



Joint Chance-Constrained Dynamic Programming

Masahiro Ono

Keio University



**Yoshiaki Kuwata
J. (Bob) Balaram**





Future Mars entry, descent & landing requires to:

- Optimize lander's control law/landing site
- Manage risk of landing failure
- Minimize rover's driving distance to science targets



$$\min_{u_{\tau:\tau+N-1}} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=\tau}^{\tau+N-1} g_k(x_k, u_k) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=\tau}^{\tau+N-1} h(x_k) \preceq 0 \right\} \geq 1 - \Delta \quad \leftarrow \text{Risk bound}$$

$$\bigwedge_{k=\tau}^{\tau+N-1} x_{k+1} = f(x_k, u_k, w_k)$$

- Significant body of work in Chance-constrained MPC
- Applications
 - Building control
 - Power grid control
 - Path planning



$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=1}^N x_k \in \mathcal{X}_k \mid x_0 \right\} \geq 1 - \Delta \quad \leftarrow \text{Risk bound}$$

$$\bigwedge_{k=0}^{N-1} x_{k+1} = f(x_k, u_k, w_k)$$

DP can deal with

- nonlinear dynamics
- nonconvex feasible region
- discrete state space
- non-Gaussian probability distribution
- state-dependent disturbance

State constraint: $x_k \in \mathcal{X}_k$

Control constraint: $u_k \in \mathcal{U}_k(x_k)$

Control policy: $\mu_k : \mathcal{X} \mapsto \mathcal{U}(x_k)$



$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=1}^N x_k \in \mathcal{X}_k \mid x_0 \right\} \geq 1 - \Delta \quad \leftarrow \text{Risk bound}$$

$$\bigwedge_{k=0}^{N-1} x_{k+1} = f(x_k, u_k, w_k)$$

DP optimizes control policy
(i.e., feedback control law)

State constraint: $x_k \in \mathcal{X}_k$

Control constraint: $u_k \in \mathcal{U}_k(x_k)$

Control policy: $\mu_k : \mathcal{X} \mapsto \mathcal{U}(x_k)$



CCDP: Challenges and Approaches **JPL**

Darts Lab

Mobility & Robotic Systems - 347

$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$
$$\text{s.t. } \mathbb{E} \left\{ c_N(x_N) + \sum_{k=0}^{N-1} c_k(x_k, \mu_k(x_k)) \right\} \leq \Delta$$

Challenge 1: Standard DP cannot solve a constrained optimization

Constrained DP/MDP (Rossman 1977, Beulter & Ross 1985, Williams et al. 2008) can be applied only to constraints in the same form as the objective function

State constraint: $x_k \in \mathcal{X}_k$

Control constraint: $u_k \in \mathcal{U}_k(x_k)$

Control policy: $\mu_k : \mathcal{X} \mapsto \mathcal{U}(x_k)$



CCDP: Challenges and Approaches **JPL**

Darts Lab

Mobility & Robotic Systems - 347

$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=1}^N x_k \in \mathcal{X}_k \mid x_0 \right\} \geq 1 - \Delta \quad \leftarrow \text{Risk bound}$$

Challenge 1: Standard DP cannot solve a constrained optimization

Approach 1: Solve the dual optimization problem

Challenge 2: Cannot obtain exact solution (mainly due to duality gap)

Approach 2: Obtain error bound

State constraint: $x_k \in \mathcal{X}_k$

Control constraint: $u_k \in \mathcal{U}_k(x_k)$

Control policy: $\mu_k : \mathcal{X} \mapsto \mathcal{U}(x_k)$



CCDP: Dual Reformulation



Darts Lab

Mobility & Robotic Systems - 347

$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=1}^N x_k \in \mathcal{X}_k \mid x_0 \right\} \geq 1 - \Delta$$

Indicator function

$$I_k(x_k) = \begin{cases} 1 & (x_k \notin \mathcal{X}_k) \\ 0 & (\text{Otherwise}) \end{cases}$$



CCDP: Dual Reformulation



$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \Pr \left\{ \bigwedge_{k=1}^N x_k \in \mathcal{X}_k \mid x_0 \right\} = 1 - \Pr \left\{ \bigvee_{k=1}^N x_k \notin \mathcal{X}_k \mid x_0 \right\}$$

Indicator function

$$I_k(x_k) = \begin{cases} 1 & (x_k \notin \mathcal{X}_k) \\ 0 & (\text{Otherwise}) \end{cases}$$

De Morgan's law

$$\geq 1 - \sum_{k=1}^N \Pr \{x_k \notin \mathcal{X}_k \mid x_0\}$$

Boole's ineq.

$$= 1 - \sum_{k=1}^N \mathbb{E}\{I_k(x_k) \mid x_0\}$$

$$= 1 - \mathbb{E} \left\{ \sum_{k=1}^N I_k(x_k) \mid x_0 \right\}$$

Risk at *k*th time step

$$\Pr \{x_k \notin \mathcal{X}_k \mid x_0\}$$

1 time step



CCDP: Dual Reformulation



Darts Lab

Mobility & Robotic Systems - 347

$$\min_{\mu} \mathbb{E} \left\{ g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k)) \right\}$$

$$\text{s.t.} \quad \mathbb{E} \left\{ \sum_{k=1}^N I_k(x_k) \mid x_0 \right\} \leq \Delta \quad \leftarrow \text{Risk bound}$$

Indicator function

$$I_k(x_k) = \begin{cases} 1 & (x_k \notin \mathcal{X}_k) \\ 0 & (\text{Otherwise}) \end{cases}$$

Lagrangian

$$L_k^\lambda(x_k, u_k) = \begin{cases} g_0(x_0, u_0) & (k = 0) \\ g_k(x_k, u_k) + \lambda I_k(x_k) & (k = 1 \dots N - 1) \\ g_N(x_N) + \lambda I_N(x_N) & (k = N). \end{cases}$$

λ : Dual variable (Lagrange multiplier)



CCDP: Dual Reformulation



Darts Lab

Mobility & Robotic Systems - 347

Dual optimization problem

$$\max_{\lambda \geq 0} \min_{\mu} \mathbb{E} \left\{ \sum_{k=0}^N L_k^\lambda(x_k, \mu_k(x_k)) \right\} - \lambda \Delta$$

Indicator function

$$I_k(x_k) = \begin{cases} 1 & (x_k \notin \mathcal{X}_k) \\ 0 & (\text{Otherwise}) \end{cases}$$

Lagrangian

$$L_k^\lambda(x_k, u_k) = \begin{cases} g_0(x_0, u_0) & (k = 0) \\ g_k(x_k, u_k) + \lambda I_k(x_k) & (k = 1 \dots N - 1) \\ g_N(x_N) + \lambda I_N(x_N) & (k = N). \end{cases}$$

λ : Dual variable (Lagrange multiplier)



CCDP: Dual Reformulation



Darts Lab

Mobility & Robotic Systems - 347

Dual optimization problem

$$\max_{\lambda \geq 0} \min_{\boldsymbol{\mu}} \mathbb{E} \left\{ \sum_{k=0}^N L_k^\lambda(x_k, \mu_k(x_k)) \right\} - \lambda \Delta$$

Good news 1: Control policy (primal var) can be optimized by standard DP



CCDP: Dual Reformulation



Darts Lab

Mobility & Robotic Systems - 347

Dual optimization problem

$$\max_{\lambda \geq 0} \min_{\boldsymbol{\mu}} \mathbb{E} \left\{ \sum_{k=0}^N L_k^\lambda(x_k, \mu_k(x_k)) \right\} - \lambda \Delta = q(\lambda)$$

Good news 1: Control policy (primal var) can be optimized by standard DP



CCDP: Dual Solution



Darts Lab

Mobility & Robotic Systems - 347

Dual optimization problem

$$\max_{\lambda \geq 0} q(\lambda)$$

Good news 2: Dual optimization can be solved by zero-finding

- Condition for optimality: $0 \in \partial q(\lambda)$ (∂q : subgradient)

$$\underline{r_0^\lambda(x_0) - \Delta \in \partial q(\lambda)}$$

Excess risk

Risk-to-go function

$$r_0^\lambda(x_0) := \mathbb{E} \left\{ \sum_{k=1}^N I_k(x_k) \mid x_0, \mu^\lambda \right\}$$

Probability of failure given the optimal control policy μ^λ



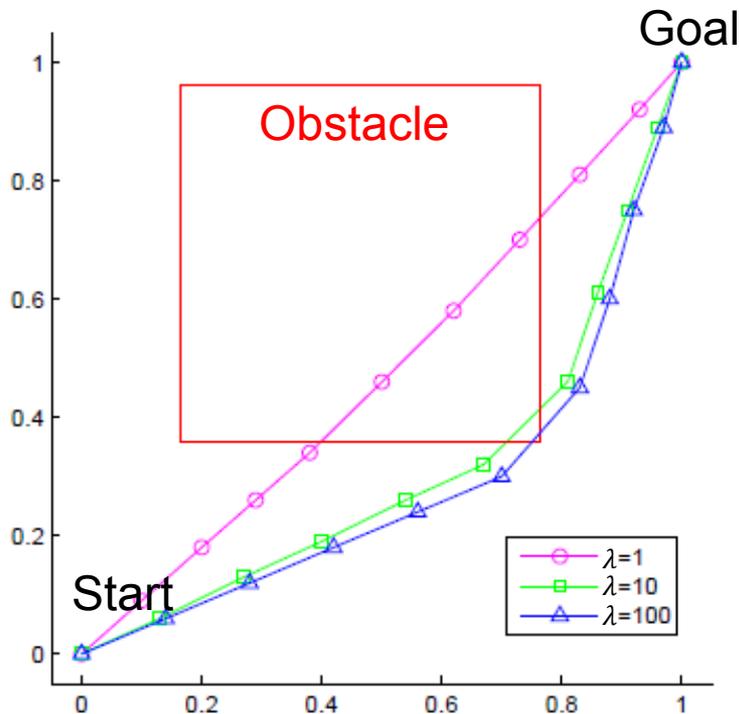
Intuition: Risk Management via Penalty **JPL**

Darts Lab

Mobility & Robotic Systems - 347

λ : Penalty of failure

$$\max_{\lambda \geq 0} \min_{\mu} \mathbb{E} \left\{ \sum_{k=0}^N (g_k(x_k, u_k) + \underbrace{\lambda I_k(x_k)}_{\text{Penalty cost of failure}}) \right\} - \lambda \Delta$$



Penalty λ	Prob. of failure
1	100.0%
10	37.5%
100	3.1%

Interpretation: CCDDP finds the level of penalty with which the risk of failure is Δ .

