Information-Theoretic Approaches to Active Sensing: Theory and Practice

Haruki Nishimura
Ph.D. Candidate in Aeronautics and Astronautics, Stanford University, USA

2nd Workshop on Informative Path Planning and Adaptive Sampling @ RSS 2019
June 22nd, 2019

MSL

Stanford University

Multi-Robot Systems Lab

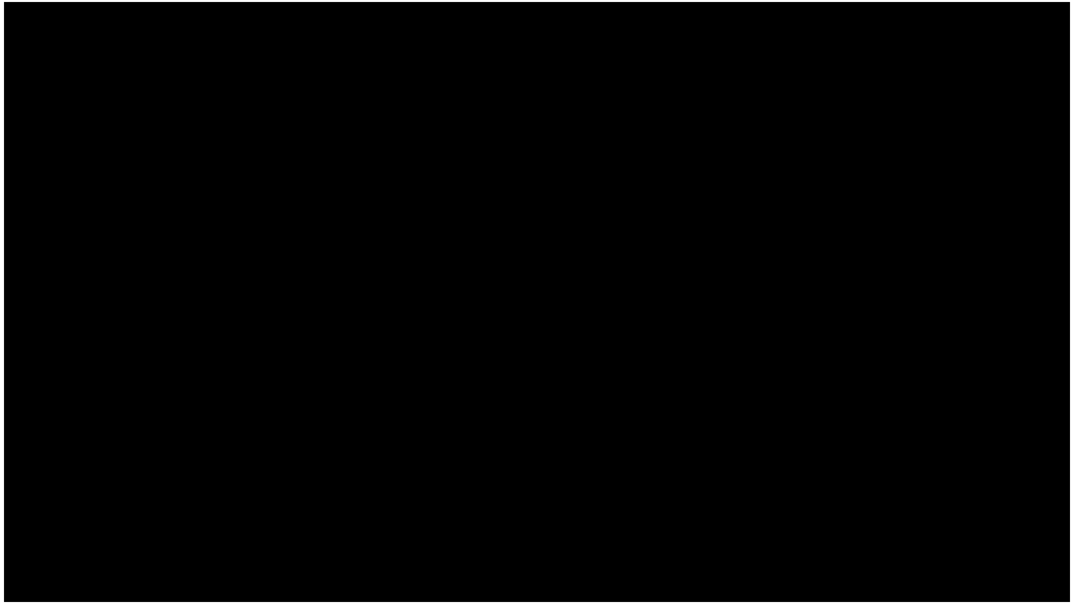

We develop
theory and algorithms for
control, planning, and estimation of
multiple mobile robots in
interactive environments.

Prof. Mac. Schwager

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing

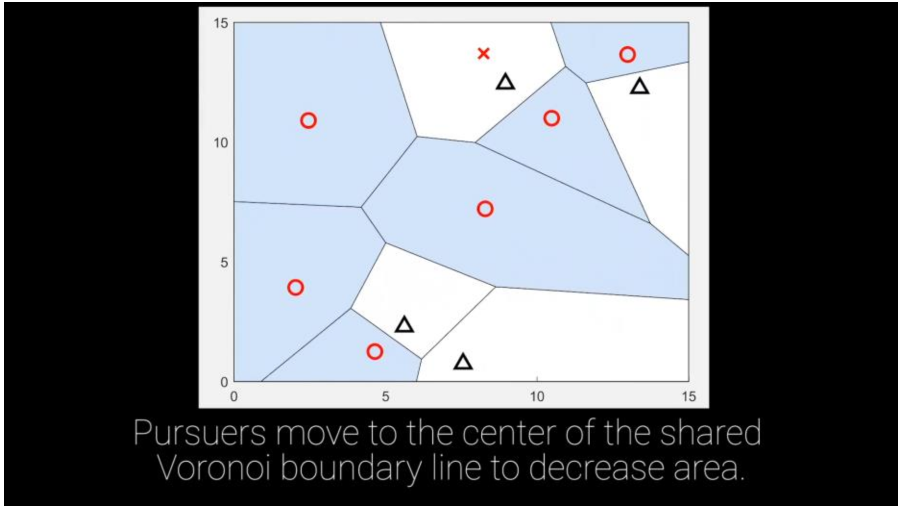

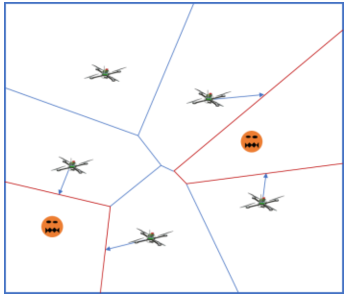

Dr. Alyssa Pierson

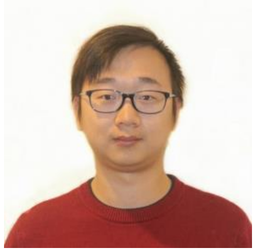

Dr. Zijian Wang

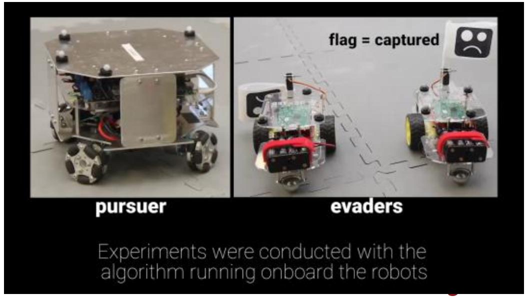

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing

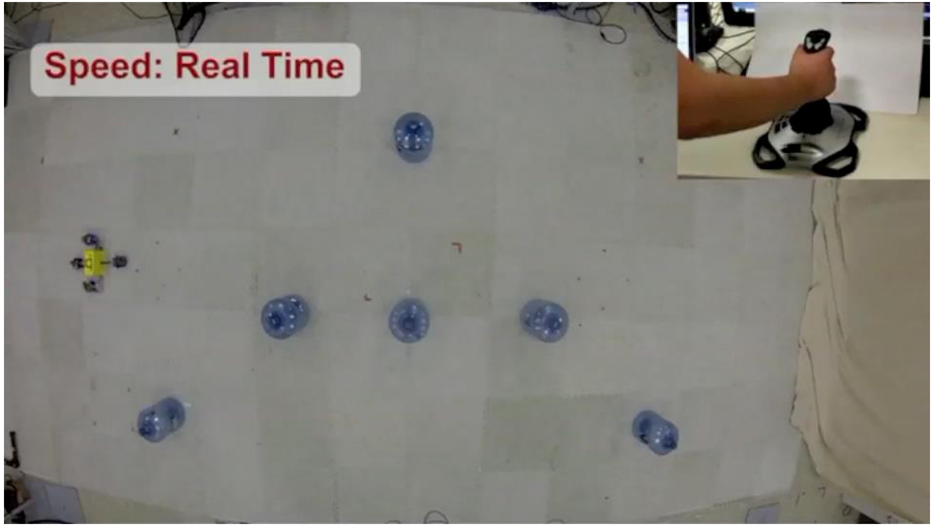

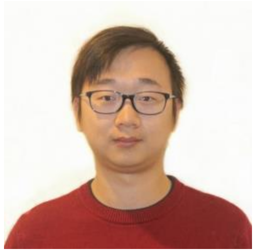

Dr. Zijian Wang

$$\sum_{i=1}^{N} F_i = ma + \mu v$$

acceleration and velocity

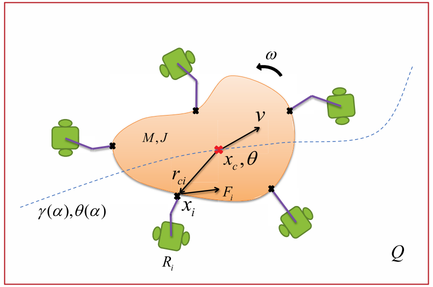

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing

Preston Culbertson

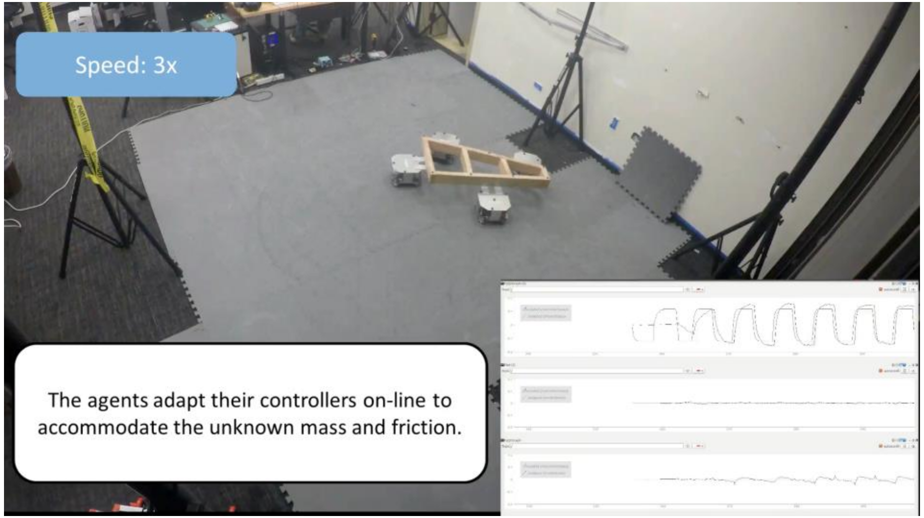

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing

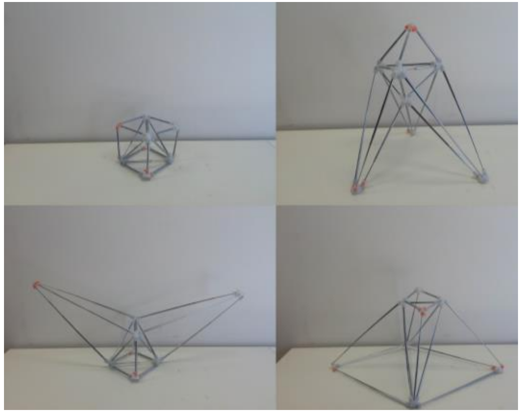

Nathan Usevitch

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing


Dr. Riccardo Spica

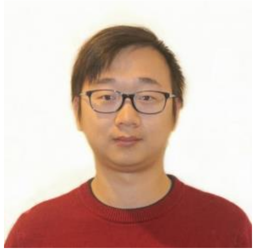
Dr. Zijian Wang

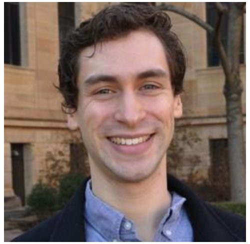
Eric Cristofalo

Multi-Robot Systems Lab

Recent projects include:

- Intercepting Rogue Robots
- Coordinating without Talking
- Distributed Shape Shifting Robots
- Game-Theoretic Planning for Racing

