Information-Geometric Path Planning: Roles of Information Theory in Motion Planning and Future Research Opportunities

Takashi Tanaka
The University of Texas at Austin
Purdue University (Jan. 2025 -)

2024 AFOSR Dynamical Systems and Control Theory Review
August 28, 2024

Outline

1. Summary of YIP "Information-Geometric Path Planning" 2020-2023
   - Review: “Rationally inattentive” path planning
   - Follow-up studies: Deep-learning-assisted motion planning

2. Outlook on a new project “Motion Planning, Partial Observability, and Quantum Mechanics: Advancing the Frontiers of Path Integral Control Theory” 2025-2027
   - Brief history of path integral control
   - Path integral control of partially observable systems

3. Deceptive path planning and hypothesis testing

Rationally Inattentive Path Planning

“Rationally Inattentive” Path Planning

Sensing is costly. Sensing cost appears in various forms:

- Computational cost
- Communication cost
- Mechanical cost
- Time cost

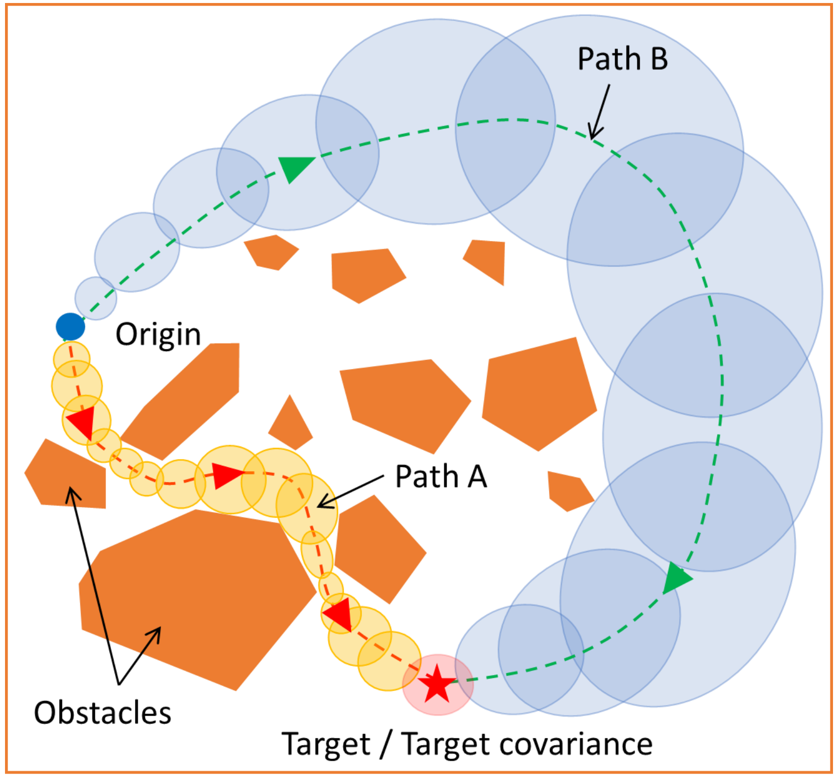

Path geometry is a key factor determining sensing cost.

- Following Path A requires precise localization.
  Expected sensing cost for path following will be higher.

How to generate a motion plan with a “simple” path geometry that requires minimum sensing to execute?

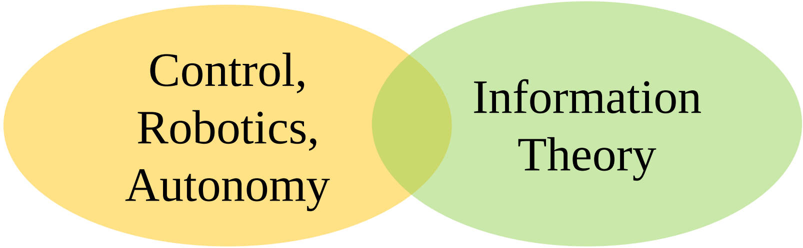

Shortest Path Problem in Gaussian Belief Space

- Robot's configuration is represented by $\mathcal{N}(x_k, P_k)$. Belief state: $b_k = (x_k, P_k)$.
- We want to compute the "shortest" collision-free path in the Gaussian belief space $\mathbb{B} = \mathbb{R}^n \times \mathbb{S}^n_+$.

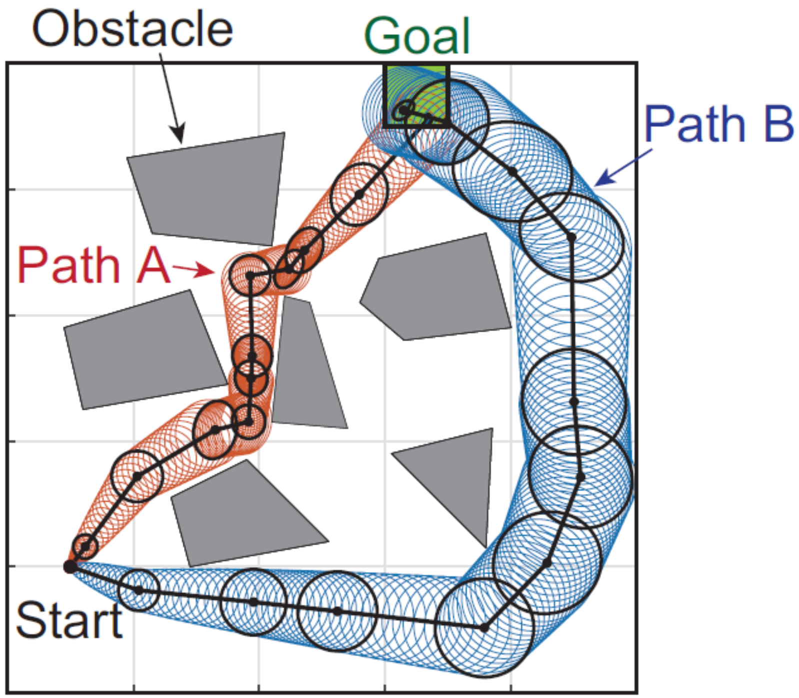

Distance Function

We define the distance from b_k to b_{k+1} to be

$$D(b_k, b_{k+1}) = D_{\text{travel}}(b_k, b_{k+1}) + \alpha D_{\text{info}}(b_k, b_{k+1})$$

- Travel cost: $D_{\text{travel}}(b_k, b_{k+1}) = \|x_k - x_{k+1}\|$.
- Information cost (entropy reduction):

$$D_{\text{info}}(b_k, b_{k+1}) = \frac{1}{2} \log\det(P_k + \|x_{k+1} - x_k\| W) - \frac{1}{2} \log\det P_{k+1}$$

where W is the natural growth of uncertainty.

$D(\cdot, \cdot)$ introduces a quasi-pseudometric (Lawvere metric) on the belief manifold $\mathbb{B}$.

Path Length

- The length of a belief path $\gamma(t) = (x(t), P(t))$ w.r.t. the partition $\mathcal{P} = (0 = t_1 < t_2 < \cdots < t_K = T)$ is defined by $c(\gamma, \mathcal{P}) = \sum_{k=1}^{K} D(\gamma(t_k), \gamma_{t_{k+1}})$.
- The length of γ is defined by $c(\gamma) = \sup_{\mathcal{P}} c(\gamma, \mathcal{P})$.
- Belief path $\gamma(t) = (x(t), P(t))$ is collision-free if

$$(x(t) - x_{\mathrm{obs}})^\top P^{-1}(t)(x(t) - x_{\mathrm{obs}}) \geq \chi^2 \quad \forall t \in [0, T], \forall x_{\mathrm{obs}} \in \mathcal{X}_{\mathrm{obs}}$$

where χ^2 is a user-defined confidence parameter.

