

Quadratic Mean-Field Games and Entropy Minimization

Guillaume Carlier ^a

Based on joint works with Jean-David Benamou, Simone Di Marino and Luca Nenna.

EDP, transport optimal et applications, Essaouira, October 2018.

^aCEREMADE, Université Paris Dauphine and MOKAPLAN (Inria-Dauphine).

Introduction

Aim of this talk is to present a new viewpoint, related to the so-called entropic interpolation (a problem which goes back to Schrödinger in the 1930's) on second-order MFGs with a quadratic Hamiltonian and a variational structure.

More precisely, as we know from the seminal works of Lasry and Lions, the second-order MFG system

$$\begin{cases} -\partial_t \varphi - \Delta \varphi + \frac{1}{2} |\nabla \varphi|^2 = \frac{\delta F}{\delta \rho}(\rho), & \varphi|_{t=1} = \frac{\delta G}{\delta \rho}(\rho_1) \\ \partial_t \rho - \Delta \rho - \operatorname{div}(\rho \nabla \varphi) = 0, & \rho|_{t=0} = \rho_0. \end{cases}$$

($t \in (0, 1)$ and $x \in \mathbf{T}^d$, could be \mathbf{R}^d as well taking some care of what's happening at infinity..) is, in some sense, the system of optimality conditions for the optimal control problem:

$$\inf_{(\rho, v)} \int_0^1 \int_{\mathbf{T}^d} \frac{1}{2} \rho |v|^2 dx dt + \int_0^1 F(\rho_t) dt + G(\rho_1)$$

subject to

$$\partial_t \rho - \Delta \rho + \operatorname{div}(\rho v) = 0, \quad \rho|_{t=0} = \rho_0.$$

This is a convex problem when F and G are local and convex. We shall assume that either F and G are convex and local (plus growth conditions) or that they are *regular*, this ensures existence of minimizers.

We shall see that this problem is a (Lagrangian) entropy minimization problem (whose discretization in time can be solved numerically efficiently).

Outline

- ① Schrödinger's problem
- ② The structure of entropic interpolation
- ③ MFGs by Entropy minimization
- ④ Numerical results

Schrödinger's problem

Let $\mathbf{T}^d := \mathbf{R}^d / 2\pi\mathbf{Z}^d$, $\Omega := C([0, 1], \mathbf{T}^d)$ and R be the reversible Wiener measure i.e. $R \in \mathcal{P}(\Omega)$ is given by

$$R = \frac{1}{(2\pi)^d} \int_{\mathbf{T}^d} \text{Law}(x + B) dx \quad (1)$$

where B is the standard Brownian motion starting at 0 (that is the Markov process whose generator is Δ on \mathbf{T}^d).

Given ρ_0 and ρ_1 probability measures on \mathbf{T}^d (with finite entropy), the Schrödinger Bridge problem between ρ_0 and ρ_1 reads

$$S(\rho_0, \rho_1) := \inf \left\{ H(Q|R) : e_{0\#}Q = \rho_0, e_{1\#}Q = \rho_1 \right\}$$

where

$$H(Q|R) := \int_{\Omega} \ln \left(\frac{dQ}{dR} \right) dQ$$

and e_t is the evaluation at time t map, $e_t(\omega) = \omega(t)$.

Seems a complicated problem but can be brought down to a static problem (C. Léonard). Denote by $\Pi(\rho_0, \rho_1)$ the set of transport plans between ρ_0 and ρ_1 i.e. the set of joint probability measures having ρ_0 and ρ_1 as marginals.

Now disintegrate Q and R with respect to

$Q_{0,1} := (e_0, e_1)_\# Q \in \Pi(\rho_0, \rho_1)$ and $R_{0,1} := (e_0, e_1)_\# R$ respectively:

$$Q = \int Q^{x,y} Q_{0,1}(dx, dy), \quad R = \int R^{x,y} R_{0,1}(dx, dy)$$

Note that $R_{0,1}$ can be identified with its density $G_{0,1}(x - y)$ which is nothing but the heat kernel and $R^{x,y}$ is the Brownian Bridge between x and y .

