Path Planning Under Initial Uncertainty

A. Dhillon, A. Farah, D. Qi,
T. Reynoso, Q. Song, A. Vladimirsky

Center for Applied Mathematics
Cornell University

Dec. 9, 2019, SCAN Seminar

Outline

Cornell University

1. Trajectories to Known Targets
2. Delayed Target Identification Time: Fixed T
3. Random Target Identification Time
4. Different Notions of Robustness

Outline

Cornell University

What Are Optimal Trajectories?

Cornell University

Optimal trajectories from $\mathbf{x}_0$ to every target:

$$\begin{cases} |\nabla u(\mathbf{x})| f(\mathbf{x}) = K(\mathbf{x}), \\ u(\mathbf{x}_0) = 0. \end{cases} \qquad u(\mathbf{x}) = u(\mathbf{x}; \mathbf{x}_0).$$

Isotropic speed: $f(\mathbf{x})$; Running cost $K(\mathbf{x})(=1)$.

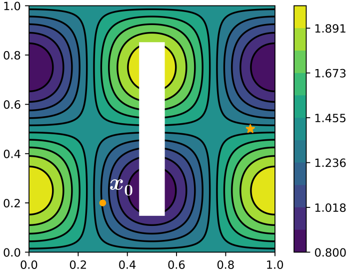

$f(\mathbf{x})$

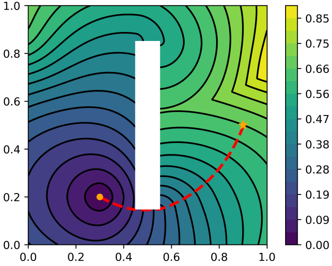

$u(\mathbf{x})$

What Are Optimal Trajectories?

$$\begin{cases} |\nabla u(\mathbf{x})| f(\mathbf{x}) = K(\mathbf{x}), \\ u(\mathbf{x}_0) = 0. \end{cases} \qquad u(\mathbf{x}) = u(\mathbf{x}; \mathbf{x}_0).$$

Optimal trajectories to **all** targets by solving only one PDE!

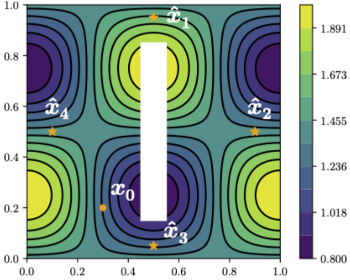

$f(\mathbf{x})$

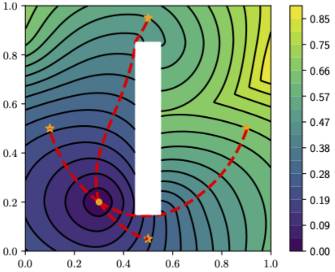

Fastest paths to 4 targets

Alternatively: Set Target as Source

Cornell University

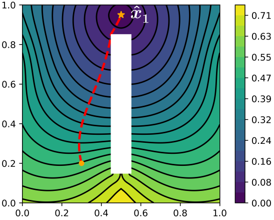

$$u_1(\mathbf{x}) = u(\mathbf{x}; \hat{\mathbf{x}}_1)$$

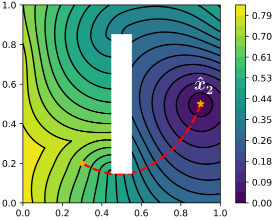

$$u_2(\mathbf{x}) = u(\mathbf{x}; \hat{\mathbf{x}}_2)$$

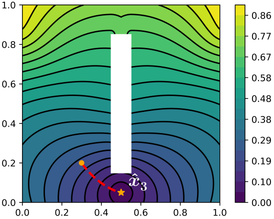

$$u_3(\mathbf{x}) = u(\mathbf{x}; \hat{\mathbf{x}}_3)$$

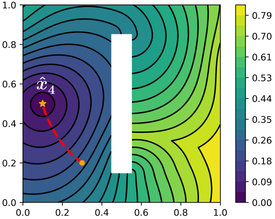

$$u_4(\mathbf{x}) = u(\mathbf{x}; \hat{\mathbf{x}}_4)$$

Outline

Cornell University

1. Trajectories to Known Targets
2. Delayed Target Identification Time: Fixed T
3. Random Target Identification Time
4. Different Notions of Robustness

True Target Revealed Later at T

Multiple likely targets $\hat{\mathbf{x}}_1, \cdots, \hat{\mathbf{x}}_4$.

