Recent progress on mean-field dynamics of Coulomb/Riesz gases

Matthew Rosenzweig

Based on joint works w/ Elias Hess-Childs (CMU), Jiaoyang Huang (UPenn), Sylvia Serfaty (NYU/Sorbonne)

September 16, 2025

“Workshop I: Embracing Stochasticity in Electrochemical Modeling,” ***Bridging the Gap: Transitioning from Deterministic to Stochastic Interaction Modeling in Electrochemistry*****, IPAM, Los Angeles, CA**

Introduction

The microscopic model I

The system of N particles of interest is described by an energy of the form

$$\mathcal{H}_N(X_N) := \frac{1}{N} \sum_{1 \le i \neq j \le N} \mathrm{g}(x_i, x_j) + \sum_{i=1}^{N} V(x_i), \qquad X_N := (x_1, \dots, x_N).$$

- The first term is the energy associated to the *interaction potential* g.
- The second term is the energy associated to the *confining potential* V (e.g. $V(x) = |x|^2$).
- The $1/N$ *mean-field* scaling of the interaction energy ensures both terms scale like N. Energy of each particle is the average of all the pairwise energies.
- No self-interaction in model
- Interaction is *long-range*

The microscopic model II

The general form of dynamics associated is the *second-order/kinetic* system of SDEs

$$\begin{cases} dx_i^t = v_i^t dt \\ m dv_i^t = -\gamma v_i^t dt + \mathbb{M}\nabla_{x_i}\mathcal{H}_N(X_N^t)dt + \sqrt{2/\beta}dW_i^t \end{cases}$$

with initial position velocities $(x_i^\circ, v_i^\circ) \in \mathbb{D}^\mathrm{d} \times \mathbb{R}^\mathrm{d}$.

- m is the *mass* of each particle, assumed to be identical.
- γ is the *friction coefficient*.
- $W_1, \ldots, W_N$ are independent d-dimensional Wiener processes. The parameter β (which may depend on N) has the interpretation of *inverse temperature*. The noise models *thermal fluctuations* in the system.
- Constant real $\mathrm{d} \times \mathrm{d}$ matrix $\mathbb{M}$ encodes type of dynamics.

When $\mathbb{M} = -\mathbb{I}$, this is the famous *Langevin system* in Newtonian mechanics.

The microscopic model III

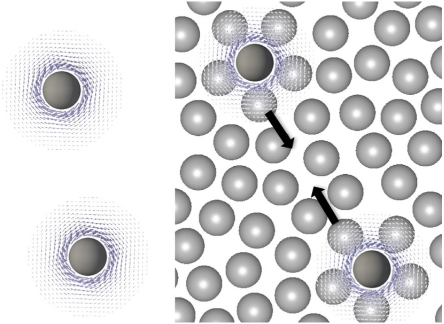

Figure: Courtesy of Alfredo Alexander-Katz, MIT

Particles moving in space according their velocities which are changing according to the pairwise forces exerted between the particles.

The microscopic model IV

In the *overdamped limit* when $m = 0$, v_i^t may be written in terms of X_N^t and W_i^t, and we obtain the *first-order* system

$$dx_i^t = \mathbb{M}\nabla\mathcal{H}_N(X_N^t)dt + \sqrt{2/\beta}dW_i^t.$$

We focus on this case because it is better understood.

In this context, the dynamics are called

- *Hamiltonian/conservative* if $\mathbb{M}$ is antisymmetric,
- *gradient/dissipative* if $\mathbb{M} = -\mathbb{I}$.

Relevance of model: Coulomb/Riesz gases I

Model case for potential $\mathrm{g}(x-y)=\mathrm{g}(x,y)$ are *log/Riesz interactions* indexed by parameter $\mathrm{s}<\mathrm{d}$:

$$\mathrm{g}(x)=\frac{1}{\mathrm{c}_{\mathrm{d},\mathrm{s}}}\begin{cases}-\log|x|, & \mathrm{s}=0\\ |x|^{-\mathrm{s}}, & \mathrm{s}\neq 0\end{cases}$$

- $\mathrm{s}=\mathrm{d}-2$ *Coulomb*
- $\mathrm{s}<\mathrm{d}-2$ *sub-Coulomb*
- $\mathrm{s}>\mathrm{d}-2$ *super-Coulomb*

Such systems are known as *Coulomb/Riesz gases.*

- Focus on the *singular* case $\mathrm{s}\geq 0$.
- Restrict to the *potential* case $\mathrm{s}<\mathrm{d}$. In the *hypersingular* case $\mathrm{s}\geq\mathrm{d}$, the potential is no longer locally integrable and the expected limiting behavior is completely different (local density vs. mean-field).
- Interactions are repulsive, meaning the energy diverges to $+\infty$ as the distance between particles tends to zero. Prevents collisions.

Relevance of model: Coulomb/Riesz gases II

Numerous applications & connections to particle systems in physics, particle methods for PDEs, finding equilibrium states for interaction energies, biological and sociological models, large neural networks, random matrix theory, approximation theory...

Refer to surveys Jabin 2014, Jabin-Wang 2017, Chaintron-Diez 2022, Golse 2022 and book Borodachov-Hardin-Saff 2019. Survey and lecture notes Lewin 2022, Serfaty 2024 specifically for the Coulomb/Riesz case.

Motivation I

Interested in the behavior of the microscopic system when is $N \gg 1$. Thinking of a gas in a finite volume, $N \sim 10^{23}$ (Avogadro's number).

In theory, one could solve the system of equations directly (if $\beta = \infty$, the dynamics are completely deterministic), and know the exact behavior of every particle.

In practice, this is too expensive, if at all feasible, given the size of N, the nonlinearity, and the coupling of the system.

Punchline: There is *too much* information in the N-particle system to do anything useful with it. Instead, one can adopt a statistical point of a view and ask what are the approximate dynamics of a single or fixed finite number of particles?

Questions of interest: mean-field limit I

What are the limiting dynamics of the empirical measure

$$\mu_N^t := \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i^t} \in \mathcal{P}(\mathbb{D}^{\mathsf{d}})$$

as $N \to \infty$?

For a test function φ and microscopic solution $X_N^t = (x_1^t, \ldots, x_N^t)$, computing

$$\frac{d}{dt}\int \varphi d\mu_N^t = \frac{1}{N}\sum_{i=1}^{N} \frac{d}{dt}\varphi(x_i^t),$$

leads us to *formally* expect that if $\mu_N^\circ \xrightarrow[N\to\infty]{} \mu^\circ$, then $\mu_N^t \xrightarrow[N\to\infty]{} \mu^t$, where μ^t is a solution to the *mean-field equation*

$$\begin{cases} \partial_t \mu - \mathrm{div}(\mu(\nabla V - \mathbb{M}\nabla \mathsf{g} * \mu)) = \frac{1}{\beta}\Delta\mu \\ \mu|_{t=0} = \mu^\circ. \end{cases}$$

When $\beta = \infty$, μ_N^t is, in fact, a weak solution to the mean-field equation.

