KU LEUVEN

Data-driven distributionally robust LQR with multiplicative noise

Peter Coppens, Mathijs Schuurmans, Panagiotis Patrinos

KU Leuven, ESAT–STADIUS Center for Dynamical Systems, Signal Processing, and Data Analytics

June 11, 2020

Introduction

Control in uncertain environments

model = reality



Control in uncertain environments

model $=$ reality



model $\neq$ reality



Control in uncertain environments

model $=$ reality



model $\neq$ reality



conservative



Control in uncertain environments

model $=$ reality



generalize

model $\neq$ reality



conservative



Control in uncertain environments

model $=$ reality



generalize

model $\neq$ reality



safety

conservative



Control in uncertain environments

model $=$ reality



generalize

model $\neq$ reality



safety

conservative



performance

Uncertain systems

financial



(bio-)chemical



climate control



power distribution



robotics/autonomous vehicles

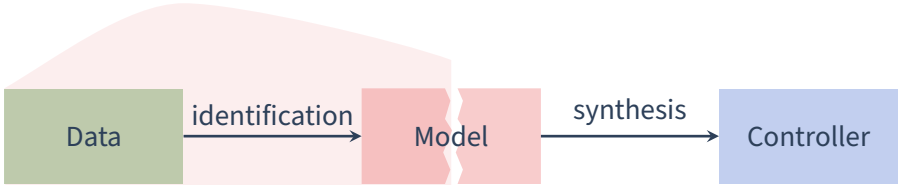


aerospace



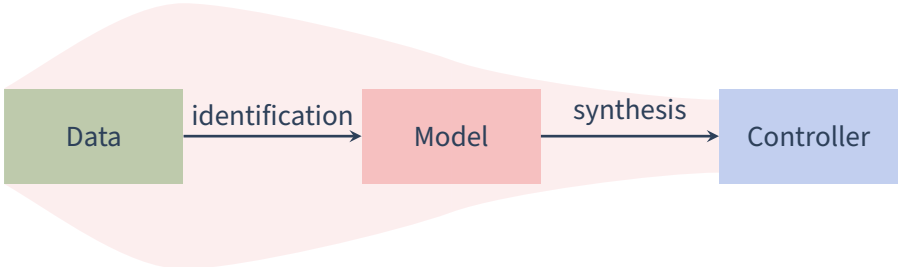
Learning to generalize

classical control identification and controller design are kept seperate.



Learning to generalize

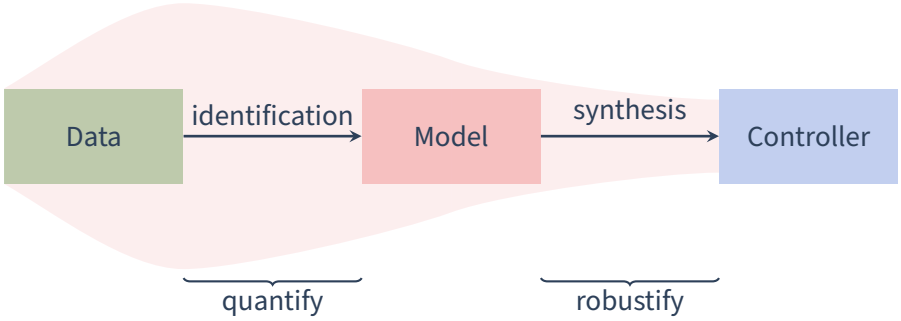
coarse id control explicitly include ambiguity in the model throughout the design[1].



[1] Benjamin Recht. "A Tour of Reinforcement Learning: The View from Continuous Control". In: *Annual Review of Control, Robotics, and Autonomous Systems* 2.1 (2019), pp. 253–279

Learning to generalize

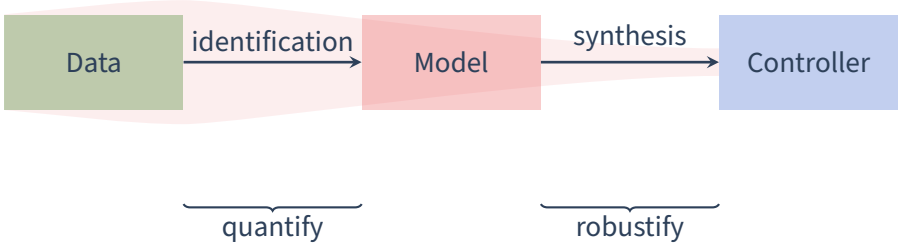
coarse id control explicitly include ambiguity in the model throughout the design[1].



[1] Recht, “A Tour of Reinforcement Learning: The View from Continuous Control”

Learning to generalize

coarse id control explicitly include ambiguity in the model throughout the design[1].



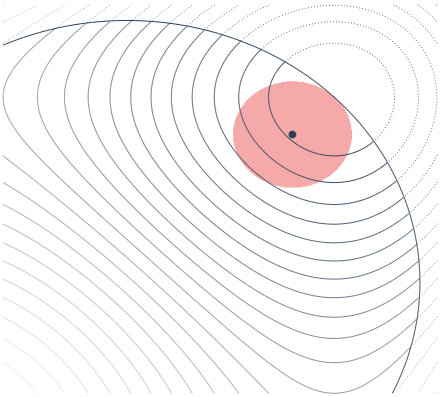
[1] Recht, “A Tour of Reinforcement Learning: The View from Continuous Control”

Robustification in Optimization

Robustify — Stochastic and robust optimization

- Robust optimization ($w \in \mathcal{W}$)

$$\begin{aligned} \min_{x} \max_{w \in \mathcal{W}} \quad & f(x, w) \\ \text{s.t.} \quad & g(x, w) \leq 0, \ \forall w \in \mathcal{W} \end{aligned}$$



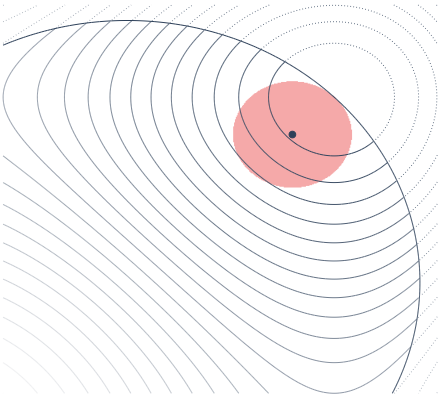
[1] Rahimian et al. 2019

Robustify — Stochastic and robust optimization

- Robust optimization ($w \in \mathcal{W}$)

$$\begin{aligned} \min_{x} \max_{w \in \mathcal{W}} \quad & f(x, w) \\ \text{s.t.} \quad & g(x, w) \leq 0, \ \forall w \in \mathcal{W} \end{aligned}$$

