Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

Haruki Nishimura

Advisor: Professor Mac Schwager

Doctoral Thesis Defense

July 22nd, 2021

MSL

Stanford University

Mobile Robots in Real-World, 2010s

Amazon Warehouse Robots, 2016


Ocado Grocery Packing Robots, 2018

https://youtu.be/hevPwkH0D0k
https://youtu.be/ssZ_8cqfBlE

Mobile Robots in Real-World, 2010s

Knightscope Security Robot, 2017

Security robot 'in critical condition' after nearly drowning on the job

— CNN, July 2017

TECH

Security Robot Suspended After Colliding With a Toddler

Robot was on duty at a shopping center in California when the accident occurred

— WSJ, July 2016

https://www.cnn.com/2017/07/18/us/security-robot-drown-trnd/index.html
https://www.wsj.com/articles/security-robot-suspended-after-colliding-with-a-toddler-1468446311?st=z82ea0cmiu2kqsf&reflink=desktopwebshare_permalink

Differences in Environments

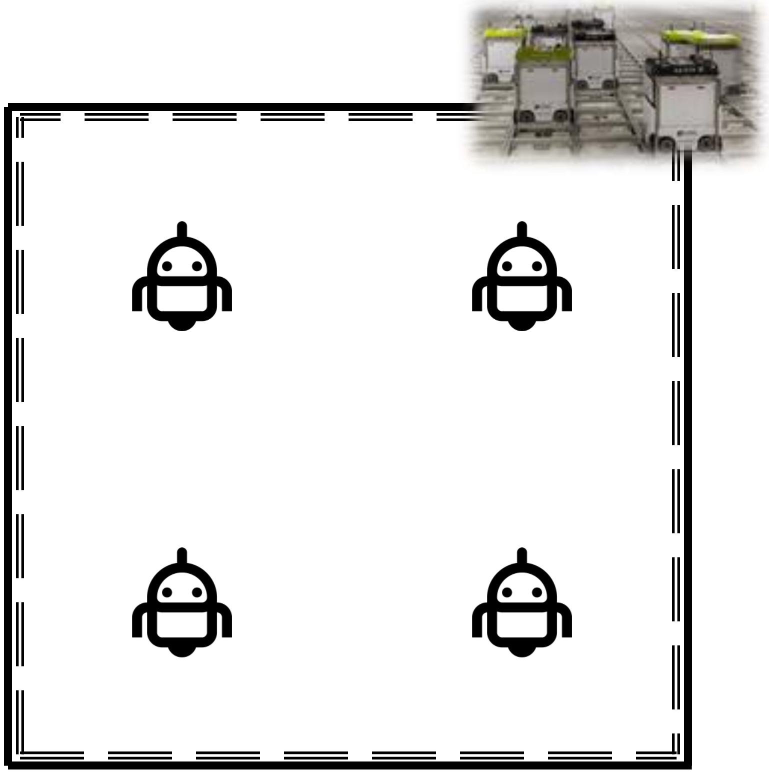

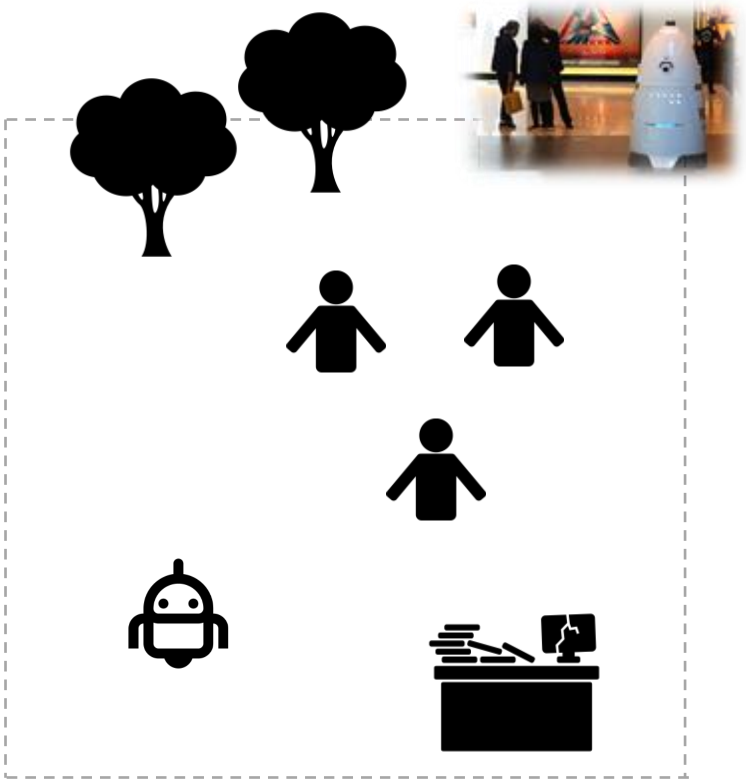

http://www.fruitnet.com/fpj/article/176666/ocado-robots-ramping-up-capacity
https://www.kqed.org/science/1943240/these-bay-area-robots-are-cool-but-they-freak-me-out-anyway

Robots in Real-World, 202X

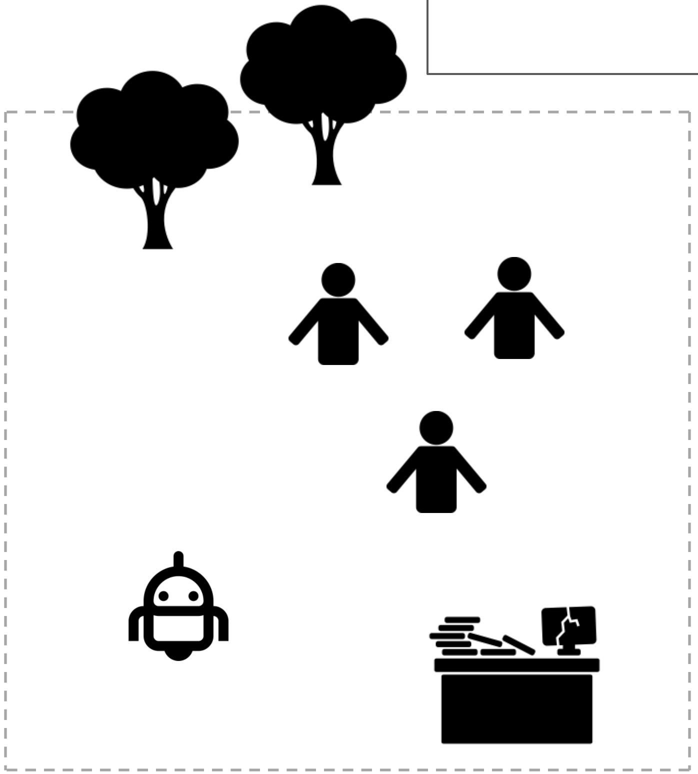

https://www.mydronelab.com/blog/drone-uses.html
https://shorturl.at/nBNO0
https://shorturl.at/uHOUX

Uncertainty in Open & Interactive Environments

1. Uncertainty about the **current state**
   - Limited perception

Robots should
actively reduce uncertainty
using onboard perception.

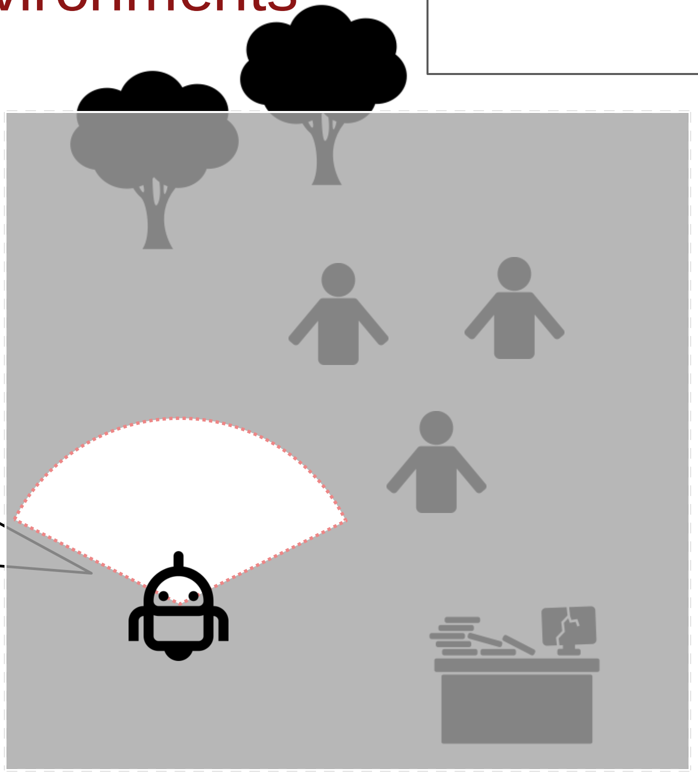

Uncertainty in Open & Interactive Environments

1. Uncertainty about the **current state**
   - Limited perception

Robots should
actively reduce uncertainty
using onboard perception.

2. Uncertainty about the **future states**
   - Inherent randomness
   - Imperfect knowledge about models

Robots should
be resilient to uncertainty
(rare but catastrophic events).

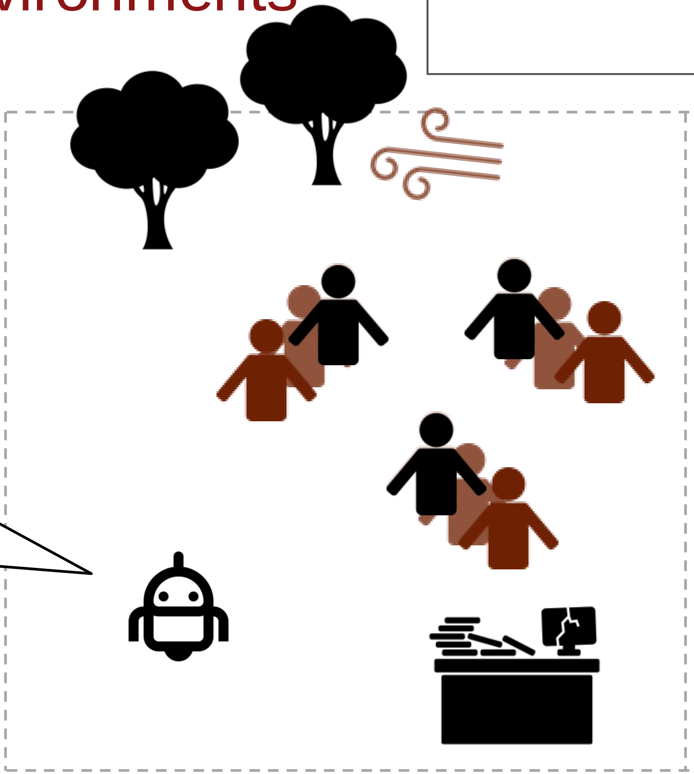

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

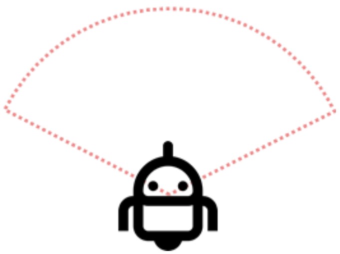

[Nishimura & Schwager, ICRA 2018]
[Nishimura & Schwager, WAFR 2018]
[Nishimura & Schwager, IJRR 2021]

Resilience to Randomness

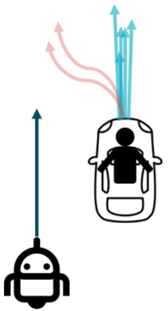

[Nishimura, Ivanovic, Gaidon, Pavone & Schwager, IROS 2020]

Resilience to Imperfect Models

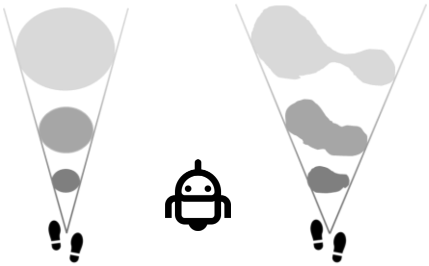

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

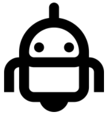

$$\begin{aligned} &\underset{\text{control policy}}{\text{minimize}} && J\left(\sum_{t=0}^{T} \text{Cost}_t, \text{Distr.}\right) \\ &\text{subject to} && \text{Stochastic Dynamics} \\ & && \text{Actuation Limits} \end{aligned}$$

