same finite mass. We denote by $\mathsf{cpl}(\mathbf{m}_0, \mathbf{m}_1)$ the set of all *couplings* between $\mathbf{m}_0$ and $\mathbf{m}_1$, i.e. the set of all measures on $\mathbb{R}^\mathsf{d} \times \mathbb{R}^\mathsf{d}$ having $\mathbf{m}_0, \mathbf{m}_1$ as marginals. We recall that for $p \in [1, +\infty)$, the "usual" p-Wasserstein distance between $\mathbf{m}_0$ and $\mathbf{m}_1$ is defined as:

$$\mathrm{W}_p(\mathbf{m}_0, \mathbf{m}_1) := \left(\inf_{\Pi \in \mathsf{cpl}(\mathbf{m}_0, \mathbf{m}_1)} \int_{\mathbb{R}^\mathsf{d} \times \mathbb{R}^\mathsf{d}} |x-y|^p \mathsf{d}\Pi(x,y) \right)^{1/p}.$$

The notion of Wasserstein distance between stationary point processes was introduced in [EHJM25]. First, we need a suitable cost function: for two point configurations $\mathbf{X}_0, \mathbf{X}_1$, we let $\mathsf{cpl}(\mathbf{X}_0, \mathbf{X}_1)$ be the set of all Radon measures on $\mathbb{R}^\mathsf{d} \times \mathbb{R}^\mathsf{d}$ with marginals $\mathbf{X}_0$ and $\mathbf{X}_1$. Then we define the cost function w_p as a "Wasserstein cost per unit volume" between $\mathbf{X}_0$ and $\mathbf{X}_1$:

$$\mathsf{w}_p(\mathbf{X}_0, \mathbf{X}_1) := \left(\inf_{q \in \mathsf{cpl}(\mathbf{X}_0, \mathbf{X}_1)} \limsup_{n \to \infty} \frac{1}{n^\mathsf{d}} \int_{\Lambda_n \times \mathbb{R}^2} |x-y|^p \mathsf{d}q(x,y) \right)^{\frac{1}{p}}, \text{ with } \Lambda_n := \left[-\frac{n}{2}, \frac{n}{2}\right]^\mathsf{d}.$$

Definition 3.3.1. Next, if $\mathsf{P}^0, \mathsf{P}^1$ are stationary point processes, their p-Wasserstein distance is defined in [EHJM25] by:

$$\mathcal{W}_p(\mathsf{P}^0, \mathsf{P}^1) := \left(\inf_{Q \in \mathsf{Cpl}_s(\mathsf{P}^0, \mathsf{P}^1)} \int \mathsf{w}_p^p(\mathbf{X}_0, \mathbf{X}_1) \, \mathsf{d}Q(\mathbf{X}_0, \mathbf{X}_1) \right)^{\frac{1}{p}},$$

where $\mathsf{Cpl}_s(\mathsf{P}^0, \mathsf{P}^1)$ denotes the set of all couplings between P^0 and P^1 which are invariant under the diagonal action of $\mathbb{R}^\mathsf{d}$.

These definitions extend to more general translation-invariant random measures. We are mostly interested in the case $p = 2$, and we denote by $\mathtt{Wass}_2$ the property of having a finite $\mathcal{W}_2$ distance to the Lebesgue measure - or equivalently to a stationary lattice, as it itself is at finite distance of Lebesgue.

The link between those notions

Thanks to [EHJM25, Prop. 2.11], finiteness of the 2-Wasserstein distance to a stationary lattice is equivalent to the fact that $\mathbf{X}^\bullet$ can be written as a *perturbed lattice* $\mathbf{PL}$ with perturbations *whose second moment is finite*, as the one considered above. In this context, the main result of [DFHL24], presented in the previous section, thus says that $\mathtt{Wass}_2$ implies $\mathtt{HU}$ in dimension 2. In [LRY24], a partial converse was stated as follows: if a two-dimensional point process is hyperuniform ($\mathtt{HU}$) and if its two-point correlation function satisfies *"a mild integrability condition"*, then it is at finite distance to Lebesgue ($\mathtt{Wass}_2$).
In [HL25], with M. Huesmann, we clarified the link between hyperuniformity, Wasserstein distance, and Coulomb energy, by showing that:

$$\mathtt{Coul} \implies \mathtt{Wass}_2 \, (\implies \mathtt{HU} \text{ by [DFHL24]}), \qquad \mathtt{HU} \not\Longrightarrow \mathtt{Wass}_2, \quad \mathtt{Wass}_2 \not\Longrightarrow \mathtt{Coul}.$$

Note that both converse implications are proven to be false in general. However, they are "almost" true, because the counterexamples have to be of a very peculiar form. In order to explain what we mean by that, we introduce the following property:

Definition 3.3.2 ($\star$-hyperuniformity [HL25])**.** We say that a stationary point process satisfies $\mathtt{HU}_\star$ when the rescaled number variance $\sigma(r)$ (see (3.1.1)) goes to 0 as $r \to \infty$ *in such a way that*:

$$\sum_{n\geq 0} \sigma(2^n) < +\infty. \tag{3.3.1}$$

We call condition (3.3.1) a "slight reinforcement of hyperuniformity" because any point process such that $\sigma(r) \to 0$ as some negative power of r, or even as $(\log r)^{-c}$ with $c > 1$, satisfies $\mathtt{HU}_\star$. However, $\mathtt{HU}_\star$ *is strictly stronger than* $\mathtt{HU}$ since there exists point processes such that $\lim_{r\to\infty} \sigma(r) = 0$ but $\liminf_{r\to\infty} \sigma(r) \log r > 0$.