➢ strong guarantees, conservative



[1] Rahimian et al. 2019

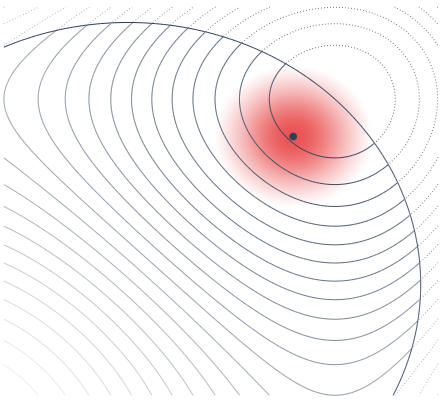
Robustify — Stochastic and robust optimization

- Robust optimization ($w \in \mathcal{W}$)

$$\begin{aligned} \min_{x} \max_{w \in \mathcal{W}} \quad & f(x, w) \\ \text{s.t.} \quad & g(x, w) \leq 0, \ \forall w \in \mathcal{W} \end{aligned}$$

- Stochastic optimization ($w \sim \mathbb{P}_w$)

$$\begin{aligned} \min_{x} \quad & \mathbb{E}\left[f(x, w)\right] \\ \text{s.t.} \quad & \mathbb{P}\left[g(x, w) \leq 0\right] \geq 1 - \delta \end{aligned}$$



[1] Rahimian et al. 2019

Robustify — Stochastic and robust optimization

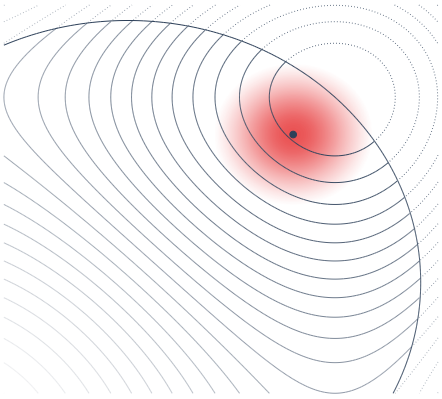
- Robust optimization ($w \in \mathcal{W}$)

$$\min_{x} \max_{w \in \mathcal{W}} \quad f(x, w)$$
$$\text{s.t.} \quad g(x, w) \leq 0, \ \forall w \in \mathcal{W}$$

- Stochastic optimization ($w \sim \mathbb{P}_w$)

$$\min_{x} \quad \mathbb{E}\left[f(x, w)\right]$$
$$\text{s.t.} \quad \mathbb{P}\left[g(x, w) \leq 0\right] \geq 1 - \delta$$

> less guarantees, optimal



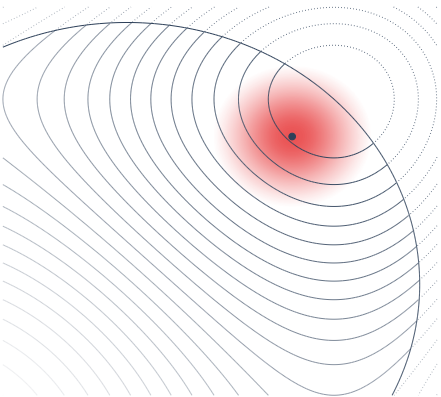
[1] Rahimian et al. 2019

Robustify — Distributionally robust optimization

- Distributionally robust ($\xi \sim \mathbb{P}_w$, $\mathbb{P}_w \in \mathcal{A}$)

$$\min_{x} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w\left[f(x, w)\right]$$

$$\text{s.t.} \qquad \mathbb{P}_w\left[g(x, w) \leq 0\right] \geq 1 - \delta, \ \forall \mathbb{P}_w \in \mathcal{A}$$



[1]Rahimian et al. 2019, [2]Abadeh et al. 2015

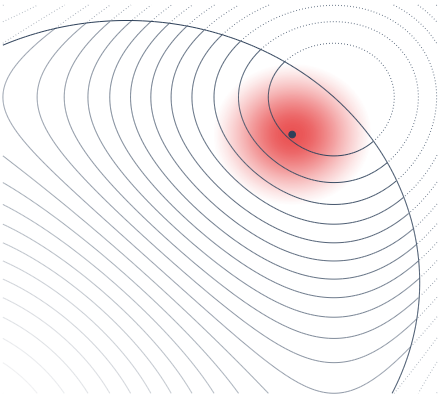
Robustify — Distributionally robust optimization

- Distributionally robust ($\xi \sim \mathbb{P}_w$, $\mathbb{P}_w \in \mathcal{A}$)

$$\min_{x} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w\left[f(x, w)\right]$$

$$\text{s.t.} \qquad \mathbb{P}_w\left[g(x, w) \leq 0\right] \geq 1 - \delta, \ \forall \mathbb{P}_w \in \mathcal{A}$$

  - ➢ conservatism decreases with data



[1]Rahimian et al. 2019, [2]Abadeh et al. 2015

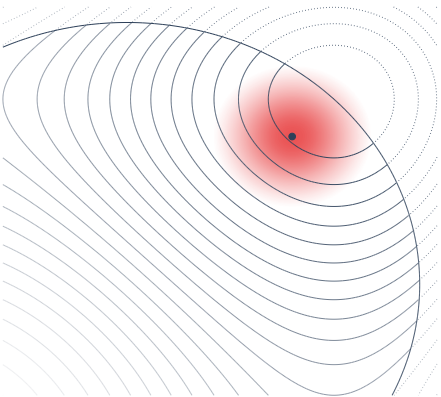
Robustify — Distributionally robust optimization

- Distributionally robust ($\xi \sim \mathbb{P}_w, \mathbb{P}_w \in \mathcal{A}$)

$$\min_{x} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w\left[f(x, w)\right]$$

$$\text{s.t.} \qquad \mathbb{P}_w\left[g(x, w) \leq 0\right] \geq 1 - \delta, \ \forall \mathbb{P}_w \in \mathcal{A}$$

  - conservatism decreases with data
  - akin to automatic regularization[2]