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Model Predictive Control

$t = 1$

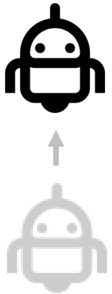

$t = 0$

$$\underset{\text{control policy}}{\text{minimize}} \quad J\left(\sum_{t=1}^{T+1} \text{Cost}_t, \text{Distr.}\right)$$

$$\text{subject to} \quad \text{Stochastic Dynamics}$$

$$\text{Actuation Limits}$$

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

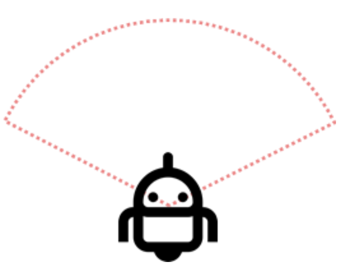

[Nishimura & Schwager, ICRA 2018]
[Nishimura & Schwager, WAFR 2018]
[Nishimura & Schwager, IJRR 2021]

Resilience to Randomness

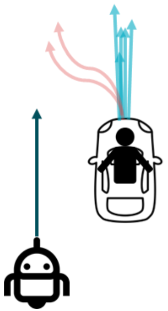

[Nishimura, Ivanovic, Gaidon, Pavone & Schwager, IROS 2020]

Resilience to Imperfect Models

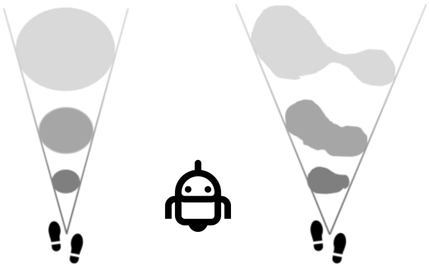

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

A Look at Human Intelligence

Humans have a natural ability to actively reduce uncertainty via perception, driven by…

1. extrinsic motivation (i.e. task-oriented)

C. Kidd and B.Y. Hayden, “The psychology and neuroscience of curiosity,” *Neuron*, 88(3), 2015, pp.449-460.
J. Gottlieb and P.Y. Oudeyer, “Towards a neuroscience of active sampling and curiosity,” *Nature Reviews Neuroscience*, 19(12), 2018, pp.758-770.

https://youtu.be/vt3E45tjJao

A Look at Human Intelligence

Humans have a natural ability to actively reduce uncertainty via perception, driven by…

1. Extrinsic motivation (i.e. task-oriented)

2. Intrinsic motivation (i.e. curiosity)

C. Kidd and B.Y. Hayden, “The psychology and neuroscience of curiosity,” *Neuron*, 88(3), 2015, pp.449-460.
J. Gottlieb and P.Y. Oudeyer, “Towards a neuroscience of active sampling and curiosity,” *Nature Reviews Neuroscience*, 19(12), 2018, pp.758-770.

 https://youtu.be/tsLD_UGYJ6Y

Computational Model – Belief Space Planning

Uncertainty reduction can be posed as a minimization problem based on probabilistic (Bayesian) inference.

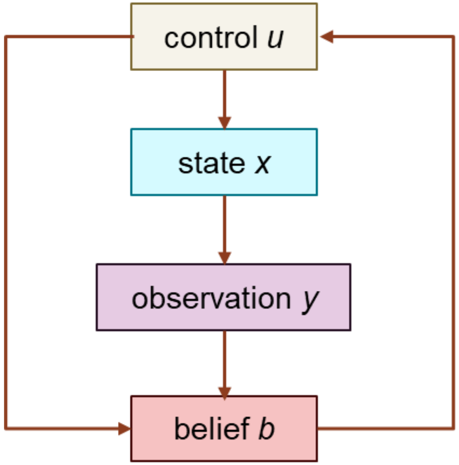

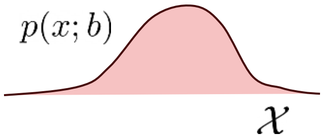

Belief: b

State Transition: $p(x' \mid x, u)$

Observation: $p(y' \mid x')$

Belief Dynamics: $\boxed{b' = g(b, u, y')}$

$$\boxed{p(x';b') = \eta p(y' \mid x') \int_{\mathcal{X}} p(x' \mid x, u) p(x;b) dx}$$

S.C.H. Yang et al., “Theoretical perspectives on active sensing,” *Current Opinion in Behavioral Sciences*, 11, 2016. pp.100-108.
J. Gottlieb and P.Y. Oudeyer, “Towards a neuroscience of active sampling and curiosity,” in *Nature Reviews Neuroscience*, 19(12), 2018, pp.758-770.

Computational Model – Belief Space Planning

Uncertainty reduction can be posed as a minimization problem based on probabilistic (Bayesian) inference.

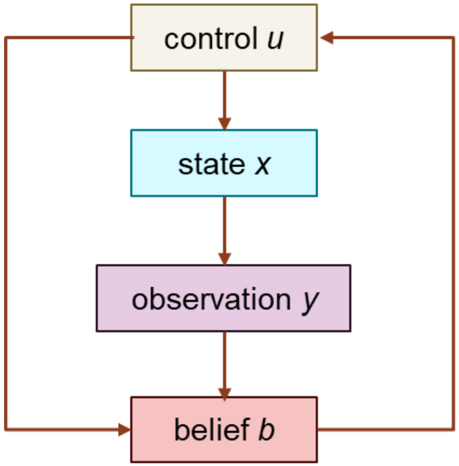

Policy: $u = \pi(b)$

Belief-based Cost: $J(u, b)$

$J(u, b) = \mathbb{E}_{b(x)}\left[c(x, u)\right]$ ⇨ Task-oriented

$J(u, b) = \mathcal{H}(b)$ ⇨ Intrinsically-motivated

(e.g. Entropy)

History of Belief Space Planning in Robotics

Partially Observable Markov Decision Processes (Task-oriented)

- K. J. Åström, "Optimal control of Markov processes with incomplete state information," *Journal of Mathematical Analysis and Applications,* 10, **1965**, pp. 174-205.

- L. P. Kaelbling et al., "Planning and acting in partially observable stochastic domains," *Artificial Intelligence,* 101(1-2), **1998**, pp. 99-134.

History of Belief Space Planning in Robotics

Partially Observable Markov Decision Processes (Task-oriented)

- K. J. Åström, "Optimal control of Markov processes with incomplete state information," *Journal of Mathematical Analysis and Applications,* 10, **1965**, pp. 174-205.
- L. P. Kaelbling et al., "Planning and acting in partially observable stochastic domains," *Artificial Intelligence,* 101(1-2), **1998**, pp. 99-134.

Active Perception (Intrinsically-motivated)

- J. M. Tenenbaum, "Accommodation in computer vision," Ph.D. Thesis, Computer Science Department Report, No. CS182, Stanford University, **1970**.
- R. Bajcsy, "Active perception," in *Proc. of the IEEE*, 76(8), **1988**, pp. 966-1005.

Example Robotic Applications

1. Task-oriented Uncertainty Reduction

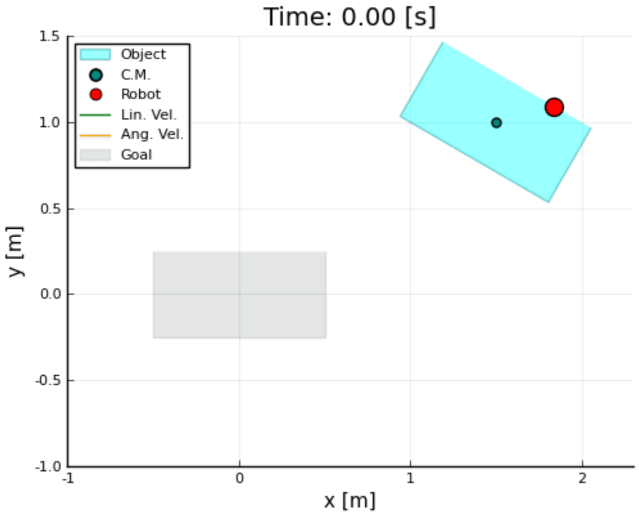

Uncertainty: Parameters of plate dynamics
Perception: Noisy pos., vel., acc. measurements

2. Intrinsically-motivated Uncertainty Reduction

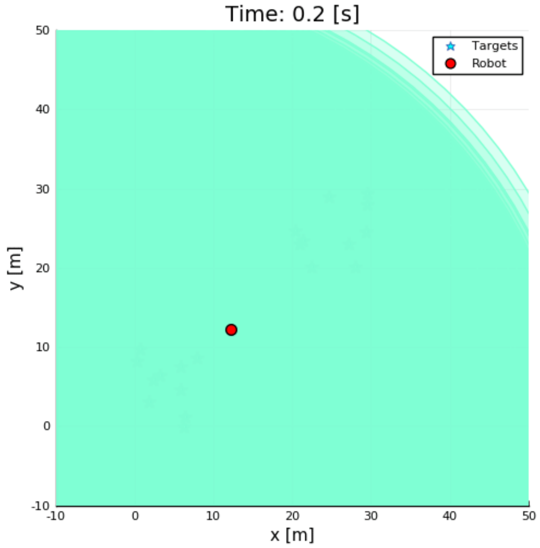

Uncertainty: Positions of moving targets
Perception: Noisy range measurements

Prior Work: Belief Space Planning

Exact optimization is **intractable** due to

1. **High-dimensionality** and **continuity** of belief space
2. **Stochasticity** in future observations

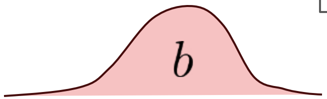

Greedy

- F. Bourgault et al., "Information based adaptive robotic exploration," in *Proc. IROS,* 2002, pp. 540-545.

Tree Search (Generic POMDP or Belief Space Planner)

- A. Couëtoux et al., "Continuous Upper Confidence Trees," in *Proc. LION,* 2011, pp. 443-445.
- A. Somani et al., "Despot: online pomdp planning with regularization," in *Proc. NeurIPS*, 2013, pp. 1772-1780.
- Z. Sunberg and M.J. Kochenderfer, "Pomcpow: an online algorithm for pomdps with continuous state, action, and observation spaces," in *Proc. ICAPS,* 2018, pp. 259-263.

Local Trajectory Optimization

- R. Platt et al., "Belief space planning assuming maximum likelihood observations," in *RSS,* 2010.
- J. van den Berg et al., "Motion planning under uncertainty using iterative local optimization in belief space," *IJRR*, 31(11), 2012, pp. 1263-1278.
- M. Rafieisakhaei et al., "T-lqg: closed-loop belief space planning via trajectory-optimized lqg," in *Proc. ICRA*, 2017, pp. 649-656.

Greedy Approach

Toy Problem: 1D Manipulation, *without Uncertainty*

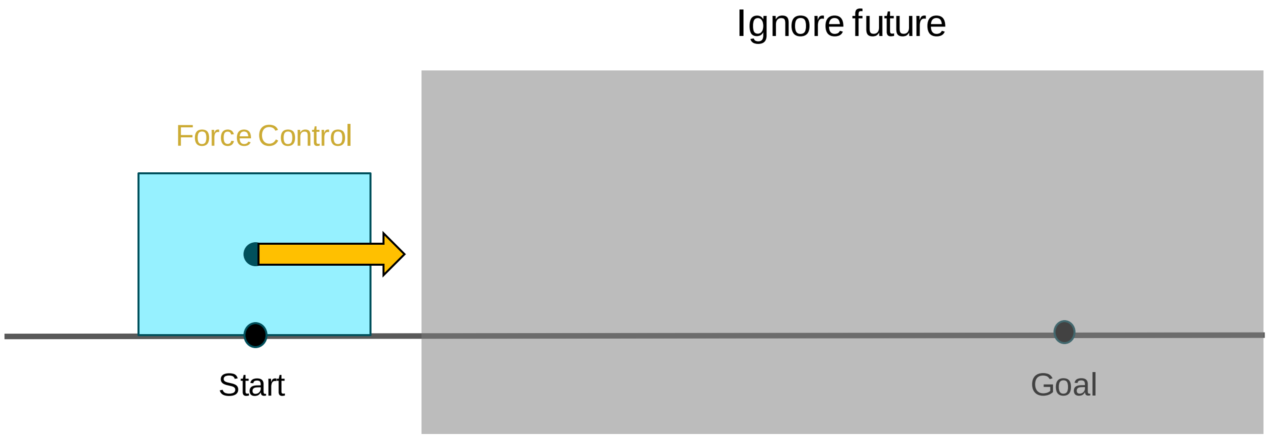

Prior Work: Belief Space Planning

Greedy

- Computationally efficient
- **Ignores long-term effect** of current action.