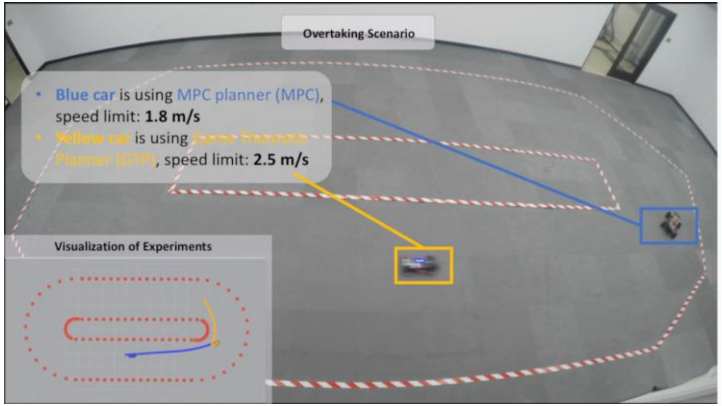

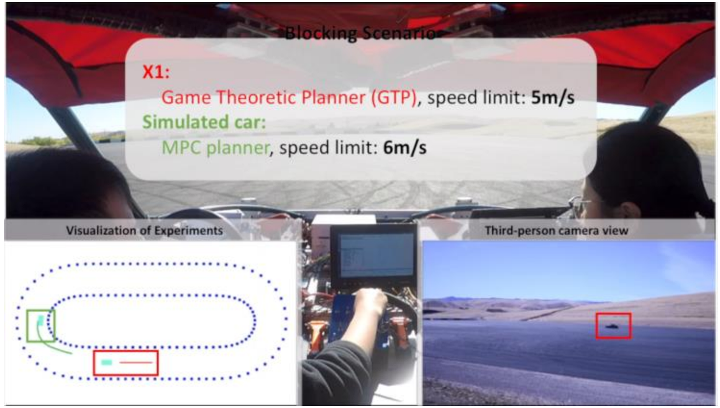

Mingyu Wang

Today's Main Focus

Algorithms for Information-Theoretic Active Sensing:

Rapidly-Exploring Random Cycles (RRC/RRC*)

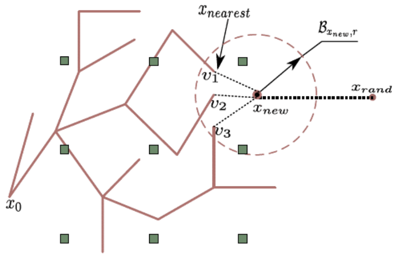

- Discrete-Time Linear Gaussian Systems
- Persistent Surveillance/Monitoring

[X. Lan and M. Schwager, IEEE T-RO 2016]

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW)

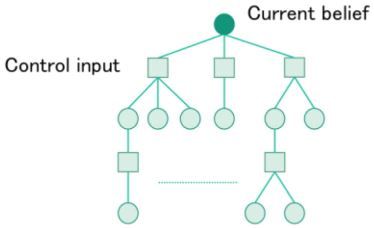

- Discrete-Time Nonlinear Non-Gaussian Systems
- Active Intent Inference

[H. Nishimura and M. Schwager, ICRA 2018]

Sequential Action Control for Belief Space Planning (SACBP)

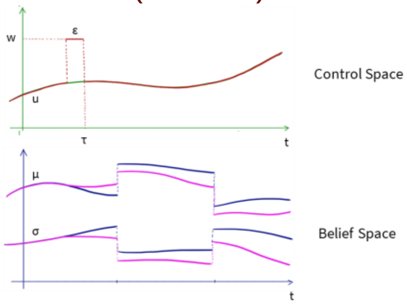

- Continuous-Time Nonlinear Non-Gaussian Systems
- Multi-Target Tracking, etc.

[H. Nishimura and M. Schwager, WAFR 2018]

Information-Theoretic Active Sensing

1. Model a time-varying phenomenon as a partially-observable stochastic process.

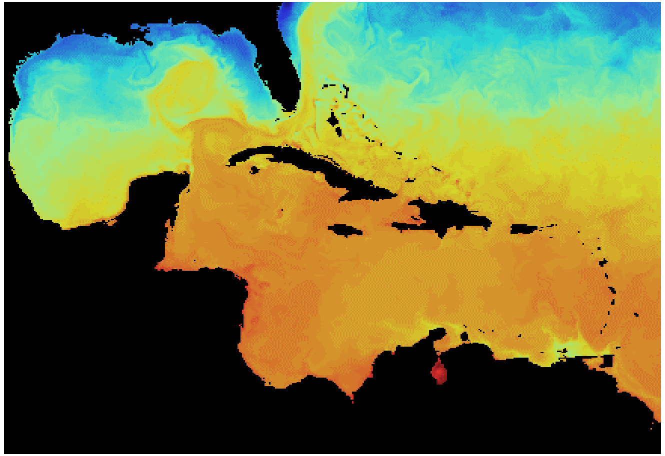

Spatio-Temporal Field in the Caribbean Sea Surface

Hidden Markov Model (HMM) with Linear Gaussian dynamics

$$\phi_t(q) = C(q)a_t$$

$$a_{t+1} = Aa_t + \omega_t \qquad \omega_t \sim \mathcal{N}(0, Q)$$

$$y_t = \phi(x_t) + \nu_t \qquad \nu_t \sim \mathcal{N}(0, R)$$

[X. Lan and M. Schwager, IEEE T-RO 2016]

Information-Theoretic Active Sensing

2. Track the "belief state" using Bayesian filtering techniques.

The Kalman Filter (online estimation)

$$\hat{\phi}_{t+1} = A\hat{\phi}_t + A\Sigma_t C(p_t)^T \left(C(p_t)\Sigma_t C(p_t)^T + R\right)^{-1} \times (y_t - C(p_t)\hat{\phi}_t)$$

$$\Sigma_{t+1} = A\Sigma_t A^T - A\Sigma_t C(x_t)^T (C(x_t)\Sigma_t (C(x_t)^T + R)^{-1} C(x_t)\Sigma_t A^T + Q$$

Hidden Markov Model (HMM) with Linear Gaussian dynamics

$$\phi_t(q) = C(q)a_t$$

$$a_{t+1} = Aa_t + \omega_t \qquad \omega_t \sim \mathcal{N}(0, Q)$$

$$y_t = \phi(x_t) + \nu_t \qquad \nu_t \sim \mathcal{N}(0, R)$$

[X. Lan and M. Schwager, IEEE T-RO 2016]

Information-Theoretic Active Sensing

3. Devise an algorithm to optimize an information-theoretic cost associated with the evolving belief states.

Rapidly-Exploring Random Cycles (RRC/RRC*)

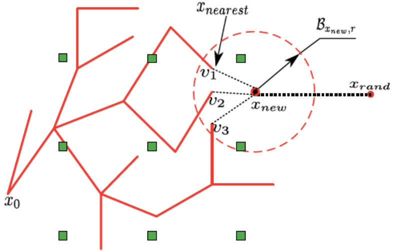

- Linear Gaussian Systems
- Persistent Surveillance/Monitoring

[X. Lan and M. Schwager, IEEE T-RO 2016]

Winner of 2016 King-Sun Fu Memorial IEEE Transactions on Robotics Best Paper Award

Dr. Xiaodong Lan

Rapidly Exploring Random Cycles (RRC/RRC*)

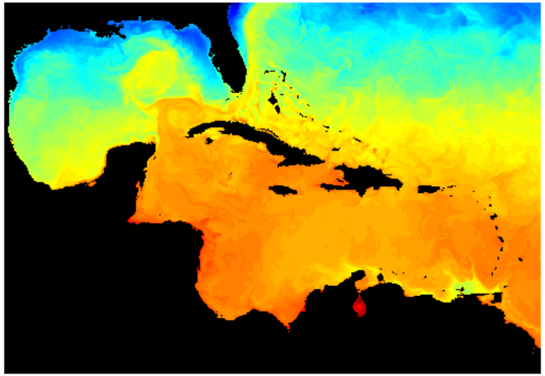

Spatio-Temporal Field

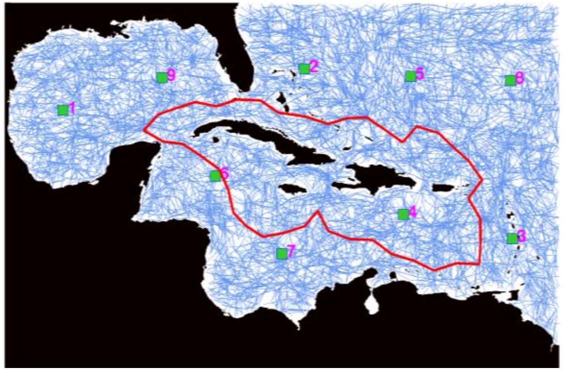

Single Robot Periodic Trajectory

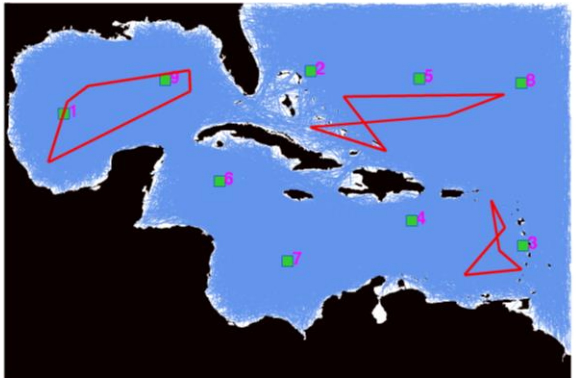

Multi-Robot Periodic Trajectory

Rapidly Exploring Random Cycles (RRC/RRC*)

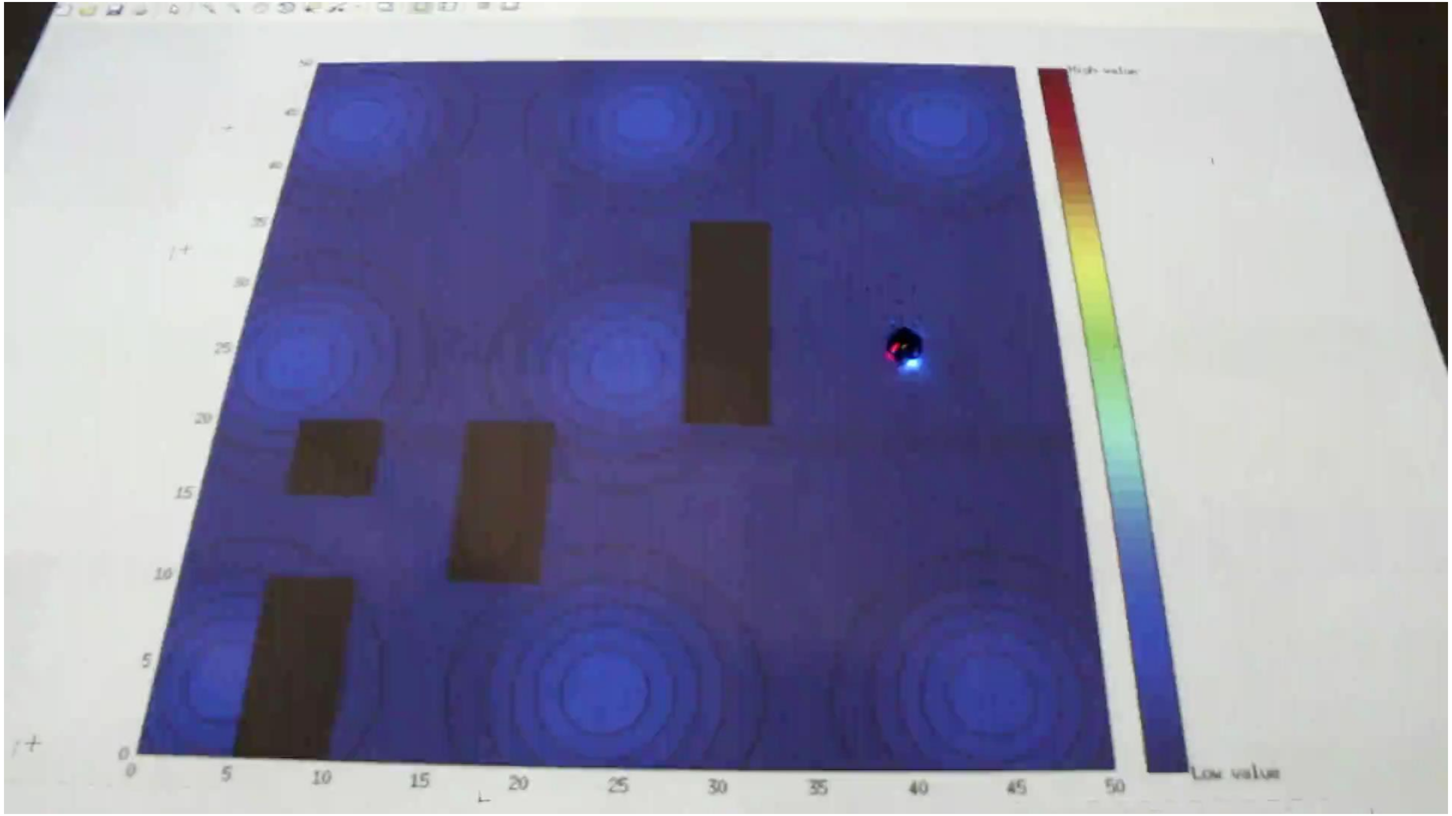

Tracking the Belief State

The Kalman Filter
(online estimation)

$$\hat{\phi}_{t+1} = A\hat{\phi}_t + A\Sigma_t C(p_t)^T \left(C(p_t)\Sigma_t C(p_t)^T + R\right)^{-1} \times (y_t - C(p_t)\hat{\phi}_t)$$

$$\Sigma_{t+1} = A\Sigma_t A^T - A\Sigma_t C(x_t)^T (C(x_t)\Sigma_t (C(x_t)^T + R)^{-1} C(x_t)\Sigma_t A^T + Q$$

