SAND2023-07554C

Sandia National Laboratories

A Bayesian approach to designing experiments that account for risk

RISE: Robust, Interpretable, Scalable, Efficient

Z Fundamental Science Workshop
Data Science and Machine Learning Breakout Session

Bekah White‡, John Jakeman‡, Bart van Bloemen Waanders‡, Drew Kouri‡, and Alen Alexanderian†

‡ Sandia National Laboratories Center for Computing Research

† North Carolina State University Department of Mathematics

U.S. Department of Energy, NNSA

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Optimal experimental design (OED) predicts *a-priori* what data is most informative to collect

Sensor placement

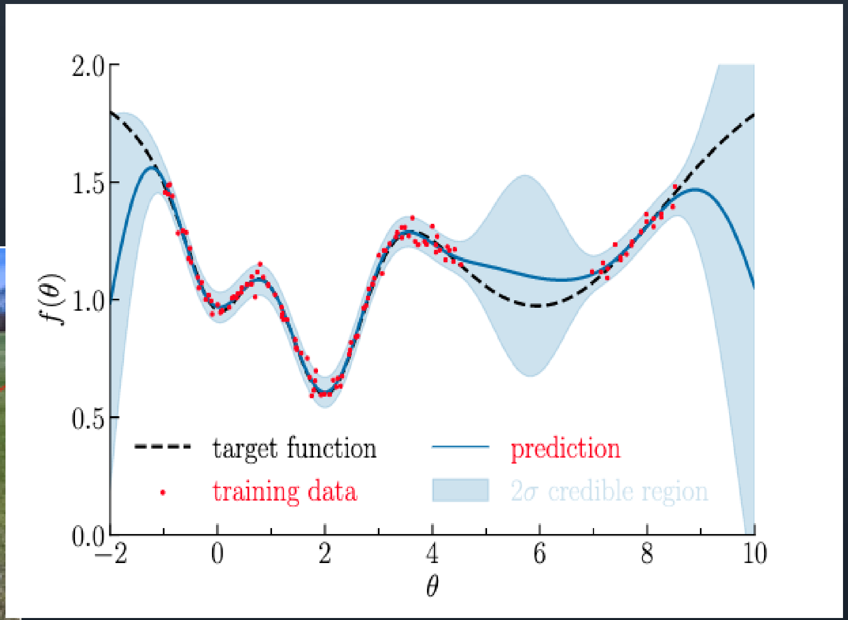

Training surrogate model†

Important tool in model development

- **Data is expensive**
- **Limited resources**

† F. Leclercq. Bayesian optimization for likelihood-free cosmological inference. 2018

‡T. Filkin, N. Sliusar, M. Ritzowski, and M. Huber-Humer. Unmanned aerial vehicles for operational monitoring of landfills. 2021

Bayesian optimal experimental design (OED) relies on physics-based models to predict what data is informative

Data model:

$$y = \text{model} + \text{noise}$$

$$= f(\boldsymbol{\theta}) + \text{noise}$$

depends on

Uncertain model parameters

The experimental design

† F. Leclercq. Bayesian optimization for likelihood-free cosmological inference. 2018

‡T. Filkin, N. Sliusar, M. Ritzowski, and M. Huber-Humer. Unmanned aerial vehicles for operational monitoring of landfills. 2021

Bayesian optimal experimental design (OED) relies on physics-based models to predict what data is informative

Data model:

$$y = \text{model} + \text{noise}$$

$$= f(\boldsymbol{\theta}) + \text{noise}$$

depends on

Uncertain model parameters

The experimental design

- Where to place sensors
- What frequency ranges to interrogate a system
- Where input loads to provide a system

† F. Leclercq. Bayesian optimization for likelihood-free cosmological inference. 2018

‡T. Filkin, N. Sliusar, M. Ritzowski, and M. Huber-Humer. Unmanned aerial vehicles for operational monitoring of landfills. 2021

Bayesian optimal experimental design minimizes uncertainty associated with the posterior distribution

