Control, Planning, and Inference for robotic systems under uncertainty

Hongzhe Yu

Ph.D. Candidate, Georgia Tech

February 14, 2025

Outline

- Introduction
  - Motivations and Challenges

- Variational Inference Motion Planning
  - Formulation
  - Method and Results

- Covariance Steering for Hybrid Systems
  - Covariance steering for linear stochastic system with hybrid transitions

- Path Integral Control for Hybrid Systems
  - Nonlinear stochastic systems
  - Method and Results

Uncertainties in Robotic Systems

- Environment
  - Unpredictable scenarios: highways, private homes (service robots)
- Sensors
  - Limited range and resolution
  - Noise
- Robots
  - Control actuation noise
  - Wear-and-tear
- Models
  - Modeling error

Probabilistic Robotics

- Motivation
  - Better performance under the above-mentioned uncertainties

- Conjecture
  - *A robot that carries a notion of its own uncertainty and that acts accordingly is superior to one that does not*

- Downside / Trade-off
  - Computational inefficiency
    - Much higher dimensional problem, much more decision variables
  - Approximate
    - Describe the true uncertainty using distributions with compact parameter set

Tools in probability

- Trajectory / path distribution
  - Distribution induced by Stochastic Differential Equations (SDEs)

$$dX_t = A_t X_t\, dt + \mathrm{a_t dt} + \mathrm{B_t dW_t}$$

$$d\mathbb{P}^0 \approx \hat{p}(\mathbf{X}) \triangleq p(X_1|X_0)\dots p(X_N|X_{N-1})$$

- KL-divergence
  - Measure the distance between two distributions (q, p)

$$\boxed{\mathrm{KL}\left(q \parallel p\right) \triangleq \mathbb{E}_q\left[\log \frac{q}{p}\right]}$$

- Bayes' view of probability
  - Prior, likelihood, and posterior

$$p(\boldsymbol{X}|Z) \propto p(\boldsymbol{X})p(Z|\boldsymbol{X})$$

- Importance Sampling
  - Sample from a difficult-to-sample distribution

$$\mathbb{E}_{\mathbb{P}_0}\left[f\right] = \mathbb{E}_{\mathbb{P}_u}\left[f \times \frac{d\mathbb{P}_0}{d\mathbb{P}_u}\right]$$

Variational Inference Motion Planning

Probability in path distribution space

- Linear SDE

$$dX_t = A_t X_t\, dt + \mathrm{a}_t \mathrm{dt} + \mathrm{B}_t \mathrm{dW}_t$$

Linear dynamics **Gaussian noise**

- Trajectory **prior** over the joint states $\boldsymbol{X}$
  - Linear transformations of Gaussian **remains** a Gaussian

$$p(\boldsymbol{X}) \propto exp\left(-\frac{1}{2}\|\boldsymbol{X}-\boldsymbol{\mu}\|_{\boldsymbol{K}^{-1}}^2\right)$$: **A Gaussian distribution!**

- Planning task **likelihood**:

$$p(Z|\boldsymbol{X}) \propto exp(-V(\boldsymbol{X}))$$ **: A Penalty term**

- Motion Planning **posterior**: $p(\boldsymbol{X}|Z) \propto p(\boldsymbol{X})p(Z|\boldsymbol{X})$
  - $p(\boldsymbol{X}|Z) \propto exp\left(-\frac{1}{2}\|\boldsymbol{X}-\boldsymbol{\mu}\|_{\boldsymbol{K}^{-1}}^2 - V(\boldsymbol{X})\right)$

Gaussian Variational Inference Motion Planning

$$\mathrm{KL}\left(q \parallel p\right) \triangleq \mathbb{E}_q\left[\log \frac{q}{p}\right]$$

- Objective: $min_{q\in\mathcal{Q}} KL\big(q(\boldsymbol{X}) \parallel p(\boldsymbol{X}|Z)\big)$; q($\boldsymbol{X}$): Gaussian

$$p(\boldsymbol{X}|Z) \propto p(\boldsymbol{X})p(Z|\boldsymbol{X})$$

- Expanding:

$$\begin{aligned} q^\star &= \arg\min_{q\in\mathcal{Q}} \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X}|Z)\right) \\ &= \arg\min_{q\in\mathcal{Q}} \mathbb{E}_q[\log q(\mathbf{X}) - \log p(Z|\mathbf{X}) - \log p(\mathbf{X})] \\ &= \arg\max_{q\in\mathcal{Q}} \mathbb{E}_q\left[\log p(Z|\mathbf{X})\right] - \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X})\right) \end{aligned}$$

KL between Gaussians

Expectation of Nonlinear function

- Algorithm
  - Natural Gradient Descent (NGD)

- Complexity
  - Sparse Factor Graph
  - Linear to discretization, polynomial to state dimensions
  - Can be Parallelized

NGD on Sparse Factor Graph

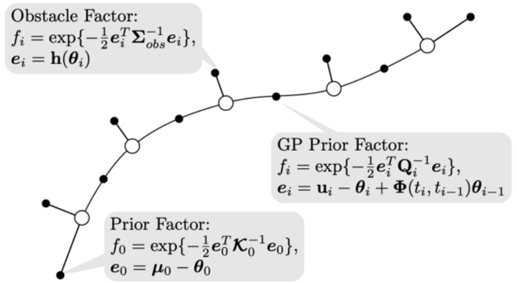

- NGD Gradients

$$\psi(\mathbf{X}) \triangleq -\log p(\mathbf{X}|Z)$$

$$\frac{\partial \mathcal{J}(q)}{\partial \mu_\theta} = \Sigma_\theta^{-1}\mathbb{E}\left[(\mathbf{X}-\mu_\theta)\psi\right] \qquad \frac{\partial^2 \mathcal{J}(q)}{\partial \mu_\theta \partial \mu_\theta^T} = \Sigma_\theta^{-1}\mathbb{E}\left[(\mathbf{X}-\mu_\theta)(\mathbf{X}-\mu_\theta)^T\psi\right]\Sigma_\theta^{-1} - \Sigma_\theta^{-1}\mathbb{E}\left[\psi\right]$$

