SIMULATION OF THE TWO-DIMENSIONAL UNIFORM ELECTRON GAS.

RODRIGUE LELOTTE

ABSTRACT. In this document, we study both from a theoretical and numerical perspective the *Multi-marginal Optimal Transport* (MOT) problem with transportation cost given by the *Coulomb cost* in two space-dimensions, and where the N marginals are all equal to some probability measure μ, that is

$$\inf_{\pi \in \Pi^N(\mu)} \mathbb{E}_{(X_1,\dots,X_N)\sim\pi} \left[- \sum_{1 \leqslant i < j \leqslant N} \log |X_i - X_j| \right].$$

This problem is of paramount interest in a variety of physical applications (see Section 3), for instance in defining the *Uniform Electron Gas* (UEG), an important model in condensed matter physics, which serves as the fundamental building block of the *Local Density Approximation* in *Density Functional Theory*, a successful set of computational methods in quantum chemistry. More generically, this problem is essential in the study of the two-dimensional *Coulomb gases*, a very relevant topic in mathematical physics, which shares a substantial link with random matrix theory. In the recent years, MOT has drawn a convincing attention across the mathematical literature (and outside of it), ranging from the field of pure mathematics to applied statistical learning. As such, our investigations could have potential offsprings in a large collection of mathematical fields.

We first present general theoretical results for this problem (Section 2) and its entropy-regularized version (Section 4). The main difficulty with the Coulomb cost in two space-dimensions is its divergence at infinity. Several recent results have considered costs which diverge at the origin, but few seems to cover both divergences. We thus had to adapt many existing works to the logarithmic case. In fact we will most often present our own take on the problem, providing proofs that differ from the ones proposed in the literature. We then present an original numerical scheme in order to solve the entropy-regularized version of this problem, together with some theoretical results. This iterative scheme solves a kind of inverse problem for the dual (Kantorovich) potential using Monte Carlo simulation. We then conduct numerical experiments (Section 5) using parallel computing. The results, reminiscent of an important conundrum in mathematical physics regarding the equality of the UEG with another model, namely *Jellium*, are promising, and will lead to further investigations.

1. INTRODUCTION

Optimal Transport (OT) is a very active field of research in mathematics [87, 98, 99]. In the recent years, this field has gained considerable attention, especially from a numerical perspective [76], because of its high potential to be applied in many practical situations. Recall that, given two (Polish) spaces $\mathscr{X}$ and $\mathscr{Y}$ equipped with two probability measures $\mu \in \mathcal{P}(\mathscr{X})$ and $\nu \in \mathcal{P}(\mathscr{Y})$, together with some

Date: November 9, 2020.

abstract cost function $c : \mathscr{X} \times \mathscr{Y} \to \mathbb{R} \cup \{\infty\}$, the *Kantorovich formulation* of OT reads

$$K_c(\mu, \nu) = \inf_{\pi \in \Pi(\mu,\nu)} \mathbb{E}_{(X,Y)\sim\pi} \left[c(X, Y)\right], \tag{K}$$

where $\Pi(\mu, \nu)$ is the set of *couplings* of μ and ν, that is the set of all probability measures $\pi \in \mathcal{P}(\mathscr{X} \times \mathscr{Y})$ such that the first (resp. second) marginal of π is μ (resp. ν). That is, given any continuous, real-valued and bounded function $\phi \in C_b(\mathscr{X})$ (resp. $\psi \in C_b(\mathscr{Y})$), we have $\mathbb{E}_{(X,Y)\sim\pi}[\phi(X)] = \mathbb{E}_{X\sim\mu}[\phi(X)]$ (resp. $\mathbb{E}_{(X,Y)\sim\pi}[\psi(Y)] = \mathbb{E}_{Y\sim\nu}[\psi(Y)]$). The Kantorovich formulation to OT is a relaxation of the historical *Monge formulation* of OT, which reads

$$M_c(\mu, \nu) = \inf_{T \in \Pi_d(\mu,\nu)} \mathbb{E}_{X\sim\mu} \left[c(X, T(X))\right], \tag{M}$$

where $\Pi_d(\mu, \nu)$ is the set of all measurable maps $T : \mathscr{X} \to \mathscr{Y}$ such that $T_\#\mu = \nu$, where the *push-forward* measure $T_\#\mu$ is defined for every measurable set $A \subset \mathscr{Y}$ as $T_\#\mu(A) := \mu(T^{-1}(A))$. That is, intuitively speaking, we want to find a map T which *transports* μ onto ν while minimizing a cost of transportation defined by c. Given any $T \in \Pi_d(\mu, \nu)$, one can consider the measure $(\mathrm{Id}, T)_\#\mu$ as a *deterministic* element of $\Pi(\mu, \nu)$, in the sense that no *mass-splitting* occurs. Hence, one necessarily has $K_c(\mu, \nu) \leqslant M_c(\mu, \nu)$, but the converse statement need not be true without further assumptions (*e.g.* c is continuous).

In the very recent years, *Multi-marginal Optimal Transport* (MOT) has begun to attract considerable attention, due to a wide variety of emerging applications outside of mathematics, such as economics, finance, physics and image processing (see [72] for a rather detailed review and citations therein). MOT is nothing but an extension of the classical OT to a multi-marginal setting: that is, given N (Polish) spaces $\mathscr{X}_1, \dots, \mathscr{X}_N$ equipped with probability measures $\mu_1 \in \mathcal{P}(\mathscr{X}_1), \dots, \mu_N \in \mathcal{P}(\mathscr{X}_N)$, together with some abstract cost function $c : \mathscr{X}_1 \times \cdots \times \mathscr{X}_N \to \mathbb{R} \cup \{\infty\}$, the Kantorovich formulation of MOT reads

$$\inf_{\pi \in \Pi(\mu_1,\dots,\mu_N)} \mathbb{E}_{(X_1,\dots,X_N)\sim\pi} \left[c(X_1, \dots, X_N)\right], \tag{mK}$$

where evidently $\Pi(\mu_1, \dots, \mu_N)$ is the set of all probability measures $\pi \in \mathcal{P}(\mathscr{X}_1 \times \cdots \times \mathscr{X}_N)$ such that the i-th marginal of π is μ_i for every $i = 1, \dots, N$. Likewise, one can also consider the corresponding Monge formulation of MOT, which reads

$$\inf_{(T_2,\dots,T_N)\in\Pi_d(\mu_1,\dots,\mu_N)} \mathbb{E}_{X\sim\mu_1} \left[c(X, T_2(X), \dots, T_N(X))\right], \tag{mM}$$

where the infimum runs over the measurable maps $T_i : \mathscr{X}_1 \to \mathscr{X}_i$ such that $T_{i\#}\mu_1 = \mu_i$ for all $i = 2, \dots, N$. Note that under rather mild assumptions on the cost function c (*e.g.* lower semi-continuous and bounded from below), it is not difficult to prove that the infimum in (mK) (and subsequently in (K)) is attained. Nevertheless, proving the existence (or non-existence) of deterministic solutions, that is solutions to (M) or (mM), is a much more delicate question, especially in a multi-marginal setting.

In physics, MOT appears in a variety of applications. For instance, computing the *Coulomb* (or *electrostatic*) *energy* of a quantum system of N charged particles $\mathbf{x}_1, \dots, \mathbf{x}_N$ in $\mathbb{R}^d$ with fixed *electronic density* $\rho(\mathbf{r})$, defined as the expectation value of the charge density at locus $\mathbf{r}$, amounts to solve a MOT problem of the

form (mK) with N marginals. Given a N-particle probability $\mathbb{P} \in \mathcal{P}(\mathbb{R}^{dN})$, the corresponding electrostatic energy $\mathfrak{C}(\mathbb{P})$ is given by the expectation value

$$\mathfrak{C}(\mathbb{P}) = \mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}}\left[\sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l)\right],$$

where c_d is the Coulomb potential in dimension d (see Remark 1). The system being prescribed an electronic density ρ (with $\rho(\mathbf{r}) \geqslant 0$ for every $\mathbf{r}$ and $\int_{\mathbb{R}^d}\rho(\mathbf{r})\mathrm{d}\mathbf{r} = N$), we only consider the N-particle probabilities $\mathbb{P}$ such that $\rho_{\mathbb{P}} = \rho$, where the electronic density $\rho_{\mathbb{P}}$ associated to $\mathbb{P}$ is evidently defined as the sum of one-particle marginals

$$\rho_{\mathbb{P}}(\mathbf{r}) := \int_{\mathbb{R}^{d(N-1)}} \mathrm{d}\mathbb{P}(\mathbf{r},\mathbf{x}_2,\dots,\mathbf{x}_N) + \cdots + \int_{\mathbb{R}^{d(N-1)}} \mathrm{d}\mathbb{P}(\mathbf{x}_1,\mathbf{x}_2,\dots,\mathbf{r}).$$

As dictated by the general laws of physics, the particles want to occupy the (not necessarily unique) classical state $\mathbb{P}$ yielding the lowest possible electrostatic energy, so that we are left to solve the following minimization problem

$$\inf_{\substack{\mathbb{P}\in\mathcal{P}(\mathbb{R}^{dN}),\\ \rho_{\mathbb{P}}=\rho}} \left\{\mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}}\left[\sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l)\right]\right\}.$$

Because of the nature of quantum measurements, the particles are usually considered to be *indistinguishable*, which mathematically translates into the fact that the admissible quantum states $\mathbb{P}$ must be symmetric with respect to the $\mathbf{x}_i$'s. With this symmetry constraint, the electronic density is proportional to the one-particle marginal of $\mathbb{P}$,

$$\rho_{\mathbb{P}}(\mathbf{r}) = N\int_{\mathbb{R}^{d(N-1)}} \mathrm{d}\mathbb{P}(\mathbf{r},\mathbf{x}_2,\dots,\mathbf{x}_N).$$

Therefore, at the exception of the symmetry constraint imposed on $\mathbb{P}$, we recognize the Kantorovich formulation to MOT as defined at (mK) with all marginals equal to $\frac{\rho}{N}$ and cost function c_d given by the *Coulomb cost*

$$c_d(\mathbf{x}_1,\dots,\mathbf{x}_N) = \sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l).$$

In fact, the symmetry constraint plays no role whatsoever, since for given any $\mathbb{P} \in \mathcal{P}(\mathbb{R}^{dN})$, the symmetric probability $\widetilde{\mathbb{P}} = \frac{1}{N!}\sum_{\sigma\in\mathfrak{S}_N}\sigma_{\#}\mathbb{P}$ yields the same cost as $\mathbb{P}$ because of the symmetry of the Coulomb cost. Therefore, the problem exactly rewrites as

$$\inf_{\mathbb{P}\in\Pi(\frac{\rho}{N},\dots,\frac{\rho}{N})} \left\{\mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}}\left[\sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l)\right]\right\}.$$

In this document, we will be interested both in the theoretical and numerical aspects of this problem in two space-dimensions ($d = 2$). Ultimately, our goal is to show the equivalence between two models used by physicists in condensed matter physics, namely the *uniform electron gas* and the *jellium*, that we will introduce in full details later on.

Remark 1. In its full generality, the Coulomb potential can be derived in any dimension of space from *Gauss's law*, which states that the total electric flux through

any closed surface is proportional to the total electric charge enclosed by the surface. In its differential form, this law reads $\nabla \cdot \vec{\mathbf{E}}(\mathbf{r}) = |\mathbb{S}^{d-1}|\rho(\mathbf{r})/\epsilon_0$, where $\vec{\mathbf{E}}$ is the electric field generated by the charge distribution ρ and ϵ_0 is the dielectric constant. Under the assumption that the electric field is curl-free, which is the case in the absence or near-absence of a time-varying magnetic field, $\vec{\mathbf{E}}$ is a conservative vector field, that is there exists a potential $\mathbf{U}$ such that $\vec{\mathbf{E}} = -\nabla \mathbf{U}$, which therefore verifies the Poisson's equation $\Delta \mathbf{U}(\mathbf{r}) = -\rho(\mathbf{r})/\epsilon_0$. The Coulomb potential c_d is then defined as the radial distribution which solves the fundamental equation $\Delta \mathsf{c}_d(\mathbf{r}) = -|\mathbb{S}^{d-1}|\delta(\mathbf{r})$, where δ is the Dirac distribution. One can easily check that

$$\mathsf{c}_d(\mathbf{r}) = \begin{cases} -\log|\mathbf{r}|, & d = 2, \\ \frac{1}{|\mathbf{r}|^{d-2}}, & d \geqslant 3. \end{cases}$$

Remark 2. From a mathematical perspective, the Coulomb costs c_d in dimension $d \geqslant 3$ are particular cases of a broader class of costs, stemming from the so-called *Riezs potentials* $\mathsf{c}_{d,s}(\mathbf{r}) = |\mathbf{r}|^{-s}$ where $0 < s < d$. Many results for the Coulomb costs c_d remains valid for the *Riesz costs* $c_{d,s}$, evidently defined as

$$c_{d,s}(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{1 \leqslant k < l \leqslant N} \mathsf{c}_{d,s}(\mathbf{x}_k - \mathbf{x}_l).$$

Remark 3. Compared to higher dimensions, the Coulomb potential c_2 in two-dimension of space is not bounded from below, which, as we will see, makes the theoretical analysis a little bit harder. Therefore, we might expect somewhat eccentric behaviors for two-dimensional systems arising from the very long-range nature of this potential. Note that, in spite of its theoretical interest, the two-dimensional Coulomb potential does arise "naturally" in a three-dimensional setting. Indeed, consider an infinitely long and thin charged wire supported on $\mathcal{W} = \{(0, 0, z) \in \mathbb{R}^3 : z \in \mathbb{R}\}$ with constant linear charge density σ. Given a point $\mathbf{r} \in \mathbb{R}^3$ such that $\mathrm{d}(\mathbf{r}, \mathcal{W}) = r > 0$, by considering the cylinder of length l and radius r centered on $\mathcal{W}$ and applying Gauss's law, it holds that the magnitude of the electric field is given by $E = \frac{\sigma}{2\pi\varepsilon_0 r}$, so that the corresponding potential is proportional $-\log(r)$. We also note that two-dimensional (and one-dimensional) *log gases* also occur naturally in the theory of random matrices (see Remark 5) or for the *fractional quantum Hall effect*.

From a numerical perspective, solving (mK) is a challenging problem due to the high dimension of the space of transport plans $\Pi(\mu_1, \dots, \mu_N)$. Indeed, traditional numerical approximations are usually impractical, for the number of degrees of freedom grows exponentially with the number of marginals N when using straightforward numerical discretizations, so that only very small systems can be considered in practice. Nonetheless, because of the emerging importance of MOT in a wide variety of concrete applications, novel numerical methods have emerged in the literature [6, 67], and a rather promising advance was recently proposed in [25], which could potentially lead to circumvent the infamous curse of dimensionality. Note that the problem is even more challenging when one is to consider Coulomb-type costs (or for that matters, Riesz-type costs) because of their unusual features, that is both their repulsive nature and the fact that they carry singularities on their diagonals, which might give rise to numerical instabilities. To alleviate for those impracticalities, two natural strategies come into play.

Duality theory. A first strategy would be to look for a meaningful dual formulation to (mK). Recall that one can usually gain a lot of insightful knowledge about the structure of the original problem by studying the dual formulation, and this is especially veracious when the dual has a meaningful interpretation. Following the same formal strategy as in the classical OT case, a natural dual to (mK) reads

$$\sup_{(\varphi_1,\dots,\varphi_N)\in\Phi_c(\mu_1,\dots,\mu_N)} \left\{\mathbb{E}_{X_1\sim\mu_1}[\varphi_1(X_1)] + \cdots + \mathbb{E}_{X_N\sim\mu_N}[\varphi_N(X_N)]\right\},$$

where the φ_i's, so-called *Kantorovich potentials*, are such that $\varphi_1(x_1) + \cdots + \varphi_N(x_N) \leqslant c(x_1,\dots,x_N)$. Remark that this constraint is highly non-linear and non-compact, which makes the problem difficult even though the objective is linear. When all the marginals are equal to μ, we will write $\Phi_c(\mu) := \Phi_c(\mu,\dots,\mu)$, and in this case, under the assumption that the cost c is symmetric, we can also assume that all the potentials are the same, *i.e.* $\varphi_1 = \cdots = \varphi_N = \varphi$.

As usual, solving the dual yields a lower bound on the primal solution. An important question is whether or not *strong duality* holds, that is whether or not the converse statement is true. It is a rather old result established in a pleasant generality in [46] that the duality gap vanishes when c is either lower semi-continuous, bounded from below and bounded above by a direct sum of μ_i-integrable functions, or that c is simply uniformly bounded without demanding any continuity hypotheses. In either cases, this result does not apply to Coulomb-type costs because of their singularities. The hypotheses were later weakened in [5], where strong duality is shown to hold for any lower semi-continuous cost c bounded from below, by using a relaxed formulation to the primal problem, though their strategy does not allow to prove the existence of maximizers for the dual problem. In [27], this shortcoming was eventually corrected, providing a full duality theory for the Coulomb-type cost in dimension $d = 3$ together with bound estimates on the dual maximizers, and in [13] the result was extended to the cases $d \geqslant 3$ (and to Riesz-type costs $c_{d,s}$), together with regularity estimates on the dual maximizers. In two space-dimensions, where the Coulomb cost is not bounded from below, it was only very recently that a proof of strong duality was given [36], completing the overall picture.

Entropic regularization. A second strategy would be to work with a regularized version of the MOT problem. Indeed, the current successes of OT did not stem from the ability to solve (K) from a numerical point-of-view, which altogether remains a very complicated problem, but rather came from the rather astute idea of adding an entropic barrier to the objective in order to enforce the positivity (and uniqueness) of the transport plan. This idea was first considered from a multi-marginal point-of-view in [1] and [7], where the problem reads

$$\inf_{\pi\in\Pi(\mu_1,\dots,\mu_N)} \left\{\mathbb{E}_{(X_1,\dots,X_N)\sim\pi}\left[c(X_1,\dots,X_N)\right] + \beta^{-1}\mathrm{Ent}(\pi|\otimes_{i=1}^N \mu_i)\right\},$$

where $\beta > 0$ is a regularization parameter, which should formally map the regularized problem to the original one as $\beta \to \infty$, and where, given two measures α and β, the (negative) *entropy of* α *relative to measure* β (also called the *Kullback-Leibler divergence*) is defined as

$$\mathrm{Ent}(\alpha|\beta) = \begin{cases} \mathbb{E}_{Z\sim\alpha}\left[\log\left(\frac{\mathrm{d}\alpha}{\mathrm{d}\beta}(Z)\right)\right], & \text{if } \alpha \ll \beta, \\ +\infty, & \text{otherwise.} \end{cases}$$

While adding the entropy makes the objective strongly convex, the main interest of this method lies in the simplicity of its implementation. Indeed, in a discretized setting, it leads to the famous *iterative proportional fitting procedure* [28], also known as *Sinkhorn algorithm* [94] or *Sinkhorn-Knopp algorithm* [95], which consists in nothing but a simple alternate minimization scheme, which itself translates into iterations that are simple matrix-vector products. Altogether, the algorithm can easily be made parallel and fully-vectorized, making it particularly suitable for execution on GPU's [26].