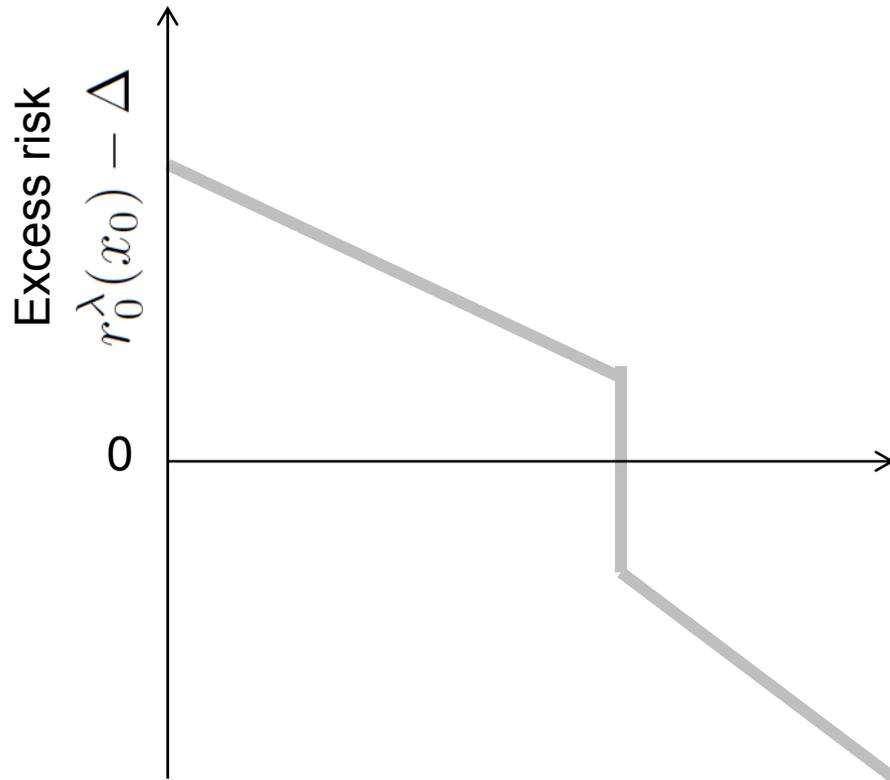


Algorithm

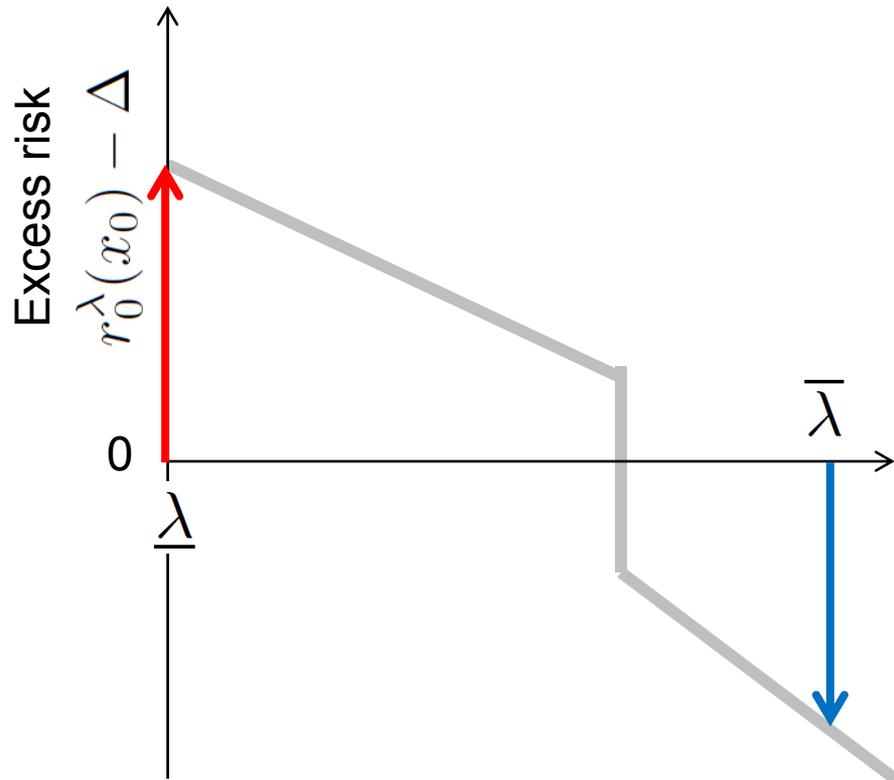


Darts Lab

Mobility & Robotic Systems - 347



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- **Initialize** $[\underline{\lambda} \ \bar{\lambda}]$
- **while** $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- **end while**