Establishing the mean-field limit refers to proving this convergence.

Questions of interest: propagation of chaos I

Suppose the initial positions $X_N^\circ = (x_1^\circ, \dots, x_N^\circ)$ are μ°-iid:

$$f_N^\circ = (\mu^\circ)^{\otimes N}.$$

What is the limiting behavior as $N \to \infty$ of the law f_N^t of the positions $X_N^t = (x_1^t, \dots, x_N^t)$ of the particles at time t?

If μ^t is the solution of the mean-field equation with initial datum μ°, does it hold that

$$f_N^t \approx (\mu^t)^{\otimes N} \qquad \text{as } N \to \infty \text{ ?}$$

Propagation of chaos refers to the asymptotic factorization of k-point marginals $f_{N,k}^t \rightharpoonup (\mu^t)^{\otimes k}$.[1]

Known that mean-field convergence and propagation of chaos are closely related; qualitatively, they are equivalent Hauray-Mischler 2014.

[1] Recall that the k-point marginal $f_{N,k} := \int_{(\mathbb{R}^d)^{N-k}} f_N(\cdot, x_{k+1}, \dots, x_N) dx_{k+1} \cdots dx_N$.

Questions of interest: generation of chaos I

Related notion of **generation of chaos**.[2] Even when the initial law f_N° is not μ°-chaotic, one still has that f_N is μ^t-chaotic as $t \to \infty$ and $N \to \infty$.

$$f_{N,k}^t - (\mu^t)^{\otimes k} = o_N(1) + o_t(1),$$

where $o_N(1) \to 0$ as $N \to \infty$ uniformly in t and $o_t(1) \to 0$ as $t \to \infty$ and only depends on $f_{N,k}^\circ - (\mu^\circ)^{\otimes k}$.

Questions of interest: generation of chaos II

For (repulsive) overdamped Langevin dynamics ($\mathbb{M} = -\mathbb{I}$), one expects the law f_N^t weakly converges as $t \to \infty$ to the *canonical Gibbs ensemble*

$$d\mathbb{P}_{N,\beta}(X_N) = \frac{1}{Z_{N,\beta}} e^{-\beta \mathcal{H}_N(X_N)} dX_N.$$

with Hamiltonian $\mathcal{H}_N$.

The mean-field density μ^t should weakly converge to the *thermal equilibrium measure* μ_β, which is the minimizer among probability measures of the mean-field free energy

$$\mathcal{E}_\beta(\mu) := \frac{1}{2}\int_{(\mathbb{R}^d)^2} \mathsf{g}(x,y)\,d\mu^{\otimes 2}(x,y) + \int_{\mathbb{R}^d} V d\mu + \frac{1}{\beta}\int_{\mathbb{R}^d} \log \mu\, d\mu.$$

If V grows sufficiently fast at infinity, then $\mathcal{E}_\beta$ has a unique minimizer, which is characterized by the existence of a constant $c_\beta \in \mathbb{R}$ such that

$$\mathsf{g} * \mu_\beta + V + \frac{1}{\beta}\log \mu_\beta = c_\beta \quad \text{in } \mathbb{R}^d.$$

In all log/Riesz cases, it's known that $\mathbb{P}_{N,\beta}^{(k)} \rightharpoonup \mu_\beta^{\otimes k}$.

Questions of interest: generation of chaos III

Figure: Large Particle Number and Large Time Limits

$$
\begin{tikzcd}
N\text{-particle law } f_N^t \arrow[r, "N \to \infty"] \arrow[d, "t \to \infty"'] \arrow[dr, "{N \to \infty,\ t \to \infty}"] & \text{mean-field law } \mu^t \arrow[d, "t \to \infty"] \\
N\text{-particle Gibbs measure } \mathbb{P}_{N,\beta}^V \arrow[r, "N \to \infty"] & \text{thermal equilibrium measure } \mu_\beta
\end{tikzcd}
$$

---

[2] This term was recently coined by Jani Lukkarinen.

Questions of interest: convergence of fluctuations

The mean-field limit is a *LLN* result; it describes the *leading-order* behavior of the empirical measure.

The next-order behavior is captured by the (unnormalized) *fluctuation field*

$$\mathrm{Fluc}_N^t := \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i^t} - \mu^t,$$

which is a random distribution (in the sense of Schwartz).

What is the size σ_N^{-1} of Fluc_N^t?

Does the normalized fluctuation $\sigma_N \mathrm{Fluc}_N^t$ converge to a solution of some (S)PDE? Can we characterize that solution? Is it a Gaussian process?

Affirmatively answering the last question amounts to proving a *CLT*. Even more interesting if a non-CLT holds, i.e. fluctuations not Gaussian.

Previous results: mean-field convergence/propagation of chaos

- Coupling method Sznitman 1991, Hauray-Jabin 2015, Boers-Pickl 2016, Lazarovici-Pickl 2017, Graß 2021, Guillin-Le Bris-Monmarché 2021,...
- Wasserstein stability Braun-Hepp 1977, Dobrushin 1979, Neunzert-Wick 1974, Hauray 2009, Carrillo-Choi-Hauray 2014,...
- Relative entropy Jabin-Wang 2016, 2018, Wynter 2021, Guillin-Le bris-Monmarché 2021, Feng-Wang 2023-2024, Cai-Feng-Gong-Wang 2024
- Control of microscopic dynamics and compactness for well-chosen point configurations Goodman-Hou-Lowengrub 1990, Schochet 1996,...
- Displacement convexity for Wasserstein gradient flow Carrillo-Ferreira-Precioso 2012, Berman-Onnheim 2015,...
- Compactness via diffusion Osada 1985-1987, Rogers-Shi 1993, Cépa-Lepingle 1997, Fournier-Hauray-Mischler 2014, Wang-Zhao-Zhu 2022,...
- Stability for BBGKY hierarchy Lacker 2021, Han 2022, Jabin-Poyato-Soler 2021, Bresch-Jabin-Soler 2022, Lacker-Le Flem 2023, Hess Childs-Rowan 2023, S. Wang 2024
- Cumulant hierachies Bresch-Duerinckx-Jabin 2024
- **Modulated energy/free energy method** Duerinckx 2016, Serfaty 2020, R. 2020-2021, Nguyen-R.-Serfaty 2021, R.-Serfaty 2021, Ménard 2023, BenPorat-Carrillo-Jabin 2025 / Bresch-Jabin-Wang 2019-2020, Chodron de Courcel-R.-Serfaty 2023, R.-Serfaty 2023-2024

Modulated energy

Modulated-energy method I

$$\mathsf{F}_N(X_N,\mu) := \int_{(\mathbb{R}^{\mathsf{d}})^2\setminus\triangle} \mathsf{g}(x-y)d\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}-\mu\Big)^{\otimes 2}(x,y).$$

Total interaction of system of N discrete charges at x_i against neutralizing background of charge μ, with self-interaction of points (infinite, if $\mathsf{g}(x,x)=\infty$) removed.