- Target probabilities: $\hat{\mathbf{p}} = (0.2, 0.2, 0.3, 0.3)$.
- Certainty time: the true target is identified at T.

Question: What should we do until the time T?

True Target Revealed Later at T

Cornell University

Multiple likely targets $\hat{\mathbf{x}}_1, \cdots, \hat{\mathbf{x}}_4$.

- Target probabilities: $\hat{\mathbf{p}} = (0.2, 0.2, 0.3, 0.3)$.
- Certainty time: the true target is identified at T.

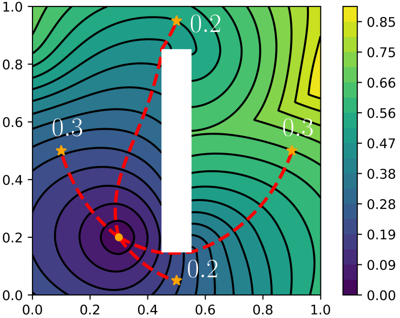

Question: What should we do until the time T?

Minimizing Expected Time-to-target

Cornell University

Expected time after **fixed** T:

$$q(\mathbf{x}) = \mathbb{E}[u(\mathbf{x};\hat{\mathbf{x}})] = \sum_{i=1}^{4} \hat{p}_i \, u_i(\mathbf{x}).$$

It suffices to select a waypoint $\mathbf{s} \in \underset{\{\mathbf{x}:\, u(\mathbf{x};\mathbf{x}_0)\leq T\}}{\arg\min} q(\mathbf{x})$!

Minimizing Expected Time-to-target

Expected time after **fixed** T:

$$q(\mathbf{x}) = \mathbb{E}[u(\mathbf{x};\hat{\mathbf{x}})] = \sum_{i=1}^{4} \hat{p}_i \, u_i(\mathbf{x}).$$

It suffices to select a waypoint $\mathbf{s} \in \underset{\{\mathbf{x}:\, u(\mathbf{x};\mathbf{x}_0)\leq T\}}{\arg\min} \; q(\mathbf{x})$!

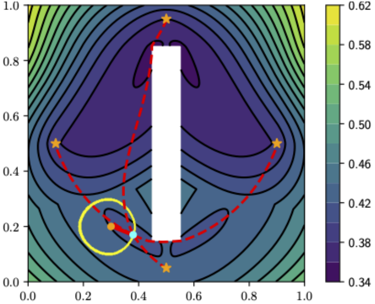

$T = 0.08$

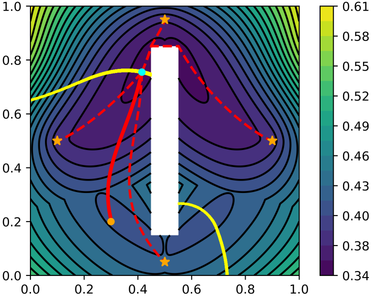

$T = 0.4$

Avoid a Not-Yet-Known Storm [SB16]

Cornell University

- Three possible storm locations with $\hat{\mathbf{p}} = (\hat{p}_1, \hat{p}_2, \hat{p}_3)$.
- Certainty time: the true location is revealed at T.
- Steady speed: $f(\mathbf{x}) = 1$. Safe region: $K(\mathbf{x}) = 1$.
- Storm region: $K_i(\mathbf{x}) = 1 + \alpha[1 - (\mathbf{x} - \hat{\mathbf{x}}_i)^\top A(\mathbf{x} - \hat{\mathbf{x}}_i)]^\gamma$.

It suffices to select a waypoint $\mathbf{s} \in \underset{\{\mathbf{x}:\, u(\mathbf{x};\mathbf{x}_0) \leq T\}}{\arg\min} q(\mathbf{x})$!

