12/2/2025 AMPD UP Presentation

I. Mean field games for generative modeling

A. Review: Mean field games

1. we have a continuum of players who get to pick their velocity v and move according to

$$dX_t = v(X_t)\,dt + \sigma(X_t)\,dB_t$$

a. player at x,t wants to pick v to optimize

$$J(x,t) = E^{x,t}\left(M(X_T, \rho(\cdot,T)) + \int_t^T L(X_s, v(X_s)) + I(X_s, \rho(\cdot,s))\,ds\right)$$

2. trick: introduce

$$u(x,t) = \min_v J(x,t).$$

3. we used dynamic programming to deduce

Hamilton-Jacobi-Bellman
$$\begin{cases} -\partial_t u + H(x, \nabla u) - \frac{1}{2}\sum_{ij} \partial_{x_i x_j} u\,(\sigma\sigma^T)_{ij} = I(x, \rho(\cdot,t)) \\ u(x,T) = M(x, \rho(\cdot,T)) \end{cases}$$

Forward Kolmogorov/Fokker-Planck
$$\begin{cases} \partial_t \rho = -\nabla\cdot(\rho v^*) + \frac{1}{2}\sum_{ij} \partial_{x_i x_j}(\rho(\sigma\sigma^T)_{ij}) \\ \rho(x,0) = \rho_0(x) \end{cases}$$

where

a. $H(x,p) = \max_v \{-v\cdot p - L(x,v)\}$

b. $v^*(x,t) = -\nabla_p H(x, \nabla u(x,t))$ is the optimal velocity

4. plot twist: "This isn't even my final form."

a. $I(x, \rho_t) = \frac{\delta \mathcal{I}}{\delta \rho}[\rho_t](x)$

b. $M(x, \rho_T) = \frac{\delta \mathcal{M}}{\delta \rho}[\rho_T](x)$

5. real goal:

a. let $\mathcal{J}(\rho, v) = \mathcal{M}(\rho_T) + \int_0^T \mathcal{I}(\rho_s) + \int_{\mathbb{R}^d} L(x,v)\,\rho(x,s)\,dx\,ds.$

b. $\min_{\rho, v} \mathcal{J}$ s.t. $\partial_t \rho = -\nabla\cdot(\rho v) + \frac{1}{2}\sum_{ij} \partial_{x_i x_j}(\rho(\sigma\sigma^T)_{ij})$

B. Background: Calculus of variations

1. a functional is a map sending a function to a number

2. to identify minima of a functional, we perturb its input

$$\mathcal{F}(f + \delta f)$$

a. we hope to expand this as

$$\mathcal{F}[f+\delta f] = \mathcal{F}[f] + \underbrace{\delta\mathcal{F}[f](\delta f)}_{\text{linear fnal}} + \underbrace{\delta^2\mathcal{F}[f](\delta f)}_{\text{quadratic fnal}} + \cdots$$

I B2 b. if we can find a fxn g such that
$$\delta\mathcal{F}[f](\delta f) = \int \delta f\, g\, dx \quad \forall \delta f,$$
then we call g $\mathcal{F}$'s "variational derivative" and write $g = \frac{\delta \mathcal{F}}{\delta f}[f]$.

3. alternatively:
$$\delta\mathcal{F}[f](\delta f) = \lim_{h\to 0} \frac{\mathcal{F}[f+h\delta f] - \mathcal{F}[f]}{h}$$

4. example: classical mechanics

a. let $L(x,v) = \frac{m}{2}v^2 - U(x)$ for some function U

b. consider the "action" functional
$$I(x) = \int_0^t L(x,\dot{x})\, ds$$

c. Then if $\delta x(0) = 0 = \delta x(t)$,
$$I(x+\delta x) = \int_0^t L(x,\dot{x}) + \partial_x L\, \delta x + \partial_v L(\delta\dot{x}) + O(\delta x^2)\, ds$$
$$= I(x) + \int_0^t \partial_x L\, \delta x - \frac{d}{dt}\partial_v L\, \delta x\, ds + O(\delta x^2)$$

d. we conclude that
$$\delta I[x](\delta x) = \int_0^t \left(\partial_x L - \frac{d}{dt}\partial_v L\right)\delta x\, ds$$
and
$$\frac{\delta I}{\delta x}[x] = \partial_x L - \frac{d}{dt}\partial_v L$$
$$= -\partial_x U - \ddot{x} m$$

e. Postulate: motion minimizes I

f. then $0 = \frac{\delta I}{\delta x}[x]$, so $m\ddot{x} = -\partial_x U$

g. This is Newton's 2nd law for $F = -\partial_x U$.