[1]Rahimian et al. 2019, [2]Abadeh et al. 2015

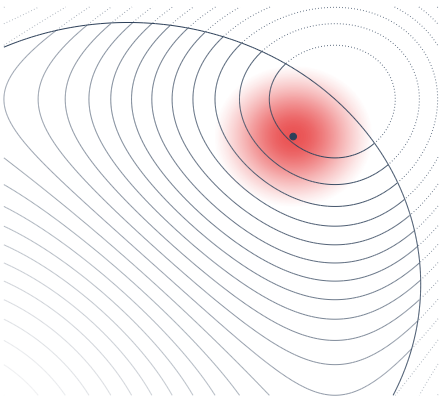
Robustify — Distributionally robust optimization

- Distributionally robust ($\xi \sim \mathbb{P}_w$, $\mathbb{P}_w \in \mathcal{A}$)

$$\min_{x} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w\left[f(x, w)\right]$$

$$\text{s.t.} \qquad \mathbb{P}_w\left[g(x, w) \leq 0\right] \geq 1 - \delta, \ \forall \mathbb{P}_w \in \mathcal{A}$$

  - conservatism decreases with data
  - akin to automatic regularization[2]
  - already applied in control before[3-5]



[1]Rahimian et al. 2019, [2]Abadeh et al. 2015, [3]Schuurmans et al. 2019, [4]Abadeh et al. 2018, [5]Yang 2018

Distributionally robust LQR

DR-LQR — The nominal problem

$$\underset{u_0,u_1,\dots}{\textbf{minimize}} \quad \mathbb{E}\left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k\right]$$

$$\textbf{subj. to} \quad x_{k+1} = \left(A_0 + \sum_{i=1}^{n_w} A_i w_k^{(i)}\right) x_k + \left(B_0 + \sum_{i=0}^{n_w} B_i w_k^{(i)}\right) u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

DR-LQR — The nominal problem

$$\underset{u_0,u_1,\dots}{\textbf{minimize}} \quad \mathbb{E}\left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k\right]$$

$$\textbf{subj. to} \quad x_{k+1} = \left(A_0 + \sum_{i=1}^{n_w} A_i w_k^{(i)}\right) x_k + \left(B_0 + \sum_{i=0}^{n_w} B_i w_k^{(i)}\right) u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

w_k i.i.d. copies of $w \sim \mathbb{P}_w^\star$, with $\mathbb{E}[w] = \mu$ and $\mathbb{E}[(w-\mu)(w-\mu)^\top] = \Sigma$

DR-LQR — The nominal problem

$$\begin{aligned} \underset{u_0,u_1,\ldots}{\textbf{minimize}} \quad & \mathbb{E}\left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k\right] \\ \textbf{subj. to} \quad & x_{k+1} = \left(A_0 + \textstyle\sum_{i=1}^{n_w} A_i w_k^{(i)}\right) x_k + \left(B_0 + \textstyle\sum_{i=0}^{n_w} B_i w_k^{(i)}\right) u_k, \quad \forall k \in \mathbb{N} \\ & x_0 \text{ given} \end{aligned}$$

Bellman equation: solution is still quadratic

$$V(x) = \inf_u \{\ell(x,u) + \mathbb{E}\left[V(x_+) \mid (x,u)\right]\}$$

DR-LQR — The nominal problem

$$\mathbb{E}\left[x_+^\top P x_+ \mid (x,u)\right] = (\mathbf{A}_0 x + \mathbf{B}_0 u)^\top \left(\begin{bmatrix} 1 & \mu^\top \\ \mu & \Sigma + \mu\mu^\top \end{bmatrix} \otimes P \right) (\mathbf{A}_0 x + \mathbf{B}_0 u)\,,$$

where $\mathbf{A}_0 = \begin{bmatrix} A_0^\top & A_1^\top & \dots & A_{n_w}^\top \end{bmatrix}^\top$ and $\mathbf{B}_0 = \begin{bmatrix} B_0^\top & B_1^\top & \dots & B_{n_w}^\top \end{bmatrix}^\top$

DR-LQR — The nominal problem

$$\underset{u_0,u_1,\dots}{\textbf{minimize}} \quad \mathbb{E}\left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k\right]$$

$$\textbf{subj. to} \quad x_{k+1} = \left(A_0 + \textstyle\sum_{i=1}^{n_w} A_i w_k^{(i)}\right) x_k + \left(B_0 + \textstyle\sum_{i=0}^{n_w} B_i w_k^{(i)}\right) u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: solution is still quadratic

$$V(x) = \inf_u \{\ell(x,u) + \mathbb{E}\left[V(x_+) \mid (x,u)\right]\}$$

➢ generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

DR-LQR — The nominal problem

properties of optimal controller:

- linear $u = Kx$ (solve SDP)

> generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

DR-LQR — The nominal problem

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing

$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

> generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

DR-LQR — The nominal problem

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing

  $$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$
- **performant** true optimum

> generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

DR-LQR — The nominal problem

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing
  $$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$
- **performant** true optimum
- only depends on μ and Σ

> generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

DR-LQR — The nominal problem

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing

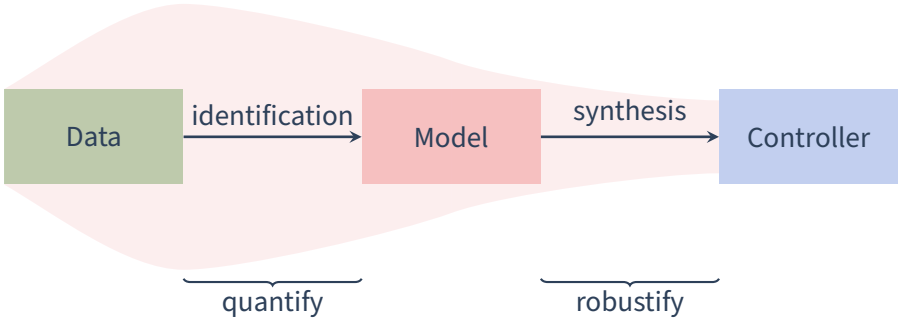
  $$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

- **performant** true optimum
- only depends on μ and Σ ➢ **what if these are uncertain?**

➢ generalized Riccati equation[1-3]