- Factorized costs:

$$\boxed{\frac{\partial \mathcal{J}(q)}{\partial \mu_\theta} = \sum_{l=1}^{L} M_\ell^T \frac{\partial \mathcal{J}_\ell(q_\ell)}{\partial \mu_\theta^\ell}} \boxed{\frac{\partial^2 \mathcal{J}(q)}{\partial \mu_\theta \partial \mu_\theta^T} = \sum_{l=1}^{L} M_\ell^T \frac{\partial^2 \mathcal{J}_\ell(q_\ell)}{\partial \mu_\theta^l (\partial \mu_\theta^l)^T} M_\ell}$$

- Prior factor: closed-form

	Closed-form Prior (36)	Sparse-grid GH	Full-grid GH
2 DOF Arm	**0.0150**	0.0180	0.0954
7 DOF WAM exp 1	**0.3086**	-	∞

TABLE II: Implementation time comparison for computing prior costs in GVI-MP, averaged over 50 runs. The trajectory consists of 50 support states. For a 7-DOF robot, a 3-degree sparse-grid (resp., full-grid) quadrature method requires $28^3 = 21,952$ (resp., $3^{28} \approx 10^{13}$) sigma points to evaluate one expectation. The closed-form expression (36) is thus indispensable in GVI-MP.

- Nonlinear collision factor: sparse GH-quadratures

Entropy-regularized Robust Motion Planning

- Equivalent: Entropy regularized Variational Motion Planning

$$
\begin{aligned}
q^\star &= \arg\min_{q \in \mathcal{Q}} \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X}|Z)\right) \\
&= \arg\max_{q \in \mathcal{Q}} \underbrace{\boxed{\mathbb{E}_q\left[\log p(\mathbf{X}|Z)\right]}}_{\text{Expected Posterior}} + \underbrace{\boxed{\mathcal{H}(q(\mathbf{X}))}}_{\text{Entropy of the proposal}}.
\end{aligned}
$$

- Trading-off `optimality’ and `robustness’

$$
q^\star = \arg\max_{q \in \mathcal{Q}} \mathbb{E}_q\left[\log p\left(\mathbf{X}|Z\right)\right] + \underbrace{\boxed{\hat{T}}}_{\text{Temperature parameter}} \mathcal{H}(q(\mathbf{X}))
$$

