Density Control, Covariance Steering, and Optimization over Distributions

Yongxin Chen

School of Aerospace Engineering

Georgia Institute of Technology

Joint work with

Tryphon Georgiou Michele Pavon

ACC workshop: Control of Distributions: Theory and Application, May 24, 2021

Control uncertainty

optimally drive a system from one uncertain state to another

- missile guidance
- spacecraft landing
- motion planning
- ...

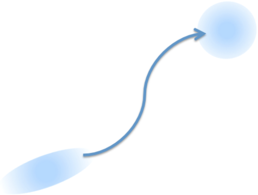

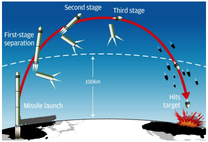

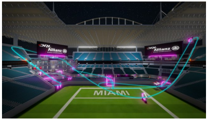

Control distribution

control the formation of a collection of systems

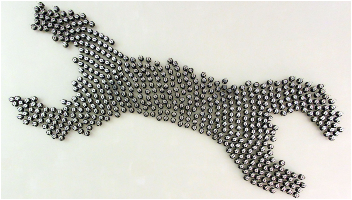

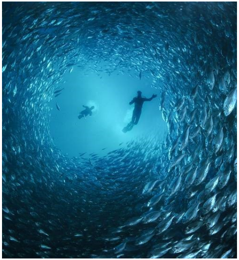

Density control

$$dX_t = f(t, X_t)dt + g(t, X_t)(u_t dt + \sqrt{\epsilon} dW_t)$$

$$X_0 \sim \rho_0$$

Find u causal, finite-energy control such that the cost

$$J(u) = \mathbb{E}\left\{\int_0^T [\frac{1}{2}|u_t|^2 + V(X_t)]dt\right\}$$

is minimized and

$$X_T \sim \rho_T$$

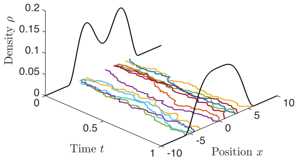

Reformulation as optimization over distributions

Uncontrolled process $\mathcal{P}^0$

$$dX_t = f(t, X_t)dt + \sqrt{\epsilon} g(t, X_t)dW_t, \quad X_0 \sim \rho_0$$

Controlled process $\mathcal{P}^u$

$$dX_t = f(t, X_t)dt + g(t, X_t)(u_t dt + \sqrt{\epsilon} dW_t), \quad X_0 \sim \rho_0$$

KL divergence between $\mathcal{P}^u$ and $\mathcal{P}^0$ in the interval $t \in [0, T]$

$$\mathrm{KL}(\mathcal{P}^u, \mathcal{P}^0) = \tfrac{1}{2\epsilon}\mathbb{E}\left\{\int_0^T |u_t|^2 dt\right\}$$

Optimization over distributions

$$\min_{\mathcal{P}^u \in \Pi(\rho_0, \rho_T)} \int d\mathcal{P}^u \left[\log \frac{d\mathcal{P}^u}{d\mathcal{P}^0} + \frac{1}{\epsilon} V\right]$$