h. Also: let $p = \partial_v L = mv$

i. note: we can solve $p = \partial_v L(x,v)$ for v, so may think of v as fxn of x and p:
$$v = v(x,p) \quad (= \tfrac{p}{m})$$

j. so introduce
$$H(x,p) = p\, v(x,p) - L(x, v(x,p)) \quad \left(= \tfrac{1}{2m}p^2 + U(x)\right)$$

i. then $\partial_x H = p\,\partial_x v - \partial_x L - \partial_v L\, \partial_x v$
$$= -\partial_x L = -\frac{d}{dt}\partial_v L = -\dot{p}$$

since we take $L(x,\dot{x})$

$$\partial_p H = v(x,p) + p\,\partial_p v - \partial_v L\, \partial_p v = v(x,p) = \dot{x}$$

k. i. so we also have the equations of motion

$$\begin{cases} \dot{x} = \partial_p H = \frac{p}{m} \\ \dot{p} = -\partial_x H = -\partial_x U \end{cases}$$

l. Finally, we claim

$$H(x,p) = \max_v \{ pv - L(x,v) \}$$

i. if such exists, must have

$$p = \partial_v L(x, v^*)$$

for the maximizing v^*

C. Background: Lagrange multipliers

1. consider

$$\min_{x \in \mathbb{R}^d} F(x) \quad \text{s.t. } G(x) = 0$$

2. envision $G(x)=0$ as a surface in $\mathbb{R}^d$, x is a particle moving on this surface following $-\nabla F$ (trying to minimize F):

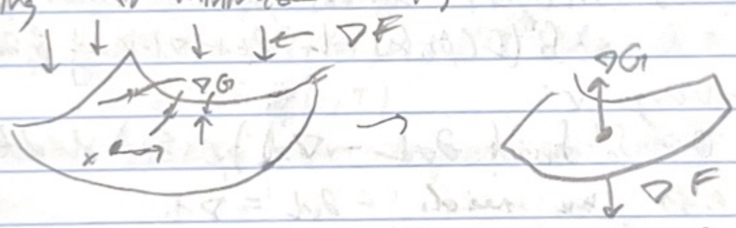

a. x will move until "normal force" from surface cancels out potential force from F:

$$\nabla F = \lambda \nabla G \quad \text{for appropriate } G$$

3. thus, our equation is

$$\nabla (F - \lambda G) = 0$$

a. or

$$\min_x \left(F(x) - \lambda G(x) \right)$$

4. For multiple constraints, get

$$\min_x \{ F(x) + \lambda_1 G_1(x) + \dots + \lambda_N G_N(x) \}$$

5. now, formally say "$N \to \infty$", have constraint $G(x,y) = 0$ $\forall y$, then have

$$\min_x \{ F(x) + \int \lambda(y) G(x,y) \, dy \}$$

6. Next, say "$d \to \infty$", x is replaced by function f, F by functional, and G by differential operator L

a. i.e., consider

$\min_f \mathcal{F}[f] \quad$ s.t. $L(f)(x) = 0$

6b. The Lagrange approach becomes

$$\min_f \left\{ \mathcal{F}[f] + \int \lambda(x) L[f](x)\, dx \right\}$$

7. More rigorous: min-max justification

D. Potential mean field games

1. recall want to solve

$$\min_{\rho, v} \mathcal{J}(\rho, v) = \min_{\rho, v} \left\{ \mathcal{M}(\rho_T) + \int_0^T \mathcal{I}(\rho_t) + \int_{\mathbb{R}^d} L(x,v)\rho(x,t)\, dx\, dt \right\}$$

$$\text{s.t.} \quad \partial_t \rho = -\nabla \cdot (\rho v) + \frac{1}{2} \sum_{ij} \partial_{x_i}\partial_{x_j}(\rho(\sigma\sigma^T)_{ij})$$

