Mean field games via probability manifold I

Wuchen Li

Mean field games summer school, June 26, 2018

Introduction

Lecture I

- Static Population games;
- Wasserstein Gradient flow on graphs;
- Entropy dissipation.

Lecture II

- Differential Population games;
- Wasserstein Hamiltonian flow on graphs;
- Schrödinger equation on graphs;
- Schrödinger bridge problems on graphs.

Games

Game contains: Players; Strategies; Payoffs.

- Players: 2;
- Strategies: $S_1 = S_2 = \{\text{Rock, Paper, Scissors}\}$;
- Payoffs: $F_1, F_2 \; : \; S_1 \times S_2 \to \{+1, 0, -1\}$.

Stag hunt

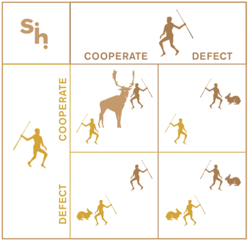

- Players: Infinity;
- Strategy set: $S = \{C, D\}$; Players form (ρ_C, ρ_D) with $\rho_C + \rho_D = 1$;
- Payoffs: $F(\rho) = (F_C(\rho), F_D(\rho))^T = A\rho$, where $A = \begin{pmatrix} 3 & 0 \\ 2 & 2 \end{pmatrix}$, meaning a deer worthing 6, a rabbit worthing 2.

Population game

Population games model the strategic interactions in large populations of small, anonymous agents.

- Strategy set

$$S = \{1, \cdots, n\} \ ;$$

- Players (Simplex)

$$\mathcal{P} = \{(\rho_i)_{i=1}^n \in \mathbb{R}^n \ : \ \sum_{i=1}^n \rho_i = 1 \ , \ \rho_i \geq 0\} \ ;$$

- Payoff function to strategy i: $F_i : \ \mathcal{P} \to \mathbb{R}$. E.g.

$$F(\rho) = (F_i(\rho))_{i=1}^n = A\rho \ , \quad \text{where} \quad A \in \mathbb{R}^{n\times n} \ .$$

Applications

Social Network, Biology species, Virus, Trading, Cancer, Congestion and many more (See Sandholm's textbook). We plan to design new dynamics to model for the evolution of a game, and study their asymptotic properties.

Nash Equilibrium

Nash Equilibrium (NE): Players have no unilateral incentive to deviate from their current strategies.
$\rho^* = (\rho_i^*)_{i=1}^n$ is a Nash equilibrium (NE) if

$$\rho_i^* > 0 \quad \text{implies that} \quad F_i(\rho^*) \geq F_j(\rho^*) \quad \text{for all } j \in S.$$

E.g. if $S = \{C, D\}$, $F_C(\rho) = 3\rho_C$, $F_D(\rho) = 2$, it is simple to check $(1,0)$, $(\frac{2}{3}, \frac{1}{3})$, $(0,1)$ are three NEs.

A particular type of game, named Potential games, are widely considered: There exists a potential $\mathcal{F} \ : \ \mathcal{P} \to \mathbb{R}$, such that

$$\frac{\partial}{\partial \rho_i} \mathcal{F}(\rho) = F_i(\rho) \ .$$

If $F(\rho) = A\rho$, consider $\mathcal{F}(\rho) = \frac{1}{2}\rho^T A\rho$, where A is a symmetric matrix.

In potential games, NE is the critical points of

$$\max_{\rho} \ \mathcal{F}(\rho) \ : \ \rho \in \mathcal{P} \ .$$

Evolutionary dynamics

In literature, people have designed many dynamics, named mean or evolutionary dynamics, to model games. Typical examples are BNN (Brown-von Neumann-Nash 1950), Best response dynamics (Gilboa-Matsui 1991), Logit (Fudenberg-Levine 1998), Smith dynamics (Smith 1983) etc.

A famous dynamics is Replicator dynamics (Taylor and Jonker 1978)

$$\frac{d\rho_i}{dt} = \rho_i(F_i(\rho) - \bar{F}(\rho)), \quad \text{where } \bar{F}(\rho) = \sum_{j \in S} \rho_j F_j(\rho).$$

In potential games, the Replicator dynamics is a gradient flow in probability set $\mathcal{P}$ w.r.t a Fisher-Rao (Shahshahani) metric (Akin (1980)).