Quantity first appeared in stat mech of Coulomb/Riesz gases Sandier-Serfaty 2015, Rougerie-Serfaty 2016, Petrache-Serfaty 2017, Leblé-Serfaty 2017-2018 as a *next-order energy*.

Analogous to a Riesz *jellium* energy, but in the physics literature, this is understood as an infinite configuration of points screened by a uniform background (i.e., $\mu\equiv 1$), so we avoid this terminology

Modulated-energy method II

This modulated energy first used in dynamics context Duerinckx 2016, Serfaty 2020

Idea: establish a Grönwall relation for $\mathsf{F}_N(X_N^t, \mu^t)$

- Method goes back to Brenier 2000; similarities with *relative-entropy method* Dafermos 1979, DiPerna 1979, Yau 1991, Saint-Raymond 2009
- Exploits a weak-strong uniqueness principle for limiting equation
- Advantages - quantitative; no need for study of microscopic dynamics
- Disadvantages - typically requires some regularity for or an *a priori* assumption on the limiting solution

Coercivity of modulated energy I

For simplicity, assume $-\Delta \mathrm{g} = \mathrm{c_d}\delta_0$ (Coulomb). Then letting $h = \mathrm{g} * f$,

$$\int_{(\mathbb{R}^{\mathrm{d}})^2} \mathrm{g}(x-y) df^{\otimes 2}(x,y) = \frac{1}{\mathrm{c_d}} \int_{\mathbb{R}^{\mathrm{d}}} |\nabla h|^2 = \mathrm{c_d} \|f\|_{\dot{H}^{-1}}^2 .$$

Unfortunately, infinite for $f = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i} - \mu$ if $\mathrm{d} \geq 2$!

Coercivity of modulated energy II

Consider the difference of electric potentials

$$h_{N,\vec{\eta}} := \mathsf{g} * \Big(\frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}^{(\eta_i)} - \mu\Big), \qquad \vec{\eta} = (\eta_1, \ldots, \eta_N) \in \mathbb{R}_+^N,$$

where $\delta_{x_i}^{(\eta_i)}$ are the smeared electric charges. If

$$\mathsf{r}_i := \frac{1}{4}\min\Big(\min_{j\neq i}|x_i - x_j|, \lambda\Big),$$

where $\lambda := (N\|\mu\|_{L^\infty})^{-1/\mathsf{d}}$ is the *microscale*, then for any $\eta_i \leq \mathsf{r}_i$,

$$\mathsf{F}_N(X_N, \mu) = \frac{1}{2\mathsf{c_d}}\Bigg(\int_{\mathbb{R}^\mathsf{d}} |\nabla h_{N,\vec{\eta}}|^2 - \frac{\mathsf{c_d}}{N^2}\sum_{i=1}^{N}\mathsf{g}(\eta_i)\Bigg) - \frac{1}{N}\sum_{i=1}^{N}\int_{\mathbb{R}^\mathsf{d}} \mathsf{f}_{\eta_i}(x - x_i)d\mu(x)$$

where $\mathsf{f}_{\eta_i} := \mathsf{g} - \mathsf{g}_{\eta_i}$ for $\mathsf{g}_{\eta_i} := \min(\mathsf{g}, \mathsf{g}(\eta_i))$.

This is the so-called *electric reformulation*, which makes it clear that the modulated energy is a *renormalized* energy.

With more difficulty, this idea of renormalizing the energy through smearing the Dirac masses can be generalized to the full range $0 \leq \mathsf{s} < \mathsf{d}$.

Modulated energy and commutators

The modulated energy identity

Suppose that $\beta = \infty$. One computes

$$\frac{d}{dt}\mathsf{F}_N(X_N^t, \mu^t)$$
$$\leq -\int_{(\mathbb{R}^\mathsf{d})^2\setminus\triangle} \left(u^t(x) - u^t(y)\right) \cdot \nabla \mathsf{g}(x-y) d\left(\frac{1}{N}\sum_{i=1}^N \delta_{x_i^t} - \mu^t\right)^{\otimes 2}(x,y),$$

where $u^t := \nabla V - \mathbb{M}\nabla\mathsf{g} * \mu^t$ is the vector field associated to the mean-field equation solution μ^t.

If we have a functional inequality of the form

$$|\mathrm{RHS}| \leq C_1(\|u^t\|)\Big(|\mathsf{F}_N(X_N^t, \mu^t)| + C_2(\|\mu^t\|)N^{-\alpha}\Big),$$

where C_1, C_2 depends on d, some norm of u^t, μ^t, respectively, and $\alpha > 0$, then Grönwall implies an estimate

$$|\mathsf{F}_N(X_N^t, \mu^t)| \leq \left(|\mathsf{F}_N(X_N^0, \mu^0)| + N^{-\alpha}\int_0^t C_2(\|\mu^\tau\|)d\tau\right)e^{\int_0^t C_1(\|u^\tau\|)d\tau}.$$

Variation by transport I

In the context of mean-field limits, essential to control quantities that correspond to differentiating F_N along a transport field:

$$\frac{d^n}{dt^n|_{t=0}}\mathsf{F}_N\Big((\mathbb{I}+tv)^{\otimes N}(X_N),(\mathbb{I}+tv)\#\mu\Big)$$
$$=\int_{(\mathbb{R}^\mathsf{d})^2\setminus\triangle}\nabla^{\otimes n}\mathsf{g}(x-y):(v(x)-v(y))^{\otimes n}d\Big(\frac{1}{N}\sum_{i=1}^N\delta_{x_i}-\mu\Big)^{\otimes 2}(x,y),$$

where $(\mathbb{I}+tv)(x)=x+tv(x)$.

Harmonic Analysis Problem: Prove functional inequalities of form

$$|\mathrm{RHS}|\leq C(\|v\|)\big(\mathsf{F}_N(X_N,\mu)+C(\|\mu\|)N^{-\alpha}\big)$$

for some $\alpha>0$. Because then this yields an estimate for the time-evolved modulated energy $\mathsf{F}_N(X_N^t,\mu^t)$ that implies mean-field convergence.

Two perspectives on proving these functional inequalities...

The stress-energy tensor perspective I

The first perspective is due Leblé-Serfaty 2018, Serfaty 2020, Serfaty 2023.