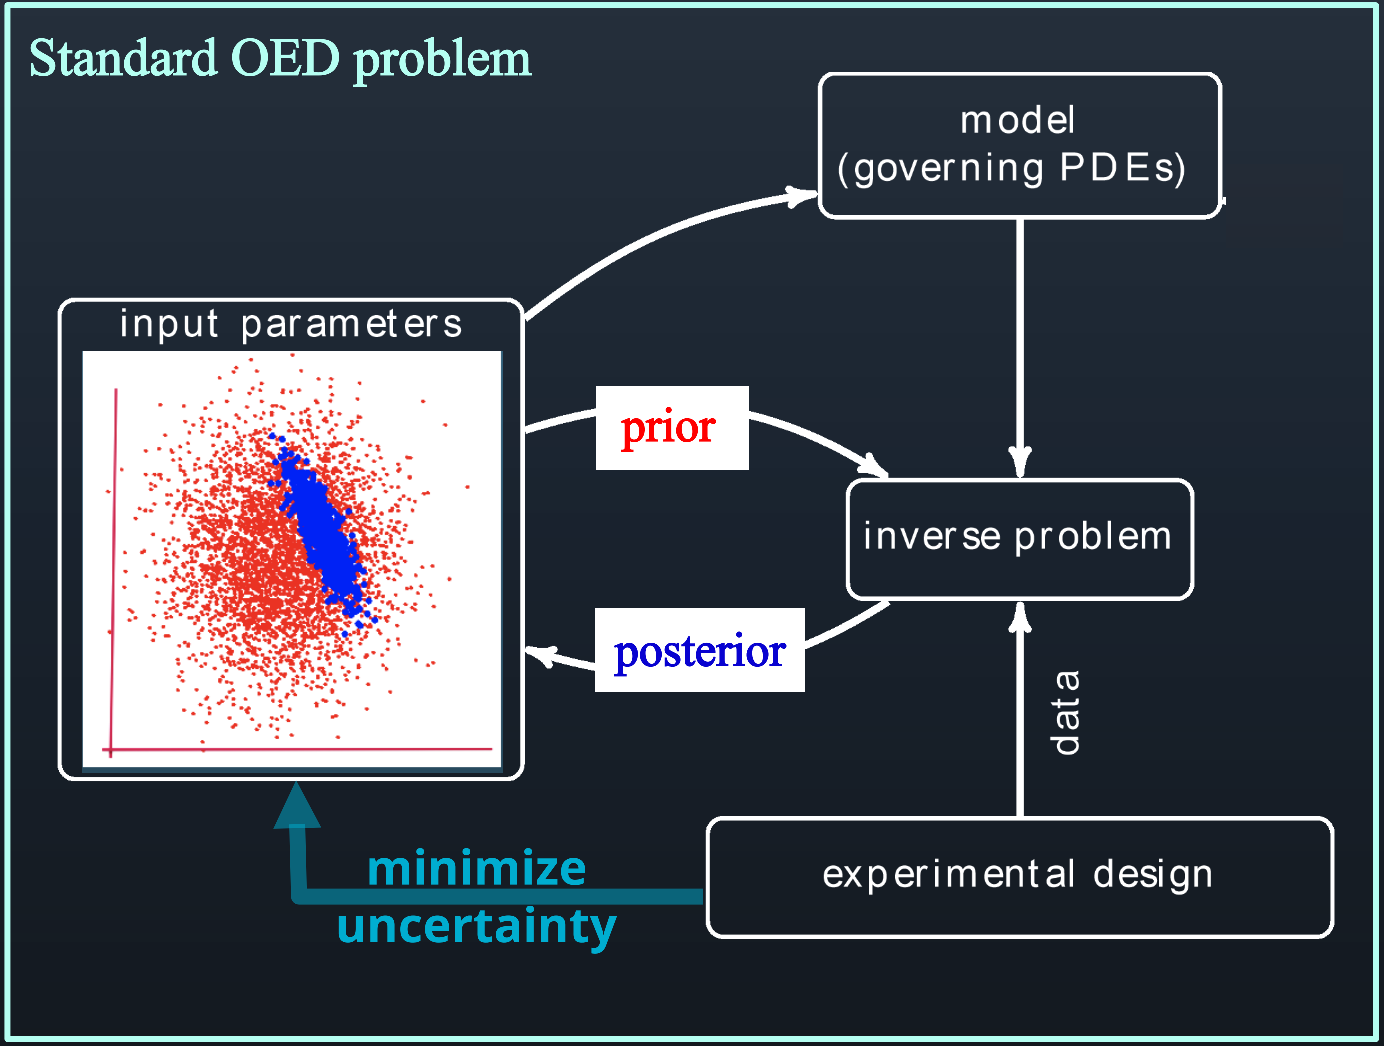

Bayes' Rule

$$\pi(\theta|y) \propto \underbrace{\pi(y|\theta)}_{\text{Likelihood}} \underbrace{\pi_{\text{pri}}(\theta)}_{\text{Prior}}$$

Bayesian optimal experimental design minimizes uncertainty associated with the posterior distribution

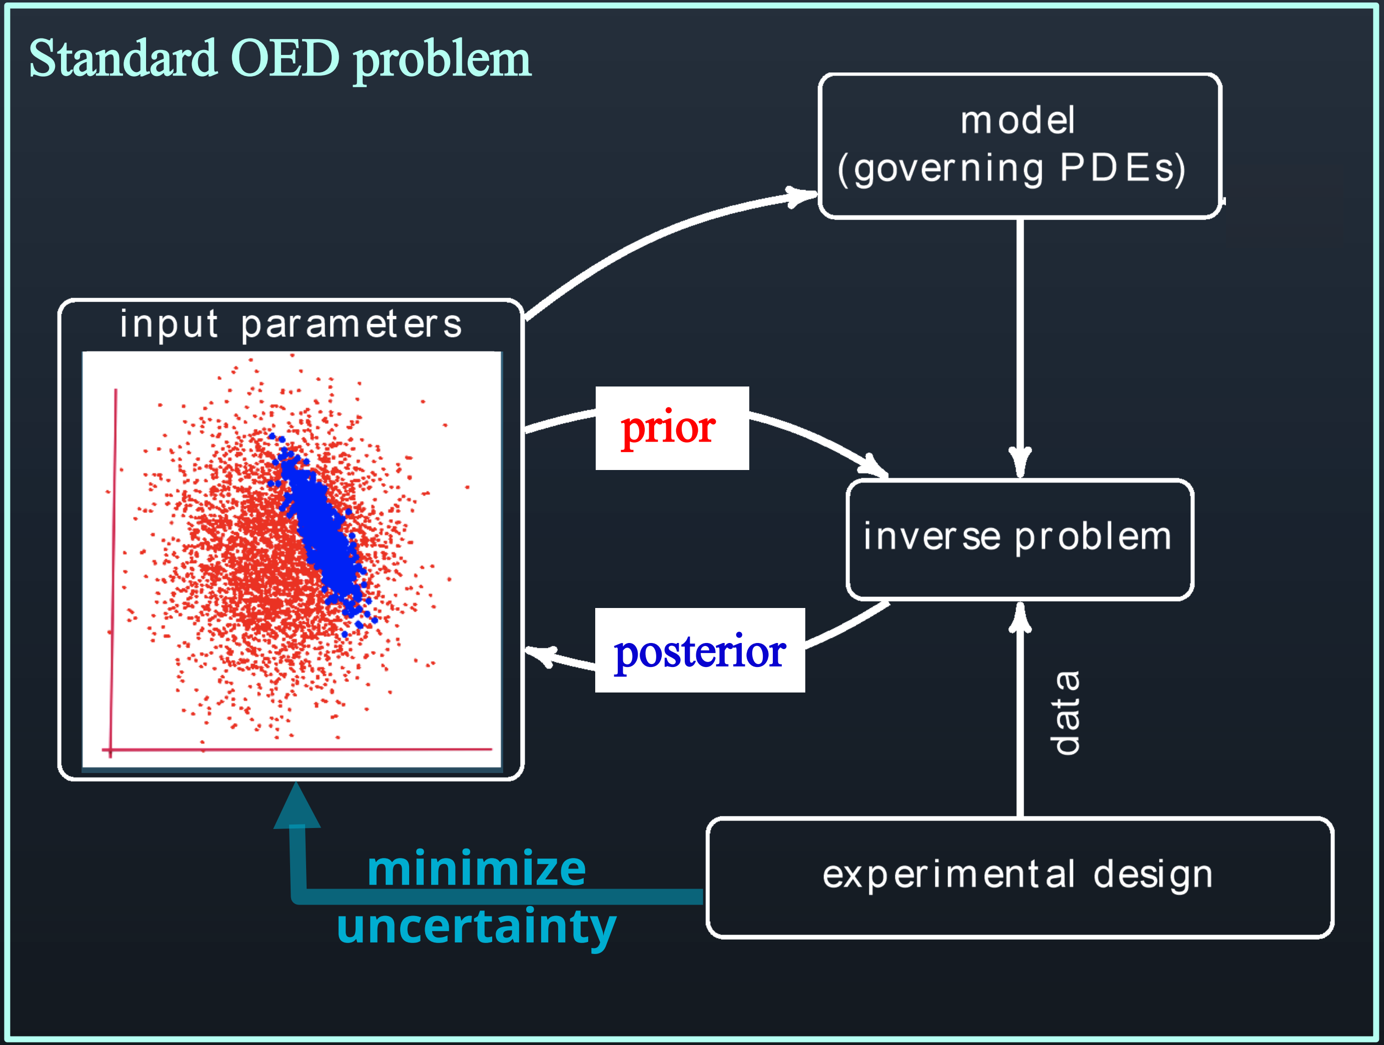

Bayes' Rule

$$\pi(\theta|\boldsymbol{y}) \propto \pi(\boldsymbol{y}\,|\theta)\,\pi_{\text{pri}}(\theta)$$