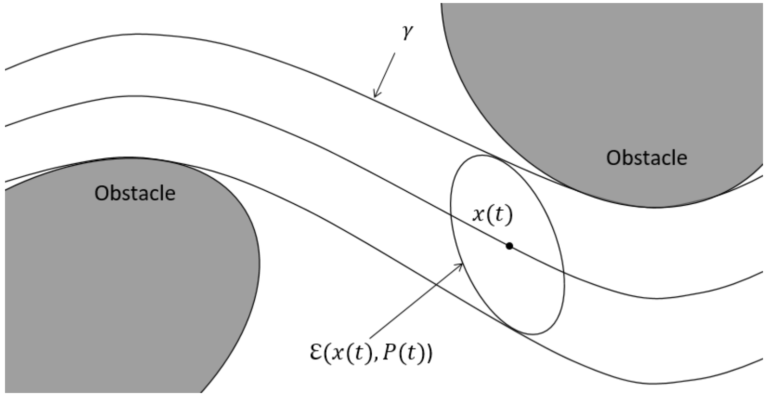

Shortest Belief Path Problem

$$
\begin{aligned}
\min_{\gamma} \quad & c(\gamma) \\
\text{s.t.} \quad & \gamma(0) = b_0, \quad \gamma(T) \in \mathcal{B}_{\text{target}} \\
& (x(t) - x_{\text{obs}})^\top P^{-1}(t)(x(t) - x_{\text{obs}}) \geq \chi^2 \;\; \forall t \in [0, T], \forall x_{\text{obs}} \in \mathcal{X}_{\text{obs}}
\end{aligned}
$$

We developed an RRT* and PRM*-based shortest path solver:

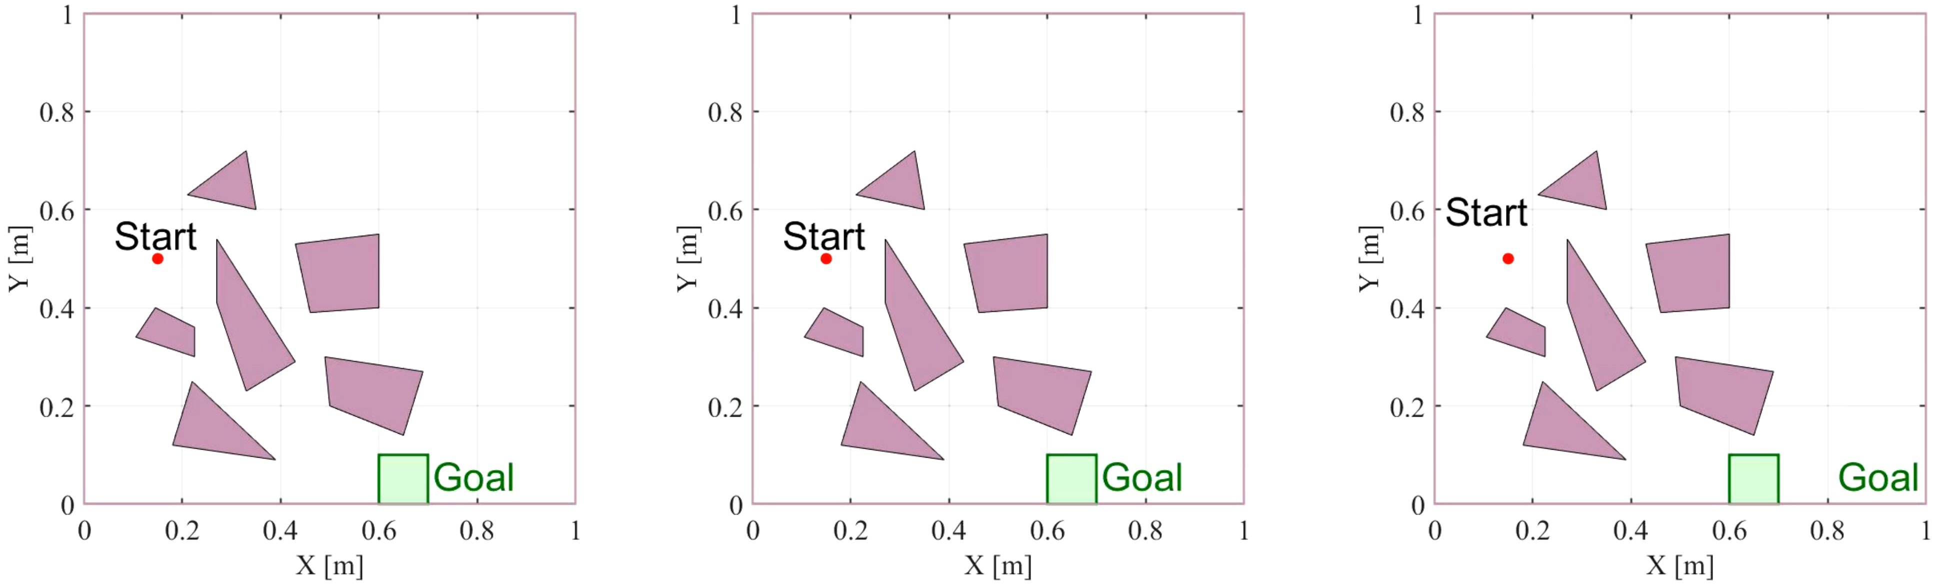

a) $\alpha = 0.0$ b) $\alpha = 0.3$ c) $\alpha = 0.7$

Case Study: Ground Robot Navigation with Event-based Sensing

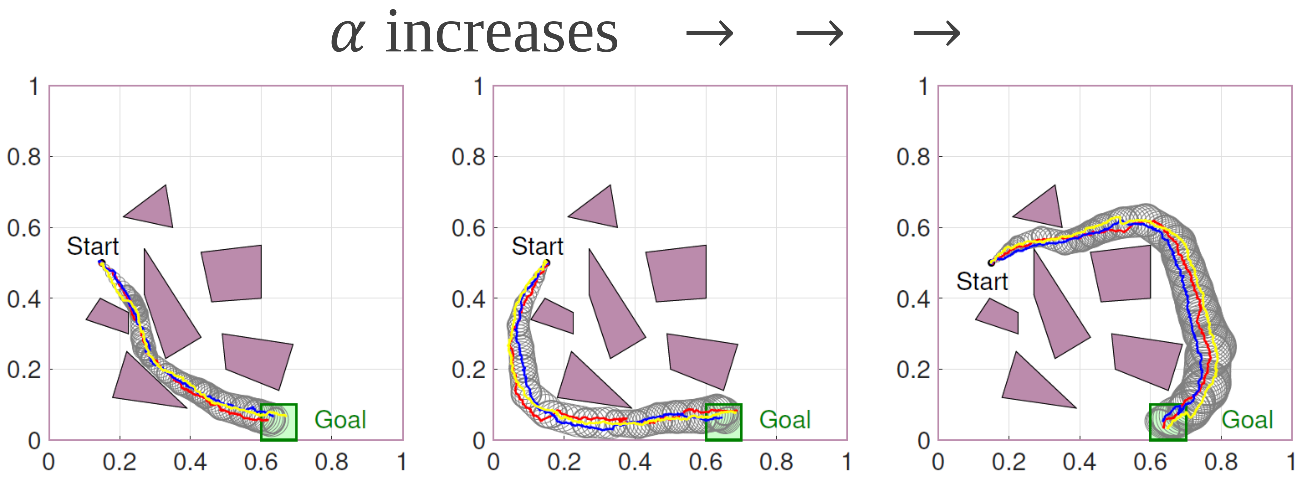

- A reference belief path $\{(x_k^{ref}, P_k^{ref})\}_{k=0,1,2,\ldots}$ is generated by the RRT*-based algorithm.
- Belief path following: Sensors and actuators are operated in real-time to follow the reference belief path.
- Event-based sensing: Location sensor is activated only when $\epsilon(\hat{x}_k, P_k)$ is not contained in $\epsilon(x_k^{ref}, P_k^{ref})$.