Since

$$H(Q|R) = H(Q_{0,1}|R_{0,1}) + \int H(Q^{x,y}|R^{x,y})Q_{0,1}(dx, dy)$$

hence the entropy minimizing strategy consists in taking $Q^{x,y} = R^{x,y}$ and $Q_{0,1}$ solution of the static problem:

$$S(\rho_0, \rho_1) := \inf \left\{ H(Q_{0,1}|R_{0,1}), Q_{0,1} \in \Pi(\rho_0, \rho_1) \right\}.$$

The optimality condition for this static problem (strictly convex minimization) reads

$$Q_{0,1}(x, y) = G_{0,1}(x - y)f_0(x)g_1(y)$$

and the fixed marginal conditions impose

$$(G_{0,1} \star g_1)f_0 = \rho_0, \quad (G_{0,1} \star f_0)g_1 = \rho_0 \quad (2)$$

this is the so-called Schrödinger system. Here the log of f_0 and g_1 are Lagrange multipliers associated to the marginal constraints, f_0 and g_1 are called Schrödinger potentials.

Existence of such potentials is not obvious at all but has been very much studied since the 1940's (Bernstein, Fortet, Beurling, Föllmer, Nussbaum, Borwein, Lewis, Léonard...). These potentials are unique up to $(f_0, g_1) \mapsto (\lambda f_0, g_1/\lambda)$.

The structure of entropic interpolation

Assume that we have a pair of Schrödinger potentials, f_0, g_1 , set then $f_1 := G_{0,1} \star f_0$, $g_0 := G_{0,1} \star g_1$ so that $\rho_0 = f_0 g_0$, $\rho_1 = f_1 g_1$. Interpolation: $\rho_t := f_t g_t$ where f_t and g_t respectively solve a forward and backward heat equation

$$\partial_t f = \Delta f, f|_{t=0} = f_0, \partial_t g = -\Delta g, g|_{t=1} = g_1. \quad (3)$$

This (Eulerian) entropic interpolation ρ solves Fokker-Planck with the drift $\nabla\varphi$ where $\varphi = -\log g$ (this is Hopf-Cole...):

$$\partial_t \rho - \Delta \rho - \operatorname{div}(\rho \nabla \varphi) = 0, \quad (4)$$

whereas φ solves the backward Hamilton-Jacobi-Bellman equation

$$-\partial_t \varphi - \frac{1}{2} \Delta \varphi + \frac{1}{2} |\nabla \varphi|^2 = 0 \quad (5)$$

At least formally, $(\rho, -\nabla\varphi)$ solves

$$\text{FP}(\rho_0, \rho_1) := \inf \left\{ \frac{1}{2} \int_0^1 \int_{\mathbf{R}^d} |v_t(x)|^2 \mu_t(x) dx dt \right\},$$

subject to

$$\partial_t \mu - \Delta \mu + \text{div}(\mu v) = 0, \quad \mu_{t=0,1} = \rho_0, \rho_1$$

And one has

$$S(\rho_0, \rho_1) = \text{FP}(\rho_0, \rho_1) + \int_{\mathbf{T}^d} \rho_0 \ln(\rho_0).$$

where we recall that

$$S(\rho_0, \rho_1) := \inf \left\{ H(Q|R) : e_{0\#}Q = \rho_0, e_{1\#}Q = \rho_1 \right\}$$

The solution of $S(\rho_0, \rho_1)$ is Markovian it is the law of the diffusion process X :

$$dX_t = -\nabla\varphi_t(X_t)dt + dW_t, \quad X_0 \sim \mu_0.$$

Conversely if Q solves the dynamic Schrödinger problem, $\rho_t := e_{t\#}Q$ is an optimal trajectory for $\text{FP}(\rho_0, \rho_1)$.

The manipulations above are slightly formal but can be made rigorous (juste under a finite entropy condition):

- By Girsanov theory (Léonard),
- By stochastic control arguments (Mikami).

Nice references (connections with Nelson's stochastic mechanics, with optimal transport, large deviations, functional inequalities....): Beurling 1960, Föllmer, 1988, Léonard 2012, 2014, Gentil, Léonard, Ripani 2016, Mikami 1991, Zambrini 1986, Jamison 1975. Incompressible fluids: Yasue 1982, Arnaudon, Cruzeiro, Léonard, Zambrini, 2017, Benamou-C.-Nenna 2017....