Tree Search (Generic POMDP or Belief Space Planner)

- A. Couëtoux et al., "Continuous Upper Confidence Trees," in *Proc. LION,* 2011, pp. 443-445.
- A. Somani et al., "Despot: online pomdp planning with regularization," in *Proc. NeurIPS*, 2013, pp. 1772-1780.
- Z. Sunberg and M.J. Kochenderfer, "Pomcpow: an online algorithm for pomdps with continuous state, action, and observation spaces," in *Proc. ICAPS,* 2018, pp. 259-263.

Local Trajectory Optimization

- R. Platt et al., "Belief space planning assuming maximum likelihood observations," in *RSS,* 2010.
- J. van den Berg et al., "Motion planning under uncertainty using iterative local optimization in belief space," *IJRR*, 31(11), 2012, pp. 1263-1278.
- M. Rafieisakhaei et al., "T-lqg: closed-loop belief space planning via trajectory-optimized lqg," in *Proc. ICRA*, 2017, pp. 649-656.

Tree Search Approach

Toy Problem: 1D Manipulation, *without Uncertainty*

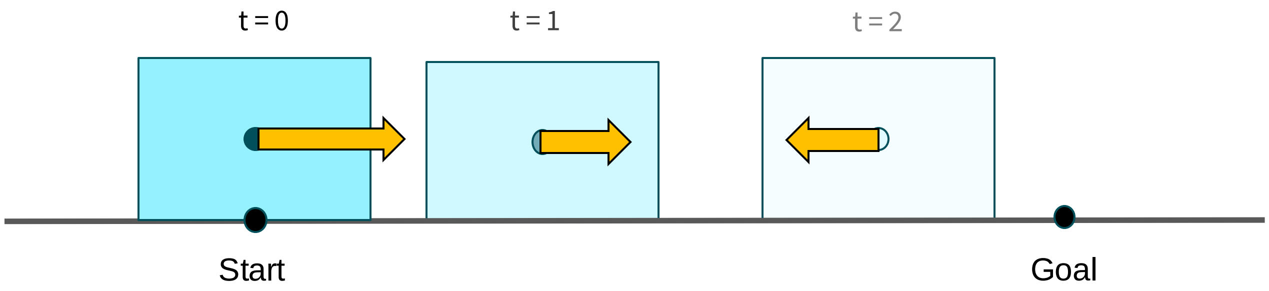

Simulate many possible scenarios, pick best one.

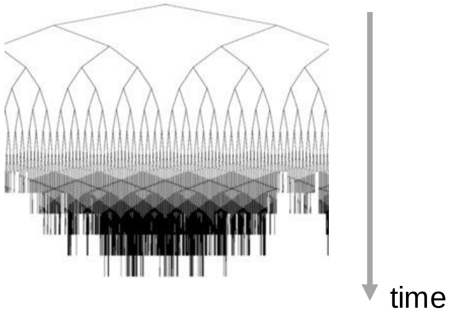

Action Space = { 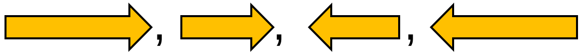 }

https://www.researchgate.net/figure/Asymmetric-tree-growth-68_fig2_235985858

Prior Work: Belief Space Planning

Greedy

- Computationally efficient
- **Ignores long-term effect** of current action.

Tree Search (Generic POMDP or Belief Space Planner)

- Asymptotic convergence to (near-) optimal policies under certain assumptions
- **Actions are noisy** with finite samples.
- **High sample complexity**

Local Trajectory Optimization

- R. Platt et al., "Belief space planning assuming maximum likelihood observations," in *RSS,* 2010.
- J. van den Berg et al., "Motion planning under uncertainty using iterative local optimization in belief space," *IJRR*, 31(11), 2012, pp. 1263-1278.
- M. Rafieisakhaei et al., "T-lqg: closed-loop belief space planning via trajectory-optimized lqg," in *Proc. ICRA*, 2017, pp. 649-656.

Local Trajectory Optimization

Toy Problem: 1D Manipulation, *without Uncertainty*

Iteratively refine space-time trajectory using physics model.

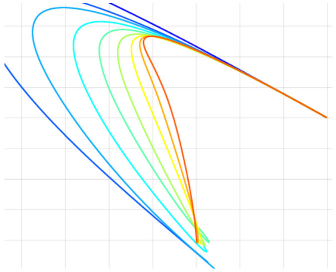

https://www.semanticscholar.org/paper/Solving-Large-scale-Quadratic-Programs-with-with-to-Miculescu-Karaman/0b32a9bbba8274db6773345a3b633cbbe4680868/figure/1

Prior Work: Belief Space Planning

Greedy

- Computationally efficient
- **Ignores long-term effect** of current action.

Tree Search (Generic POMDP or Belief Space Planner)

- Asymptotic convergence to (near-) optimal policies under certain assumptions
- **Actions are noisy** with finite samples.
- **High sample complexity**

Local Trajectory Optimization

- Convergence to locally optimal open-loop (or closed-loop) policy
- Often **ignores stochasticity** of future observations.
- **Many iterations and long computation time** to reach convergence

Our Work: Stochastic Sequential Action Control

Local Perturbation of Nominal Policy

- Computationally efficient (two-step optimization)
- Considers long-term effect & stochasticity
- Outperforms conventional methods

Greedy

- Computationally efficient
- **Ignores long-term effect** of current action.

Tree Search (Generic POMDP or Belief Space Planner)

- Asymptotic convergence to (near-) optimal policies under certain assumptions
- **Actions are noisy** with finite samples.
- **High sample complexity**

Local Trajectory Optimization

- Convergence to locally optimal open-loop (or closed-loop) policy
- Often **ignores stochasticity** of future observations.
- **Many iterations and long computation time** to reach convergence

Local Perturbation of Nominal Policy

Toy Problem: 1D Manipulation, *without Uncertainty*

Start with an imperfect nominal policy, which would fail if not modified.

Local Perturbation of Nominal Policy

Toy Problem: 1D Manipulation, *without Uncertainty*

Start with an imperfect nominal policy.

Local Perturbation of Nominal Policy

Toy Problem: 1D Manipulation, *without Uncertainty*

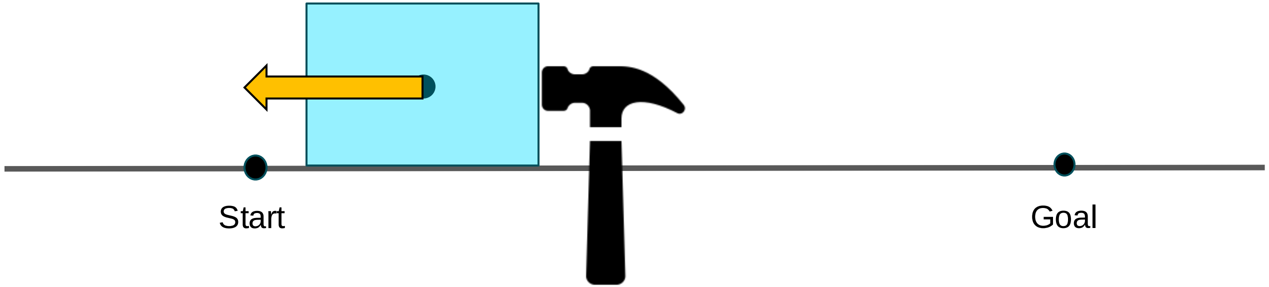

Intervene and perturb the system with an impulsive control.

Local Perturbation of Nominal Policy

Toy Problem: 1D Manipulation, *without Uncertainty*

Repeat this process fast until task done.

Impulsive Perturbation = “Mode Insertion Gradient”

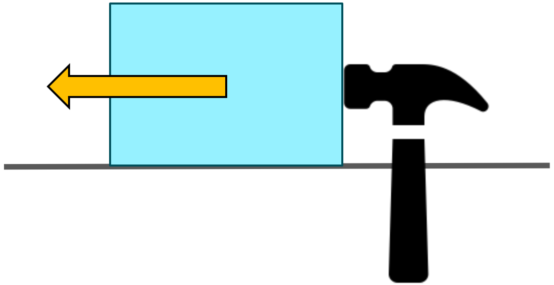

Optimal perturbation for deterministic systems is known by prior work (**S**equential **A**ction **C**ontrol).

A. Ansari and T.D. Murphey, “Sequential action control: closed-form optimal control for nonlinear and nonsmooth systems,” *T-RO*, 32(5), 2016, pp. 1196-1214.

We extend the method to stochastic hybrid systems with time-driven switching (Stochastic SAC).

Mode Insertion Gradient Optimization

Optimization can be done efficiently, in 2 steps.

1. Simulate stochastic system under (imperfect) nominal policy.

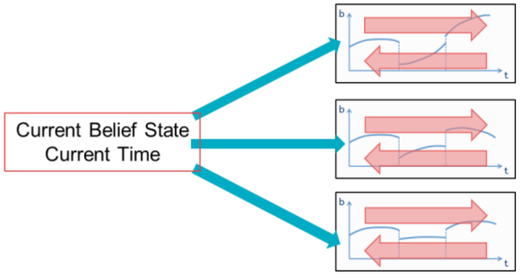

$\{y_1^{(1)}, \ldots, y_T^{(1)}\}$

$\{y_1^{(i)}, \ldots, y_T^{(i)}\}$

$\{y_1^{(N)}, \ldots, y_T^{(N)}\}$

Belief Trajectory: $b(t)$

Adjoint Trajectory: $\rho(t)$

2. Solve a quadratic minimization problem for the perturbation variable.

$$\underset{v}{\text{minimize}} \quad \frac{1}{2} v^{\mathrm{T}} R v + \mathbb{E}[\rho(\tau)]^{\mathrm{T}} H(b(\tau))(v - u(\tau))$$

MCTS-DPW (Tree Search)

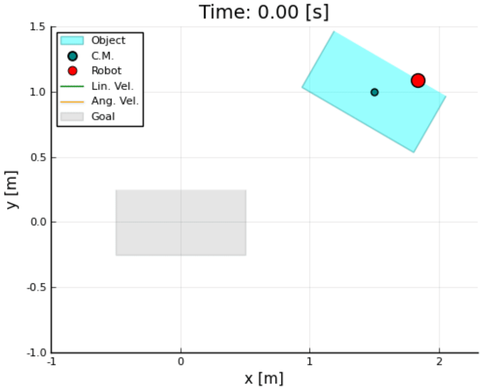

MCTS-DPW: Slade et al., IROS 2017.
T-LQG: Rafieisakhaei et al., ICRA 2017.
Belief iLQG : van den Berg et al., IJRR 2012.

Stochastic SAC (Ours)

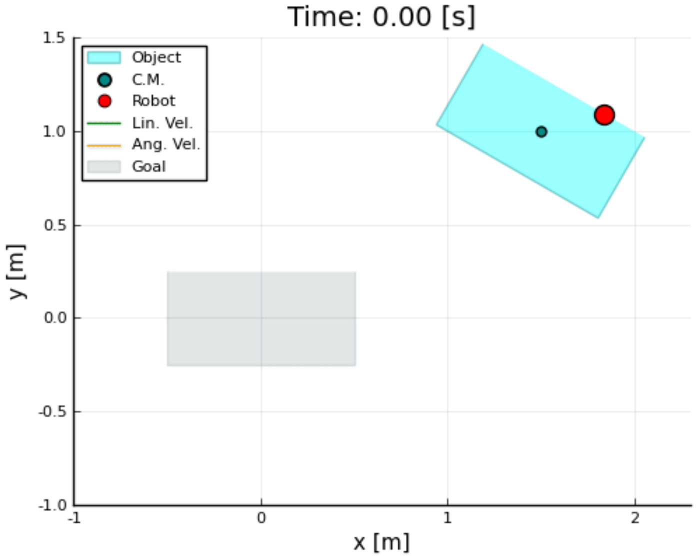

T-LQG (Traj. Opt.)

