Density Control of Interacting Agent Systems

Yongxin Chen

Georgia Tech

Dynamics and Discretization: PDEs, Sampling, and Optimization, Simons Institute, Berkeley

Swarm formation control

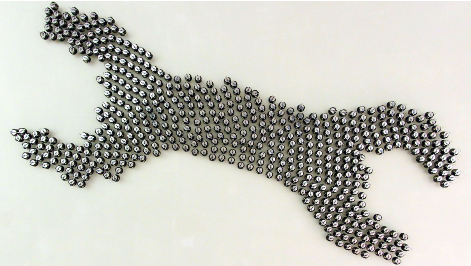

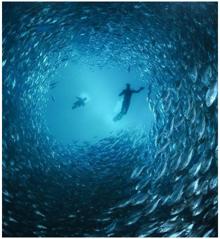

Optimal transport approach

Idea: configuration/distribution ρ_0 to configuration/distribution ρ_1

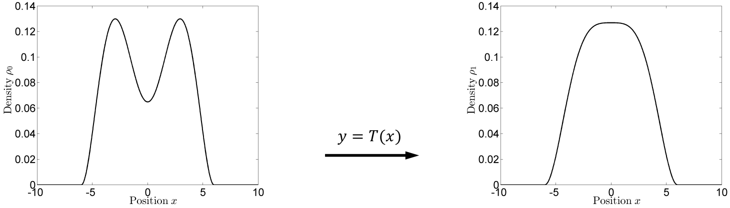

ρ_0 ρ_1

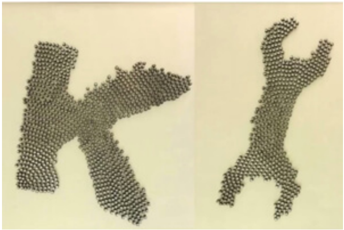

Collision avoidance interaction is ignored

Problem formulation

Given N agents

$$dX_t^i = -\frac{1}{N}\sum_{j=1}^{N} \nabla W(X_t^i - X_t^j)dt + u_t^i dt + \sqrt{\epsilon}dB_t^i,\ i = 1, \ldots, N$$

find a common strategy $u_t^i = \xi_t(X_t^i)$ that minimizes

$$\mathbb{E}\left\{\int_0^1 \frac{1}{2N}\sum_i \|\xi_t(X_t^i)\|^2 dt\right\}$$

and drives the particles/agents from initial configuration

$$\frac{1}{N}\sum_{i=1}^{N} \delta_{X_0^i} \approx \mu$$

to target configuration

$$\frac{1}{N}\sum_{i=1}^{N} \delta_{X_1^i} \approx \nu$$

Mean-field limit

McKean-Vlasov equation

$$\partial_t \rho_t + \nabla \cdot (\rho_t(-\nabla W * \rho_t + \xi_t)) - \frac{\epsilon}{2}\Delta\rho_t = 0$$

Mean-field formulation

$$\begin{aligned} \inf_{\rho,\xi} \quad & \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2}\|\xi_t(x)\|^2 \rho_t(x) dx dt \\ & \partial_t \rho_t + \nabla \cdot (\rho_t(-\nabla W * \rho_t + \xi_t)) - \frac{\epsilon}{2}\Delta\rho_t = 0 \\ & \rho_0 = \mu, \quad \rho_1 = \nu \end{aligned}$$

Special cases

Schrödinger bridge $W = 0$

$$\inf_{\rho,\xi} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2}\|\xi_t(x)\|^2 \rho_t(x) dx dt$$
$$\partial_t \rho_t + \nabla \cdot (\rho_t \xi_t) - \frac{\epsilon}{2}\Delta \rho_t = 0$$
$$\rho_0 = \mu, \quad \rho_1 = \nu$$

Optimal transport $W = 0, \epsilon = 0$

$$\inf_{\rho,\xi} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2}\|\xi_t(x)\|^2 \rho_t(x) dx dt$$
$$\partial_t \rho_t + \nabla \cdot (\rho_t \xi_t) = 0$$
$$\rho_0 = \mu, \quad \rho_1 = \nu$$

Reformulation in path measures

Path measure $\mathcal{P}$, marginal $\rho_t = (X_t)_\sharp \mathcal{P}$

$$\begin{aligned} \mathcal{P}: \quad & dX_t = -[\nabla W * \rho_t](X_t)dt + \xi_t(X_t)dt + \sqrt{\epsilon}dB_t \\ \mathcal{Q}(\mathcal{P}): \quad & dX_t = -[\nabla W * \rho_t](X_t)dt + \sqrt{\epsilon}dB_t \end{aligned}$$