Results

- Optimization Iterations

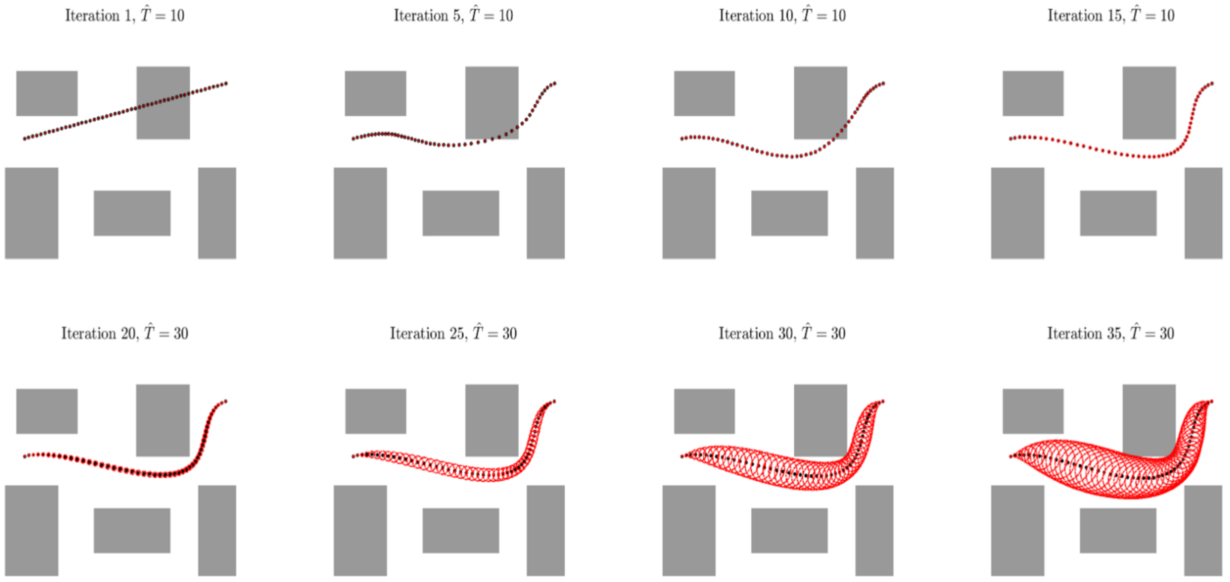

- Iteration Costs

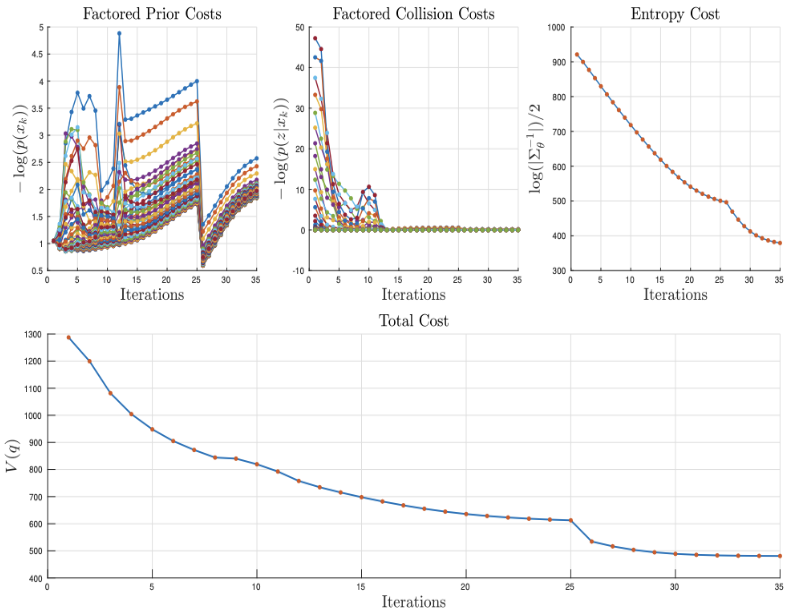

- Optimizable covariances

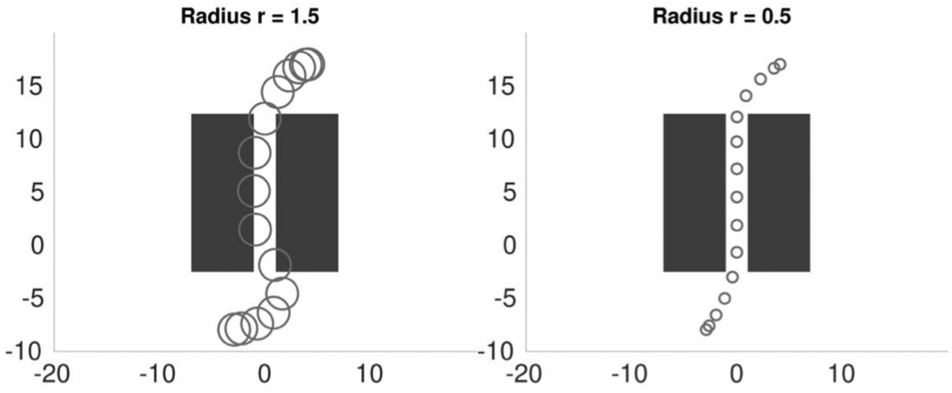

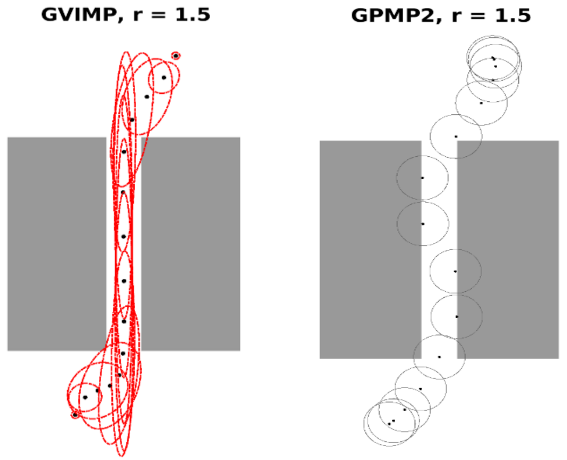

Results

- Planning in cluttered environments

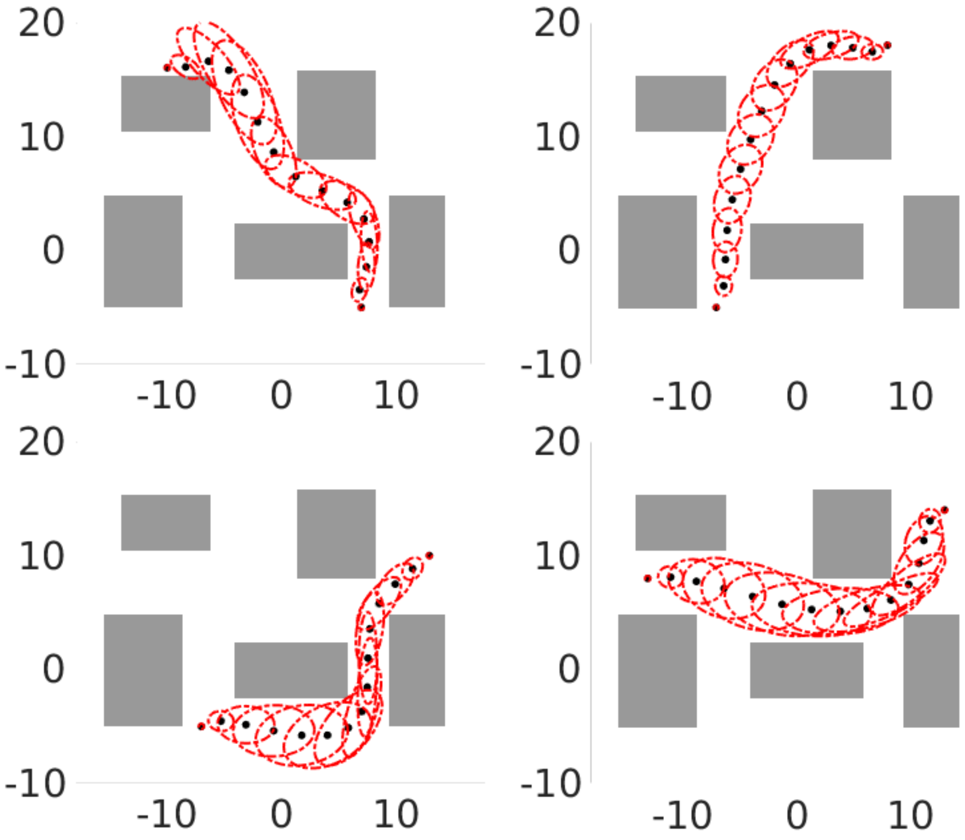

- Robust Decision-Making under uncertainties

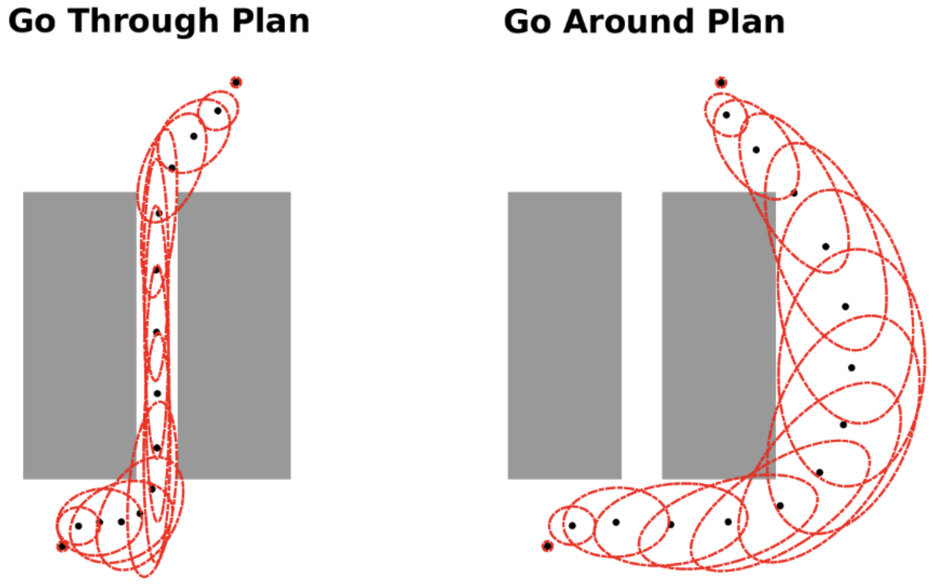

	Prior	Collision	MP	Entropy	Total
Left	**34.4583**	9.1584	**43.6168**	44.1752	87.7920
Right	42.9730	**2.0464**	45.0193	**39.9193**	**84.9387**

Inference-control duality

Girsanov's theorem

- **Prior (uncontrolled)**
  - **SDE**
    $$dX_t = A_t X_t\, dt + \mathrm{a}_t \mathrm{d}t + \mathrm{B}_t \mathrm{d}\mathrm{W}_t$$
  - **Induced path distribution**
    $$d\mathbb{P}^0$$

- **Controlled**
  - **SDE**

    Control Inputs
    $$dX_t = A_t X_t\, dt + \mathrm{a}_t \mathrm{d}t + \mathrm{B}_t u_t + \mathrm{B}_t \mathrm{d}\mathrm{W}_t$$
  - **Induced path distribution**