[1] Gravell et al. 2019, [2] Wonham 1967, [3] Costa et al. 1997

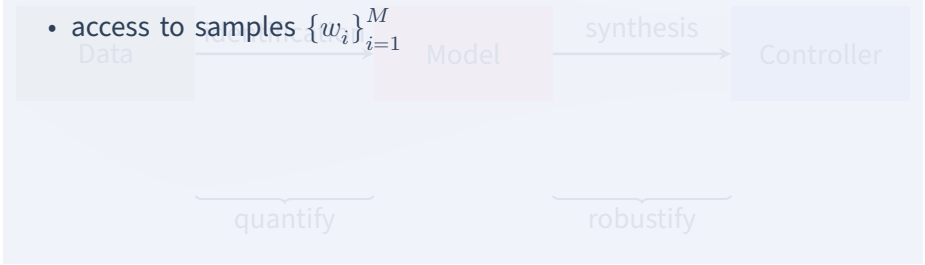
DR-LQR — Learning to generalize



DR-LQR — Learning to generalize

assumptions about the data:

- access to samples $\{w_i\}_{i=1}^{M}$



DR-LQR — Learning to generalize

assumptions about the data:

- access to samples $\{w_i\}_{i=1}^{M}$
- $\Sigma^{-1/2}(w - \mu)$ is sub-Gaussian[1]

[1] Delage et al. 2010

DR-LQR — Learning to generalize

assumptions about the data:

- access to samples $\{w_i\}_{i=1}^{M}$
- $\Sigma^{-1/2}(w - \mu)$ is sub-Gaussian[1]
- $\hat{\mu} = \frac{1}{M}\sum_{i=1}^{M} w_i$
- $\hat{\Sigma} = \frac{1}{M}\sum_{i=1}^{M}(w_i - \hat{\mu})(w_i - \hat{\mu})^\top$

[1] Delage et al. 2010

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \widehat{\mu})^\top \widehat{\Sigma}^{-1} (\mu - \widehat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \widehat{\Sigma} \end{array} \right\}$$

[1] Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \widehat{\mu})^\top \widehat{\Sigma}^{-1} (\mu - \widehat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \widehat{\Sigma} \end{array} \right\}$$

[1] Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \widehat{\mu})^\top \widehat{\Sigma}^{-1} (\mu - \widehat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \widehat{\Sigma} \end{array} \right\}$$

[1]Delage et al. 2010, [2]Hsu et al. 2012b, [3]Hsu et al. 2012a

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \widehat{\mu})^\top \widehat{\Sigma}^{-1} (\mu - \widehat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \widehat{\Sigma} \end{array} \right\}$$

- radii decrease with $1/\sqrt{\#\text{ samples}}$

[1] Delage et al. 2010, [2] Hsu et al. 2012b, [3] Hsu et al. 2012a

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \widehat{\mu})^\top \widehat{\Sigma}^{-1} (\mu - \widehat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \widehat{\Sigma} \end{array} \right\}$$

- radii decrease with $1/\sqrt{\#\text{ samples}}$
- no knowledge of true distribution required

[1]Delage et al. 2010, [2]Hsu et al. 2012b, [3]Hsu et al. 2012a

DR-LQR — The ambiguity set

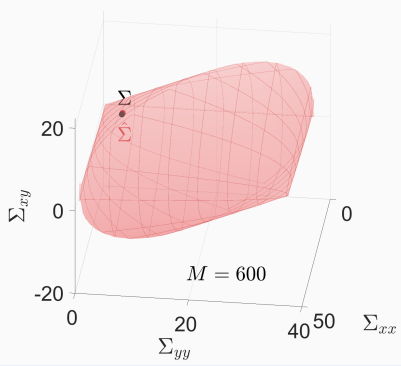
quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1 - \beta$

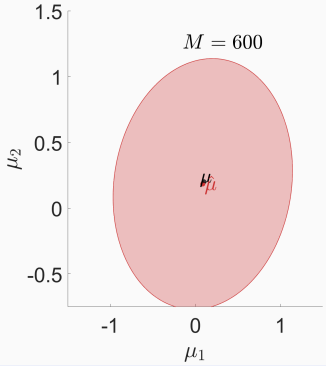
$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu - \hat{\mu})^\top \hat{\Sigma}^{-1} (\mu - \hat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta) \hat{\Sigma} \end{array} \right\}$$

- radii decrease with $1/\sqrt{\#\text{ samples}}$
- no knowledge of true distribution required
- sufficient samples required

[1]Delage et al. 2010, [2]Hsu et al. 2012b, [3]Hsu et al. 2012a

DR-LQR — The ambiguity set





[1]Delage et al. 2010, [2]Hsu et al. 2012b, [3]Hsu et al. 2012a

DR-LQR — The ambiguity set

quantization: determine r_μ and r_Σ such that $\mathbb{P}\left[\mathbb{P}_w^\star \in \mathcal{A}\right] \geq 1-\beta$

$$\mathcal{A} := \left\{ \mathbb{P}_w \;\middle|\; \begin{array}{l} (\mu-\hat{\mu})^\top \hat{\Sigma}^{-1}(\mu-\hat{\mu}) \leq r_\mu(\beta) \\ \Sigma \preceq r_\Sigma(\beta)\hat{\Sigma} \end{array} \right\}$$

synthesis: find the distributionally robust controller

$$\begin{array}{ll} \underset{u_0,u_1,\ldots}{\textbf{min}} \; \underset{\mathbb{P}_w \in \mathcal{A}}{\textbf{max}} & \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right] \\ \textbf{subj. to} & x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N} \\ & x_0 \text{ given} \end{array}$$

DR-LQR — Uncertain covariance

$$\min_{u_0,u_1,\cdots} \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \qquad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

DR-LQR — Uncertain covariance

$$\min_{u_0,u_1,\dots} \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\text{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

DR-LQR — Uncertain covariance

$$\min_{u_0,u_1,\cdots} \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \qquad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: solution is still quadratic

$$V(x) = \inf_u \left\{ \ell(x,u) + \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \mathbb{E}_w \left[V(x_+) \mid (x,u) \right] \right\}$$

DR-LQR — Uncertain covariance

$$\min_{u_0,u_1,\cdots} \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\text{subj. to} \qquad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: solution is still quadratic

$$V(x) = \inf_u \left\{ \ell(x,u) + \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \mathbb{E}_w \left[V(x_+) \mid (x,u) \right] \right\}$$

$$\max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \mathbb{E}_w \left[x_+^\top P x_+ \mid (x,u) \right] = (\mathbf{A}_0 x + \mathbf{B}_0 u)^\top \left(\begin{bmatrix} 1 & \mu^\top \\ \mu & r_\Sigma \hat{\Sigma} + \mu\mu^\top \end{bmatrix} \otimes P \right) (\mathbf{A}_0 x + \mathbf{B}_0 u)$$