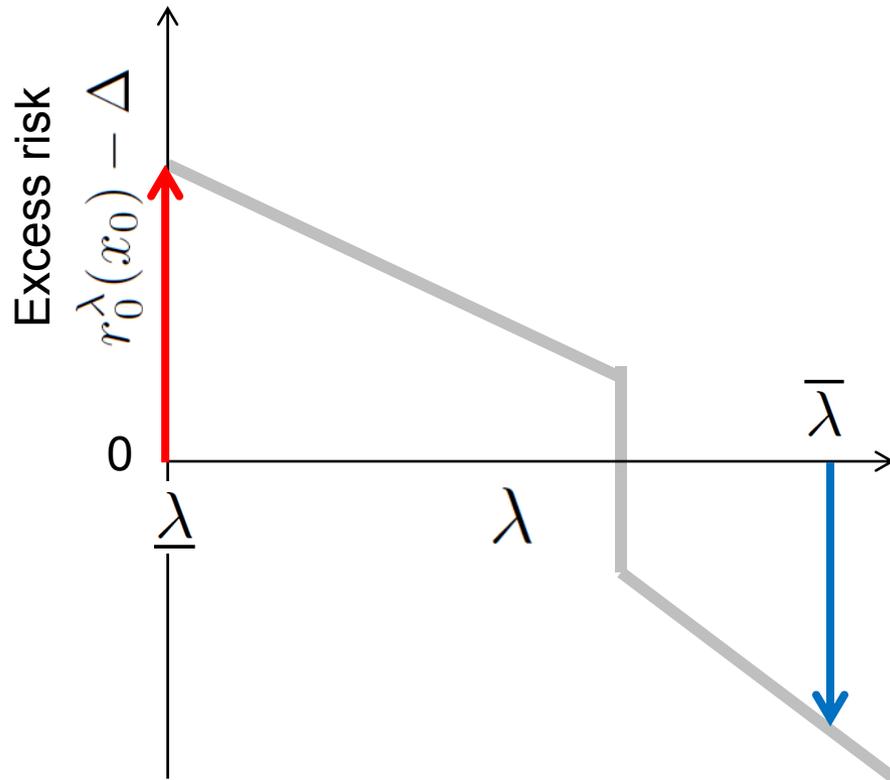


Algorithm

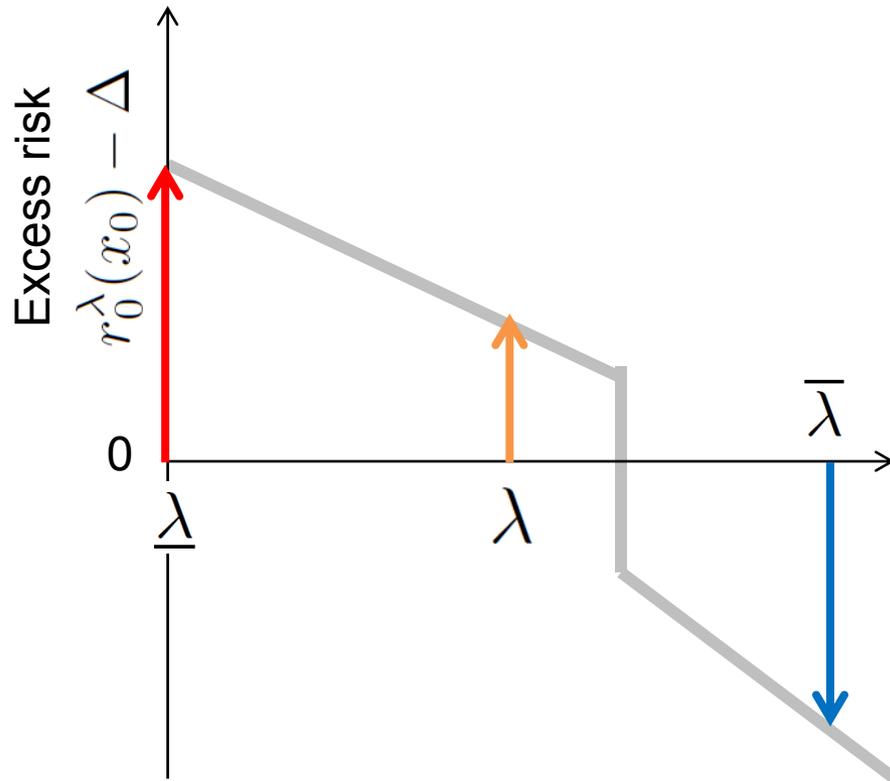


Darts Lab

Mobility & Robotic Systems - 347



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while

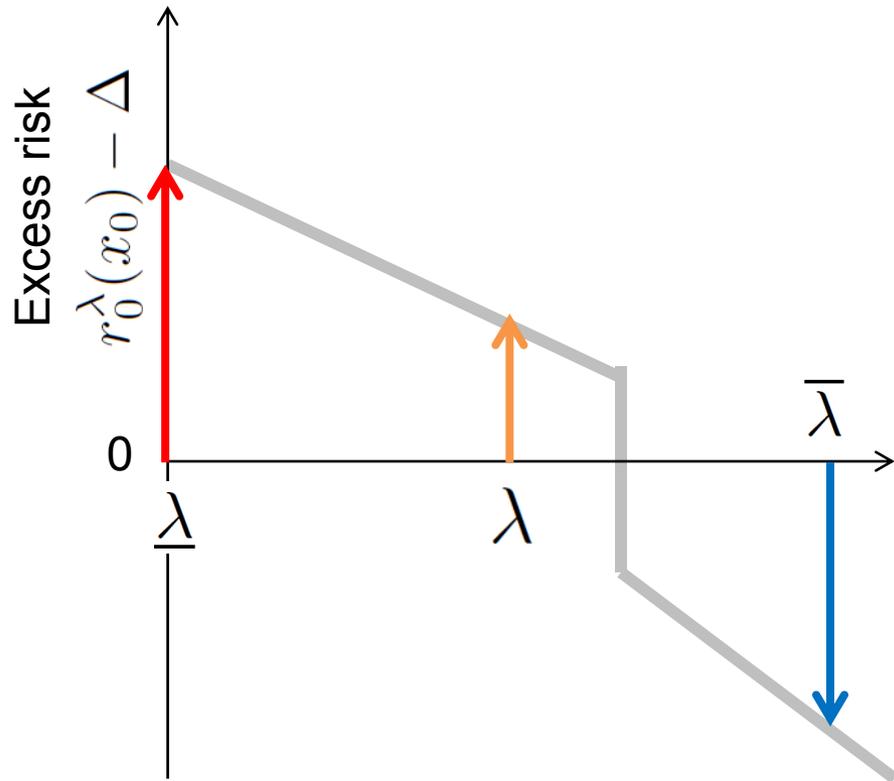


Algorithm

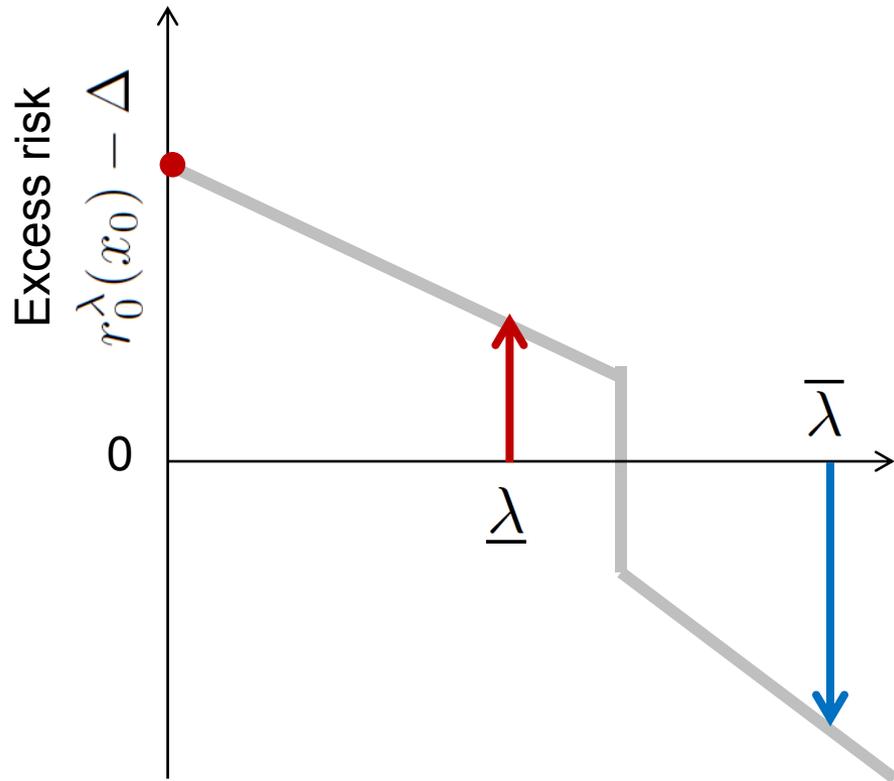


Darts Lab

Mobility & Robotic Systems - 347



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while

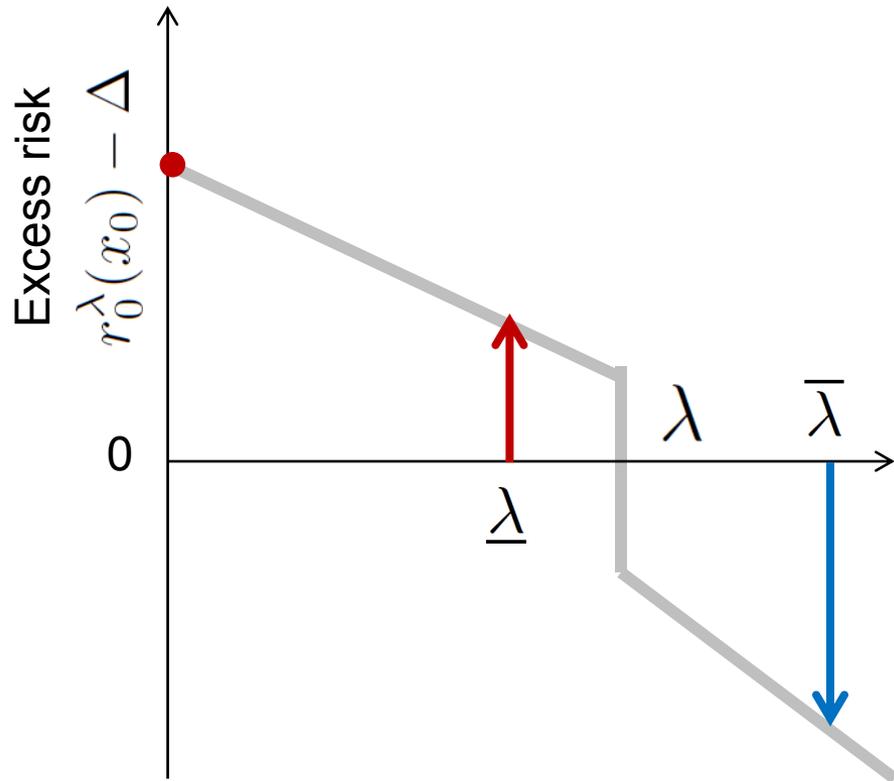


Algorithm

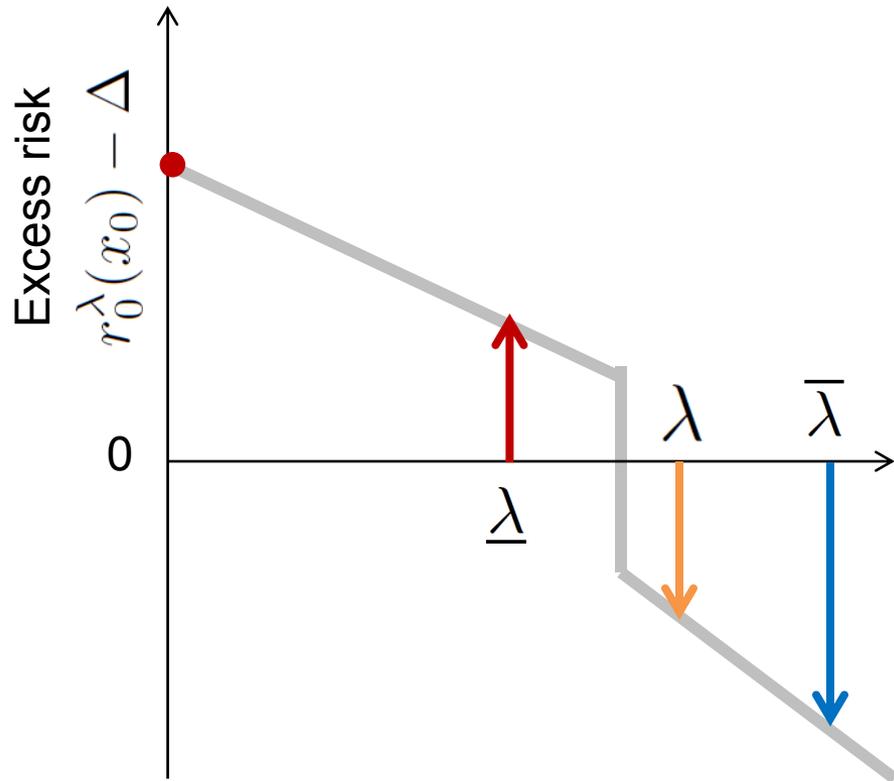