    $$d\mathbb{P}^u$$

- **Girsanov's Theorem**
  - **The ratio** $\frac{d\mathbb{P}^u}{d\mathbb{P}^0}$ **has a closed form expression**
  - **The KL divergence** $\mathrm{KL}\left(\mathbb{P}^u \parallel \mathbb{P}^0\right) = \int \log \frac{d\mathbb{P}^u}{d\mathbb{P}^0} d\mathbb{P}^u = \frac{1}{2}\mathbb{E}\left\{\int_{t_0}^{t_N} \|u_t\|^2 dt\right\}$ **equals to the expected control energy**

- **Variational Inference (Distributional Control)**

$$\boxed{q^\star = \arg\min_{q \in \mathcal{Q}} \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X}|Z)\right)}$$

$\longleftrightarrow$

- **Stochastic Control**

Control energy

$$\boxed{\begin{aligned} \min_{X(\cdot),u(\cdot)} \quad & \mathbb{E}\left\{\int_{t_0}^{t_N} \frac{1}{2}\|u_t\|^2 dt + \|\mathbf{h}(\mathbf{X})\|^2_{\Sigma_{\mathrm{obs}}} \right. \\ & \left. + \frac{1}{2}\|X_0 - \mu_0\|_{K_0^{-1}} + \frac{1}{2}\|X_N - \mu_N\|_{K_N^{-1}} \right\} \\ \text{s.t.} \quad & dX_t = A_t X_t dt + a_t dt + B_t(u_t dt + dW_t). \end{aligned}}$$