2. Lagrange approach:

$$\min_{\rho, v} \left\{ \mathcal{J} + \underbrace{\int_0^T \int_{\mathbb{R}^d} \lambda(x,t) \left(\partial_t \rho + \nabla \cdot (\rho v) - \frac{1}{2} \sum_{ij} \partial_{x_i x_j}(\rho(\sigma\sigma^T)_{ij})\right) dx\, dt}_{} \right\}$$

assuming ρ decays "quickly" in x

$$\longrightarrow \int_{\mathbb{R}^d} (\lambda\rho)\Big|_0^T dx + \int_0^T \int_{\mathbb{R}^d} \left(-\partial_t \lambda - \nabla\lambda \cdot v - \frac{1}{2}\sum_{ij} \partial_{x_i x_j} \lambda (\sigma\sigma^T)_{ij}\right) \rho\, dx\, dt$$

3. vary ρ:

$$0 = \int (M'(x, \rho_T) + \lambda(x,T))\, \delta\rho(T) - \lambda(x,0)\, \delta\rho(0)$$
$$+ \int_0^T \left(I(\rho_t; x) + L - \partial_t \lambda - \nabla\lambda \cdot v - \frac{1}{2}\sum_{ij} \partial_{x_i x_j} \lambda (\sigma\sigma^T)_{ij}\right) \delta\rho(t)\, dx\, dt$$

4. vary v:

$$0 = \int_0^T \int_{\mathbb{R}^d} (\partial_v L - \nabla\lambda)\, \rho\, \delta v\, dx\, dt$$

a. so we need $\partial_v L = \nabla \lambda$

b. recall our Hamiltonian

$$H(x,p) = \max_v \{ -v \cdot p - L(x,v) \}$$

i. if L is nice enough, this is achieved when v, p are s.t.

$$-p = \partial_v L(x, v^*)$$

c. so "if all goes as planned,"

$$-\nabla\lambda \cdot v + L = -H(x, -\nabla\lambda)$$

5. so from varying ρ, we have

$$\partial_t \lambda + H(x, -\nabla\lambda) + \frac{1}{2}\sum_{ij} \partial_{x_i x_j} \lambda (\sigma\sigma^T)_{ij} = I(\rho_t, x)$$

a. also, $\delta\rho(0) = 0$ since don't want to disturb I.C.

b. but $\delta\rho(T)$ forces

$$\lambda(x,T) = -M(x, \rho_T)$$

6. we conclude $\lambda = -u$ from before

5. Q.: Why is soln to $\min_{v} J(v, t)$ a (reg.) Lagrange multiplier for soln to $\min_{v} J$?

E. Generative modeling: reverse diffusion

1. the goal
   a. we have data points sampled from some unknown distribution π
   b. we want to produce a new point that "looks like" it was sampled from π
2. strategy
   a. let data points evolve per SDE

note to self: write densities:

$$\begin{cases} dY(s) = -f(Y(s), T-s)\,ds + \sigma(T-s)\,dW(s) \\ Y(0) \sim \pi \end{cases}$$

   b. ~~then the process~~ write Y's law as $Y(t) \sim \eta(\cdot, t)$
   c. then the process

$$\begin{cases} dX(t) = (f(X(t), t) + \sigma(t)^2 \nabla \log \eta(X, T-t))\,dt + \sigma(t)\,dW(t) \\ X(0) \sim \eta(\cdot, T) \end{cases}$$

   will "reverse" the trajectories of Y
3. ex.
   a. $f(Y, t) = Y$, $\sigma(t) = \sqrt{2}$
   b. then

$$dY = -Y\,ds + \sqrt{2}\,dW$$

   is the Ornstein Uhlenbeck process
   c. Y converges towards $N(0, 1)$
   d. then

$$dX = (X + 2\nabla \log \eta(X, T-t))\,dt + \sqrt{2}\,dW(t)$$

4. If you can learn $\nabla \log \eta(X, T-t)$, then you can generate new samples

F. Reverse diffusion as a MFG

1. consider MFG with
   a. $L(x, v) = \frac{1}{2}\left|\frac{v-f}{\sigma}\right|^2 - \nabla \cdot f$
   b. $\alpha = 0$
   c. $M(\rho) = -\mathbb{E}_\rho(\log \pi)$ (cross entropy)
2. our Hamiltonian becomes