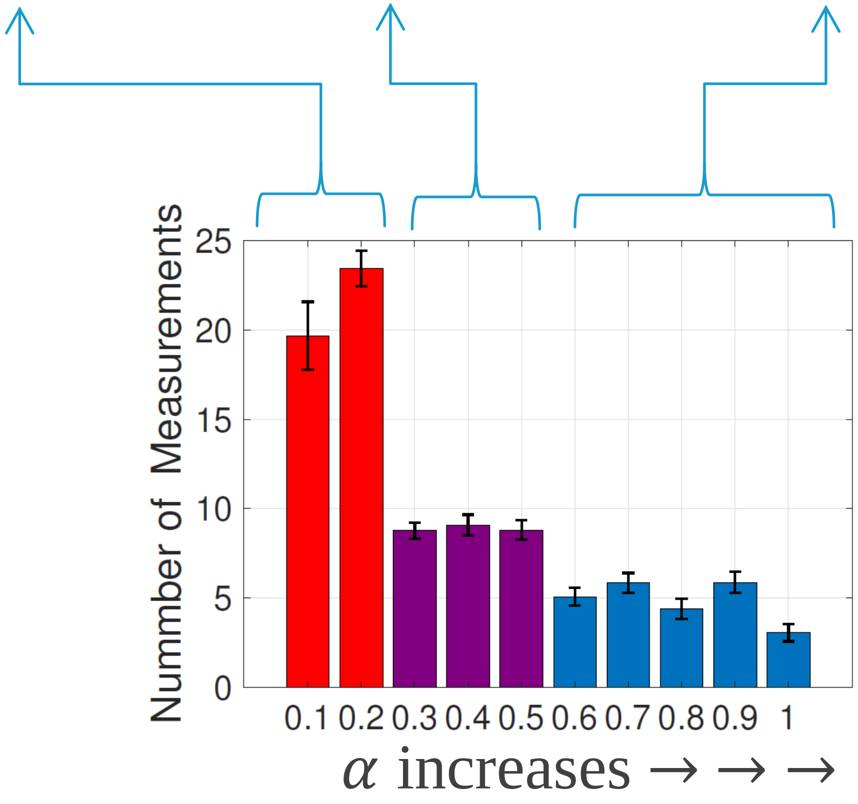

Follow-Up Work: Deep-Learning-Assisted Real-Time Planning

It is difficult to run sampling-based path planner in real-time. We consider applying a deep learning model (U-net) to quickly generate an optimal path candidate.

Input Image -- (3, 150, 225)

Indicator function of randomly generated triangular obstacles

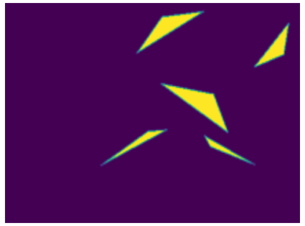

Distance function to initial position

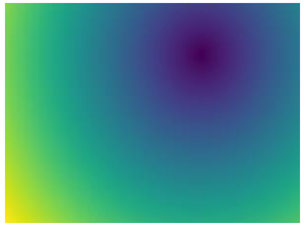

Distance function to target area

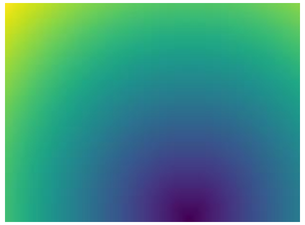

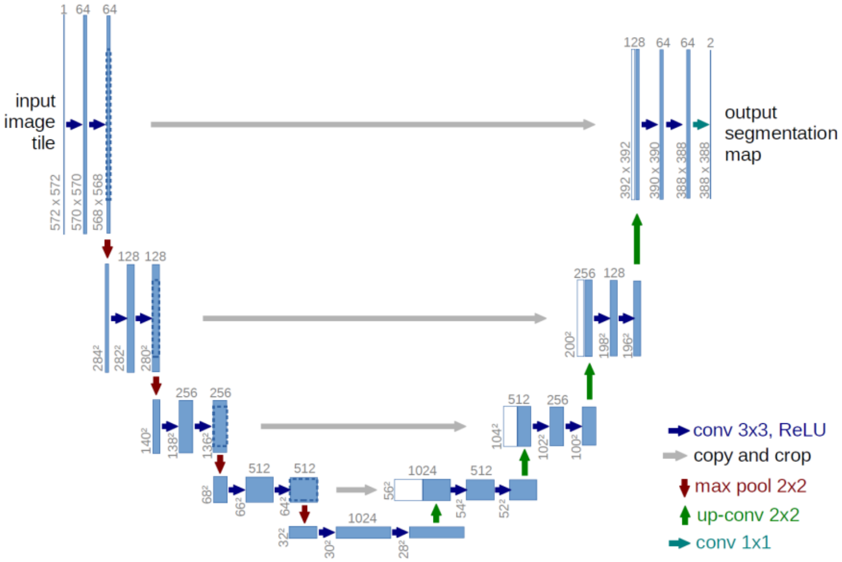

U-net [Ronneberger et al. 2015]

Output Image (Label) -- (150, 225)

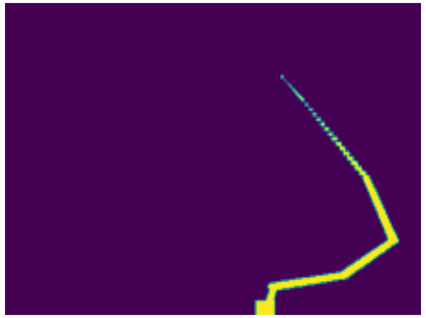

Optimal path computed by RRT*-based planner

U-Net Architecture

- A transition-down block extracts and processes high-frequency components
- A transition-up block reassembles the processed information
- Skip connections give the transition-up block to the previous unprocessed information
- The bottleneck layer transforms the leftover low-frequency components
- The final convolution refines the fully reassembled solution

This allows U-Net to access different frequency components of the input separately to process them and reconstruct the solution

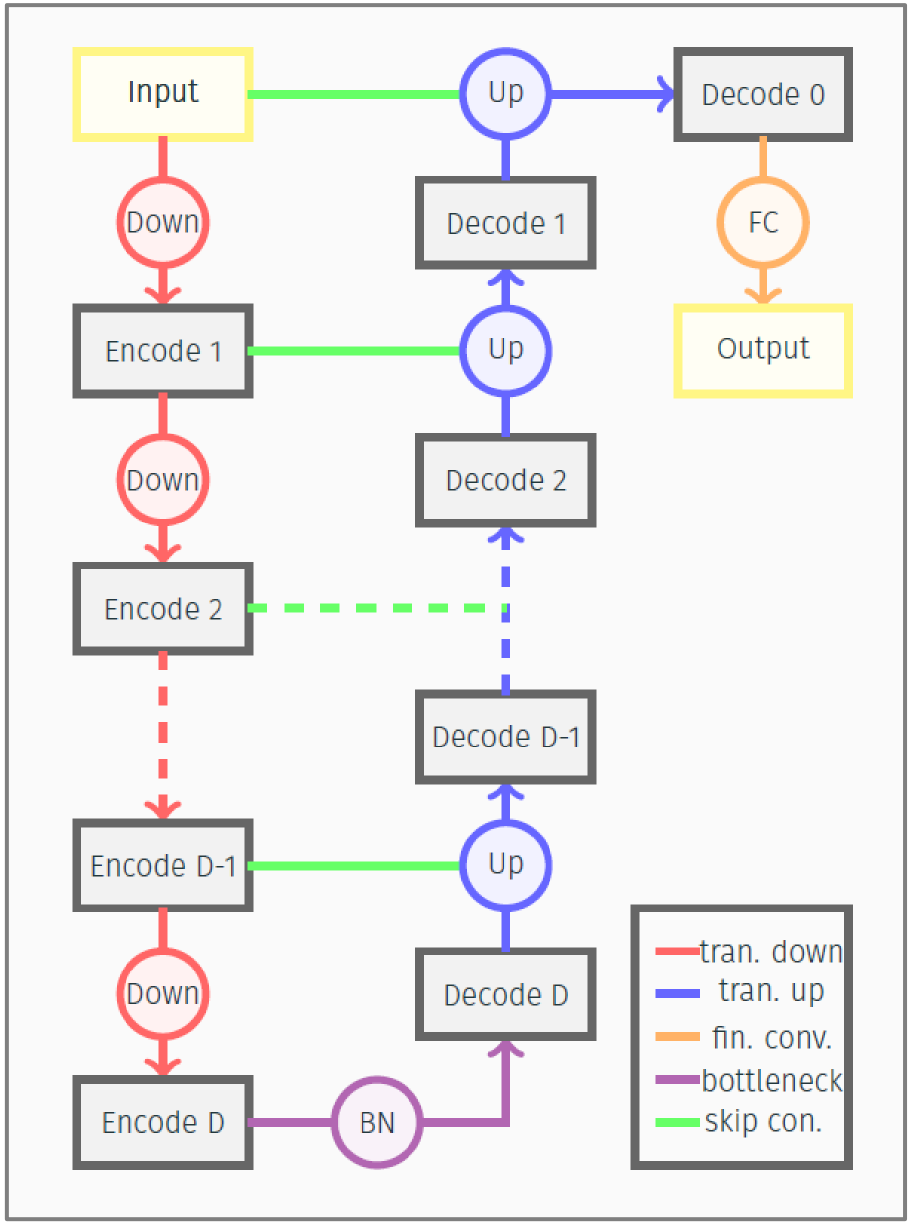

U-Net Training

The network f_θ is trained by $\min_\theta BCE(f_\theta(X), Y)$ where X: input, Y: output (label), and

$$BCE(x, y) = -\frac{1}{L}\sum_{i=1}^{L}[y_i \log(x_i) + (1 - y_i)\log(1 - x_i)]$$

- We trained a U-net model until it can predict an output image from a given input image.
- A trained U-net model is tested on unseen input images to see if it can predict the corresponding path plan.
- While the preparation of a training data set and the training of a U-net are time-consuming processes, a forward execution of the trained U-net is instantaneous.