DR-LQR — Uncertain covariance

$$\begin{aligned} \min_{u_0,u_1,\cdots} \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \quad & \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right] \\ \text{subj. to} \quad & x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N} \\ & x_0 \text{ given} \end{aligned}$$

Bellman equation: solution is still quadratic

$$V(x) = \inf_u \left\{ \ell(x,u) + \max_{\Sigma \preceq r_\Sigma \hat{\Sigma}} \mathbb{E}_w \left[V(x_+) \mid (x,u) \right] \right\}$$

- generalized Riccati equation

DR-LQR — Uncertain covariance

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing w.h.p.

$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

> generalized Riccati equation

DR-LQR — Uncertain covariance

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing w.h.p.

$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

- **performant** converges to true optimum

> generalized Riccati equation

DR-LQR — Uncertain covariance

properties of optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing w.h.p.

$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

- **performant converges to true optimum**
- only depends on μ and $r_\Sigma \hat{\Sigma}$

> generalized Riccati equation

DR-LQR — Full uncertainty

$$\min_{u_0,u_1,\dots} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

DR-LQR — Full uncertainty

$$\min_{u_0,u_1,\ldots} \max_{\mathbb{P}_w \in \mathcal{A}} \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\text{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

DR-LQR — Full uncertainty

$$\min_{u_0,u_1,\dots} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: solution **no longer quadratic**

$$V(x) = \inf_u \left\{ \ell(x,u) + \max_{\mathbb{P}_w \in \mathcal{A}} \mathbb{E}_w \left[V(x_+) \mid (x,u) \right] \right\}$$

DR-LQR — Full uncertainty

$$\begin{aligned} \min_{u_0,u_1,\dots} \max_{\mathbb{P}_w \in \mathcal{A}} \quad & \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right] \\ \textbf{subj. to} \quad & x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N} \\ & x_0 \text{ given} \end{aligned}$$

Bellman equation: approximate solution[1]

$$\begin{aligned} \underset{P}{\textbf{minimize}} \quad & x_0^\top P x_0 \\ \textbf{subj. to} \quad & x^\top P x \geq \min_u \left\{ \ell(x,u) + \max_{\mathbb{P}_w \in \mathcal{A}} \mathbb{E}_w [x_+^\top P x_+] \right\}, \quad \forall x \end{aligned}$$

[1] Boyd et al. 1994

DR-LQR — Full uncertainty

$$\min_{u_0,u_1,\dots} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: approximate solution[1]

$$\underset{P}{\textbf{minimize}} \quad x_0^\top P x_0$$

$$\textbf{subj. to} \quad x^\top P x \geq \min_u \left\{ \ell(x,u) + \max_{\mathbb{P}_w \in \mathcal{A}} \mathbb{E}_w [x_+^\top P x_+] \right\}, \quad \forall x$$

➢ solve approximate SDP[2]

[1] Boyd et al. 1994, [2] Ben-Tal et al. 2000

DR-LQR — Full uncertainty

$$\min_{u_0,u_1,\dots} \max_{\mathbb{P}_w \in \mathcal{A}} \quad \mathbb{E}_w \left[\sum_{k=0}^{\infty} x_k^\top Q x_k + u_k^\top R u_k \right]$$

$$\textbf{subj. to} \quad x_{k+1} = A(w_k)x_k + B(w_k)u_k, \quad \forall k \in \mathbb{N}$$

$$x_0 \text{ given}$$

Bellman equation: approximate solution[1]

$$\underset{P}{\textbf{minimize}} \quad x_0^\top P x_0$$

$$\textbf{subj. to} \quad x^\top P x \geq \min_u \left\{ \ell(x,u) + \max_{\mathbb{P}_w \in \mathcal{A}} \mathbb{E}_w [x_+^\top P x_+] \right\}, \quad \forall x$$

➢ solve **approximate** SDP[2]

[1] Boyd et al. 1994, [2] Ben-Tal et al. 2000

DR-LQR — Full uncertainty

properties of approximate optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing w.h.p.

$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

> solve approximate SDP[2]

[1]Boyd et al. 1994, [2]Ben-Tal et al. 2000

DR-LQR — Full uncertainty

properties of approximate optimal controller:

- linear $u = Kx$ (solve SDP)
- **safe** mean-square stabilizing w.h.p.

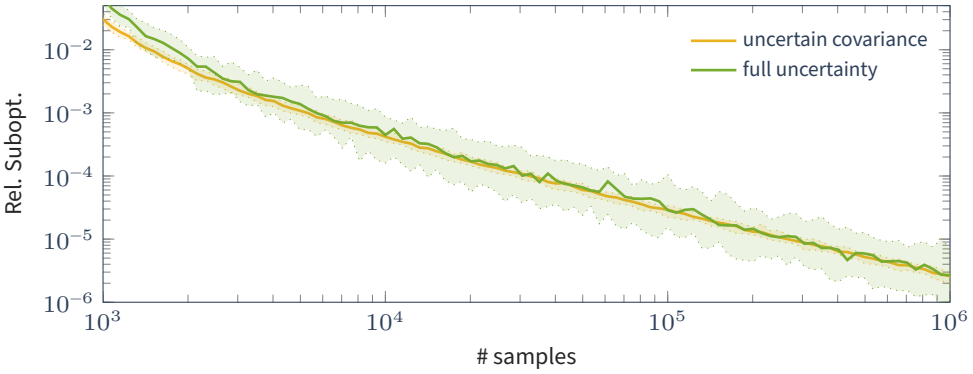
$$\mathbb{E}[x_k^\top x_k] \to 0 \text{ as } k \to \infty$$