A possible dual to the regularized problem reads

$$\sup_{\varphi_1,\dots,\varphi_N\,:\,\varphi_i\in C_b(\mathscr{X}_i)} \left\{ \sum_{i=1}^N \mathbb{E}_{X_i\sim\mu_i}[\varphi_i(X_i)] \right.$$
$$\left. - \beta^{-1}\mathbb{E}_{(X_1,\dots,X_N)\sim\mu_1\otimes\cdots\otimes\mu_N}\left[e^{-\beta\left(c(X_1,\dots,X_N)-\sum_{i=1}^N\varphi_i(X_i)-1\right)}\right]\right\},$$

and using the primal-dual relationship, one can then formally show that there exist a set of functions $\omega_1,\dots,\omega_N$, so-called *entropic weights*, such that the unique minimizer to the primal regularized problem is of the form

$$e^{-\beta c(x_1,\dots,x_N)}\left[\prod_{i=1}^N \omega_i(x_i)\right]\mathrm{d}\mu_1(x_1)\cdots\mathrm{d}\mu_N(x_N),$$

where the ω_i's are related to the dual variables φ_i's through $\omega_i(x) = e^{\beta\varphi_i(x)}$. Once again, remark that when the cost function is symmetric and all the marginals are equal, we can take $\varphi_1 = \cdots = \varphi_N = \varphi$ and subsequently all the entropic weights are the same as well.

Coming back to the problem of computing the Coulomb energy of a system at fixed density, adding the entropic regularization is equivalent to looking at the system at *positive temperature*. Indeed, let us consider a system of charged particles $\mathbf{x}_1,\dots,\mathbf{x}_N$ in thermal equilibrium with a heat reservoir at fixed temperature β^{-1}, and suppose that the particles evolve in a potential landscape V, so that the (classical) *Hamiltonian* of the system reads

$$\mathcal{H}_V(\mathbf{x}_1,\dots,\mathbf{x}_N) = \sum_{1\leqslant k<l\leqslant N} \mathtt{c}_d(\mathbf{x}_k-\mathbf{x}_l) + \sum_{i=1}^N V(\mathbf{x}_i),$$

where we have dropped the kinetic energy term because the temperature is kept fixed by the heat reservoir. According to the general postulate of statistical mechanics [38, 83], this system is represented from a microscopic point-of-view by the so-called *canonical ensemble*, that is the Gibbs measure

$$Z_\beta(V)^{-1}e^{-\beta\mathcal{H}_V(\mathbf{x}_1,\dots,\mathbf{x}_N)}\mathrm{d}\mathbf{x}_1\dots\mathrm{d}\mathbf{x}_N,$$

where $Z_\beta(V) := \int_{\mathbb{R}^{dN}} e^{-\beta\mathcal{H}_V(\mathbf{x}_1,\dots,\mathbf{x}_N)}\mathrm{d}\mathbf{x}_1\dots\mathrm{d}\mathbf{x}_N$ is the so-called *partition function*. Now, because of the fully symmetric nature of our problem, all the entropic weights are the same and equal to ω, which we can rewrite as above as $\omega(\mathbf{x}) = e^{-\beta V(\mathbf{x})}$, where V is the dual variable, which has (almost; see Remark 4 below) the physical interpretation of a potential landscape. Now, passing over in silence some technicalities that we will thoroughly explain later on, the regularized problem rephrases as :

« *Find the potential V such that the canonical ensemble has density ρ at temperature β^{-1}* », that is find V such that

$$Z_\beta(V)^{-1} \int_{\mathbb{R}^{dN}} e^{-\beta \mathcal{H}_V(\mathbf{x}_1,\dots,\mathbf{x}_N)} \mathrm{d}\rho(\mathbf{x}_1) \dots \widehat{\mathrm{d}\rho(\mathbf{x}_i)} \dots \mathrm{d}\rho(\mathbf{x}_N) = N^{-1}\rho(\mathbf{x}_i)$$

for every $i = 1, \dots, N$.

Remark 4. In the context of the aforementioned dual, the dual variable V fails to qualify as a true potential. Indeed, adding a constant to a given V will lead to a different value for the dual objective, while from a physical perspective the potential is only meaningful up to an additive constant. We will explain later on how to modify the dual to make it invariant to constant shift in the dual variable.

Solving these equations analytically is virtually impossible, so that we set ourselves on the path of solving them numerically. Once again, let us recall that the problem is difficult because of *(i)* the high dimension of the integration space, which is exponential in the number of particles, and *(ii)* the fact that the dual variable V is highly non-local, in the sense that it lives on the entire Euclidian space. We will introduce a numerical scheme to circumvent this second issue, based on the fact that, from a physical point-of-view, one might rewrite $V = \mathtt{c}_d * \rho_{\mathrm{ext}}$ where ρ_{ext} is an exterior charge distribution generating the potential V. Indeed, the variable ρ_{ext} is typically local in space (*i.e.* compactly supported), making it more suitable as dual variable from a numerical perspective.

Remark 5. In two dimension of space, the (unnormalized) Gibbs measure at temperature β^{-1} reads

$$\prod_{1 \leqslant k < l \leqslant N} |\mathbf{x}_k - \mathbf{x}_l|^\beta \prod_{i=1}^{N} e^{\beta V(\mathbf{x}_i)} \mathrm{d}\mathbf{x}_1 \dots \mathrm{d}\mathbf{x}_N .$$

When $V(\mathbf{x}) = \frac{|\mathbf{x}|^2}{4}$, this is the so-called *β-Ginibre ensemble*, famously known in the field of random matrix theory [34]. When $\beta = 1$ and $d = 1$, we have the *Gaussian Orthogonal Ensemble* (GOE) which is invariant under orthogonal conjugation, and which models Hamiltonians with time-reversal symmetry. When $\beta = 2$ (and $d = 1$) we end up with the *Gaussian Unitary Ensemble* (GUE), which is invariant under unitary conjugation, and which models Hamiltonians lacking time-reversal symmetry. Finally, when $\beta = 4$ (and $d = 1$), we have the *Gaussian Symplectic Ensemble* (GSE) which is invariant under conjugation by the symplectic group, and which models Hamiltonians with time-reversal symmetry but no rotational symmetry. Bits of theoretical analysis are available in the literature concerning the general case where β is even, though there are no simple ways to generate the corresponding ensemble, contrary to the preceding cases. In [97] and [96], the partition function is computed exactly by recognizing that the Vandermonde determinant in front of Gibbs measure is the square of a Jack polynomial and by expanding it in an appropriate monomial base, while in [85], this is done by mapping the system onto a one-dimensional chain of interacting fermions. In [40] a connection between the two methods is derived. This could potentially lead to an interesting basis of work in our case, though we do not follow this strategy.

Acknowledgements. This master's thesis was advised and funded by MATHIEU LEWIN (*CNRS & CEREMADE, Université Paris-Dauphine*), to whom I express

my sincere gratitude for allowing me to pursue a PhD under his supervision. I would also like to thank LUCA NENNA (*Université Paris-Saclay*) and VIRGINIE EHRLACHER (*CERMICS*, *Ecole des Ponts*) for fruitful discussions.

2. MOT WITH LOGARITHMIC COST AT ZERO TEMPERATURE ($\beta = \infty$)

In this section, we study the problem of computing the Coulomb energy of a system of N two-dimensional charged particles $\mathbf{x}_1, \dots, \mathbf{x}_N$ with prescribed electronic density ρ (*i.e.* recall that $\rho(\mathbf{r}) \geqslant 0$ for every $\mathbf{r}$ and $\int_{\mathbb{R}^2} \rho(\mathbf{r})\mathrm{d}\mathbf{r} = N$) at zero temperature, that is we discard entropic regularization for the moment (*i.e.* $\beta = \infty$). The problem reads

$$C(\rho) = \inf_{\mathbb{P} \in \Pi^N(\rho)} \left\{ \mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}} \left[- \sum_{1 \leqslant k < l \leqslant N} \log|\mathbf{x}_k - \mathbf{x}_l| \right] \right\}, \tag{Coul2D}$$

where we use the shorthand notation $\Pi^N(\rho) := \Pi(\frac{\rho}{N}, \dots, \frac{\rho}{N})$. Under *ad hoc* assumptions on ρ, we prove that the problem admits a minimizer (Theorem 1), and that its associated cost is finite (Theorem 3). We then show that the particles must be bounded away from each others by a positive distance on every compact set (Theorem 5), and we eventually prove that a strong duality result holds (Theorem 8).

Remark 6. Throughout this entire document, we will always suppose that ρ is absolutely continuous with respect to the Lebesgue measure $\mathrm{d}\mathbf{r}$. This actually need not be true in classical setting, where ρ could be a positive measure with singular parts. Nevertheless, our results can easily be extended to this more general case.

2.1. Finiteness of $C(\rho)$ and diagonal bounds.

Throughout what follows, we make the assumptions that

$$\int_{\mathbb{R}^2} \log(2 + |\mathbf{r}|)\rho(\mathbf{r})\mathrm{d}\mathbf{r} < \infty, \tag{H1}$$

$$\rho \in L^q(\mathbb{R}^2) \text{ for some } q > 1. \tag{H2}$$

Remark 7. A stronger statement for hypothesis (H1) would be to ask that ρ has a finite second-moment, that is $\int_{\mathbb{R}^2} |\mathbf{r}|^2\rho(\mathbf{r})\mathrm{d}\mathbf{r} < \infty$, which physically translates into the fact that the electronic density ρ has a finite *quadrupole*.

Remark 8. From a quantum perspective, we only care for those densities which are *admissible*, in the sense that there exists some quantum state $\mathbb{P}$ such that $\rho = \rho_{\mathbb{P}}$. Under the rather natural hypothesis that

$$\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_N) = |\Psi(\mathbf{x}_1, \dots, \mathbf{x}_N)|^2 \mathrm{d}\mathbf{x}_1 \dots \mathrm{d}\mathbf{x}_N$$

with $\Psi \in H^1(\mathbb{R}^{2N})$, that is that the system has finite *kinetic energy*, then it is a well-known fact that $\sqrt{\rho}$ is also in $H^1(\mathbb{R}^2)$ [58, Theorem 1.1]. Invoking Sobolev injection, it holds that $\rho \in L^q(\mathbb{R}^2)$ for every $q \in [1, +\infty)$, so that ρ verifies the hypothesis (H2).

Theorem 1. *The infimum at* (Coul2D) *is attained.*

Proof of Theorem 1. By standard arguments, we have that set $\Pi^N(\rho)$ is compact for the weak-$\star$ topology. We now want to prove that the objective $\mathbb{P} \mapsto \mathfrak{C}(\mathbb{P})$ is lower semi-continuous with respect to $\mathbb{P} \in \Pi^N(\rho)$, which will yield the existence of a minimizer to (Coul2D). Because the Coulomb cost $\mathtt{c}_2$ is not bounded from below, we need to use the following lemma:

Lemma 2. *Let $c : Z \to \mathbb{R} \cup \{+\infty\}$ be a lower semi-continuous function, where Z is an arbitrary metric space, and let $\mathscr{M} \subset \mathcal{M}_+(Z)$ be a subset of positive measures on Z. If there exists an upper semi-continuous function $h : Z \to \mathbb{R} \cup \{-\infty\}$ such that* (i) *$h \in L^1(\mathbb{P})$ for all $\mathbb{P} \in \mathscr{M}$,* (ii) *with $\int_Z h\mathrm{d}\mathbb{P} = \int_Z h\mathrm{d}\mathbb{Q}$ for all $\mathbb{P}, \mathbb{Q} \in \mathscr{M}$, and such that* (iii) *$h \leqslant c$, then the cost functional $\mathbb{P} \mapsto J_c(\mathbb{P}) := \mathbb{E}_{z\sim\mathbb{P}}[c(z)]$ is lower semi-continuous over $\mathscr{M}$ for the weak-$\star$ topology.*

Proof of Lemma 2. Recall that, if c is lower semi-continuous and bounded from below, one can find a sequence $(c_k : Z \to \mathbb{R})_k$ of bounded above and uniformly bounded from below continuous functions, such that for every $z \in Z$, the sequence $(c_k(z))_k$ increasingly converge to $c(z)$. By monotonic convergence, we have that $\sup_k J_{c_k}(\mathbb{P}) = J_c(\mathbb{P})$ for every $\mathbb{P} \in \mathcal{M}_+(Z)$. Since the functional J_{c_k} is continuous over $\mathcal{M}_+(Z)$ for the weak-$\star$ topology, it holds that J_c is lower semi-continuous as pointwise supremum of continuous functionals.

Now, in the general case, by hypothesis *(iii)*, we have that $c - h$ is lower semi-continuous and non-negative, so that we can apply what precedes. By hypothesis *(i)*, we can legally write $J_c(\mathbb{P}) = J_{c-h}(\mathbb{P}) + J_h(\mathbb{P})$ for every $\mathbb{P} \in \mathscr{M}$. The functional J_{c-h} is lower semi-continuous, and by hypothesis *(ii)*, the functional J_h is constant over $\mathscr{M}$. Therefore, J_c is lower semi-continuous over $\mathscr{M}$. □

Let us now prove the objective is indeed lower semi-continuous with respect to $\mathbb{P} \in \Pi^N(\rho)$ by using the above lemma. Denoting $\ell(t) := \min\{0, -\log(t)\}$, we define $\widetilde{h}(\mathbf{x}_1, \dots, \mathbf{x}_N) := \sum_{i=1}^N \ell(|\mathbf{x}_i|)$. Remark that the function $\widetilde{h}$ is continuous and thus trivially upper semi-continuous. Furthermore, given any $\mathbb{P} \in \Pi^N(\rho)$, we have that

$$\begin{aligned}\int_{\mathbb{R}^{2N}} |\widetilde{h}|(\mathbf{x}_1, \dots, \mathbf{x}_N)\mathrm{d}\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_N) &\leqslant \int_{\mathbb{R}^2} |\ell(|\mathbf{r}|)|\rho(\mathbf{r})\mathrm{d}\mathbf{r} \\ &= \int_{\mathbb{R}^2\setminus B(0,1)} \log(|\mathbf{r}|)\rho(\mathbf{r})\mathrm{d}\mathbf{r} < \infty,\end{aligned}$$

by hypothesis (H1), and that

$$\int_{\mathbb{R}^{2N}} \widetilde{h}(\mathbf{x}_1, \dots, \mathbf{x}_N)\mathrm{d}\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_N) = N \int_{\mathbb{R}^2} \ell(|\mathbf{r}|)\rho(|\mathbf{r}|)\mathrm{d}\mathbf{r},$$

so that the value of the integral is independent of the choice $\mathbb{P} \in \Pi^N(\rho)$. Finally, for every $(\mathbf{x}_1, \dots, \mathbf{x}_N) \in \mathbb{R}^{2N}$, we have

$$\begin{aligned} c\left(\frac{\mathbf{x}_1}{2}, \dots, \frac{\mathbf{x}_N}{2}\right) &\geqslant \sum_{1 \leqslant k < l \leqslant N} \ell\left(\frac{|\mathbf{x}_k - \mathbf{x}_l|}{2}\right) \\ &\geqslant \sum_{1 \leqslant k < l \leqslant N} \ell\left(\frac{|\mathbf{x}_k| + |\mathbf{x}_l|}{2}\right) \\ &\geqslant \sum_{1 \leqslant k < l \leqslant N} \min\{\ell(|\mathbf{x}_k|), \ell(|\mathbf{x}_l|)\} \\ &\geqslant \sum_{1 \leqslant k < l \leqslant N} \ell(|\mathbf{x}_k|) + \ell(|\mathbf{x}_l|) \\ &= (N-1)\widetilde{h}(\mathbf{x}_1, \dots, \mathbf{x}_N), \end{aligned}$$

Therefore, by defining $h(\mathbf{x}_1, \dots, \mathbf{x}_N) := (N-1)\widetilde{h}(2\mathbf{x}_1, \dots, 2\mathbf{x}_N)$, we can apply Lemma 2 and conclude that the objective is lower semi-continuous, thus yielding the existence of an optimal solution to (Coul2D). □

Remark 9. A rather important by-product of Lemma 2 is that the integral at (Coul2D) is actually *well-defined*. Indeed, because $\mathfrak{c}_2$ has no definite sign, one ought to be careful about the legitimacy of the cost integral $\mathfrak{C}(\mathbb{P})$ for a given $\mathbb{P} \in \Pi^N(\rho)$.

Let us give an alternative proof of Theorem 1, which, truth be told, in nothing but a rephrasing of the preceding arguments, but in such a manner that it sheds light on the peculiarity of the logarithmic cost. In fact, another way to look at the misfortune of the Coulomb potential in two-dimension of space is that, compared to higher dimensions, its Fourier transform is not a positive distribution over the all space (*i.e.* $|\mathbf{r}|^{-2}$ is not locally integrable at its singularity), but only on $\mathbb{R}^2 \setminus \{0\}$.

Alternative proof of Theorem 1. We want to prove that the functional C is weakly lower semi-continuous over $\Pi^N(\rho)$. To a certain extent, we will use the old strategy proposed by Onsager [70], which consists in replacing the punctual particles $\mathbf{x}_1, \dots, \mathbf{x}_N$ by disks D_r of radius $r > 0$ and subtracting a positively charged background of mass $N|D_r|$, in such a way that we can wisely use the positivity of the distribution $\widehat{-\log|\cdot|}$ away from the origin. We will then recover the original problem by letting $r \to \infty$.

Consider $\chi_r := |D_r|^{-1}\mathbb{1}_{D_r}$ where $D_r \subset \mathbb{R}^2$ is a small disk of radius $r > 0$. We have $-\log|\cdot| * \chi_r(\mathbf{r}) \leqslant -\log|\mathbf{r}|$ for every $\mathbf{r}$: indeed, we have equality by harmonicity for $|\mathbf{r}| \leqslant r$, and when $|\mathbf{r}| > r$, we have

$$-\log|\cdot| * \chi_r(\mathbf{r}) = -\log(r) + \frac{1}{2}\left(1 - \frac{|\mathbf{r}|^2}{r^2}\right) \leqslant -\log|\mathbf{r}|,$$

where the inequality follows from the fact that $1 - \frac{1}{t} \leqslant \log(t)$ for all $t > 0$. Taking the convolution with χ_r once again, it holds that $-\log|\cdot| * \chi_r * \chi_r \leqslant -\log|\cdot|$. Note that it also follows that $-\log|\cdot| * \chi_r * \chi_r \leqslant -\log|\cdot| * \chi_s * \chi_s$ when $r \leqslant s$, which we will allow us to recover the true cost functional by monotonic convergence. We then define the regularized cost c_r as

$$c_r(\mathbf{x}_1, \dots, \mathbf{x}_N) := \sum_{1 \leqslant k < l \leqslant N} -\log|\cdot| * \chi_r * \chi_r(\mathbf{x}_k - \mathbf{x}_l),$$

which verifies $c_r \leqslant c$ according to what precedes, and we write $\mathfrak{C}_r$ the associated cost functional. For two functions f and g, physically representing two charge distributions, we denote by $D(f,g)$ their *Coulomb* (or *electrostatic*) *interaction energy*, that is

$$D(f,g) := -\iint_{\mathbb{R}^2\times\mathbb{R}^2} \log|\mathbf{x}-\mathbf{y}|\, f(\mathbf{x})g(\mathbf{y})\mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}.$$

When f verifies hypothesis (H1) and $\int_{\mathbb{R}^2} f(\mathbf{x})\mathrm{d}\mathbf{x} = 0$, we can in all legitimacy use Plancherel theorem to write

$$D(f,f) = c_f \int_{\mathbb{R}^2} \frac{|\widehat{f}(\mathbf{k})|^2}{|\mathbf{k}|^2}\mathrm{d}\mathbf{k} \geqslant 0,$$

where $c_f > 0$ is a constant depending on the definition used for the Fourier transform. Now, defining $\rho_{\mathrm{ext}} := N\chi_r$ the charged background which we astutely subtract, we write

$$\begin{aligned} c_r(\mathbf{x}_1,\dots,\mathbf{x}_N) &= -\frac{1}{2}\sum_{1\leqslant k,l\leqslant N} -\log|\cdot| * \chi_r * \chi_r(\mathbf{x}_k-\mathbf{x}_l) - \frac{ND(\chi_r,\chi_r)}{2} \\ &\quad + \sum_{i=1}^N \log|\cdot| * \chi_r * \rho_{\mathrm{ext}}(\mathbf{x}_i) - \sum_{i=1}^N \log|\cdot| * \chi_r * \rho_{\mathrm{ext}}(\mathbf{x}_i) \\ &= c_r'(\mathbf{x}_1,\dots,\mathbf{x}_N) - \sum_{i=1}^N \log|\cdot| * \chi_r * \rho_{\mathrm{ext}}(\mathbf{x}_i) \\ &\qquad - \frac{ND(\chi_r,\chi_r)}{2} - \frac{D(\rho_{\mathrm{ext}}, D(\rho_{\mathrm{ext}})}{2}, \end{aligned}$$

where we have defined the cost c_r' as

$$\begin{aligned} c_r'(\mathbf{x}_1,\dots,\mathbf{x}_N) &= \frac{1}{2}D\left(\sum_{i=1}^N \chi_{r,\mathbf{x}_i} - \rho_{\mathrm{ext}}, \sum_{i=1}^N \chi_{r,\mathbf{x}_i} - \rho_{\mathrm{ext}}\right) \\ &= c_f \int_{\mathbb{R}^2} \frac{\left|\widehat{\chi_r}(\mathbf{k})\sum_{j=1}^N e^{-i\mathbf{k}\cdot\mathbf{x}_j} - \widehat{\rho_{\mathrm{ext}}}(\mathbf{k})\right|^2}{|\mathbf{k}|^2}\mathrm{d}\mathbf{k} \geqslant 0 \end{aligned}$$

with $\chi_{r,\mathbf{x}_i}(\mathbf{x}) = \chi_r(\mathbf{x}-\mathbf{x}_i)$, and where everything is well-defined because of the astute choice of ρ_{ext}. Therefore, we have

$$\mathfrak{C}_r(\mathbb{P}) = \mathfrak{C}_r'(\mathbb{P}) + \int_{\mathbb{R}^2} V(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r} - \frac{ND(\chi_r,\chi_r)}{2} - \frac{D(\rho_{\mathrm{ext}},\rho_{\mathrm{ext}})}{2},$$

where $V = -\log|\cdot| * (\rho_{\mathrm{ext}} * \chi_r)$. Note that V has at most a logarithmic growth-rate at infinity, so that its integral against ρ is well-defined. Finally, $\mathfrak{C}_r'$ is weakly lower semi-continuous since c_r' is positive and continuous, and therefore $\mathfrak{C}_r$ is also weakly lower semi-continuous. By monotone convergence, we have $\sup_{r>0}\mathfrak{C}_r(\mathbb{P}) = \mathfrak{C}(\mathbb{P})$ for every $\mathbb{P}\in\Pi^N(\rho)$, so that C is lower semi-continuous over $\Pi^N(\rho)$. □

Theorem 3. *The minimum* (Coul2D) *is finite,* i.e. $C(\rho) \in (-\infty, +\infty)$.