Hidden Markov Model (HMM)
with Linear Gaussian dynamics

$$\phi_t(q) = C(q)a_t$$

$$a_{t+1} = Aa_t + \omega_t \qquad \omega_t \sim \mathcal{N}(0, Q)$$

$$y_t = \phi(x_t) + \nu_t \qquad \nu_t \sim \mathcal{N}(0, R)$$

Covariance Matrix evolves deterministically! (as a function of robot state x)

Rapidly-Exploring Random Cycles (RRC)

Sampling-based algorithm for offline motion planning,
inspired by RRT [Lavalle and Kuffner 2001] and RRT* [Karaman and Frazzoli 2011]

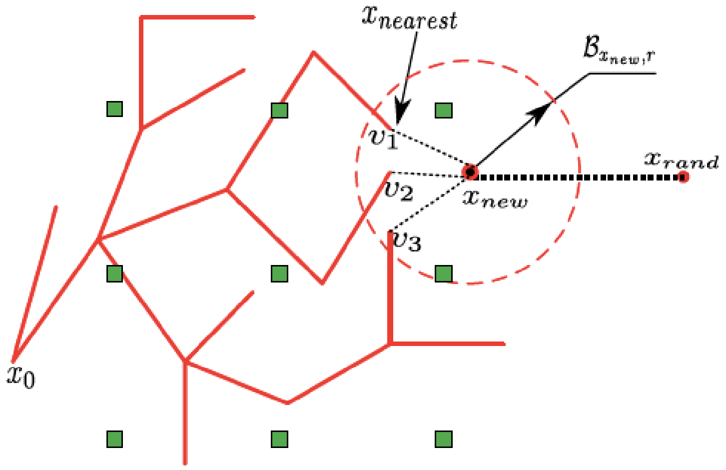

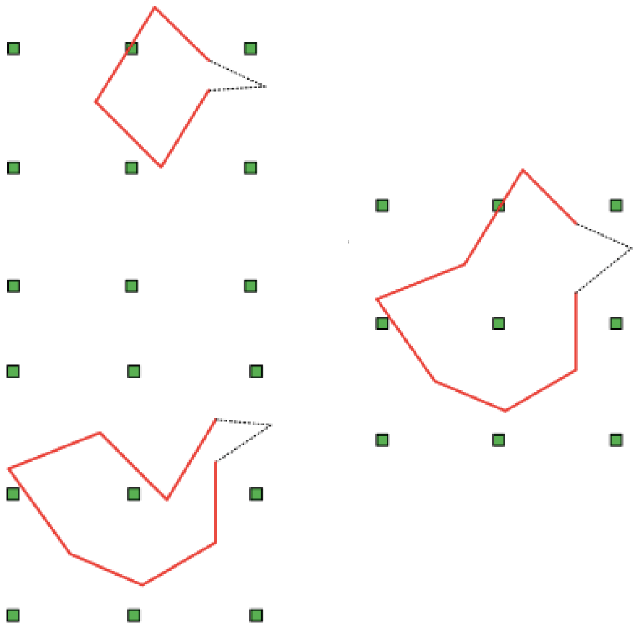

Cycle = Spanning Tree + Single Edge

Caribbean Sea Simulation Results

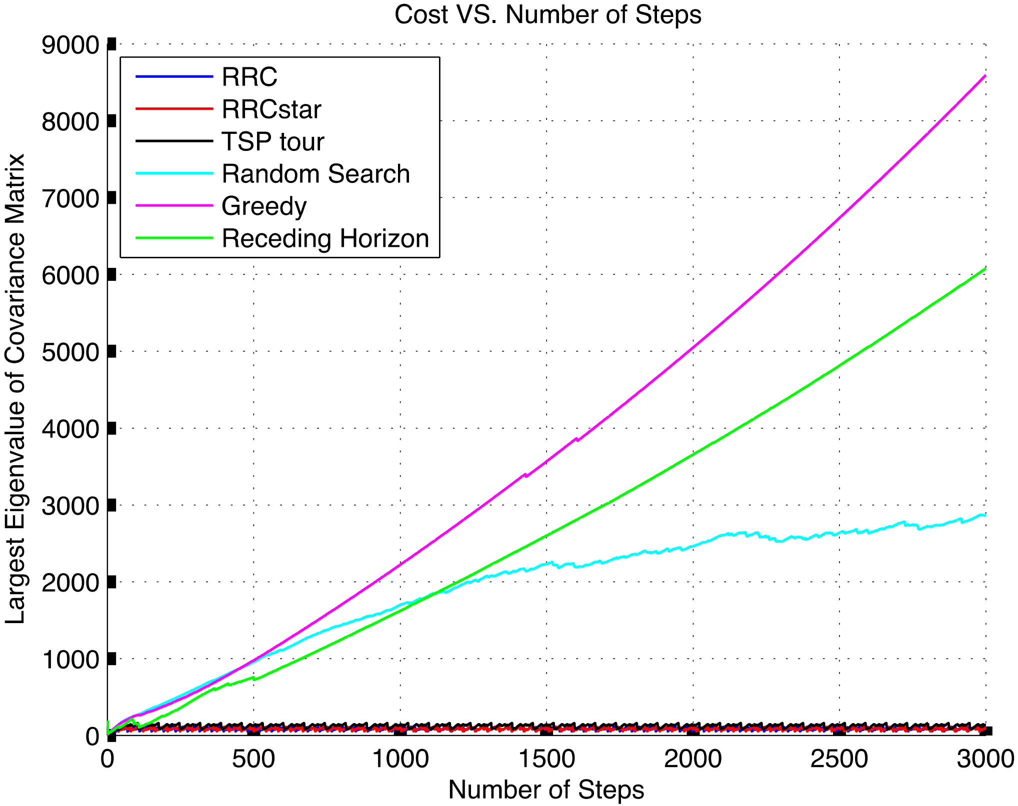

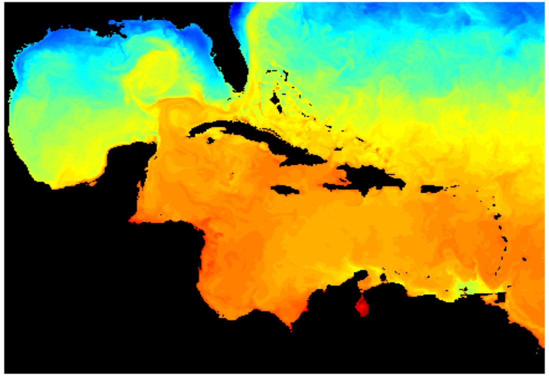

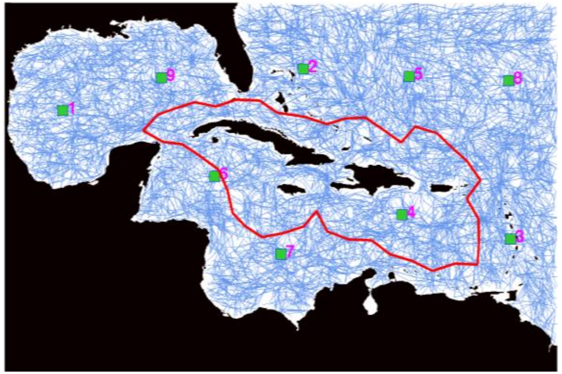

RRC* Periodic Trajectory

Multi-Robot RRC

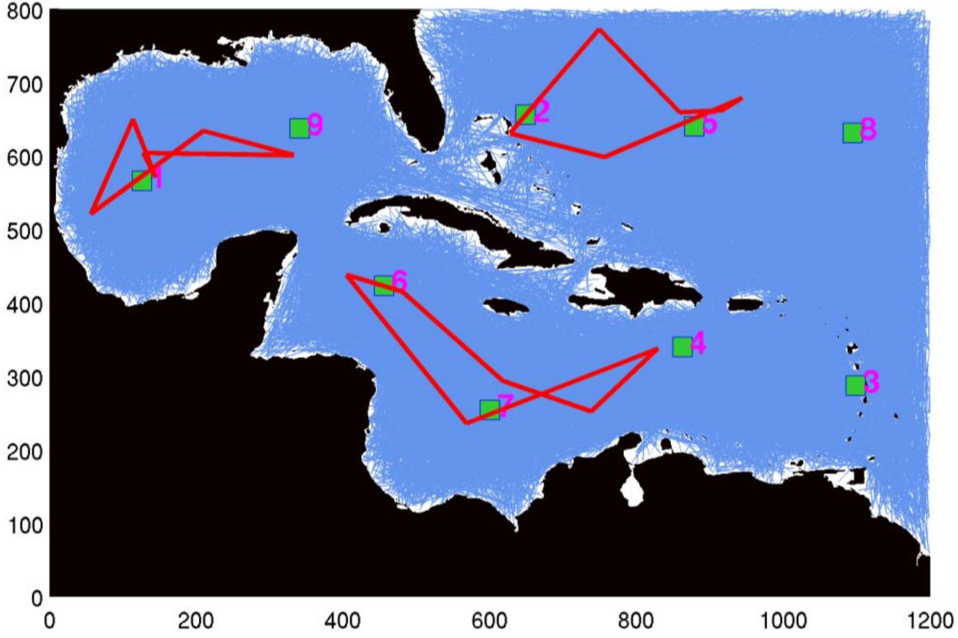

Optimal cycle in 6D space project onto 2D space

Joint Measurement

$$y_t = (y_t^1, \ldots, y_t^n)$$

Plan in joint state space

$$x_{1:T} = (x_{1:T}^1, \ldots, x_{1:T}^n)$$

Information-Theoretic Active Sensing

Rapidly-Exploring Random Cycles (RRC/RRC*)

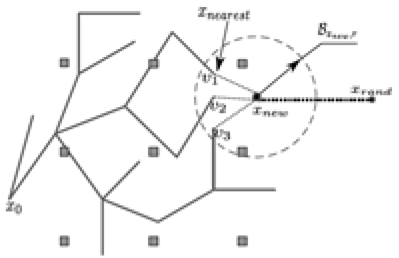

- Discrete-Time Linear Gaussian Systems
- Persistent Surveillance/Monitoring

[X. Lan and M. Schwager, IEEE T-RO 2016]

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW)

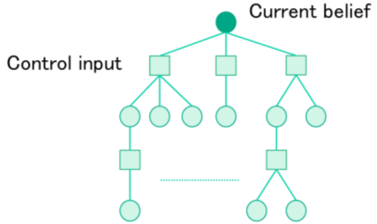

- Discrete-Time Nonlinear Non-Gaussian Systems
- Active Intent Inference

[H. Nishimura and M. Schwager, ICRA 2018]

Active Intent Inference Problem

How to resolve message ambiguity
without knowing the relative pose?