Likelihood

$$\pi(\boldsymbol{y}\,|\theta) \propto \exp\left(-\frac{1}{2}\,\|f(\theta)-\boldsymbol{y}\|_{\Gamma}^{2}\right)$$

Experimental Data

Goal-oriented approaches minimize uncertainty directly in quantities-of-interest

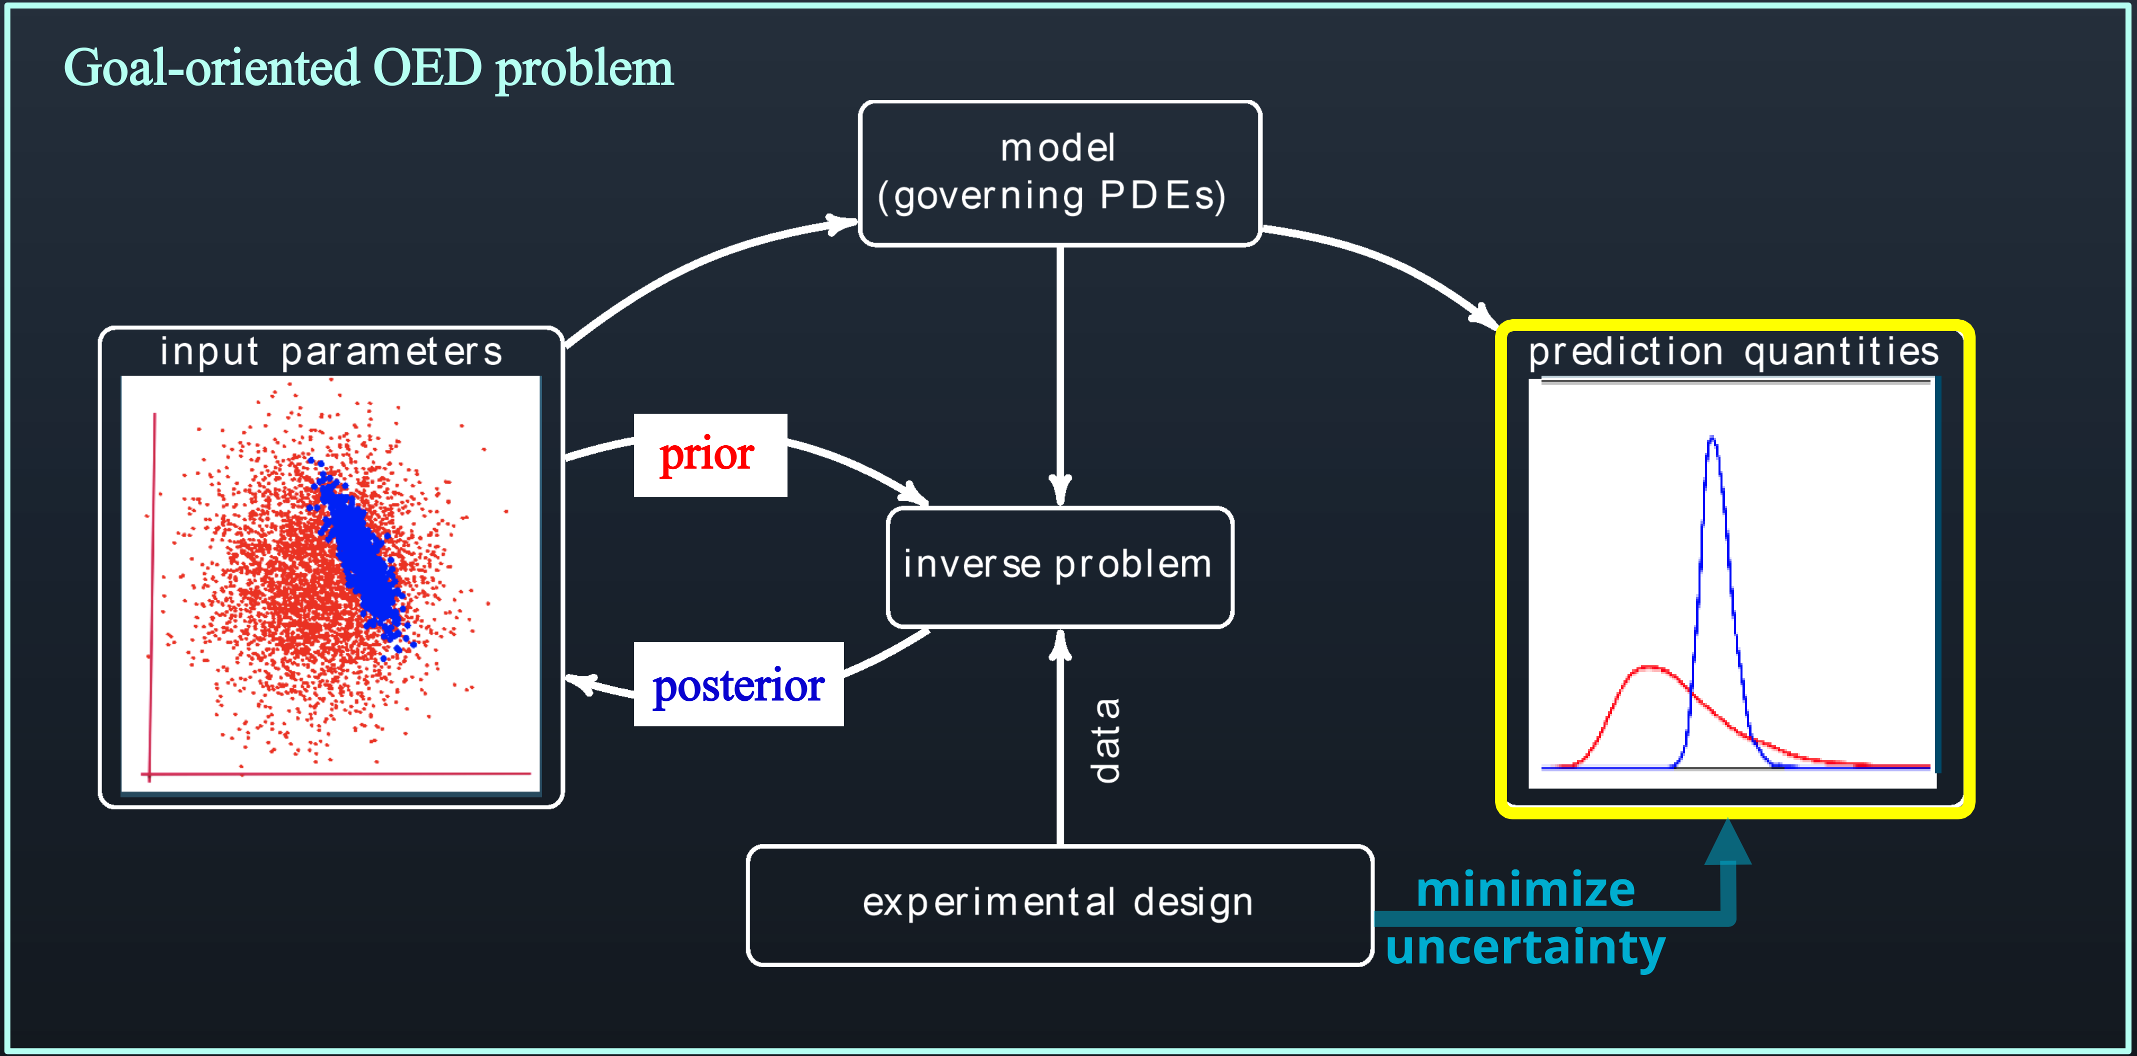

Goal-oriented OED allows us to introduce the notion of risk in experimental design

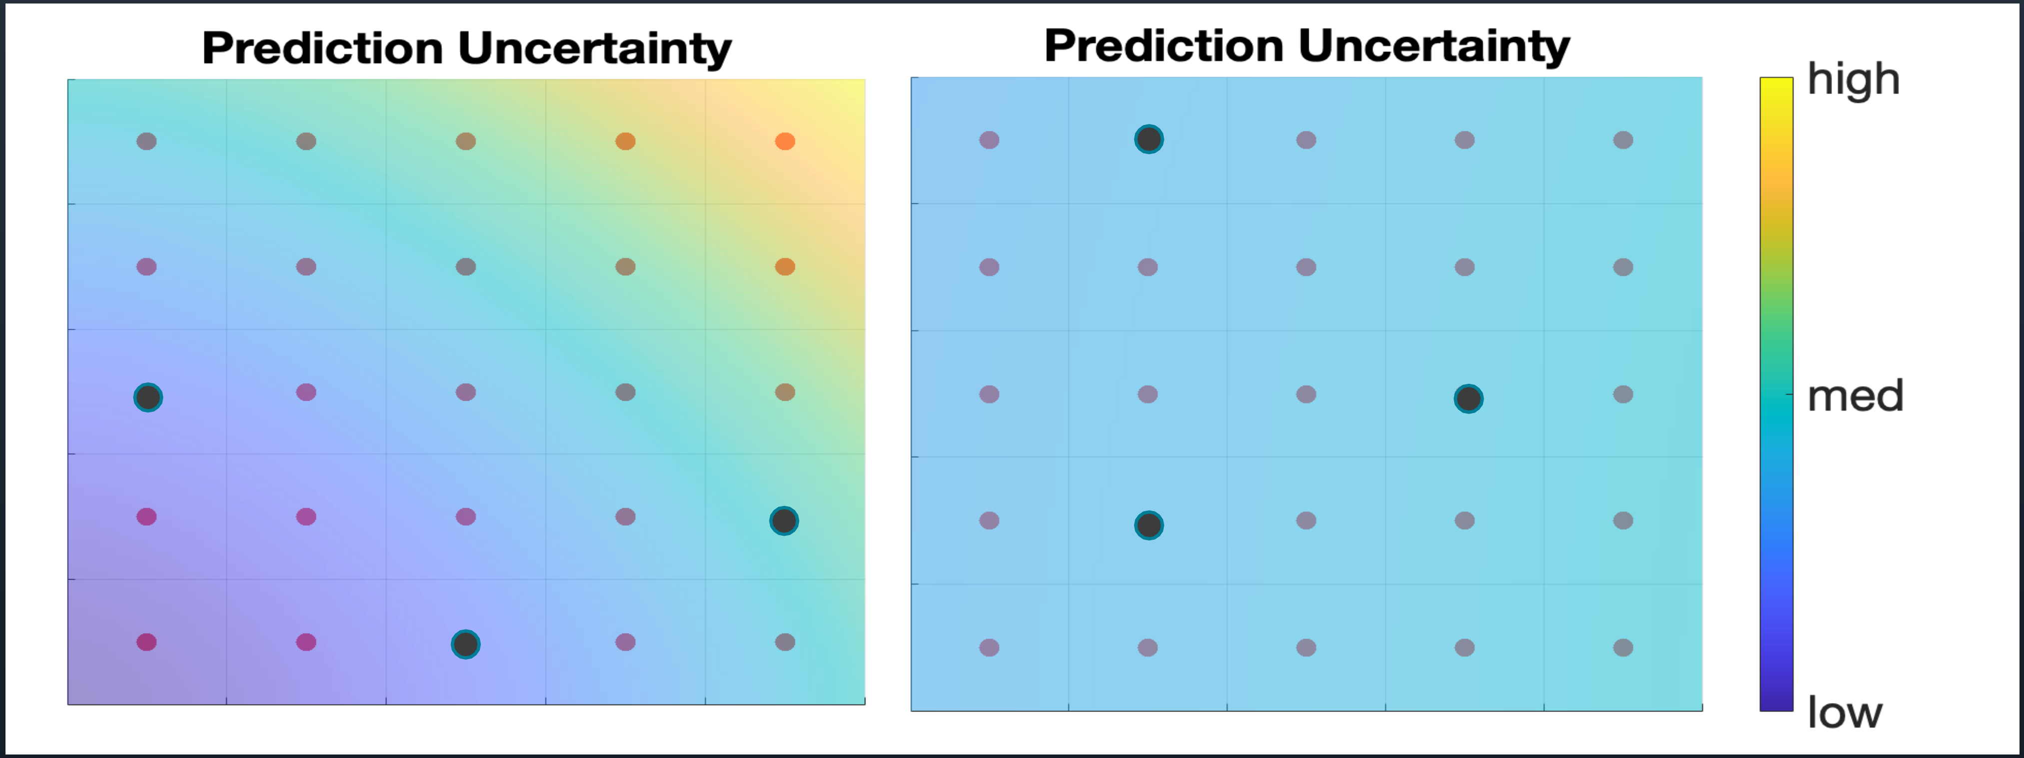

Which design is optimal for reducing prediction uncertainty?