Idea: interpret the variation expression in terms of a stress-energy tensor structure. In the Coulomb case, formally letting h_N solve $-\Delta h_N = \mathsf{c}_\mathsf{d}\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i} - \mu\Big)$, one can use integration by parts to write

$$\int_{(\mathbb{R}^\mathsf{d})^2\setminus\triangle} \Big(v(x)-v(y)\Big)\cdot\nabla \mathsf{g}(x-y) d\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i} - \mu\Big)^{\otimes 2}(x,y) \\ = \frac{1}{\mathsf{c}_{\mathsf{d},\mathsf{s}}}\int_{\mathbb{R}^\mathsf{d}} v\cdot \mathrm{div}\,[h_N,h_N]\,dx,$$

where the *stress-energy tensor* is defined by

$$[\nabla h_1,\nabla h_2]^{ij} := \partial_i h_1\partial_j h_2 + \partial_i h_1 \partial_j h_2 - \delta^{ij}(\nabla h_1\cdot\nabla h_2), \qquad 1\le i,j\le \mathsf{d}.$$

Integration by parts and Cauchy-Schwarz allow one to conclude the bound

$$\left|\int_{\mathbb{R}^\mathsf{d}} v\cdot\mathrm{div}\,[\nabla h_N,\nabla h_N]\,dx\right| \le C\|\nabla v\|_{L^\infty}\|\nabla h_N\|_{L^2}^2.$$

The stress-energy tensor perspective II

Formally, $\|\nabla h_N\|_{L^2}^2$ is the Coulomb energy of $\frac{1}{N}\sum_{i=1}^N \delta_{x_i} - \mu$; but $\|\nabla h_N\|_{L^2} = \infty$ due to the singularity of the Dirac masses.

However, the reasoning may be implemented after a <u>renormalization</u>: namely, replace δ_{x_i} with the smeared charge $\delta_{x_i}^{(\eta_i)}$ above, and apply reasoning to

$$h_{N,\vec{\eta}} := \mathrm{g} * \left(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}^{(\eta_i)} - \mu\right), \qquad \vec{\eta} = (\eta_1, \ldots, \eta_N) \in \mathbb{R}_+^N.$$

Estimate the error directly from this replacement after choosing the η_i to be a nearest-neighbor type distance of the order $N^{-\frac{1}{\mathrm{d}}}$.

The stress-tensor approach is elegant, low-tech (i.e., integration by parts), and based on local arguments.

But rigid in the sense that it seems restricted to exact Riesz potentials.

The commutator perspective I

The second perspective originates in R. 2020 (generalized in Q.H. Nguyen-R.-Serfaty 2021) on developing a new generalization of the modulated-energy method for multiplicative noise.

If $f = \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i} - \mu$, then *formally* the first variation may be rewritten

$$\int_{\mathbb{R}^{\mathsf{d}}} \Big(u\cdot(\nabla \mathsf{g} * f) - \mathsf{g} * (\mathrm{div}(uf))\Big) df(x) = \left\langle f, \left[u^i, \frac{\partial_i}{(-\Delta)^{\frac{\mathsf{d}-\mathsf{s}}{2}}}\right] f\right\rangle_{L^2}.$$

The first variation is the quadratic form associated to the <u>commutator</u> $\left[u^i, \frac{\partial_i}{(-\Delta)^{\frac{\mathsf{d}-\mathsf{s}}{2}}}\right]$. The higher-order variations can be similarly formulated in terms of iterated commutators.

After renormalization, one can then apply estimates for a class of singular integral operators known as *Calderón* d*-commutators* Calderón 1980, Christ-Journé 1987, Seeger et al. 2019, Lai 2020. The error introduced by renormalization can be estimated directly.

Sharpness of functional inequality

A natural question is the *size of the exponent* α in the functional inequality.

$$\left|\int_{(\mathbb{R}^{\mathrm{d}})^2\setminus\triangle}\Big(v(x)-v(y)\Big)\cdot\nabla \mathrm{g}(x-y)d\Big(\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}-\mu\Big)^{\otimes 2}(x,y)\right|$$
$$\leq C_1(\|\nabla v\|)\Big(\mathrm{F}_N(X_N,\mu)+C_2(\|\mu\|)N^{-\alpha}\Big).$$

More precisely, we want both terms in the RHS to be of the same order as $N\to\infty$; and we say the functional inequality is *sharp* if this is the case.

- By only counting nearest-neighbor (with typical distance of $N^{-\frac{1}{\mathrm{d}}}$) interactions, one expects F_N is at least of order $N^{\frac{\mathrm{s}}{\mathrm{d}}-1}$
- $\mathrm{F}_N \geq -CN^{\frac{\mathrm{s}}{\mathrm{d}}-1}$, where $C = C(\|\mu\|_{L^\infty}) > 0$
- Known that $\min \mathrm{F}_N$ is of order $N^{\frac{\mathrm{s}}{\mathrm{d}}-1}$ Sandier-Serfaty 2015, Rougerie-Serfaty 2015, Petrache-Serfaty 2017, Hardin-Saff-Simanek-Su 2017
- Optimal exponent $\alpha = 1-\frac{\mathrm{s}}{\mathrm{d}}$ only been shown for Coulomb case $\mathrm{s}=\mathrm{d}-2$ Leblé-Serfaty 2018, Serfaty 2020, R. 2021

Localizability

Another important aspect is the *localizability* of such estimates to take into account only the contribution to the energy of particles from a region $\Omega \subset \mathbb{D}^{\mathsf{d}}$.

Introduce the *localized (Coulomb) modulated energy*

$$\mathsf{F}_N^{\Omega}(X_N,\mu) := \frac{1}{2\mathsf{c_d}}\left(\int_{\Omega} |\nabla h_{N,\tilde{\mathsf{r}}}|^2 - \frac{\mathsf{c_d}}{N^2}\sum_{i\in I_\Omega} \mathsf{g}(\tilde{\mathsf{r}}_i)\right) \\ - \frac{1}{N}\sum_{i\in I_\Omega}\int_{\mathbb{R}^{\mathsf{d}}} \mathsf{f}_{\tilde{\mathsf{r}}_i}(x-x_i)d\mu(x),$$

where $I_\Omega := \{1 \le i \le N : x_i \in \Omega\}$. Here, $\tilde{\mathsf{r}}_i$ is a slight modification of the nearest-neighbor distance r_i to not smear points too close to $\partial\Omega$.

We think of Ω as being comparable to $\operatorname{supp} v$ for the vector field appearing in the transport estimates. Want to show

$$\Big|\int_{(\mathbb{R}^{\mathsf{d}})^2\setminus\triangle} (v(x)-v(y))\cdot\nabla\mathsf{g}(x-y)d\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}-\mu\Big)^{\otimes 2}(x,y)\Big| \\ \le C_v(\mathsf{F}_N^{\Omega}(X_N,\mu) + C\# I_\Omega N^{\frac{\mathsf{s}}{\mathsf{d}}-2}).$$

Sharp estimates in all Riesz cases I

Over the last few years, we have proven the sharp commutator estimate in all log/Riesz cases.