Schrödinger bridges

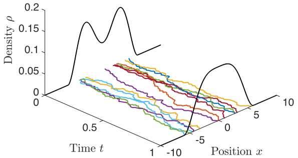

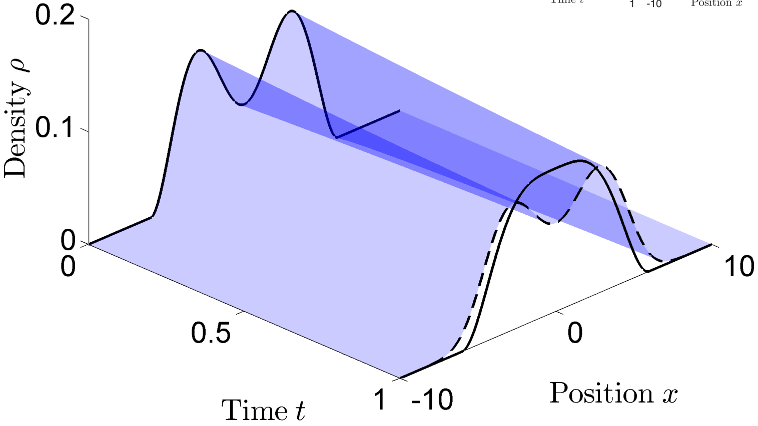

Schrödinger bridges

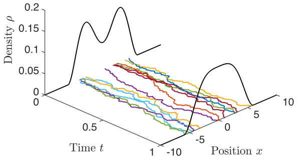

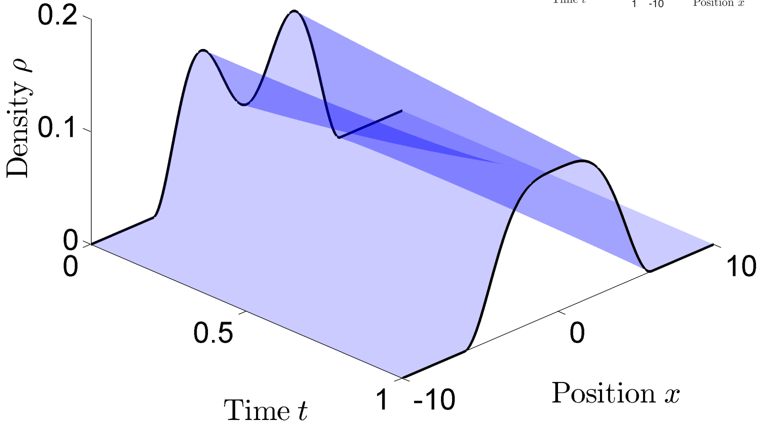

Large deviations interpretation

$\mathcal{Q}$ prior process
$\mathcal{P}$ Schrödinger bridge with marginals ρ_0 and ρ_T

Large deviations (N particles)

$$\mathrm{Prob}(\mathcal{P}) \approx \exp[-N\mathrm{KL}(\mathcal{P}, \mathcal{Q})]$$

Kullback-Leibler divergence of $\mathcal{P}$ with respect to $\mathcal{Q}$

$$\mathrm{KL}(\mathcal{P}, \mathcal{Q}) = \int \log\left(\frac{d\mathcal{P}}{d\mathcal{Q}}\right) d\mathcal{P}$$

Schrödinger bridge $\mathcal{P}$ minimizes $\mathrm{KL}(\mathcal{P}, \mathcal{Q})$
over all the processes with marginal distributions ρ_0 and ρ_T

Föllmer 88

Schrödinger bridges: solution

density interpolation

$$\begin{aligned} \rho(t,x) &= \varphi(t,x)\hat{\varphi}(t,x) \\ d\mathcal{P} &\propto \varphi(0,X_0)^{-1}\varphi(T,X_T)d\mathcal{Q} \end{aligned}$$

two-point boundary value problem ($\mathcal{L}$: generator)

$$\boxed{\begin{gathered} \frac{\partial \varphi}{\partial t}(t,x) = -\mathcal{L}\varphi(t,x) \\ \frac{\partial \hat{\varphi}}{\partial t}(t,x) = \mathcal{L}^\dagger \hat{\varphi}(t,x) \\ \varphi(0,x)\hat{\varphi}(0,x) = \rho_0(x) \\ \varphi(T,x)\hat{\varphi}(T,x) = \rho_T(x) \end{gathered}}$$

Fortet 40, Beurling 60

Jamison 74, Föllmer 88

Chen et al 15

Algorithm based on the Hilbert metric: Sinkhorn iteration

$$\begin{tikzcd}
\hat{\varphi}(0,x_0) \arrow[r, "\mathcal{L}^\dagger"] & \hat{\varphi}(T,x_T) \arrow[d, "\mathcal{D}_{\rho_T}"] \\
\varphi(0,x_0) \arrow[u, "\hat{\mathcal{D}}_{\rho_0}"] & \varphi(T,x_T) \arrow[l, "-\mathcal{L}"']
\end{tikzcd}$$