Goal-oriented OED allows us to introduce the notion of risk in experimental design

High uncertainties = poor model predictions

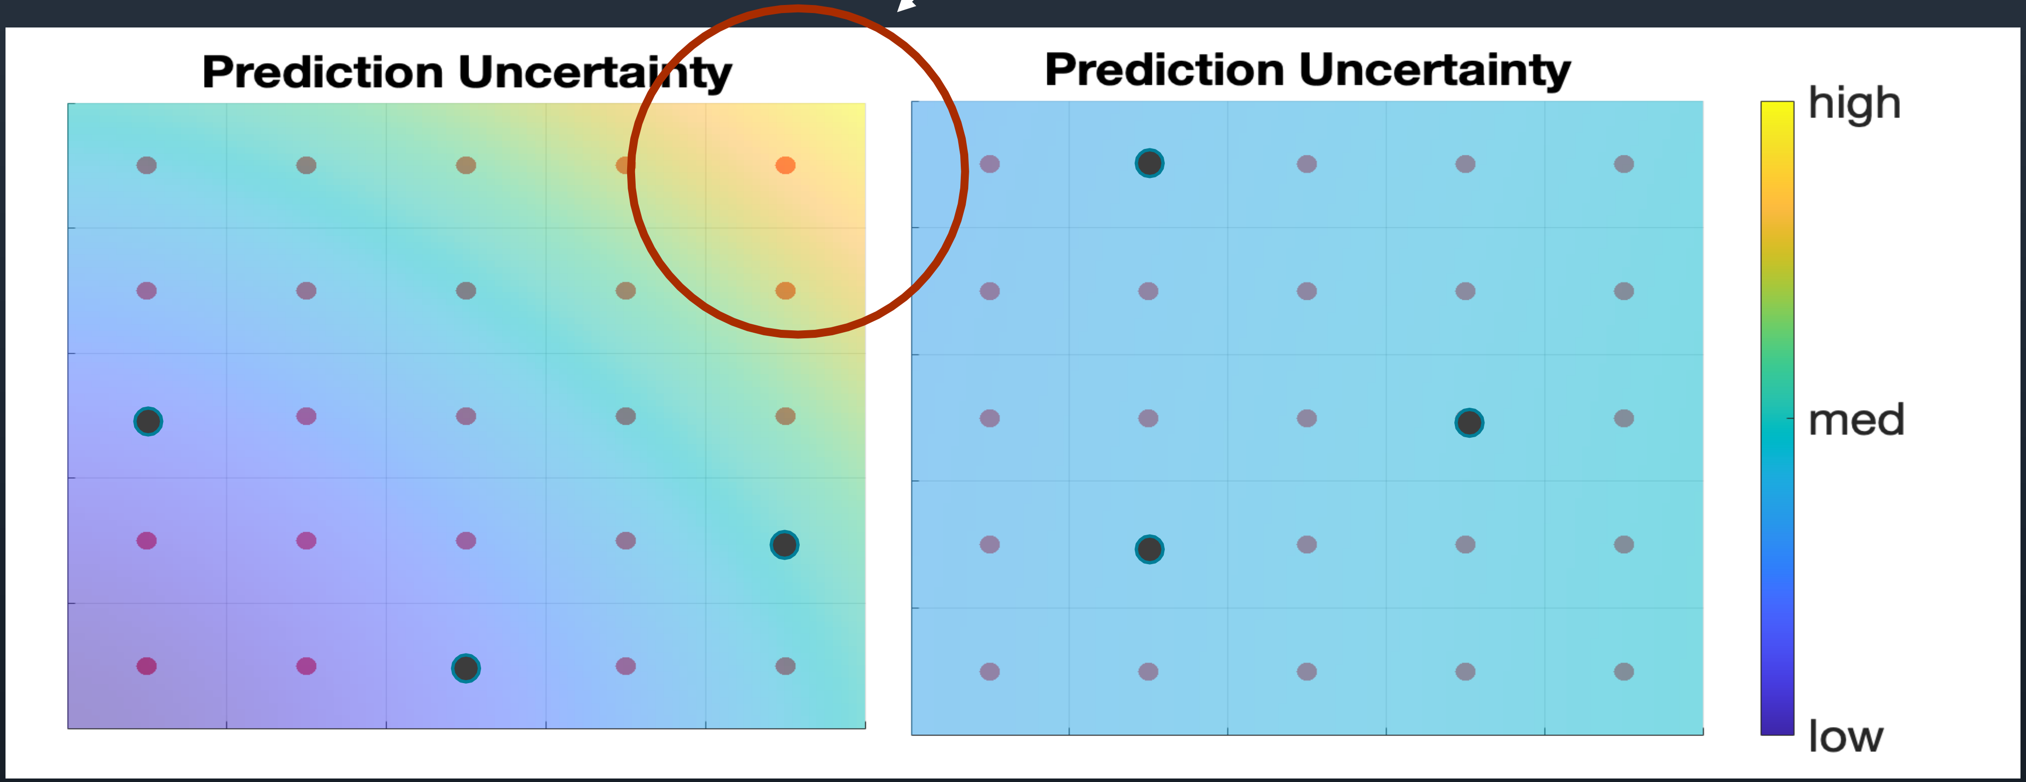

Which design is optimal for reducing prediction uncertainty?