Goal

Design the other dynamics for evolutionary games with

- Evolution only using local information in strategies;
- Gradient flow in potential games;
- Ability to include white noise perturbations.

Related materials:

Dynamical system; Optimal transport; Riemannian Geometry; Partial differential equations; Topology; Graph theory; Information theory.

Best Reply dynamics

Choose $S = \mathbb{T}^d$. Consider the model (P. Degond, J. G. Liu, C. Ringhofer, 2014.)

$$dX_t = \nabla_{X_t} F(X_t, \rho)dt + \sqrt{2\beta} dW_t \ , \quad X_t \in \mathbb{T}^d.$$

where W_t is the standard Brownian motion (or noise level) and $\rho(t, x)$ is the density function of X_t. In this case, the mean field equation refers to the evolution of density:

$$\frac{\partial \rho}{\partial t} + \nabla_x \cdot \big(\rho \nabla_x F(x, \rho)\big) = \beta \Delta_x \rho \ .$$

- Individual players change their pure strategies according to the direction that maximizes their own payoff functions most rapidly. And the Brownian motion represents uncertainties.
- In potential games, this PDE is the gradient flow w.r.t the optimal transport metric.

Density manifold

Optimal transport has a variational formulation (Benamou-Brenier 2000):

$$\inf_v \int_0^1 \mathbb{E} v(t, X_t)^2 \, dt ,$$

where $\mathbb{E}$ is the expectation operator and the infimum runs over all vector fields v_t, such that

$$\dot{X}_t = v(t, X_t) , \quad X_0 \sim \rho^0 , \quad X_1 \sim \rho^1 .$$

Under this metric, the probability set has a Riemannian geometry structure[1].

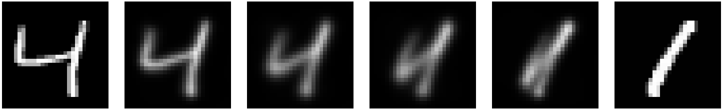

[1] John D. Lafferty: the density manifold and configuration space quantization, 1988.

Brownian motion and Entropy dissipation

The gradient flow of the entropy

$$\mathcal{H}(\rho) = \int_{\mathbb{T}^d} \rho(x) \log \rho(x) dx \ ,$$

w.r.t. optimal transport metric is:

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla \log \rho) = \Delta \rho \ .$$

Entropy dissipation:

$$\frac{d}{dt}\mathcal{H}(\rho) = \int_{\mathbb{T}^d} \log \rho \nabla \cdot (\rho \nabla \log \rho) dx = -\int_{\mathbb{T}^d} (\nabla \log \rho)^2 \rho dx \ .$$

Evolutionary games via optimal transport

Question:

Can we derive a Best-Reply dynamics on a discrete strategy set?

Answer:

Yes, we need a discrete dynamical optimal transport metric. Using this metric, we derive the gradient flow as the evolutionary dynamics for potential games.

Recent Developments:

Mielke, Maas, Chow, Zhou, Li, Huang, Erbar, Fathi, Gangbo, Mou and many more.

Basic setting

Graph with finite vertices

$G = (S, E, \omega), \quad S = \{1, \cdots, n\}, \quad \text{E is the edge set}, \quad \omega \text{ is the weight ;}$

Probability set

$$\mathcal{P} = \{(\rho_i)_{i=1}^n \mid \sum_{i=1}^n \rho_i = 1, \ \rho_i \geq 0\} \ ;$$

Noise potential:

$$\mathcal{F}(\rho) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n A_{ij} \rho_i \rho_j \quad - \quad \beta \sum_{i=1}^n \rho_i \log \rho_i \ ,$$

Interaction Potential energy   Boltzmann-Shannon entropy

where A is a given symmetric matrix and $\beta > 0$ is a given constant.

Definition I

We plan to find the discrete analog of density manifold (Maas, Mielke, Chow et.al).