$$\frac{\partial\hat{\varphi}}{\partial t}(t,x) = \mathcal{L}^\dagger\hat{\varphi}(t,x)$$

$$\varphi(T,x) = \rho_T(x)/\hat{\varphi}(T,x)$$

$$\frac{\partial\varphi}{\partial t}(t,x) = -\mathcal{L}\varphi(t,x)$$

$$\hat{\varphi}(0,x) = \rho_0(x)/\varphi(0,x)$$

Strictly contractive with respect to Hilbert metric

$$d_H(p,q) = \log\frac{M(p,q)}{m(p,q)}$$

$$\begin{aligned}
M(p,q) &:= \inf\{\lambda \mid p \le \lambda q\} \\
m(p,q) &:= \sup\{\lambda \mid \lambda q \le p\}
\end{aligned}$$

Cuturi 13, Chen et al 15

Alternative approach

unconstrained $\tilde{\mathcal{U}} := \{u \mid \text{ causal, finite-energy } \}$

constrained $\mathcal{U} := \{u \mid X_T \sim \rho_T\} \subset \tilde{\mathcal{U}}$

1. Choose h, let $\tilde{J}(u) = \mathbb{E}\left\{\int_0^T [\frac{1}{2}|u_t|^2 + V(X_t)]dt + h(X_T)\right\}$

   $\text{argmin}_{u\in\mathcal{U}} J(u) = \text{argmin}_{u\in\mathcal{U}} \tilde{J}(u)$

2. Compute unconstrained optimal control $u^\star = \text{argmin}_{u\in\tilde{\mathcal{U}}} \tilde{J}(u)$
3. Compute distribution $X_T^\star \sim \rho_T^\star$

Approach: study $h \mapsto \rho_T^\star$

Choose h

If $\rho_T^\star = \rho_T$ then $u^\star \in \mathcal{U}$

It follows $u^\star = \text{argmin}_{u\in\mathcal{U}} \tilde{J}(u)$

Density control: solution

Terminal cost:

$$h(x) = -\epsilon \log \varphi(T, x)$$

Nonlinear state feedback

$$u(t,x) = \epsilon g(t,x)^T \nabla \log \varphi(t,x)$$

$\frac{\partial \varphi}{\partial t} = -\mathcal{L}\varphi$

Controlled process

$$dX_t = f(t, X_t)dt + g(t, X_t)(\epsilon g(t, X_t)^T \nabla \log \varphi(t, X_t)dt + \sqrt{\epsilon} dW_t), \quad X_0 \sim \rho_0$$

The probability density $\rho(t,x)$ of X_t satisfies

$$\rho(t,x) = \varphi(t,x)\hat{\varphi}(t,x)$$

Chen et al 15

Linear covariance steering

$$dX_t = AX_t dt + B(u_t dt + \sqrt{\epsilon} dW_t)$$

$$X_0 \sim N(m_0, \Sigma_0)$$

(A, B) controllable

Find u causal, finite-energy control such that the cost

$$J(u) = \mathbb{E}\left\{\int_0^T \frac{1}{2}|u_t|^2 + \frac{1}{2}X_t^T Q X_t dt\right\}$$

is minimized and

$$X_T \sim N(m_T, \Sigma_T)$$

Linear covariance steering: solution

Mean and covariance can be controlled separately

$$u(t,x) = -B^T \Pi(t) x + d(t)$$

coupled Riccati equations (closed-form available)

$$\begin{aligned}
-\dot{\Pi}(t) &= A^T \Pi(t) + \Pi(t) A - \Pi(t) B B^T \Pi(t) + Q \\
-\dot{\mathrm{H}}(t) &= A^T \mathrm{H}(t) + \mathrm{H}(t) A + \mathrm{H}(t) B B^T \mathrm{H}(t) - Q \\
\epsilon \Sigma_0^{-1} &= \Pi(0) + \mathrm{H}(0) \\
\epsilon \Sigma_T^{-1} &= \Pi(T) + \mathrm{H}(T)
\end{aligned}$$