I F 2 a. $H(x,p) = \max_v \{-p\cdot v - L(x,v)\}$

$= \max_v \{-p\cdot v - \frac{1}{2}\left|\frac{v-f}{\sigma}\right|^2 + \nabla\cdot f\}$

b. this maximum is achieved at

$-p - \frac{v-f}{\sigma}\frac{1}{\sigma} = 0, \quad v = -\sigma^2 p + f$

c. so

$H(x,p) = -p\cdot(-\sigma^2 p + f) - \frac{1}{2}\sigma^2|p|^2 + \nabla\cdot f$

d. and

$\nabla_p H(x,p) = 2\sigma^2 p - f - \sigma^2 p = \sigma^2 p - f$

3. our HJB and FP eq'ns are thus

$$\begin{cases} -\frac{\partial u}{\partial t} + \frac{1}{2}\sigma^2|\nabla u|^2 - \nabla u\cdot f + \nabla\cdot f - \frac{1}{2}\sigma^2\Delta u = 0. \\ u(x,T) = -\log\pi(x) \end{cases}$$

$$\begin{cases} \partial_t\rho = -\nabla\cdot(\rho(-\sigma^2\nabla u + f)) + \frac{1}{2}\sigma^2\Delta\rho \\ \rho(x,0) = \rho_0(x) \end{cases}$$

4. Now, if $\eta(x,T-t) = e^{-u(x,t)}$, or $-\log\eta(T-t) = u(t)$ then

a. $-\frac{\partial_t\eta}{\eta} + \frac{1}{2}\sigma^2\frac{|\nabla\eta(T-t)|^2}{\eta^2} + \frac{\nabla\eta\cdot f}{\eta} + \nabla\cdot f + \frac{1}{2}\sigma^2\frac{\Delta\eta\,\eta - |\nabla\eta|^2}{\eta^2} = 0$

b. so $-\partial_t\eta + \nabla\eta\cdot f + \eta\nabla\cdot f + \frac{1}{2}\sigma^2\Delta\eta = 0$

c. $\partial_t\eta = \nabla\cdot(f\eta) + \frac{1}{2}\sigma^2\Delta\eta$

5. hence η is the law of our forward process Y

6. and then

$\partial_t\rho = -\nabla\cdot(\rho(\sigma^2\nabla\log\eta(T-t) + f)) + \frac{1}{2}\sigma^2\Delta\rho$

a. so if we ~~fro~~ pick $\rho_0(x) = \eta(x,T)$, then ρ is the law of our reverse process X

7. Hence, we can recover reverse diffusion in a totally different way

8. ~~this also~~ suggests an ~~alternative~~ objective fcn:

~~a. take~~ ~~$u = \log\eta(T-t)$~~

~~b. so~~

a. MFG ss

$$\inf_{\rho, v} \left\{ \mathbb{E}_\rho[\log \pi] + \int_0^T \int_{\mathbb{R}^d} \left(\frac{1}{2} \left| \frac{v-f}{\sigma} \right|^2 - \nabla \cdot f \right) \rho \, dx \, dt \right\}$$

b. #8b. for optimal v^*, we have $v^* = -\sigma^2 \nabla U + f$

$$= \sigma^2 \nabla \log \eta + f$$

c. so $\nabla \log \eta = \frac{v^* - f}{\sigma^2}$

d. so try to learn v^* instead of $\nabla \log \eta$ directly

9. note that

$$\partial_t \eta = \nabla \cdot (f \eta) + \frac{1}{2} \sigma^2 \Delta \eta$$

corresponds to MFG

r. also, since optimal $v^* = \sigma^2 \nabla \log \eta + f$,

we can try to learn v^* and use

$$\frac{v^* - f}{\sigma^2}$$

instead of learning $\nabla \log \eta$

a. get a handle on v^* by finding a new MFG for η, ρ.