Girsanov's theorem

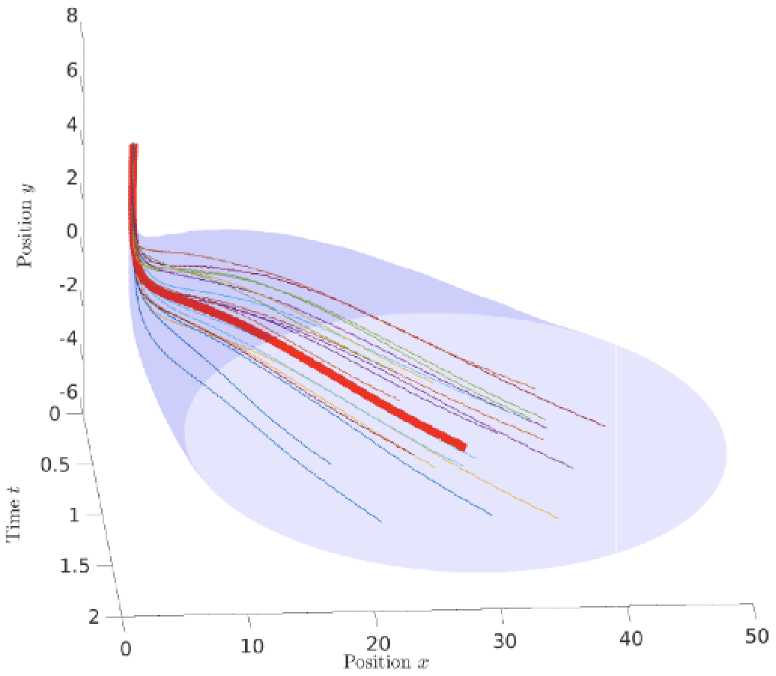

(a) Uncontrolled path distribution

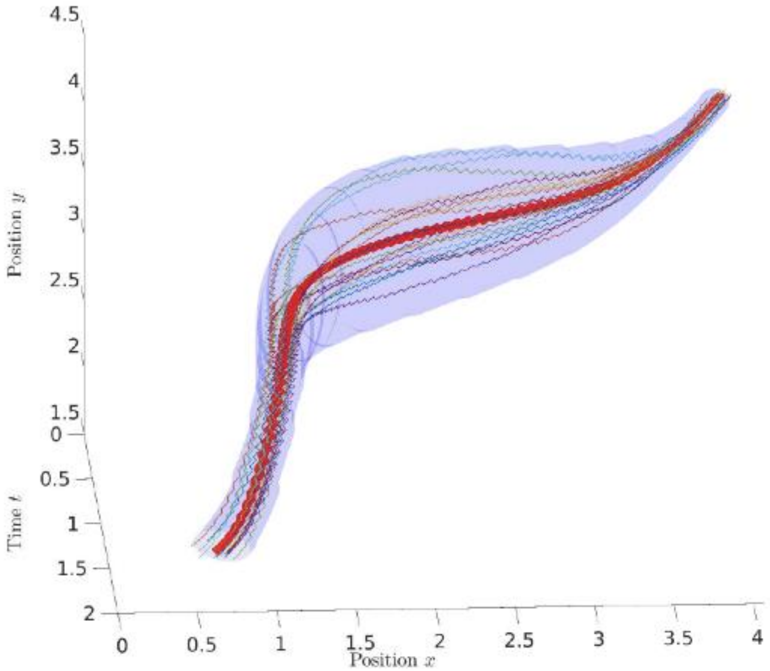

(b) Controlled path distribution

$d\mathbb{P}^0$

$d\mathbb{P}^u$

$$\boxed{\mathrm{KL}\left(\mathbb{P}^u \parallel \mathbb{P}^0\right) = \int \log \frac{d\mathbb{P}^u}{d\mathbb{P}^0} d\mathbb{P}^u = \frac{1}{2}\mathbb{E}\left\{\int_{t_0}^{t_N} \|u_t\|^2 dt\right\}}$$

Covariance Steering for Linear Stochastic Systems with Hybrid Transitions

Covariance Control

- **Goal:** Control the state **covariance** around a mean trajectory from initial to target

$$\min_u \quad \mathbb{E}\left\{\underbrace{\int_0^T [\frac{1}{2}\|u_t\|^2}_{\text{Control energy}} + V(X_t)]dt\right\}$$

$$dX_t = A_t X_t dt + a_t dt + B_t(u_t dt + \sqrt{\epsilon} dW_t)$$

$$X_0 \sim \mathcal{N}(\mu_0, K_0), \quad \boxed{X_T \sim \mathcal{N}(\mu_T, K_T)}$$

Constraint on terminal-time covariance

- **Girsanov's Theorem**
  - **The ratio** $\frac{d\mathbb{P}^u}{d\mathbb{P}^0}$ **has a closed form expression**
  - **The KL divergence equals to the expected control energy:** $\mathrm{KL}\left(\mathbb{P}^u \parallel \mathbb{P}^0\right) = \int \log \frac{d\mathbb{P}^u}{d\mathbb{P}^0} d\mathbb{P}^u = \frac{1}{2}\mathbb{E}\left\{\int_{t_0}^{t_N} \|u_t\|^2 dt\right\}$

- **Distributional Control Formulation**

$$\boxed{q^\star = \arg\min_{q\in\mathcal{Q}} \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X}|Z)\right)}$$

$\mathcal{Q}$: **The collection of distributions that have the target terminal covariance**

Enjoys a convex optimization formulation

Covariance Control for Hybrid Systems

- **Goal**
  - Control the state covariance around a mean trajectory from initial to target, through possible hybrid transitions
  - Hybrid event decided by a mean trajectory
- **Formulation**

$$\min_{u_j(t)} \mathcal{J}_H \triangleq \mathbb{E}\left\{\int_0^T [\|u_j(t)\|^2 + X_j'(t)Q_j(t)X_j(t)]dt\right\}$$

$$dX_1 = A_1(t)X_1 dt + B_1(t)(u_1 dt + \sqrt{\epsilon}dW_1),$$

$$\boxed{X_2(t^+) = \Xi X_1(t^-),}$$

$$dX_2 = A_2(t)X_2 dt + B_2(t)(u_2 dt + \sqrt{\epsilon}dW_2),$$

$$\boxed{X_1(0) \sim \mathcal{N}(m_0, \Sigma_0), \quad X_2(T) \sim \mathcal{N}(m_T, \Sigma_T),}$$