Chen et al 17

Extensions

- Different input and noise channels: $B \neq B_1$

$$dX_t = AX_t dt + Bu_t dt + B_1 dW_t$$

- Output feedback

$$\begin{aligned} dX_t &= AX_t dt + Bu_t dt + B_1 dW_t \\ dY_t &= CX_t dt + DdV_t \end{aligned}$$

- Stationary case

$$\limsup_{T\to\infty} \mathbb{E}\left\{\frac{1}{T}\int_0^T [\frac{1}{2}|u_t|^2 + \frac{1}{2}X_t^T Q X_t] dt\right\}$$

- Discrete time

$$X_{t+1} = AX_t + Bu_t + B_1 W_t$$

Chen et al, TAC, ACC, CDC 15, 16

Skelton, Beghi, Brockett, Bakolas, Tsiotras, Halder...

Nonlinear covariance steering

$$dX_t = f(t, X_t)dt + g(t, X_t)(u_t dt + \sqrt{\epsilon} dW_t)$$

$$X_0 \sim \rho_0 = N(m_0, \Sigma_0)$$

Find u causal, finite-energy control such that the cost

$$J(u) = \mathbb{E}\left\{ \int_0^T [\frac{1}{2}|u_t|^2 + V(X_t)]dt \right\}$$

is minimized and

$$X_T \sim \rho_T = N(m_T, \Sigma_T)$$

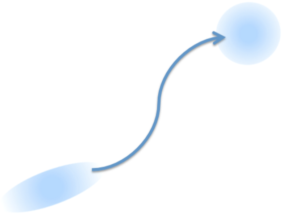

Nonlinear covariance steering

Optimization over distributions

$$\min_{\mathcal{P}^u \in \Pi(\rho_0, \rho_T)} \int d\mathcal{P}^u \left[\log \frac{d\mathcal{P}^u}{d\mathcal{P}^0} + \frac{1}{\epsilon} V\right]$$

local approximation

$$\min_{\mathcal{P}^u \in \hat{\Pi}(\rho_0, \rho_T)} \int [\frac{1}{\epsilon}\hat{V} - \log d\hat{\mathcal{P}}^0] d\mathcal{P}^u + \int d\mathcal{P}^u \log d\mathcal{P}^u$$

$\hat{V}$: quadratic approximation of V along the mean of $\mathcal{P}^u$
$\hat{\mathcal{P}}^0$: linear approximation of the uncontrolled process along
$\hat{\Pi}(\rho_0, \rho_T)$: Gaussian Markov processes in $\Pi(\rho_0, \rho_T)$

Composite optimization

$$\min_{\mathcal{P}^u \in \hat{\Pi}(\rho_0, \rho_1)} F(\mathcal{P}^u) + G(\mathcal{P}^u)$$

Mirror proximal gradient algorithm

Composite optimization (F smooth, G possibly nonsmooth)

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

Proximal gradient

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

Mirror proximal gradient (D Bregman divergence, e.g., Kullback-Leibler divergence)

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

Convergence rate $\mathcal{O}(1/k)$, objective monotonically decreasing

Nonlinear covariance steering

Optimization over distributions

$$\min_{\mathcal{P}^u \in \hat{\Pi}(\rho_0,\rho_1)} F(\mathcal{P}^u) + G(\mathcal{P}^u)$$

$$\begin{aligned} F(\mathcal{P}^u) &= \int [\frac{1}{\epsilon}\hat{V} - \log d\hat{\mathcal{P}}^0] d\mathcal{P}^u \\ G(\mathcal{P}^u) &= \int d\mathcal{P}^u \log d\mathcal{P}^u \end{aligned}$$

Proximal gradient for covariance steering (stepsize η)

$$\mathcal{P}_{k+1} = \text{argmin}_{\mathcal{P}\in\hat{\Pi}(\rho_0,\rho_1)} G(\mathcal{P}) + \frac{1}{\eta}\text{KL}(\mathcal{P},\mathcal{P}_k) + \langle \frac{\delta F}{\delta \mathcal{P}}(\mathcal{P}_k), \mathcal{P}\rangle$$

Each iteration is a linear covariance steering!