Surprisingly, in dimension 2, we found that making the ever-so-slightly stronger assumption $\mathtt{HU}_\star$ is enough to imply $\mathtt{Coul}$ - in fact $\mathtt{HU}_\star$ is *equivalent* to $\mathtt{Coul}$ - so the final chain of implications looks like this:

Theorem 11 ([HL25])**.** For stationary point processes in dimension 2, we have:

$$\mathtt{HU}_\star \iff \mathtt{Coul} \implies \mathtt{Wass}_2 \implies \mathtt{HU}, \qquad \mathtt{HU} \not\Longrightarrow \mathtt{Wass}_2, \quad \mathtt{Wass}_2 \not\Longrightarrow \mathtt{Coul}.$$

We also prove a statement of the form $\mathtt{Wass}_2 \implies \mathtt{Coul}$ for point processes that satisfy a strong, uniform density bound - which form a small, but non-empty family, including e.g. point processes with a "hardcore" constraint.

The connection with the results of [LRY24] is as follows: their "mild integrability condition" on the two-point correlation function, together with hyperuniformity, imply that the process satisfies our condition $\mathtt{HU}_\star$, i.e. has finite Coulomb energy, and thus (by our Theorem 11) has finite Wasserstein distance.

Comments on the proof.

The Coulomb energy is very closely related to the H^{-1} norm, up to important technical details. In short, if $\mathbf{m}$ is a signed measure on a finite box Λ_n, an integration by parts yields:

$$\iint_{\Lambda_n \times \Lambda_n} -\log|x-y| \mathtt{d}\mathbf{m}(x)\mathtt{d}\mathbf{m}(y) \approx \frac{1}{2\pi} \int_{\Lambda_n} |\nabla \Delta^{-1}\mathbf{m}|^2,$$

because log is the Coulomb kernel and we have $-\Delta \log|\cdot| = 2\pi\delta_0$. The left-hand side is the Coulomb interaction and the right-hand side is (reminiscent of) the H^{-1} norm of a measure - actually, one would like the vector field $\nabla\Delta^{-1}\mathbf{m}$ to have vanishing Neumann boundary conditions on $\partial\Lambda_n$.

On the other hand, there exists a link between the H^{-1} norm of the difference of two measures and the 2-Wasserstein distance between them as written down in [Led19]. This connection is best seen using the Benamou-Brenier point of view on optimal transportation, which we will not present here. It yields an inequality of the form:

$$\mathrm{W}_2^2(\mathbf{m}_0, \mathbf{m}_1) \leq \mathtt{C}\|\mathbf{m}_0 - \mathbf{m}_1\|_{H^{-1}}^2.$$

It is thus possible to connect the 2-Wasserstein distance (within a given box) between a point configuration $\mathbf{X}$ and the Lebesgue measure, to the H^{-1} norm of their difference, and ultimately to the Coulomb interaction between $\mathbf{X}$ and Leb. This is only valid in finite volume, and provided that the underlying vector fields satisfy a vanishing Neumann boundary condition. The basic idea is to then "lift" such an inequality to infinite-volume. This is made possible thanks to the screening technique of Sandier-Serfaty, using a procedure that I had already explored during my PhD in [Leb16].

The link between $\mathtt{HU}_\star$ and $\mathtt{Coul}$ is very different, it consists in moving to Fourier side and working with the "spectral measure". We rely crucially on a recent result by Sodin-Wenmann-Yakir [SWY23b, SWY23a].

Perspective: Gravitational allocation for the 2DOCP

In [NSV07], Nazarov-Sodin-Volberg construct the so-called "gradient flow" or "gravitational" allocation for the zeroes the GEF (2.2.1). This provides an explicit coupling between the Lebesgue measure and the random points, obtained by treating the points as massive "stars", and by letting each "grain" of the Lebesgue measure follow a gravitational attractive flow until it "falls" onto one of the stars. It is not hard to prove that each "star" receives a unit mass of "grains", and one turns to study its gravitational cell. Showing that those cells are typically bounded gives a "constructive" proof the $\mathtt{Wass}_2$ property.

It is natural to try to implement the strategy of [NSV07] to the 2DOCP, which bears some resemblance with the behavior of GEF zeroes. This would solve a long-standing question, open even for $\beta = 2$.

A. Appendix: Infinite-volume Coulomb energy

We let g be (a certain multiple of) the Coulomb kernel in dimension d, namely:

$$\mathsf{g}(x) := \begin{cases} -\log|x| & (\mathsf{d}=2) \\ |x|^{-(\mathsf{d}-2)} & (\mathsf{d}\geq 3) \end{cases}.$$

From finite to infinite systems

We recall the basic setup for (one-component) Coulomb systems in finite volume. Let $N \geq 1$ and let $\mathrm{X}_N = (x_1, \ldots, x_N)$ be a N-tuple of points in the box $\Lambda_N := \left[-\frac{N}{2}, \frac{N}{2}\right]^{\mathsf{d}}$. The Coulomb energy of X_N can be defined as:

$$\mathrm{H}_N(\mathrm{X}_N) := \frac{1}{2} \iint\limits_{(x,y)\in\Lambda_N\times\Lambda_N, x\neq y} \mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(x)\mathsf{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(y), \tag{A.0.1}$$

with $\mathbf{X}_N := \sum_{i=1}^N \delta_{x_i}$ a finite point configuration. The usual physical interpretation of (A.0.1) is that the N point charges $x_1, \ldots, x_N$, all living in Λ_N, interact both with each other and with a "neutralizing background" given by the Lebesgue measure on the box Λ_N. The energy H_N is bounded below by a constant times N, but it is not bounded above due to the singularity of g.