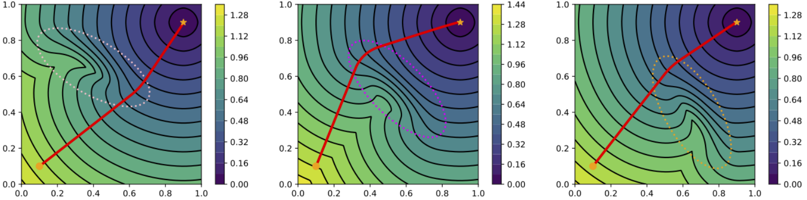

$u_1(\mathbf{x})$ $u_2(\mathbf{x})$ $u_3(\mathbf{x})$

$$q(\mathbf{x}) = \sum_{i=1,2,3} \hat{p}_i u_i(\mathbf{x})$$

Avoid a Not-yet-known Storm [SB16]

Cornell University

- Three possible storm locations with $\hat{\mathbf{p}} = (\hat{p}_1, \hat{p}_2, \hat{p}_3)$.
- Certainty time: the true location is revealed at T.
- Steady speed: $f(\mathbf{x}) = 1$. Safe region: $K(\mathbf{x}) = 1$.

It suffices to select a waypoint $\mathbf{s} \in \underset{\{\mathbf{x}:\, u(\mathbf{x};\mathbf{x}_0) \leq T\}}{\arg\min} \; q(\mathbf{x})$!

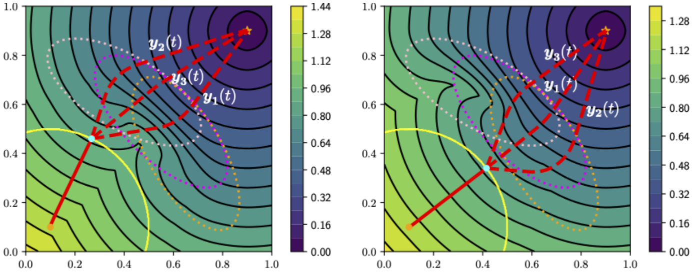

$\hat{\mathbf{p}} = (0.1, 0.8, 0.1)$ $\hat{\mathbf{p}} = (0.8, 0.1, 0.1)$

Outline

Cornell University

1 Trajectories to Known Targets

2 Delayed Target Identification Time: Fixed T

3 Random Target Identification Time

4 Different Notions of Robustness

Discrete Random Certainty Time T

Discretely distributed T: e.g. $T \in \{T_1, T_2\}$, $\mathbf{p} = (p_1, p_2)$.

$$P(T = T_1) = p_1, \quad P(T = T_2) = p_2, \quad p_1 + p_2 = 1.$$

Solve a series of time-dependent PDEs backward in time.

$t \in [T_1, T_2)$:

Target not identified at T_1 but will be revealed at T_2.

$$\begin{cases} \frac{\partial u^B}{\partial t} - |\nabla u^B| f(\mathbf{x}) + 1 = 0, \\ u^B(\mathbf{x}, T_2) = q(\mathbf{x}). \end{cases}$$

Discrete Random Certainty Time T

Discretely distributed T: $T \in \{T_1, T_2\},\ \mathbf{p} = (p_1, p_2)$.

$$P(T = T_1) = p_1,\ P(T = T_2) = p_2,\ p_1 + p_2 = 1.$$

Solve a series of time-dependent PDEs backward in time.

$t \in [0, T_1)$:

Target might be identified at T_1 with probability p_1.

$$\begin{cases} \frac{\partial u^A}{\partial t} - |\nabla u^A| f(\mathbf{x}) + 1 = 0, \\ u^A(\mathbf{x}, T_1) = p_1 q(\mathbf{x}) + p_2 u^B(\mathbf{x}, T_1). \end{cases}$$

Discrete Random Certainty Time T

Cornell University

Assumption: $T \in \{T_1, T_2\}$, $\mathbf{p} = (p_1, p_2)$.

$$P(T = T_1) = p_1, \quad P(T = T_2) = p_2, \quad p_1 + p_2 = 1.$$

The second part of trajectory is needed only if $T \neq T_1$!