Girsanov theorem

$$\mathrm{KL}(\mathcal{P}\|\mathcal{Q}(\mathcal{P})) = \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2\epsilon} \|\xi_t(x)\|^2 \rho_t(x) dx dt$$

Mean-field Schrödinger bridge

$$\begin{aligned} \min_{\mathcal{P}} \quad & \mathrm{KL}(\mathcal{P}\|\mathcal{Q}(\mathcal{P})) \\ & (X_0)_\sharp \mathcal{P} = \mu, \quad (X_1)_\sharp \mathcal{P} = \nu \end{aligned}$$

Mean-field Schrödinger bridges

Assumptions: W is twice continuously differentiable, symmetric, and has bounded Hessian

Goal: find most likely evolution of N particles

$$dX_t^i = -\frac{1}{N}\sum_{j=1}^{N} \nabla W(X_t^i - X_t^j)dt + u_t^i dt + \sqrt{\epsilon}dB_t^i,\ i = 1, \ldots, N$$

from $\frac{1}{N}\sum_{i=1}^{N} \delta_{X_0^i} \approx \mu$ to $\frac{1}{N}\sum_{i=1}^{N} \delta_{X_1^i} \approx \nu$

Large deviation principle

$$\text{Prob}(\mathcal{P}) \approx \exp(-N\ \text{KL}(\mathcal{P}\|\mathcal{Q}(\mathcal{P})))$$

Backhoff et al, 2020

Schrödinger PDE systems

HJB system

$$\partial_t \lambda + \frac{1}{2}\|\nabla\lambda\|^2 - \nabla\lambda^T \nabla W * \rho_t - \int \rho_t(y)\nabla\lambda(y)^T \nabla W(y-x)dy + \frac{\epsilon}{2}\Delta\lambda = 0$$

$$\partial_t \rho_t + \nabla\cdot(\rho_t(-\nabla W * \rho_t + \nabla\lambda)) - \frac{\epsilon}{2}\Delta\rho_t = 0$$

$$\rho_0 = \mu, \quad \rho_1 = \nu$$

Schrödinger system

$$\rho_t = \exp(-2W * \rho_t + \varphi_t + \psi_t)$$

$$\begin{aligned} \partial_t \psi_t + \frac{\epsilon}{2}\Delta\psi_t + \frac{1}{2}\|\nabla\psi_t\|^2 &= \int \nabla W(x-\tilde{x})^T(\nabla\psi_t(x) - \nabla\psi_t(\tilde{x}))\rho_t(\tilde{x})d\tilde{x} \\ -\partial_t \varphi_t + \frac{\epsilon}{2}\Delta\varphi_t + \frac{1}{2}\|\nabla\varphi_t\|^2 &= \int \nabla W(x-\tilde{x})^T(\nabla\varphi_t(x) - \nabla\varphi_t(\tilde{x}))\rho_t(\tilde{x})d\tilde{x} \end{aligned}$$

Backhoff et al, 2020

Discretization

Objective: $\mathrm{KL}(\mathcal{P}\|\mathcal{Q}(\mathcal{P}))$
Discretize $\mathcal{P}, \mathcal{Q}(\mathcal{P})$ over a grid:
time ($t_i = i/T, i = 0, 1, \ldots, T$) space ($N_x$ points each coordinate)
Discretized objective: $\mathbf{M}$ is a $T+1$ dimensional tensor

$$\langle \mathbf{C}(\mathbf{M}), \mathbf{M} \rangle + \epsilon \langle \mathbf{M}, \log \mathbf{M} \rangle$$

where

$$\langle \mathbf{M}, \log \mathbf{M} \rangle = \sum_{x_0, x_1, \ldots, x_T} M(x_0, x_1, \ldots, x_T) \log M(x_0, x_1, \ldots, x_T)$$

$$C(\mathbf{M})(x_0, x_1, \cdots, x_T) = \sum_{i=0}^{T-1} \frac{T}{2} \| x_{i+1} - x_i + \frac{1}{T} [\nabla W * P_i(\mathbf{M})](x_i) \|^2$$

Nonlinear entropy regularized multi-marginal optimal transport

Nonlinear MOT

$$\min_{\mathbf{M}} \quad \langle \mathbf{C}(\mathbf{M}), \mathbf{M} \rangle + \epsilon \langle \mathbf{M}, \log \mathbf{M} \rangle$$
$$P_0(\mathbf{M}) = \boldsymbol{\mu}, \quad P_T(\mathbf{M}) = \boldsymbol{\nu}$$

nonlinear cost tensor

$$C(\mathbf{M})(x_0, x_1, \cdots, x_T) = \sum_{i=0}^{T-1} \frac{T}{2} \| x_{i+1} - x_i + \frac{1}{T} [\nabla W * P_i(\mathbf{M})](x_i) \|^2$$