The study of systems with pairwise Coulomb interactions is an old topic. However, due to the long-range nature of the Coulomb kernel, properly defining the Coulomb energy (per unit volume) for an infinite point configuration, or for stationary point processes is challenging.

Infinite systems

The notion of a "renormalized energy" was introduced by Sandier-Serfaty [SS12], and applied to two-dimensional Coulomb systems in [SS15], with various extensions (see in particular [RS16] for higher dimensional Coulomb gases and [PS17] for an generalization to Riesz gases). It provides a/the framework to work with infinite Coulomb systems.

Before presenting this formalism, let us briefly evoke what a "naive" approach could look like. In view of the finite interaction energy (A.0.1) and of standard practice in statistical mechanics (for *short-range* interactions!) one could (try to) define the Coulomb energy per unit volume of a point configuration $\mathbf{X}$ as:

$$\lim_{N\to\infty} \frac{1}{|\Lambda_N|}\ \frac{1}{2} \iint\limits_{(x,y)\in\Lambda_N\times\Lambda_N, x\neq y} \mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X} - \mathrm{Leb}\right)(x)\mathsf{d}\left(\mathbf{X} - \mathrm{Leb}\right)(y), \tag{A.0.2}$$

or perhaps more accurately as (note the change in the domain of integration):

$$\lim_{N\to\infty} \frac{1}{|\Lambda_N|}\ \frac{1}{2} \iint\limits_{(x,y)\in\Lambda_N\times\mathbb{R}^{\mathsf{d}}, x\neq y} \mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X} - \mathrm{Leb}\right)(x)\mathsf{d}\left(\mathbf{X} - \mathrm{Leb}\right)(y). \tag{A.0.3}$$

The second expression is in fact closer to the truth, as it takes into account the effect of the whole configuration on the points in Λ_N, but precisely because of this, and the long-range nature of the Coulomb kernel, the finiteness of this expression for fixed N is very much unclear and might only be obtained through a careful control of the overall repartition of charges. Ultimately, it is simply not possible in general to rigorously derive "infinite-volume" energies of the type (A.0.3) from the finite N energy (A.0.1).

The general spirit of the works by Serfaty et al. can be phrased as an "electric approach". Instead of working directly with the pairwise interaction energy $\mathsf{g}(x-y)$ between points as in (A.0.2), (A.0.3), one puts the emphasis on "electric fields", which are intrinsically global objects whose energy is nonetheless

easy to localize. Near each point of the configuration, those fields have a singularity which needs to be taken care of, hence the adjective "renormalized".

In the rest of Section A, we recall the main lines of the "electric approach", our goal being to define the "regularized Coulomb energy" $\mathsf{Coul}_\eta(\mathbf{X}^\bullet)$ ($\eta \in (0,1)$) for a stationary point process $\mathbf{X}^\bullet$, which is presented in Definition A.0.2.

Electric approach in finite volume

The starting point of the "electric" approach to *infinite* Coulomb systems is the following consideration valid for *finite* systems. Let us introduce the electric potential Φ and the *true electric field* E generated by X_N as:

$$\Phi : x \mapsto \int_{y\in\Lambda_N} \mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(y), \quad \mathrm{E}(x) := \nabla\Phi = \int_{y\in\Lambda_N} \nabla\mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X}_N - \mathrm{Leb}\right)(y). \tag{A.0.4}$$

Note that Φ is a scalar field, while E is a (gradient) vector field. The "true electric field" satisfies the distributional identity:

$$-\mathrm{div}(\mathrm{E}) = c_{\mathsf{d}}\left(\mathbf{X}_N - \mathrm{Leb}\right), \text{ with } \mathrm{E}(x) \to 0 \text{ as } |x| \to \infty \tag{A.0.5}$$

for some constant c_{d} depending on d (the value is 2π for $\mathsf{d} = 2$). Returning to (A.0.1), integrating by parts, and neglecting the condition $\{x \neq y\}$ in the integral, one would get:

$$\mathrm{H}_N(\mathrm{X}_N) \approx -\frac{1}{2c_{\mathsf{d}}}\int_{\mathbb{R}^{\mathsf{d}}} \Phi\ \mathrm{div}(\mathrm{E}) \approx \frac{1}{2c_{\mathsf{d}}}\int_{\mathbb{R}^{\mathsf{d}}} |\mathrm{E}|^2. \tag{A.0.6}$$

The heuristic identity (A.0.6) does not make sense as such because E fails to be in L^2 due to the point charges self-interactions (which are avoided in (A.0.1) by imposing the condition $\{x \neq y\}$). However, (A.0.6) points to the fact (which is old and common knowledge in physics) that the Coulomb interaction energy can be phrased in terms of the integral of an "electric" energy density $|\mathrm{E}|^2$, where E is the right "electric field" satisfying (A.0.5).