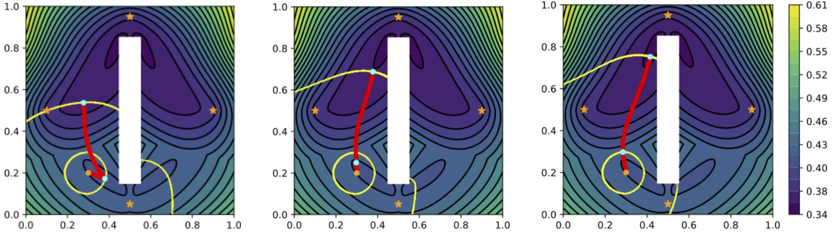

$\mathbf{p} = (0.9, 0.1)$ $\quad$ $\mathbf{p} = (0.55, 0.45)$ $\quad$ $\mathbf{p} = (0.1, 0.9)$

$$T_1 = 0.08, \; T_2 = 0.4$$

$\mathbf{p}$, $\hat{\mathbf{p}}$ Interdependence : A Drone Example

Previous example: $P(T = T_i)$ is independent of $P(u = u_i)$.

Send drones out to detect the true target:

$$P(T = T_i) = P(u = u_i \,|\, u \neq u_j,\, j \in I_{\text{visited}})$$

where I_{visited} is the set of already-ruled-out targets.

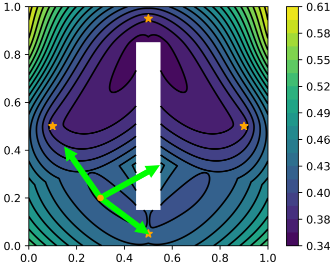

Exponentially Distributed T

Assumption: $P(T > t + \tau \mid T \geq t) = e^{-\lambda\tau}$.

Optimal trajectories can be found by solving [AV14]

$$\lambda \left(u^\lambda(\mathbf{x}) - q(\mathbf{x}) \right) + |\nabla u^\lambda(\mathbf{x})| f(\mathbf{x}) = K(\mathbf{x}).$$

- Local minima of $q(\mathbf{x})$ are the only possible waypoints!
- The waypoint might not even be reached!

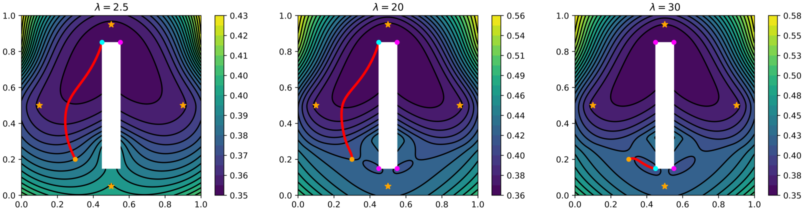

$\lambda = 2.5$ $\lambda = 20$ $\lambda = 30$

Outline

Cornell University

Optimize the Worst-Case Scenario

Select waypoint $\bar{\mathbf{s}}$ by replacing $q(\mathbf{x})$ with

$$\bar{q}(\mathbf{x}) = \max_{i=1,2,3,4} u_i(\mathbf{x}).$$

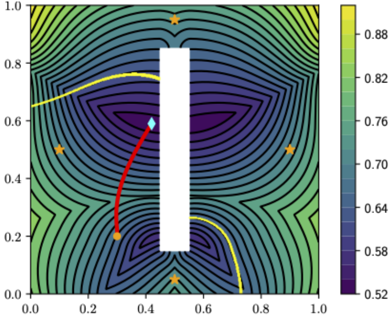

Fixed $T = 0.4$

Risk-sensitive Control

Select waypoint $\mathbf{s}_\beta$(●) minimizing

$$q_\beta(\mathbf{x}) = \mathbb{E}[e^{\beta u(\mathbf{x})}]$$

for $\beta > 0$, this penalizes large values of $u(\mathbf{x})$.