- **performant converges to true optimum**

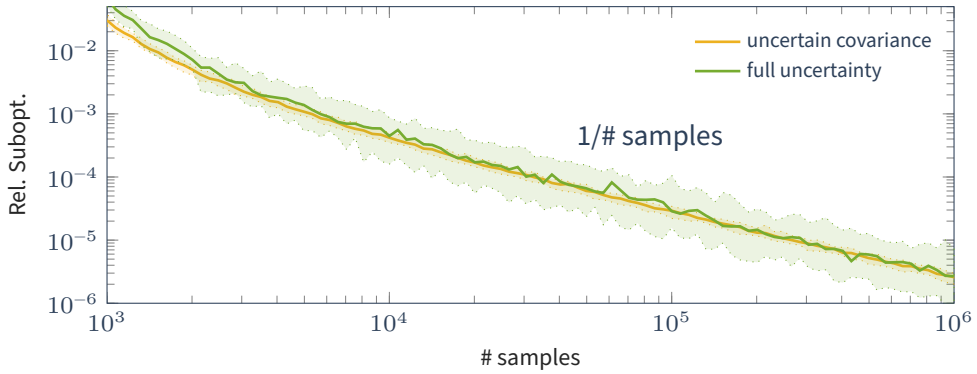
> solve approximate SDP[2]

[1]Boyd et al. 1994, [2]Ben-Tal et al. 2000

DR-LQR — Numerical result



DR-LQR — Numerical result



[1] P. Coppens and P. Patrinos. *Sample Complexity of Data-Driven Stochastic LQR with Multiplicative Uncertainty*. 2020. arXiv: 2005.12167

DR-LQR — Future work

- what to do when not enough samples?

DR-LQR — Future work

- what to do when not enough samples?
- what to do when only state measurements?

DR-LQR — Future work

- what to do when not enough samples?
- what to do when only state measurements?
- constraints, state observers, …

Thank you for your attention!

Bibliography

References i


Abadeh, Soroosh Shafieezadeh, Peyman Mohajerin Esfahani, and Daniel Kuhn. "Distributionally Robust Logistic Regression". In: *Advances in Neural Information Processing Systems 28*. 2015, pp. 1576–1584.


Abadeh, Soroosh Shafieezadeh et al. "Wasserstein distributionally robust Kalman filtering". In: *Advances in Neural Information Processing Systems*. 2018, pp. 8474–8483.


Ben-Tal, Aharon, Laurent El Ghaoui, and Arkadi Nemirovski. "Robustness". In: *Handbook of Semidefinite Programming*. Boston, MA: Springer US, 2000, pp. 139–162.


Boyd, Stephen et al. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.


Coppens, P. and P. Patrinos. *Sample Complexity of Data-Driven Stochastic LQR with Multiplicative Uncertainty*. 2020. arXiv: 2005.12167.

References ii

- Costa, Oswaldo L.V. and Carlos S. Kubrusly. “Quadratic optimal control for discrete-time infinite-dimensional stochastic bilinear systems”. In: *IMA Journal of Mathematical Control and Information* 14.4 (1997), pp. 385–399.
- Delage, Erick and Yinyu Ye. “Distributionally robust optimization under moment uncertainty with application to data-driven problems”. In: *Operations Research* 58.3 (2010), pp. 595–612.
- Gravell, Benjamin, Peyman Mohajerin Esfahani, and Tyler Summers. “Learning robust control for LQR systems with multiplicative noise via policy gradient”. In: *arXiv preprint arXiv:1905.13547* (2019). arXiv: 1905.13547.
- Hsu, Daniel, Sham M. Kakade, and Tong Zhang. “A tail inequality for quadratic forms of subgaussian random vectors”. In: *Electronic Communications in Probability* 17.52 (2012), pp. 1–6.

References iii

- Hsu, Daniel, Sham M. Kakade, and Tong Zhang. “Tail inequalities for sums of random matrices that depend on the intrinsic dimension”. In: *Electronic Communications in Probability* 17.52 (2012), pp. 1–13.
- Rahimian, H. and S. Mehrotra. *Distributionally Robust Optimization: A Review*. 2019. arXiv: 1908.05659.
- Recht, Benjamin. “A Tour of Reinforcement Learning: The View from Continuous Control”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 2.1 (2019), pp. 253–279.
- Schuurmans, Mathijs, Pantelis Sopasakis, and Panagiotis Patrinos. “Safe Learning-Based Control of Stochastic Jump Linear Systems: a Distributionally Robust Approach”. In: *2019 IEEE Conference on Decision and Control (CDC)*. IEEE. 2019, pp. 6498–6503.

References iv

- Vershynin, Roman. *High-Dimensional Probability: An Introduction with Applications in Data Science*. 2018.
- Wainwright, Martin J. *High-Dimensional Statistics*. Cambridge University Press, 2019.
- Wonham, W. Murray. “Optimal Stationary Control of a Linear System with State-Dependent Noise”. In: *SIAM Journal on Control* 5.3 (1967), pp. 486–500.
- Yang, Insoon. *Wasserstein distributionally robust stochastic control: A data-driven approach*. 2018. arXiv: 1812.09808.

Assets

- Images ‘control in uncertain environments’:
  https://flic.kr/p/<e6tYCK,939fWN,2iUoC88>
- Images ‘uncertain systems’:
  https://flic.kr/p/<k283C7,xz8h9,2hRaa8G,r3LtVQ,azfYAW,n9mn42>
- Matlab code available on request

Appendix

Quantification procedure

Quantify – Sub-Gaussian

Definition (Sub-Gaussian)

Random variable X is sub-Gaussian[1] if, for some $\sigma \geq 0$,

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] \leq e^{\sigma^2\lambda^2/2}, \qquad \forall \lambda \in \mathbb{R}$$

examples[2]:

- Gaussian random variables $\sigma^2 = \mathbb{E}[(X - \mathbb{E}[X])^2]$,
- Bernoulli random variable $\sigma = 1/\sqrt{\log(2)}$,
- random variables with bounded support $\sigma = \text{ess sup}\,|X|/\sqrt{\log(2)}$.

[1] Wainwright 2019, [2] Vershynin 2018

Quantify – Sub-Gaussian

Definition (Sub-Gaussian)

Random variable X is sub-Gaussian[1] if, for some $\sigma \geq 0$,

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] \leq e^{\sigma^2\lambda^2/2}, \qquad \forall \lambda \in \mathbb{R}$$

examples[2]:

- Gaussian random variables $\sigma^2 = \mathbb{E}[(X - \mathbb{E}[X])^2]$,
- Bernoulli random variable $\sigma = 1/\sqrt{\log(2)}$,
- random variables with bounded support $\sigma = \text{ess sup}\,|X|/\sqrt{\log(2)}$.