Electric fields: global, local, screened

We refer to vector fields in $\cap_{p\in[1,2)} L^p_{loc}(\mathbb{R}^{\mathsf{d}}, \mathbb{R}^{\mathsf{d}})$ as (global) "electric fields". If $\mathbf{X}$ is a point configuration in $\mathbb{R}^{\mathsf{d}}$, and E is an electric field, we say that E is compatible with $\mathbf{X}$ when the following identity holds:

$$-\mathrm{div}(\mathrm{E}) = c_{\mathsf{d}}\left(\mathbf{X} - \mathrm{Leb}\right). \tag{A.0.7}$$

We also work with *local* electric fields: if Λ is a square in $\mathbb{R}^{\mathsf{d}}$ and $\mathbf{X}$ a point configuration in Λ, we say that $\mathrm{E} \in \cap_{p\in[1,2)} L^p_{loc}(\Lambda, \mathbb{R}^{\mathsf{d}})$ is a (local) electric field compatible with $\mathbf{X}$ when (A.0.7) holds in Λ. To be clear, by (A.0.7) we mean that

$$\int \varphi(x)\mathsf{d}\left(\mathbf{X} - \mathrm{Leb}\right)(x) = -\frac{1}{c_{\mathsf{d}}}\int \nabla\varphi \cdot \mathrm{E}$$

for all $\varphi \in C^\infty_c(\mathbb{R}^{\mathsf{d}})$ (global electric fields), or for all $\varphi \in C^\infty_c(\Lambda)$ (local electric fields in Λ). We say that a local electric field E is *screened* when it has Neumann boundary conditions i.e. its normal component satisfies

$$\mathrm{E} \cdot \vec{n} = 0 \text{ along } \partial\Lambda.$$

The importance of screened fields Screened fields are of outmost importance because they can be glued together in the following sense: if Λ^1, Λ^2 are two squares with one side in common, $\mathbf{X}^1, \mathbf{X}^2$ are two point configurations in Λ^1, Λ^2 and $\mathrm{E}^1, \mathrm{E}^2$ are two local electric fields compatible with $\mathbf{X}^1, \mathbf{X}^2$, then defining:

$$\mathbf{X}^t := \mathbf{X}^1 + \mathbf{X}^2, \quad \mathrm{E}^t := \mathrm{E}^1 1_{\Lambda_1} + \mathrm{E}^2 1_{\Lambda_2}$$

yields a "total" point configuration and a "total" electric field, but *in general,* E^t *is not compatible with* $\mathbf{X}^t$. However, if $\mathrm{E}^1, \mathrm{E}^2$ are both screened, or more generally if *their normal components agree on the common side of* Λ_1, Λ_2, then E^t is compatible with $\mathbf{X}^t$.

The role of abstract fields An important remark is that we deal with "abstract" fields E, not necessarily of the form (cf. (A.0.4)):

$$\mathrm{E}(x) \stackrel{?}{=} \int_{\mathbb{R}^{\mathsf{d}}} \nabla \mathsf{g}(x-y)\mathsf{d}\left(\mathbf{X}-\mathrm{Leb}\right)(y),$$

which could be a candidate to solve (A.0.7), but whose convergence is unclear (note that $\nabla \mathsf{g}(x) \propto \frac{1}{|x|^{\mathsf{d}-1}}$ is far from being integrable at infinity). Considering the whole family of compatible electric fields is a fundamental feature of the approach by Sandier-Serfaty. The price to pay is that in general we do not know "who E is", besides that it solves (A.0.7).

Spreading out charges and regularizing point configurations.

To handle the singularity of the electric field near a point charge, one introduces a short-scale regularization, or "truncation" of the field, and a corresponding operation of "spreading out" the point charges. Let us start with the latter, as it is simpler to describe. The usual way to proceed is as follows (see [RS16, Sec. 2.1]):

1. Take some radial non-negative "bump function" χ supported on the unit ball, with mass 1, and for $\eta > 0$ let
$$\chi_\eta := \frac{1}{\eta^{\mathsf{d}}}\chi\left(\frac{\cdot}{\eta}\right).$$
In general, the choice of the bump function χ is not really important and does not affect the limiting objects or the values of the limiting quantities as $\eta \to 0$. A common choice is to take $\chi = 1_{\mathrm{B}_1}$, which would work for us here, but we may also take χ to be smooth. We fix such a χ for the rest of the paper.
2. If $x \in \mathbb{R}^{\mathsf{d}}$, define $\delta_x^{(\eta)}$ as a Dirac mass "spread-out" according to χ_η, namely as the measure with density $\chi_\eta(\cdot - x)$. If we think of a Dirac mass δ_x as a point charge placed at x, then $\delta_x^{(\eta)}$ is a "spread-out" charge at x.
This operation goes back at least to Onsager [Ons39], see e.g. [LL01, Sec. 9.7] for a study of its properties in the context of Coulomb energies. In short:
   a) By Newton's theorem, the field generated by a spread-out charge coincides with the field generated by the original charge outside the support of the "spread" (here B_η).
   b) Spreading out all charges slightly decreases the pairwise interaction energy between them.
   c) After this operation, a single charge self-interaction becomes *finite* - but it blows up proportionally to $\mathsf{g}(\eta)$ as $\eta \to 0$, as one could expect.
3. If $\mathbf{X}$ is a point configuration, replacing each Dirac mass δ_x by its "spread-out" version $\delta_x^{(\eta)}$ for $x \in \mathbf{X}$ yields a "regularized" version of $\mathbf{X}$, denoted by $\mathbf{X}_\eta$.