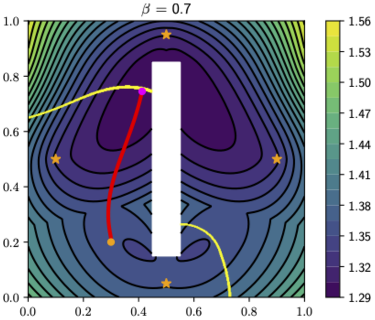

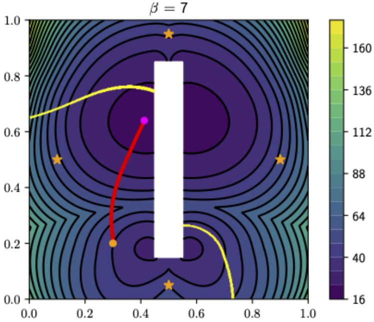

$\beta \to 0, \ \mathbf{s}_\beta \to \mathbf{s}.$ $\beta \to \infty, \ \mathbf{s}_\beta \to \bar{\mathbf{s}}.$

Distributionally Robust Waypoint

Select a waypoint $\mathbf{s}_\delta$ (●) minimizing

$$\tilde{q}_\delta(\mathbf{x}) = \max_{\tilde{\mathbf{p}} \in B_\delta(\hat{\mathbf{p}})} \mathbb{E}_{\tilde{\mathbf{p}}}[u(\mathbf{x})].$$

where $B_\delta(\hat{\mathbf{p}})$ is the δ-Wasserstein-distance ball centered at $\hat{\mathbf{p}}$.

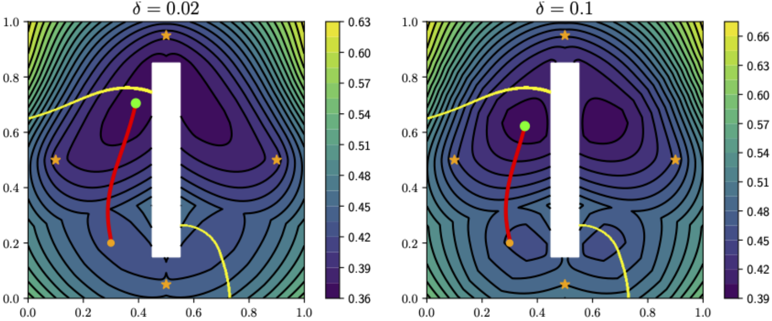

$\delta \to 0,\ \mathbf{s}_\delta \to \mathbf{s}.$ $\delta \to 1,\ \mathbf{s}_\beta \to \bar{\mathbf{s}}.$

Optimize the Average, Bound the Worst

Minimize $q(\mathbf{x})$ with a constraint $\bar{q}(\mathbf{x}) \leq C$ [EGSV12].

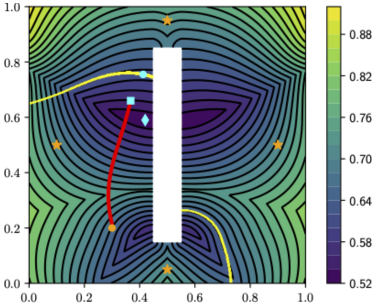

$C = 0.56$

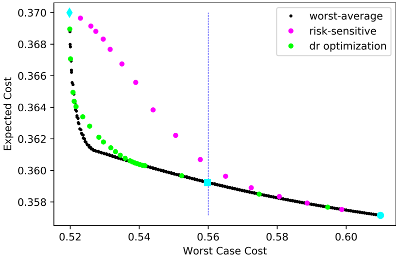

$(\bar{q}, q)$-Pareto front

Risk-sensitive (•) and distributionally robust (•) waypoints can be far from $(\bar{q}, q)-$Pareto optimal!

Chance Constrained Optimization

Risk: probability of $u(\mathbf{x}) > C$ is

$$r(\mathbf{x}) = \sum_{i=1}^{m} \hat{p}_i \, \chi_{\{u_i(\mathbf{x})>C\}}(\mathbf{x}).$$

Risky waypoints can still be used if we control their "frequency"!