Darts Lab

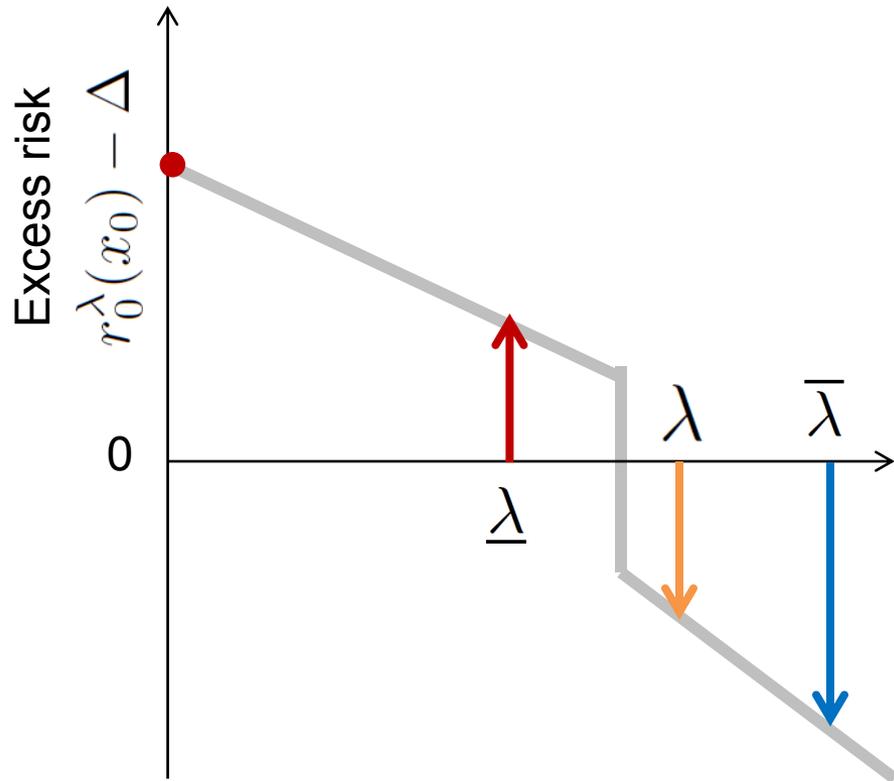
Mobility & Robotic Systems - 347



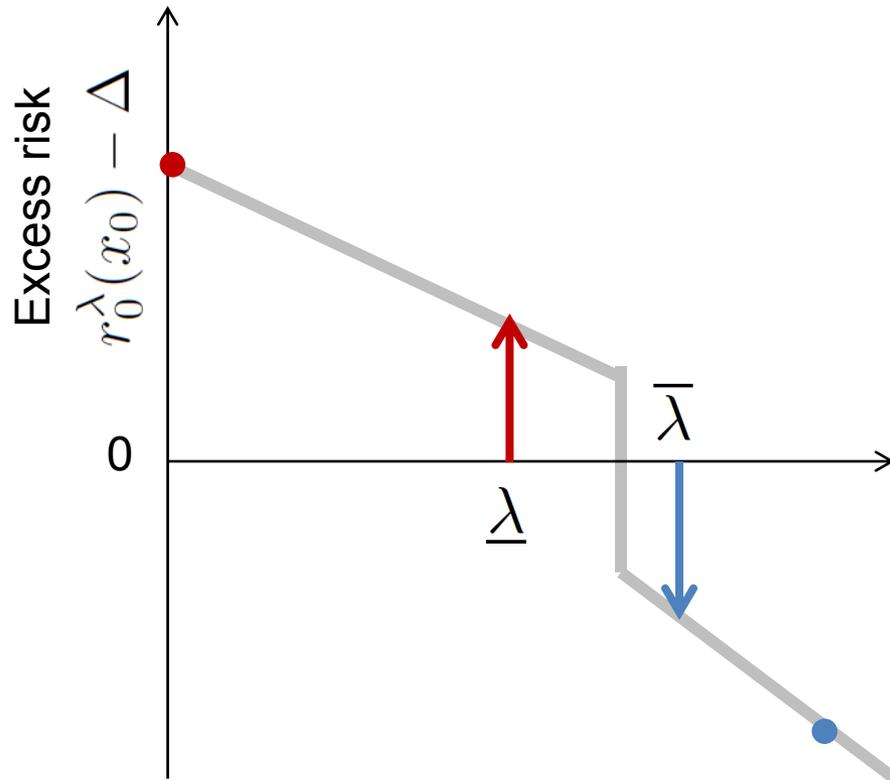
- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



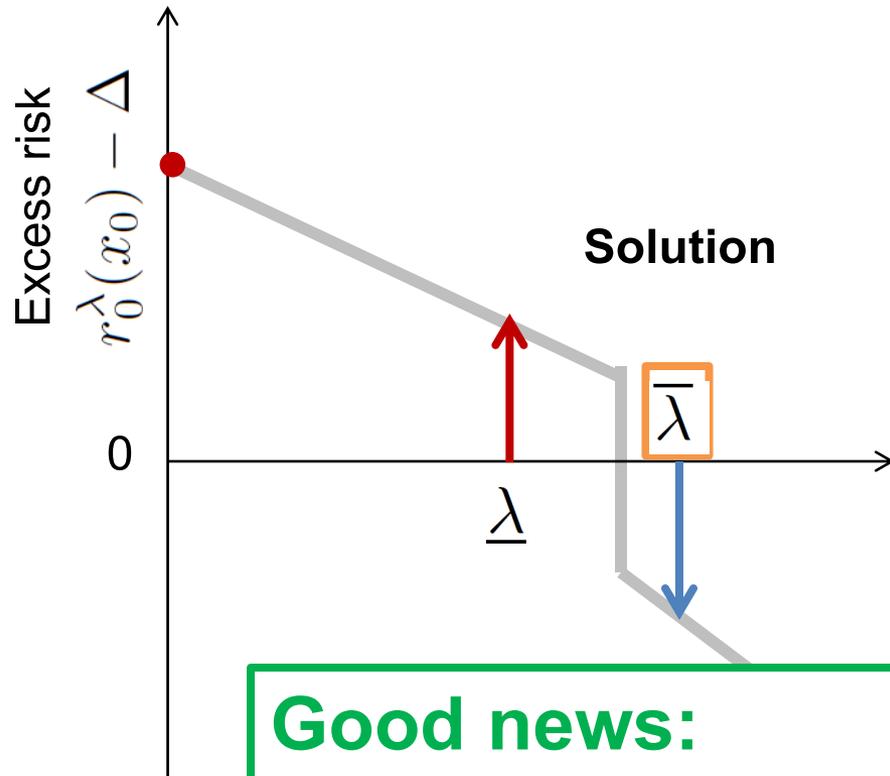
- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\lambda \leftarrow \lambda$

Good news:

- Solution is on the “safe side”
- Guaranteed to satisfy the risk constraint

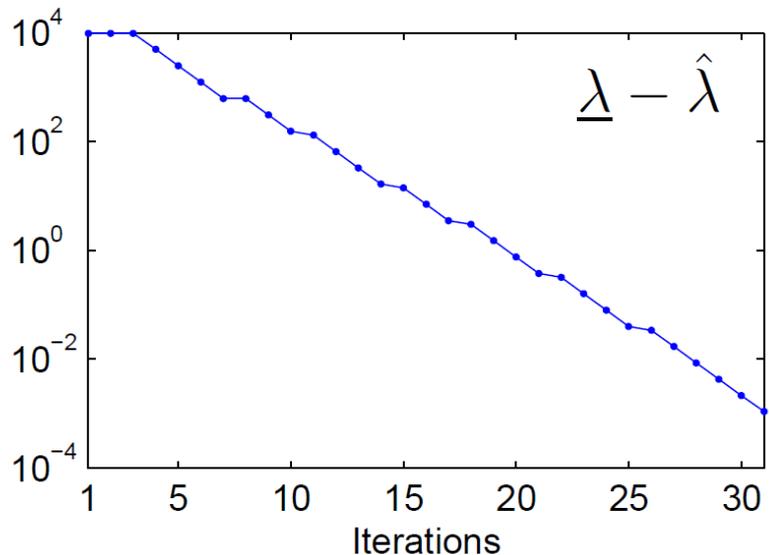
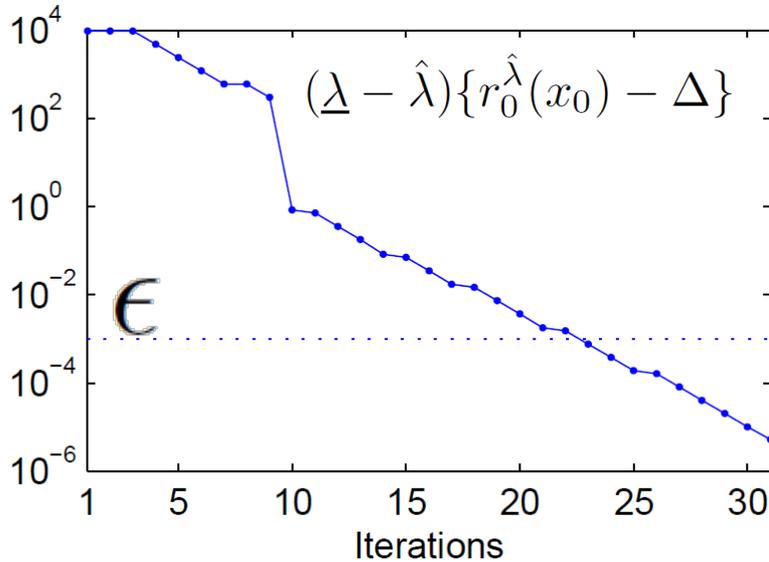