Nonlinear covariance steering

Gaussian Markov approximation of $\mathcal{P}_k^u$ (with mean z_t^k)

$$dX_t = A_k(t)X_t dt + a_k(t)dt + \sqrt{\epsilon} g(t, z_t^k)dW_t$$

Gaussian Markov approximation of $\mathcal{P}^0$

$$dX_t = \hat{A}_k(t)X_t dt + \hat{a}_k(t)dt + \sqrt{\epsilon} g(t, z_t^k)dW_t$$

Optimal policy to linear covariance steering

$$K_k(t)X_t + d_k(t)$$

Proximal gradient iteration

$$\begin{aligned} A_{k+1}(t) &= \frac{1}{1+\eta}[\eta A_k(t) + \hat{A}_k(t)] + g(t, z_t^k)K_k(t) \\ a_{k+1}(t) &= \frac{1}{1+\eta}[\eta a_k(t) + \hat{a}_k(t)] + g(t, z_t^k)d_k(t) \end{aligned}$$

Thermodynamical systems: Cooling of oscillators

Gravitational wave detector

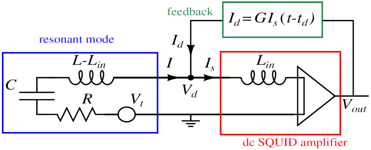

Zendri, Bonaldi,

Conti et al 09

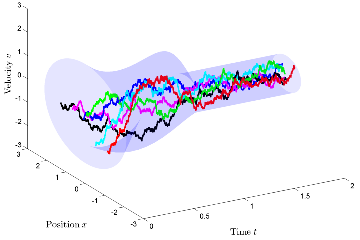

Chen et al 15

Control of collective dynamics

$N = 20000$ agents with dynamics

$$dX_t^i = X_t^i dt + u_t^i dt + dW_t^i$$

marginals $\rho_0 \sim N[1, 4], \quad \rho_1 \sim N[-4, 1]$

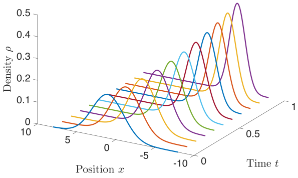

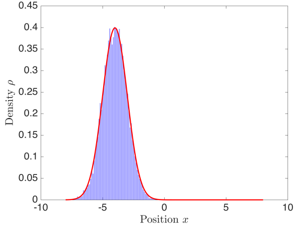

Motion planning for nonlinear systems

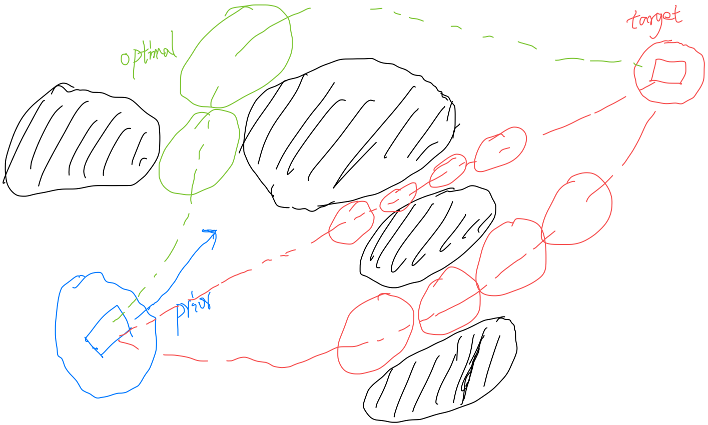

Takeaway:

- Density/covariance control is an optimization over path distributions
- Some linear covariance steering can be solved in closed-form
- Nonlinear covariance steering can be solved iteratively

References

1. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint
2. Optimal steering of a linear stochastic system to a final probability distribution, Part I, II, III
3. Covariance steering for nonlinear control-affine systems
4. Optimal transport in systems and control
5. Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge
6. Controlling uncertainty: Schrödinger's inference method and the optimal steering of probability distributions

Thank you for your attention!