Probabilistic strategy: select a pdf θ over a grid of n points.

$$\begin{aligned}
\text{minimize} \quad & \sum_{j=1}^{n} \theta_j q(\mathbf{x}_j) \\
\text{subject to} \quad & \sum_{j=1}^{n} r(\mathbf{x}_j)\theta_j \leq \epsilon, \\
& \sum_{j=1}^{n} \theta_j = 1, \quad \theta_j \geq 0, \; j = 1, \cdots, n.
\end{aligned}$$

Chance Constrained Optimization

Risk: probability of $u(\mathbf{x}) > C$ is

$$r(\mathbf{x}) = \sum_{i=1}^{m} \hat{p}_i \, \chi_{\{u_i(\mathbf{x}) > C\}}(\mathbf{x}).$$

Risky waypoints can still be used if we control their "frequency"!

Probabilistic strategy: select a pdf θ over a grid of n points.

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^{n} \theta_j q(\mathbf{x}_j) \\ \text{subject to} \quad & \sum_{j=1}^{n} r(\mathbf{x}_j)\theta_j \leq \epsilon, \\ & \sum_{j=1}^{n} \theta_j = 1, \quad \theta_j \geq 0, \; j = 1, \cdots, n. \end{aligned}$$

Chance Constrained Optimization

Cornell University

- An LP problem, but can be solved in $\mathcal{O}(nm)$ time for $\forall\epsilon$!
- **Theorem:** $\exists$ an optimal θ^* that uses only 2 waypoints regardless of n and m.

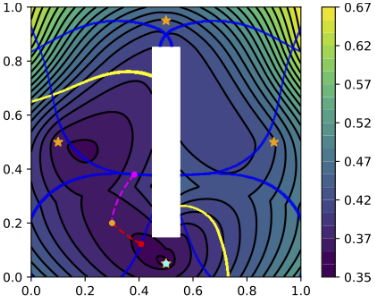

$\hat{\mathbf{p}} = (0.18, 0.18, 0.35, 0.29)$

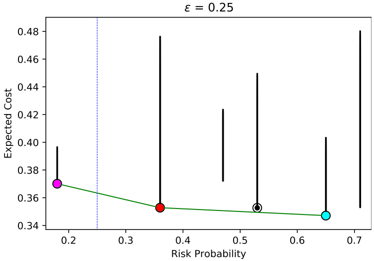

$(r, q)-$pairs

$$C = 0.365, \quad \epsilon = 0.25, \quad \theta_{\color{red}\bullet} \approx 0.39, \quad \theta_{\color{magenta}\bullet} \approx 0.61.$$

References I

Cornell University

June Andrews and Alexander Vladimirsky.
Deterministic control of randomly-terminated processes.
Interfaces and Free Boundaries, 16, 10 2014.

A. Dhillon, A. Farah, D. Qi, T. Reynoso, Q. Song, and A. Vladimirsky.
Path planning under initial uncertainty.
In preparation, 2019.

Stefano Ermon, Carla Gomes, Bart Selman, and Alexander Vladimirsky.
Probabilistic planning with non-linear utility functions and worst-case guarantees.
In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 965–972. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

Alexander V Sadovsky and Karl D Bilimoria.
Risk-hedged approach for re-routing air traffic under weather uncertainty.
In *16th AIAA Aviation Technology, Integration, and Operations Conference*, page 3601, 2016.

Appendix

Cornell University

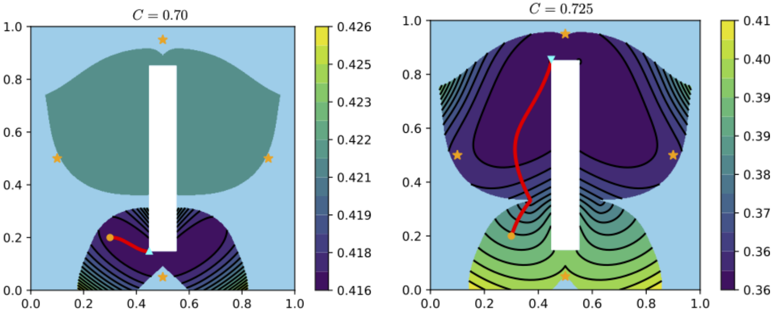

Paper in preparation [DFQ+19]!