➢ **Sub-Gaussian vector** Z: for every $u \in \mathcal{S}_d$ we have $u^\top Z$ sub-Gaussian.

[1] Wainwright 2019, [2] Vershynin 2018

Quantify – Sub-Gaussian

Definition (Sub-Gaussian)

Random variable X is sub-Gaussian[1] if, for some $\sigma \geq 0$,

$$\mathbb{E}[e^{\lambda(X-\mathbb{E}[X])}] \leq e^{\sigma^2\lambda^2/2}, \qquad \forall \lambda \in \mathbb{R}$$

examples[2]:

- Gaussian random variables $\sigma^2 = \mathbb{E}[(X - \mathbb{E}[X])^2]$,
- Bernoulli random variable $\sigma = 1/\sqrt{\log(2)}$,
- random variables with bounded support $\sigma = \text{ess sup}\,|X|/\sqrt{\log(2)}$.

➢ **Sub-Exponential**: $|\lambda| < 1/\alpha$.

[1] Wainwright 2019, [2] Vershynin 2018

Quantify — Sum of random vectors/matrices

$$\mathcal{A} = \left\{ \mathbb{P}_w \mid (\hat{\mu} - \mu)\hat{\Sigma}^{-1}(\hat{\mu} - \mu) \leq r_\mu,\ \Sigma \preceq r_\Sigma \hat{\Sigma} \right\},$$

- first consider isotropic case: $\mathbb{E}[w_I] = 0$, $\mathbb{E}[w_I w_I^\top] = I$.

Quantify — Sum of random vectors/matrices

mean uncertainty: $w_I \in \mathbb{R}^d \sim \mathsf{subG}(\sigma) \quad \succ \quad \hat{\mu}_I \sim \mathsf{subG}_d(\sigma/\sqrt{M})$.

$$\mathbb{P}\left[\|\hat{\mu}_I\|_2 \leq \frac{\sigma^2}{M}\left(n_w + 2\sqrt{n_w \log(1/\beta)} + 2\log(1/\beta)\right)\right] \geq 1 - \beta$$

> first consider isotropic case: $\mathbb{E}[w_I] = 0, \mathbb{E}[w_I w_I^\top] = I$

Daniel Hsu, Sham M. Kakade, and Tong Zhang. "A tail inequality for quadratic forms of subgaussian random vectors". In: *Electronic Communications in Probability* 17.52 (2012), pp. 1–6

Quantify — Sum of random vectors/matrices

covariance uncertainty: $w_I \in \mathbb{R}^d \sim \mathsf{subG}(\sigma)$.

$$\mathbb{P}\left[\|\hat{\Sigma}_I - I\|_2 \leq \frac{\sigma^2}{1-2\epsilon}\left(\sqrt{\frac{32q(\beta,\epsilon,n_w)}{M}} + \frac{2q(\beta,\epsilon,n_w)}{M},\right)\right] \geq 1-\beta,$$

where $q(\beta,\epsilon,n_w) = n_w \log(1+1/\epsilon) + \log(2/\beta)$.

Daniel Hsu, Sham M. Kakade, and Tong Zhang. "Tail inequalities for sums of random matrices that depend on the intrinsic dimension". In: *Electronic Communications in Probability* 17.52 (2012), pp. 1–13

Quantify — Final ambiguity set

$$\mathcal{A} = \left\{ \mathbb{P}_w \mid (\hat{\mu} - \mu)\hat{\Sigma}^{-1}(\hat{\mu} - \mu) \leq r_\mu,\ \Sigma \preceq r_\Sigma \hat{\Sigma} \right\},$$

- first consider isotropic case: $\|\hat{\mu}_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

- first consider isotropic case: $\|\hat{\mu}_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$

Erick Delage and Yinyu Ye. “Distributionally robust optimization under moment uncertainty with application to data-driven problems”. In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma$$

- first consider isotropic case: $\|\hat{\mu}\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$

Erick Delage and Yinyu Ye. “Distributionally robust optimization under moment uncertainty with application to data-driven problems”. In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$(1 - t_\Sigma)I \preceq \hat{\Sigma}_I \preceq (1 + t_\Sigma)I$$

Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$(1 - t_\Sigma) I \preceq \hat{\Sigma}_I$$

- first consider isotropic case: $\|\mu_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$

Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$I \preceq \frac{1}{(1 - t_\Sigma)} \hat{\Sigma}_I$$

- first consider isotropic case: $\|\hat{\mu}_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$

Erick Delage and Yinyu Ye. “Distributionally robust optimization under moment uncertainty with application to data-driven problems”. In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$\Sigma^{1/2} I \Sigma^{1/2} \preceq \frac{1}{(1 - t_\Sigma)} \Sigma^{1/2} \hat{\Sigma}_I \Sigma^{1/2}$$

Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$\Sigma^{1/2} I \Sigma^{1/2} \preceq \frac{1}{(1 - t_\Sigma)} \Sigma^{1/2} \hat{\Sigma}_I \Sigma^{1/2}$$

- $w_I = \Sigma^{-1/2} w$

Erick Delage and Yinyu Ye. "Distributionally robust optimization under moment uncertainty with application to data-driven problems". In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

illustration: for $\mu = 0$.

$$\Sigma \preceq \frac{1}{(1 - t_\Sigma)} \hat{\Sigma}$$

- first consider isotropic case: $\|\mu_I\|_2 \leq$ and $\|\Sigma_I - I\|_2 \leq t_\Sigma(\beta)$

Erick Delage and Yinyu Ye. “Distributionally robust optimization under moment uncertainty with application to data-driven problems”. In: *Operations Research* 58.3 (2010), pp. 595–612

Quantify — Final ambiguity set

$$\mathcal{A} = \left\{ \mathbb{P}_w \mid (\hat{\mu} - \mu)\hat{\Sigma}^{-1}(\hat{\mu} - \mu) \leq r_\mu,\ \Sigma \preceq r_\Sigma \hat{\Sigma} \right\},$$

- first consider isotropic case: $\|\hat{\mu}_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$
- bound:

$$r_\Sigma(\beta) = 1/(1 - t_\Sigma(\beta/2) - t_\mu(\beta/2)) \quad \text{and} \quad r_\mu(\beta) = t_\mu(\beta/2)/(1 - t_\Sigma(\beta/2) - t_\mu(\beta/2)).$$

Quantify — Final ambiguity set

$$\mathcal{A} = \left\{ \mathbb{P}_w \mid (\hat{\mu} - \mu)\hat{\Sigma}^{-1}(\hat{\mu} - \mu) \leq r_\mu,\ \Sigma \preceq r_\Sigma \hat{\Sigma} \right\},$$

- first consider isotropic case: $\|\hat{\mu}_I\|_2 \leq t_\mu(\beta)$ and $\|\hat{\Sigma}_I - I\|_2 \leq t_\Sigma(\beta)$
- bound:

$$r_\Sigma(\beta) = 1/(1 - t_\Sigma(\beta/2) - t_\mu(\beta/2)) \quad \text{and} \quad r_\mu(\beta) = t_\mu(\beta/2)/(1 - t_\Sigma(\beta/2) - t_\mu(\beta/2)).$$

- only valid after M sufficiently large $(t_\Sigma(\beta/2) + t_\mu(\beta/2) < 1)$!