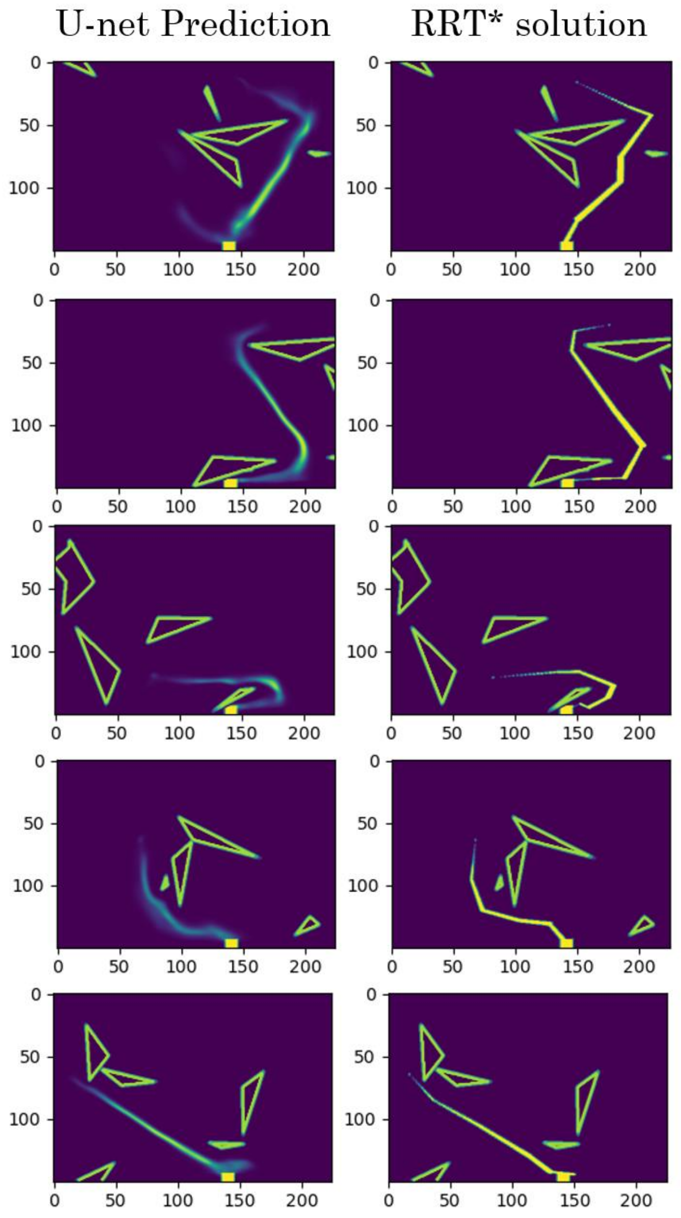

Generalizability

How does the trained U-net perform on unseen problems?

- We trained U-net only using triangle-shaped obstacles.
- We tested the trained U-net in environments with circle-shaped obstacles.

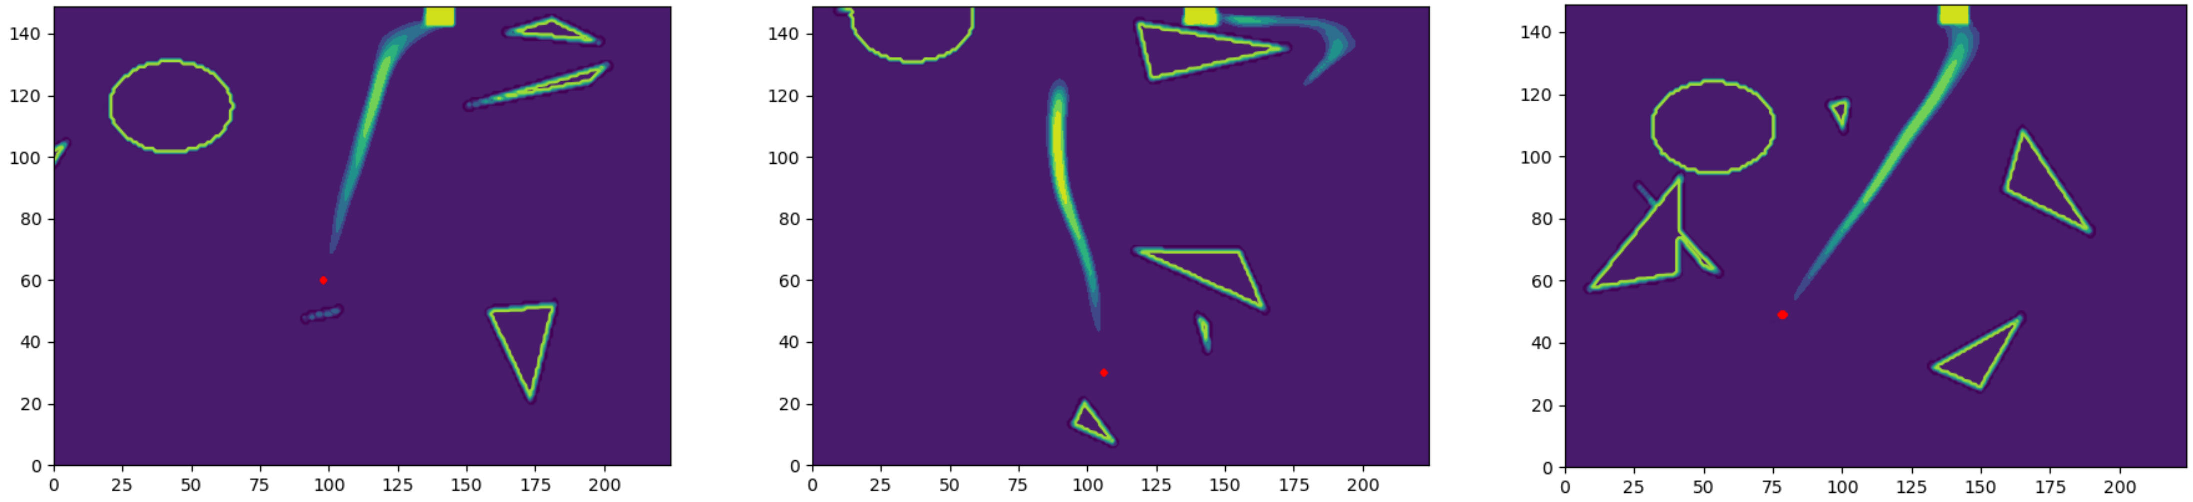

Future work: Quantitative analysis of the robustness against unseen data.

Major Outcomes of YIP

“Rationally Inattentive” path planning

- A. Pedram, “Information-Theoretic Path Planning and Navigation” Ph.D. dissertation, The University of Texas at Austin, 2023.
- A. Pedram, R. Funada and T. Tanaka, “Gaussian Belief Space Path Planning for Minimum Sensing Navigation,” IEEE Transactions on Robotics, vol. 39, no. 3, pp. 2040-2059, June 2023.

Simultaneous Perception-Action Design

- M. Hibbard, T. Tanaka, and U. Topcu, “Simultaneous Perception-Action Design via Invariant Finite Belief Sets,” Automatica, vol. 155, pp. 111140, Sep. 2023.

Publication

2 Ph.D. Theses, 4 Journal publications, 5 Journal papers under review, 13 conference papers

What’s Next?

Through this YIP project, we witnessed various roles that information theory can play in motion planning problems.