Formally, we have defined two unambiguous operations:

$$\delta_x \to \delta_x^{(\eta)} := \chi_\eta * \delta_x \text{ (for } x \in \mathbb{R}^{\mathsf{d}}), \quad \mathbf{X} \to \mathbf{X}_\eta := \sum_{x\in\mathbf{X}} \delta_x^{(\eta)}.$$

Let us emphasize that $\mathbf{X} \to \mathbf{X}_\eta$ is an operation at the level of Radon measures on $\mathbb{R}^{\mathsf{d}}$. The image of a point configuration $\mathbf{X}$ is no longer purely atomic, it is a locally finite measure with locally bounded density with respect to the Lebesgue measure (the density of one regularized Dirac is bounded by $\|\chi\|_\infty \eta^{-\mathsf{d}}$ but the density of $\mathbf{X}_\eta$ depends also on the number of points per unit volume).

Truncation of the field. We now study a corresponding operation at the level of electric fields. With χ chosen as before and for $\eta > 0$, let f_η be the mean-zero solution to:

$$-\Delta \mathsf{f}_\eta = c_{\mathsf{d}}\left(\delta_0^{(\eta)} - \delta_0\right), \quad \mathsf{f}_\eta \equiv 0 \text{ outside } \mathrm{B}_\eta,$$

which is given by $\mathsf{f}_\eta : x \mapsto \max\left(0, \log\frac{\eta}{|x|}\right)$. If E is an electric field (global or local) compatible with $\mathbf{X}$, we define the "truncated field" E_η as:

$$\mathrm{E}_\eta := \mathrm{E} + \sum_{x\in\mathbf{X}} \nabla \mathsf{f}_\eta(x - \cdot). \tag{A.0.8}$$

Claim A.0.1 (Properties of the truncated field)**.** *It is straightforward that:*

1. *If* E *is compatible with* $\mathbf{X}$ *in the sense of* (A.0.7)*, then* E_η *is compatible with* $\mathbf{X}_\eta$ *in the sense:*

$$-\mathrm{div}(\mathrm{E}_\eta) = c_{\mathsf{d}}\,(\mathbf{X}_\eta - \mathrm{Leb})\,. \tag{A.0.9}$$

2. *If* E *is a gradient vector field, then so is* E_η*.*
3. *(Local version.) If* E *is compatible with* $\mathbf{X}$ *in* Λ *and* $\min_{x\in\mathbf{X}} \mathrm{dist}(x,\partial\Omega) > \eta$*, then* (A.0.9) *holds in* Λ*. Moreover, if* E *is "screened", then so is* E_η*.*

Stationary random electric fields

We call a random electric field $\mathrm{E}^\bullet$ stationary when its law is invariant under translations $\mathrm{E} \to \mathrm{E}(\cdot - t)$ for all $t \in \mathbb{R}^{\mathsf{d}}$. We say that a random electric field $\mathrm{E}^\bullet$ is compatible with a point process $\mathbf{X}^\bullet$ when the random measure $\frac{1}{c_{\mathsf{d}}}\mathrm{div}(\mathrm{E}^\bullet) + \mathrm{Leb}$ and $\mathbf{X}^\bullet$ have the same distribution. Note however that if $\mathbf{X}^\bullet$ is stationary, it does not imply that so is $\mathrm{E}^\bullet$.

Regularized Coulomb energy

Finally, we may define the notion that will be relevant for us, namely "finite regularized Coulomb energy".

Definition A.0.2. Let $\mathbf{X}^\bullet$ be a stationary point process. We say that it has *finite regularized Coulomb energy* when there exists $\eta \in (0,1]$ and a *stationary* random electric field $\mathrm{E}^\bullet$ compatible with $\mathbf{X}^\bullet$ such that, using the truncation/regularization presented in (A.0.8):

$$\mathbb{E}\left[|\mathrm{E}^\bullet_\eta(0)|^2\right] < +\infty.$$

Since stationary, compatible electric fields are not unique, we define:

$$\mathsf{Coul}_\eta(\mathbf{X}^\bullet) := \inf_{\mathrm{E}^\bullet \text{ stat., comp. with } \mathbf{X}^\bullet} \mathbb{E}\left[|\mathrm{E}^\bullet_\eta(0)|^2\right].$$

Renormalized energy

For completeness, we conclude by recalling the definition of the "renormalized" energy of Sandier-Serfaty. The renormalized Coulomb energy of a stationary random electric field $\mathrm{E}^\bullet$ is obtained by taking the limit $\eta \to 0$ of $\mathbb{E}\left[|\mathrm{E}^\bullet_\eta(0)|^2\right]$ in a renormalized fashion, namely by substracting a divergent part of order $\mathsf{g}(\eta)$ (the following limit as η to 0 exists by "almost"-monotonicity, see [PS17, Sec. 2.3]):

$$\mathsf{Coul}(\mathrm{E}^\bullet) := \lim_{\eta\to 0} \left(\mathbb{E}\left[|\mathrm{E}^\bullet_\eta(0)|^2\right] - c_{\mathsf{d}}\mathsf{g}(\eta)\right),$$

and the Coulomb energy of a point process is then again phrased as an infimum:

$$\mathsf{Coul}(\mathbf{X}^\bullet) := \inf_{\mathrm{E}^\bullet \text{ stat., comp. with } \mathbf{X}^\bullet} \mathsf{Coul}(\mathrm{E}^\bullet).$$

The fact that this is the "correct" notion of an infinite-volume Coulomb energy is supported by a result of Γ-convergence connecting the finite-volume energy H_N (as defined in (A.0.1)) to Coul. We refer e.g. to [LS17, Sec. 2.7 & 3.] for more details.