Semidefinite programs

SDPs — Nominal problem

solution exists when system is mean square stabilizable.

- **Riccati equation:**

$$P = Q + F(P) - H(P)^\top (R + G(P))^{-1} H(P)$$

- **optimal controller:**

$$K = -(R + G(P))^{-1} H(P)$$

$F(P) := \mathbf{A}_0^\top (\Sigma_0 \otimes P) \mathbf{A}_0$, $G(P) := \mathbf{B}_0^\top (\Sigma_0 \otimes P) \mathbf{B}_0$, $H(P) := \mathbf{B}_0^\top (\Sigma_0 \otimes P) \mathbf{A}_0$ and $\Sigma_0 = \begin{bmatrix} 1 & \mu^\top \\ \mu & \Sigma + \mu\mu^\top \end{bmatrix}$

SDPs — Nominal problem

solution exists when system is mean square stabilizable.

Semidefinite program:

$$\underset{P}{\text{minimize}} \quad -\mathbf{Tr}\, P$$

$$\text{subj. to} \quad \begin{bmatrix} Q - P + F(P) & H(P)^\top \\ H(P) & R + G(P) \end{bmatrix} \succeq 0,$$

$$P \succeq 0.$$

- **Riccati equation:**
- **optimal controller:**

$F(P) := \mathbf{A}_0^\top(\Sigma_0 \otimes P)\mathbf{A}_0$, $G(P) := \mathbf{B}_0^\top(\Sigma_0 \otimes P)\mathbf{B}_0$, $H(P) := \mathbf{B}_0^\top(\Sigma_0 \otimes P)\mathbf{A}_0$ and $\Sigma_0 = \begin{bmatrix} 1 & \mu^\top \\ \mu & \Sigma + \mu\mu^\top \end{bmatrix}$

SDPs — Unknown covariance

solution

Semidefinite program: only change covariance

- Riccati equation:
- optimal controller:

$$\begin{aligned} \underset{P}{\text{minimize}} \quad & -\mathbf{Tr}\, P \\ \text{subj. to} \quad & \begin{bmatrix} Q - P + F(P) & H(P)^\top \\ H(P) & R + G(P) \end{bmatrix} \succeq 0, \\ & P \succeq 0. \end{aligned}$$

$F(P) := \mathbf{A}_0^\top(\Sigma_0 \otimes P)\mathbf{A}_0$, $G(P) := \mathbf{B}_0^\top(\Sigma_0 \otimes P)\mathbf{B}_0$, $H(P) := \mathbf{B}_0^\top(\Sigma_0 \otimes P)\mathbf{A}_0$ and $\Sigma_0 = \begin{bmatrix} 1 & \mu^\top \\ \mu & r_\Sigma \hat{\Sigma} + \mu\mu^\top \end{bmatrix}$

SDPs — Full uncertainty

$$\begin{aligned}
\underset{W,V,S,L}{\textbf{maximize}} \quad & \mathbf{Tr}\, W \\
\textbf{subj. to} \quad & \begin{bmatrix} S & r_\mu H_1^\top & r_\mu H_2^\top & \cdots & r_\mu H_{n_w}^\top \\ r_\mu H_1 & L & & & \\ r_\mu H_2 & & L & & \\ \vdots & & & \ddots & \\ r_\mu H_{n_w} & & & & L \end{bmatrix} \succeq 0, \\
& \begin{bmatrix} W - \sqrt{2}S & (\mathbf{A}W + \mathbf{B}V)^\top & (\hat{A}W + \hat{B}V)^\top & W^\top Q^{\frac{1}{2}} & V^\top R^{\frac{1}{2}} \\ \mathbf{A}W + \mathbf{B}V & (r_\Sigma \hat{\Sigma})^{-1} \otimes W & & & \\ \hat{A}W + \hat{B}V & & W - \sqrt{2}L & & \\ Q^{\frac{1}{2}} W & & & I_{n_x} & \\ R^{\frac{1}{2}} V & & & & I_{n_u} \end{bmatrix} \succeq 0,
\end{aligned}$$

$\hat{A} = A_0 + \sum_{i=1}^{n_w} A_i \hat{\mu}^{(i)}$, $\hat{B} = B_0 + \sum_{i=1}^{n_w} B_i \hat{\mu}^{(i)}$, $\mathbf{A} = \begin{bmatrix} A_1^\top & \cdots & A_{n_w}^\top \end{bmatrix}^\top$, $\mathbf{B} = \begin{bmatrix} B_1^\top & \cdots & B_{n_w}^\top \end{bmatrix}^\top$ and $H_i = \sum_{j=1}^{n_w} [\hat{\Sigma}^{1/2}]_{ji} (A_j W + B_j V)$

SDPs — Full uncertainty

solution

- **approximation of value function:**

$$V(x) = x^\top W^{-1} x$$

- **optimal controller:**

$$K = VW^{-1}$$

$\hat{A} = A_0 + \sum_{i=1}^{n_w} A_i \hat{\mu}^{(i)}$, $\hat{B} = B_0 + \sum_{i=1}^{n_w} B_i \hat{\mu}^{(i)}$, $\mathbf{A} = \left[A_1^\top \ \cdots \ A_{n_w}^\top \right]^\top$, $\mathbf{B} = \left[B_1^\top \ \cdots \ B_{n_w}^\top \right]^\top$ and $H_i = \sum_{j=1}^{n_w} [\hat{\Sigma}^{1/2}]_{ji} (A_j W + B_j V)$