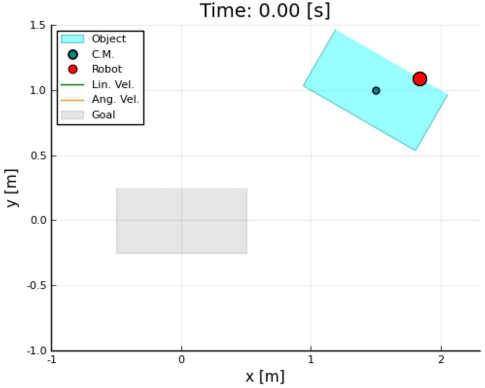

Belief iLQG (Traj. Opt.)

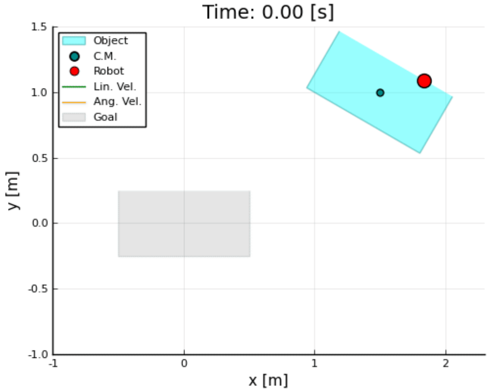

Results

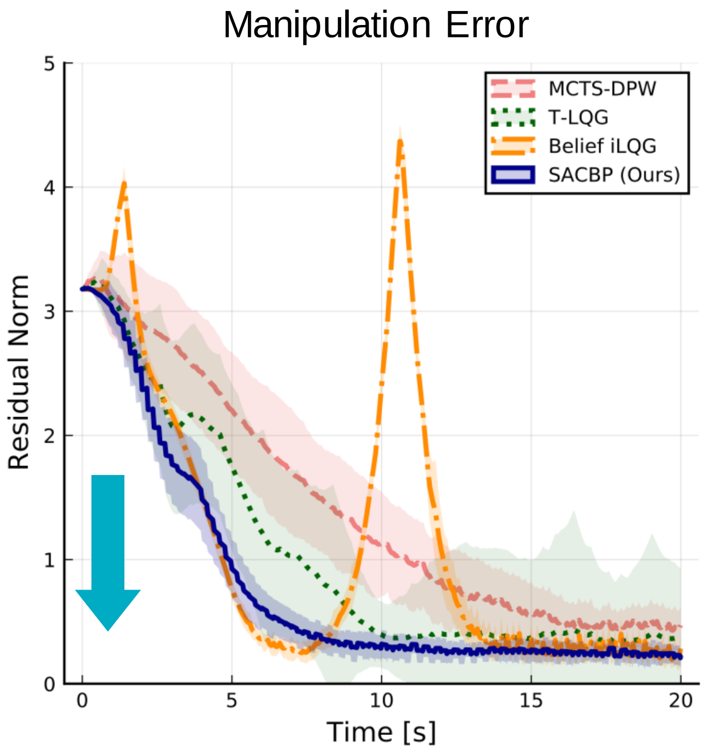

MCTS-DPW: Slade et al., IROS 2017.
T-LQG: Rafieisakhaei et al., ICRA 2017.
Belief iLQG : van den Berg et al., IJRR 2012.

Stochastic SAC achieves the best performance under a limited time budget.

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

Formulated as
Belief Space Planning

Proposed Stochastic SAC

- Computationally efficient
- Handles stochasticity
- Considers long-term effect

Resilience to Randomness

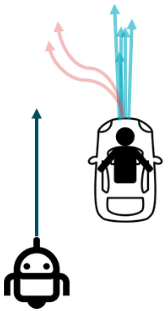

[Nishimura, Ivanovic, Gaidon, Pavone & Schwager, IROS 2020]

Resilience to Imperfect Models

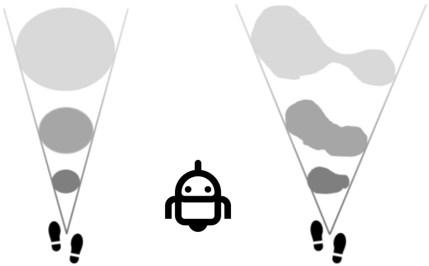

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

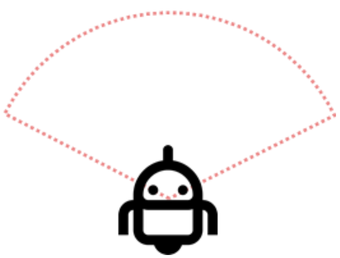

[Nishimura & Schwager, ICRA 2018]
[Nishimura & Schwager, WAFR 2018]
[Nishimura & Schwager, IJRR 2021]

Resilience to Randomness

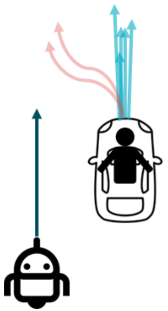

[Nishimura, Ivanovic, Gaidon, Pavone & Schwager, IROS 2020]

Resilience to Imperfect Models

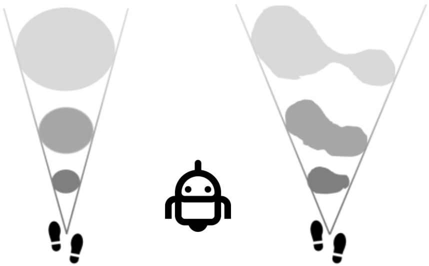

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

“Rare but Catastrophic” Events

Tesla-Subaru Accident on Highway, UT, USA

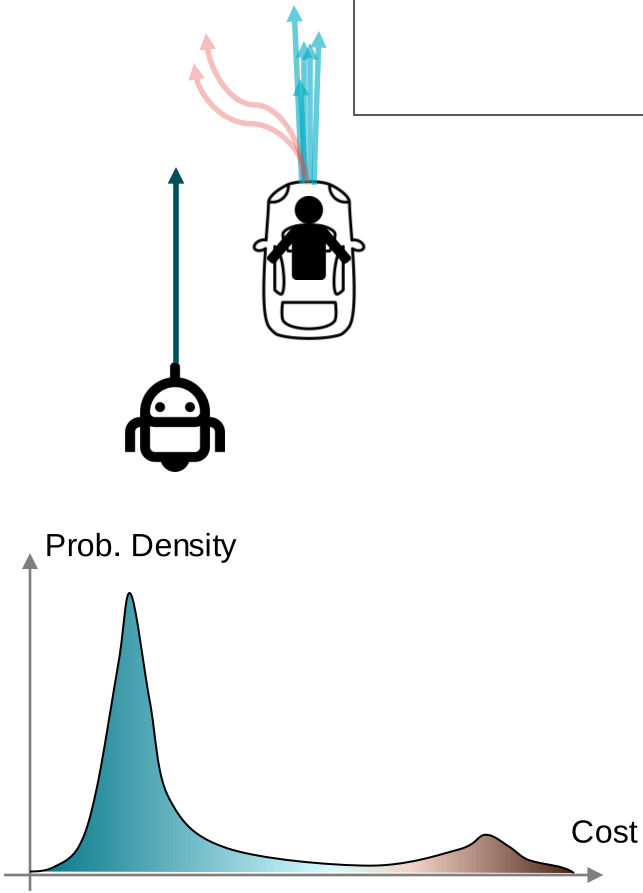

https://youtu.be/iCCmyDLN49k
https://insideevs.com/news/355036/video-tesla-model-3-crash-dashcam-police/

Stanford University

Risk-Aware Planning for Interactive Navigation

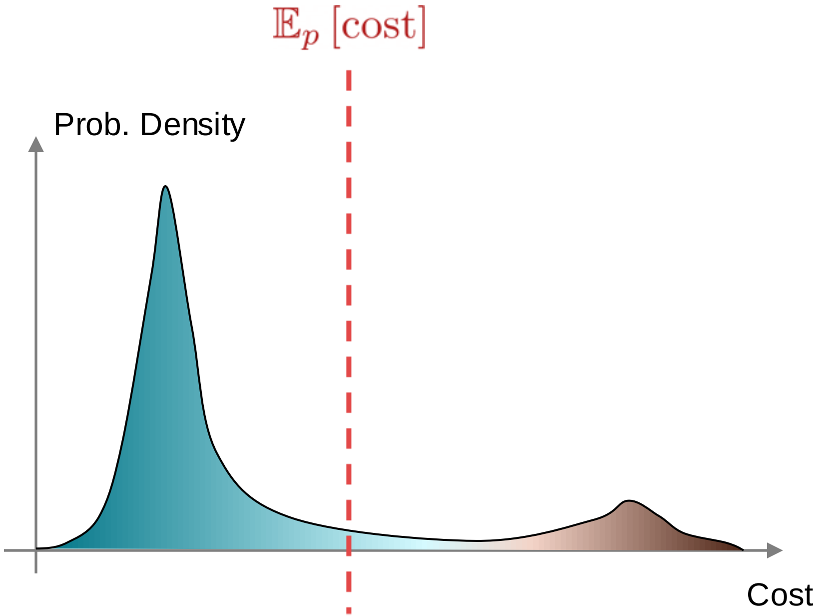

Optimizing the expected value does not suffice.

How should robots plan their future motion under **multi-modal & long-tailed distributions**?

Prior Work: Risk-Aware Planning & Robot Navigation

Deep Reinforcement Learning

- M. Everett et al., "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *Proc. IROS,* 2018, pp. 3052-3059.
- C. Chen et al., "Crowd-robot interaction: crowd-aware robot navigation with attention-based deep reinforcement learning," in *Proc. ICRA*, 2019, pp. 6015-6022.

Chance-Constrained Planning

- L. Blackmore et al., "Chance-constrained optimal path planning with obstacles," *T-RO,* 27(6), 2011, pp. 1080-1094.
- A. Wang et al., "Non-Gaussian chance-constrained trajectory planning for autonomous vehicles under agent uncertainty," *RA-L*, 5(4), 2020, pp. 6041-6048.

Conditional Value at Risk (CVaR) Optimization

- Y. Chow et al., "Risk-sensitive and robust decision making: a cvar optimization approach," in *Proc. NeurIPS,* 2015, pp. 1522-1530.
- S. Samuelson and I. Yang, "Safety-aware optimal control of stochastic systems using conditional value-at-risk," in *Proc. ACC,* 2018, pp. 6285-6290.

Risk-Sensitive Optimal Control

- P. Whittle, "Risk-sensitive linear/quadratic/gaussian control," *Advances in Applied Probability*, 13(4), 1981, pp. 764–777.
- F. Farshidian and J. Buchli, "Risk-sensitive, nonlinear optimal control: iterative linear exponential-quadratic optimal control with gaussian noise," *arXiv preprint arXiv:1512.07173*, 2015.

Prior Work: Risk-Aware Planning & Robot Navigation

Deep Reinforcement Learning

- No online computation needed
- **Not risk-aware**
- **Training with real human-robot interaction is hard**

Chance-Constrained Planning

- L. Blackmore et al., "Chance-constrained optimal path planning with obstacles," *T-RO,* 27(6), 2011, pp. 1080-1094.
- A. Wang et al., "Non-Gaussian chance-constrained trajectory planning for autonomous vehicles under agent uncertainty," *RA-L*, 5(4), 2020, pp. 6041-6048.

Conditional Value at Risk (CVaR) Optimization

- Y. Chow et al., "Risk-sensitive and robust decision making: a cvar optimization approach," in *Proc. NeurIPS,* 2015, pp. 1522-1530.
- S. Samuelson and I. Yang, "Safety-aware optimal control of stochastic systems using conditional value-at-risk," in *Proc. ACC,* 2018, pp. 6285-6290.

Risk-Sensitive Optimal Control

- P. Whittle, "Risk-sensitive linear/quadratic/gaussian control," *Advances in Applied Probability*, 13(4), 1981, pp. 764–777.
- F. Farshidian and J. Buchli, "Risk-sensitive, nonlinear optimal control: iterative linear exponential-quadratic optimal control with gaussian noise," *arXiv preprint arXiv:1512.07173*, 2015.