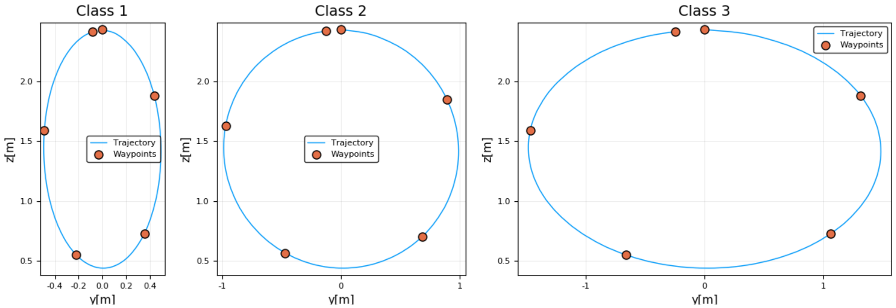

Trajectory Class = "Message" or "Intent"

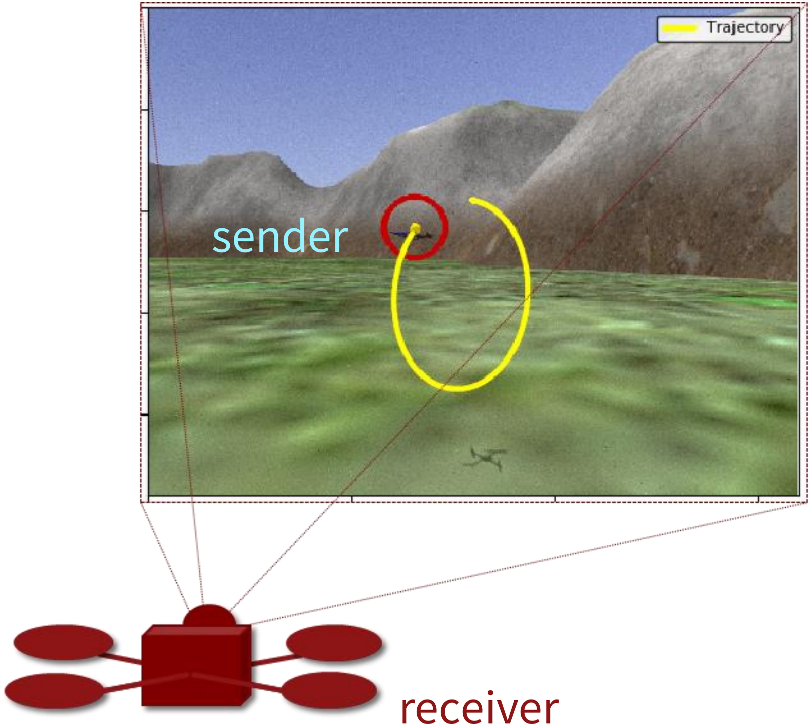

The receiver actively changes its viewing positions to disambiguate between trajectory hypotheses while estimating the relative pose.

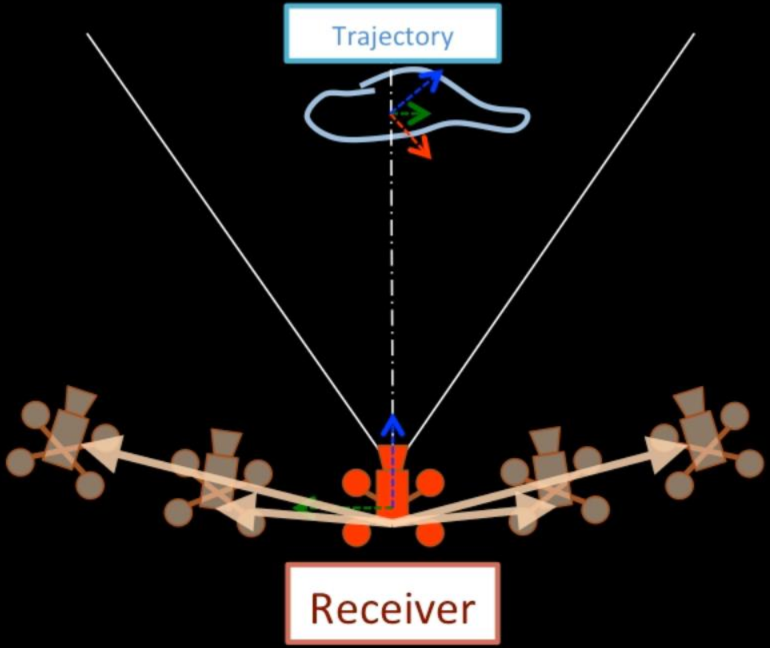

Intent Inference Process Model

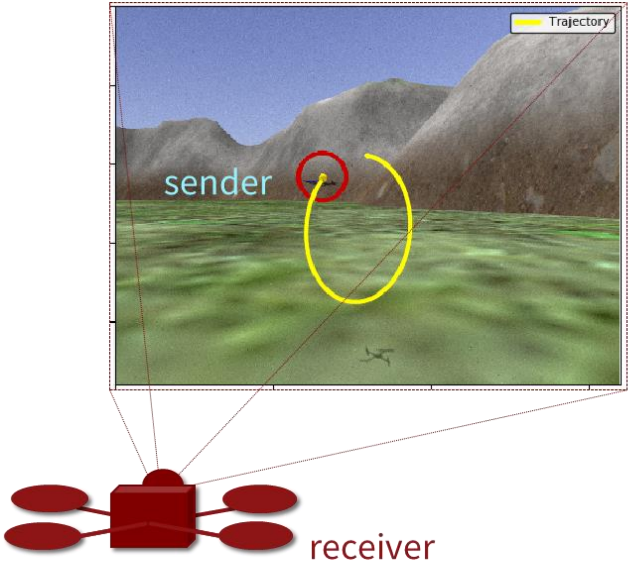

Hidden Markov Model (HMM)
with Categorical Latent Variable *m*

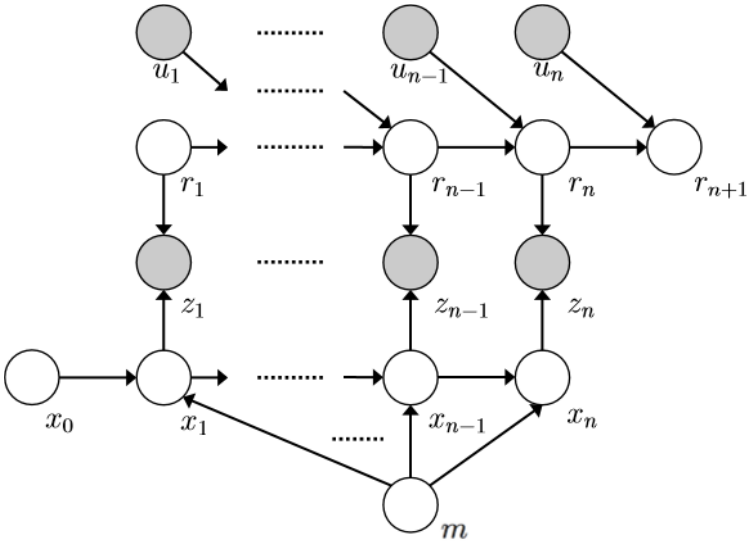

Tracking the Belief State

Multi-Hypothesis Extend Kalman Filter (MHEKF)

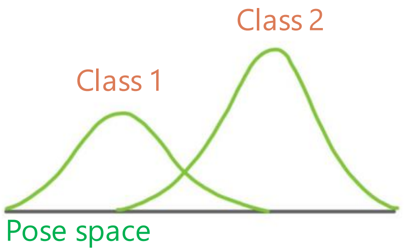

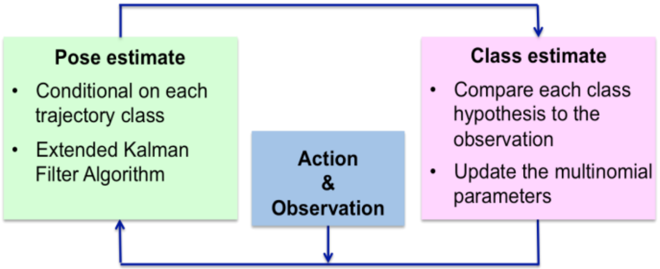

Hidden Markov Model (HMM) with Categorical Latent Variable *m*

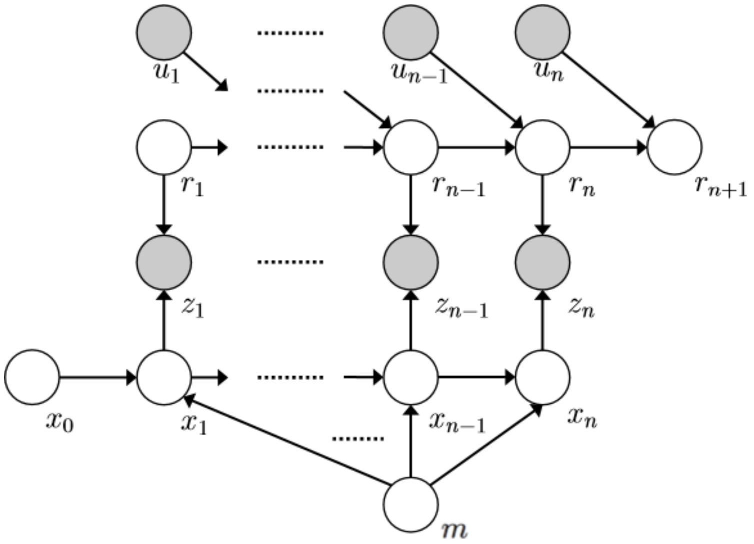

Dynamic Programming in Belief Space

Cost Objective: Entropy of trajectory class distribution in T steps

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW) [Couëtoux et al. 2011] in the belief space

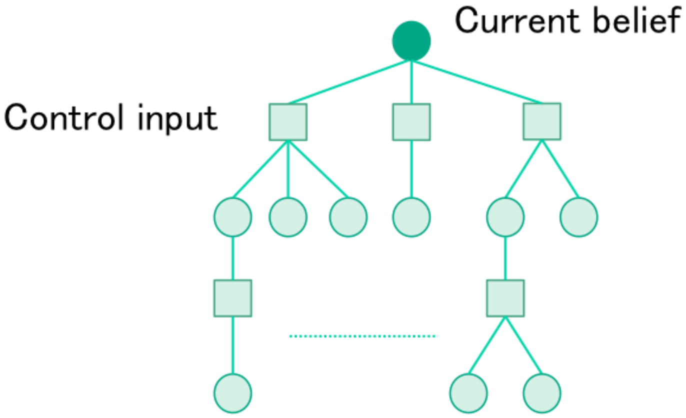

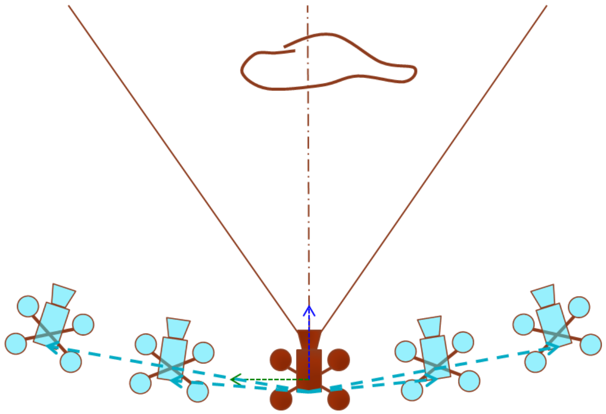

Dynamic Programming in Belief Space

Cost Objective: Entropy of trajectory class distribution in T steps

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW) [Couëtoux et al. 2011] in the belief space

