9/10/2025 AMPD UP Presentation

I. Mean field games for generative modeling

A. Background: stochastic differential equations

1. Brownian motion

a. ODE's: $\dot{x} = F(x)$

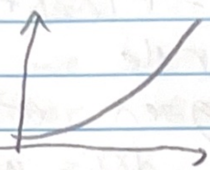

b. in real life, though, looks like

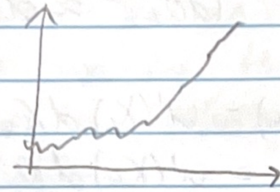

c. Hence, we would like

$$\dot{x} = F(x) + \xi \leftarrow \text{white noise}$$

d. To work w/ "ξ", we introduce Brownian motion: "$\dot{B}_t = \xi$"

i. motion whose path is determined solely by white noise

e. Brownian motion is defined by properties

i. $B(0) = 0$ a.s.

ii. $B(t_1), B(t_2) - B(t_1), \ldots, B(t_n) - B(t_{n-1})$ independent for all $t_1 < \cdots < t_n$

iii. $B(t) - B(s) \sim \mathcal{N}(0, t-s)$ for all $0 \le s \le t$.

f. justification:

i. say $\xi_i = 1$ or -1 w/ prob. $1/2$, $1/2$

ii. consider $S_n = \sum_1^n \xi_i$

iii. jumps occur every Δt and go distance Δx

iv. Then displacement is $X(t) = \Delta x \, S_{t/\Delta t}$

v. and $Var(X(t)) = (\Delta x)^2 \, Var(S_{t/\Delta t}) = (\Delta x)^2 \frac{t}{\Delta t} D^2$ where $D^2 = var(\xi_i)$

vi. to get Brownian motion, need $\Delta x, \Delta t \to 0$. For finite, non zero variance, need $\frac{(\Delta x)^2}{\Delta t} = \text{const.}$

g. We then may define

$$\int \sigma \, dB_t$$

i. w/ Riemann sums

h. and we model white noise with

$$dX_t = F(X_t) + \sigma \, dB_t$$

2. Itô's rule

a. ~~treat~~ given ~~a~~

$$dX_t = \mu \, dt + \sigma \, dB_t,$$

and a fcn u, what is $du(X_t)$?

b. Trick: $(dB_t)^2 \approx dt$

c. so

$$du(X_t) = u'(X_t)\, dX_t + \tfrac{1}{2} u''(X_t)\,(dX_t)^2 + \cdots$$

$$= u'(X_t)\, dX_t + \tfrac{1}{2} u''(X_t)\left(\mu^2 dt^2 + \mu\sigma \, dt \, dB_t + \sigma^2 (dB_t)^2\right) + \cdots$$

more generally:

$du(X_t)$

$= \nabla u \cdot dX$

$+ \frac{1}{2} \frac{\partial}{\partial x_i \partial x_j} u \, (\sigma\sigma^T)_{ij} \, dt$

d. only need 1st order terms:

$$du(X_t) = u'(X_t)\, dX_t + \tfrac{1}{2} u''(X_t)\, \sigma^2 \, dt$$

3. Fokker Planck / forward Kolmogorov / continuity equation

a. What equation does X_t's law $\rho(x,t)$ obey?

b. strategy: weak formulation

c. For any test fcn φ,

$$\frac{d}{dt} E(\varphi(X_t)) = E(d\varphi(X_t))$$

$$= E(\varphi' \, dx + \tfrac{1}{2}\varphi'' \sigma^2 \, dt)$$

$$= E(\varphi'(\mu \, dt + \sigma \, dB_t) + \tfrac{\varphi''}{2}\sigma^2 \, dt)$$

$$\frac{d}{dt}\int \varphi(x)\rho(x,t)\,dx = \int (\varphi'\mu + \tfrac{\varphi''}{2}\sigma^2)\,\rho(x,t)\,dx$$

$$= \int \varphi(x)\,\partial_t \rho \, dx = \int (-\partial_x(\mu\rho) + \tfrac{1}{2}\partial_x^2(\sigma^2\rho))\,\varphi \, dx$$

d. conclusion:

$$\partial_t \rho = -\partial_x(\mu\rho) + \tfrac{1}{2}\partial_x^2(\sigma^2\rho)$$

i. general:

$$\partial_t \rho = -\nabla \cdot (\mu\rho) + \tfrac{1}{2}\sum \frac{\partial^2}{\partial x_i \partial x_j}((\sigma\sigma^T)_{ij}\,\rho)$$