Prior Work: Risk-Aware Planning & Robot Navigation

Deep Reinforcement Learning

- No online computation needed
- **Not risk-aware**
- **Training with real human-robot interaction is hard**

Chance-Constrained Planning

Constraint: $\mathbb{P}(X \in \mathcal{X}_{\text{safe}}) \geq 1 - \Delta$

Conditional Value at Risk (CVaR) Optimization

Objective: $\text{CVaR}_{p,\alpha}(\text{cost})$

$$= \mathbb{E}_p \left[\text{cost} \mid \text{cost} \geq \text{VaR}_{p,\alpha}(\text{cost})\right]$$

Risk-Sensitive Optimal Control

Objective: $R_{p,\theta}(\text{cost})$

$$= \frac{1}{\theta} \log \mathbb{E}_p \left[\exp(\theta \cdot \text{cost})\right]$$

- **Limited to unimodal (e.g. Gaussian) distributions**, **linear systems**, and/or **discrete problems.**

  Exception is Wang et al. (2020), but can be overly conservative for interaction with many agents.

A. Wang et al., "Non-Gaussian chance-constrained trajectory planning for autonomous vehicles under agent uncertainty," RA-L, 5(4), 2020, pp. 6041-6048.

Our Work: Risk-Sensitive Sequential Action Control

Extension of Stochastic SAC to Risk-Sensitive Optimal Control

- Nonlinear systems
- Arbitrary distributions
- Scalable to interaction with ~50 humans

Chance-Constrained Planning

Constraint: $\mathbb{P}(X \in \mathcal{X}_{\text{safe}}) \geq 1 - \Delta$

Conditional Value at Risk (CVaR) Optimization

Objective: $\text{CVaR}_{p,\alpha}(\text{cost})$

$$= \mathbb{E}_p\left[\text{cost} \mid \text{cost} \geq \text{VaR}_{p,\alpha}(\text{cost})\right]$$

Risk-Sensitive Optimal Control

Objective: $R_{p,\theta}(\text{cost})$

$$= \frac{1}{\theta} \log \mathbb{E}_p\left[\exp(\theta \cdot \text{cost})\right]$$

- **Limited to unimodal (e.g. Gaussian) distributions**, **linear systems**, and/or **discrete problems.**

  Exception is Wang et al. (2020), but can be overly conservative for interaction with many agents.

System Model

Stochastic Hybrid System with Time-Driven Switching

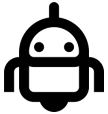

$$\dot{x}(t) = f\left(x(t)\right) + H\left(x(t)\right) u(t)$$

- Continuous-Time
- Deterministic
- Control-Affine

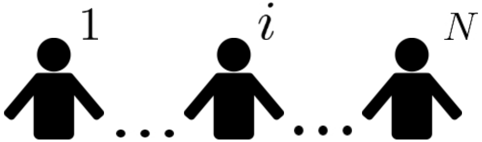

$$p^i_{k+1} = p^i_k + y^i_k \quad \{y^{1:N}_{1:T}\} \sim p$$

- Discrete-Time
- Stochastic
- Arbitrary Distribution

Entropic Risk Objective

$$R_{p,\theta}(J) = \frac{1}{\theta}\log \mathbb{E}_p\left[\exp(\theta J)\right]$$

$$\approx \underbrace{\mathbb{E}_p[J]}_{\text{Mean}} + \frac{\theta}{2}\underbrace{\mathrm{Var}_p(J)}_{\text{Variance}}$$

J : {Collision, Tracking, Control} Cost

θ : Risk-Sensitivity Parameter

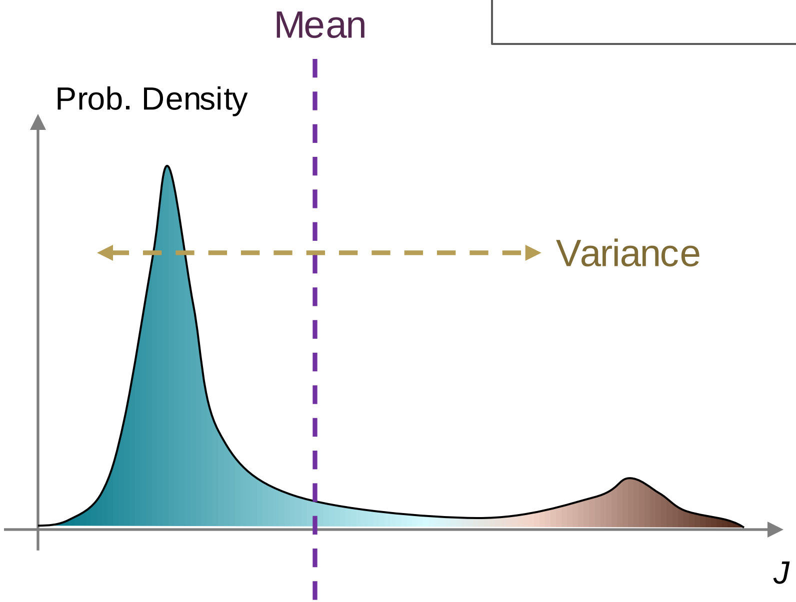

Risk-Neutral ⟷ Risk-Sensitive

$\theta = 0$ ⟷ $\theta > 0$

Mode Insertion Gradient for Entropic Risk

Expected Mode Insertion Gradient (Stochastic SAC):

$$\frac{1}{2}v^{\mathrm{T}}Rv + \boxed{\mathbb{E}[\rho(\tau)]^{\mathrm{T}}}H(x(\tau))(v - u(\tau))$$

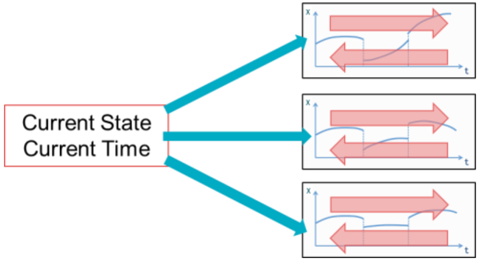

Risk-Sensitive Mode Insertion Gradient (this work):

$$\frac{1}{2}v^{\mathrm{T}}Rv + \boxed{\frac{\mathbb{E}[\exp(\theta J)\rho(\tau)]^{\mathrm{T}}}{\mathbb{E}[\exp(\theta J)]}}H(x(\tau))(v - u(\tau))$$

State Trajectory: $x(t)$

Adjoint Trajectory: $\rho(t)$

Generative Behavior Prediction: Trajectron++

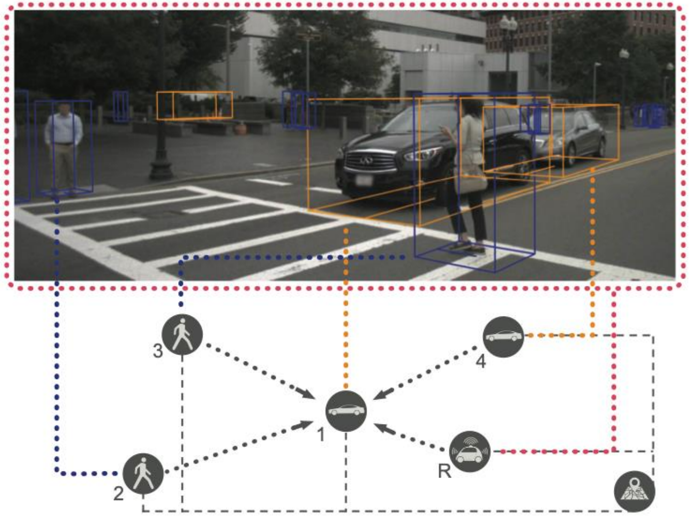

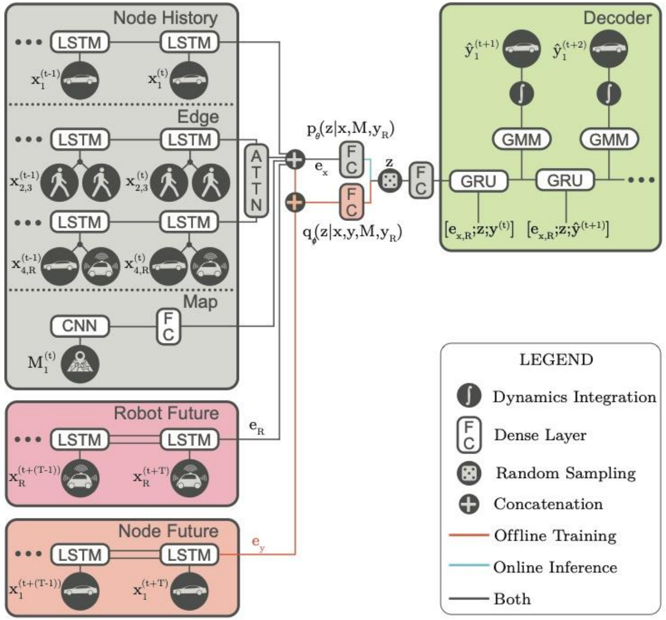

T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: dynamically-feasible trajectory forecasting with heterogeneous data," in *ECCV,* 2020.