- (super-)Coulomb R.-Serfaty 2022
- (sub-)Coulomb HessChilds-R.-Serfaty 2024

Consequently, we have following applications:

- Optimal $N^{\frac{s}{d}-1}$ rate of convergence for the mean-field limit measured in terms of $\mathsf{F}_N(X_N^t, \mu^t)$ for all Riesz cases. Previously, the optimal rate was only known for the Coulomb case $\mathsf{s} = \mathsf{d} - 2$ Leblé-Serfaty 2018, Serfaty 2021, R. 2021.
- Optimal scaling for supercritical mean-field limits (Lake/anelastic equation) R.-Serfaty 2024.

Sharp estimates in all Riesz cases II

Proofs use both aforementioned perspectives, but are rather different. In both cases, prove estimates for regular test functions (i.e. not $\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}-\mu$) and then separately handle the renormalization.

(super-)Coulomb

- View commutators as solutions of elliptic PDE and develop an L^2-L^∞ regularity theory.
- Use Caffarelli-Silvestre extension to handle the nonlocality when $\mathrm{s} \neq \mathrm{d}-2$.

(sub-)Coulomb

- Introduce a new potential truncation based on a wavelet-type representation of the Riesz potential as an average over scales of approximate δ-functions.
- Apply Kato-Ponce type estimates which we then average over.

Sharp estimates in all Riesz cases III

(super-)Coulomb estimate is essentially as good as it gets.

- Sharp estimates for the higher variations.
- Localizable
- Depends on v only through $\|\nabla v\|_{L^\infty}$; cannot be meaningfully weakened (counterexample to $\|\nabla v\|_{BMO}$).

(sub-)Coulomb estimate has room for improvement.

- Only for first variation
- Not localizable
- Has an additional regularity dependence on v that seems suboptimal (cf. Q.H. Nguyen-R.-Serfaty 2021)

Modulated (free) energy at positive temperature

Modulated energy at positive temperature I

R.-Serfaty 2021 developed a stochastic version of the modulated-energy method based on

$$\mathbb{E}\Big(|\mathsf{F}_N(X_N^t, \mu^t)|\Big).$$

Works for either conservative or repulsive dissipative dynamics if $\mathsf{s} < \mathsf{d} - 2$ when $\beta_N \gtrsim 1$. Purely pathwise; requires no randomness of the initial data.

Exploiting the convergence of the mean-field density to equilibrium, can even obtain estimates which are *uniform-in-time*!

Previous works on uniform-in-time convergence (e.g., Malrieu 2003, Cattiaux-Guillin-Malrieu 2008, Salem 2018, Durmus-Eberle-Guillin-Zimmer 2020, Arnaudon-Del Moral 2020, Delarue-Tse 2021, Delgadino-Gvalani-Pavliotis-Smith 2023,...) impose strong convexity and/or regularity assumptions on the interaction potential g.

Modulated energy at positive temperature II

Method breaks down for $\mathsf{s} \geq \mathsf{d} - 2$, essentially because $\Delta \mathsf{g}$ is no longer a nonpositive locally integrable function.

Bad term

$$\frac{1}{\beta_N} \int_{(\mathbb{D}^{\mathsf{d}})^2 \setminus \triangle} \Delta \mathsf{g}(x-y) d\Big(\frac{1}{N} \sum_{i=1}^{N} \delta_{x_i^t} - \mu^t\Big),$$

which, through the coercivity property of modulated energy, is nonpositive up to vanishing error if $\mathsf{s} < \mathsf{d} - 2$.

Also, doesn't work for $\mathsf{s} = \mathsf{d} - 2$ (cf. Guillin-Le bris-Monmarché 2021 for $\mathsf{d} = 2$, conservative).

What's missing is an entropic term whose time derivative kills off this bad term.

This leads us to...

Modulated free energy I

In contrast, the *modulated free energy* of Bresch-Jabin-Wang 2019-2020,

$$E_N(f_N, \mu) := \frac{1}{\beta} H_N(f_N|\mu^{\otimes N}) + \mathbb{E}_{f_N}\left[\mathsf{F}_N(X_N, \mu)\right],$$

is well-suited to studying overdamped Langevin dynamics at positive temperature for all $\mathsf{s} < \mathsf{d}$.

Combines

- (normalized) relative entropy/KL divergence

$$H_N(f_N|\mu^{\otimes N}) := \frac{1}{N} \int_{(\mathbb{D}^{\mathsf{d}})^N} \log \frac{f_N}{\mu^{\otimes N}} d\mu^{\otimes N}$$

- (average) modulated energy $\mathbb{E}_{f_N}\left[\mathsf{F}_N(X_N, \mu)\right]$.

Now the points $X_N \sim f_N \in \mathcal{P}((\mathbb{D}^{\mathsf{d}})^N)$, and we take the expectation of the modulated energy $\mathsf{F}_N(X_N, \mu)$.

Enlightening to re-express the modulated free energy in a different form...

Modulated free energy II

Given a probability density μ, we can define the *modulated Gibbs measure* and *modulated partition function*

$$d\mathbb{Q}_{N,\beta}(\mu) := \frac{1}{K_{N,\beta}(\mu)} e^{-\beta N F_N(X_N,\mu)} d\mu^{\otimes N}(X_N),$$
$$K_{N,\beta}(\mu) := \int_{(\mathbb{D}^{\mathsf{d}})^N} e^{-\beta N F_N(X_N,\mu)} d\mu^{\otimes N}(X_N).$$

Using the explicit form of $\mathbb{Q}_{N,\beta}(\mu)$, we may rewrite

$$E_N(f_N,\mu) = \frac{1}{\beta}\left(H_N(f_N|\mathbb{Q}_{N,\beta}(\mu)) - \frac{\log K_{N,\beta}(\mu)}{N}\right).$$

Up to an error determined by the modulated partition function, the modulated free energy is another relative entropy!

Dissipation of modulated free energy I

Crucial computation of Bresch et al. (reformulated by R.-Serfaty 2023),

$$\frac{d}{dt}E_N(f_N^t,\mu^t) \le -\frac{1}{\beta^2} I_N(f_N^t|\mathbb{Q}_{N,\beta}(\mu^t))$$
$$-\frac{1}{2}\mathbb{E}_{f_N^t}\left[\int_{(\mathbb{D}^{\mathsf{d}})^2\setminus\triangle}(u^t(x)-u^t(y))\cdot\nabla \mathsf{g}(x-y)d\left(\frac{1}{N}\sum_{i=1}^N\delta_{x_i}-\mu^t\right)^{\otimes 2}(x,y)\right],$$

where the velocity field associated to the mean-field dynamics is

$$u^t := \frac{1}{\beta}\nabla\log\mu^t + \nabla V + \nabla\mathsf{g}*\mu^t = \frac{1}{\beta}\nabla\log\frac{\mu^t}{\mu_\beta} + \nabla\mathsf{g}*(\mu^t-\mu_\beta).$$

- The first term on the RHS is the modulated Fisher information, which is ≤ 0, and may be discarded. In some cases ($\mathsf{d}=1, \mathsf{s}\ge\mathsf{d}-2$), this term may be exploited through a *uniform LSI* for $\mathbb{Q}_{N,\beta}(\mu^t)$ to establish generation of chaos R.-Serfaty 2023.
- The second term on the RHS can be controlled by the modulated energy itself through the commutator inequality, allowing to close a Grönwall loop. This is what is done in Bresch-Jabin-Wang 2019 and uniformly in time in Chodron de Courcel-R.-Serfaty 2023, R.-Serfaty 2024.