Saltation Matrix for jump dynamics

- Linearizing the nonlinear jump dynamics
  - Nonlinear reset map
  - Instantaneous changes in the state
- Direct way: Jacobian of the reset map, low precision
- Saltation matrix: Considering the total variation of state and time in the jump events

$$\Xi_{(\mathrm{I},\mathrm{J})} := \mathrm{D}_x R^- + \frac{\left(F_\mathrm{J}^+ - \mathrm{D}_x R^- F_\mathrm{I}^- - \mathrm{D}_t R^-\right) \mathrm{D}_x g^-}{\mathrm{D}_t g^- + \mathrm{D}_x g^- F_\mathrm{I}^-}$$

- **Covariance propagation** at hybrid event:

$$\boxed{\Sigma(t^+) = \Xi_{(\mathrm{I},\mathrm{J})} \Sigma(t^-) \Xi_{(\mathrm{I},\mathrm{J})}^T}$$

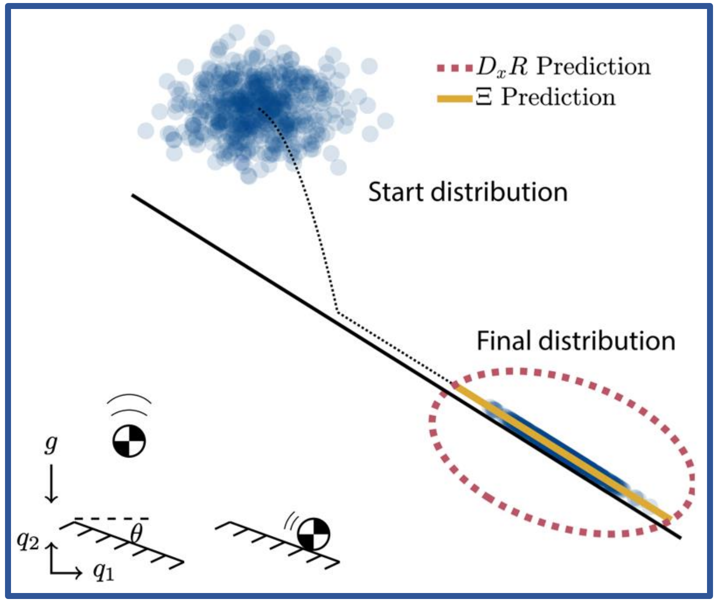

Hybrid Covariance Steering

- **Distributional Control Formulation for covariance steering**

$$\boxed{q^\star = \arg\min_{q \in \mathcal{Q}} \mathrm{KL}\left(q(\mathbf{X}) \parallel p(\mathbf{X}|Z)\right)}$$

$\mathcal{Q}$: **The collection of distributions that have the target terminal covariance**

$$\min_{u_j(t)} \mathcal{J}_H \triangleq \mathbb{E}\left\{\int_0^T [\|u_j(t)\|^2 + X_j'(t)Q_j(t)X_j(t)]dt\right\}$$

$$dX_1 = A_1(t)X_1 dt + B_1(t)(u_1 dt + \sqrt{\epsilon} dW_1), \quad (21a)$$

$$\boxed{X_2(t^+) = EX_1(t^-), \quad (21b)}$$

$$dX_2 = A_2(t)X_2 dt + B_2(t)(u_2 dt + \sqrt{\epsilon} dW_2), \quad (21c)$$

$$X_1(0) \sim \mathcal{N}(m_0, \Sigma_0), \quad X_2(T) \sim \mathcal{N}(m_T, \Sigma_T), \quad (21d)$$

SDP Formulation

$$\boxed{\mathcal{J}_j = \mathrm{KL}\left(\mathbb{P}_{(t_0^j, T_j)} \parallel \mathbb{P}^*_{(t_0^j, T_j)}\right)}$$

$$\boxed{\Sigma^+ = E\Sigma^- E'.}$$

- This objective function is convex in the optimization variables!
  - SDP program
  - Dimension linear to the number of hybrid events, instead of time discretization