Algorithm



Darts Lab

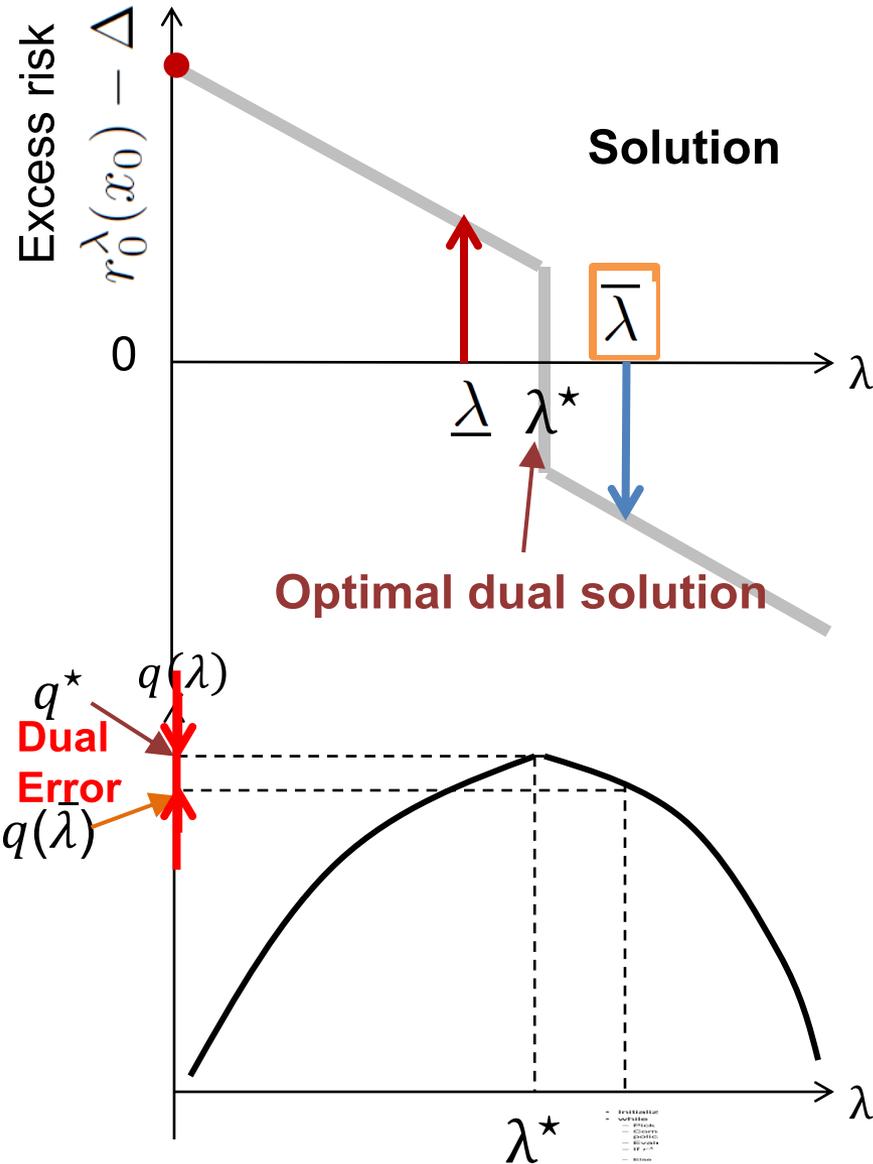
Mobility & Robotic Systems - 347



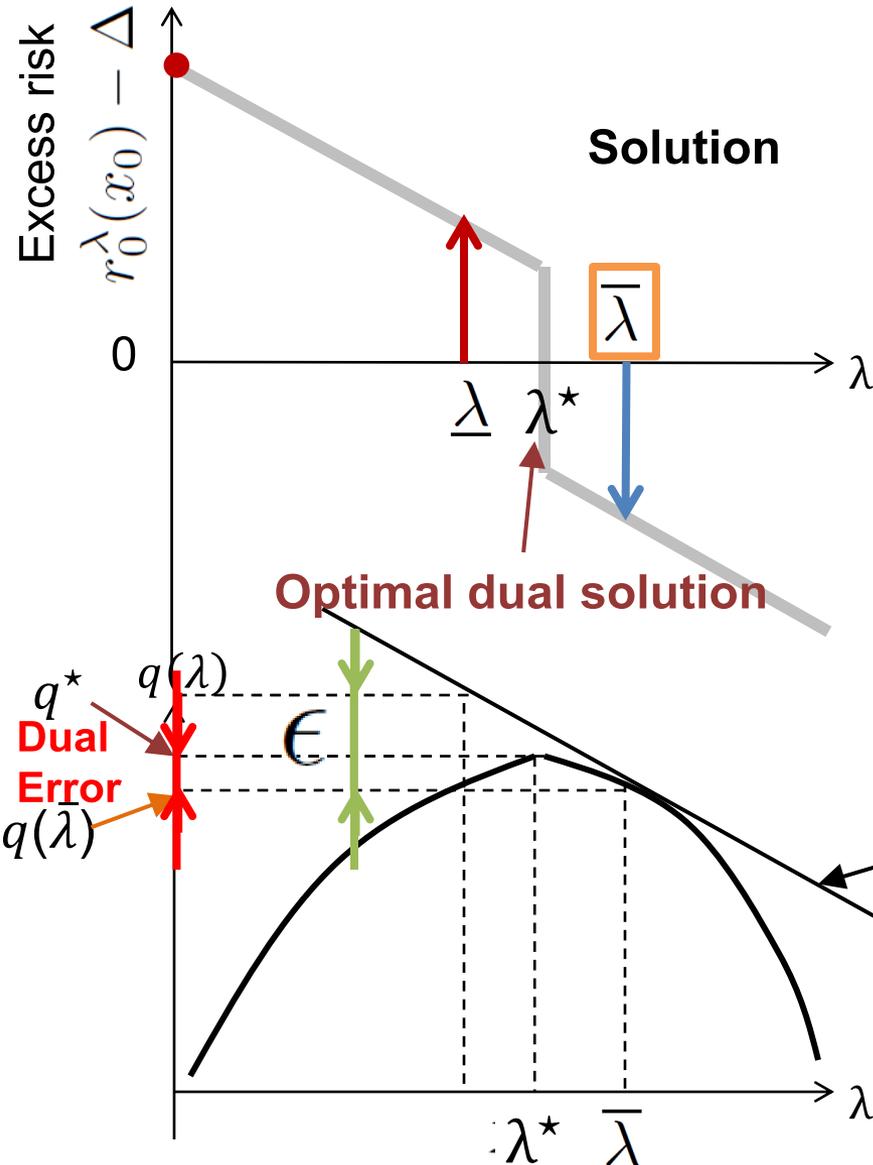
- Initialize $[\underline{\lambda} \bar{\lambda}]$
- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \bar{\lambda})$

Good news:

- Dual optimization converges exponentially
- Typically requires 10-20 iterations
- # of iterations does not grow with problem size
 - No exponential blowup



- Initialize $[\underline{\lambda} \ \bar{\lambda}]$
- **while** $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$
 - Pick $\lambda \in (\underline{\lambda} \ \bar{\lambda})$
 - Compute optimal control policy μ^λ by solving a DP
 - Evaluate risk r^λ
 - If $r^\lambda - \Delta > 0$
 - $\underline{\lambda} \leftarrow \lambda$
 - Else
 - $\bar{\lambda} \leftarrow \lambda$
- end while



- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$

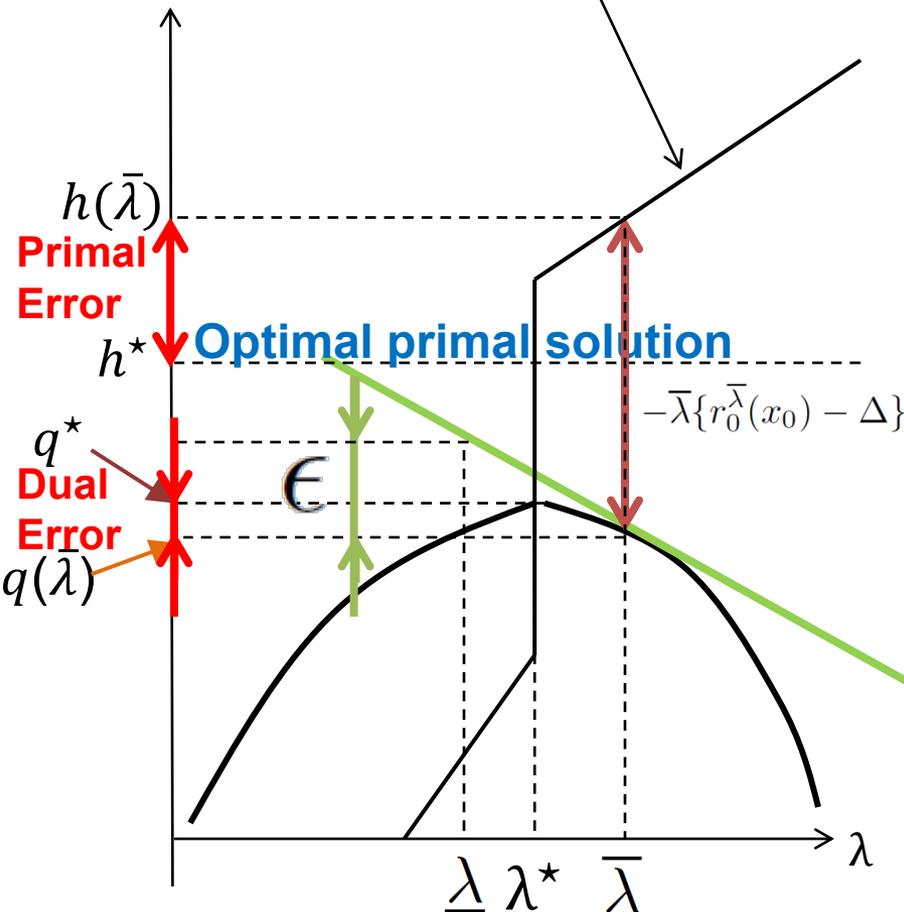
Theorem 1

$$q^* - q(\bar{\lambda}) \leq \epsilon$$

Cost with the optimal policy given λ

$$h(\lambda) = q(\lambda) - \lambda\{r_0^\lambda(x_0) - \Delta\}$$

- while $(\underline{\lambda} - \bar{\lambda})\{r_0^{\bar{\lambda}}(x_0) - \Delta\} > \epsilon$



Theorem 1

$$q^* - q(\bar{\lambda}) \leq \epsilon$$

Theorem 2

$$h(\bar{\lambda}) - h^* \leq -\bar{\lambda}\{r_0^{\bar{\lambda}}(x_0) - \Delta\}$$

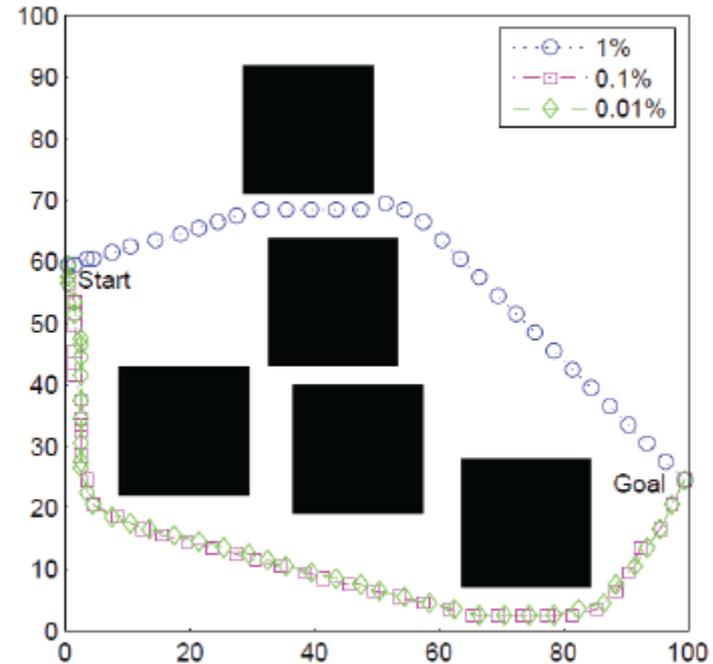
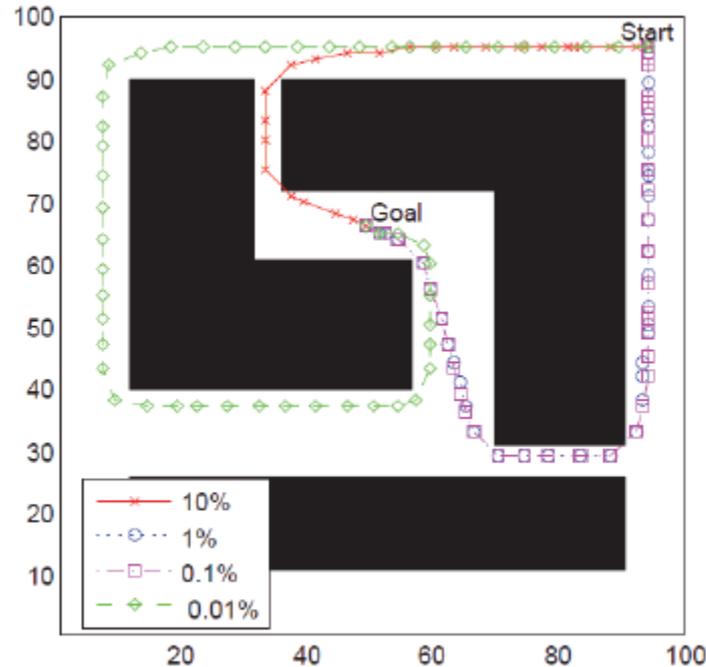


Application 1: Path Planning



Darts Lab

Mobility & Robotic Systems - 347



Dynamics

$$x_{k+1} = x_k + u_k + w_k$$

$$\|u_k\|_2 \leq d_k, \quad w_k \sim \mathcal{N}(0, \Sigma_k)$$

State space: 100 x 100 grid

Cost: expected path length from start to goal

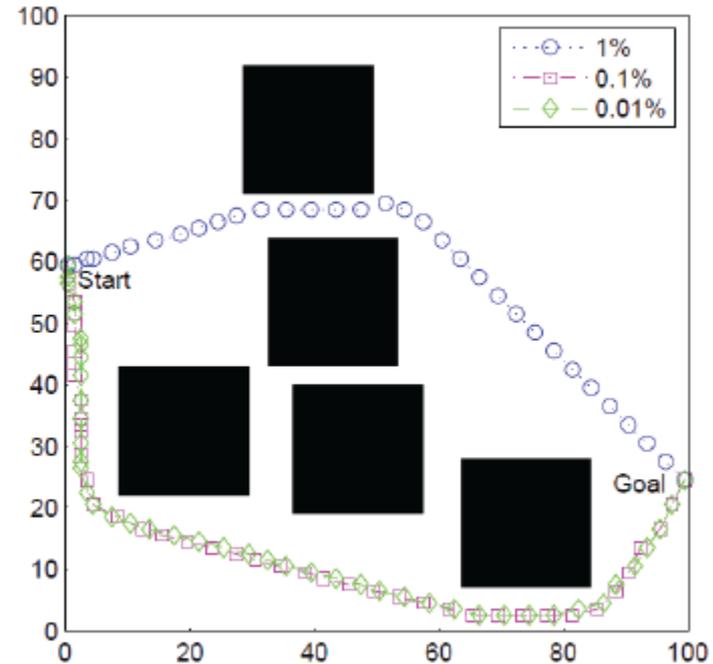
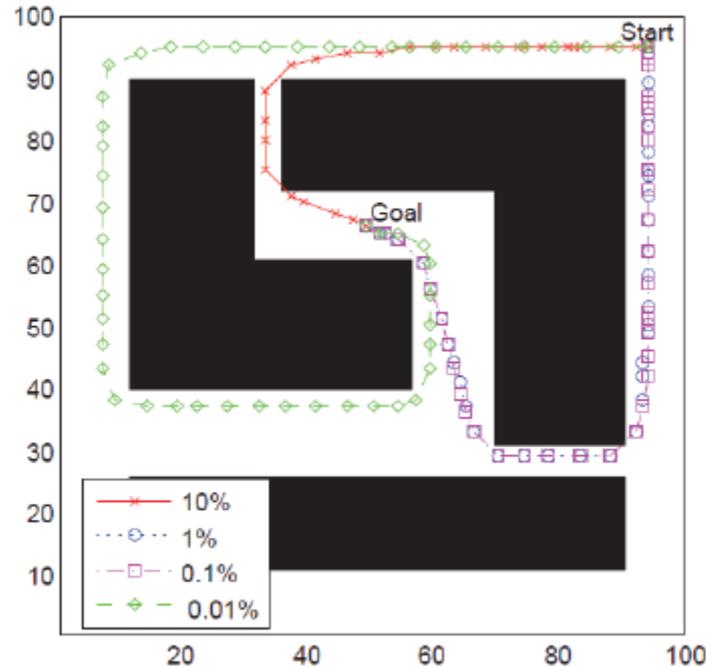


Application 1: Path Planning



Darts Lab

Mobility & Robotic Systems - 347



Risk bound	Cost	Length of nominal path	Computation time [sec]
$\Delta = 1\%$	0.88551 ± 0.0176	82.58 ± 25.10	17.2 ± 5.4
$\Delta = 0.1\%$	0.88555 ± 0.0175	86.74 ± 27.94	14.9 ± 4.6
$\Delta = 0.01\%$	0.88556 ± 0.0175	87.21 ± 28.11	12.2 ± 4.7