B. Mean field games

1. We have players whose states evolve according to
$$dX_t = v(X_t,t)\,dt + \sigma\,dB_t$$
a. σ is given
b. we can pick v
2. we assume we have coly many players whose densities are ~~distributed~~ given by $\rho(x,t)$ at time t
a. by Fokker planck, we have
$$\partial_t \rho = -\nabla\cdot(v\rho) + \tfrac{1}{2}\sum \partial_{x_i x_j}((\sigma\sigma^T)_{ij}\,\rho)$$
3. at state x and time t, player wants to pick $v(x,t)$ to minimize

terminal cost
$$J_{x,t}(v) = E^{x,t}\big(M(X(T),\rho(\cdot,T)) + \int_t^T L(X_s, v(X_s,s)) + I(X_s,\rho(\cdot,s))\,ds\big)$$
state cost

interaction/competition cost

~~4. Richard Bellman: Dynamic Programming~~
4. introduce
$$U(x,t) = \min_v\{E^{x,t}(\int_t^T L(X_s, v(X_s)) + I(X_s,\rho)\,ds + M(X(T),\rho))\}$$

5. Richard Bellman: Dynamic Programming
a. U satisfies for small dt
$$U(x,t) = \min_{v}\{E^{x,t}(\int_t^{t+dt} L(X_s,v) + I(X_s,\rho)\,ds + U(X_{t+dt}, t+dt))\}$$
~~b. Taylor series~~
6. Ito:
$$dU = \frac{\partial U}{\partial t}dt + \nabla U\cdot dX + \tfrac{1}{2}\sum \partial_{x_i x_j} U\,(\sigma\sigma^T)_{ij}\,dt$$
c. so
$$E(U(X_{t+dt}, t+dt)) = E(U(X_t) + \int_t^{t+dt} \frac{\partial U}{\partial t} + \nabla U\cdot v + \tfrac{1}{2}\sum \partial_{x_i x_j} U\,(\sigma\sigma^T)_{ij}\,dt + 0)$$

183d-50

$$0 = \min_v \{ E^{x,t} (\int_t^{t+dt} \underset{\wedge}{} \frac{\partial u}{\partial t} + \nabla u \cdot v + \tfrac{1}{2} \partial_{x_i} \partial_{x_j} u \, (\sigma\sigma^T)_{ij}) ds \}$$

$L(x,v) + I(x,\rho) +$

e. divide by dt, take $dt \to 0$:

$$0 = \min_v \{ E^{x,t} (\underset{\uparrow}{} \frac{\partial u}{\partial t} + \nabla u \cdot v + \tfrac{1}{2} \partial_{x_i} \partial_{x_j} u \, (\sigma\sigma^T)_{ij} \}$$

$L(x,v) + I(x,\rho) +$

f. all the X_t's are replaced by x, E goes away

g. then rearrange:

this can be dropped since only looking at $v(x,t)$ at one specific (x,t)!

$$-\frac{\partial u}{\partial t} + \max_v \{ -\nabla u \cdot v - L(x,v) \}$$
$$- \tfrac{1}{2} \partial_{x_i x_j} u \, (\sigma\sigma^T)_{ij} = I(x,\rho)$$

h. also,

$$u(x,T) = \min_v \{ 0 + M(T,\rho) \} = M(T,\rho)$$

i. let

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

6. Our Hamilton-Jacobi-Bellman equations are then:

$$\text{HJB} \begin{cases} -\frac{\partial u}{\partial t} + H(x, \nabla u) - \tfrac{1}{2} \sum \partial_{x_i x_j} u \, (\sigma\sigma^T)_{ij} = I(x,\rho) \\ u(x,T) = M(T,\rho) \end{cases}$$

7. we need a way to recover optimal v

a. for H to have max, we expect
$$-p - \nabla_v L = 0$$

b. or $p = -\nabla_v L(x, v^*)$

c. <u>assume</u> this can be solved for
$$v^* = v^*(x,p)$$

d. So
$$H = -p \cdot v^* - L(x, v^*)$$

e. we note
$$\nabla_p H = -v^* - p \cdot \nabla_p v^* - \nabla_v L(x,v^*) \cdot \nabla_p v^*$$
$$= -v^* + \nabla_v L \cdot \nabla_p v^* - \nabla_v L \cdot \nabla_p v^*$$
$$= -v^*$$

f. Therefore,

$v^*(x,t) = -\nabla_p H(x, \nabla u(x,t))$

† 9. So our full sys. of eq'ns is

HJB $\begin{cases} -\frac{\partial u}{\partial t} + H(x, \nabla u) - \frac{1}{2}\sum \partial_{x_i x_j} u \, (\sigma\sigma^T)_{ij} = I(x, \rho_t) \\ u(x,T) = M(T, \rho_T) \end{cases}$