Proof of Theorem 3. By Fubini theorem, we have

$$-\int_{|\mathbf{x}-\mathbf{y}|\leqslant 1}\log|\mathbf{x}-\mathbf{y}|\rho(\mathbf{x})\rho(\mathbf{y})=\int_{\mathbf{x}\in\mathbb{R}^2}\left(\int_{\mathbf{y}\in B(0,1)}-\log|\mathbf{y}|\rho(\mathbf{x}+\mathbf{y})\right)\rho(\mathbf{x}),$$

so that by hypothesis (H2), we can use Hölder's inequality both for the inner and outer integrals, yielding

$$-\int_{|\mathbf{x}-\mathbf{y}|\leqslant 1}\log|\mathbf{x}-\mathbf{y}|\rho(\mathbf{x})\rho(\mathbf{y})\leqslant\|\rho\|_{L^1(\mathbb{R}^2)}\|\rho\|_{L^q(\mathbb{R}^2)}\|\log|\cdot|\|_{L^{q'}(B(0,1))}$$

where q' is the conjugate exponent of q. Therefore, one has

$$\mathfrak{C}((\rho/N)^{\otimes N})\leqslant\frac{N(N-1)}{2}\|\rho\|_{L^1(\mathbb{R}^2)}\|\rho\|_{L^q(\mathbb{R}^2)}\|\log|\cdot|\|_{L^{q'}(B(0,1))},$$

yielding that $C(\rho)<\infty$. Now, given any $\mathbb{P}\in\Pi^N(\rho)$, we have

$$-\int_{\mathbb{R}^{2N}}\log|\mathbf{x}_k-\mathbf{x}_l|\mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N)\geqslant-\frac{2}{N}\int_{\mathbb{R}^2}|\log(2|\mathbf{r}|)|\rho(\mathbf{r})\mathrm{d}\mathbf{r},$$

so that, by using hypothesis (H2) to control the integral at the origin and hypothesis (H1) to control the integral at infinity, we obtain a uniform lower bound to $\mathfrak{C}(\mathbb{P})$ for every $\mathbb{P}\in\Pi^N(\rho)$, thus proving that $C(\rho)>-\infty$. □

Given an optimal transport plan $\mathbb{P}\in\Pi^N(\rho)$, we now turn our attention towards the structure of its support $\mathrm{supp}(\mathbb{P})$. Because of the repulsive nature of the Coulomb cost, we intuitively expect that the particles $\mathbf{x}_1,\dots,\mathbf{x}_N$ cannot get infinitely close to each others, that is, mathematically, that there must exist some $\alpha:=\alpha(\rho)>0$ such that $\mathrm{supp}(\mathbb{P})\subset\mathbb{R}^{2N}\setminus\Delta_\alpha$, where for any $r>0$, we have defined the *r-diagonal* Δ_r as

$$\Delta_r:=\left\{(\mathbf{x}_1,\dots,\mathbf{x}_N)\in\mathbb{R}^{2N}:\exists k\neq l\text{ s.t. }|\mathbf{x}_k-\mathbf{x}_l|\leqslant r\right\}.$$

Remark 10. For the Coulomb cost in dimension $d=3$, this statement was first proved in a slightly weaker form in [27], and later strengthened and generalized to a broader class of costs, among which the Riesz-type costs, in [13]. It is claimed without proof in [37] that this result also applies for the logarithmic cost. In what follows, we provide the details of this argument on every compact set. This is altogether a sufficient statement when considering compactly supported densities (see Remark 11).

For any $\alpha>0$, because $\rho\in L^1(\mathbb{R}^2)$, one can find a threshold $\beta_\rho(\alpha)>0$ such that $\sup_{\mathbf{x}\in\mathbb{R}^2}\int_{B(\mathbf{x},\beta_\rho(\alpha))}\rho(\mathbf{r})\mathrm{d}\mathbf{r}<\alpha$. We will need the following simple lemma:

Lemma 4 ([13])**.** *Let $\beta:=\beta_\rho\left((N-1)^{-2}\right)$. Given $\mathbb{P}\in\Pi^N(\rho)$ and $\mathbf{x}_1=(x_1^1,\dots,x_1^N)\in\mathbb{R}^{2N}$, there exist $\mathbf{x}_2,\dots,\mathbf{x}_N\in\mathbb{R}^{2N}$ such that, whenever $i\neq j$ and $k\neq l$, we have*

$$\beta<|x_k^i-x_l^j|,\quad\forall k,l,i,j=1,\dots,N.$$

Proof of Lemma 4. Let us denote by $\mathbb{R}^{2(N-1)}\otimes_\alpha A$ the Cartesian product on N factors the α-th of which is A while all the others are copies of $\mathbb{R}^2$. We have that

$$\mathbb{P}\left(\mathbb{R}^{2(N-1)}\otimes_\alpha B(x_k^i,\beta)\right)<\frac{1}{N(N-1)^2},$$

for all α, k, i. Then, for any $j = 2, \dots, N$, we trivially have that

$$\mathbb{P}\left(\bigcup_{k=1}^{j-1}\bigcup_{i=1}^{N}\bigcup_{\alpha\neq i}\mathbb{R}^{2(N-1)}\otimes_\alpha B(x_k^i,\beta)\right) < \frac{j-1}{N-1}.$$

Since $\mathbb{P}$ is a probability measure, it holds that there exists $\mathbf{x}_j$ such that

$$\mathbf{x}_j \in \mathrm{supp}(\mathbb{P}) \setminus \left(\bigcup_{k=1}^{j-1}\bigcup_{i=1}^{N}\bigcup_{\alpha\neq i}\mathbb{R}^{2(N-1)}\otimes_\alpha B(x_k^i,\beta)\right),$$

and by construction the $\mathbf{x}_j$'s verify the desired property. □

Theorem 5. *Given any compact set $K \subset \mathbb{R}^{2N}$, there exists $\alpha := \alpha(\rho, K) > 0$ such that, for every minimizer $\mathbb{P}$ to (Coul2D), one has* $\mathrm{supp}(\mathbb{P}) \cap K \subset \mathbb{R}^{2N} \setminus \Delta_\alpha$.

Proof of Theorem 5. Let $\mathbb{P}$ be an optimal solution to (Coul2D). Without loss of generality, we place ourselves on the closed ball $B_R := \overline{B(0,R)} \subset \mathbb{R}^{2N}$ with $R > 0$, and denote $\mathrm{supp}_R(\mathbb{P}) = \mathrm{supp}(\mathbb{P}) \cap B_R$. Suppose there exists some $\mathbf{x}_1 \in \Delta_0 \cap \mathrm{supp}_R(\mathbb{P})$, and choose $\mathbf{x}_2, \dots, \mathbf{x}_N \in \mathrm{supp}(P)$ as in the above lemma. In what follows, given $\mathbf{x} = (x^1, \dots, x^N) \in \mathbb{R}^{2N}$ and $r > 0$, we denote $Q(\mathbf{x}, r)$ the Cartesian product $B(x^1, r) \times \cdots \times B(x^N, r)$. Let us then define $\mathbb{P}_k = \mathbb{P}_{|Q(\mathbf{x}_k, r)}$ for every $k = 1, \dots, N$ where we fix some radius $r > 0$. Because the points $\mathbf{x}_1, \dots, \mathbf{x}_N$ belong to the support of $\mathbb{P}$, it is possible to find $\lambda_1, \dots, \lambda_N \in (0, 1]$ such that

$$\lambda_1|\mathbb{P}_1| = \cdots = \lambda_N|\mathbb{P}_N|.$$

We then decompose $\mathbb{P}$ as $\mathbb{P} = \lambda_1\mathbb{P}_1 + \cdots + \lambda_N\mathbb{P}_N + \mathbb{P}_R$ where $\mathbb{P}_R$ is a remainder. The cost of $\mathbb{P}$ is bounded below by

$$\mathfrak{C}(\mathbb{P}) \geqslant \mathfrak{C}(\mathbb{P}_R) - \sum_{i=1}^{N}\sum_{k=1}^{N}\sum_{k<l\leqslant N}\log\Big(|x_i^k - x_i^l| + 2r\Big)\lambda_i|\mathbb{P}_i|. \tag{1}$$

For every $k = 1, \dots, N$, let us denote $\nu_k^1, \dots, \nu_k^N$ the marginals of $\lambda_k\mathbb{P}_k$, and by defining $\widetilde{\mathbb{P}}_k = \nu_1^{k\,(\mathrm{modN})} \times \cdots \times \nu_N^{k+N-1\,(\mathrm{modN})}$ for every $k = 1, \dots, N$. Let us then define $\widetilde{\mathbb{P}} := \widetilde{\mathbb{P}}_1 + \cdots + \widetilde{\mathbb{P}}_N + \mathbb{P}_R$. Given any Borel set $A \subset \mathbb{R}^2$, and for every $\alpha = 1, \dots, N$, we have

$$\widetilde{\mathbb{P}}\left(\mathbb{R}^{2(N-1)}\otimes_\alpha A\right) = \sum_{k=1}^{N}\nu_\alpha^k(A) + \mathbb{P}_R\left(\mathbb{R}^{2(N-1)}\otimes_\alpha A\right),$$

and because $\nu_\alpha^k(A) = \lambda_k\mathbb{P}_k(\mathbb{R}^{2(N-1)}\otimes_\alpha A)$, it follows that $\widetilde{\mathbb{P}}$ has the same marginals as $\mathbb{P}$, that is $\widetilde{\mathbb{P}} \in \Pi^N(\rho)$. Moreover $|\widetilde{\mathbb{P}_k}| = \lambda_k|\mathbb{P}_k|$ for every $i = 1, \dots, N$. We therefore have the upper bound

$$\mathfrak{C}(\widetilde{\mathbb{P}}) \leqslant \mathfrak{C}(\mathbb{P}_R) - \sum_{i=1}^{N}\sum_{k=1}^{N}\sum_{k<l\leqslant N}\log\Big(|x^k_{k+i-1\,(\mathrm{modN})} - x^l_{l+i-1\,(\mathrm{modN})}| - 2r\Big)\lambda_i|\mathbb{P}_i|. \tag{2}$$

We want to select $r > 0$ such that $\mathfrak{C}(\widetilde{\mathbb{P}}) < \mathfrak{C}(\mathbb{P})$. But from the fact that $\mathbf{x}_1 \in \Delta_0$, it holds that the RHS of (1) diverges to $-\infty$ as $r \to 0$, when the RHS of (2) remains bounded, so that such a r exists, contradicting the minimality of $\mathbb{P}$. Therefore Δ_0

does not intersect $\mathrm{supp}_R(\mathbb{P})$, which means that, by compactness, the two sets are separated by a positive distance. □

Remark 11. Suppose that ρ is compactly supported, and let $\delta := \mathrm{diam}(\mathrm{supp}(\rho)) < \infty$. Looking closely at the above proof, it is easy to see that for any optimal solution $\mathbb{P}$ and any $r < \alpha(\rho)$, we have $\Delta_r \cap \mathrm{supp}(\mathbb{P}) = \emptyset$, where

$$\alpha(\rho) := \log^{-1}\left[\frac{N^2(N-1)}{2}(\log(\beta) - \log(\delta)) + \log(\delta)\right].$$

Remark 12. It later came to our attention that in [21], it is proved for a wide variety of repulsive costs, among which the logarithmic cost, the stronger statement that the support of any optimal transport plan is supported away from the diagonals, without supposing, as we did in the preceding remark, the compactness of the support of the targeted marginals.

2.2. **Duality theory.** We now prove that a strong duality result holds for our problem, that is $C(\rho) = D(\rho)$, with

$$D(\rho) = \sup_{\varphi \in \Phi_{c_2}(\rho)} \int_{\mathbb{R}^2} \varphi(\mathbf{r}) \mathrm{d}\rho(\mathbf{r}),$$

Recall that the misfortune of the logarithmic cost arises from both the singularities that it carries on its diagonals and the fact that it is not bounded from below. Indeed, if we could disregard the singularities, strong duality would immediately follows from the very general results presented in [46], and because of the non-boundedness at infinity, the result established at Theorem 5, which allows us to work away from the diagonals, is unfortunately not strong enough since it only works when we restrict our attention on compact sets.

In order to prove that strong duality holds, we will work with the truncated potential $\mathtt{c}_{2,\alpha}(\mathbf{r}) := \min\{\mathtt{c}_2(\mathbf{r}), -\log(\alpha)\}$ and the corresponding cost $c_{2,\alpha}$, that is

$$c_{2,\alpha}(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{1 \leqslant k < l \leqslant N} \mathtt{c}_{2,d}(\mathbf{x}_k - \mathbf{x}_l),$$

and we will use some limiting arguments as $\alpha \to 0$. We will write $\mathfrak{C}_\alpha(\mathbb{P})$ and $C_\alpha(\rho)$ the (self-explanatory) associated functionals. With this cost, the results in [46] apply, so that

$$C_\alpha(\rho) = D_\alpha(\rho) = \sup_{\varphi \in \Phi_{c_{2,\alpha}}(\rho)} \int_{\mathbb{R}^2} \varphi(\mathbf{r}) \mathrm{d}\rho(\mathbf{r})$$

In what follows, we will work on the set $\rho \in \mathscr{R}$, where $\mathscr{R} = L^1(\log(2 + |\mathbf{r}|)\mathrm{d}\mathbf{r})$, that is exactly the set of ρ verifying hypothesis (H1). The weak topology on $\mathscr{R}$ is defined as follows : we say that a sequence $\{\rho_n\}_n$ in $\mathscr{R}$ converge weakly to $\rho \in \mathscr{R}$, which we write $\rho_n \rightharpoonup \rho$, if

$$\int_{\mathbb{R}^2} \rho_n(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r} \xrightarrow[n\to\infty]{} \int_{\mathbb{R}^2} \rho(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r}, \quad \forall \varphi \in \mathscr{R}',$$

where $\mathscr{R}'$ denotes the topological dual of $\mathscr{R}$ (*i.e.* those functions $\varphi \in L^\infty_{\mathrm{loc}}(\mathbb{R}^2)$ such that $\varphi(\mathbf{r})/\log(2 + |\mathbf{r}|) \in L^\infty(\mathbb{R}^2)$). Then, from the monotonicity of the integral, we see that we can restrict our attention to Kantorovich potentials for the problem

D_α which belong to the dual $\mathscr{R}'$, so that

$$D_\alpha(\rho) =: \sup_{\varphi\in\mathscr{R}'}\left\{\int_{\mathbb{R}^2}\varphi(\mathbf{r})\mathrm{d}\rho(\mathbf{r}) : \varphi(\mathbf{x}_1)+\cdots+\varphi(\mathbf{x}_N)\leqslant c_\alpha(\mathbf{x}_1,\ldots,\mathbf{x}_N)\right\}.$$

As mentioned earlier, strong duality was proved for the logarithmic cost in [36]. We do not follow their proof, and instead we present our own take on the problem, which we believe to be simpler.

Lemma 6. *For every $\alpha > 0$, the functional C_α is weakly lower semi-continuous over $\mathscr{R}$.*

Proof of Lemma 6. By duality, it holds that $C_\alpha(\rho) = D_\alpha(\rho)$ for every $\rho \in \mathscr{R}$. Considering a sequence $\{\rho_n\}_n$ in $\mathscr{R}$ such that $\rho_n \rightharpoonup \rho \in \mathscr{R}$, we have

$$\begin{aligned}
\liminf_{n\to\infty} D_\alpha(\rho_n) &= \liminf_{n\to\infty} \sup_{\varphi\in\mathscr{R}':\varphi\leqslant c_\alpha}\int_{\mathbb{R}^2}\rho_n(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r}\\
&\geqslant \sup_{\varphi\in\mathscr{R}':\varphi\leqslant c_\alpha}\liminf_{n\to\infty}\int_{\mathbb{R}^2}\rho_n(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r}\\
&= \sup_{\varphi\in\mathscr{R}':\varphi\leqslant c_\alpha}\int_{\mathbb{R}^2}\rho(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r}\\
&= D_\alpha(\rho),
\end{aligned}$$

where we have used the self-explanatory shorthand notation $\varphi \leqslant c_\alpha$. □

Lemma 7. *For every $\rho \in \mathscr{R}$, we have* $\sup_{\alpha>0} C_\alpha(\rho) = C(\rho)$.

Proof of Lemma 7. For every $\alpha > 0$, let $\mathbb{P}_\alpha$ be a minimizer for $C_\alpha(\rho)$. The sequence $\{\mathbb{P}_\alpha\}_\alpha$ being tight, there exists a converging subsequence for the weak-$\star$ topology, which we do not relabel for simplicity, $\mathbb{P}_\alpha \rightharpoonup \mathbb{P} \in \Pi^N(\rho)$ as $\alpha \to 0$. Given $\beta > \alpha$, we have $C_\alpha(\rho) \geqslant \mathfrak{C}_\beta(\mathbb{P}_\alpha)$. Because $\mathfrak{C}_\beta$ is lower semi-continuous for the weak-$\star$ topology (by the argument as Theorem 1), it holds that

$$\liminf_{\alpha\to 0} C_\alpha(\rho) \geqslant \mathfrak{C}_\beta(\mathbb{P}).$$

Letting $\beta \to 0$, by monotonic convergence we have that $\mathfrak{C}_\beta(\mathbb{P}) \to \mathfrak{C}(\mathbb{P})$, so that $\liminf_{\alpha\to 0} C_\alpha(\rho) \geqslant \mathfrak{C}(\mathbb{P})$, yielding the desired result.