It depends on how much risk you are willing to take

Goal-oriented OED allows us to introduce the notion of risk in experimental design

Classical approaches offer two choices:
minimize

Average uncertainty

Worst-case uncertainty

Our goal is to create a more flexible framework for accounting for risk preferences in nonlinear Bayesian OED problems

1. Introduce risk measures

2. Show how risk measures are used in Bayesian OED

3. Computational examples

The average-value-at-risk (AVaR) measures tail statistics

AVaR:

$$\mathcal{R}[X] := \mathrm{AVaR_p}[X] = \frac{1}{1-p}\int_{q_p}^{\infty} x\,\pi(x)dx$$

p-quantile

Rockafellar, R.T, and Uryasev, S., The fundamental risk quadrangle in risk management, optimization and statistical estimation, *Surveys in Operations Research and Management Science,* 2013

The average-value-at-risk (AVaR) measures tail statistics

AVaR:

$$\mathcal{R}[X] := \mathrm{AVaR}_p[X] = \frac{1}{1-p}\int_{q_p}^{\infty} x\,\pi(x)dx$$

p-quantile

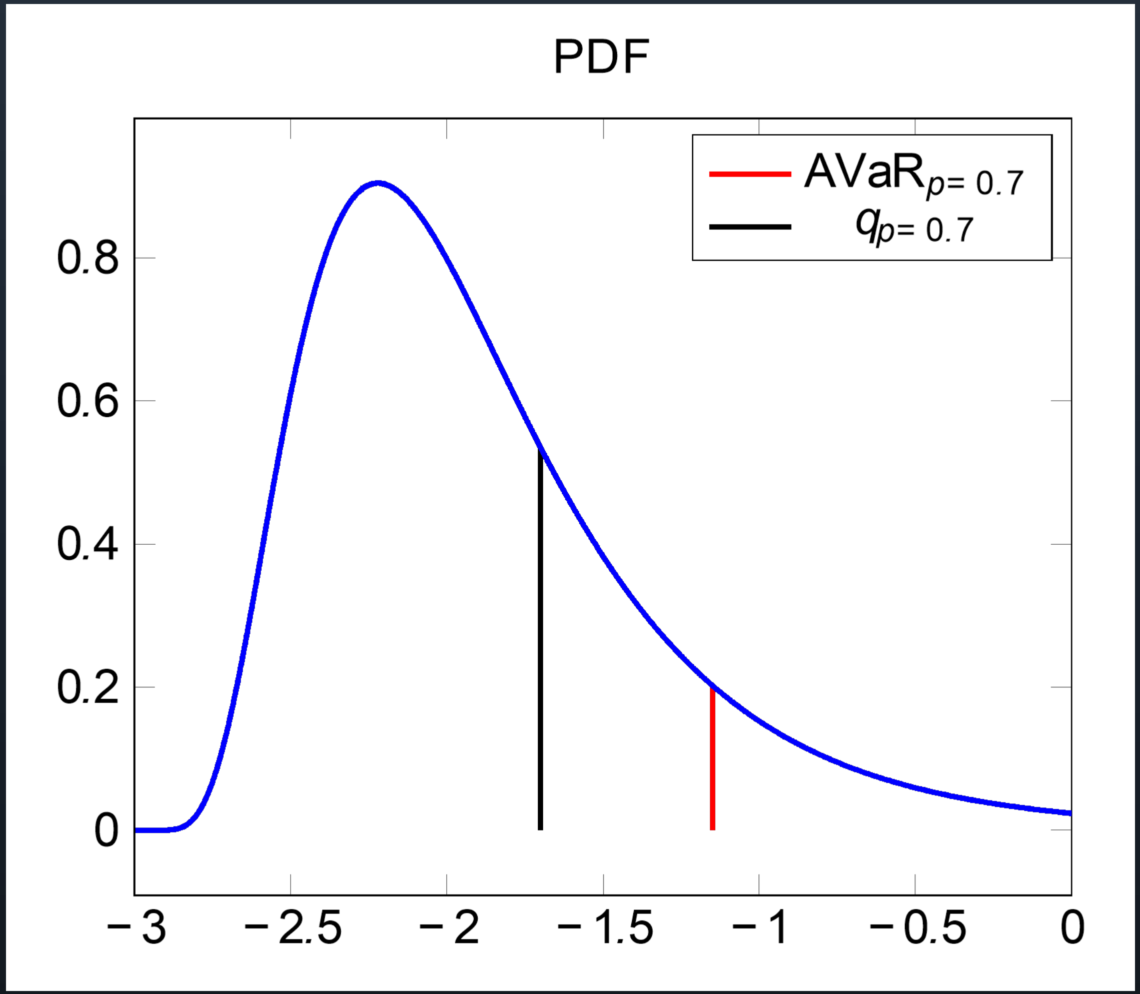

Expected deviations

The average-value-at-risk (AVaR) measures tail statistics

AVaR:

$$\mathcal{R}[X] := \mathrm{AVaR_p}[X] = \frac{1}{1-p}\int_{q_p}^{\infty} x\,\pi(x)dx$$