Summary of mean-field limit

Modulated (free) energy and commutator estimates have allowed to prove the mean-field limit/propagation of chaos for all log/Riesz cases with fixed or vanishing temperature for both conservative and dissipative dynamics.

Sharp error $N^{\frac{s}{d}-1}$ has been obtained in all log/Riesz cases.

Can even obtain rates of convergence that are uniform in time in cases where there is a dissipative/diffusive effect either through the noise or through the interaction.

There are extensions even to mildly singular (e.g. $\log$) attractive interactions Bresch-Jabin-Wang 2020, Chodron de Courcel-R.-Serfaty 2023.

Fluctuations around the mean-field limit

Goal

Let X_N^t be a solution of N-particle system

$$dx_i = \Big(\frac{1}{N}\sum_{1\le j\le N: j\ne i} \mathbb{M}\nabla \mathsf{g}(x_i^t - x_j^t) - \nabla V(x_i^t)\Big)dt + \sqrt{2/\beta_N}\, dW_i^t$$

and let $\mu_{\beta_N}^t$ be the solution of the mean-field PDE

$$\partial_t \mu_{\beta_N} = \operatorname{div}\Big(\mu_{\beta_N}(\nabla V - \mathbb{M}\nabla \mathsf{g} * \mu_{\beta_N})\Big) + \frac{1}{\beta_N}\Delta \mu_{\beta_N}.$$

Find the appropriate normalization σ_N such that the normalized fluctuation

$$\sigma_N \mathrm{Fluc}_N^t = \sigma_N\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i^t} - \mu_{\beta_N}^t\Big)$$

converges in a suitable sense to a limiting distribution-valued process η as $N \to \infty$, assuming initial convergence of $\sigma_N \mathrm{Fluc}_N^0$ to some suitably regular distribution η^0.

Will see that the limiting process η determined by an SPDE corresponding to the linearization of the mean-field equation around μ_{β_∞} driven by a spacetime white noise.

Previous results

Classical work Tanaka-Hitsuda 1981, Tanaka 1984, Sznitman 1985, Fernandez-Méléard 1997, Lancelotti 2009 — nonsingular g.

Bender 2008, Unteberger 2018, J. Huang-Landon 2019 — Dyson Brownian Motion (DBM), i.e. $d = 1$, $s = 0$, $\beta_N = N$, $\sigma_N = N$

Wang-Zhao-Zhu 2021 — $d \geq 2$, $s = 0$, $\beta_N = O(N)$ conservative case (i.e. $\mathbb{M}$ antisymmetric), $\sigma_N = \sqrt{N}$.

We are unaware of any work treating the fluctuations of Riesz flows for $s > 0$ or even the gradient flow for the $s = 0$ case for $d \geq 2$.[3]

There is work covering (for varying scalings β_N) the fluctuations of linear statistics of the Gibbs measure $\mathbb{P}_{N,\beta_N}$ in the log/Riesz case for $d \geq 2$.

- Coulomb $d = 2$ Bodineau-Guionnet 1999/Grotto-Romito 2020, Leblé-Serfaty 2018/Bauerschmidt-Bourgade-Nikula-Yau 2019
- Coulomb $d = 2$, $d = 3$ (conditional result) Serfaty 2020
- Riesz $d = 1$ Boursier 2021

[3] Very recent work of Cecchi-Nikolaev 2025 treats the Coulomb case for $d = 2, 3$ when $\beta_N \propto 1$.

Evolution of linear statistics I

Given test function φ°, want to compute the linear statistic $\langle \varphi^\circ, \mathsf{Fluc}_N^t \rangle$.

Idea: use essentially method of characteristics to study the initial fluctuation Fluc_N° tested against an *evolved* test function.

Consider the solution $f^{t',t}$, the solution of the backwards Cauchy problem

$$\begin{cases} -\partial_{t'} f^{t',t} = \mathcal{L}_{\beta_N, \mu_{\beta_N}^{t'}} f^{t',t} \\ f^{t,t} = \varphi^\circ. \end{cases}$$

where

$$\mathcal{L}_{\beta,\nu}\phi := \beta\Delta\phi + \nabla V \cdot \nabla\phi + \nabla\phi \cdot \mathbb{M}\nabla \mathsf{g} * \nu - \mathrm{div}\mathsf{g} * (\nu\mathbb{M}^*\nabla\phi),$$

which corresponds to the L^2 adjoint of the linearization of the RHS of the mean-field PDE around the state ν. We call this equation the *adjoint PDE*.

Evolution of linear statistics II

This choice removes the linear evolution of Fluc_N^t. By Itô's formula, integration by parts, and some algebraic manipulation,

$$\langle \varphi^\circ, \sigma_N \mathrm{Fluc}_N^t \rangle = \langle f^{0,t}, \sigma_N \mathrm{Fluc}_N^\circ \rangle + \sqrt{\frac{2\sigma_N^2}{N^2 \beta_N}} \sum_{i=1}^N \int_0^t \nabla f^{t',t}(x_i^{t'}) \cdot dW_i^{t'}$$

$$+ \frac{\sigma_N}{2} \int_0^t \int_{(\mathbb{D}^{\mathrm{d}})^2 \setminus \triangle} (\nabla f^{t',t}(x) - \nabla f^{t',t}(y)) \cdot \mathbb{M} \nabla \mathrm{g}(x-y) d(\mathrm{Fluc}_N^{t'})^{\otimes 2}(x,y) dt'.$$

Let's understand the convergence of each of the three terms on the RHS.

Convergence of linear statistics I

Suppose for some sequence κ_N, we have $\kappa_N \mathrm{Fluc}_N^\circ \xrightharpoonup{N\to\infty} \eta^\circ$, where η° is a $H^{-\gamma}$-valued random distribution.

Assuming $f^{0,t}$ is suitably regular,

$$\langle f^{0,t}, \kappa_N \mathrm{Fluc}_N^\circ \rangle \xrightharpoonup[N\to\infty]{} \langle f^{0,t}, \eta^\circ \rangle.$$

Convergence of linear statistics II

Quadratic variation of process $\sqrt{\frac{2\sigma_N^2}{N^2\beta_N}}\sum_{i=1}^N\int_0^{(\cdot)}\nabla f^{t',t}(x_i^{t'})\cdot dW_i^{t'}$ given by

$$\frac{2\sigma_N^2}{N^2\beta_N}\sum_{i=1}^N\int_0^{(\cdot)}|\nabla f^{t',t}(x_i^{t'})|^2dt'$$

$$=\frac{2\sigma_N^2}{N\beta_N}\int_0^{(\cdot)}\langle|\nabla f^{t',t}|^2,\mu_{\beta_N}^{t'}\rangle dt'+\frac{2\sigma_N^2}{N\beta_N}\int_0^{(\cdot)}\langle|\nabla f^{t',t}|^2,\mathrm{Fluc}_N^s\rangle dt'$$

- If $\lim_N\frac{\kappa_N^2}{N\beta_N}=:\frac{1}{\beta_*}\in[0,\infty)$, then choose $\sigma_N=\kappa_N$.
- If $\lim_N\frac{\kappa_N^2}{N\beta_N}=\infty$, then choose $\sigma_N=\sqrt{N\beta_N}$.