Bibliography

[AHM11] Y. Ameur, H. Hedenmalm, and N. Makarov. Fluctuations of eigenvalues of random normal matrices. *Duke Math. J.*, 159(1):31–81, 2011.

[AJ81] A. Alastuey and B. Jancovici. On the classical two-dimensional one-component coulomb plasma. *J. Phys.*, 42(1):1–12, 1981.

[Ame18] Y. Ameur. Repulsion in low temperature β-ensembles. *Comm. Math. Phys.*, 359(3):1079–1089, 2018.

[AS21] S. Armstrong and S. Serfaty. Local laws and rigidity for Coulomb gases at any temperature. *Ann. Probab.*, 49(1), 2021.

[BBNY19] R. Bauerschmidt, P. Bourgade, M. Nikula, and H.-T. Yau. The two-dimensional Coulomb plasma: quasi-free approximation and central limit theorem. *Adv. Theor. Math. Phys.*, 23(4):841–1002, 2019.

[Bec87] J. Beck. Irregularities of distribution. I. *Acta Math.*, 159:1–49, 1987.

[Ber19] R. J. Berman. Sharp deviation inequalities for the 2d coulomb gas and quantum hall states, i. *arXiv preprint arXiv:1906.08529*, 2019.

[BEY14] P. Bourgade, L. Erdős, and H.-T. Yau. Universality of general β-ensembles. *Duke Math. J.*, 163(6):1127–1190, 2014.

[BL18] F. Bekerman and A. Lodhia. Mesoscopic central limit theorem for general β-ensembles. 2018.

[BLS18] F. Bekerman, T. Leblé, and S. Serfaty. Clt for fluctuations of β-ensembles with general potential. *Electron. J. Probab.*, 23(none), January 2018.

[BNQ19] A. I. Bufetov, P. P. Nikitin, and Y. Qiu. On number rigidity for Pfaffian point processes. *Mosc. Math. J.*, 19(2):217–274, 2019.

[BS17] A. I. Bufetov and T. Shirai. Quasi-symmetries and rigidity for determinantal point processes associated with de Branges spaces. *Proc. Japan Acad.*, 93(1):1–5, 2017.

[CC83] P. Choquard and J. Clerouin. Cooperative phenomena below melting of the one-component two-dimensional plasma. *Phys. Rev. Lett.*, 50(26):2086, 1983.

[CHM18] D. Chafai, A. Hardy, and M. Maïda. Concentration for coulomb gases and coulomb transport inequalities. *J. Funct. Anal.*, 275(6):1447–1483, 2018.

[CLWH82] J. M. Caillol, D. Levesque, J. J. Weis, and J. P. Hansen. A monte carlo study of the classical two-dimensional one-component plasma. *J. Stat. Phys.*, 28:325–349, 1982.

[CN18] R. Chhaibi and J. Najnudel. Rigidity of the Sine-β process. *Electron. Commun. Probab.*, 23, 2018.

[Cos21] S. Coste. Order, fluctuations, rigidities. 2021.

[DE02] I. Dumitriu and A. Edelman. Matrix models for beta ensembles. *J. Math. Phys.*, 43(11):5830–5847, 2002.

[DFHL24] D. Dereudre, D. Flimmel, M. Huesmann, and T. Leblé. (non)-hyperuniformity of perturbed lattices. *arXiv preprint arXiv:2405.19881*, 2024.

[DHLM21] D. Dereudre, A. Hardy, T. Leblé, and M. Maïda. DLR equations and rigidity for the Sine-beta process. *Comm. Pure Appl. Math.*, 74(1):172–222, 2021.

[DMPS95] A. B. De Monvel, L. Pastur, and M. Shcherbina. On the statistical mechanics approach in the random matrix theory: integrated density of states. *J. Stat. Phys.*, 79(3):585–611, 1995.

[Dys62] F. J. Dyson. Statistical theory of the energy levels of complex systems. i. *J. Math. Phys.*, 3(1):140–156, 1962.

[EH15] M. Erbar and M. Huesmann. Curvature bounds for configuration spaces. *Calc. Var. Partial Differential Equations*, 54(1):397–430, 2015.

[EHJM25] M. Erbar, M. Huesmann, J. Jalowy, and B. Müller. Optimal transport of stationary point processes: Metric structure, gradient flow and convexity of the specific entropy. *Journal of Functional Analysis*, 289(4):110974, 2025.

[EHL21] M. Erbar, M. Huesmann, and T. Leblé. The one-dimensional log-gas free energy has a unique minimizer. *Comm. Pure Appl. Math.*, 74(3):615–675, January 2021.

[FL21] M. Fenzl and G. Lambert. Precise Deviations for Disk Counting Statistics of Invariant Determinantal Processes. *Int. Math. Res. Not. IMRN*, 2022(10):7420–7494, 01 2021.

[For10] P. J. Forrester. *Log-gases and random matrices (LMS-34)*. Princeton University Press, 2010.

[FP81] J. Fröhlich and C. Pfister. On the absence of spontaneous symmetry breaking and of crystalline ordering in two-dimensional systems. *Comm. Math. Phys.*, 81:277–298, 1981.

[FP86] J. Fröhlich and C.-E. Pfister. Absence of crystalline ordering in two dimensions. *Comm. Math. Phys.*, 104:697–700, 1986.

[Gab04] A. Gabrielli. Point processes and stochastic displacement fields. *Phys. Rev. E*, 70(6):066131, 2004.