Robot-Future-Conditional Prediction

Robot Going Up

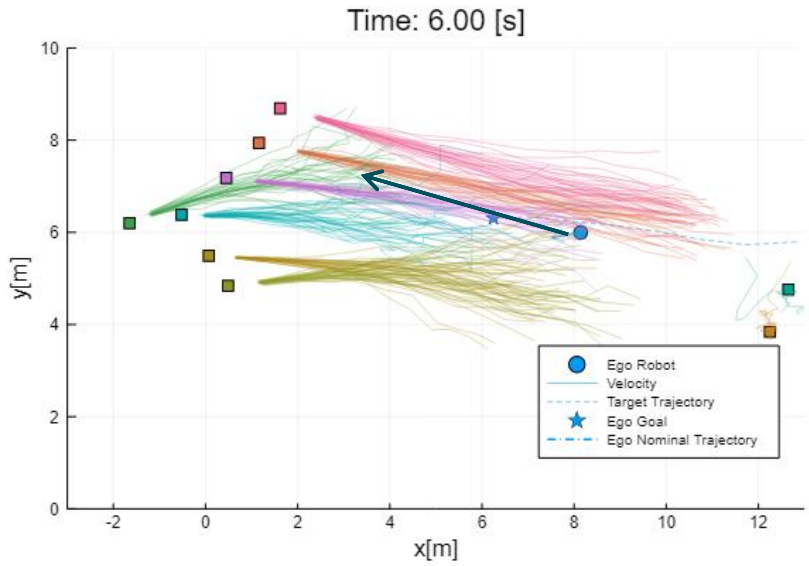

Robot Going Down

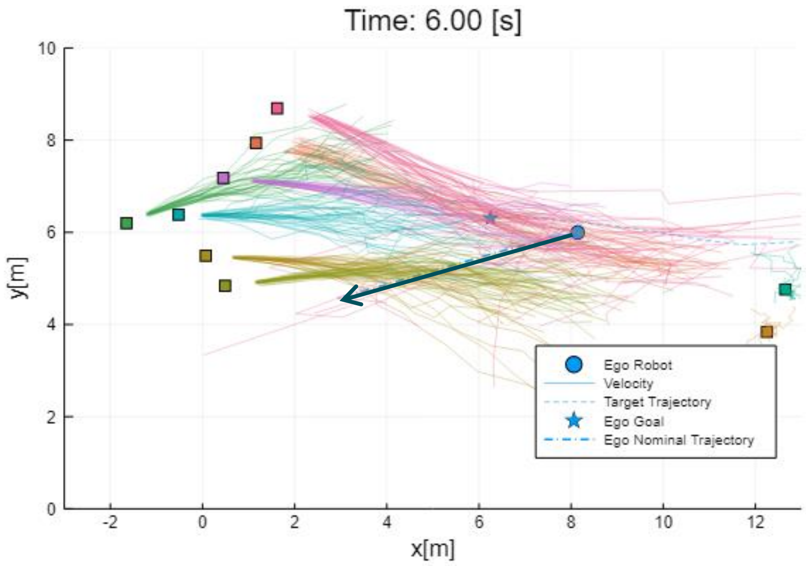

Simulation Benchmark for Risk-Neutral Robot

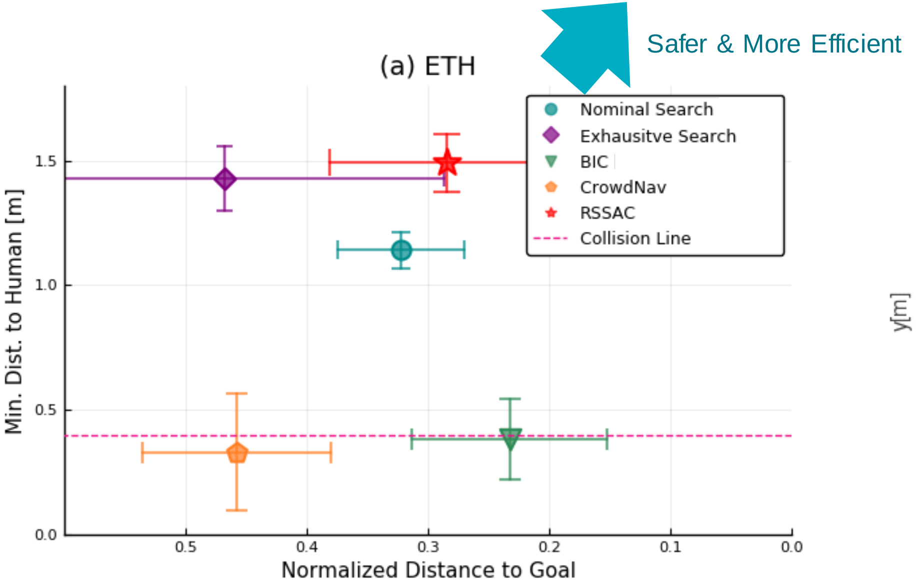

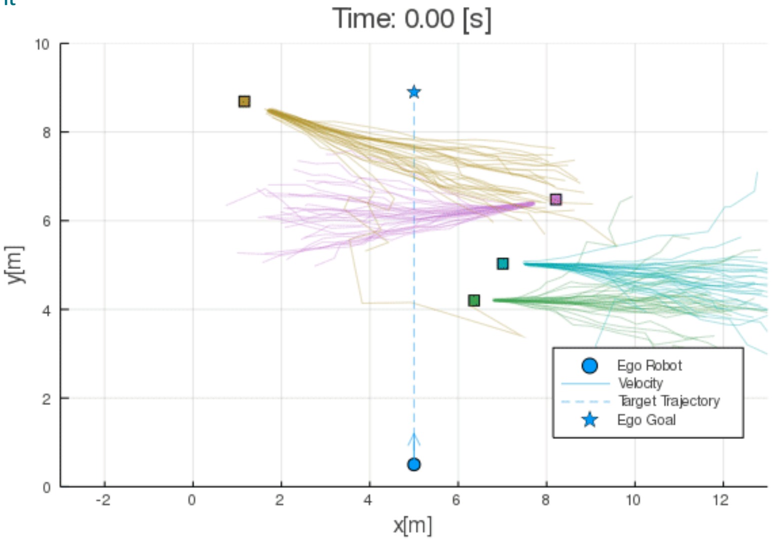

Exhaustive Search: E. Schmerling et al., “Multimodal probabilistic model-based planning for human-robot interaction,” in *Proc. ICRA*, 2018, pp. 3399-3406.

BIC: Wang et al., “Safe distributed lane change maneuvers for multiple autonomous vehicles using buffered input cells.” in *Proc. ICRA,* 2018, pp. 4678-4684.

CrowdNav: C. Chen et al., “Crowd-robot interaction: crowd-aware robot navigation with attention-based deep reinforcement learning,” in *Proc. ICRA*, 2019, pp. 6015-6022.

Risk-Sensitivity and Navigation Behavior

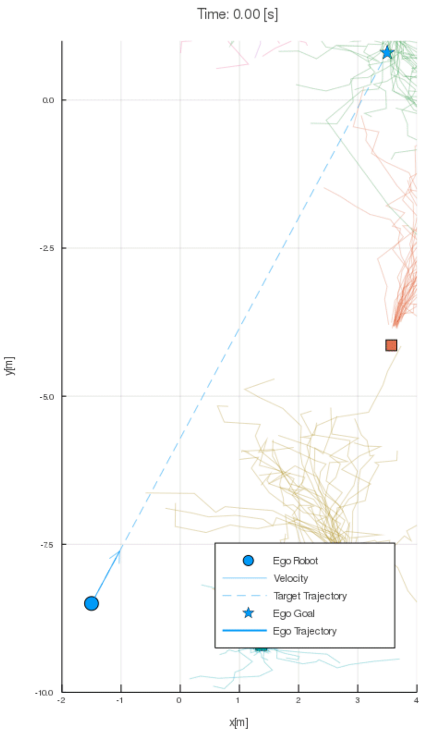

θ = 0.0 (Risk-Neutral)

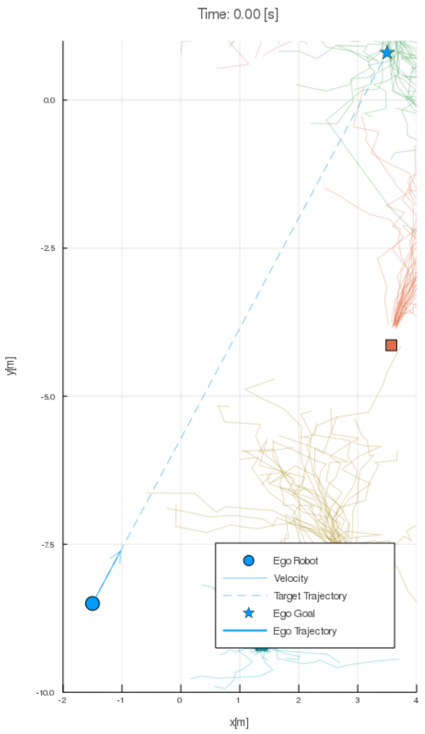

θ = 1.0 (Risk-Sensitive)

Yielding behavior naturally emerges from Risk-Sensitivity.

Risk-Sensitivity and Safety

θ = 0.0 (Risk-Neutral)

θ = 1.0 (Risk-Sensitive)

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

Formulated as
Belief Space Planning

Proposed Stochastic SAC

- Computationally efficient
- Handles stochasticity
- Considers long-term effect

Resilience to Randomness

Formulated as
Risk-Sensitive Optimal Control

Proposed Risk-Sensitive SAC

- Nonlinear systems
- Arbitrary distributions
- Scalable to interaction with ~50 humans

Resilience to Imperfect Models

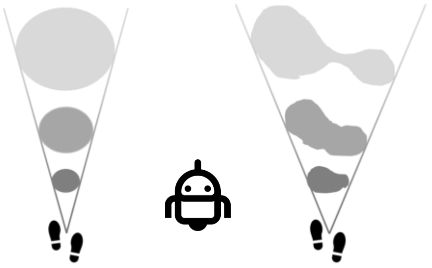

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

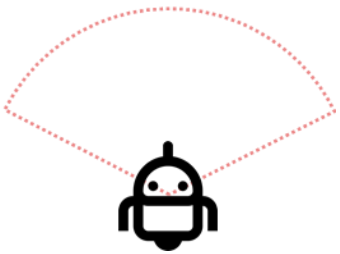

[Nishimura & Schwager, ICRA 2018]
[Nishimura & Schwager, WAFR 2018]
[Nishimura & Schwager, IJRR 2021]

Resilience to Randomness

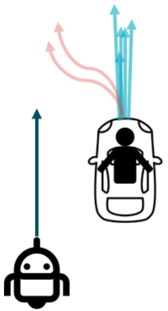

[Nishimura, Ivanovic, Gaidon, Pavone & Schwager, IROS 2020]

Resilience to Imperfect Models

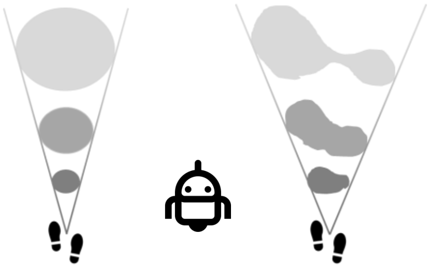

[Nishimura, Mehr, Gaidon & Schwager, RA-L 2021]

Distributional Model Mismatch

Model Distribution: $q(w)$

True Distribution: $p(w)$

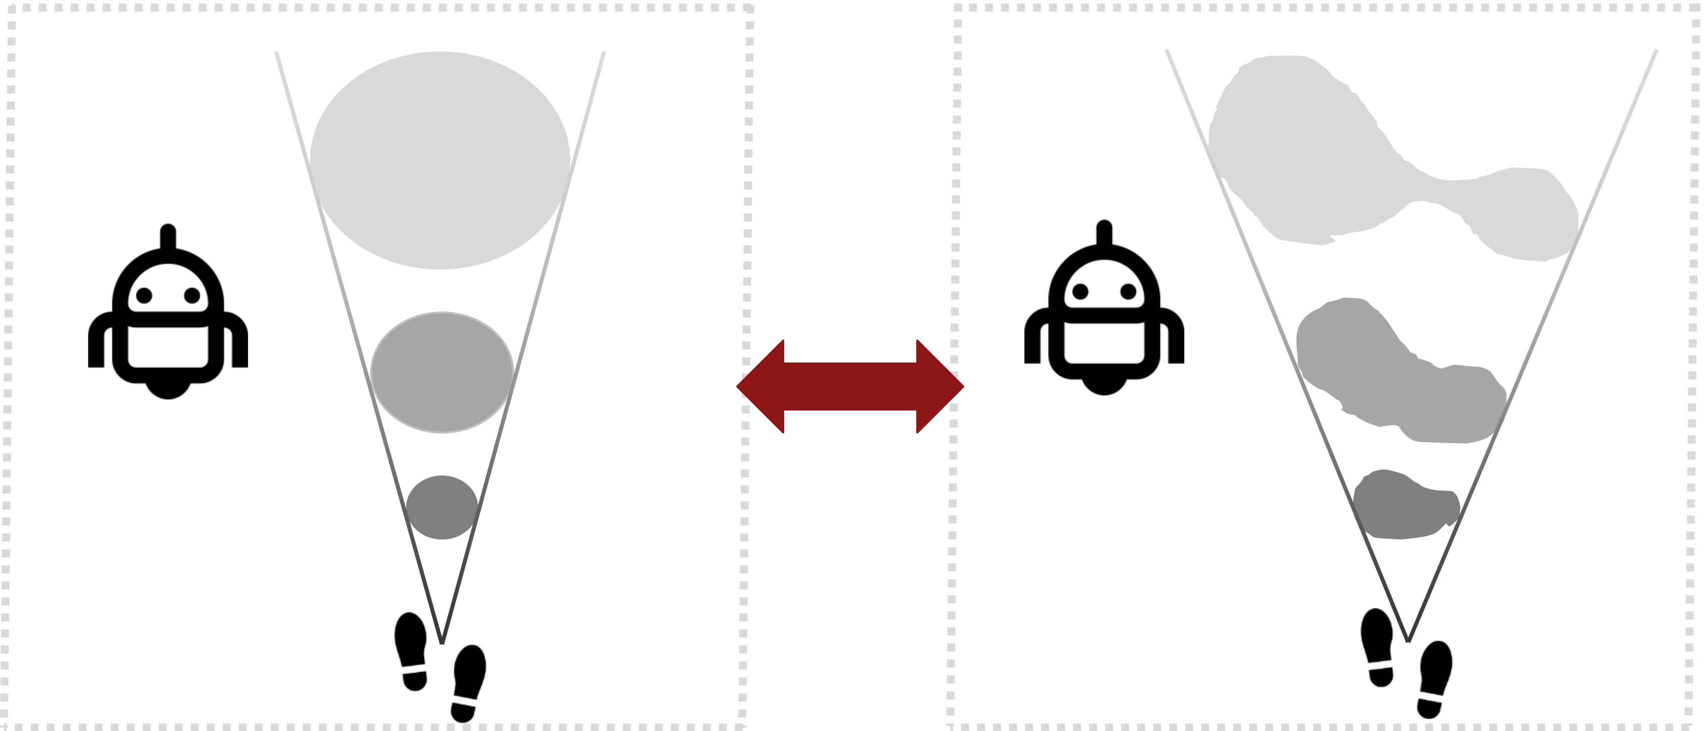

Distributional Model Mismatch

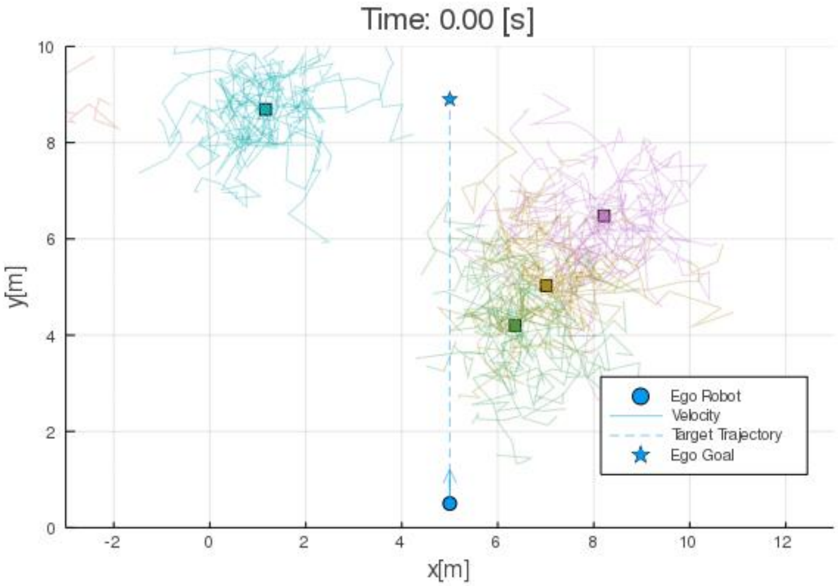

Imperfect models alone can lead to disastrous failure.