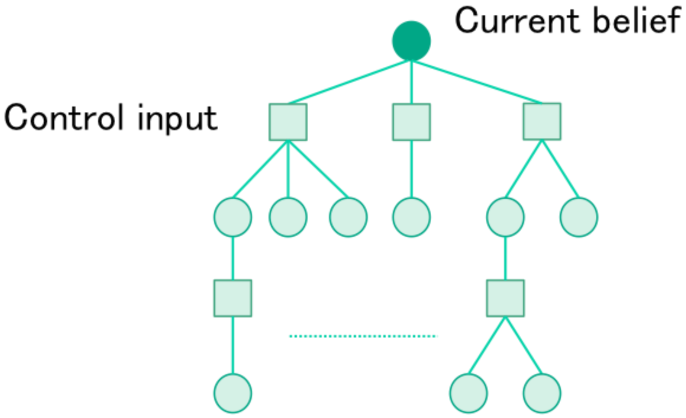

- DPW makes MCTS effective in continuous state spaces
- Online search for receding-horizon execution

Simulation Results in ROS-Gazebo

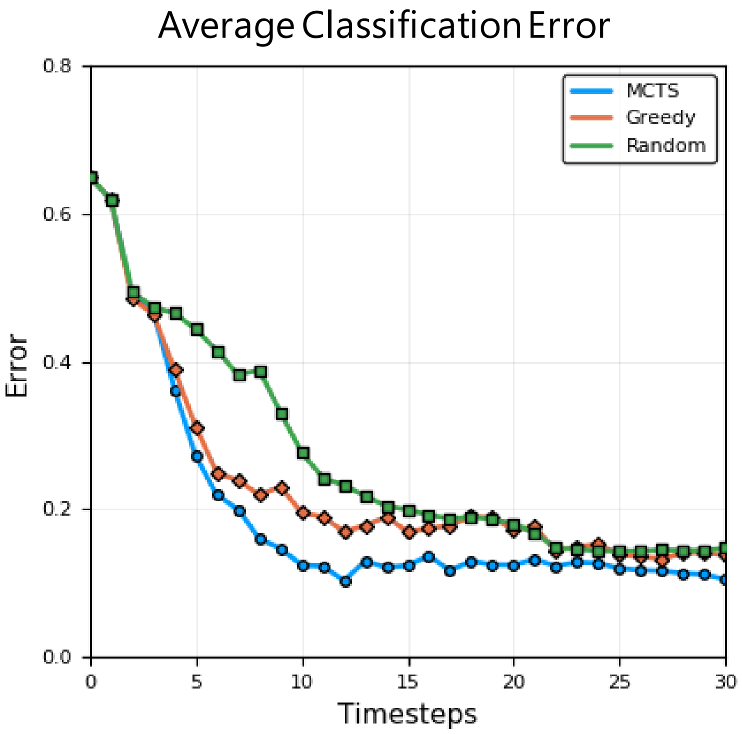

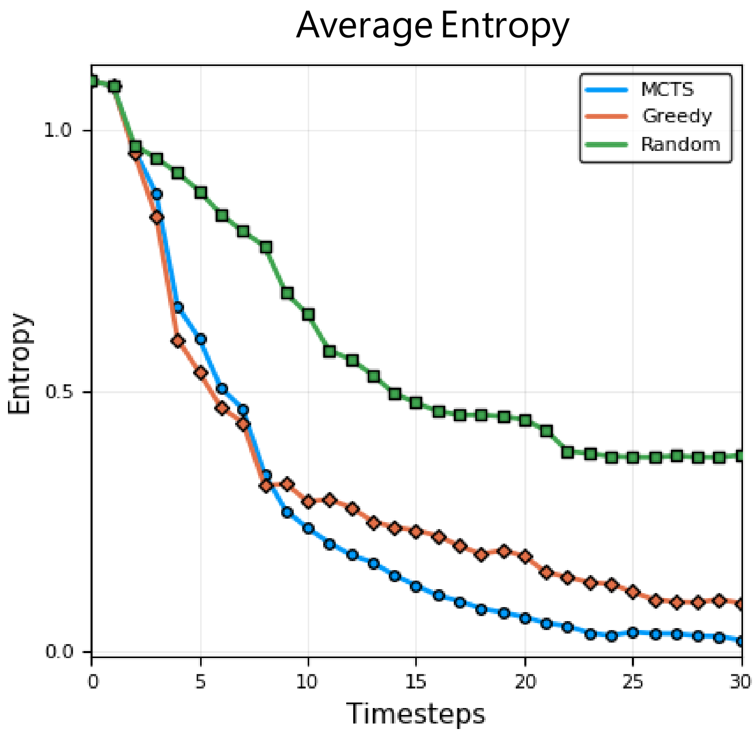

Information-Theoretic Active Sensing

Rapidly-Exploring Random Cycles (RRC/RRC*)

[X. Lan and M. Schwager, IEEE T-RO 2016]

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW)

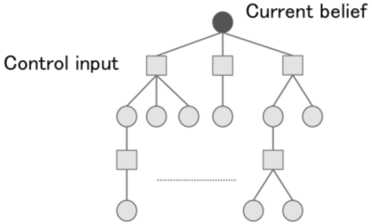

- Discrete-Time Nonlinear Non-Gaussian Systems
- Active Intent Inference

[H. Nishimura and M. Schwager, ICRA 2018]

Sequential Action Control for Belief Space Planning (SACBP)

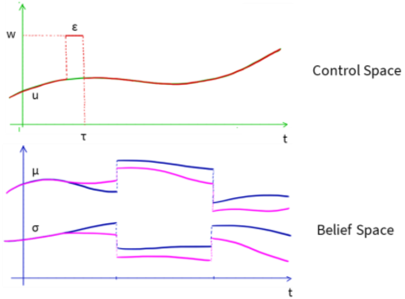

- Continuous-Time Nonlinear Non-Gaussian Systems
- Multi-Target Tracking, etc.

[H. Nishimura and M. Schwager, WAFR 2018]

Continuous-Time Belief Space Planning

Active Multi-Target Tracking with Range-only Measurements

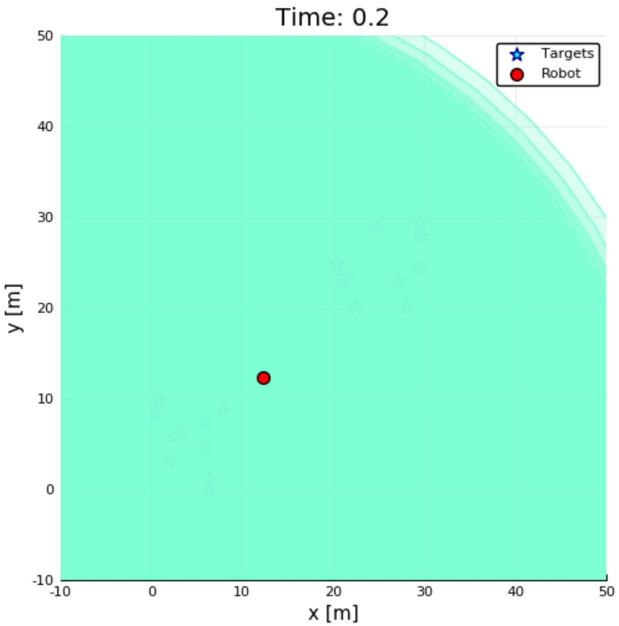

Continuous-Time Belief Space Planning

Active Multi-Target Tracking with Range-only Measurements

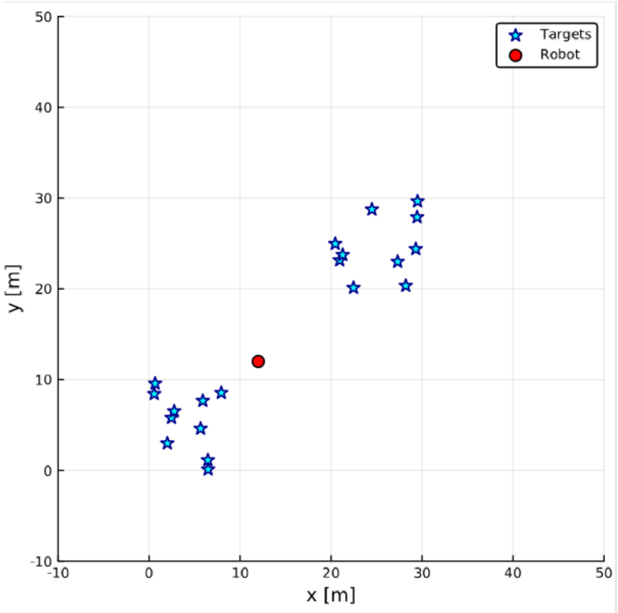

Object Manipulation under Model Uncertainty

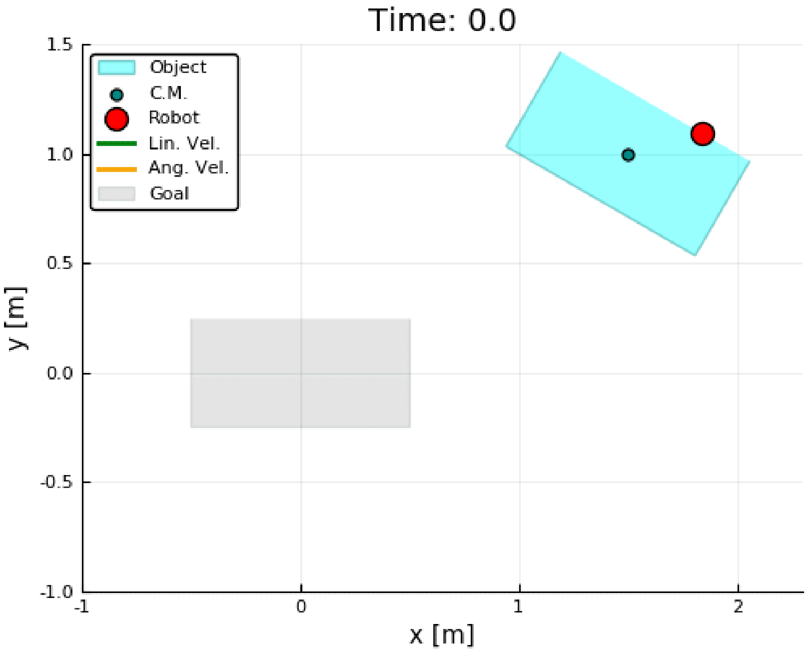

Key Idea

View belief dynamics as a "hybrid system with time-driven switching"