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

C. Background: calculus of variations

1. say $\mathcal{F}$ maps fons to real numbers
2. consider

$\delta\mathcal{F} = \mathcal{F}[f + \delta f] - \mathcal{F}[f]$

3. our hope is we can rewrite this in the form

$$\int g \, \delta f \, dx + O(\delta f^2)$$

a. then we call g $\mathcal{F}$'s variational derivative at f and we write

$\frac{\delta \mathcal{F}}{\delta f}[f](x) = g(x)$

4. Alt: $\frac{\delta \mathcal{F}}{\delta f}[f](x)$ is the fon such that for any δf,

$$\lim_{h \to 0} \frac{\mathcal{F}[f + h\delta f] - \mathcal{F}[f]}{h} = \int \frac{\delta \mathcal{F}}{\delta f}[f] \, \delta f \, dx.$$

a. in L^2, exists by Riesz representation

5. to minimize $\mathcal{F}$, we need

$\frac{\delta \mathcal{F}}{\delta f} = 0$

D. Background: Lagrange multipliers

1. say we want to minimize $F(x)$ such that $G(x) = 0$ holds
2. think of $G(x) = 0$ as a surface and x as a particle following field $-\nabla F(x)$ on this surface

a. will keep moving until normal force of surface and field force cancel:

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

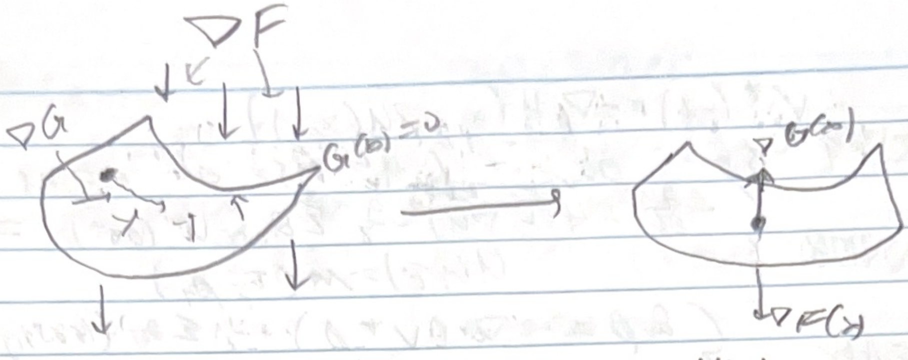

so λ called a Lagrange multiplier

3. What if we have a vector constraint?
   a. $\nabla F(x) + \lambda_1 \nabla G_1(x) + \cdots + \lambda_n \nabla G_n(x) = 0$
   b. or = min: $F(x) + \lambda \cdot G(x)$

4. Now, say we want to minimize $\mathcal{F}[f]$ subject to $L[f]=0$ - differential operator

~~5. we use a Lagrange multiplier function - that enforces the "~~

5. now, our vector constraint $\vec{\lambda}$ becomes a fn as $n \to \infty$ = $\lambda(x)$

6. so we try to minimize

$\lambda \cdot G$ goes to integral $\longrightarrow$ $\mathcal{F}[f] + \int \lambda L[f]\, dx$

7. Alt: min max

$$\min_f \{\mathcal{F}[f] \text{ s.t. } L[f]=0\}$$
$$= \min_f \max_\lambda \{\mathcal{F}[f] + \int \lambda L[f] \,dx\}$$

E. Potential mean field games

1. say $F = \frac{\delta \mathcal{F}}{\delta \rho}$, $M = \frac{\delta \mathcal{M}}{\delta \rho}$

2. define "real" goal

$$J(v,\rho) = \mathcal{M}(\rho(\cdot,T)) + \int_0^T \mathcal{F}(\rho(\cdot,t)) + \int\int_{\mathbb{R}^d} L(x, v(x,t))\,\rho(x,t)\,dx\,dt$$

   a. minimize this for ρ s.t.

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (v\rho) + \frac{1}{2} \sum \partial_{x_i x_j}((\sigma\sigma^T)_{ij}\,\rho)$$

3. Lagrange multiplier?

$$\int\int \lambda\left(\partial_t \rho + \nabla \cdot (v\rho) - \frac{1}{2} \sum \partial_{x_i x_j}((\sigma\sigma^T)_{ij}\,\rho)\right) dx\,dt$$
$$= \int\int \left(-\partial_t \lambda - \nabla\lambda \cdot v - \frac{1}{2} \sum \partial_{x_i x_j} \lambda\, (\sigma\sigma^T)_{ij}\right) \rho \,dx\,dt$$
$$+ \int \lambda(T)\rho(T)\,dx - \int \lambda(0)\rho(0)\,dx = \Lambda(\lambda, \rho, v)$$

4. variation:

with $\delta\rho(x,0)=0$ (don't mess up I.C.)