R. Cheng et al., "Limits of probabilistic safety guarantees when considering human uncertainty," in *Proc. ICRA*, 2021.

"No model is perfect, but some are useful."

— every roboticist

Set of Possible Models – Ambiguity Set

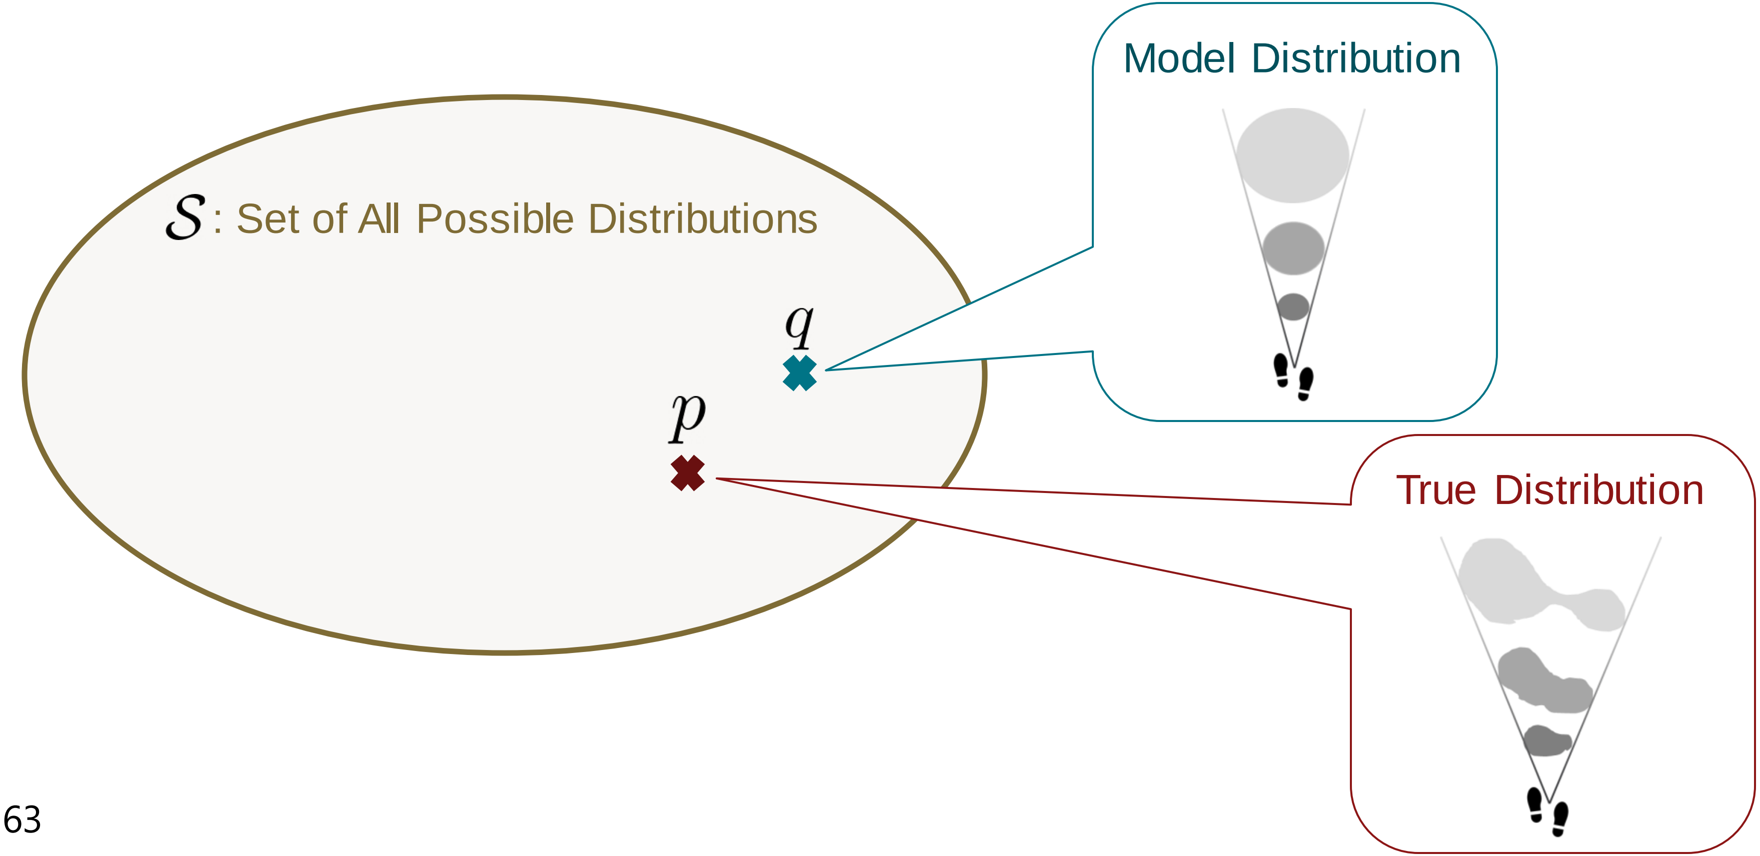

Set of Possible Models – Ambiguity Set

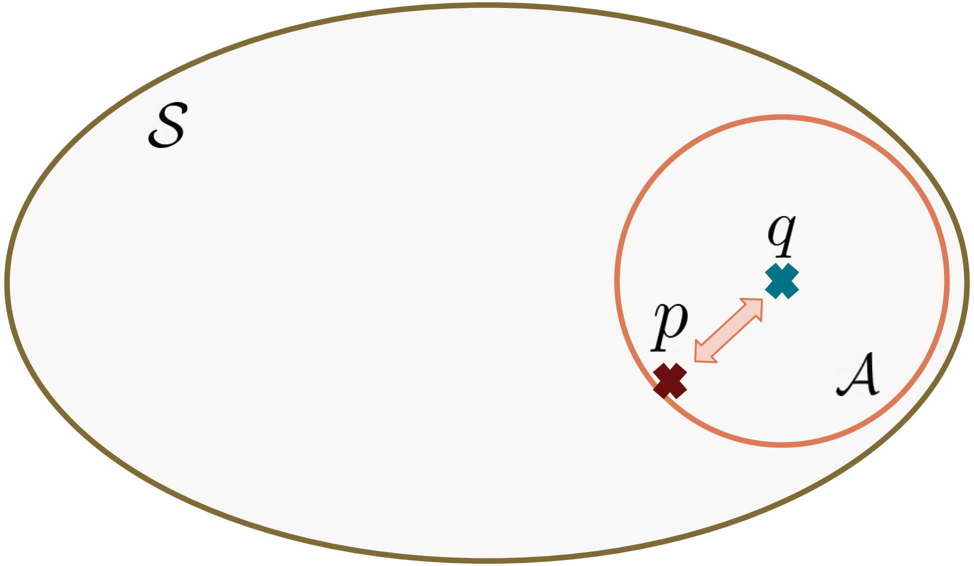

Prior Work: Distributionally Robust Control

Key Idea: Planning against a worst-case distribution out of the ambiguity set.

Moment-based Ambiguity Set

- B.P.G. Van Parys et al., "Distributionally robust control of constrained stochastic systems," *TAC*, 61(2), 2016, pp. 430-442.
- S. Samuelson and I. Yang, "Data-driven distributionally robust control of energy storage to manage wind power fluctuations, in *Proc. CCTA,* 2017, pp. 199-204.

Wasserstein Metric-based Ambiguity Set

- A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion planning and control with safety constraints using conditional value-at-risk," in *Proc. ICRA*, 2020, pp. 490-496.

f-divergence-based Ambiguity Set

- I. R. Petersen et al., "Minimax optimal control of stochastic uncertain systems with relative entropy constraints," *TAC*, 45(3), 2000, pp. 398-412.
- A. Sinha et al., "Formulazero: distributionally robust online adaptation via offline population synthesis," in *Proc. ICML,* 2020, pp. 8992-9004.

Prior Work: Distributionally Robust Control

Planning is against a worst-case distribution out of the ambiguity set.

Moment-based Ambiguity Set

- Only need moments such as mean and variance.
- Often **overly conservative**

Wasserstein Metric-based Ambiguity Set

- A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion planning and control with safety constraints using conditional value-at-risk," in *Proc. ICRA*, 2020, pp. 490-496.

f-divergence-based Ambiguity Set

- I. R. Petersen et al., "Minimax optimal control of stochastic uncertain systems with relative entropy constraints," *TAC*, 45(3), 2000, pp. 398-412.
- A. Sinha et al., "Formulazero: distributionally robust online adaptation via offline population synthesis," in *Proc. ICML*, 2020, pp. 8992-9004.

Prior Work: Distributionally Robust Control

Planning is against a worst-case distribution out of the ambiguity set.

Moment-based Ambiguity Set

- Only need moments such as mean and variance.
- Often **overly conservative**

Wasserstein Metric-based Ambiguity Set

f-divergence-based Ambiguity Set

- Existing solution methods are **not for nonlinear systems with continuous distributions**.

Our Work: Risk Auto-Tuning Iterative LQR

KL-divergence-based Ambiguity Set & Risk-Sensitive Optimal Control

- Based on theory developed by Petersen et al. (2000).
- Nonlinear Systems
- Continuous Distributions
- Locally-optimal feedback policy

I. R. Petersen et al., "Minimax optimal control of stochastic uncertain systems with relative entropy constraints," *TAC*, 45(3), 2000, pp. 398-412.

Moment-based Ambiguity Set

- Only need moments such as mean and variance.
- Often **overly conservative**

Wasserstein Metric-based Ambiguity Set

f-divergence-based Ambiguity Set

- Existing solution methods are **not for nonlinear systems with continuous distributions**.

Distributional Robustness and Risk-Sensitivity

$$\min_{\mathcal{K}\in\Lambda} \boxed{\max_{p\in\mathcal{A}}} \mathbb{E}_p[J]$$

$$\mathcal{A} = \{p \in \mathcal{S} : \mathbb{D}_{\mathrm{KL}}(p\|q) \leq d\}$$

p : True distribution (unknown)

$\mathcal{A}$: Ambiguity set (known)

Λ : Feedback Policy Class (known)

Petersen et al. (2000)
Lagrange Duality & Variational Representation of KL-Divergence

$$\underline{\min_{\theta\in\Gamma}} \boxed{\left(\min_{\mathcal{K}\in\Lambda} R_{q,\theta}(J)\right)} + \frac{d}{\theta}$$

$R_{q,\theta}(\cdot)$: Entropic Risk Objective

θ : Risk-Sensitivity Parameter

Distributional robustness yields Risk-aware Planning with optimal risk-sensitivity.