- Minimum-information (“rationally inattentive”) motion planning
- Maximum information path planning
- Information sharing in multi-agent path planning
- Path integral control
- Deceptive control and hypothesis testing

Path Integral Control

Outlook on a new project "Motion Planning, Partial Observability, and Quantum Mechanics: Advancing the Frontiers of Path Integral Control Theory"

Path Integral Control (PIC)

A control algorithm inspired by the path-integral formulation of quantum mechanics

Physical system

Equation-based modeling

Nonlinear/stochastic/uncertain
Continuous/discrete/hybrid
High/infinite dimensional
Full/partial observability

$$\dot{x} = f(x, u)$$

Control Policy Synthesis

$$\boxed{\dot{x} = f(x, u)}$$

$$\boxed{u = k(x)}$$

Control Policy

Digital Twin

Simulator Building

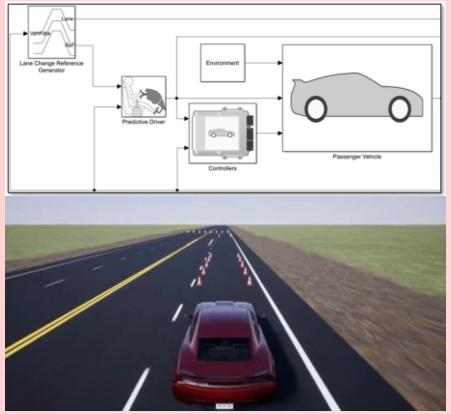

Image: Matlab Vehicle Dynamics Blockset

Real-time Simulation

Monte-Carlo sampling of open-loop (uncontrolled) trajectories

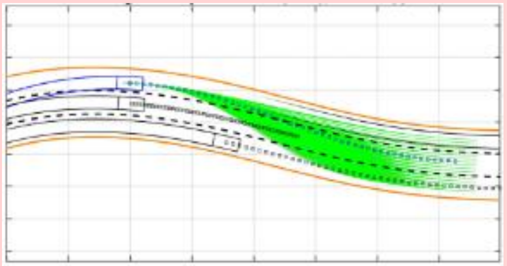

Image: Peng et al. 2022

Path integral control

Background: Simulator Development Became Easy and Quick

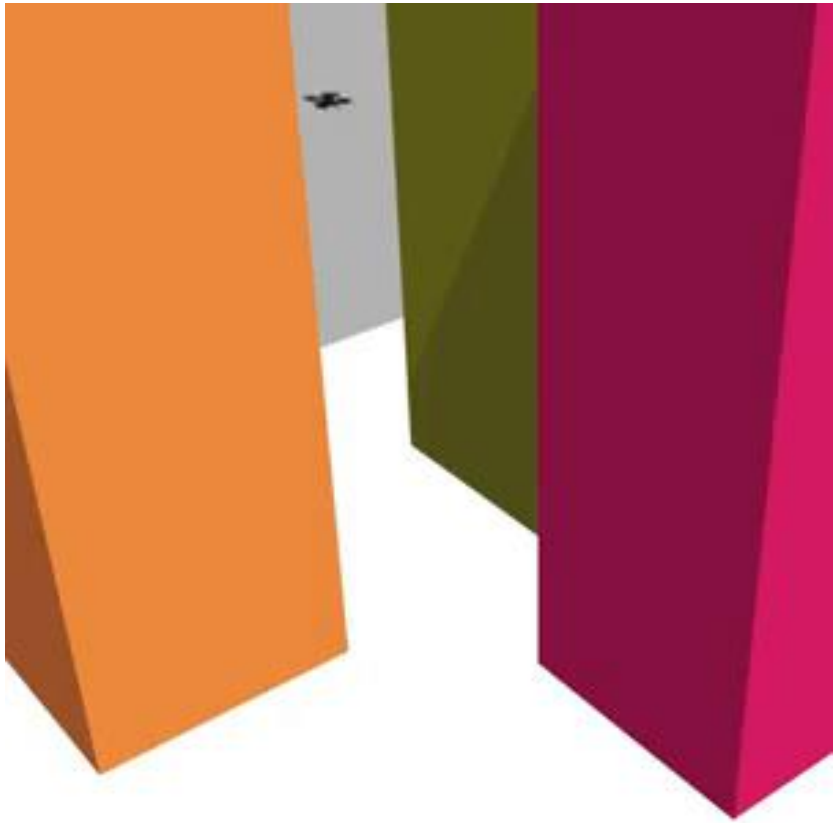

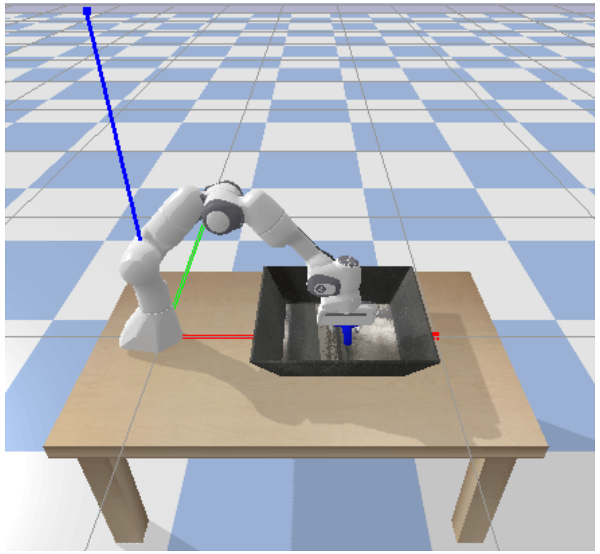

Open AI Gym [Abdeetedal, https://www.etedal.net/2020/04/pybullet-panda.html]

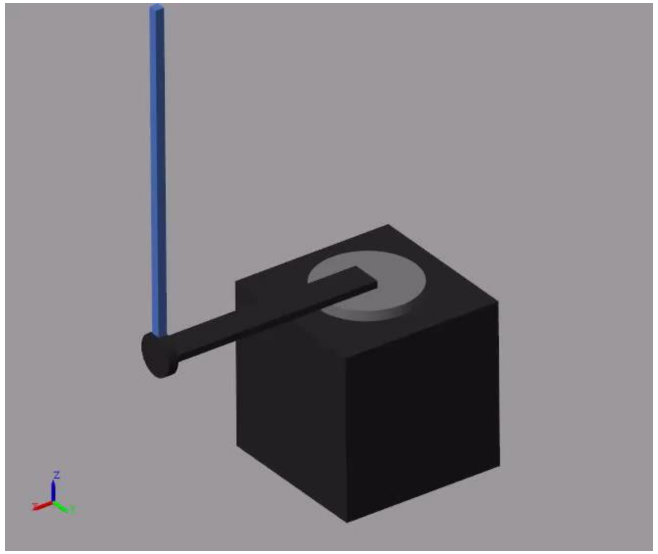

Matalb Simscape
[Suh and Tanaka, Lecture material 2019]

- Building a robotic simulator is often easier than writing down an analytical equation of motion.
- PIC can compute the optimal control input using simulators only.

Gazebo simulator
[Pedram, Funada and Tanaka, TRO 2022]

Brief History of Path Integral Control

- The origin of PIC can be traced back to the stochastic variational treatment of quantum mechanics [1-4]
- [4,5] identified the class of SOC problems in which the associated HJB equation coincides with the linear Schrödinger equation.
- [6,7] noticed the semi-classical limit ($\hbar \to 0$) of the Schrödinger equation approximates the HJB equation corresponding to a broader class of SOC problems.
- [6,7] invoked the Feynman-Kac formula to numerically evaluate the solution of Schrödinger equations using Monte Carlo (Metropolis-Hastings) algorithm.
- In separate threads: Risk-sensitive control, distributionally robust control against relative entropy ambiguity set, linearly solvable MDPs, …

1. E. Nelson, Derivation of the Schrödinger equation from Newtonian mechanics, Physical Review, 1966.
2. H. Rosenbrock, A variational principle for quantum mechanics, Physics Letters A, 1985.
3. H. Rosenbrock, A stochastic variational treatment of quantum mechanics, Proceedings of the Royal Society of London. 1995.
4. F. Guerra and L. M. Morato, Quantization of dynamical systems and stochastic control theory, Physical review D, 1983.
5. K. Yasue, Stochastic calculus of variations, Journal of Functional Analysis, 1981.
6. T. Itami, Optimization of nonlinear control systems based on the principle of superposition (in Japanese), Transactions of the Society of Instrument and Control Engineers, 2001.
7. T. Itami, Nonlinear optimal control via Monte-Carlo evaluation of path integrals (in Japanese), Transactions of the Institute of Systems, Control and Information Engineers, 2003.

Modern Path Integral Control

Under the deterministic system assumption, the SOC can be solved exactly by Monte Carlo as $N \to \infty$.