Proximal gradient algorithm

Composite optimization (F smooth, G possibly nonsmooth)

$$\min_{y \in \mathcal{Y}} F(y) + G(y)$$

Proximal gradient

$$y^{k+1} = \text{argmin}_{y \in \mathcal{Y}}\, G(y) + \frac{1}{2\eta}\|y - (y^k - \eta \nabla F(y^k))\|^2$$

Non-Euclidean proximal gradient (D: Bregman divergence, e.g., Kullback-Leibler divergence)

$$y^{k+1} = \text{argmin}_{y \in \mathcal{Y}}\, G(y) + \frac{1}{\eta} D(y, y^k) + \langle \nabla F(y^k), y \rangle$$

Proximal gradient for nonlinear MOT

Composite optimization $\mathsf{F}(\mathbf{M}) = \langle \mathbf{C}(\mathbf{M}), \mathbf{M} \rangle$

$$\min_{\mathbf{M} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \mathsf{F}(\mathbf{M}) + \epsilon \langle \mathbf{M}, \log \mathbf{M} \rangle$$

Set $\mathsf{D}(\cdot, \cdot) = \mathrm{KL}(\cdot \| \cdot)$

$$\mathbf{M}_{k+1} = \mathrm{argmin}_{\mathbf{M} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \langle \nabla \mathsf{F}(\mathbf{M}_k), \mathbf{M} \rangle + \frac{1}{\eta} \mathrm{KL}(\mathbf{M} \| \mathbf{M}_k) + \epsilon \langle \mathbf{M}, \log \mathbf{M} \rangle$$

Each iteration is an entropy regularized MOT

$$\mathbf{M}_{k+1} = \mathrm{argmin}_{\mathbf{M} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \langle \nabla \mathsf{F}(\mathbf{M}_k) - \frac{1}{\eta} \log \mathbf{M}_k, \mathbf{M} \rangle + (\epsilon + \frac{1}{\eta}) \langle \mathbf{M}, \log \mathbf{M} \rangle$$

$\nabla \mathsf{F}(\mathbf{M}) = \mathbf{C}(\mathbf{M}) + \mathsf{E}(\mathbf{M})$ with $\mathsf{E}(\mathbf{M})(x_0, x_1, \ldots, x_T) = \sum_{i=0}^{T-1} \mathsf{E}_i(x_i)$

$$\mathsf{E}_i(y) = \sum_{x_i, x_{i+1}} \nabla W(x_i - y)^T [x_{i+1} - x_i + \frac{1}{T} \nabla W * P_i(\mathbf{M})] P_{i,i+1}(\mathbf{M})(x_i, x_{i+1})$$

Graphical optimal transport

Multi-marginal optimal transport

$$\min_{\mathbf{M}\in\mathbb{R}_+^{d_1\times\cdots\times d_J}} \langle \mathbf{C}, \mathbf{M}\rangle + \epsilon\langle \mathbf{M}, \log \mathbf{M}\rangle$$

$$\text{subject to } P_j(\mathbf{M}) = \boldsymbol{\mu}_j, \text{ for } j \in \Gamma$$

with graphical cost (each α is a subset of $\{1, 2, \ldots, J\}$)

$$C(\mathbf{x}) = \sum_{\alpha\in F} C_\alpha(\mathbf{x}_\alpha)$$

$J = 6, \Gamma = \{2, 3\}, F = \{\{1, 2\}, \{1, 4\}, \{1, 5\}, \{2, 4\}, \{2, 5\}, \{4, 5\}, \{4, 6\}, \{3, 5\}\}$