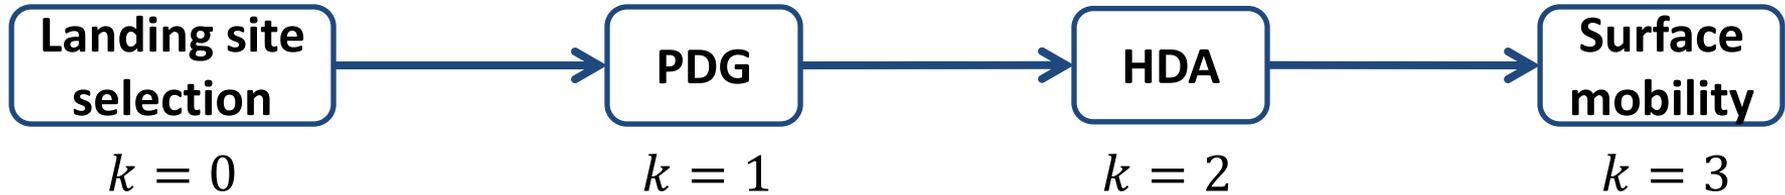


Application 2: Mars Entry, Descent & Landing

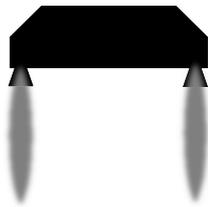


Darts Lab

Mobility & Robotic Systems - 347



Powered descent guidance



Hazard detection & avoidance



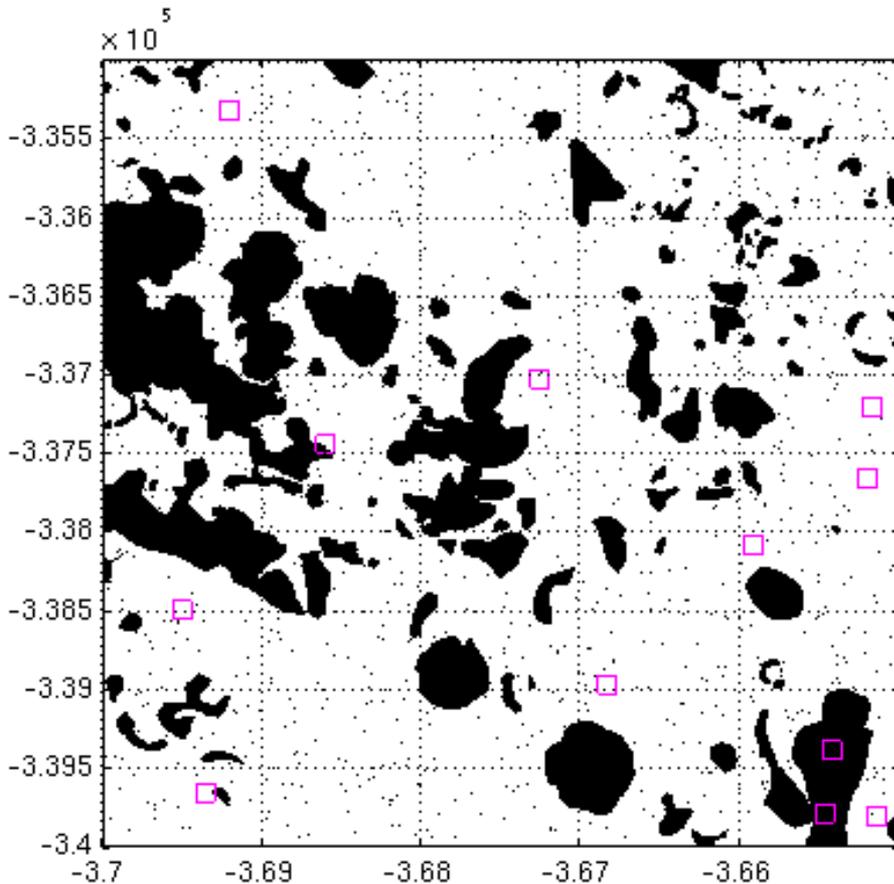
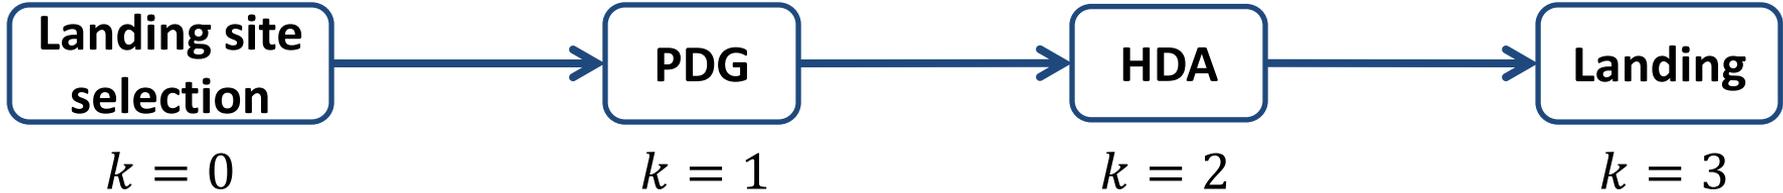