p-quantile

$$p = 0 \Rightarrow \mathrm{E}[X]$$
$$p \to 1 \Rightarrow \sup[X]$$

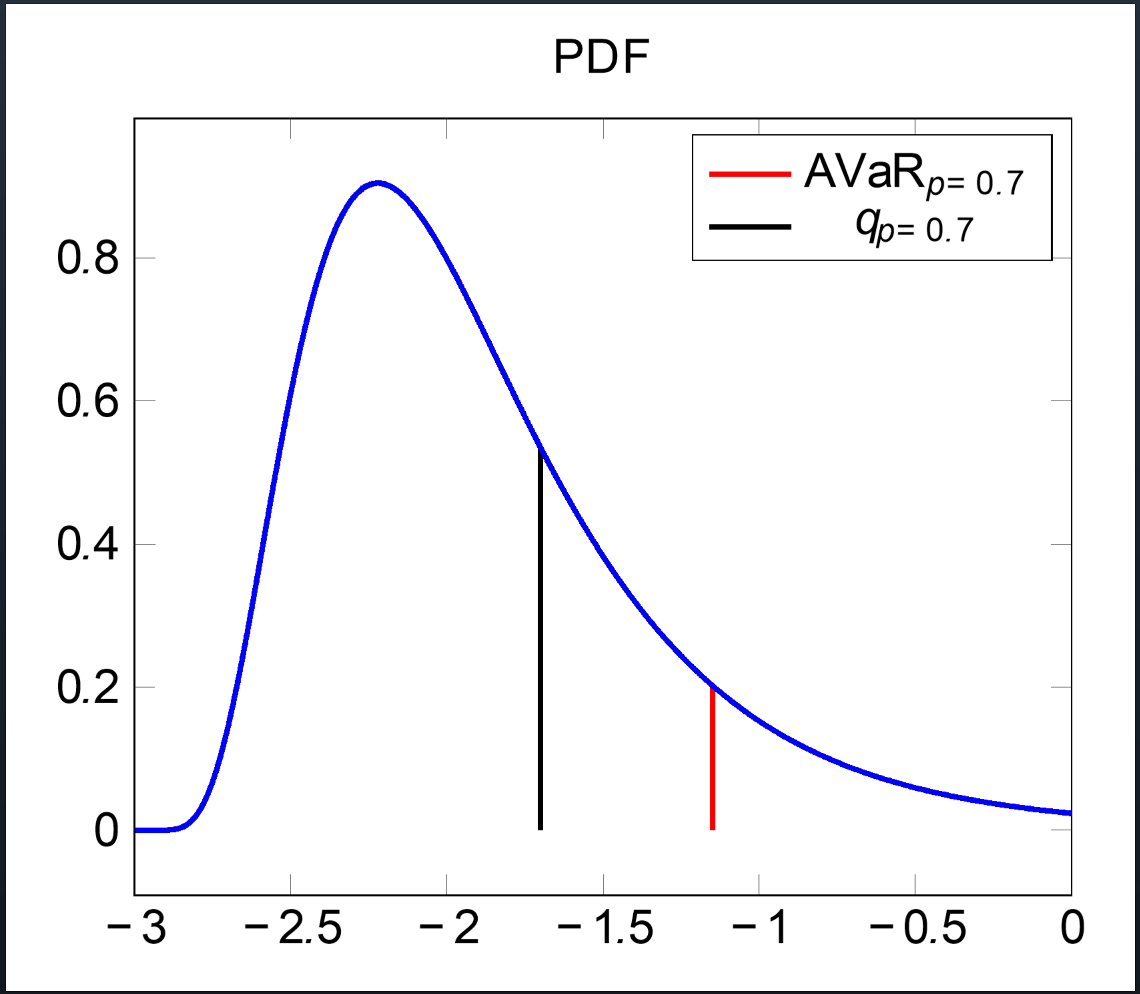

Expected deviations

The average-value-at-risk (AVaR) measures tail statistics

AVaR:

$$\mathcal{R}[X] := \mathrm{AVaR}_{\mathrm{p}}[X] = \frac{1}{1-p}\int_{q_p}^{\infty} x\,\pi(x)dx$$

p-quantile

$$p = 0 \Rightarrow \mathrm{E}[X]$$
$$p \to 1 \Rightarrow \sup[X]$$

Nonlinear interpolation between minimizing the average versus worst-case prediction uncertainty across the domain

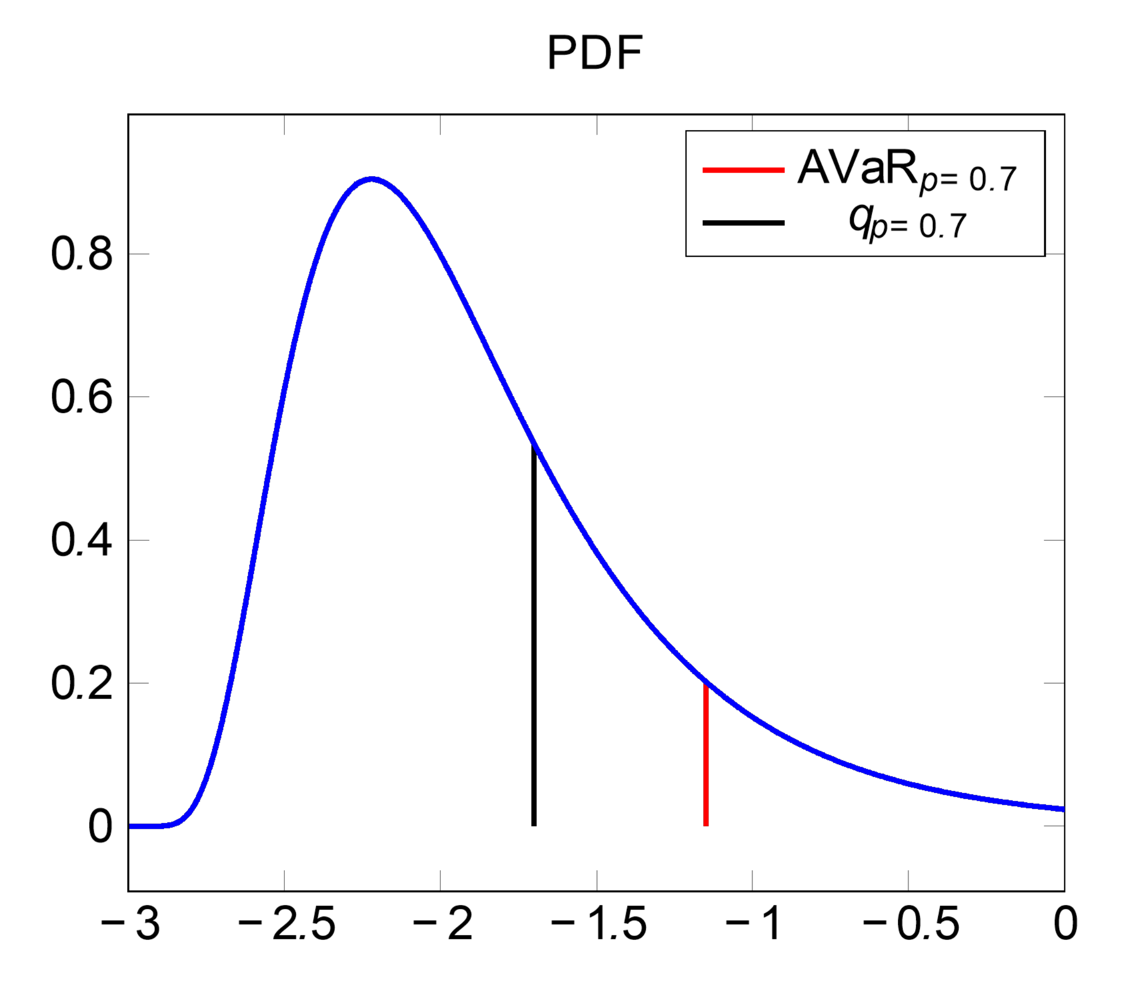

Expected deviations

Kouri, Drew P., John D. Jakeman, and J. Gabriel Huerta. "Risk-Adapted Optimal Experimental Design." *SIAM/ASA Journal on Uncertainty Quantification* 10, no. 2 (June 30, 2022): 687–716. https://doi.org/10.1137/20M1357615.