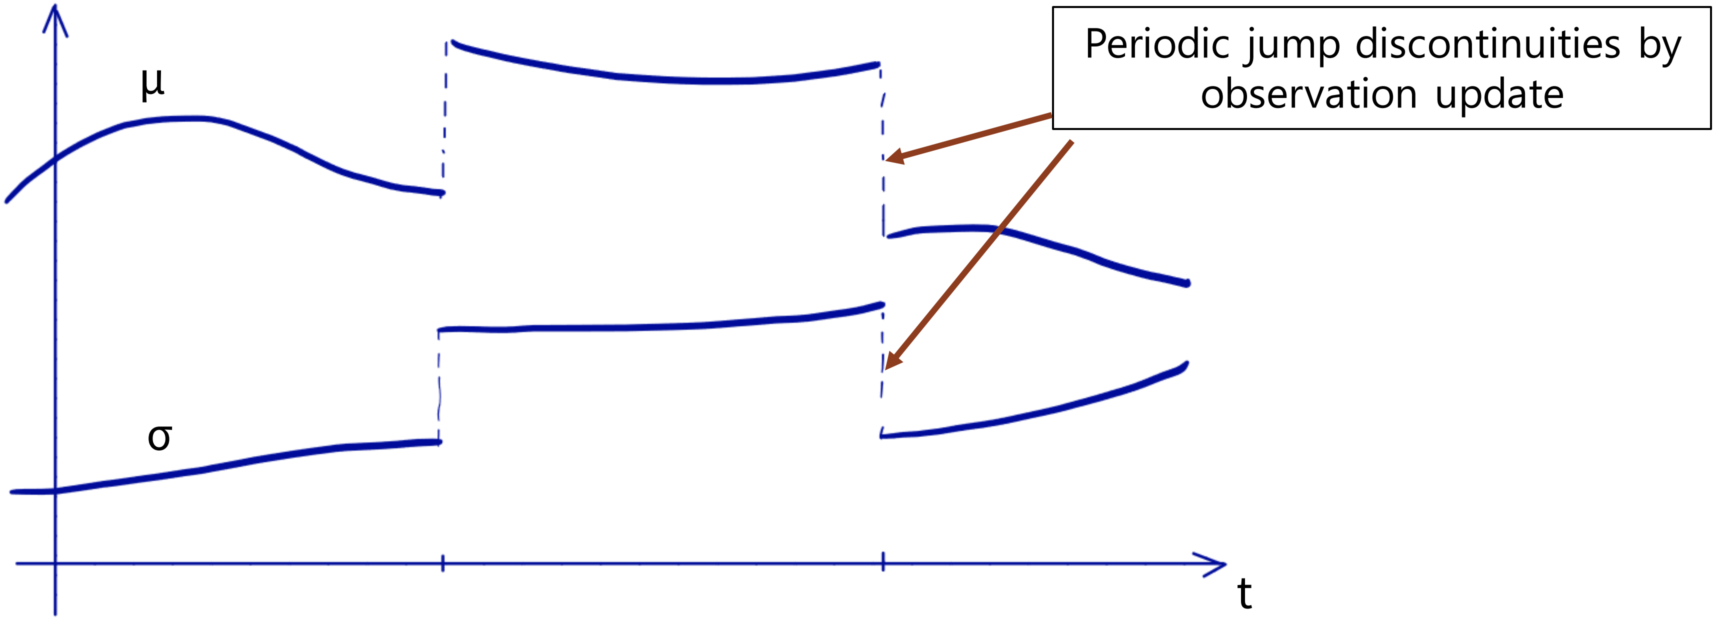

Example: 1D Gaussian Belief Case (Continuous-Discrete Gaussian Filter)

Key Idea

Optimally "perturb" a given nominal control policy

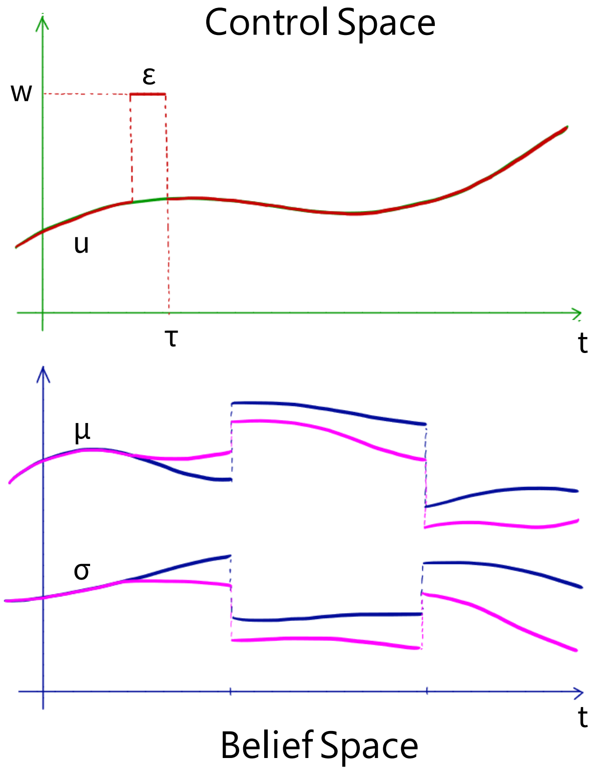

$$u_w(t) = \begin{cases} w & \text{if } t \in [\tau - \epsilon, \tau] \\ u_n(t) & \text{otherwise.} \end{cases}$$

This short perturbation has a non-myopic effect on the resulting belief state trajectory

Optimal Perturbation for Hybrid Systems

Sequential Action Control [Ansari and Murphey, 2016]

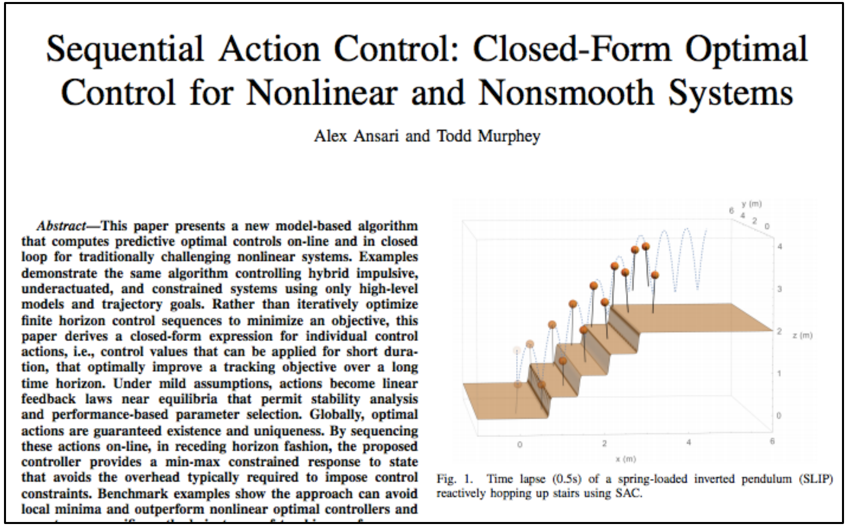

Sequential Action Control: Closed-Form Optimal Control for Nonlinear and Nonsmooth Systems

Alex Ansari and Todd Murphey

Abstract—This paper presents a new model-based algorithm that computes predictive optimal controls on-line and in closed loop for traditionally challenging nonlinear systems. Examples demonstrate the same algorithm controlling hybrid impulsive, underactuated, and constrained systems using only high-level models and trajectory goals. Rather than iteratively optimize finite horizon control sequences to minimize an objective, this paper derives a closed-form expression for individual control actions, i.e., control values that can be applied for short duration, that optimally improve a tracking objective over a long time horizon. Under mild assumptions, actions become linear feedback laws near equilibria that permit stability analysis and performance-based parameter selection. Globally, optimal actions are guaranteed existence and uniqueness. By sequencing these actions on-line, in receding horizon fashion, the proposed controller provides a min-max constrained response to state that avoids the overhead typically required to impose control constraints. Benchmark examples show the approach can avoid local minima and outperform nonlinear optimal controllers and

Fig. 1. Time lapse (0.5s) of a spring-loaded inverted pendulum (SLIP) reactively hopping up stairs using SAC.

Original SAC Algorithm:

- Physical Systems
- Deterministic State Transitions

SACBP (Ours):

- **Belief Systems**
- **Stochastic State Transitions**

Problem Formulation

Cost Functional:

$$J(b, u) = \int_{t_0}^{t_f} c\,(b(t), u(t))\, dt + h(b(t_f))$$

Belief Dynamics:

$$\begin{cases} b(t_k^+) = g(b(t_k^-), y_k) \\ \dot{b}(t) = f(b(t), u(t)) & t \in [t_k^+, t_{k+1}^-] \end{cases}$$

Perturbed Control:

$$u_w(t) = \begin{cases} w & \text{if } t \in [\tau - \epsilon, \tau] \\ u_n(t) & \text{otherwise.} \end{cases}$$

Optimization Problem:

$$\underset{w,\tau}{\text{minimize}} \quad \mathbb{E}_{(y_1,\ldots,y_k)} \left[\frac{\partial +}{\partial \epsilon} J(b, u_w) \Big|_{\epsilon=0} \right]$$

$$\text{subject to} \quad \begin{cases} b(t_k^+) = g(b(t_k^-), y_k) & k \in \{1, \ldots, K\} \\ \dot{b}(t) = f(b(t), u(t)) & t \in [t_k^+, t_{k+1}^-] \end{cases}$$

$$u_w(t) = \begin{cases} w & \text{if } t \in [\tau - \epsilon, \tau] \\ u_n(t) & \text{otherwise.} \end{cases}$$

Perturbation Theory of Differential Equations

Given $y_1, \dots, y_K$, we can compute the conditional cost value using:

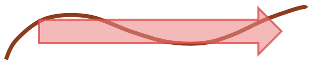

Nominal Belief Trajectory

$$\begin{cases} b_n(t_k^+) = g(b_n(t_k^-), y_k) \\ \dot{b}_n(t) = f(b_n(t), u_n(t)) \quad t \in [t_k^+, t_{k+1}^-] \end{cases}$$

Nominal Adjoint Trajectory

$$\begin{cases} \rho(t_k^-) = \frac{\partial}{\partial b} g(b_n(t_k^+), y_k)^{\mathrm{T}} \rho(t_k^+) \\ \dot{\rho}(t) = -\frac{\partial}{\partial b} c(b_n(t), u_n(t)) - \frac{\partial}{\partial b} f(b_n(t), u_n(t))^{\mathrm{T}} \rho(t) \end{cases}$$

Optimization Problem

$$\underset{w,\tau}{\text{minimize}} \quad \mathbb{E}_{(y_1,\dots,y_k)} \left[\frac{\partial +}{\partial \epsilon} J(b, u_w) \Big|_{\epsilon=0} \right]$$

$$\text{subject to} \quad \begin{cases} b(t_k^+) = g(b(t_k^-), y_k) & k \in \{1, \dots, K\} \\ \dot{b}(t) = f(b(t), u(t)) & t \in [t_k^+, t_{k+1}^-] \end{cases}$$

$$u_w(t) = \begin{cases} w & \text{if } t \in [\tau - \epsilon, \tau] \\ u_n(t) & \text{otherwise.} \end{cases}$$