Mars EDL Problem Formulation



Darts Lab

Mobility & Robotic Systems - 347



Dynamics

$$x_{k+1} = x_k + u_k + w_k$$

$$\|u_k\|_2 \leq d_k, \quad w_k \sim \mathcal{N}(0, \Sigma_k)$$

State space: 2000 x 2000 grid

Cost: expected driving distance to visit N science target after landing

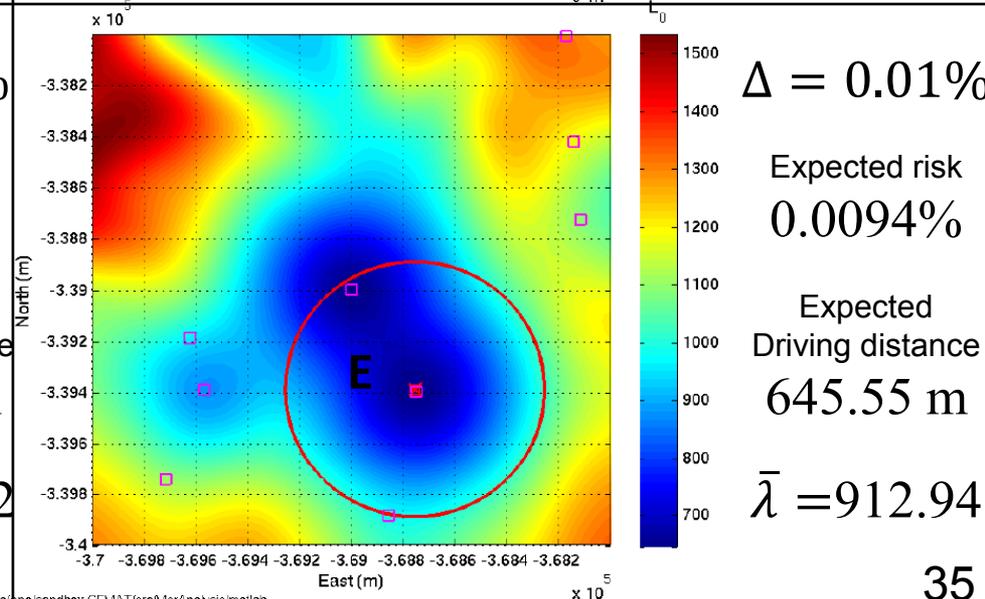
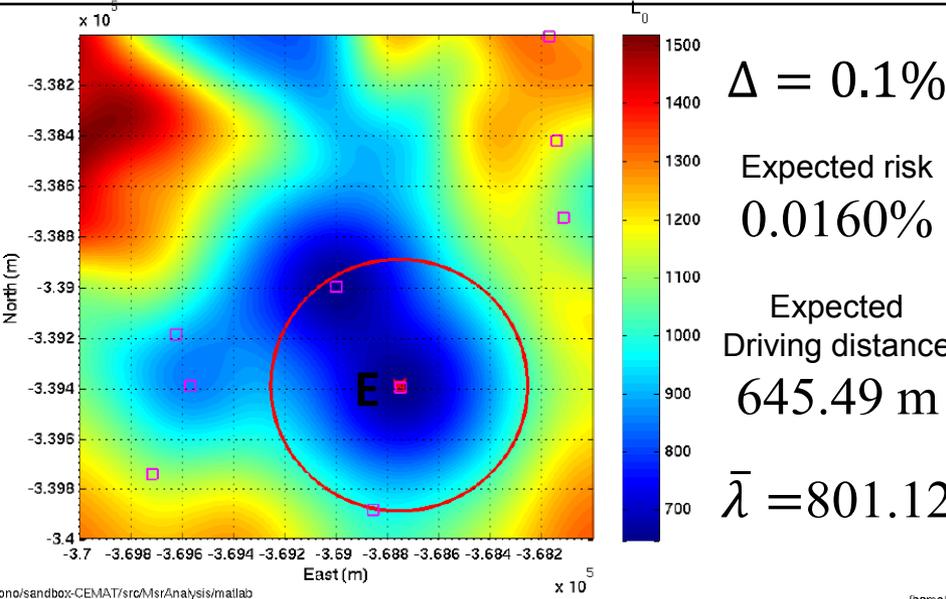
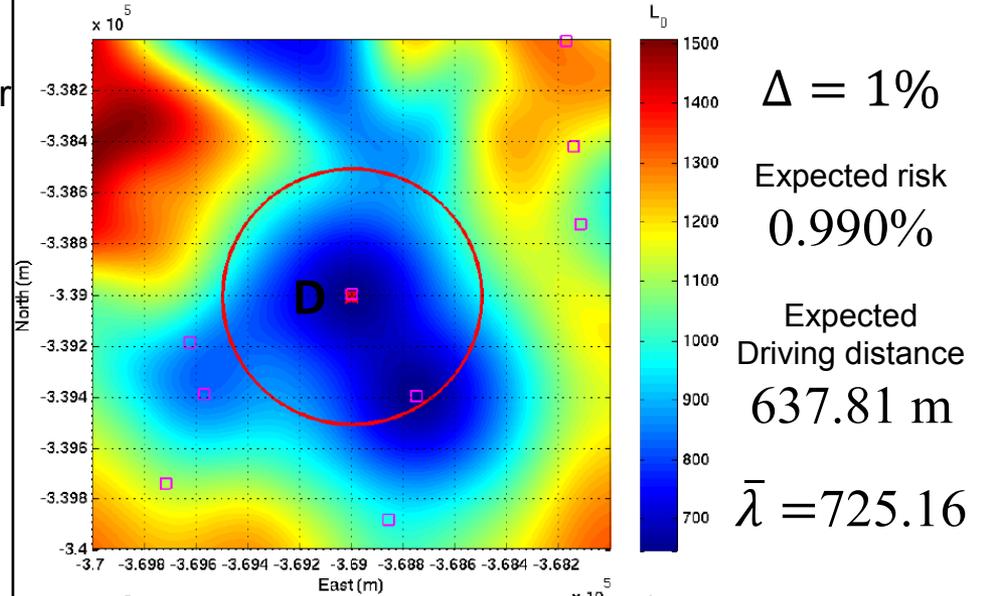
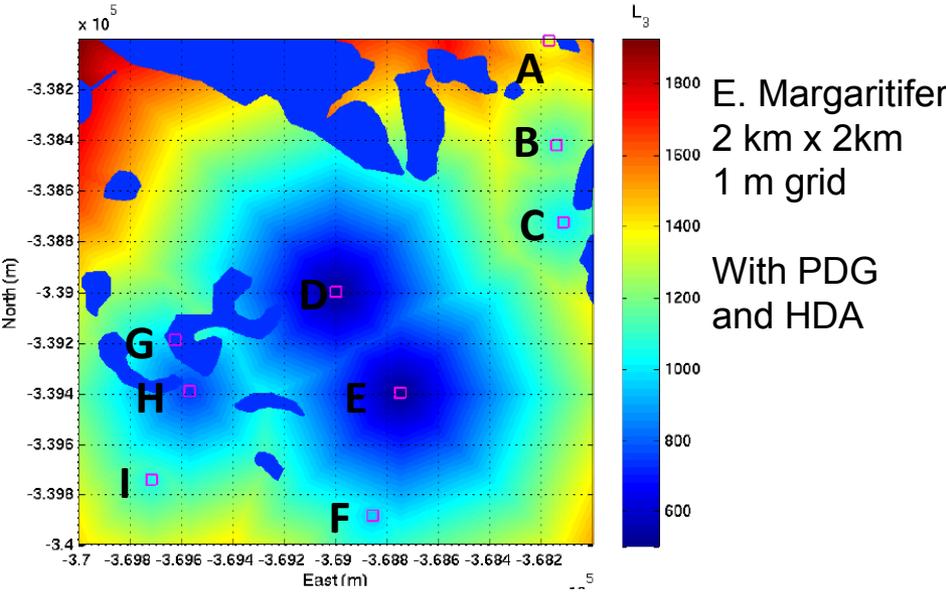


Simulation on E. Margaritifer



Darts Lab

Mobility & Robotic Systems - 347





Bound on Solution Error

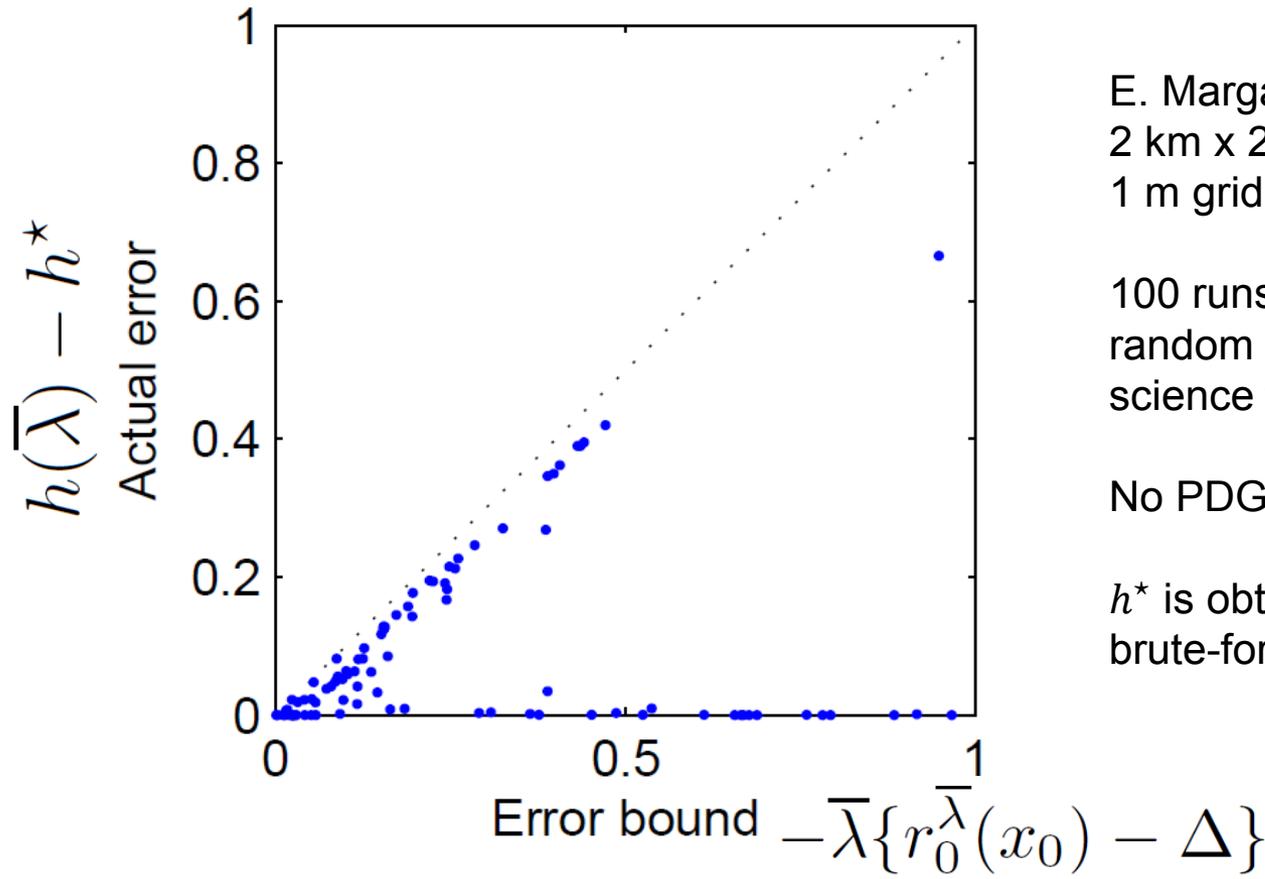


Darts Lab

Mobility & Robotic Systems - 347

Theorem 2

$$h(\bar{\lambda}) - h^* \leq -\bar{\lambda}\{r_0^{\bar{\lambda}}(x_0) - \Delta\}$$



E. Margaritifer
2 km x 2km
1 m grid

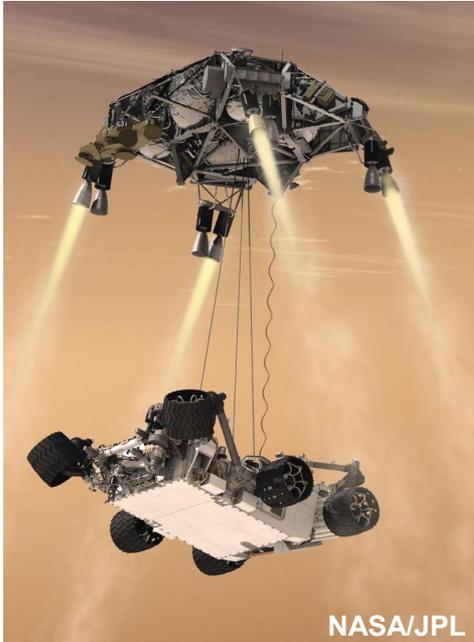
100 runs with
random location of
science targets

No PDG, no HAD

h^* is obtained by a
brute-force method



Conclusion



- Developed a CCDP algorithm that can
 - Make optimal sequential decisions under uncertainty
 - Minimize expected cost within a user-specified risk bound
 - Provide bounds on the solution error
- Demonstrated the CCDP algorithm on
 - Path planning problem
 - Future Mars entry, descent, and landing scenario



Questions



Darts Lab

Mobility & Robotic Systems - 347