□

Theorem 8. *Strong duality holds,* i.e. *$C(\rho) = D(\rho)$, for every $\rho \in \mathscr{R}$.*

Proof of Theorem 8. Let us consider the functional $\rho \in \mathscr{R} \mapsto C(\rho)$. As a pointwise infimum of linear functionals, it holds that C is convex. Moreover, from Lemma 6 and Lemma 7, we have that C is weakly lower semi-continuous. Furthermore, since $\mathscr{R} \cap L^q \neq \emptyset$, it holds by Theorem 3 that C is not identically $+\infty$. Therefore C is a so-called *regular convex function*, and by Fenchel's theorem [92, Theorem 5.17], it holds that

$$C(\rho) = (C^*)^*(\rho) = \sup_{\varphi\in\mathscr{R}}\left\{\int_{\mathbb{R}^2}\rho(\mathbf{r})\varphi(\mathbf{r})\mathrm{d}\mathbf{r} - C^*(\varphi)\right\},$$

where the star $*$ denotes the Legendre transform. It is then straightforward to check that C^* is the convex indicator of the $\{\varphi \leqslant c\}$, so that $C(\rho) \leqslant D(\rho)$, and recall that the converse inequality trivially holds. □

Remark 13. It later came to our attention that one can prove the weak lower-semicontinuity of $\rho \mapsto C(\rho)$ over $\mathscr{R}$ with a much more shorter argument. Indeed, let us take $\rho_n \rightharpoonup \rho \in \mathscr{R}$, and let $n' \to \infty$ be a subsequence such that $\liminf C(\rho_n) = \lim C(\rho_{n'})$. We can extract from n' a subsequence n'' such that the minimizers $\mathbb{P}_{n''}$ to $C(\rho_{n''})$ weakly converge to $\mathbb{P} \in \Pi^N(\rho)$. Then, one has

$$\liminf_{n\to\infty} C(\rho_n) = \lim_{n''\to\infty} \mathfrak{C}(\mathbb{P}_{n''}) \geqslant \liminf_{n''\to\infty} \mathfrak{C}(\mathbb{P}_{n''}) \geqslant \mathfrak{C}(\mathbb{P}) \geqslant C(\rho).$$

3. Coulomb gases

In this section, we present some physical applications to motivate the importance of optimal transport in physics, especially from a multi-marginal point-of-view. In the first subsection, we look at the regime $N \to \infty$, that this where we let the number of particles in the physical system becomes very large and formally grows to infinity. This regime is evidently of paramount interest in statistical mechanics, in the so-called *thermodynamic limit* [43, 83], where the macroscopic properties of a system of particles can be recovered since most thermal fluctuations arising from the stochastic nature of the microscopic modelization are erased in the limit. We formally present two basic models, namely *jellium* and the *uniform electron gas*, state some of their fundamental properties, and eventually discuss their equivalence to the first leading order. In the second subsection, we give a short introduction to *density-functional theory*, both from a quantum and classical perspective, and show how some problems arising in this field can be formulated through the scope of (M)OT.

3.1. The regime $N \to \infty$.

Jellium [59, 101] (also referred to as the *one-component plasma* (OCP)) is one of the most fundamental models in condensed matter physics. It has been used in a variety of physical applications, among which *density-functional theory* (DFT), where it serves as the fundamental building block of the so-called *local density approximation* (LDA) (see below). More generically, it is used as a rather crude model of delocalized electrons in metals and semiconductors, where it can qualitatively reproduce features of true materials (*e.g.* plasmons, crystallization etc.). Despite its apparent simplicity, the jellium model remains an open, challenging problem, especially from a mathematical perspective. For instance, jellium is often presumed to be equivalent to another model, namely the *uniform electron gas* (UEG) which arises as a multi-marginal optimal transport problem, but it was only very recently that a rigorous proof of this statement was given in the mathematical literature, first in [24] with a rather long and technical proof, and subsequently in [54] using a much shorter argument. This argument relies upon building an astute trial state for the optimal transport problem by modifying the *floating Wigner crystal* in such a way that its boundary is melted to a thin layer of incompressible fluid. The rationale remains valid in any dimension except in the two-dimensional case. It what follows, we introduce both models from a classical perspective, that is, without the kinetic energy.

3.1.1. *Classical jellium.*

Consider N electrons with positions $\mathbf{x}_1, \dots, \mathbf{x}_N$ in $\mathbb{R}^d$ interacting with each other through the Coulomb potential $\mathtt{c}_d$. The main idea of the jellium model is to get rid of the complications associated with the structure of the host material in which the delocalized electrons roam by replacing the actual structure of the background (*e.g.* the atomic lattice) by a homogeneous jelly-like

continuum $\Lambda_N \subset \mathbb{R}^d$ of positive charge. By scaling, we may assume that the density of the background is normalized, *i.e.* $\rho = 1$. We furthermore require that the overall system be neutral, that is $|\Lambda_N| = N$. The jellium energy is then defined as

$$\mathcal{E}_{\text{jel}}(\Lambda_N; \mathbf{x}_1, \dots, \mathbf{x}_N) := \sum_{1 \leqslant i < j \leqslant N} \mathsf{c}_d(\mathbf{x}_i - \mathbf{x}_j) - \sum_{i=1}^{N} \int_{\Lambda_N} \mathsf{c}_d(\mathbf{x}_i - \mathbf{y}) \mathrm{d}\mathbf{y} + \frac{1}{2} \iint_{\Lambda_N \times \Lambda_N} \mathsf{c}_d(\mathbf{y} - \mathbf{z}) \mathrm{d}\mathbf{y} \mathrm{d}\mathbf{z}. \quad (3)$$

The first term represents the particle-particle interaction energy, while the second accounts for the interaction energy between the background and the particles, and finally the third term is the background self-energy. We may then minimize this energy with respect to the positions $\mathbf{x}_i$'s. Note that it does not matter whether or not we constrain the particles to stay inside the background Λ_N, for after minimization they will always end up there. This is because the energy is harmonic with respect to each $\mathbf{x}_i$ on $\Lambda_N \setminus \{\mathbf{x}_1, \dots, \widetilde{\mathbf{x}_i}, \dots, \mathbf{x}_N\}$. The *jellium ground-state* e_{jel} is then defined as

$$e_{\text{jel}} := \lim_{\Lambda_N \nearrow \mathbb{R}^d} \min_{\mathbf{x}_1, \dots, \mathbf{x}_N} \frac{\mathcal{E}_{\text{jel}}(\Lambda_N; \mathbf{x}_1, \dots, \mathbf{x}_N)}{N}.$$

Under some natural conditions of the boundary of Λ_N (*i.e.* it must be regular enough in the sense that its area is negligible compared to $|\Lambda_N|$), the limit can be proved to exist and to be independent of the shape of the background [59]. For instance, we may think of $\Lambda_N = N^{1/d} \Lambda$ with Λ a fixed open convex set of unit volume. It is a famous conjecture, the so-called *crystallization conjecture*, that we shall expect the electrons to crystallize on a lattice whose geometry depends on the dimension d. For instance, in dimension three, it is conjectured that the electrons should crystallize on a bcc lattice [100, 101], though the longed-for proof of this statement is still currently missing. If this were proved to hold, we would have $e_{\text{jel}} = \zeta_{\text{bcc}}(1)$ where $\zeta_{\text{bcc}}(s)$ is the *Epstein zeta function* of the bcc lattice [10, 20]. In dimension two, it is proved in [86] that if crystallization occurs, then the particles must form a hexagonal lattice. In dimension one, the conjecture has been fully answered in [48], where it is proved that at zero temperature the particles form a lattice of unit step, or, for that matters, of step $1/\rho$ when the density is not normalized, and at positive temperature, it is proved that the particle density is periodic. Note that this result has also been extended to the quantum case in [11]. The crystallization conjecture has also been answered in dimension eight and twenty-four in [19, 75]. We refer the reader to the furnished review proposed in [10] for further discussions on the crystallization conjecture.

3.1.2. *Uniform Electron Gas.* The uniform electron gas was rigorously defined in [53, 55] and is obtained by assuming that the density of the electrons is exactly constant over a large domain Λ_N which grows such as to cover the whole space. That is, contrary to the jellium model, the compensating background which constrains the particles to remain bounded together is formally replaced by a constraint

on their electronic density. The *indirect Coulomb energy* of a given density ρ is defined as

$$\mathcal{E}_{\mathrm{ind}}(\rho) := \inf_{\substack{\mathbb{P}\in\mathcal{P}(\mathbb{R}^{dN}),\\ \rho_{\mathbb{P}}=\rho}} \left\{ \mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}} \left[\sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l) \right] \right\} - \frac{1}{2}\iint_{\mathbb{R}^d\times\mathbb{R}^d} \mathsf{c}_d(\mathbf{y}-\mathbf{z})\rho(\mathbf{y})\rho(\mathbf{z})\mathrm{d}\mathbf{y}\mathrm{d}\mathbf{z},$$

that is, the total Coulomb energy as previously defined to which we have subtracted the classical interaction energy of the density ρ, so that we only care for the remaining quantum energy arising from the *exchange* and *correlation* effects.The ground-state energy per unit volume e_{UEG} is then defined as

$$e_{\mathrm{UEG}} := \lim_{\Lambda_N\nearrow\mathbb{R}^d} \frac{\mathcal{E}_{\mathrm{ind}}(\mathbb{1}_{\Lambda_N})}{|\Lambda_N|}.$$

It is proved in [53] that under the same technical hypotheses on the domain Λ_N as for the jellium model, this limit exists and is independent of the shape of Λ_N. In fact, one can recover the value e_{UEG} as it arises as a limit of slowly varying densities. Indeed, it is proved in the three-dimensional case in [53] that given any fixed density ρ with $\int_{\mathbb{R}^3}\rho(\mathbf{r})\mathrm{d}\mathbf{r}=1$, one has

$$\lim_{N\to\infty} \frac{\mathcal{E}_{\mathrm{ind}}(\rho(\cdot/N^{1/3}))}{N} = e_{\mathrm{UEG}}\int_{\mathbb{R}^3}\rho(\mathbf{r})^{4/3}\mathrm{d}\mathbf{r}.$$

The interpretation is the following: if we think of splitting the space $\mathbb{R}^3$ using a tiling made of cubes of side length $1 \ll \ell \ll N^{1/3}$, we see that $\rho(\cdot/N^{1/3})$ is essentially constant in each of these cubes. The local energy in each cube can therefore be replaced by $e_{\mathrm{UEG}}(\rho_k)^{4/3}$ where ρ_k is the average value of ρ over the k-th cube. However, the total energy is evidently not local and there are interactions between the different cubes, so that proving this limit therefore demands to show that these interaction energies do not appear at the leading order.

For any N-particle probability measure $\mathbb{P}$ such that $\rho_{\mathbb{P}}=\mathbb{1}_{\Lambda_N}$,we have

$$\begin{aligned}\int_{\mathbb{R}^{dN}} \sum_{1\leqslant k<l\leqslant N} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l)\mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N) - \frac{1}{2}\iint_{\Lambda_N\times\Lambda_N}\mathsf{c}_d(\mathbf{y}-\mathbf{z})\mathrm{d}\mathbf{y}\mathrm{d}\mathbf{z}\\ = \int_{\mathbb{R}^{dN}} \mathcal{E}_{\mathrm{jel}}(\Lambda_N;\mathbf{x}_1,\dots,\mathbf{x}_N)\mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N)\\ \geqslant \min_{\mathbf{x}_1,\dots,\mathbf{x}_N} \mathcal{E}_{\mathrm{jel}}(\Lambda_N;\mathbf{x}_1,\dots,\mathbf{x}_N),\end{aligned}$$

so that after optimizing over all the probability measures $\mathbb{P}$ having the right one-particle density and passing to the thermodynamic limit $N\to\infty$, it holds that $e_{\mathrm{UEG}}\geqslant e_{\mathrm{jel}}$.

The classical UEG has been the object of many recent numerical works, based on methods from OT [39, 88, 89, 90, 91]. Note that the classical UEG has also been used to get numerical bounds on the best constant in the *Lieb-Oxford inequality* [52, 57, 60, 61, 68, 81].

3.1.3. *Floating Wigner Crystal.* Let us now describe the idea behind floating Wigner crystal in *three-dimension of space*, and why it fails to be a good trial state for the UEG. First, suppose for the moment that we have a proof that the jellium is crystallized in a bcc lattice, so that $e_{\text{jel}} = \zeta_{\text{bcc}}(1)$. To prove that $e_{\text{UEG}} \leqslant \zeta_{\text{bcc}}(1)$, an intuitive strategy would be to average the positions of the particles in the bcc lattice by translating this lattice over some given volume. Formally, let $\mathcal{L}$ be the bcc lattice, with *Wigner-Seitz* unit cell Q centered at 0, such that $|Q| = 1$ [80]. We place the particles on the intersection of the lattice $\mathcal{L}$ with a large cube C so that $\mathcal{L} \cap C = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ are the corresponding locii of the particles. We then take

$$\Lambda_N = \bigcup_{i=1}^{N} (Q + \mathbf{x}_i)$$

the union of the cells centered at the particles. The floating Wigner crystal [9, 30, 66] is then obtained by taking the Dirac distribution of the N particles, then translating by an amount $\mathbf{q} \in Q$ and integrating over the unit cell Q, which leads to the N-particle probability

$$\widetilde{\mathbb{P}} = \int_Q \delta_{\mathbf{x}_1+\mathbf{q}} \otimes \cdots \otimes \delta_{\mathbf{x}_N+\mathbf{q}} \mathrm{d}\mathbf{q},$$

where one evidently has $\rho_{\widetilde{\mathbb{P}}} = \mathbb{1}_{\Lambda_N}$. The indirect energy of this state is given by

$$\frac{1}{2N} \sum_{1 \leqslant k < l \leqslant N} \frac{1}{|\mathbf{x}_k - \mathbf{x}_l|} - \frac{1}{2N} \iint_{\Lambda_N \times \Lambda_N} \frac{\mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}}{|\mathbf{x} - \mathbf{y}|},$$

and in [52], it has been proved to converge to

$$\zeta_{\text{bcc}}(1) + \frac{2\pi}{3} \int_Q |\mathbf{x}|^2 \mathrm{d}\mathbf{x}.$$

As one remarks, a positive shift appears in the thermodynamic limit. Indeed, the difference $\mathbb{1}_{\Lambda_N - \mathbf{q}} - \mathbb{1}_{\Lambda_N}$ describes a monopole layer in a neighborhood of the surface that produces an electric potential felt by all the particles in the system. That is, moving the particles away from the center of the unit cells is not at all energetically favorable, since this creates a large excess of negative charges on one side and a corresponding excess of background charge on the opposite side, giving rise to the positive shift which survives in the thermodynamic limit (see Fig. 1).

To avoid the positive shift, the idea proposed in [54] is to immerse the crystal in a thin layer of incompressible fluid of unit density. That is, we choose a large cubic container C such that $\Lambda_N + Q \subset C$, such that the volume of the fluid $|C \setminus \Lambda_N| = M$ is an integer. We will furthermore demand that M be negligible compared to N, so that the fluid layer around the floating crystal will have a vanishing energy per unit volume in the thermodynamic limit. We consider the new trial state composed of the N particle on the floating crystal, translated as before by $q \in Q$, together with the M other particles forming a fluid around the crystal, that is

$$\mathbb{P} = \int_Q \delta_{\mathbf{x}_1+q} \otimes \cdots \otimes \delta_{\mathbf{x}_N+q} \otimes \left(\frac{\mathbb{1}_{C \setminus (\Lambda_N + q)}}{M} \right)^{\otimes M} \mathrm{d}q.$$

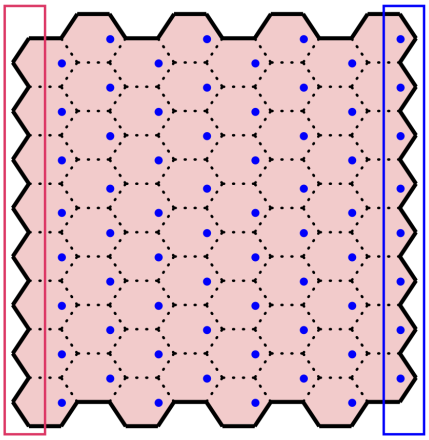

FIGURE 1. A two-dimensional picture of the jellium model from [56]. The blue dots represent the electrons, which are placed on a finite subset of a lattice $\mathcal{L}$. The red-colored set is the union of the corresponding unit cells Q and it represents a uniform background charge distribution of opposite charge. The indirect energy of the floating Wigner crystal is obtained after integrating the position of the lattice over the unit cell Q. When the lattice is not centered, this results in an excess of point charges on one side and an excess of background charge on the other side, indicated by the two rectangles.

The density of this trial state is given by $\rho_{\mathbb{P}} = \mathbb{1}_C + \mathbb{1}_{\Lambda_N} - \mathbb{1}_Q * \mathbb{1}_{\Lambda_N}$. Straightforward algebraic computations yields that

$$\int_{\mathbb{R}^{3(N+M)}} \sum_{1\leqslant k<l\leqslant N+M} \frac{1}{|\mathbf{x}_k-\mathbf{x}_l|} \mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N) = \sum_{1\leqslant k<l\leqslant N} \frac{1}{|\mathbf{x}_k-\mathbf{x}_l|} - \sum_{i=1}^{N} \int_{\Lambda_N} \frac{\mathrm{d}\mathbf{y}}{|\mathbf{x}_i-\mathbf{y}|} + D(\mathbb{1}_{\Lambda_N}) + D(\mathbb{1}_C) + 2D(\mathbb{1}_C, \mathbb{1}_{\Lambda_N} - \mathbb{1}_{\Lambda_N} * \mathbb{1}_Q) - \frac{1}{M}\int_Q D(\mathbb{1}_{C\setminus(\Lambda_N+\mathbf{q})})\mathrm{d}\mathbf{q},$$

where we have used the shorthand notation $D(f) := D(f,f)$ where D the Coulomb interaction energy as defined in the previous section. The first line corresponds to the jellium energy $\mathcal{E}_{\mathrm{jel}}(\Lambda_N;\mathbf{x}_1,\dots,\mathbf{x}_N)$ of the crystal, while the second corresponds to the quantity $D(\rho_{\mathbb{P}}) - D(\mathbb{1}_{\Lambda_N} - \mathbb{1}_{\Lambda_N} * \mathbb{1}_Q)$, as can be readily seen from the bilinearity of the functional D. Therefore, the indirect energy of this new trial state rewrites as

$$\mathcal{E}_{\mathrm{jel}}(\Lambda_N;\mathbf{x}_1,\dots,\mathbf{x}_N) - D(\mathbb{1}_{\Lambda_N} - \mathbb{1}_{\Lambda_N} * \mathbb{1}_Q) - \frac{1}{M}\int_Q D(\mathbb{1}_{C\setminus(\Lambda_N+\mathbf{q})})\mathrm{d}\mathbf{q}.$$

Using the fact that the inequality $D(f) \geqslant 0$ holds in dimension three of space, it follows that

$$\mathcal{E}_{\mathrm{ind}}(\rho_{\mathbb{P}}) \leqslant \mathcal{E}_{\mathrm{jel}}(\Lambda_N;\mathbf{x}_1,\dots,\mathbf{x}_N)$$

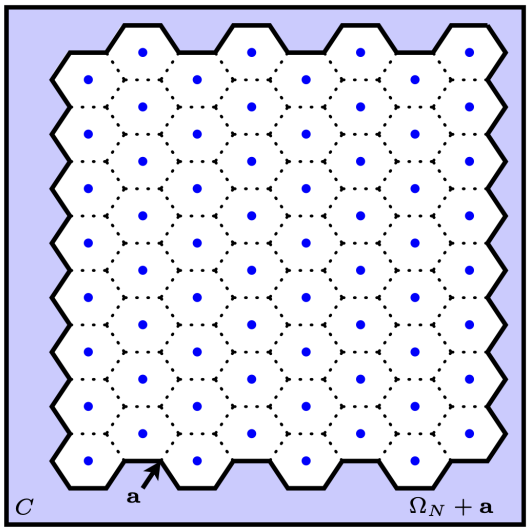

FIGURE 2. A two-dimensional picture of the modified floating crystal from [54] (where Ω_N is our Λ_N). The blue dots represents the electrons centered in the unit cells Q. When the whole crystal block Λ_N get displaced by an amount $\mathbf{q} \in Q$, the fluid gets displaced to fill the remaining space $C \setminus (\Lambda_N + \mathbf{q})$.