Monte Carlo Sampling

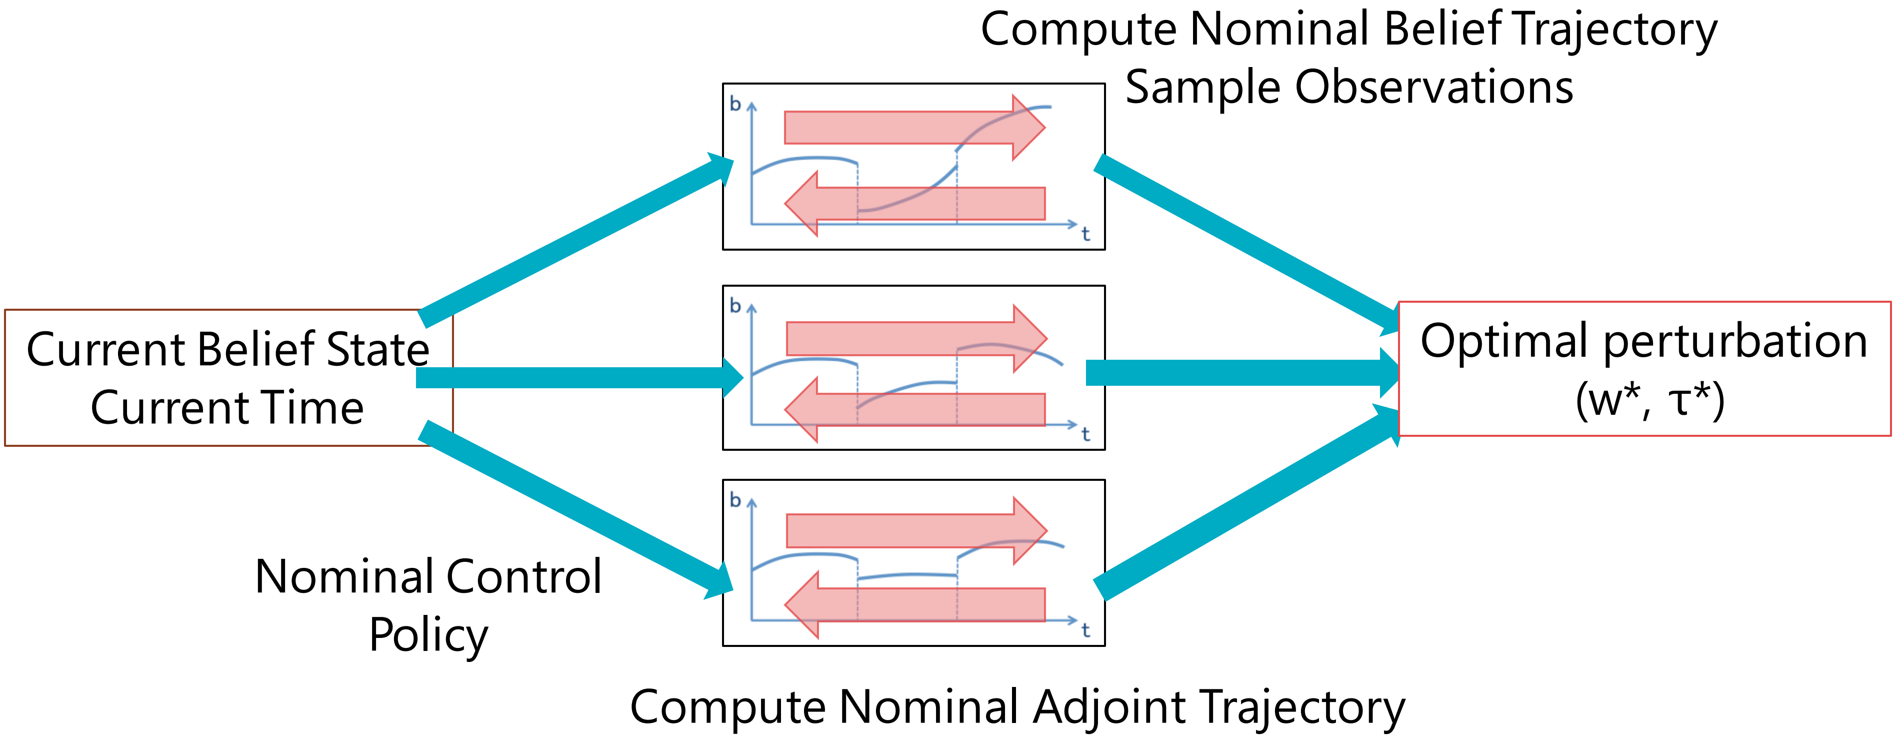

SACBP Algorithm at a Glance

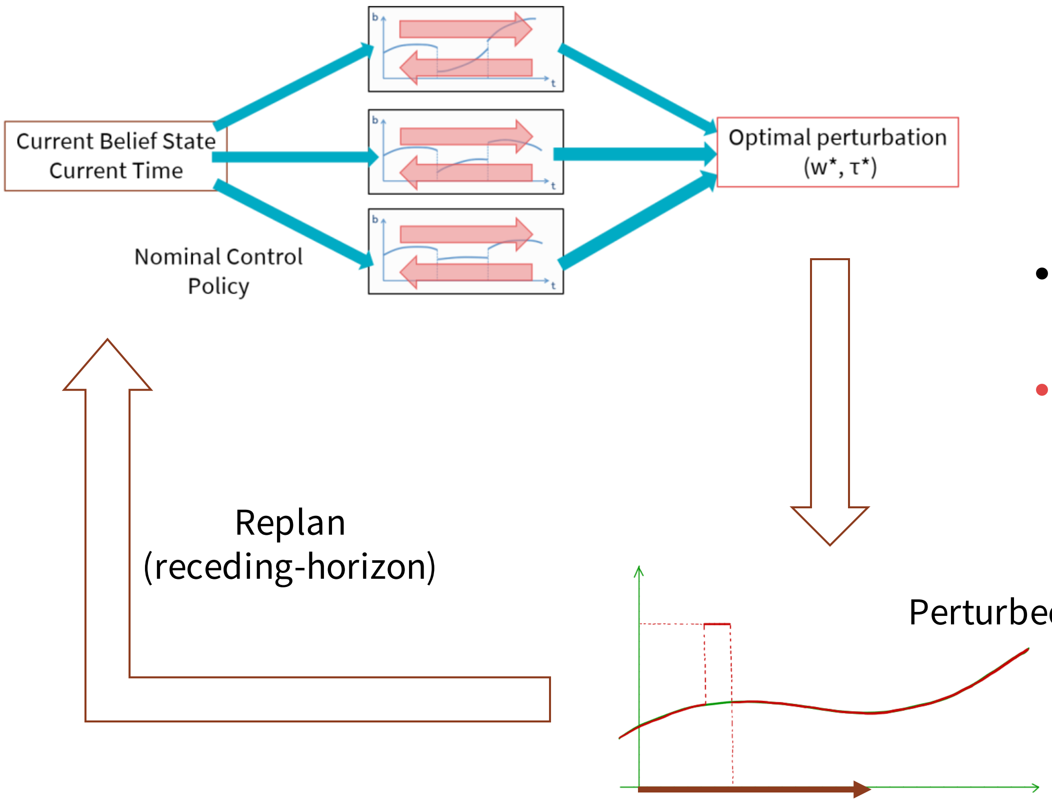

- Highly parallelizable
- Near real-time computation (with naïve implementation)

Theoretical Aspects

Property of SACBP:

In expectation
performs no worse than
the given nominal policy,
with an appropriate choice of ε.

[To be submitted to IJRR]

Optimization Problem:

$$\underset{w,\tau}{\text{minimize}} \quad \mathbb{E}_{(y_1,\dots,y_k)}\left[\frac{\partial +}{\partial \epsilon} J(b, u_w)\Big|_{\epsilon=0}\right]$$

$$\text{subject to} \quad \begin{cases} b(t_k^+) = g(b(t_k^-), y_k) & k \in \{1,\dots,K\} \\ \dot{b}(t) = f(b(t), u(t)) & t \in [t_k^+, t_{k+1}^-] \end{cases}$$

$$u_w(t) = \begin{cases} w & \text{if } t \in [\tau - \epsilon, \tau] \\ u_n(t) & \text{otherwise.} \end{cases}$$