First, it is natural to define a *vector field on a graph*

$$v = (v_{ij})_{(i,j)\in E}, \quad \text{satisfying } v_{ij} = -v_{ji}.$$

Given a potential $\Phi = (\Phi_i)_{i=1}^n$, a gradient vector field refers

$$\nabla_G \Phi_{ij} = \sqrt{\omega_{ij}}(\Phi_i - \Phi_j).$$

Definition II

We next define an *inner product* of two vector fields v^1, v^2:

$$(v^1, v^2)_\rho := \frac{1}{2} \sum_{(i,j)\in E} v^1_{ij} v^2_{ij} \theta_{ij}(\rho);$$

and a *divergence* of a vector field v at $\rho \in \mathcal{P}$:

$$\mathrm{div}_G(\rho v) := \Big(- \sum_{j\in N(i)} \sqrt{\omega_{ij}} v_{ij} \theta_{ij}(\rho) \Big)_{i=1}^n.$$

Here θ represents the probability weight on the edge. E.g. θ_{ij} is given by a upwind scheme:

$$\theta_{ij}(\rho) = \begin{cases} \rho_i/d_i & \text{if } \frac{\partial}{\partial \rho_i}\mathcal{F}(\rho) > \frac{\partial}{\partial \rho_j}\mathcal{F}(\rho),\ j \in N(i); \\ \rho_j/d_j & \text{if } \frac{\partial}{\partial \rho_i}\mathcal{F}(\rho) < \frac{\partial}{\partial \rho_j}\mathcal{F}(\rho),\ j \in N(i); \\ \frac{\rho_i/d_i + \rho_j/d_j}{2} & \text{if } \frac{\partial}{\partial \rho_i}\mathcal{F}(\rho) = \frac{\partial}{\partial \rho_j}\mathcal{F}(\rho),\ j \in N(i). \end{cases}$$

where $d_i = \frac{\sum_{j\in N(i)} \omega_{ij}}{\sum_{(i,j)\in E} \omega_{ij}}$ is the volume form on the node. θ_{ij} has the other choices.

Optimal transport on a graph

Definition

For any ρ^0, $\rho^1 \in \mathcal{P}$, define the Wasserstein metric $W\colon \mathcal{P} \times \mathcal{P} \to \mathbb{R}$ by

$$W(\rho^0,\rho^1)^2$$
$$=: \inf_{v_t} \{ \int_0^1 (v_t, v_t)_{\rho_t} dt \colon \frac{d\rho_t}{dt} + \mathrm{div}_G(\rho_t v_t) = 0,\ \rho(0) = \rho^0,\ \rho(1) = \rho^1 \}.$$

Hodge decomposition on graphs

Consider a Hodge decomposition of vector field $v = (v_{ij})_{(i,j)\in E}$ on a graph G

$$v_{ij} = \underbrace{\nabla_G \Phi_{ij}}_{\text{Gradient}} + \underbrace{u_{ij}}_{\text{Divergence free}}$$

where $\Phi \in \mathbb{R}^n$ and $u \in \mathbb{R}^{|E|}$, with

$$\nabla_G \Phi_{ij} = \sqrt{\omega_{ij}}(\Phi_i - \Phi_j),$$

and $u_{ij} = -u_{ji}$ satisfying $\mathrm{div}_G(\rho u) = 0$, i.e.

$$\sum_{j\in N(i)} u_{ij}\theta_{ij} = 0.$$

In this case,

$$(v, v)_\rho = (\nabla_G \Phi, \nabla_G \Phi)_\rho + (u, u)_\rho.$$

Effect of Hodge decomposition

Lemma

The discrete Wasserstein metric is equivalent to

$$W(\rho^0, \rho^1)^2 = \inf_{\nabla_G \Phi} \int_0^1 (\nabla_G \Phi, \nabla_G \Phi)_\rho \, dt \ ,$$

where the infimum is taken among all discrete potential vector fields $\nabla_G \Phi$*, such that*

$$\frac{d\rho}{dt} + div_G(\rho \nabla_G \Phi) = 0 \ , \ \rho(0) = \rho^0, \ \rho(1) = \rho^1 \ .$$

$(\mathcal{P}, W)$ has a Riemannian geometry structure.