The density $\rho_{\mathbb{P}}$ is equal to 1 inside Λ_N, at a distance at least equal to the diameter of the unit cell Q from the boundary of Λ_N, and equal to 0 outside of C. Finally, it varies between 0 and 2 in the intermediate region. Using the fact that M is negligible compared to N, we make use of the relaxed formulation in [53] to conclude that

$$\lim_{N\to\infty} \frac{\mathcal{E}_{\text{ind}}(\rho_{\mathbb{P}})}{N} = e_{\text{UEG}},$$

eventually yielding that $e_{\text{UEG}} \leqslant \zeta_{\text{bcc}}(1)$ as claimed.

As it should now start to become apparent, the major difficulty encountered when trying to construct a good trial state for the UEG is mainly due to the boundary, namely to the fact that we work with a finite piece of material in the physical space $\mathbb{R}^d$. If one were to work on the torus, as is often done in practical calculations, these difficulties disappear. This is the rationale used in [54], where a third model is introduced, namely the *periodic jellium*, which is formally obtained when we repeat periodically a jellium configuration in the whole space and compute its energy per unit volume e_{per}. Then, using an argument based on the modified floating Wigner crystal as introduced above, the authors were able to prove that $e_{\text{UEG}} \leqslant e_{\text{per}}$. The equivalence follows from the fact that $e_{\text{per}} = e_{\text{jel}}$ [24].

3.2. **Applications to DFT.** *Density-functional theory* (DFT) [12, 29, 31, 71, 79] is one of the most widely used methods to conduct computations in condensed matter physics and quantum chemistry, for instance to give reasonably accurate quantitative and qualitative descriptions of the electronic structure of atoms, molecules, crystals and surfaces. Being able to carry out those computations in a realistic time-frame stands out as a key issue in a wide range of scientific industries, advocating for the overall importance of the field. DFT dates back to the seminal

work of Hohenberg and Kohn [44], and was later rigorously formalized from a mathematical point-of-view by Levy [50] and Lieb [58]. It entirely relies on the fundamental property that the *ground-state energy* E_0 of a many-electron quantum system can be obtained by minimizing some functional of the electronic density ρ. In what follows, let us give a self-contained primer on quantum DFT. We introduce the *Kohn-Sham approach* (KS-DFT) and the *Strictly Correlated Electrons approach* (SCE-DFT), and show that the latter is explicitly formulated as an MOT problem. We then discuss the importance of the UEG model in DFT by introducing the so-called *Local Density Approximation* (LDA). Finally, we give a rough introduction to classical DFT, where (M)OT plays a central role.

3.2.1. *A primer on quantum DFT.* For simplicity, let us only consider N electrons of dimension d together with some time-independent external potential V, with interaction potential given by the Coulomb potential c_d, and let us just remark that DFT applies to a wide range of other quantum particles and interaction potentials. Recall that the resulting quantum system is completely described by the operator H_N^V, the so-called *Hamiltonian*, defined as

$$H_N^V = \sum_{i=1}^{N} \left(-\frac{\hbar^2}{2} \Delta_{\mathbf{x}_i} + V(\mathbf{x}_i) \right) + \sum_{1 \leqslant k < l \leqslant N} \mathsf{c}_d(\mathbf{x}_k - \mathbf{x}_l),$$

where $\Delta_{\mathbf{x}_i}$ denotes the Laplacian with respect to the variable $\mathbf{x}_i$, and where $\hbar$ is the famous *Planck's constant* accounting for the quantization of the classical energy. Note that we neglect the *spin* variable for convenience. Let us denote by $Q(H_N^V)$ the *form domain* of the operator H_N^V (see [82, Section VIII.6]), which is a subspace of the Hilbert space $L^2(\mathbb{R}^{dN})$ or, to be more precise, of the subspace $L^2_{\mathrm{a}}(\mathbb{R}^{dN}) := \bigwedge_{i=1}^{N} L^2(\mathbb{R}^d)$ of antisymmetric wave functions, since the electrons are *fermionic* particles as they obey *Pauli exclusion principle*. Of paramount interest is the problem of identifying (if any) the *ground-state energy* of the system, that is the quantity

$$E_0 := \inf_{\substack{\Psi \in Q(H_N^V) \\ \|\Psi\|_{L^2(\mathbb{R}^{dN})} = 1}} \langle \Psi, H_N^V \Psi \rangle,$$

where the quantity $\langle \Psi, H_N^V \Psi \rangle$ is understood in the sense of *quadratic forms*. Solving this problem analytically is virtually infeasible, at the exception of very peculiar system (*e.g.* Hydrogen atom, harmonic oscillator etc.), and straightforward numerical schemes are usually impractical because of the high dimensionality of the problem. The main idea of DFT is then to replace the minimization over the wave function Ψ by a two-step minimization where we first minimize over the density ρ and then over all the wave functions having that prescribed density, that is

$$\inf_{\Psi} (\cdots) = \inf_{\rho} \inf_{\substack{\Psi \\ \rho_\Psi = \rho}} (\cdots).$$

We are then required to identify the set of N*-representable electronic densities*, that is those functions which arise as electronic densities of admissible quantum states, which according to [58], reduces to the set of positive function $\rho \in L^1(\mathbb{R}^d)$ of total mass $\int_{\mathbb{R}^d} \rho(\mathbf{r}) \mathrm{d}\mathbf{r} = N$ such that $\sqrt{\rho} \in H^1(\mathbb{R}^d)$. Therefore, the minimization

problem rewrites as

$$E_0 = \inf_{\substack{\sqrt{\rho}\in H^1(\mathbb{R}^d)\\ \int_{\mathbb{R}^d}\rho(\mathbf{r})\mathrm{d}\mathbf{r}=N}} \left\{ F_{LL}(\rho) + \int_{\mathbb{R}^d} V(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r} \right\}.$$

where the functional F_{LL}, the so-called *Levy-Lieb functional*, is defined as

$$F_{LL}(\rho) = \inf_{\substack{\Psi\in L^2_{\mathrm{a}}(\mathbb{R}^{dN})\\ \|\Psi\|_{L^2(\mathbb{R}^{dN})}=1\\ \rho_\Psi=\rho}} \left\{ \frac{\hbar^2}{2}\sum_{i=1}^{N}\int_{\mathbb{R}^{dN}} |\nabla_{\mathbf{x}_i}\Psi(\mathbf{x}_1,\ldots,\mathbf{x}_N)|^2 \mathrm{d}\mathbf{x}_1\ldots\mathrm{d}\mathbf{x}_N \right.$$
$$\left. + \sum_{1\leqslant k<l\leqslant N}\int_{\mathbb{R}^{dN}} \mathsf{c}_d(\mathbf{x}_k-\mathbf{x}_l)|\Psi(\mathbf{x}_1,\ldots,\mathbf{x}_N)|^2\mathrm{d}\mathbf{x}_1\ldots\mathrm{d}\mathbf{x}_N \right\}.$$

This quantity corresponds to the lowest possible (kinetic plus electron-electron interaction) energy of a quantum system having the prescribed density ρ. Note that it is proved in [58] that this infimum is attained.

This universal functional is the central object of DFT, since knowing it would allow one to compute the ground-state energy of a system with any external potential V. Nevertheless, very little is known about this functional, so that in practice, one needs to use approximations. The very core of practical DFT lies within the ability to provide approximations as accurate and efficient as possible.

3.2.2. *Kohn-Sham approach.* One famous route is the one proposed by Kohn and Sham [47], based on the assumption that there exists a reference non-interacting system of fermions which has exactly the same ground-state density as the one from the initial, physical system. This roughly amounts to purporting that the kinetic energy largely dominates electron-electron interactions. Formally, we take N orthonormal functions $\Phi = (\varphi_1,\ldots,\varphi_N)$ describing N fictitious uncorrelated electrons to build the desired density through the formula

$$\rho_\Phi(\mathbf{r}) := \sum_{i=1}^{N}|\varphi_i(\mathbf{r})|^2 = \rho_\Psi$$

where Ψ is a *Slater determinant*, that is

$$\Psi(\mathbf{x}_1,\ldots,\mathbf{x}_N) = \det[(\varphi_i(\mathbf{x}_j))_{i,j}]/\sqrt{N!}.$$

For a density ρ with $\int_{\mathbb{R}^d}\rho(\mathbf{r})\mathrm{d}\mathbf{r} = N$, we introduce the lowest kinetic energy of Slater determinants

$$T_S(\rho) := \inf_{\substack{\varphi_1,\ldots,\varphi_N\in H^1(\mathbb{R}^d)\\ \langle\varphi_i,\varphi_j\rangle=\delta_{ij}\\ \rho_\Phi=\rho}} \frac{\hbar^2}{2}\sum_{i=1}^{N}\int_{\mathbb{R}^d}|\nabla\varphi_i(\mathbf{x}_1,\ldots,\mathbf{x}_N)|^2\mathrm{d}\mathbf{x}_1\ldots\mathrm{d}\mathbf{x}_N.$$

We then add and subtract T_S from F_{LL}, which allows us to rewrite the N-particle ground-state as

$$E_0 = \inf_{\substack{\varphi_1,\dots,\varphi_N \in H^1(\mathbb{R}^d) \\ \langle \varphi_i, \varphi_j \rangle = \delta_{ij}}} \left\{ \frac{\hbar^2}{2} \sum_{i=1}^{N} \int_{\mathbb{R}^d} |\nabla \varphi_i(\mathbf{x}_1, \dots, \mathbf{x}_N)|^2 \mathrm{d}\mathbf{x}_1 \dots \mathrm{d}\mathbf{x}_N \right.$$
$$\left. + \int_{\mathbb{R}^d} V(\mathbf{r}) \rho_{\boldsymbol{\Phi}}(\mathbf{r}) \mathrm{d}\mathbf{r} + \frac{1}{2} \iint_{\mathbb{R}^d \times \mathbb{R}^d} \mathsf{c}_d(\mathbf{x} - \mathbf{y}) \rho_{\boldsymbol{\Phi}}(\mathbf{x}) \rho_{\boldsymbol{\Phi}}(\mathbf{y}) \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y} + E_{\mathrm{xc}}(\rho_{\boldsymbol{\Phi}}) \right\},$$

where E_{xc}, the so-called *exchange-correlation energy*, is exactly defined as the difference

$$E_{\mathrm{xc}}(\rho) := F_{LL}(\rho) - T_S(\rho) - \frac{1}{2} \iint_{\mathbb{R}^d \times \mathbb{R}^d} \mathsf{c}_d(\mathbf{x} - \mathbf{y}) \rho_{\boldsymbol{\Phi}}(\mathbf{x}) \rho_{\boldsymbol{\Phi}}(\mathbf{y}) \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y},$$

that is E_{xc} has no explicit definition other than being the remaining quantity needed to recover the Levy-Lieb functional, and it is thus an interesting question to find a way to study E_{xc} directly, without interpreting it as a difference.

3.2.3. *Strictly Correlated Electrons DFT.* In the case where the system displays strong correlation, another route to follow is the so-called *Strictly Correlated Electrons approach* (SCE-DFT), first proposed in [91], which stands as the exact counterpart of KS-DFT: the reference system is purported to have infinite electronic correlation and zero kinetic energy. The roughly amounts to first minimizing the electron-electron interaction and then correct with the remaining energy, that is we write

$$F_{LL}(\rho) = \inf_{\substack{\Psi \in L^2_{\mathrm{a}}(\mathbb{R}^{dN}) \\ \|\Psi\|_{L^2(\mathbb{R}^{dN})} = 1 \\ \rho_\Psi = \rho}} \left\{ \int_{\mathbb{R}^{dN}} \sum_{1 \leqslant k < l \leqslant N} \mathsf{c}_d(\mathbf{x}_k - \mathbf{x}_l) \mathrm{d}|\Psi(\mathbf{x}_1, \dots, \mathbf{x}_N)|^2 \right\}$$
$$+ E_{kd}(\rho)$$

where E_{kd} is the *kinetic-decorrelation energy*, accounting for the remaining energy. Let us mention that SCE-DFT is an exact theory at the semi-classical limit, as proved when $N = 2$ in [22], when $N = 3$ in [8], and eventually generalized to the cases $N \geqslant 4$ in [23, 51]. As it should now become apparent, we remark that SCE-DFT exactly consists in solving a MOT problem with transport cost associated with the pairwise interaction potential (here with the Coulomb potential).

3.2.4. *Local density approximation.* The universal functional F_{LL} defined in the previous section allow in principle to describe any quantum system of electrons, though, as noticed before, this functional is of course not known exactly, so that one must find reliable and efficient approximations. The most widely used of these approximations is the so-called the *Local Density Approximation* (LDA) [29, 44, 47, 63, 71, 73] . The LDA is often considered as "*the mother of all approximations*" [74] and it yields surprisingly good results, even in cases where the density is not at all slowly varying [63, 71]. Its successors involving gradient corrections are even better and have become the standard in DFT calculations. In what follows, we present LDA in three dimension of space.

The functional F_{LL} is not local at all, since two electrons at different locii are always entangled and, furthermore, the Coulomb potential has a very long range

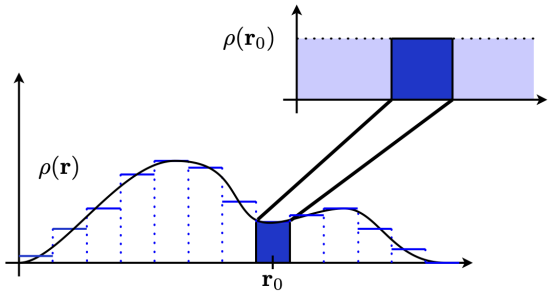

FIGURE 3. From [56]. The Levy-Lieb energy (with the classical Coulomb energy subtracted) is replaced by the sum of the energies per unit volume of an infinite uniform gas with the local density $\rho(\mathbf{r}_0)$, times the volume $d\mathbf{r}$.

so that electrons interact even when they are far apart. In the LDA, one makes the assumption that the only non-local part is the classical Coulomb energy of the density ρ and one approximates the rest of the energy by a local function of ρ, that is, the integral of a function f depending only on the value $\rho(\mathbf{r})$ at $\mathbf{r}$:

$$F_{LL}(\rho) \approx \frac{1}{2}\iint_{\mathbb{R}^3\times\mathbb{R}^3} \frac{\rho(\mathbf{x})\rho(\mathbf{y})}{|\mathbf{x}-\mathbf{y}|}d\mathbf{x}d\mathbf{y} + \int_{\mathbb{R}^3} f(\rho(\mathbf{r}))d\mathbf{r}.$$

The function f is usually taken to be the UEG itself, so that the approximation becomes exact when ρ is constant over a very large domain. The idea behind the LDA is depicted at Fig. 3. After subtraction of the classical Coulomb energy of ρ, one splits the space into small boxes of volume $d\mathbf{r}$ and assumes that the remaining energy is the sum of the local energies in each box. In each little box, one replaces the density by a constant. One does not use the energy of the constant function in the small box, but rather the energy per unit volume of an infinite system having the corresponding uniform density, multiplied by the volume $d\mathbf{r}$ of the small box.

The LDA was first justified rigorously in [55], where it is proved that there exists some constant $C > 0$ such that for every density ρ (verifying some integrability hypotheses) and for every $\varepsilon > 0$, we have

$$\left|F_{LL}(\rho) - \frac{1}{2}\iint_{\mathbb{R}^3\times\mathbb{R}^3} \frac{\rho(\mathbf{x})\rho(\mathbf{y})}{|\mathbf{x}-\mathbf{y}|}d\mathbf{x}d\mathbf{y} - \int_{\mathbb{R}^3} e_{\mathrm{UEG}}(\rho(\mathbf{r}))d\mathbf{r}\right|$$
$$\leqslant \varepsilon\int_{\mathbb{R}^3}(\rho(\mathbf{r})+\rho(\mathbf{r})^2)d\mathbf{r} + \frac{C(1+\varepsilon)}{\varepsilon}\int_{\mathbb{R}^3}|\nabla\sqrt{\rho}(\mathbf{r})|^2d\mathbf{r} + \frac{C}{\varepsilon^{15}}\int_{\mathbb{R}^3}|\nabla\sqrt{\rho}(\mathbf{r})|^4d\mathbf{r},$$

where F_{LL} is the *grand-canonical* version of the Levy-Lieb functional (*i.e.* the number of particles is not fixed), though it is expected that the exact same result holds for the canonical functional as defined above.

3.2.5. *Classical DFT.* So far, we have introduced DFT from a quantum point-of-view, and shown how methods from (M)OT were inherently present in the bigger picture. Here, let us introduce the classical counterpart to DFT, which reads as follows: « *Given a density of particles ρ at some fixed finite temperature β^{-1}, classical DFT is about finding the external potential V for which the equilibrium measure of a system of particles has density ρ* ».

For an infinite translation-invariant gas, the density is constant over the whole space, so that V need also be constant, with $V = -\mu$, where μ is called the *chemical potential*. In the usual method of statistical mechanics, the problem of finding the chemical potential μ which stems a given density ρ is the usual *equivalence of ensemble* [83]. At small *activity* $z = e^{\beta\mu} \ll 1$, the observables of the system, including the density, can usually be expanded in terms of z, using the so-called *Mayer* or *cluster expansion* as first proposed in [64]. Inverting the function $\rho = \rho(z) = z + \dots$ provides expansions in terms of ρ (*i.e. Virial expansion* [69]), as first proposed in [49] (see also [83, Chapter 4]).

The goal of classical DFT is to extend this rationale to inhomogeneous systems. Most practical applications are for infinite systems (*i.e.* $\int_{\mathbb{R}^d} \rho(\mathbf{r})\mathrm{d}\mathbf{r} = \infty$), for instance ρ can be equal to two different constants over two half space (and smooth in the transition region) so as to describe the interface between two phases. Nevertheless, the theory is also important for finite systems (*i.e.* $\int_{\mathbb{R}^d} \rho(\mathbf{r})\mathrm{d}\mathbf{r} < \infty$). As a matter of fact, understanding infinite systems often starts by looking at a finite system and then taking the thermodynamic limit. For infinite systems it evidently should not matter whether the problem is settled *canonically* or *grand-canonically*, so that we usually work grand-canonically since the algebra is usually simpler.