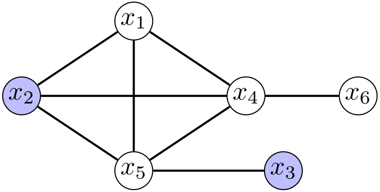

Examples

Wasserstein barycenter

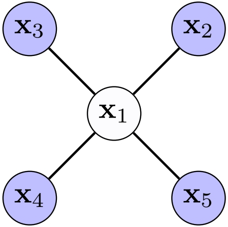

Wasserstein geodesic

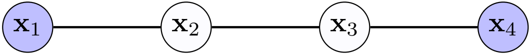

Wasserstein spline

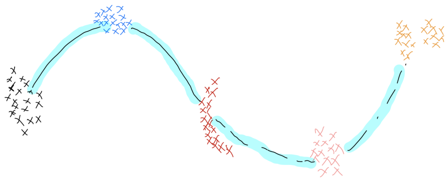

Cubic spline in Wasserstein space

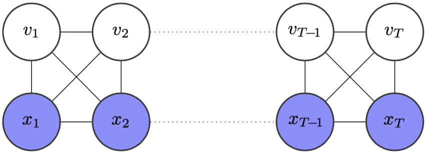

Graphical structure for density control

Cost at each iteration

$$
\begin{aligned}
\tilde{\mathbf{C}}_k &= \nabla F(\mathbf{M}_k) - \frac{1}{\eta}\log \mathbf{M}_k \\
&= \mathbf{C}(\mathbf{M}) + \sum_{i=0}^{T-1} E_i(x_i) - \frac{1}{\eta}\log \mathbf{M}_k
\end{aligned}
$$

with

$$
C(\mathbf{M})(x_0, x_1, \cdots, x_T) = \sum_{i=0}^{T-1} \frac{T}{2} \| x_{i+1} - x_i + \frac{1}{T}[\nabla W * P_i(\mathbf{M})](x_i) \|^2
$$

How to utilize the graphical structure of the cost?

Sinkhorn algorithm

MOT

$$\min_{\mathbf{M}\in\mathbb{R}^{d_1\times\cdots\times d_J}} \langle \mathbf{C}, \mathbf{M}\rangle + \epsilon\langle \mathbf{M}, \log \mathbf{M}\rangle$$
$$\text{subject to } P_j(\mathbf{M}) = \boldsymbol{\mu}_j, \text{ for } j \in \Gamma$$

Optimal solution $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$

$$\begin{aligned}
\mathbf{K} &= \exp(-\mathbf{C}/\epsilon) \\
\mathbf{U} &= \mathbf{u}_1 \otimes \mathbf{u}_2 \otimes \cdots \otimes \mathbf{u}_J \\
\mathbf{u}_j &= \begin{cases} \exp\left(-\frac{1}{J} - \frac{\boldsymbol{\lambda}_j}{\epsilon}\right), & \text{if } j \in \Gamma \\ \exp\left(-\frac{1}{J}\right)\mathbf{1}, & \text{otherwise} \end{cases}
\end{aligned}$$

Sinkhorn algorithm as dual block ascent

Dual problem

$$\max_{\{\boldsymbol{\lambda}_j, j \in \Gamma\}} -\epsilon \langle \mathbf{K}, \mathbf{U} \rangle - \sum_{j \in \Gamma} \boldsymbol{\lambda}_j^T \boldsymbol{\mu}_j$$

Sinkhorn algorithm

1. Compute $\mathbf{U} = \mathbf{u}_1 \otimes \mathbf{u}_2 \otimes \cdots \otimes \mathbf{u}_J$
2. Update $\mathbf{u}_j$ as $\mathbf{u}_j \leftarrow \mathbf{u}_j \odot \boldsymbol{\mu}_j ./ P_j(\mathbf{K} \odot \mathbf{U})$

Advantages

- global linear convergent
- easily parallelizable

bottleneck: projection $P_j(\mathbf{K} \odot \mathbf{U})$ ($O(N_x^J)$ complexity)

Bayesian inference

Bayesian marginal inference

$$p(x_j) = \sum_{x_1,\ldots,x_{j-1},x_{j+1},\ldots,x_J} p(\mathbf{x})$$

- Brutal force computation is expensive
- There are efficient ways to compute the marginals such as variable elimination and belief propagation

Projection as Bayesian inference

Bottleneck: Projection $P_j(\mathbf{K} \odot \mathbf{U})$

$$\mathbf{K} = \exp(-\mathbf{C}/\epsilon), \quad \mathbf{U} = \mathbf{u}_1 \otimes \mathbf{u}_2 \otimes \cdots \otimes \mathbf{u}_J$$