Probability manifold

Denote

$$\frac{d\rho}{dt} = -\mathrm{div}_G(\rho\nabla_G\Phi) = L(\rho)\Phi \ .$$

With this discrete Wasserstein metric, $\mathcal{P}$ forms a Riemannian manifold

$$\inf_{\rho(t)}\{\int_0^1 \dot\rho^T L(\rho)^{-1}\dot\rho dt\colon \rho(0)=\rho^0,\ \rho(1)=\rho^1\} \ .$$

Here $L(\rho)\in\mathbb{R}^{|V|\times|V|}$ is the weighted Laplacian matrix

$$L(\rho) = -\mathrm{div}_G(\rho\nabla_G) = -D^T\Theta(\rho)D \ ,$$

where $D\in\mathbb{R}^{|E|\times|V|}$ is a discrete gradient matrix, $D^T\in\mathbb{R}^{|V|\times|E|}$ is a discrete divergence matrix, and $\Theta(\rho)\in\mathbb{R}^{|E|\times|E|}$ is a diagonal weight matrix

$$\Theta(\rho)_{(i,j)\in E,(k,l)\in E} = \begin{cases}\theta_{ij} & \text{if } (i,j)=(k,l)\in E \ ; \\ 0 & \text{otherwise} \ .\end{cases}$$

Gradient flow in Riemannian manifold

The gradient flow in abstract form

$$\frac{d\rho}{dt} = \mathrm{grad}_W \mathcal{F}(\rho) \ ,$$

where the gradient is defined by

- Tangency:

$$\mathrm{grad}_W \mathcal{F}(\rho) \in \mathrm{T}_\rho \mathcal{P} = \left\{ (\sigma_i)_{i=1}^n \colon \sum_{i=1}^n \sigma_i = 0 \right\} \ .$$

- Duality:

$$\mathrm{grad}_W \mathcal{F}(\rho)^T L(\rho)^{-1} \sigma = \mathrm{d}\mathcal{F}(\rho) \cdot \sigma, \quad \text{for any } \sigma \in \mathrm{T}_\rho \mathcal{P} \ ,$$

where $\mathrm{d}\mathcal{F}(\rho) = (\frac{\partial}{\partial \rho_i} \mathcal{F}(\rho))_{i=1}^n$.

Derivation

Theorem

Given a potential game with a strategy graph $G = (S, E, \omega)$, a payoff matrix A. Then the gradient (ascent) flow of the free energy $\mathcal{F}(\rho)$ on $\mathcal{P}$ with respect to W is

$$\frac{d\rho}{dt} = L(\rho) d_\rho \mathcal{F}(\rho).$$

I.e.

$$\begin{aligned} \frac{d\rho_i}{dt} = & \sum_{j \in N(i)} \omega_{ij} \rho_j / d_j [F_i(\rho) - F_j(\rho) + \beta \log \frac{\rho_j}{\rho_i}]_+ \\ & - \sum_{j \in N(i)} \omega_{ij} \rho_i / d_i [F_j(\rho) - F_i(\rho) + \beta \log \frac{\rho_i}{\rho_j}]_+ \ . \end{aligned} \tag{1}$$

Asymptotical behavior

Theorem

For any initial condition $\rho^0 \in \mathcal{P}_+(S)$, (1) *has a unique solution* $\rho(t) : [0, \infty) \to \mathcal{P}_+(S)$.

(i) *The free energy* $\mathcal{F}(\rho)$ *is a Lyapunov function of* (1)*;*

(ii) *If* $\lim_{t\to\infty} \rho(t)$ *exists, call it* ρ^∞*, then* ρ^∞ *is one of the possible Gibbs measures, i.e.*

$$\rho_i^\infty = \frac{1}{K} e^{\frac{F_i(\rho^\infty)}{\beta}}, \quad K = \sum_{i=1}^{n} e^{\frac{F_i(\rho^\infty)}{\beta}} \quad \textit{for all } i \in S.$$

Entropy dissipation on population games

What is the rate of convergence to a Gibbs measure?

Motivation

- Entropy dissipation: Carrillo, McCann and Villani's work[2] for nonlinear Fokker-Planck equations on $\mathbb{T}^d$;
- Gradient flows: dynamical systems viewpoint!

[2] Carrillo, McCann and Villani, "Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates", 2003.