I. R. Petersen et al., “Minimax optimal control of stochastic uncertain systems with relative entropy constraints,” *TAC*, 45(3), 2000, pp. 398-412.

Bilevel Optimization for Locally-Optimal Policy

$$\min_{\theta \in \Gamma} \left(\min_{K \in \Lambda} R_{q,\theta}(J) \right) + \frac{d}{\theta}$$

q : Gaussian distribution (known)

Intractable to achieve global optimality for nonlinear systems!

Bilevel Optimization for Locally-Optimal Policy

$$\min_{\theta\in\Gamma} \left(\min_{\mathcal{K}\in\Lambda} R_{q,\theta}(J) \right) + \frac{d}{\theta}$$

Risk Auto-Tuning Iterative LQR (RAT iLQR)

Inner-Loop Problem (Risk-Sensitive Optimal Control)

$$\min_{\mathcal{K}\in\Lambda} R_{q,\theta}(J)$$

$$x_{k+1} = f(x_k, u_k) + g(x_k, u_k)w_k \quad w_k \sim q(w)$$

$$u_k = \mathcal{K}(k, x_k)$$

iterative LEQG Algorithm

$$\mathcal{K}^*(k, x) = L_k(x - \bar{x}_k) + l_k$$

Outer-Loop Problem

$$\min_{\theta\in\Gamma} R^*_{q,\theta}(J) + \frac{d}{\theta}$$

$$\Gamma = \{\theta > 0 : R^*_{q,\theta}(J) < \infty\}$$

Cross Entropy Method

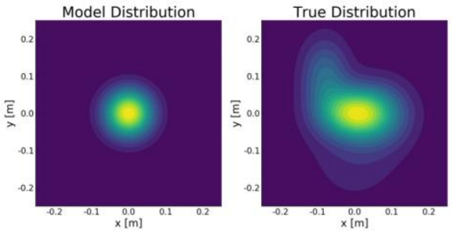

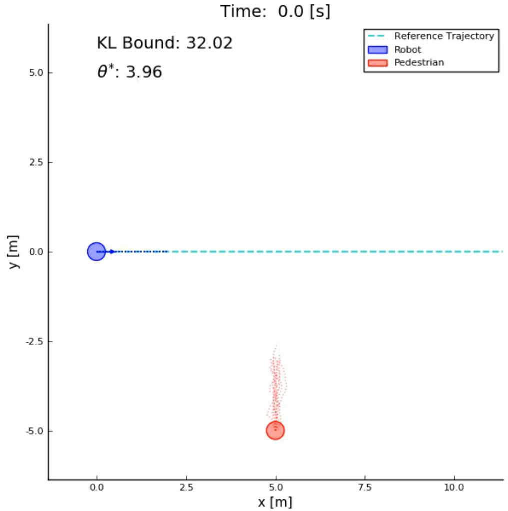

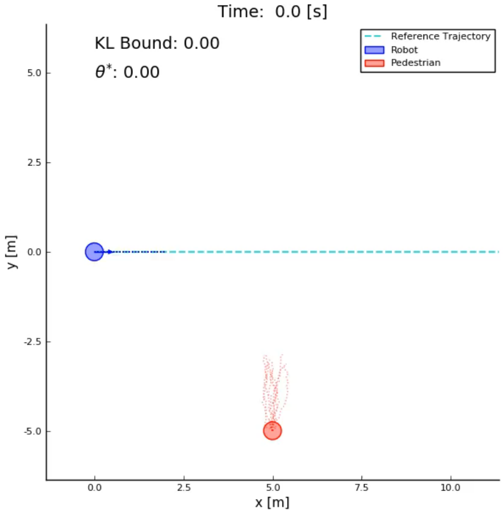

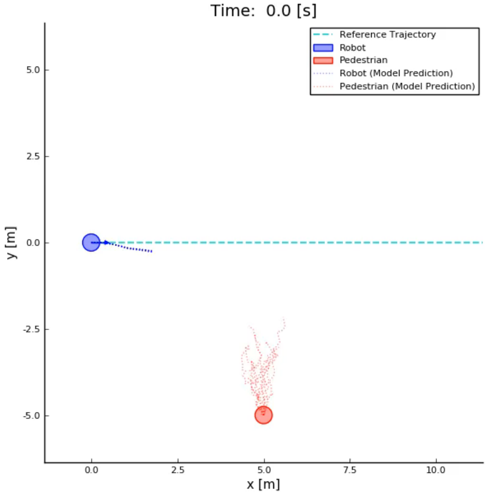

RAT iLQR (Ours)

0/30 Collisions

iLQG

1/30 Collisions

PETS

4/30 Collisions

Benefits of Risk Auto-Tuning

Conventional Risk-Sensitive Optimal Control

- No absolute scale
- Task-dependency

RAT iLQR

- No need for manual tuning

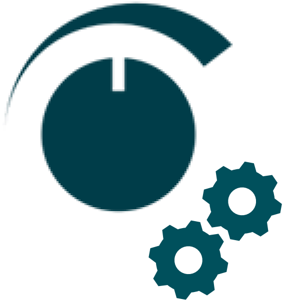

Efficiency of Risk Auto-Tuning

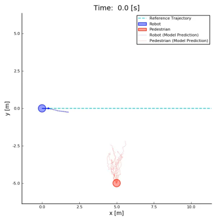

Conventional Risk-Sensitive Optimal Control

0/30 Collisions

Avg. Tracking Error: 0.38

RAT iLQR

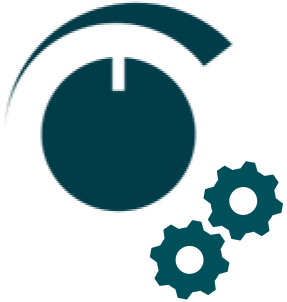

0/30 Collisions

Avg. Tracking Error: 0.32

Online Trajectory Planning Algorithms for Robotic Systems under Uncertainty in Interactive Environments

- Model-based approach when possible
- Probabilistic treatment of uncertainty

Active Reduction of Uncertainty

Formulated as
Belief Space Planning

Proposed Stochastic SAC

- Computationally efficient
- Handles stochasticity
- Considers long-term effect
- Outperforms prior methods

Resilience to Randomness

Formulated as
Risk-Sensitive Optimal Control

Proposed Risk-Sensitive SAC

- Nonlinear systems
- Arbitrary distributions
- Scalable to interaction with ~50 humans

Resilience to Imperfect Models

Formulated as
Distributionally Robust Control

Proposed RAT iLQR

- Based on Risk-Sensitive Control
- Nonlinear Systems
- Continuous Distributions
- Locally-optimal feedback policy

Planning Module ≠ Autonomy Stack

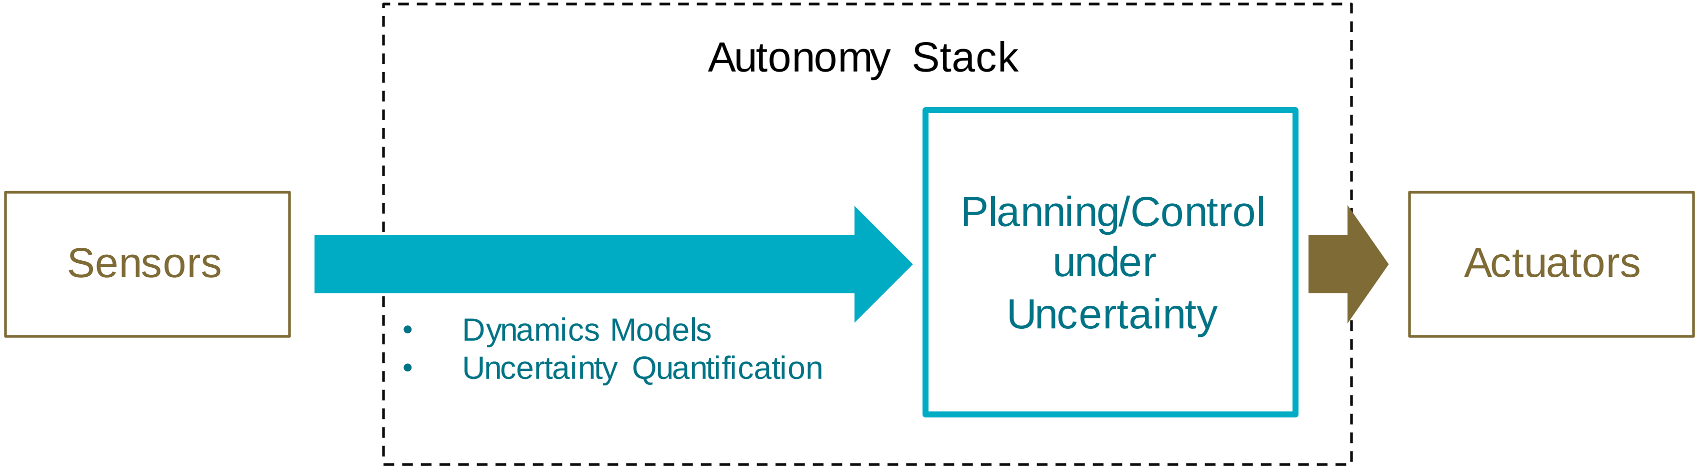

Planning as Part of Data-Driven Systems

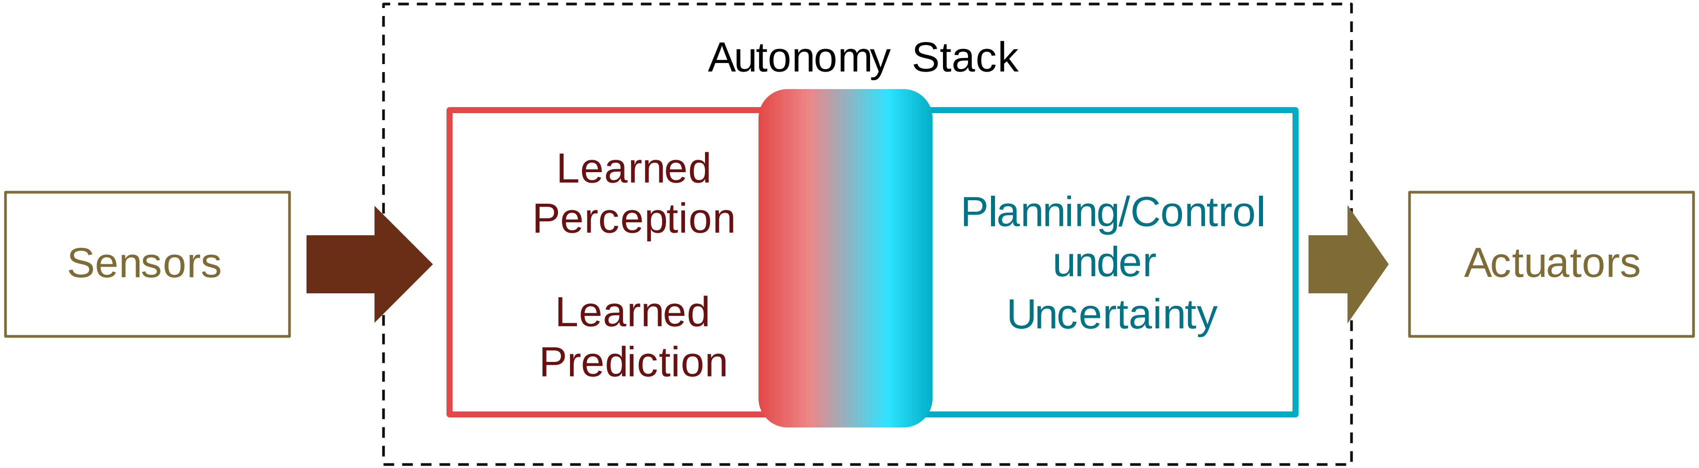

Q & A

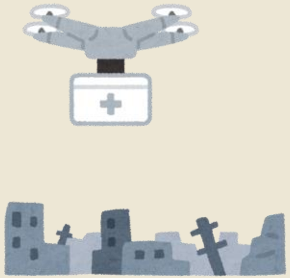

Thank you!