$$J(\rho+\delta\rho)+\Lambda(\rho+\delta\rho)-J(\rho)-\Lambda(\rho)$$
$$= \int \frac{\delta m}{\delta\rho}[\rho(T)]\,\delta\rho(T)\,dx$$
$$+ \int_0^T \int \frac{\delta cl}{\delta\rho}[\rho(t)]\,\delta\rho(t)\,dx$$
$$+ L\,\delta\rho\,dx\,dt$$
$$+ \iint \left(-\partial_t\lambda - \nabla\lambda\cdot v - \tfrac{1}{2}\sum \partial_{x_i}\partial_{x_j}\lambda\,(\sigma\sigma^T)_{ij}\right)\delta\rho\,dx\,dt$$
$$+ \int \lambda(T)\,\delta\rho(T)\,dx$$

5. So we require
$$\frac{\delta cl}{\delta\rho}[\rho]+L-\partial_t\lambda-\nabla\lambda\cdot v-\tfrac{1}{2}\sum\partial_{x_i}\partial_{x_j}\lambda\,(\sigma\sigma^T)_{ij}=0$$
$$\lambda(T)=-\frac{\delta m}{\delta\rho}[\rho(T)]$$

6. next, ~~note that~~
$$\cancel{H L(x,v)-\nabla\lambda\cdot v = H(x,-\nabla\lambda)}$$
vary v:
$$\iint \nabla_v L\cdot\delta v\,\rho\,dx\,dt-\iint\nabla\lambda\cdot\delta v\,\rho\,dx\,dt=0$$

7. conclude
$$\nabla_v L=\nabla\lambda$$

8. Finally, consider same H as before.
a. recall we found optimality cond
$$p=-\nabla_v L$$
b. we conclude $H(x,-\nabla\lambda)=\nabla\lambda\cdot v^*-L(x,v^*)$

9. Finally, rearrange
$$\frac{\delta cl}{\delta\rho}[\rho]=-L+\partial_t\lambda+\nabla\lambda\cdot v+\tfrac{1}{2}\sum\partial_{x_i}\partial_{x_j}\lambda\,(\sigma\sigma^T)_{ij}$$
$$=\partial_t\lambda+H(x,-\nabla\lambda)+\tfrac{1}{2}\sum\partial_{x_i}\partial_{x_j}\lambda\,(\sigma\sigma^T)_{ij}$$

10. we conclude that
$$\lambda=-u$$
from before

11. we've recovered the original eq'ns

12. Q: What's going on here?

a. Why does optimizing a problem with $\frac{\delta cl}{\delta\rho}$ ~~obst~~ optim and $\mathcal{E}^{int}$ also optimize a problem with cl?

F. The paper

1. goal: generative modeling

1 & 2. There are many ways to do this w/ MFG

3. Table:

Method	$\mathcal{M}(\rho)$	$\ell(\rho)$	$L(x,v)$
Normalizing flow	$E_\rho(\log\rho - \log\rho_{ref})$	0	0
Score-based	$-E_\rho \log \pi$	0	$\frac{1}{2}\lvert v\rvert^2 - \nabla\cdot f$
Gradient flows ($a \to 0$)	$\mathcal{F}(\rho)e^{-T/\varepsilon}$	$\frac{\varepsilon}{2}e^{-T/\varepsilon}\mathcal{F}(\rho)$	$\frac{1}{2}\lvert v\rvert^2 e^{-t/\varepsilon}$

Method	Dynamics	$H(x,p)$	Optimal v^*
Normalizing flow	$dx = v\,dt$	$\sup_{v} -p^T v$	$-\nabla_p H(x,\nabla U)$
Score-based	$dx = (f+\sigma v)\,dt + \sigma\,dW_t$	$-f^T p + \frac{\lvert\sigma p\rvert^2}{2} + \nabla\cdot f$	$\sigma\nabla\log\eta_t$
Gradient flow ($a \to 0$)	$dx = v\,dt$	$\frac{1}{2}e^{t/\varepsilon}\lvert p\rvert^2$	$-\frac{\delta\mathcal{F}}{\delta\rho}$

6. Score-based

1. likelihood estimation

a. empirical likelihood of π

$$\prod_i \pi(X=x_i) = \prod_{x_i}\pi(X=x_i)^{\#\text{ times } x_i \text{ occurs}}$$
$$= \prod_{x_i}\pi(X=x_i)^{N\cdot p(X=x_i)}$$

b. take log, divide by N — want to maximize
$\sum p(X=x_i)\log\pi(X=x_i)$

c. continuous version:

maximize $\int p(x)\log\pi(x)\,dx$

$\iff$ minimize $\underbrace{-\int p(x)\log\pi(x)\,dx}_{\text{cross entropy}}$