Let us formally describe the model for finite systems. In the canonical case, as already mentioned earlier in this note, we seek to solve the following equation

$$\rho(\mathbf{r}) = \frac{e^{-\beta V(\mathbf{r})}N}{Z_N(V)} \int_{\mathbb{R}^{d(N-1)}} e^{-\beta E_N(\mathbf{r},\mathbf{x}_2,\dots,\mathbf{x}_N) - \beta\sum_{i=2}^N V(\mathbf{x}_i)} \mathrm{d}\mathbf{x}_2 \dots \mathrm{d}\mathbf{x}_N,$$

where $Z_N(V)$ is the partition function that has already been defined above and where E_N is the N-particle energy, for instance

$$E_N(\mathbf{x}_1, \dots, \mathbf{x}_N) := \sum_{1 \leqslant k < l \leqslant N} w(\mathbf{x}_k - \mathbf{x}_l)$$

where w is the pairwise interaction potential (*e.g.* $w = \mathtt{c}_d$). As mentioned in the introduction of this note, this is nothing else but the entropic regularization of the MOT problem, with the transport cost defined by E_N. In the grand-canonical case, the equation reads

$$\rho(\mathbf{r}) = \frac{e^{-\beta V(\mathbf{r})}}{Z(V)} \Bigg(e^{-\beta E_1(\mathbf{r})} + \sum_{n \geqslant 2} \frac{1}{(n-1)!} \int_{\mathbb{R}^{d(n-1)}} e^{-\beta E_n(\mathbf{r},\mathbf{x}_2,\dots,\mathbf{x}_N) - \beta\sum_{i=2}^n V(\mathbf{x}_i)} \mathrm{d}\mathbf{x}_2 \dots \mathrm{d}\mathbf{x}_n \Bigg)$$

with the grand-canonical partition function

$$Z(V) := 1 + \sum_{n \geqslant 1} \frac{1}{n!} \int_{\mathbb{R}^{dn}} e^{-\beta E_n(\mathbf{x}_1,\mathbf{x}_2,\dots,\mathbf{x}_N) - \beta\sum_{i=1}^n V(\mathbf{x}_i)} \mathrm{d}\mathbf{x}_1 \dots \mathrm{d}\mathbf{x}_n$$

Remark 14. For simplicity, we have assumed that $E_0 = 0$. Note that, in practice, one often has that $E_1 = 0$ and that $E_n \geqslant -Bn$ for all $n \neq 2$, where $B > 0$ is a constant (*i.e.* the so-called *H-stability*, see [83]). Therefore, we readily get the bounds

$$1 + \int_{\mathbb{R}^d} e^{-\beta V(\mathbf{r})}\mathrm{d}\mathbf{r} \leqslant Z(V) \leqslant \exp\left(e^{\beta B} \int_{\mathbb{R}^d} e^{-\beta V(\mathbf{r})}\mathrm{d}\mathbf{r} \right)$$

and

$$\exp\left(-e^{\beta B}\int_{\mathbb{R}^d} e^{-\beta V(\mathbf{r})}\mathrm{d}\mathbf{r}\right) e^{-\beta V(\mathbf{r})} \leqslant \rho(\mathbf{r}) \leqslant \exp\left(-e^{\beta B}\int_{\mathbb{R}^d} e^{-\beta V(\mathbf{r})}\mathrm{d}\mathbf{r}\right) e^{-\beta V(\mathbf{r})} e^{\beta B}.$$

We see that we should require the quantity $\int_{\mathbb{R}^d} e^{-\beta V(\mathbf{r})}\mathrm{d}\mathbf{r}$ to be finite. Moreover, when ρ vanishes, we have $V = +\infty$, which is very intuitive from a physical point-of-view. As such, it is often convenient to write $e^{-\beta V} = e^{-\beta U}\rho$ and to work with the variable U instead of V.

Besides the equations above, the sought-for potential V can be determined through variational methods, as one can prove that it solves the maximization problem

$$\sup_V \left\{-\beta^{-1}\log Z_N(V) - \int_{\mathbb{R}^d} \rho(\mathbf{r})V(\mathbf{r})\mathrm{d}\mathbf{r}\right\}$$

in the canonical case (as well as in the grand-canonical case, by replacing $Z_N(V)$ by $Z(V)$). The uniqueness of V (modulo additive constants in the canonical case) is usually called the *Hohenberg-Kohn-Mermin theorem* since the argument is similar to the one for quantum systems [44, 65].

The existence and uniqueness of the potential V yielding the prescribed density for finite systems is proved in the canonical case in [16] under certain hypotheses. The grand-canonical case is more convoluted because the number of particles can be arbitrarily large [15]. For infinite systems, we refer the reader to [45]. More generally, we refer the reader to the references [3, 4, 32, 33, 41, 77, 93, 102, 103] for further discussions on classical DFT.

Note that, for atomic gases, the interaction w is not known and should be determined empirically. A problem of the same flavor as DFT, the so-called *inverse Henderson problem* [35, 42], consists in finding V and w which provide a given one-particle density ρ and pair density $\rho^{(2)}$.

4. MOT WITH LOGARITHMIC COST AT FINITE TEMPERATURE ($\beta > 0$)

In this section, we study the problem of computing the Coulomb energy of a system of two-dimensional N charged particles $\mathbf{x}_1, \dots, \mathbf{x}_N$ with prescribed electronic density ρ at finite temperature, that is with entropic regularization (*i.e.* $\beta > 0$),

$$C_\beta(\rho) = \inf_{\mathbb{P}\in\Pi^N(\rho)} \left\{\mathbb{E}_{(\mathbf{x}_1,\dots,\mathbf{x}_N)\sim\mathbb{P}}\left[-\sum_{1\leqslant k<l\leqslant N}\log|\mathbf{x}_k-\mathbf{x}_l|\right] + \beta^{-1}\mathrm{Ent}(\mathbb{P}|\rho^{\otimes N})\right\}. \quad \text{(Coul2DReg)}$$

Let us formally derive a meaningful dual to our problem. As usual, we can make the explicit constraint $\mathbb{P}\in\mathbb{P}^N(\rho)$ into an implicit constraint by adding the convex indicator $\mathbb{1}^{\mathrm{cvx}}_{\Pi^N(\rho)}$ to the objective (*i.e.* $\mathbb{1}^{\mathrm{cvx}}_{\Pi^N(\rho)}(\mathbb{P}) = 0$ if $\mathbb{P}\in\Pi^N(\rho)$ and $+\infty$ otherwise) and by taking the infimum over all probability measures $\mathbb{P}$. But, by

using the fact that

$$\mathbb{1}^{\mathrm{cvx}}_{\Pi^N(\rho)}(\mathbb{P}) = \sup_{\varphi\in C_b(\mathbb{R}^2)} \left\{ \int_{\mathbb{R}^2} \varphi(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r} - \int_{\mathbb{R}^{2N}} \left(\sum_{i=1}^N \varphi(\mathbf{x}_i) \right) \mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N) \right\},$$

we can formally swap the infimum and the supremum, and introduce

$$D_\beta(\rho) = \sup_{\varphi\in C_b(\mathbb{R}^2)} \left\{ E_\beta(\varphi) + \int_{\mathbb{R}^2} \varphi(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r} \right\} \tag{dCoul2Dreg}$$

where we have defined as

$$E_\beta(\varphi) := \inf_{\mathbb{P}\in\mathcal{P}(\mathbb{R}^{2N})} \left\{ \int_{\mathbb{R}^{2N}} \left(c(\mathbf{x}_1,\dots,\mathbf{x}_N) - \sum_{i=1}^N \varphi(\mathbf{x}_i) \right) \mathrm{d}\mathbb{P}(\mathbf{x}_1,\dots,\mathbf{x}_N) + \beta^{-1}\mathrm{Ent}(\mathbb{P}|\rho^{\otimes N}) \right\}.$$

It is then an easy application of Jensen's inequality that we have the so-called *Gibbs variational principle*, which says that

$$E_\beta(\varphi) = -\beta^{-1}\log\left(\int_{\mathbb{R}^{2N}} e^{-\beta\mathcal{H}_{-\varphi}(\mathbf{x}_1,\dots,\mathbf{x}_N)} \mathrm{d}\rho^{\otimes N}(\mathbf{x}_1,\dots,\mathbf{x}_N) \right),$$

where the infimum is attained for the Gibbs measure $\mathbb{P}_{\beta,-\varphi}$ associated with the Hamiltonian $\mathcal{H}_{-\varphi}$ and the reference measure $\rho^{\otimes N}$, that is,

$$\mathbb{P}_{\beta,-\varphi} = Z_\beta(-\varphi)^{-1} e^{-\beta\mathcal{H}_{-\varphi}(\mathbf{x}_1,\dots,\mathbf{x}_N)} \mathrm{d}\rho^{\otimes N}(\mathbf{x}_1,\dots,\mathbf{x}_N).$$

Remark 15. By defining the functional E_β, we took the infimum over all probability measures $\mathbb{P} \in \mathcal{P}(\mathbb{R}^{2N})$, while in the literature the infimum is usually taken over all positive measures $\mathbb{P} \in \mathcal{M}_+(\mathbb{R}^{2N})$. Doing so, we certainly end up with a different dual, but which has the advantage that the dual variable φ is only meaningful up to an additive constant, just like a true physical potential.

Under the assumption that strong duality holds, that is $C_\beta(\rho) = D_\beta(\rho)$, solving (Coul2DReg) is therefore equivalent to finding the (unique up to an additive constant) potential $\varphi^\star$ which solves (dCoul2Dreg). As we will now soon show, $\varphi^\star$ is the potential which solves for every $i = 1,\dots,N$ the equations

$$\int_{\mathbb{R}^{2(N-1)}} \mathbb{P}_{\beta,-\varphi}(\mathbf{x}_1,\dots,\mathbf{x}_N)\mathrm{d}\mathbf{x}_1\dots\widehat{\mathrm{d}\mathbf{x}_i}\dots\mathrm{d}\mathbf{x}_N = N^{-1}\rho(\mathbf{x}_i) \tag{eqCoul2D}$$

For the rest of the discussion, we make the assumption that

$$\int_{\mathbb{R}^2} (1+|\mathbf{r}|^2)\rho(\mathbf{r})\mathrm{d}\mathbf{r} < \infty, \tag{H3}$$

Remark 16. We demand the stronger hypothesis (H3) compared to (H1) for purely didactic reasons. Indeed, this allows us to give *a* proof of the lower semi-continuity of the relative entropy functional (Lemma 9). But as explained below in Remark 17, this hypothesis is actually superfluous, so that, altogether, it is only necessary to demand that ρ verifies hypothesis (H1) as previously. Let us also note that in the dual formulation (dCoul2Dreg), under hypothesis (H3), we can replace the constraint $\varphi \in C_b(\mathbb{R}^2)$ by the constraint that $\varphi \in L^\infty_{\mathrm{loc}}(\mathbb{R}^2)$ with $\varphi(\mathbf{r})/(1+|\mathbf{r}|^2) \in L^\infty(\mathbb{R}^2)$. This gives us the *ad hoc* topology to work on for the proof of strong duality (Theorem 12).

4.1. **Finiteness of $C_\beta(\rho)$ and consistence as $\beta \to \infty$.**

Lemma 9. *Let α and β be Borel probability measures over $\mathbb{R}^n$, and let the sequence $\{\alpha_k\}_k$ be such that $\alpha_k \rightharpoonup \alpha$ as $k \to \infty$ tightly. Moreover, suppose there exists a constant $C > 0$ so that*

$$\int_{\mathbb{R}^n} (1 + |\mathbf{r}|^2) \mathrm{d}\alpha_k(\mathbf{r}) < C, \quad \forall k.$$

Then, we have

$$\liminf_{k\to\infty} \mathrm{Ent}(\alpha_k|\beta) \geqslant \mathrm{Ent}(\alpha|\beta).$$

Proof of Lemma 9. We suppose that $\alpha_k \ll \beta$ for all k, otherwise the result is trivial. Let $B_R = B(0, R) \subset \mathbb{R}^n$ be a ball of radius $R > 0$. We have

$$\begin{aligned}\int_{\mathbb{R}^n} \log\left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right) \mathrm{d}\alpha_k(\mathbf{r}) &= \int_{B_R} \log\left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right) \mathrm{d}\alpha_k(\mathbf{r}) \\ &+ \int_{(\mathbb{R}^n\setminus B_R)\cap\left\{\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}\geqslant 1\right\}} \log\left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right) \mathrm{d}\alpha_k(\mathbf{r}) \\ &+ \int_{(\mathbb{R}^n\setminus B_R)\cap\left\{\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}< 1\right\}} \log\left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right) \mathrm{d}\alpha_k(\mathbf{r}).\end{aligned}$$

The two first terms are weakly lower semi-continuous according respectively to [62, Theorem B.33] and [2, Theorem 2.34]. Let us now take care of the last term. For any $\varepsilon \in (0, 1)$, there exists a constant $K_\varepsilon > 0$ such that $\log(t) \geqslant -K_\varepsilon t^{-\varepsilon}$ for all $t > 0$, so that

$$\begin{aligned}-K_\varepsilon \int_{\mathbb{R}^n\setminus B_R} \left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right)^{-\varepsilon} \mathrm{d}\alpha_k(\mathbf{r}) \\ \leqslant \int_{(\mathbb{R}^n\setminus B_R)\cap\left\{\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}< 1\right\}} \log\left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right) \mathrm{d}\alpha_k(\mathbf{r}) \leqslant 0.\end{aligned}$$

By taking $\varepsilon < \frac{1}{2}$ and using Hölder inequality, it holds

$$\begin{aligned}\int_{\mathbb{R}^n\setminus B_R} \left(\frac{\mathrm{d}\alpha_k}{\mathrm{d}\beta}(\mathbf{r})\right)^{-\varepsilon} \mathrm{d}\alpha_k(\mathbf{r}) &\leqslant \left(\int_{\mathbb{R}^n\setminus B_R} (1 + |\mathbf{r}|^2) \mathrm{d}\alpha_k(\mathbf{r})\right)^{1-\varepsilon} \left(\frac{1}{1+R^2}\right)^{1-\varepsilon} \\ &\leqslant C^{1-\varepsilon} \left(\frac{1}{1+R^2}\right)^{1-\varepsilon} := \delta(R),\end{aligned}$$

with $\delta(R)$ vanishing as R goes to infinity. Therefore, letting $R \to \infty$ yields the desired result, that is,

$$\liminf_{k\to\infty} \mathrm{Ent}(\alpha_k|\beta) \geqslant \mathrm{Ent}(\alpha|\beta).$$

□

Remark 17. In the above lemma, we have proved the lower semi-continuity for the weak-$\star$ topology of the relative entropy functional in the first variable under the hypothesis of bounded second-moments. In fact, one can prove the much stronger statement that $\mathrm{Ent}(\cdot|\cdot)$ is *jointly* lower semi-continuous without demanding any boundedness hypothesis, that is, given any sequences of probability measures

$\{\alpha_k\}_k$ and $\{\beta_k\}_k$ such that $\alpha_k \rightharpoonup \alpha$ and $\beta_k \rightharpoonup \beta$ as $k \to \infty$, where α and β are probability measures as well, it holds that

$$\liminf_{k\to\infty} \mathrm{Ent}(\alpha_k|\beta_k) \geqslant \mathrm{Ent}(\alpha|\beta).$$

A proof of this statement can be found in [78], where an equivalent definition of the relative entropy is used. Let us give a different (rough sketch of a) proof.

Theorem 10. *The infimum at* (Coul2DReg) *is attained.*

Proof of Theorem 10. From Theorem 1, it holds that the cost functional is lower semi-continuous. According to Lemma 9, the relative entropy is lower semi-continuous under uniformly bounded second-moments. But for any $\mathbb{P} \in \Pi^N(\rho)$, then by hypothesis (H3) we have

$$\int_{\mathbb{R}^{2N}} |(\mathbf{x}_1, \dots, \mathbf{x}_n)|^2 \mathrm{d}\mathbb{P}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \int_{\mathbb{R}^2} |\mathbf{r}|^2 \rho(\mathbf{r}) \mathrm{d}\mathbf{r} < \infty,$$

yielding the sought-for uniform bound. □

Remark 18. If ρ verifies the hypotheses (H2) and (H3), it is immediate that $C_\beta(\rho)$ is finite using Theorem 3.

Theorem 11. *We have* $\lim_{\beta\to\infty} C_\beta(\rho) = C(\rho)$*, so that the minimizers of* (Coul2DReg) *converge to a minimizer of* (Coul2D) *as* $\beta \to \infty$ *(up to a subsequence).*

Proof of Theorem 11. Let us roughly sketch a proof of this statement. Morally, the difficulty arises from the fact that an optimal transport plan for the unregularized problem $C(\rho)$ need not have finite relative entropy with respect to the reference measure $\rho^{\otimes N}$ (*i.e.* it might typically live on some submanifold of $\mathbb{R}^{2N}$). The idea is therefore to *regularize* this optimal plan without changing too much its cost. More formally, because the relative entropy is a positive quantity, we always have $C_\beta(\rho) \geqslant C(\rho)$. If we consider a transport plan $\mathbb{P} \in \Pi^N(\rho)$ with *finite relative entropy*, that is $\mathrm{Ent}(\mathbb{P}|\rho^{\otimes N}) < \infty$, we have $\limsup_{\beta\to\infty} C_\beta(\rho) \leqslant \mathfrak{C}(\mathbb{P})$. Therefore, we obtain

$$\limsup_{\beta\to\infty} C_\beta(\rho) \leqslant C'(\rho),$$

where the quantity $C'(\rho)$ is defined as

$$C'(\rho) := \inf_{\substack{\mathbb{P}\in\Pi^N(\rho),\\ \mathrm{Ent}(\mathbb{P}|\rho^{\otimes N})<\infty}} \mathfrak{C}(\mathbb{P}).$$

We trivially have that $C'(\rho) \geqslant C(\rho)$, and it remains to prove that the converse inequality holds, that is $C'(\rho) \leqslant C(\rho)$, which will yield that $\lim_{\beta\to\infty} C_\beta(\rho) = C(\rho)$. To do so, we use the strategy proposed in [14] in the case where the cost c is the Euclidean distance. Given an optimal transport plan $\mathbb{P}^\star$ for $C(\rho)$, we approach $\mathbb{P}^\star$ using the so-called *block approximation*, that is, roughly speaking, by decomposing the ambient space $\mathbb{R}^{2N}$ into a partition of small hypercubes $C_{i,\ell} = c_{i,\ell,1} \times \cdots \times c_{i,\ell,N}$, where $c_{i,\ell,j} \subset \mathbb{R}^2$ are squares whose areas vanish as $\ell \to 0$, and by approximating $\mathbb{P}^\star$ in each hypercube by the re-weighted tensor measure

$$\mathbb{P}^\star_{i,\ell} := \frac{\mathbb{P}(C_{i,\ell})}{\frac{\rho}{N}(c_{i,\ell,1}) \cdots \frac{\rho}{N}(c_{i,\ell,N})} \frac{\rho}{N}\Big|_{c_{i,\ell,1}} \otimes \cdots \otimes \frac{\rho}{N}\Big|_{c_{i,\ell,N}}.$$

That is, we consider $\mathbb{P}^\star_\ell := \sum_i \mathbb{P}^\star_{i,\ell}$, and we expect that $\mathbb{P}^\star_\ell$ shall become a good approximate of $\mathbb{P}^\star$ as $\ell \ll 1$. For the Coulomb cost is not as sympathetic as the

Euclidian cost, one needs to be careful, especially because of the singularities on the diagonals. Nevertheless, because the optimal transport $\mathbb{P}^\star$ is supported away from the diagonals, we eventually swerve away from this unfortunate technicality, and it is proved in [37] that $\lim_{\ell\to 0}\mathfrak{C}(\mathbb{P}^\star_\ell)=\mathfrak{C}(\mathbb{P}^\star)$ and that $\mathbb{P}^\star_\ell$ has finite relative entropy. Therefore, one has $C'(\rho)\leqslant\mathfrak{C}(\mathbb{P}^\star_\ell)$, and passing to the limit $\ell\to 0$ yields $C'(\rho)\leqslant\mathfrak{C}(\mathbb{P}^\star)=C(\rho)$, concluding the proof. □

4.2. **Strong duality for $C_\beta(\rho)$.** Let us now prove that strong duality holds. Let $\mathcal{Q}:=L^1((1+|\mathbf{r}|^2)\mathrm{d}\mathbf{r})$, that is, the set of ρ's verifying hypothesis (H3), and we write

$$D_\beta(\rho):=\sup_{\varphi\in\mathcal{Q}'}\left\{E_\beta(\varphi)+\int_{\mathbb{R}^2}\varphi(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r}\right\},$$

where $\mathcal{Q}'$ is the topological dual of $\mathcal{Q}$, which consists, as already noticed at Remark 16, of those functions $\varphi\in L^\infty_{\mathrm{loc}}(\mathbb{R}^2)$ such that $\varphi(\mathbf{r})/(1+|\mathbf{r}|^2)\in L^\infty(\mathbb{R}^2)$, endowed with the natural weak topology.