MFGs by Entropy minimization

In the entropic interpolation, as in the optimal transport problem, we prescribe the two end-points measures $e_0 \# Q = \rho_0$ and $e_1 \# Q = \rho_1$. Instead, we can penalize/put constraints on the marginals $((e_t) \# Q)_{t \in [0,1]}$. Important example: incompressibility i.e. $e_t \# Q = (2\pi)^{-d} \mathcal{L}^d$, such Q 's are called generalized incompressible flows, following the seminal contributions of Yann Brenier. See the recent work of Arnaudon, Cruzeiro, Léonard, Arnaudon, Benamou-C.-Nenna.

Consider the (Eulerian) variational formulation of the MFG system we started with:

$$\text{MFG} := \inf_{\rho, v} \int_0^1 \int_{\mathbf{T}^d} \frac{1}{2} \rho |v|^2 dx dt + \int_0^1 F(\rho_t) dt + G(\rho_1)$$

subject to

$$\partial_t \rho - \Delta \rho + \text{div}(\rho v) = 0, \quad \rho|_{t=0} = \rho_0.$$

Consider now its (Lagrangian) Schrödinger version

$$\text{SMFG} := \inf_{Q : e_0 \# Q = \rho_0} H(Q|R) + \int_0^1 F(e_t \# Q) dt + G(e_1 \# Q)$$

Equivalence with Mean-Field Games

Theorem 1 *Under the assumptions that F and G are either convex or regular and $\int_{\mathbf{T}^d} \rho_0 \ln(\rho_0) < +\infty$ we have*

$$\text{SMFG} = \text{MFG} + \int_{\mathbf{T}^d} \rho_0 \ln(\rho_0)$$

and when Q solves SMFG, $\rho_t := e_{t\#}Q$ solves MFG.

Ideas of the proof (which does not use fine Girsanov-like arguments). First discretize both problems in time, $h = 1/N$ a time step discretization, given $(\mu_0, \dots, \mu_N) \in \mathcal{P}(\mathbf{T}^d)^{N+1}$ set

$$J_N(\mu_0, \dots, \mu_N) := \frac{1}{N} \sum_{i=1}^N F(\mu_i) + G(\mu_N)$$

Define

$$\text{FP}_h(\mu, \nu) := \inf_{\rho, v: \partial_t \rho - \Delta \rho + \text{div}(\rho v) = 0} \int_0^h \int_{\mathbf{T}^d} \frac{1}{2} \rho |v|^2 \quad : \quad \rho|_{t=0, h} = \mu, \nu \}$$

And consider the time-discretization of MFG

$$\text{MFG}_h := \inf_{\mu_0 = \rho_0, \mu_1, \dots, \mu_N} \sum_{i=0}^{N-1} \text{FP}_h(\mu_{i+1}, \mu_i) + J_N(\mu_0, \dots, \mu_N)$$

As well as the time-discretization of MFG:

$$\text{SMFG}_h := \inf_{\mu_0 = \rho_0, \mu_1, \dots, \mu_N} S_N(\mu_0, \dots, \mu_N) + J_N(\mu_0, \dots, \mu_N)$$

where S_N is given by the multi-marginal Schrödinger problem

$$S_N(\mu_0, \dots, \mu_N) := \inf \{ H(Q|R) \mid e^{\frac{k}{N}} \# Q = \mu_k, k = 0, \dots, N \}.$$

We first show that

$$S_N(\mu_0, \dots, \mu_N) = \sum_{i=0}^{N-1} \text{FP}_h(\mu_{i+1}, \mu_i) + \int_{\mathbf{T}^d} \mu_0 \ln(\mu_0).$$

and then prove a Γ -convergence results for both approximations MFG_h and SMFG_h .

Idea of the proof for the identity (noisy Benamou-Brenier formula)

$$S_N(\mu_0, \dots, \mu_N) = \sum_{i=0}^{N-1} \text{FP}_h(\mu_{i+1}, \mu_i) + \int_{\mathbf{T}^d} \mu_0 \ln(\mu_0).$$

is by induction on N , $N = 1$ is classical (Chen, Pavon, Georgiou, Gentil, Léonard, Ripani, Gigli, Tamanini). The induction step is by using the Markovianity of the reference measure R which forces entropy minimizers with marginal constraints to be Markov themselves giving some additivity of the minimal entropy.