- Derivation by [1]

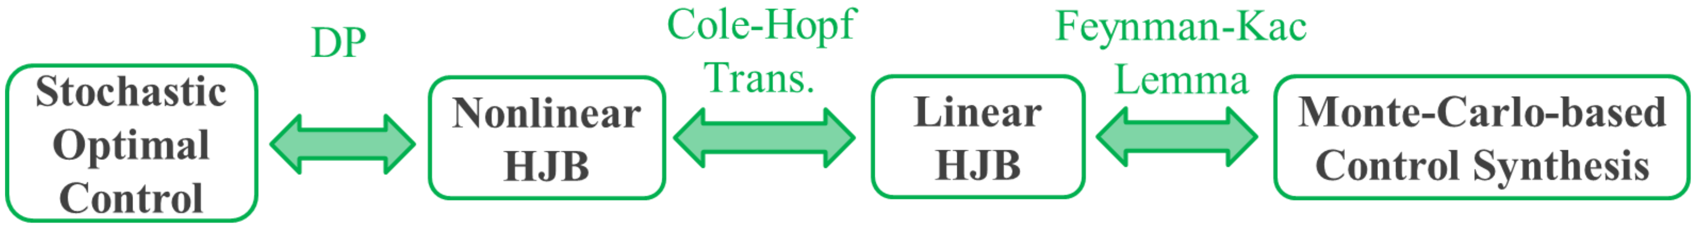

- Derivation by [2]

Removal of deterministic system assumption has been considered in [3-5].

1. H. Kappen, Path integrals and symmetry breaking for optimal control theory, Journal of Statistical Mechanics, 2005.
2. E. Theodorou and E. Todorov, Relative entropy and free energy dualities: Connections to path integral and KL control, CDC 2012
3. S. Satoh, H. J. Kappen, and M. Saeki, An iterative method for nonlinear stochastic optimal control based on path integrals, TAC, 2016.
4. G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, Information-theoretic model predictive control: Theory and applications to autonomous driving, TRO 2018
5. S. Levine, Reinforcement learning and control as probabilistic inference: Tutorial and review, arXiv:1805.00909, 2018.

Proposed Research (2025-2027)

Thrust 1: Path Integral Control for Spatial Navigation

- Task 1-1: Chance-constrained SOC and strong duality.
- Task 1-2: Minimum sensing navigation.

Thrust 2: Path Integral Method and Partially Observable Systems.

- Task 2-1: Removal of the deterministic system assumption.
- Task 2-2: KLD-regularized POMDP.

Thrust 3: Quantum Theoretic Perspectives of Path Integral Control.

- Task 3-1: Classical optimal control problems and Schrödinger equation.
- Task 3-2: Quantum optimal control via the path integral method.

Path Integral Control of Partially Observable Systems

$P(x_{t+1}\|x_t, u_t)$	State transition probability
$P(x_0)$	Initial state distribution
$P(y_t\|x_t)$	Observation model

$C(x_t, u_t)$	Running cost
$\psi(x_T)$	Terminal cost
$R(u_t\|y_{0:t}, u_{0:t-1})$	Baseline (reference) policy
$Q(u_t\|y_{0:t}, u_{0:t-1})$	Control policy to be designed

Joint probability measures of the trajectories induced by policies Q and R:

$$Q(x_{0:T}, y_{0:T-1}, u_{0:T-1}) = P(x_0)\prod_{t=0}^{T-1} P(y_t|x_t)Q(u_t|y_{0:t}, u_{0:t-1})P(x_{t+1}|x_t, u_t)$$

$$R(x_{0:T}, y_{0:T-1}, u_{0:T-1}) = P(x_0)\prod_{t=0}^{T-1} P(y_t|x_t)R(u_t|y_{0:t}, u_{0:t-1})P(x_{t+1}|x_t, u_t).$$

Partially observable KL control: $$\min_{\{Q(u_t|y_{0:t},u_{0:t-1})\}_{t=0}^{T-1}} \mathbb{E}^Q\left[\sum_{t=0}^{T-1} C(x_t, u_t) + \psi(x_T)\right] + \lambda D(Q\|R)$$

Belief Space Representation

Define the belief state $b_t(x_t) := Q(x_t|y_{0:t}, u_{0:t-1})$. By the Bayes formula, we have

$$b_{t+1}(x_{t+1}) = \frac{P(y_{t+1}|x_{t+1}) \int_{\mathcal{X}_t} P(x_{t+1}|x_t, u_t) b_t(x_t)}{\int_{\mathcal{X}_{t+1}} P(y_{t+1}|x_{t+1}) \int_{\mathcal{X}_t} P(x_{t+1}|x_t, u_t) b_t(x_t)}.$$

This formula shows that b_t is a controlled Markov process (controlled by u_t).

Equivalent problem in belief space: $\min_{\{Q(u_t|b_t)\}_{t=0}^{T-1}} \mathbb{E}^Q \left[\sum_{t=0}^{T-1} \left\{ \tilde{C}(b_t, u_t) + \lambda \log \frac{Q(u_t|b_t)}{R(u_t|b_t)} \right\} + \tilde{\psi}(b_T) \right]$

Define the value function

$$J_t(b_t) := \inf_{\{Q(u_k|b_k)\}_{k=t}^{T-1}} \left[\sum_{k=t}^{T-1} \int_{\mathcal{U}_k} Q(u_k|b_k) \left\{ \tilde{C}(b_k, u_k) + \lambda \log \frac{Q(u_k|b_k)}{R(u_k|b_k)} \right\} + \tilde{\psi}(b_T) \right].$$

Dynamic Programming in Belief Space

$$J_t(b_t) = \inf_{Q(u_t|b_t)} \int_{\mathcal{U}_t} Q(u_t|b_t) \left\{ \tilde{C}(b_t, u_t) + \int_{b_{t+1}} \tilde{P}(b_{t+1}|b_t, u_t) J_{t+1}(b_{t+1}) + \lambda \log \frac{Q(u_t|b_t)}{R(u_t|b_t)} \right\}$$

(Bellman's Equation)

$$= -\lambda \log \left\{ \int_{\mathcal{U}_t} \exp\left(-\frac{\tilde{C}(b_t, u_t)}{\lambda} \right) \exp\left(-\frac{1}{\lambda} \int_{\mathcal{B}_{t+1}} J_{t+1}(b_{t+1}) \tilde{P}(b_{t+1}|b_t, u_t) \right) R(u_t|b_t) \right\}$$

(Free energy lower bound)

$$\geq -\lambda \log \left\{ \int_{\mathcal{U}_t} \int_{\mathcal{B}_{t+1}} \exp\left(-\frac{\tilde{C}(b_t, u_t)}{\lambda} \right) \exp\left(-\frac{J_{t+1}(b_{t+1})}{\lambda} \right) \tilde{P}(b_{t+1}|b_t, u_t) R(u_t|b_t) \right\}$$

(Jensen's inequality)

$$= -\lambda \log \mathbb{E}^R \exp\left(-\frac{1}{\lambda} \tilde{C}_{t:T}(b_{t:T}, u_{t:T-1}) \right)$$

(Recursive substitutions)

$$\approx -\lambda \log \left\{ \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\frac{1}{\lambda} \tilde{C}_{t:T}(b_{t:T}(i), u_{t:T-1}(i)) \right) \right\}$$

(Monte Carlo)

If $\tilde{P}(b_{t+1}|b_t, u_t)$ is a point mass (deterministic transition), then Jensen's inequality holds with equality.

Deceptive Path Planning and Hypothesis Testing

Covert Vehicle Misguidance and Its Detection

J. Bhatti and T. E. Humphreys, “Hostile control of ships via false GPS signals: Demonstration and detection,” NAVIGATION: Journal of the Institute of Navigation, 2017.