Numerical Results

- Bouncing ball

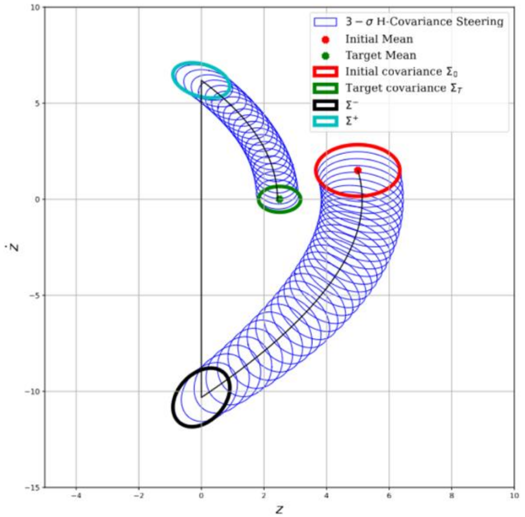

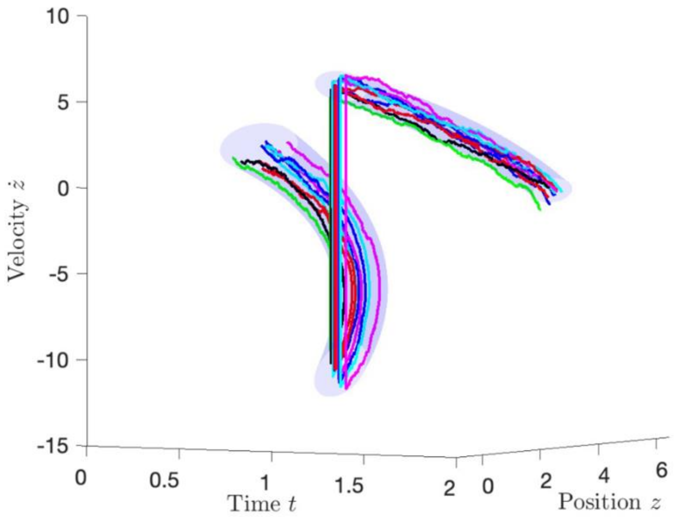

- SLIP: Linearized system

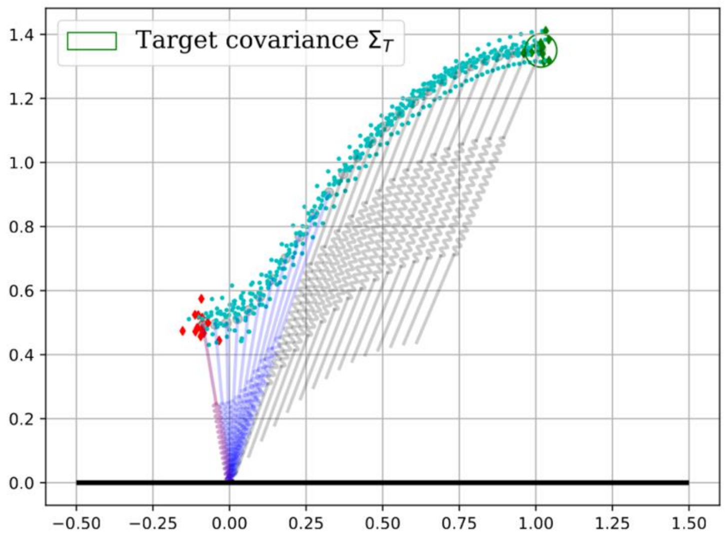

Path Integral Control for Hybrid Dynamical Systems

Stochastic Systems with Hybrid Transitions

- Hybrid Dynamical Systems
  - Co-existence of continuous flows and discrete jump events
  - Modes, fields, flows, guards, jumps (resets)

- Consider **nonlinear** SDE as the flow
  - Flow in j^{th} mode:

$$dX_t^j = F_j(t, X_t^j)dt + \sigma_j(t, X_t^j)(u^j(t, X_t^j)dt + \sqrt{\epsilon}dW_t^j)$$

  - Guard and reset (from j^{th} mode to k^{th} mode):

$$g_{jk}(t^-, X_j(t^-)) \leq 0,\ X_j(t^-) \in I_j$$
$$t^+ = t^-,\ \ X_k(t^+) = R_{jk}(X_j(t^-))$$

Problem formulation

- Objective function

$$\min_u \mathcal{J}_H \triangleq \mathbb{E}\left[\sum_{j=0}^{N_J}\int_{t_j^+}^{t_{j+1}^-}\left(V(t,X_t^j)+\frac{1}{2}\|u_t^j\|^2\right)dt+\Psi_T\right]$$

- Girsanov's theorem with hybrid transitions
  - Remains similar to the smooth case

$$\frac{d\mathbb{P}^u}{d\mathbb{P}^0}=\exp\left(\sum_{j=0}^{N_J}\int_{t_j^+}^{t_{j+1}^-}-\frac{\|u_t^j\|^2}{2\epsilon}dt+\frac{1}{\sqrt{\epsilon}}(u_t^j)'dW_t^j\right)$$

- Optimal controlled distribution: known form

$$d\mathbb{P}^*=\frac{\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H\right)d\mathbb{P}^0}{\int\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H\right)d\mathbb{P}^0}=\frac{\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H\right)}{\mathbb{E}_{\mathbb{P}^0}\left[\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H\right)\right]}d\mathbb{P}^0.$$

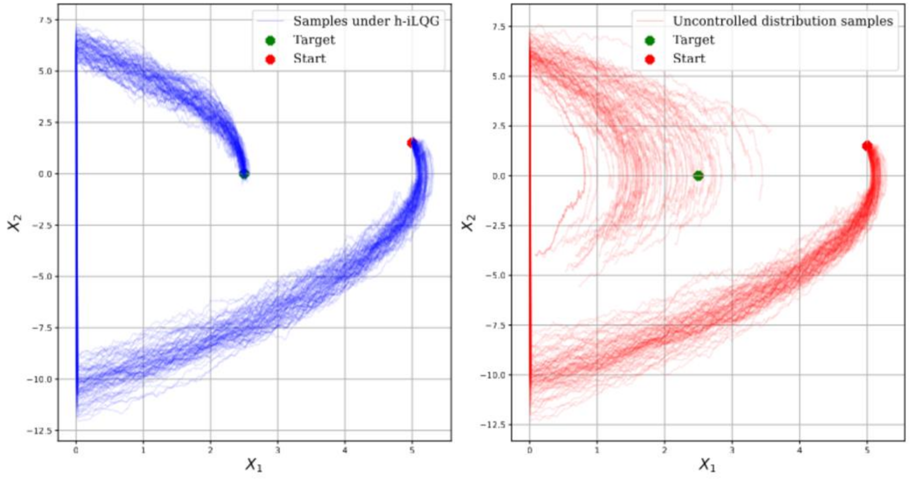

controlled $d\mathbb{P}^u$

uncontrolled $d\mathbb{P}^0$

Path integral control: Cross-entropy method

- Cross-entropy method

$$u^* = \arg\min_u \mathrm{KL}\left(\boxed{\mathbb{P}^*} \,\|\, \boxed{\mathbb{P}^u}\right)$$