[Geo99] H.-O. Georgii. Translation invariance and continuous symmetries in two-dimensional continuum systems. In *Mathematical results in statistical mechanics (Marseilles, 1998)*, pages 53–69. World Sci. Publ., River Edge, NJ, 1999.

[Gin65] J. Ginibre. Statistical ensembles of complex, quaternion, and real matrices. *J. Math. Phys.*, 6(3):440–449, 1965.

[GJL02] A. Gabrielli, M. Joyce, and F. S. Labini. Glass-like universe: Real-space correlation properties of standard cosmological models. *Phys. Rev. D*, 65(8):083523, 2002.

[GL17] S. Ghosh and J. L. Lebowitz. Fluctuations, large deviations and rigidity in hyperuniform systems: a brief survey. *Indian J. Pure Appl. Math.*, 48(4):609–631, 2017.

[GP17] S. Ghosh and Y. Peres. Rigidity and tolerance in point processes: Gaussian zeros and ginibre eigenvalues. *Duke Math. J.*, 166(10):1789–1858, 2017.

[GS75] P. Gacs and D. Szász. On a problem of Cox concerning point processes in $\mathbb{R}^k$ of "controlled variability". *Ann. Probab.*, 3(4):597–607, 1975.

[GZ19] D. García-Zelada. Concentration for coulomb gases on compact manifolds. *Electron. Commun. Probab.*, 24, 2019.

[HKPV06] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág. Determinantal processes and independence. *Probab. Surv.*, 3:206–229, 2006.

[HL25] M. Huesmann and T. Leblé. The link between hyperuniformity, coulomb energy and wasserstein distance to lebesgue for two-dimensional point processes. *Probability and Mathematical Physics*, 7(1):123–173, 2025.

[HS13] M. Huesmann and K.-T. Sturm. Optimal transport from Lebesgue to Poisson. *Ann. Probab.*, 41(4):2426–2478, 2013.

[Hue16] M. Huesmann. Optimal transport between random measures. *Ann. Inst. Henri Poincaré Probab. Stat.*, 52(1):196–232, 2016.

[HV17] D. Holcomb and B. Valkó. Overcrowding asymptotics for the sine_β process. 2017.

[JLM93] B. Jancovici, J. L. Lebowitz, and G. Manificat. Large charge fluctuations in classical Coulomb systems. *J. Stat. Phys.*, 72:773–787, 1993.

[Joh98] K. Johansson. On fluctuations of eigenvalues of random Hermitian matrices. *Duke Math. J.*, 91(1):151–204, 1998.

[KS09] R. Killip and M. Stoiciu. Eigenvalue statistics for CMV matrices: from Poisson to clock via random matrix ensembles. *Duke Math. J.*, 146(3):361–399, 2009.

[KT18] J. Kim and S. Torquato. Effect of imperfections on the hyperuniformity of many-body systems. *Phys. Rev. B*, 97(5):054105, 2018.

[KVV12] E. Kritchevski, B. Valkó, and B. Virág. The scaling limit of the critical one-dimensional random schrödinger operator. *Comm. Math. Phys.*, 314(3), 2012.

[Leb83] J. Lebowitz. Charge fluctuations in coulomb systems. *Phys. Rev. A*, 27(3):1491, 1983.

[Leb16] T. Leblé. Logarithmic, Coulomb and Riesz energy of point processes. *J. Stat. Phys.*, 162(4):887–923, 2016.

[Leb19] T. Leblé. Clt for fluctuations of linear statistics in the sine-beta process. *Int. Math. Res. Not. IMRN*, 2021(8):5676–5756, February 2019.

[Leb24] T. Leblé. DLR equations, number-rigidity and translation-invariance for infinite-volume limit points of the 2DOCP. *arXiv preprint arXiv:2410.04958*, 2024.

[Leb26] T. Leblé. The two-dimensional one-component plasma is hyperuniform. *Duke Math. J.*, 175(5):763–901, 2026.

[Led19] M. Ledoux. On optimal matching of Gaussian samples. *J. Math. Sci.*, 238(4):495–522, March 2019.

[Lew22] M. Lewin. Coulomb and Riesz gases: The known and the unknown. *J. Math. Phys.*, 63(6), 2022.

[LL01] E. H. Lieb and M. Loss. *Analysis*, volume 14. American Mathematical Soc., 2001.

[LRY19] E. H. Lieb, N. Rougerie, and J. Yngvason. Local incompressibility estimates for the laughlin phase. *Comm. Math. Phys.*, 365(2):431–470, 2019.

[LRY24] R. Lachièze-Rey and D. Yogeshwaran. Hyperuniformity and optimal transport of point processes. *arXiv preprint arXiv:2402.13705*, 2024.

[LS17] T. Leblé and S. Serfaty. Large deviation principle for empirical fields of log and Riesz gases. *Invent. Math.*, 210(3):645–757, 2017.

[LS18] T. Leblé and S. Serfaty. Fluctuations of two dimensional Coulomb gases. *Geom. Funct. Anal.*, 28(2):443–508, 2018.

[LWL00] D. Levesque, J.-J. Weis, and J. Lebowitz. Charge fluctuations in the two-dimensional one-component plasma. *J. Stat. Phys.*, 100(1):209–222, 2000.

[Mar88] P. A. Martin. Sum rules in charged fluids. *Rev. Modern Phys.*, 60(4):1075, 1988.