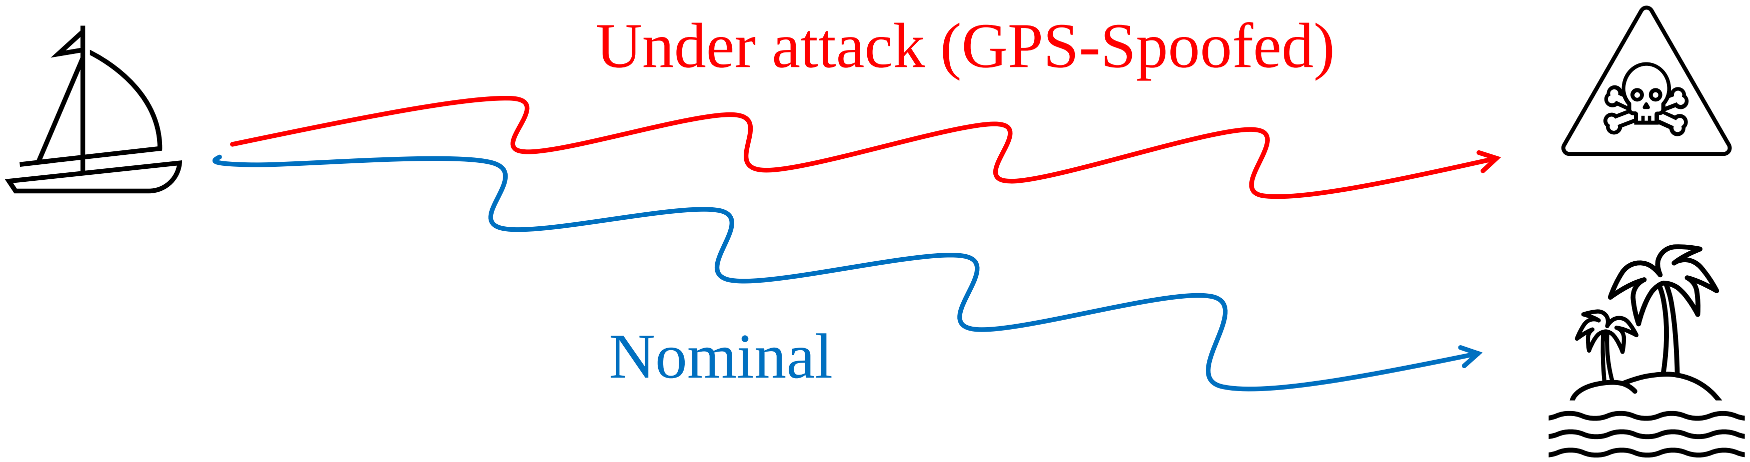

Inspired by the GPS spoofing demonstration, we formulate a stochastic zero-sum game to analyze the competition between

- Attacker, who tries to misguide the vehicle to an unsafe region covertly, and
- Detector, who tries to detect the attack signal based on the observed trajectory of the vehicle.

System Model

Vehicle trajectory with attack input $\theta(t)$

$$dx(t) = \theta(t)dt + dw(t), x(0) = 0, 0 \leq t \leq T$$

Terminal state:

$$x(T) > Td : \text{ Unsafe}$$
$$x(T) \leq Td : \text{ Safe}$$

Fix an open-loop attack policy $\theta(t), 0 \leq t \leq T$. The probability of the terminal state being unsafe

$$\gamma(\theta) := \Phi\left(\frac{1}{\sqrt{T}}\int_0^T \theta(t)dt - \sqrt{T}d\right).$$

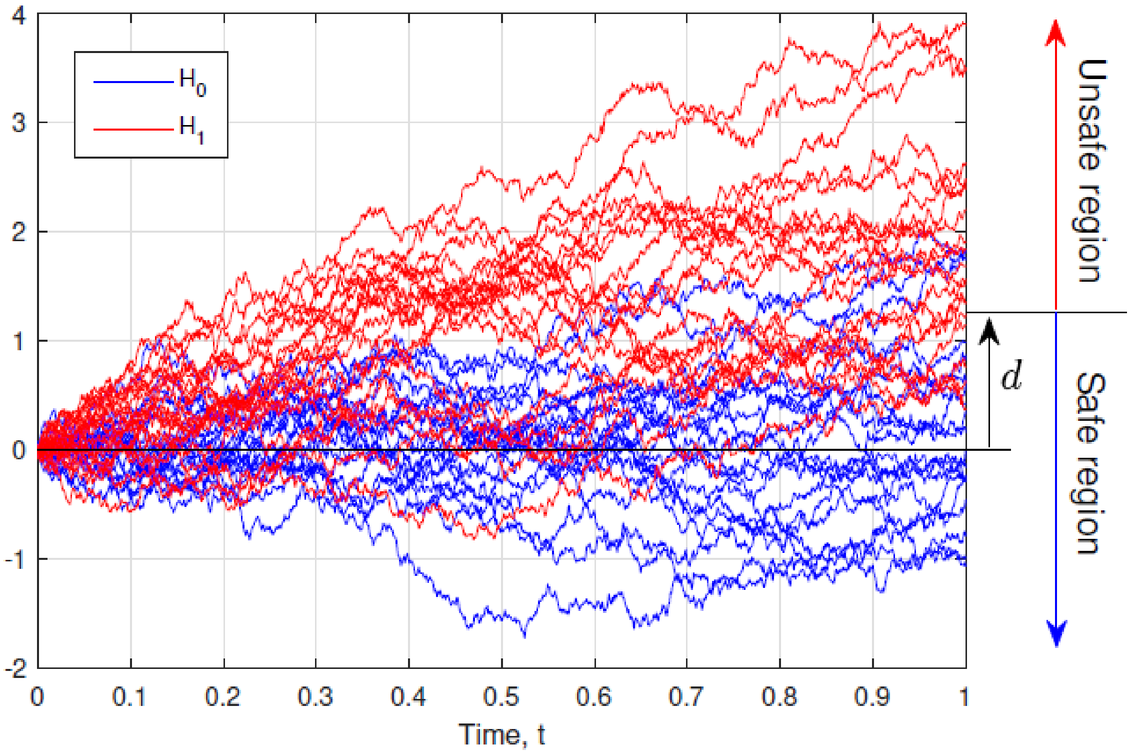

The attacker wants $\gamma(\theta) \geq c(> \frac{1}{2})$, or $\int_0^T \theta(t)dt \geq \sqrt{T}\Phi^{-1}(c) + Td$.

Hypothesis Testing

$$H_0 : \theta(t) = 0 \ \ \forall t \in [0, T] \quad \text{(Null hypothesis)}$$

$$H_1 : \int_0^1 \theta(t)dt \geq \sqrt{T}\Phi^{-1}(c) + Td \quad \text{(Alternative)}$$

The role of the detector is to design a hypothesis testing algorithm $\phi : C[0, T] \to \{0, 1\}$ such that

$$\phi(x) = \begin{cases} 0 & H_0 \text{ is accepted} \\ 1 & H_1 \text{ is accepted.} \end{cases}$$

The quality of a testing algorithm ϕ is measured in terms of:

$$\alpha(\phi) := \Pr\{\phi(x) = 1 \mid H_0 \text{ is true}\} \quad \text{(Probability of a false alarm)}$$

$$\beta(\theta, \phi) := \Pr\{\phi(x) = 0 \mid H_1 \text{ is true}\} \quad \text{(Probability of a detection failure)} .$$

Consider a mini-max game and its dual:

$$\boxed{p^* = \min_{\phi:\alpha(\phi)\leq\epsilon} \max_{\theta:\gamma(\theta)\geq c} \beta(\theta, \phi) \quad \text{and} \quad d^* = \max_{\theta:\gamma(\theta)\geq c} \min_{\phi:\alpha(\phi)\leq\epsilon} \beta(\theta, \phi)}$$

Neyman-Pearson Lemma

We want to find:

- The saddle point policy (θ^*, ϕ^*)
- The value of the game $\beta(\theta^*, \phi^*)$ and its error exponent as a function of T

Lemma: For a fixed θ, the testing algorithm $\phi : C[0,T] \to \{0,1\}$ that minimizes $\beta(\theta, \phi)$ subject to the constraint $\alpha(\phi) \leq \epsilon$ is given by

$$\phi(x) = \begin{cases} 0 & \text{if } \frac{d\mu}{d\mu_\theta}(x) \leq \lambda^* \\ 1 & \text{if } \frac{d\mu}{d\mu_\theta}(x) > \lambda^* \end{cases} \tag{1}$$

where the likelihood ratio is computed by Girsanov's theorem

$$\frac{d\mu}{d\mu_\theta}(x) = \exp\left\{ \int_0^t \theta(s) dx(s) - \frac{1}{2} \int_0^t \theta^2(s) ds \right\} \tag{2}$$

and $\lambda^* > 0$ is a constant satisfying $\alpha(\phi) = \epsilon$.