Theorem 12. *For every $\beta>0$, strong duality holds,* i.e. *$C_\beta(\rho)=D_\beta(\rho)$, for every $\rho\in\mathcal{Q}$.*

Proof of Theorem 12. The proof is similar to that of Theorem 8. We have $D_\beta(\rho)=(-E_\beta)^*(\rho)$ by definition of the Legendre transform. But $E_\beta(\varphi)=-C^*_\beta(\varphi)$ by definition of the functional E_β, so that $D_\beta(\rho)=(C^*_\beta)^*(\rho)$. The functional C_β is convex as a pointwise infimum of the strictly convex functional $\mathfrak{C}_\beta$. Furthermore, it is not identically $+\infty$ as $\mathcal{Q}\cap L^q\neq\emptyset$. Finally, C_β is weakly lower semi-continuous over $\mathcal{Q}$ for the same reasons as before (see Remark 13). Therefore C_β is a so-called regular convex function, and by Fenchel's theorem, it holds that $C_\beta(\rho)=(C^*_\beta)^*(\rho)$ for every $\rho\in\mathcal{Q}$, yielding strong duality. □

Let us now build a bridge between the optimal solutions of our dual (dCoul2Dreg) and the equations (eqCoul2D). This will allow us to shift our attention to the equations themselves, which offer a more pragmatic hold on the problem compared to the variational formulation. Note that this bridge should not come to the reader as an unsettling statement, for it is already a well-known result of the literature in the discrete case. Note that, we make the simplifying hypothesis that ρ is compactly supported, though this need not be the weakest required hypothesis, as explained below.

Theorem 13. *Suppose that ρ is compactly supported, and let $\varphi^\star$ be an optimal solution to* (dCoul2Dreg)*. Then, we have*

$$\int_{\mathbb{R}^{2(N-1)}}\mathbb{P}_{\beta,-\varphi^\star}(\mathbf{x}_1,\dots,\mathbf{x}_N)\mathrm{d}\mathbf{x}_1\dots\widehat{\mathrm{d}\mathbf{x}_i}\dots\mathrm{d}\mathbf{x}_N=N^{-1}\rho(\mathbf{x}_i)$$

for every $\mathbf{x}_i\in\Lambda_N$, and every $i=1,\dots,N$. Moreover, the unique minimizer $\mathbb{P}_\beta$ to (Coul2DReg) *is exactly $\mathbb{P}_{\beta,-\varphi^\star}$.*

Proof of Theorem 13. To prove the first statement of the theorem, it formally suffices to differentiate the dual with respect to the potential φ. We refer the reader to [18, Theorem 4.7], where the proof is provided in the case where ρ is compactly supported for the usual dual of (Coul2DReg), but the proof is unchanged for our alternative dual formulation (dCoul2Dreg). In order to compute the directional derivative of the functional $\varphi\mapsto Z_\beta(\varphi)$, one typically needs to use the dominated

convergence theorem, and this is where the hypothesis that ρ be compactly supported comes handy, though we could work with weaker hypotheses (*e.g.* ρ has some sort of exponential decay). To prove the second statement, we first remark that $\mathbb{P}_{\beta,-\varphi^\star} \in \Pi^N(\rho)$ by definition. Then, straightforward computations give that

$$\mathfrak{C}(\mathbb{P}_{\beta,-\varphi^\star}) + \beta^{-1}\mathrm{Ent}(\mathbb{P}_{\beta,-\varphi^\star}|\rho^{\otimes N}) = E_\beta(\varphi^\star) + \int_{\mathbb{R}^2} \varphi^\star(\mathbf{r})\rho(\mathbf{r})\mathrm{d}\mathbf{r},$$

where of the right-hand side of the above equation is exactly $D_\beta(\rho)$ by optimality of $\varphi^\star$, so that by strong duality, the desired result follows. □

4.3. **Toward the UEG.** From now on, we turn our attention towards the UEG, as defined in the previous section. As such, we will always work with $\rho = \mathbb{1}_{\Lambda_N}$ where $\Lambda_N := N^{1/2}\Lambda$, with $\Lambda := B(0,1)$ the unit ball of unit radius in $\mathbb{R}^2$. Working with a ball makes the theoretical analysis easier because of the radially symmetry , though we expect, at least from a qualitative point-of-view, similar results for domains which are not too pathological.

As explained in the introduction of this document, the dual variable φ does not necessarily make up for a good candidate as the optimization variable from a numerical perspective. Heuristically, φ is interpreted as a true physical potential which binds the particles together. As such, instead of working on the potential landscape itself, we might consider looking for a distribution of charge which generates it. That is, from now on, our dual variable will be a distribution of charge ρ_{ext}, and we will write $U^{\rho_{\mathrm{ext}}}$ the associated potential, that is

$$U^{\rho_{\mathrm{ext}}}(\mathbf{r}) := \mathtt{c}_2 * \rho_{\mathrm{ext}}(\mathbf{r}) = -\int_{\mathbb{R}^2} \log(\mathbf{r}-\mathbf{z})\rho_{\mathrm{ext}}(\mathbf{z})\mathrm{d}\mathbf{z}$$

Theorem 14. *If ρ_{ext} is an optimal charge distribution, that is $U^{\rho_{\mathrm{ext}}}$ is optimal for* (dCoul2Dreg)*, then one can always suppose that* $\mathrm{supp}(\rho_{\mathrm{ext}}) \subset \Lambda_N$.

Proof of Theorem 14. Notice that the objective in (dCoul2Dreg) is strictly concave up to an additive constant. Because of the radial symmetry of the domain Λ_N, it then holds that ρ_{ext} must be a radially symmetric Borel measure. By the *Lebesgue decomposition theorem*, we can write

$$\rho_{\mathrm{ext}} = \rho_{\mathrm{pp}} + \rho_{\mathrm{sing}} + \rho_{\mathrm{abs}}\mathrm{d}\mathbf{r},$$

where ρ_{pp} is the *pure point* part, ρ_{sing} is the *singular* part, and where ρ_{abs} is the *absolutely continuous* (with respect to the Lebesgue measure $\mathrm{d}\mathbf{r}$) part. Then, because of the radially symmetric nature of ρ_{ext}, it must be that $\rho_{\mathrm{pp}} = \omega_0\delta_0$ where $\omega_0 \in \mathbb{R}$ is some weight and that ρ_{sing} is supported on circles centered at the origin, that is, $\rho_{\mathrm{sing}} = \sum_{i\geqslant 1} \frac{\omega_i}{2\pi r_i}\sigma_{C(r_i)}$, where $\sigma_{C(r_i)}$ is the Lebesgue measure on the circle $C(r_i)$ centered at the origin of radius $r_i > 0$. Now, it is a well-known that a uniformly charged circle generates a constant potential on its inside, so that we can always remove the singular parts outside of Λ_N. Likewise, if there is an absolutely continuously distributed charge on some annulus outside of Λ_N, the potential felt by any particle inside Λ_N is also constant, so that one can always remove this part as well. Therefore, we can always suppose that $\mathrm{supp}(\rho_{\mathrm{ext}}) \subset \Lambda_N$.

□

Remark 19. It's not very complicated to prove that we can always suppose $\rho_{\text{sing}} = 0$. Indeed, if some charge is supported on a circle inside the domain Λ_N, retracting this circle toward the origin always gives rise to a higher energy for the dual (dCoul2Dreg).

Remark 20. If we suppose that ρ_{abs} verifies hypothesis (H2), that is, $\rho_{\text{abs}} \in L^q(\mathbb{R}^2)$ for some $q > 1$, then one can also prove that $\omega_0 = 0$, so that any optimal charge distribution ρ_{ext} can be supposed to be absolutely continuous with respect to the Lebesgue measure. Indeed, let us write $\rho_{\text{ext}} = \omega_0\delta_0 + \rho'$, so that $U^{\rho_{\text{ext}}}(\mathbf{r}) = -\omega_0 \log|\mathbf{r}| + U^{\rho'}(\mathbf{r})$. Using (eqCoul2D), we have for every $\mathbf{r} \in \Lambda_N$ (with $N = 2$ for simplicity)

$$Z_\beta(\rho_{\text{ext}})^{-1} \frac{1}{|\mathbf{r}|^{\beta\omega_0}} e^{\beta U^{\rho'}(\mathbf{r})} \int_{\Lambda_N} |\mathbf{r}-\mathbf{z}|^\beta e^{U^{\rho_{\text{ext}}}(\mathbf{z})} d\mathbf{z} = N^{-1},$$

from which it immediately follows, by evaluating the equation at $\mathbf{r} = 0$, that one must have $\omega_0 = 0$. Note that the hypothesis on ρ_{abs} assures us that $U^{\rho'}(0)$ is finite.

In what follows, we continue making the hypothesis that ρ_{abs} verifies (H2). As such, ρ_{ext} is absolutely continuous and we do not write ρ_{abs} anymore, according to what precedes. Note that we refer the reader to the extensive monograph [84] for a thorough introduction of logarithmic potential theory. In the next theorem, we prove that ρ_{ext} must positive. To do so, we use of the fact that one can retrieve the charge distribution ρ_{ext} which generates the potential $U^{\rho_{\text{ext}}}$ as $\rho_{\text{ext}} = -(2\pi)^{-1}\Delta U^{\rho_{\text{ext}}}$ (see [84, Theorem 1.3]). Of course, this requires the potential to have continuous second partial derivatives, which is the case when $\beta > 2$ by dominated convergence using (eqCoul2D).

Theorem 15. *ρ_{ext} is positive.*

Proof of Theorem 15. For simplicity, let us suppose that $N = 2$, the proof being similar for general N, only more cumbersome to write. For every $\mathbf{x} \in \text{Int}(\Lambda_N)$, we have

$$-\nabla U^{\rho_{\text{ext}}}(\mathbf{x}) = \frac{N}{Z(\rho_{\text{ext}})} e^{U^{\rho_{\text{ext}}}(\mathbf{x})} \int_{\Lambda_N} (\mathbf{x}-\mathbf{y})|\mathbf{x}-\mathbf{y}|^{\beta-2} e^{U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y}. \tag{4}$$

Therefore, one has

$$\begin{aligned} -\Delta U^{\rho_{\text{ext}}}(\mathbf{x}) = \beta \frac{N}{Z(\rho_{\text{ext}})} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{x})} \Bigg[& \int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\beta-2} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y} \\ & - \frac{N}{Z(\rho_{\text{ext}})} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{x})} \left\| \int_{\Lambda_N} (\mathbf{x}-\mathbf{y})|\mathbf{x}-\mathbf{y}|^{\beta-2} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y} \right\|^2 \Bigg]. \end{aligned}$$

Using the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\beta-1} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y} &= \int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\frac{\beta}{2}-1} e^{\frac{\beta}{2} U^{\rho_{\text{ext}}}(\mathbf{y})} |\mathbf{x}-\mathbf{y}|^{\frac{\beta}{2}} e^{\frac{\beta}{2} U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y} \\ &\leqslant \sqrt{\int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\beta-2} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y}} \sqrt{\int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\beta} e^{\beta U^{\rho_{\text{ext}}}(\mathbf{y})} d\mathbf{y}} \end{aligned}$$

Using the equation, it follows that

$$\left\| \int_{\Lambda_N} (\mathbf{x}-\mathbf{y})|\mathbf{x}-\mathbf{y}|^{\beta-2} e^{\beta U^{\rho_{\mathrm{ext}}}(\mathbf{y})} \mathrm{d}\mathbf{y} \right\|^2 \leqslant \int_{\Lambda_N} |\mathbf{x}-\mathbf{y}|^{\beta-2} e^{\beta U^{\rho_{\mathrm{ext}}}(\mathbf{y})} \mathrm{d}\mathbf{y} \times \frac{Z(\rho_{\mathrm{ext}})}{N} e^{-\beta U^{\rho_{\mathrm{ext}}}(\mathbf{x})},$$

and the result follows. □

5. Numerical experiments

In this last section, we present some numerical experiments for the UEG in two space-dimensions. As explained in Section 3 in the three-dimensional case, the floating Wigner crystal does not make up for a good candidate as a trial state for the UEG, while the modified version introduced in [54], where the crystal is immersed in a thin layer of incompressible fluid, certainly is. The question remains, « *shall we expect a similar paradigm in the two-dimensional case ?* ». Answering this question numerically poses a lot of difficulties, for it requires one to compute the optimal transport plan with rather great *quantitative* precision, in order to either confirm or infirm that the optimal solution displays fluid-like particle-particle correlations on the boundary of the system. This behavior shall especially become apparent when $N \gg 1$, which evidently withholds our investigation because of the high computational burden of considering a large number of particles. Therefore, in this document, we shall only care whether or not « *something is happening on the boundary* ». We only investigate the numerics from a *qualitative* point-of-view. This is altogether a first step, which ultimately should pave the way towards more precise results.

5.1. **Numerical scheme.** We will suppose that the density of the electrons is exactly constant over the square $\Lambda_N = [0, \sqrt{N}]^2$, and we look at the system at finite temperature β^{-1}, that is with entropic regularization. We want to solve the dual formulation (dCoul2Dreg) to our problem, using the strategy presented in the previous section, that is by working with the charge distribution ρ_{ext} instead of the corresponding potential. Note that we choose the domain Λ_N to be a square instead of a disk to simplify the implementation of our numerical scheme. The heuristic of this scheme is the following: we decompose our domain into small square cells C_i's, that is $\Lambda_N = \sqcup_i C_i$, and we want to approach the optimal charge distribution ρ_{ext} by $\sum_i \nu_i \mathbb{1}_{C_i}$, where the ν_i's are the optimization variables, representing the (constant and positive) charge distributions in the cells C_i's. We start with the jellium, that is $\nu_{i,(0)} = 1$ for every i. At iteration t, we approximate the one-particle density $\rho_{i,(t)}$ of $\mathbb{P}_{\beta,-U^{\rho_{(t)}}}$ in each cell, at the center of the cell, where $\rho_{(t)} = \sum_i \nu_{i,(t)} \mathbb{1}_{C_i}$. Note that to robustify the procedure, we might select a random point in the cell instead of its center. If $\rho_{i,(t)} < N^{-1}$, we want to correct for the lack of electrons by adding a small amount $\varepsilon_{(t)} > 0$ of positive charge, that is $\nu_{i,(t+1)} := \nu_{i,(t)} + \varepsilon_{(t)}$, and inversely by subtracting this amount when $\rho_{i,(t)} > N^{-1}$. We stop the algorithm when the error $\sum_i |\rho_{i,(t)} - N^{-1}|$ is as small as it can get.

Remark 21. We take the step $\varepsilon_{(t)}$ to decrease slowly in the sense that $\sum_t \varepsilon_{(t)} = +\infty$ and $\sum_t \varepsilon_{(t)}^2 < +\infty$, following a well-known *thumb rule*. Note that we could

imagine more sophisticated strategies, such as keeping track of an approximate gradient with respect to the weights v_i's.

The most difficult part, that is the one which carries the entire the computational burden, is of course the approximation of the one-particle density $\rho_{i,(t)}$ in each cell. This requires to compute both the partition function $Z_\beta(-U^{\rho_{(t)}})$ and the integral depending on the remaining $N-1$ two-dimensional variables. This integral depends on the $U^{\rho_{(t)}}$, which itself is obtained from $\rho_{(t)}$ by computing a convolution. We therefore proceed as follows: the partition function and the one-particle density are computed with *Monte-Carlo simulation* (MC). To compute the potential $U^{\rho_{(t)}}$ at the location $\mathbf{r} \in C_i$, we use the decomposition

$$U^{\rho_{(t)}}(\mathbf{r}) = -v_i \int_{C_i} \log|\mathbf{r}-\mathbf{z}|\mathrm{d}\mathbf{z} - \sum_{j\neq i} v_j \int_{C_j} \log|\mathbf{r}-\mathbf{z}|\mathrm{d}\mathbf{z}.$$

The first integral is computed using a straightforward MC simulation, while the other integrals are approximated as follows: given c_j the center of the cell C_j, we write for all $\mathbf{z} \in C_j$ the Taylor expansion

$$\begin{aligned}\log|\mathbf{r}-\mathbf{z}| = \log|\mathbf{r}-c_j| &+ \sum_{\alpha=1,2} \left(\nabla_{\mathbf{x}=|\mathbf{r}-c_j|}\log|\cdot|\right)_\alpha (c_j-\mathbf{z})_\alpha \\ &+ \frac{1}{2}\sum_{\alpha,\beta=1,2} \left(\nabla^2_{\mathbf{x}=|\mathbf{r}-c_j|}\log|\mathbf{r}-c_j|\right)_{\alpha,\beta} (c_j-\mathbf{z})_\alpha (c_j-\mathbf{z})_\beta + \dots\end{aligned}$$

so that, by defining the *monopole* m, the *dipoles* d_α's, the *quadrupoles* $q_{\alpha,\beta}$'s etc. as

$$m := |C_j|, \quad d_\alpha := \int_{C_j} (c_j-\mathbf{z})_\alpha \mathrm{d}\mathbf{z}, \quad q_{\alpha,\beta} := \int_{C_j} (c_j-\mathbf{z})_\alpha (c_j-\mathbf{z})_\beta \mathrm{d}\mathbf{z}, \quad \cdots,$$

we can write the *multipole expansion* around c_j as

$$\begin{aligned}\int_{C_j} \log|\mathbf{r}-\mathbf{z}|\mathrm{d}\mathbf{z} = m\log|\mathbf{r}-c_j| &+ \sum_{\alpha=1,2} d_\alpha \left(\nabla_{\mathbf{x}=|\mathbf{r}-c_j|}\log|\cdot|\right)_\alpha \\ &+ \frac{1}{2}\sum_{\alpha,\beta=1,2} q_{\alpha,\beta}\left(\nabla^2_{\mathbf{x}=|\mathbf{r}-c_j|}\log|\mathbf{r}-c_j|\right)_{\alpha,\beta} + \dots\end{aligned}$$

We then remark that, because of the symmetric geometry of C_j, some moments vanish, for instance $d_\alpha = 0$ for $\alpha = 1,2$ and $q_{\alpha,\beta} = 0$ when $\alpha \neq \beta$, simplifying the computations.

Remark 22. We believe that it might be possible to express the electric potential defined by the first integral using some special functions, as it is done in [17] for a uniformly charged square in two dimensions for the Coulomb potential $\mathtt{c}_3$. This would allow us to swerve away from the computational burden of using a MC simulation.