Optimal controlled distribution

Proposal distribution

- Optimal controller as an **expectation** estimation using samples

$$u_t^* = \frac{\sqrt{\epsilon} \times \boxed{\mathbb{E}_{\mathbb{P}^0}} \left[\Delta W_t \exp\left(-\frac{1}{\epsilon}\mathcal{L}_H(t)\right)\right]}{\Delta t \times \boxed{\mathbb{E}_{\mathbb{P}^0}} \left[\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H(t)\right)\right]}$$

- Advantages
  - No need to solve HJB with hybrid guards and resets
  - Parallel Sampling on GPU
- Potential issue: High variance (need many samples)

Importance Sampling using H-iLQR

- Optimal controller as an **expectation** estimation using samples

$$u_t^* = \frac{\sqrt{\epsilon} \times \boxed{\mathbb{E}_{\mathbb{P}^0}} \left[\Delta W_t \exp\left(-\frac{1}{\epsilon}\mathcal{L}_H(t)\right)\right]}{\Delta t \times \boxed{\mathbb{E}_{\mathbb{P}^0}} \left[\exp\left(-\frac{1}{\epsilon}\mathcal{L}_H(t)\right)\right]}$$

- Potential issue: Low variance (need many samples)
- Importance sampling

$$\mathbb{E}_{\mathbb{P}_0}[f] = \mathbb{E}_{\mathbb{P}_u}\left[f \times \frac{d\mathbb{P}_0}{d\mathbb{P}_u}\right]$$

- Girsanov theorem

$\frac{d\mathbb{P}^u}{d\mathbb{P}^0}$ is known (expression)

- Optimal controller under a proposal (controlled) distribution

$$u_t^* = u_t + \frac{\sqrt{\epsilon} \times \boxed{\mathbb{E}_{\mathbb{P}^u}} \left[\exp\left(-\frac{1}{\epsilon}\mathcal{S}_H^u(t)\right) \Delta \tilde{W}_t\right]}{\Delta t \times \boxed{\mathbb{E}_{\mathbb{P}^u}} \left[\exp\left(-\frac{1}{\epsilon}\mathcal{S}_H^u(t)\right)\right]}$$

Results and observations

- Hybrid transition affect the quality of the feedback controller

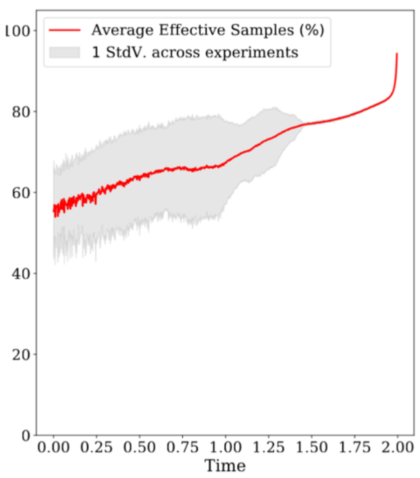

(b) Effective samples, ball.

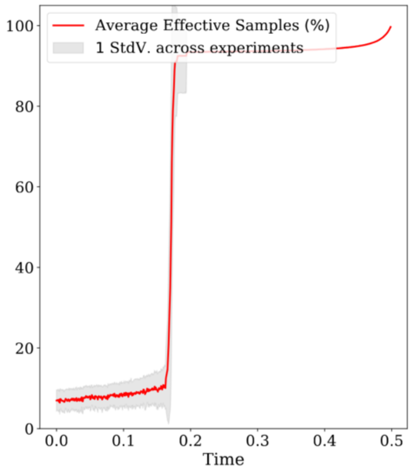

(d) Effective samples, SLIP.

- H-iLQR improves the sample efficiency in H-PI

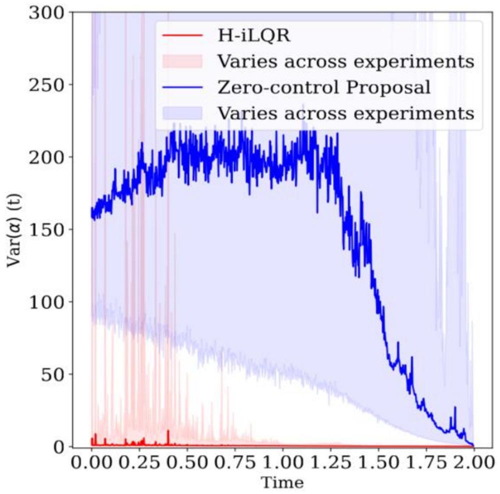

(a) Variance.

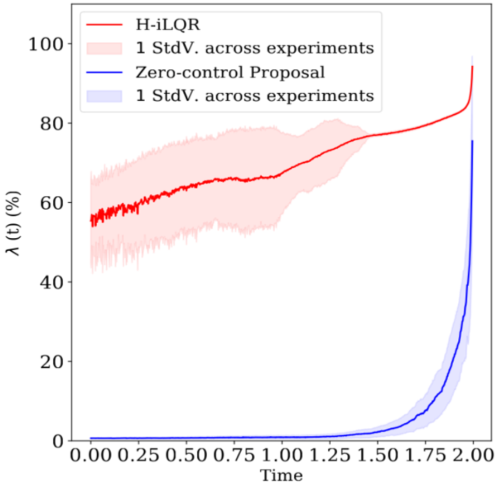

(b) Effective Samples.

- H-PI obtains optimal control

		H-iLQR	H-PathIntegral	Improved (%)
Expected Cost	Bouncing Ball	60.11	57.86	3.74
	SLIP	0.2459	0.2164	**12.00**
H-iLQR 10% tail	Bouncing Ball	141.65	115.40	**11.46**
	SLIP	0.9451	0.2616	**63.29**

Table 5.1: Expected cost improvement and the statistics on the experiments where H-iLQR has the highest 10% costs.

Thank you!

Q & A