Probabilistic graphical model

$$\begin{aligned} p(\mathbf{x}) &= \prod_{\alpha \in F} \exp(-C_\alpha/\epsilon)(\mathbf{x}_\alpha) \prod_{j \in V} u_j(\mathbf{x}_j) \\ &= \frac{1}{Z} \prod_{\alpha \in F} \psi_\alpha(\mathbf{x}_\alpha) \prod_{j \in V} \phi_j(x_j) \end{aligned}$$

$$\boxed{P_j(\mathbf{K} \odot \mathbf{U}) = p(x_j)}$$

Sinkhorn belief propagation

Apply belief propagation to compute $P_j(\mathbf{K} \odot \mathbf{U}) = p(x_j)$

Sinkhorn belief propagation

1. update
$$m_{i\to j}(x_j) \propto \sum_{x_i} \psi_{ij}(x_i, x_j) \frac{\mu_i(x_i)}{m_{j\to i}(x_i)}$$
2. update all the messages on the path from i to i_{next} with
$$m_{i\to j}(x_j) \propto \sum_{x_i} \psi_{ij}(x_i, x_j) \prod_{k\in N(i)\setminus j} m_{k\to i}(x_i)$$
3. $i = i_{\text{next}} \in \Gamma$

Proximal Sinkhorn belief propagation algorithm

For $k = 1, 2, 3, \ldots$

1. Compute $\tilde{\mathbf{C}}_k = \mathbf{C}(\mathbf{M}_k) + \mathsf{E}(\mathbf{M}_k) - \frac{1}{\eta} \log \mathbf{M}_k$
2. Solve

$$\mathbf{M}_{k+1} = \operatorname{argmin}_{\mathbf{M} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})} \langle \tilde{\mathbf{C}}_k, \mathbf{M} \rangle + (\epsilon + \frac{1}{\eta}) \langle \mathbf{M}, \log \mathbf{M} \rangle$$

   using the Sinkhorn Belief Propagation algorithm

Complexity of each iteration $\mathcal{O}(N_x^2 T)$

Extensions

- more general cost

$$\int_0^1 \int_{\mathbb{R}^d} [\frac{1}{2}\|\xi_t(x)\|^2 + V(x)]\rho_t(x) dx dt$$

- multiple species $\ell = 1, \ldots, L$

$$dX^i_{\ell,t} = -\frac{1}{N}\sum_{m=1}^{L}\sum_{j=1}^{N} \nabla W_{\ell m}(X^i_{\ell,t} - X^j_{m,t})dt + b_\ell(X^i_{\ell,t})dt$$
$$+\sigma(u^i_{\ell,t}dt + \sqrt{\epsilon}dB^i_{\ell,t}), \quad i = 1, \ldots, N$$

Numerical example

Dynamics

$$dX_t^i = -\frac{1}{N}\sum_{j=1}^{N} \nabla W(X_t^i - X_t^j)dt + u_t^i dt + \sqrt{\epsilon}dB_t^i,\ i = 1, \ldots, N$$

$$W(x) = \frac{\beta}{|x|^\alpha + \delta}, \quad \epsilon = 0.1$$

marginals

$$\mu = \mathcal{N}(-0.4, 0.2)$$

$$\nu = \mathcal{N}(0.4, 0.2)$$

Numerical example

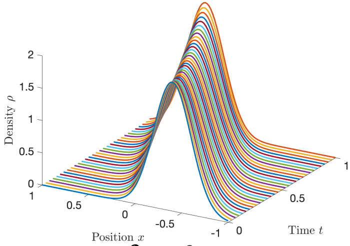

$\beta = 0$

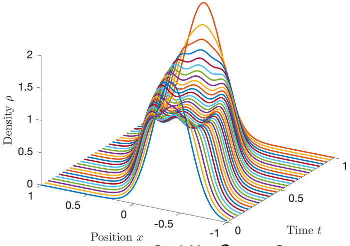

$\alpha = 0.15, \beta = 2$

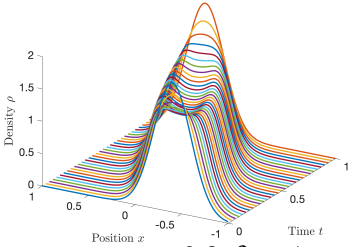

$\alpha = 0.2, \beta = 1$

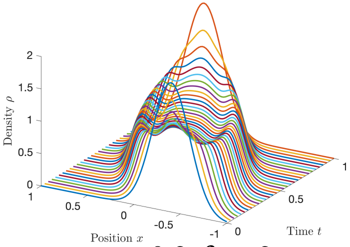

$\alpha = 0.2, \beta = 2$

Future directions

- Particle-based methods
- Neural network based algorithm, gradient flow
- Connections to consensus based sampling

References

1. Backhoff et al, The mean field Schrödinger problem: ergodic behavior, entropy estimates and functional inequalities
2. Chen, Density Control of Interacting Agent Systems
3. Singh et al, Multi-marginal Optimal Transport and Probabilistic Graphical Models
4. Haasler et al, Multi-marginal Optimal Transport with a Tree-structured Cost and the Schrödinger Bridges Problem

Thank you for the attention!