5.2. **Experiments.** Our experiments were conducted on the clusters of the *CEREMADE* laboratory, each cluster being equipped with 40 CPUs *Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz*. Our code was written using the `Julia` programming language, which we believe to be quite appropriate for the task. The code was parallelized and launched over the 40 CPUs at our disposal, allowing us to get

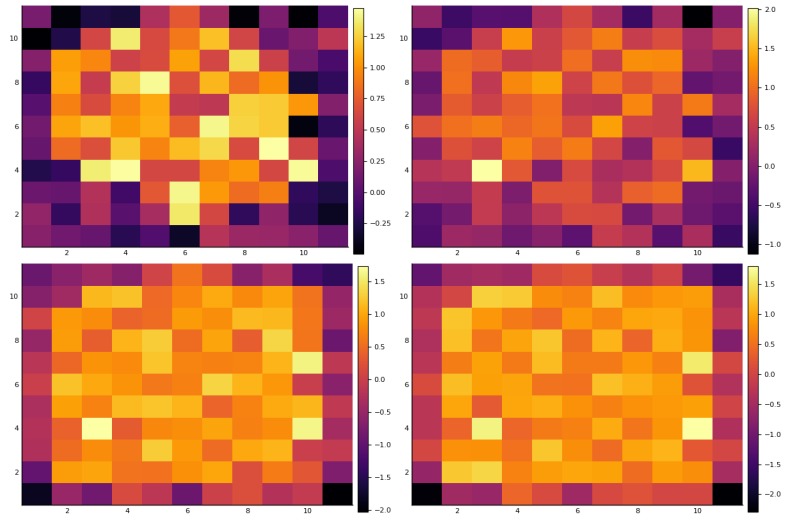

FIGURE 4. Parameters: $N = 5$, $M = 11$, $\beta = 0.01$. We display several iterations of $\rho_{(t)}$, and we see that the density of charge distribution decreases on the boundary of the system.

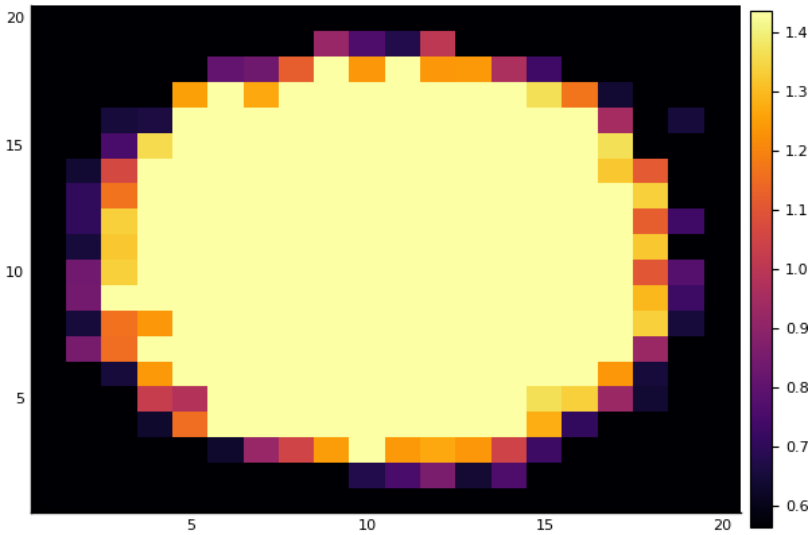

FIGURE 5. Parameters: $N = 5$, $M = 20$, $\beta = 0.01$. After many iterations, it becomes apparent that the density of the external charge distribution ρ_{ext} is much lower on the boundary of the system compared to the center.

rather correct qualitative results in a few hours ($\sim 4h$). In Fig. 4, we displayed a color-map of $\rho_{(t)}$ at several iterations, for $N = 5$ particles with a domain discretized into M^2 square cells with $M = 11$. As expected, we do see timidly appearing something at the boundary, where the approximate optimal external charge distribution decreases. In Fig. 5, by increasing the domain discretization to $M = 20$, this becomes clearly apparent. Reminiscent of Section 3, we know that the number of particles M constituting the neighboring fluid should be such

that $M \ll N$, where N is the number of particles on the floating Wigner crystal, so that the fluid layer has a vanishing energy per unit volume in the thermodynamic limit. Our experiments tend to confirm something of that type.

Remark 23. Throughout all of our experiments, the inverse temperature is fixed at $\beta = 0.01$, which altogether is a rather deceiving value. Indeed, we were not able to increase β without *underflowing* our computations. In the coming future, we will be implementing the usual trick of carrying the computations into the log-domain, we should allow us to increase β into the domain of much higher values.

REFERENCES

[1] Agueh, M. and Carlier, G. (2011). Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924.

[2] Ambrosio, L., Fusco, N., and Pallara, D. (2000). *Functions of bounded variation and free discontinuity problems*, volume 254. Clarendon Press Oxford.

[3] Baus, M. (1987). Statistical mechanical theories of freezing: An overview. *Journal of statistical physics*, 48(5-6):1129–1146.

[4] Baus, M. and Lutsko, J. F. (1991). Statistical mechanical theories of freezing: Where do we stand? *Physica A: Statistical Mechanics and its Applications*, 176(1):28–36.

[5] Beiglböck, M., Léonard, C., and Schachermayer, W. (2009). A general duality theorem for the monge–kantorovich transport problem. *arXiv preprint arXiv:0911.4347*.

[6] Benamou, J.-D., Carlier, G., and Nenna, L. (2016). A numerical method to solve multi-marginal optimal transport problems with coulomb cost.

[7] Benamou, J.-D., Carlier, G., and Nenna, L. (2019). Generalized incompressible flows, multi-marginal transport and sinkhorn algorithm. *Numerische Mathematik*, 142(1):33–54.

[8] Bindini, U. and De Pascale, L. (2017). Optimal transport with coulomb cost and the semiclassical limit of density functional theory. *Journal de l'École polytechnique—Mathématiques*, 4.

[9] Bishop, R. and Lührmann, K. (1982). Electron correlations. ii. ground-state results at low and metallic densities. *Physical Review B*, 26(10):5523.

[10] Blanc, X. and Lewin, M. (2015). The crystallization conjecture: a review. *arXiv preprint arXiv:1504.01153*.

[11] Brascamp, H. and Lieb, E. (2002). Some inequalities for gaussian measures and the long-range order of the one-dimensional plasma. In *Inequalities*, pages 403–416. Springer.

[12] Burke, K. and Wagner, L. O. (2013). Dft in a nutshell. *International Journal of Quantum Chemistry*, 113(2):96–101.

[13] Buttazzo, G., Champion, T., and De Pascale, L. (2018). Continuity and estimates for multi-marginal optimal transportation problems with singular costs. *Applied Mathematics & Optimization*, 78(1):185–200.

[14] Carlier, G., Duval, V., Peyré, G., and Schmitzer, B. (2017). Convergence of entropic schemes for optimal transport and gradient flows. *SIAM Journal on Mathematical Analysis*, 49(2):1385–1418.

[15] Chayes, J. and Chayes, L. (1984). On the validity of the inverse conjecture in classical density functional theory. *Journal of statistical physics*, 36(3-4):471–488.

[16] Chayes, J., Chayes, L., and Lieb, E. H. (1984). The inverse problem in classical statistical mechanics. *Communications in Mathematical Physics*, 93(1):57–121.

[17] Ciftja, O. (2010). Coulomb self-energy and electrostatic potential of a uniformly charged square in two dimensions. *Physics Letters A*, 374(7):981–983.

[18] Clason, C., Lorenz, D. A., Mahler, H., and Wirth, B. (2020). Entropic regularization of continuous optimal transport problems. *Journal of Mathematical Analysis and Applications*, page 124432.

[19] Cohn, H., Kumar, A., Miller, S. D., Radchenko, D., and Viazovska, M. (2019). Universal optimality of the e_8 and leech lattices and interpolation formulas. *arXiv preprint arXiv:1902.05438*.

[20] Coldwell-Horsfall, R. A. and Maradudin, A. A. (1960). Zero-point energy of an electron lattice. *Journal of Mathematical Physics*, 1(5):395–404.

[21] Colombo, M., Di Marino, S., and Stra, F. (2019). Continuity of multimarginal optimal transport with repulsive cost. *SIAM Journal on Mathematical Analysis*, 51(4):2903–2926.

[22] Cotar, C., Friesecke, G., and Klüppelberg, C. (2013). Density functional theory and optimal transportation with coulomb cost. *Communications on Pure and Applied Mathematics*, 66(4).
[23] Cotar, C., Friesecke, G., and Klüppelberg, C. (2018). Smoothing of transport plans with fixed marginals and rigorous semiclassical limit of the hohenberg–kohn functional. *Archive for Rational Mechanics and Analysis*, 228(3):891–922.
[24] Cotar, C. and Petrache, M. (2017). Equality of the jellium and uniform electron gas next-order asymptotic terms for coulomb and riesz potentials. *arXiv preprint arXiv:1707.07664*.
[25] Coyaud, R., Ehrlacher, V., Lombardi, D., and Alfonsi, A. (2019). Approximation of optimal transport problems with marginal moments constraints.
[26] Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300.
[27] De Pascale, L. (2015). Optimal transport with coulomb cost. approximation and duality. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1643–1657.
[28] Deming, W. E. and Stephan, F. F. (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444.
[29] Dreizler, R. M. and Gross, E. K. (2012). *Density functional theory: an approach to the quantum many-body problem*. Springer Science & Business Media.
[30] Drummond, N., Radnai, Z., Trail, J., Towler, M., and Needs, R. (2004). Diffusion quantum monte carlo study of three-dimensional wigner crystals. *Physical Review B*, 69(8):085116.
[31] Engel, E. and Dreizler, R. M. (2013). *Density functional theory*. Springer.
[32] Evans, R. (1979). The nature of the liquid-vapour interface and other topics in the statistical mechanics of non-uniform, classical fluids. *Advances in Physics*, 28(2):143–200.
[33] Evans, R. (1992). Density functionals in the theory of nonuniform fluids. *Fundamentals of inhomogeneous fluids*, 1:85–176.
[34] Forrester, P. J. (2010). *Log-gases and random matrices (LMS-34)*. Princeton University Press.
[35] Frommer, F., Hanke, M., and Jansen, S. (2019). A note on the uniqueness result for the inverse henderson problem. *Journal of Mathematical Physics*, 60(9):093303.
[36] Gerolin, A., Kausamo, A., and Rajala, T. (2019). Duality theory for multi-marginal optimal transport with repulsive costs in metric spaces. *ESAIM: Control, Optimisation and Calculus of Variations*, 25:62.
[37] Gerolin, A., Kausamo, A., and Rajala, T. (2020). Multi-marginal entropy-transport with repulsive cost. *Calculus of Variations and Partial Differential Equations*, 59:1–20.
[38] Gibbs, J. W. (1902). *Elementary principles in statistical mechanics: developed with especial reference to the rational foundations of thermodynamics*. C. Scribner's sons.
[39] Gori-Giorgi, P. and Seidl, M. (2010). Density functional theory for strongly-interacting electrons: perspectives for physics and chemistry. *Physical Chemistry Chemical Physics*, 12(43):14405–14419.
[40] Grimaldo, J. A. M. and Téllez, G. (2015). Relations among two methods for computing the partition function of the two-dimensional one-component plasma. *Journal of Statistical Physics*, 160(1):4–28.
[41] Hansen, J.-P. and McDonald, I. R. (1990). *Theory of simple liquids*. Elsevier.
[42] Henderson, R. (1974). A uniqueness theorem for fluid pair correlation functions. *Physics Letters A*, 49(3):197–198.
[43] Hill, T. L. (1994). *Thermodynamics of small systems*. Courier Corporation.
[44] Hohenberg, P. and Kohn, W. (1964). Inhomogeneous electron gas. *Physical review*, 136(3B):B864.
[45] Jansen, S., Kuna, T., and Tsagkarogiannis, D. (2019). Virial inversion and density functionals. *arXiv preprint arXiv:1906.02322*.
[46] Kellerer, H. G. (1984). Duality theorems for marginal problems. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 67(4):399–432.
[47] Kohn, W. and Sham, L. J. (1965). Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133.
[48] Kunz, H. (1974). The one-dimensional classical electron gas. *Annals of Physics*, 85(2):303–335.
[49] Lebowitz, J. and Penrose, O. (1964). Convergence of virial expansions. *Journal of Mathematical Physics*, 5(7):841–847.

[50] Levy, M. (1979). Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v-representability problem. *Proceedings of the National Academy of Sciences*, 76(12):6062–6065.

[51] Lewin, M. (2018). Semi-classical limit of the levy–lieb functional in density functional theory. *Comptes Rendus Mathématiques*, 356(4).

[52] Lewin, M. and Lieb, E. H. (2015). Improved lieb-oxford exchange-correlation inequality with a gradient correction. *Physical Review A*, 91(2):022507.

[53] Lewin, M., Lieb, E. H., and Seiringer, R. (2018). Statistical mechanics of the uniform electron gas. *Journal de l'École polytechnique—Mathématiques*, 5:79–116.

[54] Lewin, M., Lieb, E. H., and Seiringer, R. (2019a). Floating wigner crystal with no boundary charge fluctuations. *Physical Review B*, 100(3).

[55] Lewin, M., Lieb, E. H., and Seiringer, R. (2019b). The local density approximation in density functional theory. *Pure and Applied Analysis*, 2(1).

[56] Lewin, M., Lieb, E. H., and Seiringer, R. (2019c). Universal functionals in density functional theory. *arXiv preprint arXiv:1912.10424*.

[57] Lieb, E. H. (1979). A lower bound for coulomb energies. *Physics Letters A*, 70(5-6):444–446.

[58] Lieb, E. H. (1983). Density functionals for coulomb systems. *International Journal of Quantum Chemistry*, 24(3).

[59] Lieb, E. H. and Narnhofer, H. (1975). The thermodynamic limit for jellium. *Journal of Statistical Physics*, 12(4):291–310.

[60] Lieb, E. H. and Oxford, S. (1981). Improved lower bound on the indirect coulomb energy. *International Journal of Quantum Chemistry*, 19(3):427–439.

[61] Lieb, E. H. and Seiringer, R. (2010). *The stability of matter in quantum mechanics*. Cambridge University Press.

[62] Lott, J. and Villani, C. (2009). Ricci curvature for metric-measure spaces via optimal transport. *Annals of Mathematics*, pages 903–991.

[63] Lundqvist, S. and March, N. H. (2013). *Theory of the inhomogeneous electron gas*. Springer Science & Business Media.

[64] Mayer, J. E. and Montroll, E. (1941). Molecular distribution. *The Journal of Chemical Physics*, 9(1):2–16.

[65] Mermin, N. D. (1965). Thermal properties of the inhomogeneous electron gas. *Physical Review*, 137(5A):A1441.

[66] Mikhailov, S. A. and Ziegler, K. (2002). Floating wigner molecules and possible phase transitions in quantum dots. *The European Physical Journal B-Condensed Matter and Complex Systems*, 28(1):117–120.

[67] Nenna, L. (2016). *Numerical methods for multi-marginal optimal transportation*. PhD thesis.

[68] Odashima, M. M. and Capelle, K. (2007). How tight is the lieb-oxford bound? *The Journal of chemical physics*, 127(5):054106.

[69] Onnes, H. K. (1991). Expression of the equation of state of gases and liquids by means of series.". In *Through Measurement to Knowledge*, pages 146–163. Springer.

[70] Onsager, L. (1939). Electrostatic interaction of molecules. *Journal of Physical Chemistry*, 43(2):189–196.

[71] Parr, R. G. (1980). Density functional theory of atoms and molecules. In *Horizons of quantum chemistry*, pages 5–15. Springer.

[72] Pass, B. (2014). Multi-marginal optimal transport: theory and applications.

[73] Perdew, J. P. and Kurth, S. (2003). Density functionals for non-relativistic coulomb systems in the new century. In *A primer in density functional theory*, pages 1–55. Springer.

[74] Perdew, J. P. and Schmidt, K. (2001). Jacob's ladder of density functional approximations for the exchange-correlation energy. In *AIP Conference Proceedings*, volume 577, pages 1–20. American Institute of Physics.

[75] Petrache, M. and Serfaty, S. (2020). Crystallization for coulomb and riesz interactions as a consequence of the cohn-kumar conjecture. *Proceedings of the American Mathematical Society*, 148(7):3047–3057.

[76] Peyré, G. and Cuturi, M. (2019). Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6).

[77] Plischke, M. and Bergersen, B. (2006). *Equilibrium statistical physics*. World Scientific Publishing Company.

[78] Posner, E. (1975). Random coding strategies for minimum entropy. *IEEE Transactions on Information Theory*, 21(4):388–391.
[79] Pribram-Jones, A., Gross, D. A., and Burke, K. (2015). Dft: A theory full of holes? *Annual review of physical chemistry*, 66:283–304.
[80] Quinn, J. J. and Yi, K.-S. (2009). *Solid state physics: principles and modern applications*. Springer Science & Business Media.
[81] Räsänen, E., Pittalis, S., Capelle, K., and Proetto, C. (2009). Lower bounds on the exchange-correlation energy in reduced dimensions. *Physical review letters*, 102(20):206406.
[82] Reed, M. (2012). *Methods of modern mathematical physics: Functional analysis*. Elsevier.
[83] Ruelle, D. (1999). *Statistical mechanics: Rigorous results*. World Scientific.
[84] Saff, E. B. and Totik, V. (2013). *Logarithmic potentials with external fields*, volume 316. Springer Science & Business Media.
[85] Šamaj, L. and Percus, J. (1995). A functional relation among the pair correlations of the two-dimensional one-component plasma. *Journal of statistical physics*, 80(3-4):811–824.
[86] Sandier, E. and Serfaty, S. (2012). From the ginzburg-landau model to vortex lattice problems. *Communications in Mathematical Physics*, 313(3):635–743.
[87] Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63).
[88] Seidl, M. (1999). Strong-interaction limit of density-functional theory. *Physical Review A*, 60(6):4387.
[89] Seidl, M., Di Marino, S., Gerolin, A., Nenna, L., Giesbertz, K. J., and Gori-Giorgi, P. (2017). The strictly-correlated electron functional for spherically symmetric systems revisited. *arXiv preprint arXiv:1702.05022*.
[90] Seidl, M., Gori-Giorgi, P., and Savin, A. (2007). Strictly correlated electrons in density-functional theory: A general formulation with applications to spherical densities. *Physical Review A*, 75(4):042511.
[91] Seidl, M., Perdew, J. P., and Levy, M. (1999). Strictly correlated electrons in density-functional theory. *Physical Review A*, 59(1):51.
[92] Simon, B. (2011). *Convexity: an analytic viewpoint*, volume 187. Cambridge University Press.
[93] Singh, Y. (1991). Density-functional theory of freezing and properties of the ordered phase. *Physics Reports*, 207(6):351–444.
[94] Sinkhorn, R. (1964). A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics*, 35(2):876–879.
[95] Sinkhorn, R. and Knopp, P. (1967). Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348.
[96] Téllez, G. and Forrester, P. J. (1999). Exact finite-size study of the 2d ocp at $\gamma = 4$ and $\gamma = 6$. *Journal of Statistical Physics*, 97(3-4):489–521.
[97] Téllez, G. and Forrester, P. J. (2012). Expanded vandermonde powers and sum rules for the two-dimensional one-component plasma. *Journal of Statistical Physics*, 148(5):824–855.
[98] Villani, C. (2003). *Topics in optimal transportation*. Number 58. American Mathematical Soc.
[99] Villani, C. (2008). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
[100] Wigner, E. (1934). On the interaction of electrons in metals. *Physical Review*, 46(11):1002.
[101] Wigner, E. (1938). Effects of the electron interaction on the energy levels of electrons in metals. *Transactions of the Faraday Society*, 34:678–685.
[102] Wu, J. (2008). Density functional theory for liquid structure and thermodynamics.
[103] Yang, A. J., Fleming III, P. D., and Gibbs, J. H. (1976). Molecular theory of surface tension. *The Journal of Chemical Physics*, 64(9):3732–3747.