Entropy dissipation

Theorem (Entropy dissipation)

If the Gibbs measure ρ^∞ is a strict maximizer of $\mathcal{F}(\rho)$, then there exists a constant $C > 0$, such that

$$\mathcal{F}(\rho^\infty) - \mathcal{F}(\rho(t)) \le e^{-Ct}(\mathcal{F}(\rho^\infty) - \mathcal{F}(\rho^0)) \ .$$

Idea of Proof

The speed of convergence comes from comparing the ratio between the first and second derivative of $\mathcal{F}(\rho(t))$ along with the ODE. If one can find a constant $C > 0$, such that

$$\frac{d^2}{dt^2}\mathcal{F}(\rho(t)) \geq -C\frac{d}{dt}\mathcal{F}(\rho(t)) \ ,$$

holds for all $t \geq 0$. Then by integrating the above formula in $[t, +\infty]$, one obtains

$$\frac{d}{dt}[\mathcal{F}(\rho^\infty) - \mathcal{F}(\rho(t))] \geq -C[\mathcal{F}(\rho^\infty) - \mathcal{F}(\rho(t))] \ .$$

Proceed with the Gronwall's inequality, the result is proved.

Proof

In our case, the first derivative of energy along the gradient flow is

$$\frac{d}{dt}\mathcal{F}(\rho(t)) = F(\rho)^T\dot{\rho} = F(\rho)^T L(\rho)F(\rho) = \dot{\rho}^T L^{-1}(\rho)\dot{\rho} \ ,$$

while the second derivative forms

$$\frac{d^2}{dt^2}\mathcal{F}(\rho(t)) = 2\ \dot{\rho}^T d_{\rho\rho}\mathcal{F}(\rho)\dot{\rho} + \dot{\rho}^T L^{-1}(\rho)L(\dot{\rho})L^{-1}(\rho)\dot{\rho} \ .$$

Compare $\frac{d}{dt}\mathcal{F}(\rho(t))$ with $\frac{d^2}{dt^2}\mathcal{F}(\rho(t))$ to find

$$C := \inf_{\rho\in B(\rho^0)} \ \underbrace{\frac{2\dot{\rho}^T d_{\rho\rho}\mathcal{F}(\rho)\dot{\rho}}{\dot{\rho}^T L^{-1}(\rho)\dot{\rho}}}_{\text{Quadratic}} + \underbrace{\frac{\dot{\rho}^T L^{-1}(\rho)L(\dot{\rho})L^{-1}(\rho)\dot{\rho}}{\dot{\rho}^T L^{-1}(\rho)\dot{\rho}}}_{\text{Cubic}} \ .$$

Hessian operator at Gibbs measure

Let

$$\lambda_{\mathcal{F}}(\rho) = \min_{\Phi\in\mathbb{R}^n} \sum_{(i,j)\in E}\sum_{(k,l)\in E} h_{ij,kl}(\Phi_i - \Phi_j)\theta_{ij}(\Phi_k - \Phi_l)\theta_{kl}$$

s.t.

$$\sum_{(i,j)\in E}(\Phi_i - \Phi_j)^2\theta_{ij} = 1.$$

Here

$$h_{ij,kl} = \left(\frac{\partial^2}{\partial\rho_i\partial\rho_k} + \frac{\partial^2}{\partial\rho_j\partial\rho_l} - \frac{\partial^2}{\partial\rho_i\partial\rho_l} - \frac{\partial^2}{\partial\rho_j\partial\rho_k}\right)\mathcal{F}(\rho)\ .$$

This rate connects with Yano formula[3], which is related to Ricci curvature in geometry.

[3] Kentaro Yano, "On Harmonic and Killing Vector Fields", 38-45, Annals of Mathematics, 1958.

Stag Hunt

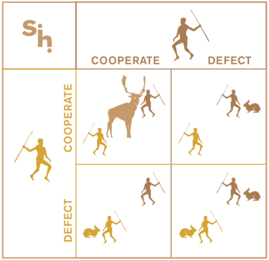

- Strategy set $\{C, D\}$;
- Players $\rho = (\rho_C, \rho_D)^T$;
- Payoff $F(\rho) = A\rho$ with $A = \begin{pmatrix} 3 & 0 \\ 2 & 2 \end{pmatrix}$.

Stag Hunt

We draw the vector field of the Fokker-Planck equation. Different noise levels lead to different NEs.