Saddle Point Policy

Theorem: The following pair of policies form a (unique) saddle point of the zero-sum game:

$$\theta^*(t) = \bar{\theta} := \frac{1}{\sqrt{T}}\Phi^{-1}(c) + d \;\; \forall t \in [0, T] \tag{1}$$

$$\phi^*(x) = \begin{cases} 0 & \text{if } x(T) \leq \sqrt{T}\Phi^{-1}(1-\epsilon) \\ 1 & \text{if } x(T) > \sqrt{T}\Phi^{-1}(1-\epsilon). \end{cases} \tag{2}$$

Moreover, the value of the game is $\beta(\theta^*, \phi^*) = \Phi(\Phi^{-1}(1-\epsilon) - \Phi^{-1}(c) - \sqrt{T}d)$.

Observation:

- The max-min policy (the most covert attack) is a constant bias injection $\theta(t) = \bar{\theta}$, where the constant $\bar{\theta}$ is chosen to be the smallest value satisfying $\gamma(\theta) \geq c$.
- The minimax policy $\phi^*(x)$ (i.e., the most powerful hypothesis test) is a likelihood ratio test, and only examines the final value $x(T)$ of the observed sample path x.

Error Exponent and Finite Sample Analysis

Under the constraint $\alpha(\phi) \leq \epsilon$, we have $\beta(\theta^*, \phi^*) \to 0$ as $T \to \infty$, i.e., the hypothesis testing problem becomes easier as the horizon length T grows. How does $\beta(\theta^*, \phi^*)$ behave as a function of T?

$$-\log \beta(\theta^*, \phi^*) \geq T\bar{D}(\mu \| \mu_{\theta^*}) + \sqrt{T}\sqrt{\bar{V}(\mu \| \mu_{\theta^*})}\Phi^{-1}(\epsilon) + \text{const.}$$

where $\bar{D}(\mu \| \mu_{\theta^*}) = \frac{1}{2}\bar{\theta}^2$ is the relative entropy rate and $\bar{V}(\mu \| \mu_{\theta^*}) = \bar{\theta}^2$ is the variance rate.

Observation:

- The first-order term is reminiscent of the classical Stein's lemma.
- The second-order term provides a tighter approximation in the regime of finite T [K. Li, Annals of Statistics, 2014].
- Further improvements are available [Lungu and Kontoyiannis, ISIT 2024].

See Tanaka et al. "Covert Vehicle Misguidance and Its Detection: A Hypothesis Testing Game over Continuous-Time Dynamics" 2024 for more details.

Generalization

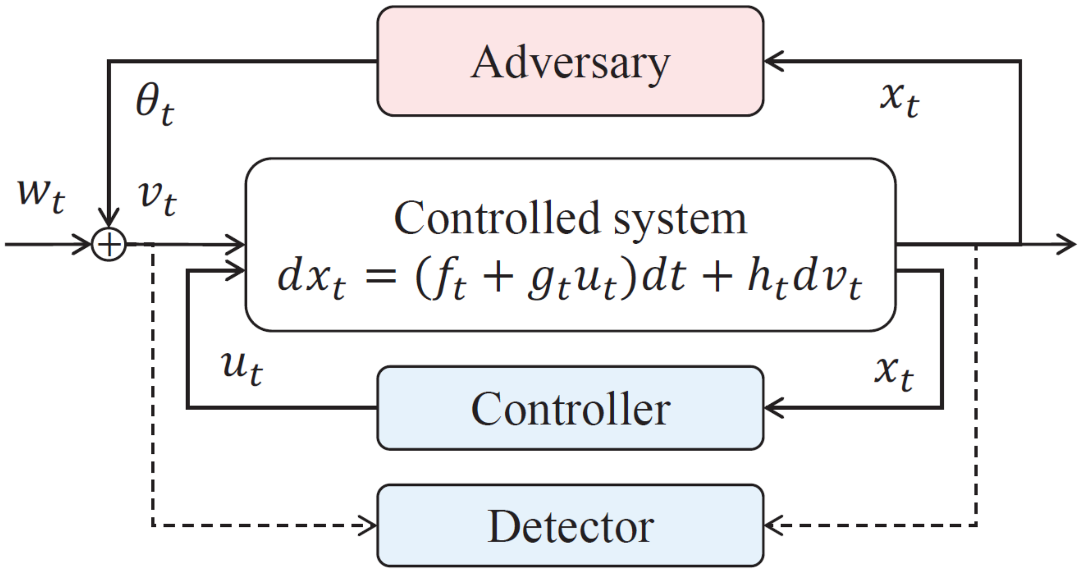

$$dv_t = dw_t \quad \text{(No attack)}$$
$$dv_t = \theta_t dt + dw_t \quad \text{(Under attack).}$$

- More general nonlinear dynamics
- Probability of successful attack ⇒ More general cost functions
- Controller/detector can now apply a legitimate control input u_t to combat with the noise w_t and the potential attack input θ_t.

$$\min_u \max_\theta \mathbb{E}^P \left[\int_0^T \ell(x_t, u_t) dt \right] - \lambda \times (\textit{Attack Detectability})$$

Stealthy Attack on Control Systems

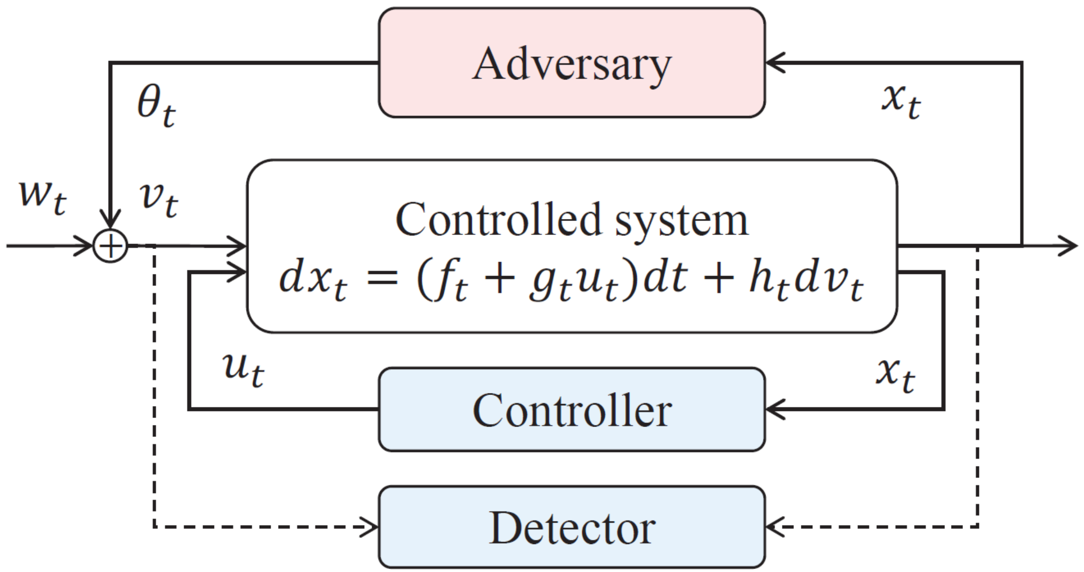

[Bai, Pasqualetti & Gupta "Data-Injection Attacks in Stochastic Control Systems: Detectability and Performance Tradeoffs“ 2017] introduced the KL divergence as the stealthiness measure (justified by Stein's Lemma).

KLD-constrained optimal control problem for the optimal stealthy attack synthesis.

$$\min_u \max_\theta \mathbb{E}^P \left[\int_0^T \ell(x_t, u_t) dt \right] - \lambda D(P||Q)$$

Is this problem computationally tractable?

Minimax KL Control, Risk-sensitive Control, Nonlinear H_∞ Control, and Path Integrals

Risk-sensitive control

$$\min_u \lambda \log \mathbb{E}^Q \left[\exp\left(\frac{1}{\lambda} \int_0^T \ell(x_t, u_t) dt \right) \right]$$

Legendre Duality

Minimax KL control

$$\min_u \max_\theta \mathbb{E}^P \left[\int_0^T \ell(x_t, u_t) dt \right] - \lambda D(P||Q)$$

- Jacobson 1973.
- Petersen, James, & Dupuis 2000.

Girsanov Theorem

Nonlinear H_∞ control

$$\min_u \max_\theta \mathbb{E}^P \int_0^T \left\{ c(x_t) + \frac{1}{2} u_t^T R u_t - \frac{\lambda}{2} \|\theta_t\|^2 \right\}$$

See Patil, Karabag, Tanaka, Topcu "Simulator-Driven Deceptive Control via Path Integral Approach" CDC 2023 for more details.

Q & A