In all cases $\sigma_N:=\min(\sqrt{N\beta_N},\kappa_N)$, and by the martingale CLT,

$$\sqrt{\frac{2\sigma_N^2}{N^2\beta_N}}\sum_{i=1}^N\int_0^\tau\nabla f^{t',t}(x_i^{t'})\cdot dW_i^{t'}\xrightarrow[N\to\infty]{}\sqrt{\frac{2}{\beta_*}}\int_0^\tau\langle\nabla f^{t',t},\sqrt{\mu_{\beta_\infty}^{t'}}\xi_{t'}\rangle dt',$$

where ξ is a $\mathbb{R}^d$-valued spacetime white noise.

Convergence of linear statistics III

Finally, by our commutator functional inequality,

$$\frac{\sigma_N}{2}\int_0^t \left| \int_{(\mathbb{D}^{\mathsf{d}})^2\setminus\triangle} (\nabla f^{t',t}(x) - \nabla f^{t',t}(y)) \cdot \mathbb{M}\nabla \mathsf{g}(x-y) d(\mathrm{Fluc}_N^{t'})^{\otimes 2}(x,y) \right| dt' \\ \leq C\sigma_N \int_0^t \|\nabla f^{t',t}\|_* \left(\mathsf{F}_N(X_N^{t'}, \mu_{\beta_N}^{t'}) + \mathsf{o}_N^{t'}\right) dt',$$

where

$$\mathsf{o}_N^s = \begin{cases} O(N^{\frac{\mathsf{s}}{\mathsf{d}}-1}), & \mathsf{s} \neq 0 \\ O(\frac{\log N}{N}), & \mathsf{s} = 0. \end{cases}$$

View this term as an error, which requires that

$$\lim_{N\to\infty} \sigma_N \int_0^t \mathbb{E}\left[\mathsf{F}_N(X_N^{t'}, \mu_{\beta_N}^{t'}) + \mathsf{o}_N^{t'}\right] dt' = 0.$$

Convergence of linear statistics IV

Putting everything together, we have shown J. Huang-R.-Serfaty 2022, under the assumptions imposed above, that for each $t \geq 0$,

$$\langle \phi^\circ, \sigma_N \mathrm{Fluc}_N^t \rangle \xrightarrow[N\to\infty]{} \langle f^{0,t}, \eta^\circ \rangle (1 - \mathbf{1}_{\lim_{N\to\infty} \frac{\kappa_N}{N\beta_N} = \infty}) + \sqrt{\frac{2}{\beta_*}} \int_0^t \langle \nabla f^{t',t}, \sqrt{\mu_{\beta_\infty}^{t'}} \xi_{t'} \rangle dt'.$$

Recalling the definition of $f^{0,t}$, the RHS may be rewritten as

$$\langle \phi^\circ, \eta^t \rangle,$$

where η^t is the unique martingale solution of the SPDE

$$\begin{cases} \partial_t \eta = \mathcal{L}^*_{\beta_\infty, \mu_{\beta_\infty}} \eta - \sqrt{\frac{2}{\beta_*}} \mathrm{div}(\sqrt{\mu_{\beta_\infty}} \xi) \\ \eta|_{t=0} = \eta^\circ (1 - \mathbf{1}_{\lim_{N\to\infty} \frac{\kappa_N}{\beta_N N} = \infty}). \end{cases}$$

To close argument, need solutions $f^{t',t}$ of the adjoint PDE which have L^∞ Hessian, possibly more. Also want regularity estimates which are uniform in the temperature. Non-trivial PDE problem, particularly in $\mathbb{R}^d$...

Interpretation of the result

Result says that if

- the size of the initial fluctuations is large enough depending on $N^{\frac{s}{d}-1}$,
- the temperature is low enough depending on σ_N^2/N

then the size of the time-evolved fluctuations remains the same order.

Moreover,

- at low temperature, the initial fluctuation is merely transported by the flow of the linearization about the mean-field density, i.e. the noise contribution vanishes in the limit;
- at the "critical" scaling, the linear transport picks up an additional source term, i.e. the noise persists in the limit.

At high enough temperature, the noise creates much larger fluctuations than that of the initial data. The contribution of the initial data vanishes in the limit.

If the initial fluctuations are very small, one cannot normalize so that the commutator is negligible while the initial data persists in the limit.

Comments on theorem

- Regularity assumptions for mean-field solution μ_{β_N} are always satisfied, thanks to many contributions in the PDE literature.
- In fact, proof yields a rate of convergence, e.g. in Wasserstein distance.
- Can also show convergence of joint statistics.
- Our theorem includes the purely deterministic case. I.e. $\mathrm{Law}(X_N^\circ) = \delta_{Y_N}$, for some $Y_N \in \mathbb{D}^N$, (no randomness of initial data) and $\beta = \infty$ (no randomness in dynamics). This case seems to have been missed in previous work.
- With more work, it is possible to show, for regimes of s, σ_N, β_N, that

$$\sigma_N \mathrm{Fluc}_N \xrightarrow[N\to\infty]{Law} \eta \quad \text{in} \quad \mathcal{P}(C([0,T], H^{-\gamma})),$$

  for γ sufficiently large. In some cases depending on s, β, and $\mathbb{D}$, this convergence holds over on $[0,\infty)$.
- Using the special properties of the 1D $\log$, can treat the case $\beta_N \propto N$ (Dyson Brownian Motion) via bootstrap arugment.

Verifying the conditions I

Let us verify the allowable scalings and ranges of s for some examples.

For independent $x_i \sim \mu^\circ$, the usual CLT implies that

$$\sqrt{N}\mathrm{Fluc}_N^\circ \xrightarrow[N\to\infty]{Law} \eta^\circ,$$

where η° is the centered Gaussian field with

$$\langle \phi, \eta^\circ \rangle \sim \mathcal{N}(0, \langle \phi^2, \mu^\circ \rangle - \langle \phi, \mu^\circ \rangle^2).$$

I.e. $\sigma_N = \sqrt{N}$.