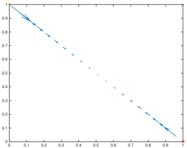

(c) $\beta = 5$

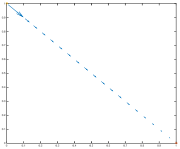

(d) $\beta = 0.5$

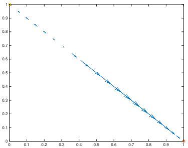

(e) $\beta = 0.1$

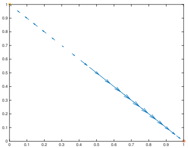

(f) $\beta = 0$

Rock-Scissors-Paper

- Strategy set $\{r, s, p\}$; Players $\rho = (\rho_r, \rho_s, \rho_p)^T$;
- Payoff $F(\rho) = A\rho$ with payoff matrix $A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$.

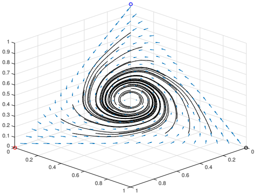

(g) $\beta = 0$

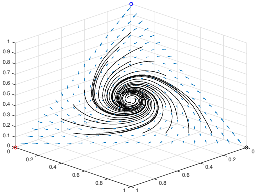

(h) $\beta = 0.1$

Bad Rock-Scissors-Paper

- Payoff $F(\rho) = A\rho$ with payoff matrix $A = \begin{pmatrix} 0 & -2 & 1 \\ 1 & 0 & -2 \\ -2 & 1 & 0 \end{pmatrix}$.

We demonstrate a Hopf Bifurcation. If β is large, there is a unique equilibrium around $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. If β is small, a limit cycle exists.

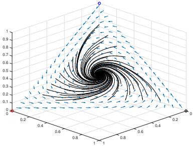

(i) $\beta = 0.5$

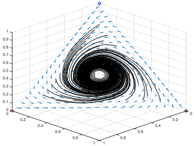

(j) $\beta = 0.1$

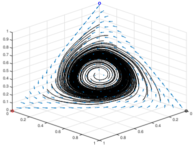

(k) $\beta = 0$

Optimal transport+Dynamical system

Boltzman-Shannon entropy

$$\mathcal{H}(\rho) = \sum_{i=1}^{n} \rho_i \log \rho_i \Rightarrow \mathrm{Hess}_{\mathbb{R}^n} \mathcal{H}(\rho) = \mathrm{diag}(\frac{1}{\rho_i}).$$

Thus our asymptotically dissipation rate forms

$$\lambda_{\mathcal{H}}(\rho) = \min\{\sum_{i=1}^{n} \frac{1}{\rho_i} (\mathrm{div}_G(\rho \nabla_G \Phi)|_i)^2 \colon \sum_{(i,j)\in E} (\Phi_i - \Phi_j)^2 \theta_{ij} = 1\} > 0.$$

Linear Entropy+ Yano formula

Consider

$$\mathcal{H}(\rho) = \int_{\mathcal{M}} \rho(x) \log \rho(x) dx,$$

whose Gibbs measure is a uniform measure. Then

$$\begin{aligned} &(\mathrm{Hess}_{\mathcal{P}(\mathcal{M})} \mathcal{H} \cdot \nabla\Phi, \nabla\Phi)_{\rho^*} \\ = &\int_{\mathcal{M}} [\mathrm{Ric}(\nabla\Phi, \nabla\Phi) + \mathrm{tr}(D^2\Phi^T D^2\Phi)]\rho^*(x) dx \\ = &\int_{\mathcal{M}} [\nabla \cdot (\rho^* \nabla\Phi)]^2 \frac{1}{\rho^*(x)} dx. \end{aligned}$$

The first equality is well known derived through **Bochner's formula**, while the second equality is new. It shows Yano's formula.

Reference

Cédric Villani
Optimal transport: Old and new, 2008.

B. Frieden
Science from Fisher Information: A Unification, 2004.

Shui-Nee Chow, Wuchen Li and Haomin Zhou
Entropy dissipation on finite graphs, DCDS, series A, 2018.

Shui-Nee Chow, Wuchen Li, Jun Lu and Haomin Zhou
Population games and discrete optimal transport, 2017.

Shui-Nee Chow, Wuchen Li and Haomin Zhou
A discrete Schrödinger equation via optimal transport, 2017.

Wilfrid Gangbo, Wuchen Li and Chenchen Mou.
Geodesic of Minimal Length in the Set of Probability Measures on Graph, 2018.

Wuchen Li.
Geometry of probability simplex via optimal transport, 2018.

Thanks.