Results: Active Multi-Target Tracking

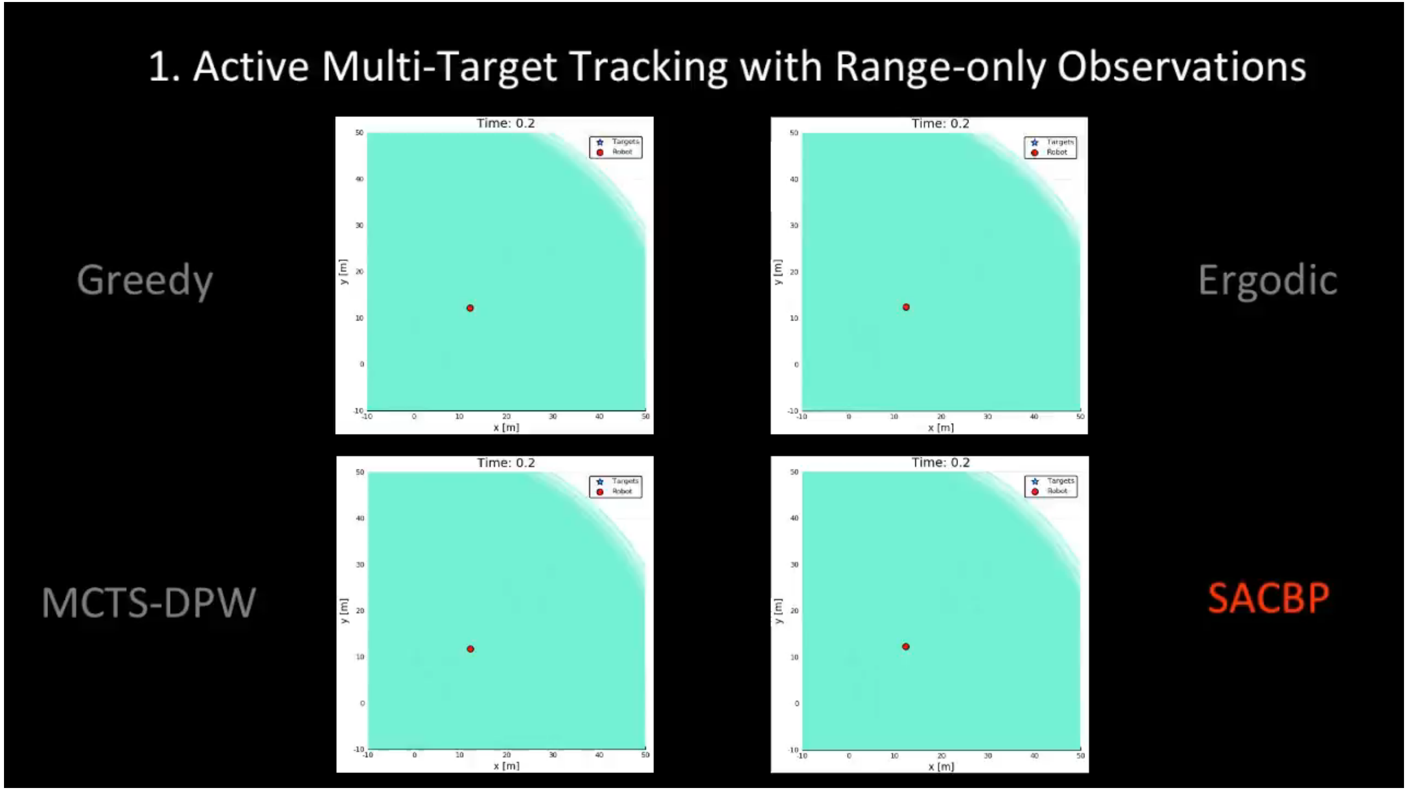

Results: Active Multi-Target Tracking

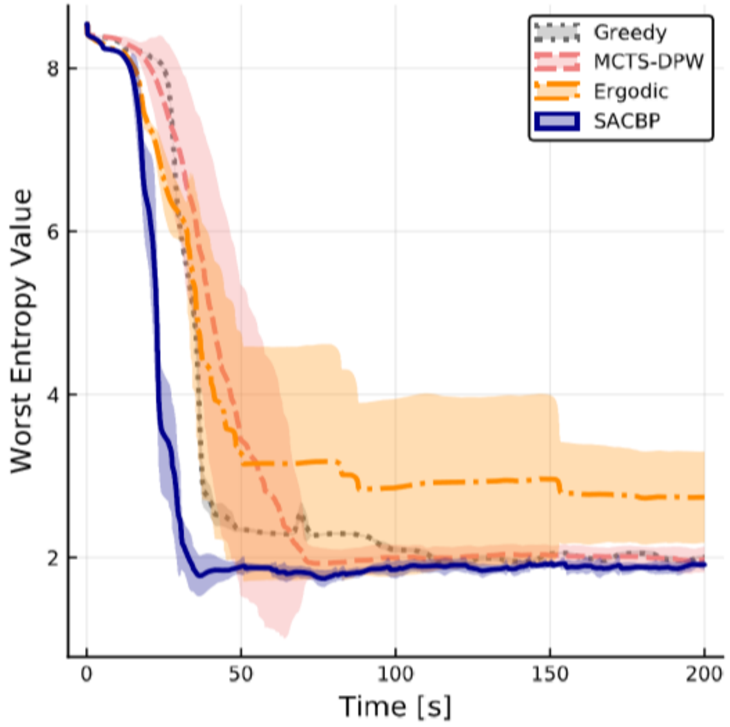

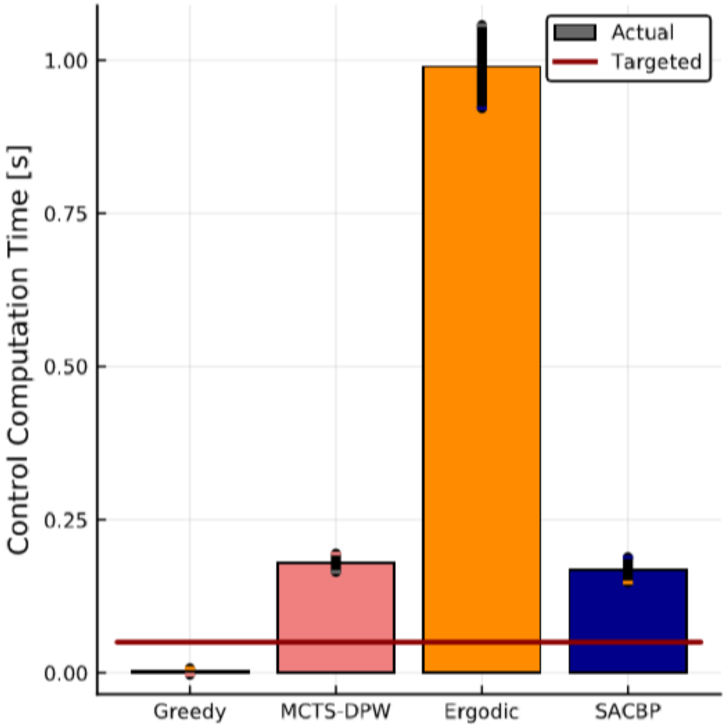

Results: Manipulation under Model Uncertainty

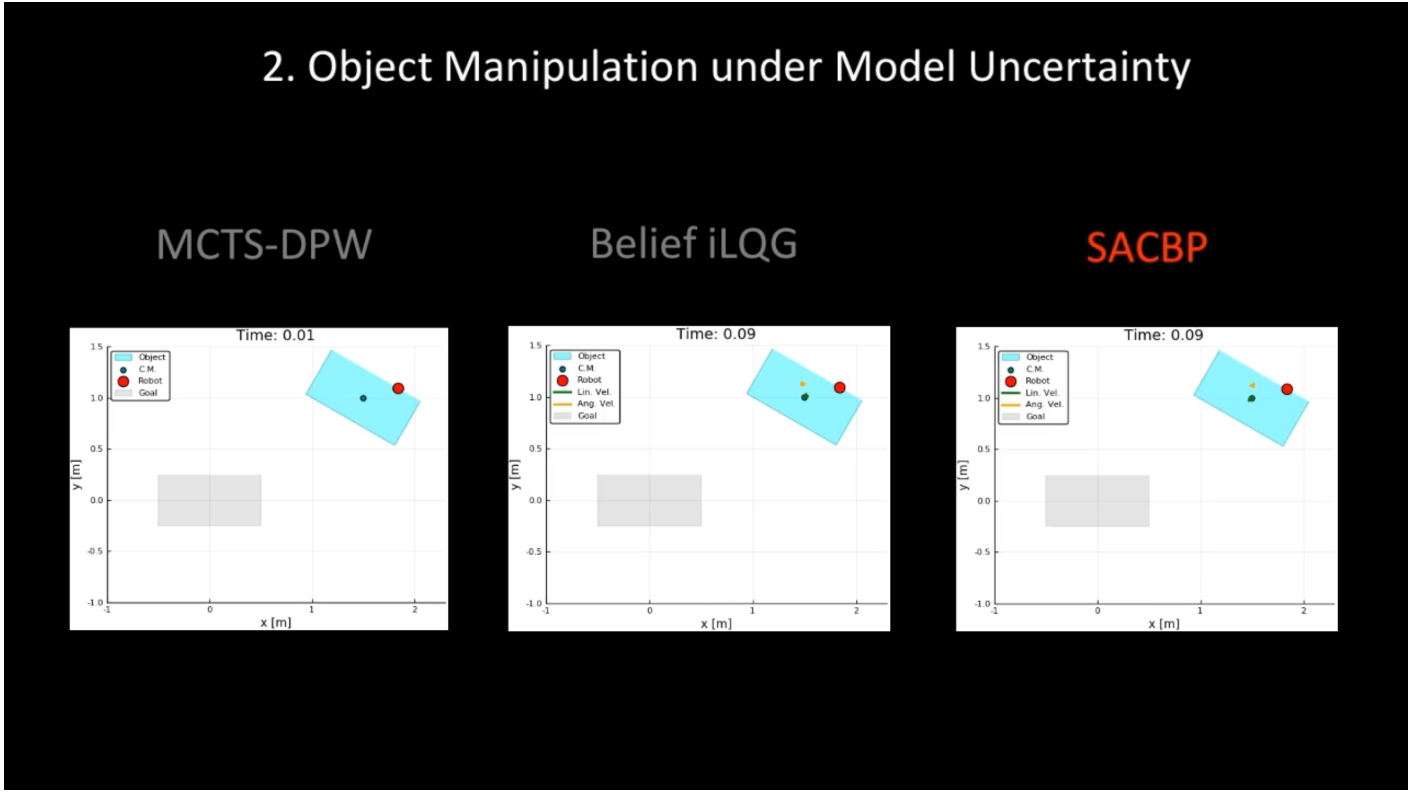

Results: Manipulation under Model Uncertainty

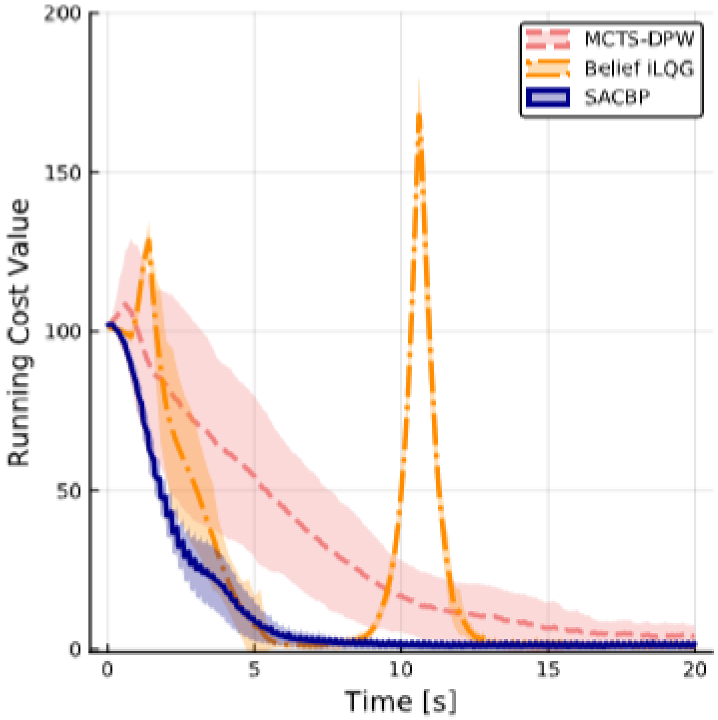

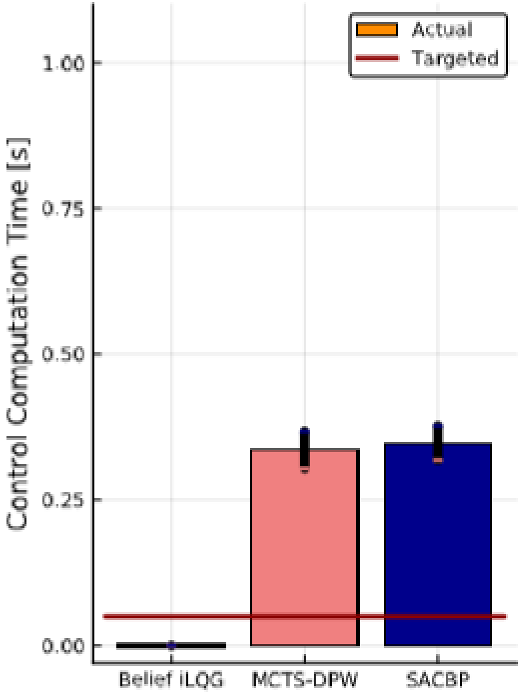

Summary

Algorithms for Information-Theoretic Active Sensing:

Rapidly-Exploring Random Cycles (RRC/RRC*)

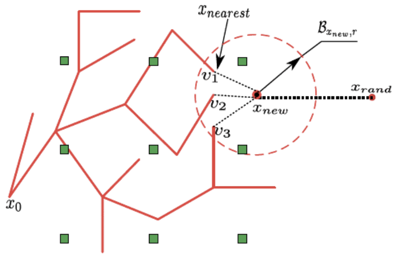

- Discrete-Time Linear Gaussian Systems
- Persistent Surveillance/Monitoring

[X. Lan and M. Schwager, IEEE T-RO 2016]

Monte Carlo Tree Search with Double-Progressive Widening (MCTS-DPW)

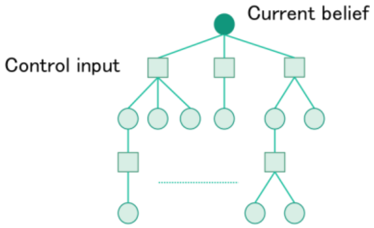

- Discrete-Time Nonlinear Non-Gaussian Systems
- Active Intent Inference

[H. Nishimura and M. Schwager, ICRA 2018]

Sequential Action Control for Belief Space Planning (SACBP)

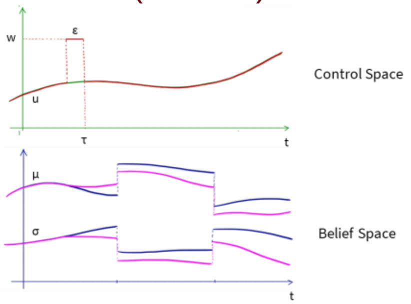

- Continuous-Time Nonlinear Non-Gaussian Systems
- Multi-Target Tracking, etc.

[H. Nishimura and M. Schwager, WAFR 2018]

References

- X. Lan and M. Schwager, “Rapidly-exploring Random Cycles: Persistent Estimation of Spatio-temoral Fields with Multiple Sensing Robots,” *IEEE Transactions on Robotics*, 2016, vol. 32, no. 5, pp. 1230-1244, **TRO King-Sun Fu Memorial Best Paper Award**.

- H. Nishimura and M. Schwager, "Active Motion-Based Communication for Robots with Monocular Vision", *2018 IEEE International Conference on Robotics and Automation (ICRA)*. Brisbane, Australia, pp. 2948-2955, 2018

- H. Nishimura and M. Schwager, "SACBP: Belief Space Planning for Continuous-Time Dynamical Systems via Stochastic Sequential Action Control", *The 13th International Workshop on the Algorithmic Foundations of Robotics (WAFR)*. Mérida, México, 2018.

Information-Theoretic Approaches to Active Sensing: Theory and Practice

Haruki Nishimura
Ph.D. Candidate in Aeronautics and Astronautics, Stanford University, USA
Contact: hnishimura@stanford.edu

Independent Administrative Institution Japan Student Services Organization

Stanford University