[McC97] R. McCann. A Convexity Principle for Interacting Gases. *Adv. Math.*, 128(1):153 – 179, 1997.

[MY80] P. A. Martin and T. Yalcin. The charge fluctuations in classical Coulomb systems. *Phys. Rev. A*, (22):435–463, 1980.

[Nak14] F. Nakano. Level statistics for one-dimensional schrödinger operators and gaussian beta ensemble. *J. Stat. Phys.*, 156(1):66–93, 2014.

[NSV07] F. Nazarov, M. Sodin, and A. Volberg. Transportation to random zeroes by the gradient flow. *Geom. Funct. Anal.*, 17(3):887–935, 2007.

[NSV08] F. Nazarov, M. Sodin, and A. Volberg. The jancovici–lebowitz–manificat law for large fluctuations of random complex zeroes. *Comm. Math. Phys.*, 284(3):833–865, 2008.

[Ons39] L. Onsager. Electrostatic interaction of molecules. *J. Phys. Chem.*, 43(2):189–196, 1939.

[OS08] H. Osada and T. Shirai. Variance of the linear statistics of the ginibre random point field. In *RIMS Kôkyûroku Bessatsu: Proceedings of RIMS Workshop on Stochastic Analysis and Applications*, volume B6, pages 193–200, 2008.

[Pei23] L. Peilen. Local laws and a mesoscopic CLT for β-ensembles. *Comm. Pure Appl. Math.*, 77(4):2452–2567, October 2023.

[PG23] D. Padilla-Garza. Concentration inequality around the thermal equilibrium measure of coulomb gases. *J. Funct. Anal.*, 284(1):109733, 2023.

[PS17] M. Petrache and S. Serfaty. Next order asymptotics and renormalized energy for Riesz interactions. *J. Inst. Math. Jussieu*, 16(3):501–569, 2017.

[Ric07] T. Richthammer. Translation-invariance of two-dimensional Gibbsian point processes. *Comm. Math. Phys.*, 274:81–122, 2007.

[RRV11] J. Ramirez, B. Rider, and B. Virág. Beta ensembles, stochastic airy spectrum, and a diffusion. *J. Amer. Math. Soc.*, 24(4):919–944, 2011.

[RS16] N. Rougerie and S. Serfaty. Higher-dimensional coulomb gases and renormalized energy functionals. *Comm. Pure Appl. Math.*, 69(3):519–605, 2016.

[Ser23] S. Serfaty. Gaussian fluctuations and free energy expansion for Coulomb gases at any temperature. In *Annales de l'Institut Henri Poincare (B) Probabilites et statistiques*, volume 59, pages 1074–1142, 2023.

[Ser24] S. Serfaty. Lectures on Coulomb and Riesz gases. *arXiv preprint arXiv:2407.21194*, 2024.

[Shi06] T. Shirai. Large deviations for the fermion point process associated with the exponential kernel. *J. Stat. Phys.*, 123(3):615–629, 2006.

[Sim14] B. Simon. The statistical mechanics of lattice gases, volume i. In *The Statistical Mechanics of Lattice Gases, Volume I.* Princeton University Press, 2014.

[SM76] R. R. Sari and D. Merlini. On the ν-dimensional one-component classical plasma: the thermodynamic limit problem revisited. *J. Stat. Phys.*, 14(2):91–100, 1976.

[SS12] É. Sandier and S. Serfaty. From the Ginzburg-Landau model to vortex lattice problems. *Comm. Math. Phys.*, 313(3):635–743, 2012.

[SS15] E. Sandier and S. Serfaty. 2D Coulomb gases and the renormalized energy. *Ann. Probab.*, 43(4):2026–2083, 2015.

[SWY23a] M. Sodin, A. Wennman, and O. Yakir. The random weierstrass zeta function ii. fluctuations of the electric flux through rectifiable curves. *J. Stat. Phys.*, 190(10):164, 2023.

[SWY23b] M. Sodin, A. Wennman, and O. Yakir. The random Weierstrass Zeta function i: Existence, uniqueness, fluctuations. *J. Stat. Phys.*, 190(10):166, 2023.

[Tho24] E. Thoma. Overcrowding and separation estimates for the coulomb gas. *Comm. Pure Appl. Math.*, 77(7):3227–3276, 2024.

[Tor18] S. Torquato. Hyperuniform states of matter. *Phys. Rep.*, 745:1–95, 2018.

[TS03] S. Torquato and F. H. Stillinger. Local density fluctuations, hyperuniformity, and order metrics. *Phys. Rev. E*, 68(4):041113, 2003.

[VV09] B. Valkó and B. Virág. Continuum limits of random matrices and the Brownian carousel. *Invent. Math.*, 177(3):463–508, 2009.

[VV17] B. Valkó and B. Virág. The Sine$_\beta$ operator. *Invent. Math.*, 209(1):275–327, 2017.

[VV20] B. Valkó and B. Virág. Operator limit of the circular beta ensemble. *Ann. Probab.*, 48(3):1286–1316, 2020.

[VV22] B. Valkó and B. Virág. The many faces of the stochastic zeta function. *Geom. Funct. Anal.*, 32(5):1160–1231, 2022.

[ZW06] A. Zabrodin and P. Wiegmann. Large-n expansion for the 2d dyson gas. *J. Phys. A*, 39(28):8933, 2006.

Life set me larger—problems—
Some I shall keep—to solve
Till Algebra is easier—
Or simpler proved—above—

Emily Dickinson