Verifying the conditions II

Estimates for the mean-field limit imply that if $\mathsf{s} < \mathsf{d}-2$, then

$$\sup_{0\le t'\le t}\int_0^t \mathbb{E}[\mathsf{F}_N(X_N^{t'},\mu_{\beta_N}^{t'}) + \mathsf{o}_N^{t'}]dt' \le C(t,\mu_{\beta_N}^\circ)\mathbb{E}[\mathsf{F}_N(X_N^\circ,\mu_{\beta_N}^\circ) + \mathsf{o}_N^\circ]$$

$$\sim \left[\frac{\log N}{N}\mathbf{1}_{\mathsf{s}=0} + N^{\frac{\mathsf{s}}{\mathsf{d}}-1}\right].$$

For $\sigma_N\times$ RHS to vanish as $N\to\infty$, need $\mathsf{s} < \frac{\mathsf{d}}{2}$; Coulomb only if $\mathsf{d}\le 3$.

If $\mathsf{d}-2\le \mathsf{s} < \mathsf{d}$, then this analysis may be extended to show that

$$\sup_{0\le t'\le t}\int_0^t \mathbb{E}[\mathsf{F}_N(X_N^{t'},\mu_{\beta_N}^{t'}) + \mathsf{o}_N^{t'}]dt' \lesssim \sigma_N\left(\frac{\log N}{N}\mathbf{1}_{\mathsf{s}=0} + N^{\frac{\mathsf{s}}{\mathsf{d}}-1} + \frac{N^{4+s}}{\beta_N}\right).$$

If $0\le \mathsf{s} < \mathsf{d}$ and $\mathbb{M} = -\mathbb{I}$, then using modulated free energy,

$$\sup_{0\le t'\le t}\int_0^t \mathbb{E}[\mathsf{F}_N(X_N^{t'},\mu_{\beta_N}^{t'}) + \mathsf{o}_N^{t'}]dt' \lesssim E_{N,\beta_N}((\mu^\circ)^{\otimes N},\mu^\circ) + \mathsf{o}_N^\circ$$

$$= \mathbb{E}\Big[\mathsf{F}_N(X_N^\circ,\mu^\circ) + \mathsf{o}_N^\circ\Big].$$

Modulated Gibbs measures and fluctuations I

For fixed $\mu \in \mathcal{P}(\mathbb{D}^{\mathrm{d}})$, recall the modulated Gibbs measure

$$d\mathbb{Q}_{N,\beta}(\mu) := \frac{e^{-\beta N \mathsf{F}_N(X_N,\mu)}}{\mathsf{K}_{N,\beta}(\mu)} d\mu^{\otimes N},$$

where $\mathsf{K}_{N,\beta}(\mu)$ is the modulated partition function. This includes the case of the canonical Gibbs ensemble $\mathbb{P}_{N,\beta}$, taking $\mu = \mu_\beta$.

- $\beta = 0$ (infinite temperature, iid) – $\mathbb{Q}_{N,0}(\mu) = \mu^{\otimes N}$.
- $\beta = \infty$ (zero temperature, ground state) – $\mathbb{Q}_{N,\infty}(\mu)$ is (formally) a probability measure supported on the set of minimizers of energy $X_N \mapsto \mathsf{F}_N(X_N, \mu)$.

One can produce a range of σ_N by considering initial data $X_N^\circ \sim \mathbb{Q}_{N,\beta_N}(\mu)$ and tuning β_N...

Modulated Gibbs measures and fluctuations II

Conjecture: Suppose that $X_N \sim \mathbb{Q}_{N,\beta_N}(\mu)$. Then the fluctuations $\frac{1}{N}\sum_{i=1}^N \delta_{x_i} - \mu$ have normalization size $\sigma_N = \max(\sqrt{\beta_N N}, \sqrt{N})$ (and are Gaussian), at least when $\beta_N N^{\frac{s}{d}-1} \lesssim 1$.

- **Punchline:** Order (i.e. lowering temperature) creates smaller fluctuations.
- Can prove conjecture using Stein's method in the case $\beta_N N^{\frac{2s}{d}-1} = o(1)$ as $N \to \infty$, which should be interpreted as a "high-temperature result."
- $\sqrt{N}$ is the normalization for iid particles; shows that $\beta_N \lesssim 1$ is not so different than iid case (complete disorder)
- Clear that normalization $\sqrt{\beta_N N}$ cannot hold all the way to zero temperature (i.e. $\beta_N = N^\alpha$ and let $\alpha \to \infty$); otherwise would imply the ground state fluctuation is zero...

Final thoughts

Summary I

Modulated (free) energy and these commutator-type functional inequalities are powerful tools for studying the large N behavior of these systems.

- Mean-field limits Serfaty 2020, Duerinckx-Serfaty 2020, Bresch-Jabin-Wang 2019-2020, R. 2020-2022, Q.H. Nguyen-R.-Serfaty 2021, Chodron de Courcel-R.-Serfaty 2023, Ménard 2023, Ben Porat-Carrillo-Jabin 2025
- Second-order FIs also have applications to mean-field limits with special kinds of multiplicative noise R. 2020
- Supercritical mean-field limits Han Kwan-Iacobelli 2020, R. 2021, R.-Serfaty 2024
- Also used to study scaling limits of quantum systems Golse-Paul 2020, R. 2021, Ben Porat 2022
- Central limits for fluctuations/cumulant bounds Leblé-Serfaty 2018, Serfaty 2023, Peilen-Serfaty, J. Huang-R.-Serfaty, R.-Serfaty, Cecchi-Nikolaev 2025
- Dynamical LDPs Hess Childs 2023

There is a crucial interplay between temperature, singularity of the interaction, timescales, and rates of convergence!

Despite progress, a number of questions remain...

Elephant in the room: kinetic systems I

What about the original (underdamped) Langevin system?

$$\begin{cases} \dot{x}_i = v_i \\ \dot{v}_i = -\frac{1}{N} \sum_{\substack{1 \le j \le N \\ j \neq i}} \nabla \mathsf{g}(x_i - x_j), \end{cases}$$

for which the expected mean-field PDE is the *Vlasov equation*

$$\begin{cases} \partial_t f + v \cdot \nabla_x f + E \cdot \nabla_v f = 0 \\ \rho = \int_{\mathbb{R}^\mathsf{d}} df(\cdot, v) \\ E = -\nabla \mathsf{g} * \rho, \end{cases} \qquad (x, v) \in \Omega \times \mathbb{R}^\mathsf{d}.$$

- Duerinckx-Serfaty 2020 treated the monokinetic case $f(x, v) = \rho(x)\mathrm{d}(v - u(x))$, which is amenable to the modulated energy method.
- Bresch-Jabin-Soler 2022 log case with noise (Vlasov-Fokker-Planck); only covers $\mathsf{d} = 2$ Coulomb case.
- Bresch-Duerinckx-Jabin 2024 $\nabla \mathsf{g} \in L^2$ w/ or w/o noise; barely misses $\mathsf{d} = 2$ Coulomb case.
- Most physically relevant $\mathsf{d} = 3$ Coulomb case open for non-monokinetic data!

The End

Thank you for your attention!