Γ -convergence is more involved, main ingredients are a priori estimates (in particular on the Fisher information) on finite-energy curves of measures, suitable smoothing approximations and Prokhorov's like arguments on the path space.

By these arguments, in fact we have a general formula for every fixed suitable (finite energy) curve of measures $t \mapsto \mu_t$:

$$\inf\{H(Q|R) : e_{t\#}Q = \mu_t\} = \int_{\mathbf{T}^d} \mu_0 \ln(\mu_0) + \inf_v \frac{1}{2} \int_0^T \int_{\mathbf{T}^d} |v|^2 \mu$$

where on the right-hand side the drift v is constrained by

$$\partial_t \mu - \Delta \mu + \operatorname{div}(\mu v) = 0.$$

Note that there are infinitely many marginal constraints in the entropy minimization problem.

Numerical results

Entropy minimization (as a regularization of OT) is particularly appealing from a numerical viewpoint. Start with a discrete OT problem, LP

$$\min_{\gamma \in \Pi(\mu, \nu)} \sum_{i,j} c_{ij} \gamma_{ij}$$

where $\gamma \in \Pi(\mu, \nu)$ reads

$$\gamma_{ij} \geq 0, \quad \sum_j \gamma_{ij} = \mu_i, \quad \sum_i \gamma_{ij} = \nu_j$$

Entropic regularization, $\varepsilon > 0$,

$$\inf_{\gamma \in \Pi(\mu, \nu)} \sum_{i,j} c_{ij} \gamma_{ij} + \varepsilon \sum_{i,j} \gamma_{ij} \log(\gamma_{ij}). \quad (6)$$

The approximation (6) is the same as the Kullback-Leibler projection of the Gibbs kernel on the set of transport plans

$$\inf \text{KL}(\gamma|\theta) \text{ where } \theta_{ij} := e^{-\frac{c_{ij}}{\varepsilon}}.$$

where

$$\text{KL}(\gamma|\theta) := \sum_{ij} \gamma_{ij} \log \left(\frac{\gamma_{ij}}{\theta_{ij}} \right)$$

The projection of θ is characterized by

$$\gamma_{ij} = a_i b_j \theta_{ij}$$

with a_i, b_j positive and such that the marginal constraints are met i.e.

$$a_i = \frac{\mu_i}{\sum_j b_j \theta_{ij}} := R_i(b), \quad b_j = \frac{\nu_j}{\sum_i a_i \theta_{ij}} := S_j(a)$$

which can be rewritten as a fixed point-problem on a only: find a in the positive cone such that $a = Ta$ (with $T = R \circ S$).

$$d_H(a, a') := \log \left(\frac{\max_i \frac{a_i}{a'_i}}{\min_i \frac{a_i}{a'_i}} \right), \quad (a, a') \in (0, +\infty)^N.$$

for which T is a contraction.

Sinkhorn scaling algorithm: iterate the map T . Linear convergence.

Note that computations are completely parallelizable, only $I + J$ coefficients to store at each step. Old idea, huge literature:

- optimization (Bregman, Bauschke, Combettes, Lewis, Cominetti, San-Martin...),
- probability and statistics (Csisczar, Dykstra, Léonard, Föllmer, Rüschendorf...),
- tropical algebra, Perron-Frobenius theory (Gaubert, Nussbaum, Borwein, Chen, Pavon, Georgiou...).

Sinkhorn algorithm was remis au goût du jour in an OT framework only recently (Marco Cuturi, Alfred Galichon). Can easily be adapted to more general problems: multi-marginals, barycenters, partial transport See Benamou, C., Cuturi, Nenna, Peyré. Highly paralelizable. Special case of the quadratic cost, convolution with the heat kernel.

Generalized incompressible flows. Simple periodic 1d case.
Fixed final (Lagrangian) configuration of the fluid, somehow
extreme case where F enforces ρ to be uniform at each time.