Risk measures provided alternative statistics to compute experimental design objective functions

Optimal design

$$\underbrace{\xi^*}_{\text{Optimal design}} = \min_{\xi} \underbrace{U(\xi)}_{\text{Objective function}}$$

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Vector valued quantify-of-interest

- Model prediction at every point in a domain

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Measure of uncertainty of deviation

- Variance
- KL-divergence

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E\left[E_y\left[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]\right]\right]$$

Expectation with respect to the likely data

- $\boldsymbol{y} = f(\boldsymbol{\theta}) + \text{ noise}$

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Average over the domain

$$\boldsymbol{x} \in \boldsymbol{\Omega}$$

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Statistics that can be replaced with risk measures

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E\left[E_y\left[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]\right]\right]$$

Replace average prediction uncertainty with the average-value-at-risk

Numerical example: Modeling the concentration of a contaminant over a 2D domain

Steady state advection-diffusion:

$$-\nabla \cdot \left(a(\boldsymbol{x}, \boldsymbol{\theta})\nabla u\right) + b\nabla u = f, \text{in } \Omega = [0, 1] \cup [0,1]$$

$$u(\boldsymbol{x}) = 0, \quad \text{on } \Gamma_D$$

$$\nabla u(\boldsymbol{x}) = 0, \text{on } \Gamma_N = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$$

Numerical example: Modeling the concentration of a contaminant over a 2D domain

Steady state advection-diffusion:

$$-\nabla \cdot \left(a(\boldsymbol{x}, \boldsymbol{\theta})\nabla u\right) + b\nabla u = f \text{ in } \Omega = [0,1] \cup [0,1]$$

$$u(\boldsymbol{x}) = 0, \quad \text{on } \Gamma_D$$

$$\nabla u(\boldsymbol{x}) = 0 \text{ on } \Gamma_N = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$$

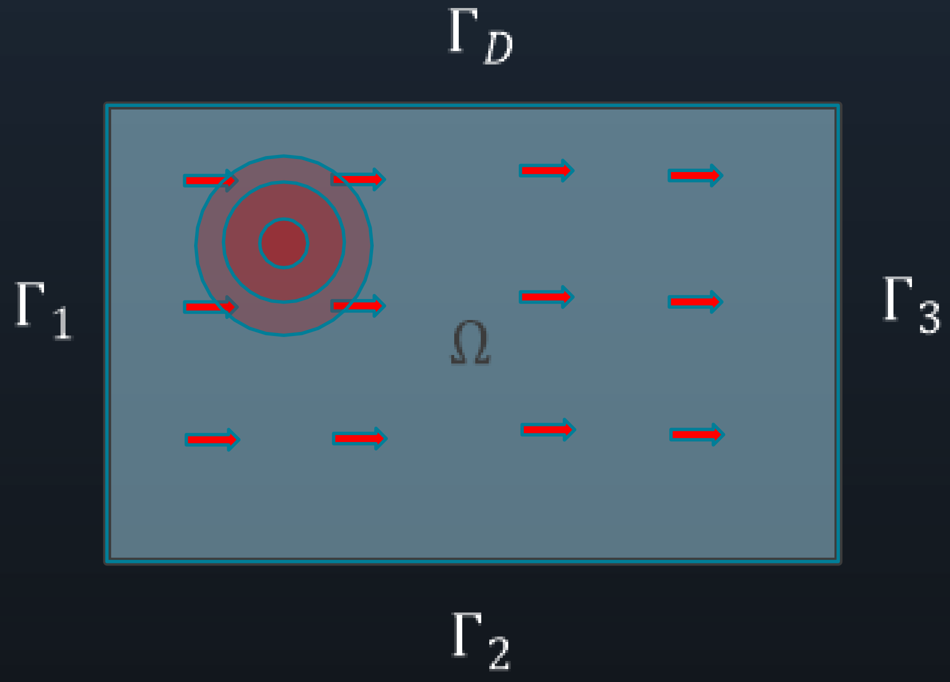

Numerical example: Modeling the concentration of a contaminant over a 2D domain

Steady state advection-diffusion:

$$-\nabla \cdot \left(a(\boldsymbol{x}, \boldsymbol{\theta})\nabla u\right) + b\nabla u = f, \text{in } \Omega = [0, 1] \cup [0,1]$$

$$u(\boldsymbol{x}) = 0, \quad \text{on } \Gamma_D$$

$$\nabla u(\boldsymbol{x}) = 0, \text{on } \Gamma_N = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$$

Diffusion

$$a(\boldsymbol{x}, \boldsymbol{\theta}) = \exp[\theta_1 \sin(x_1\pi)\sin(x_2\pi) + \theta_2 \cos(3/2x_1\pi)\cos(3/2x_2\pi)]$$

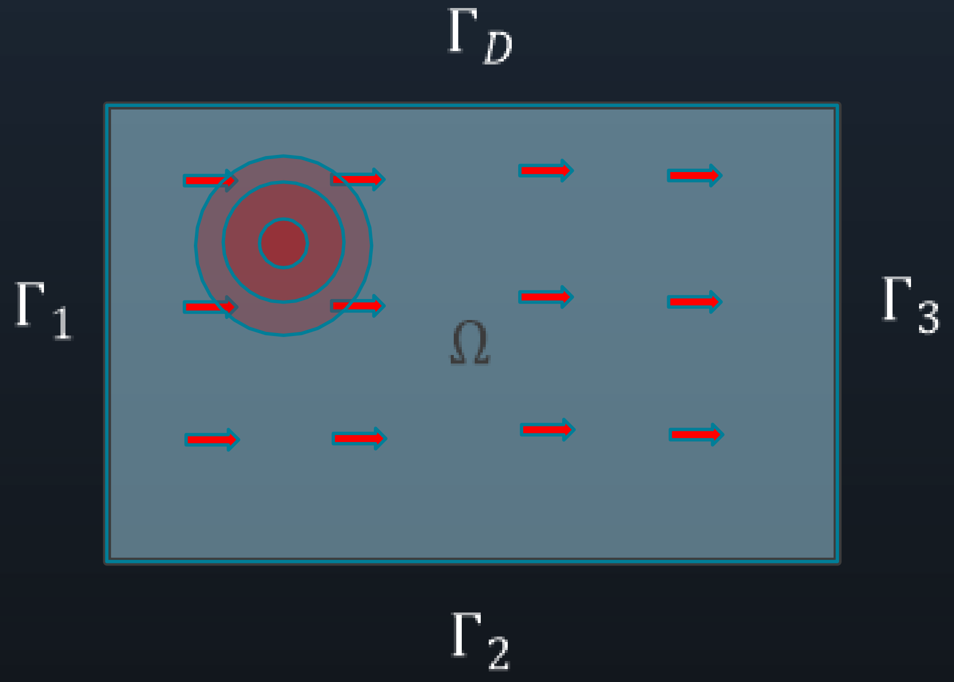

Numerical example: Modeling the concentration of a contaminant over a 2D domain

Steady state advection-diffusion:

$$-\nabla \cdot \left(a(\boldsymbol{x}, \boldsymbol{\theta})\nabla u\right) + b\nabla u = f \text{ in } \Omega = [0,1] \cup [0,1]$$

$$u(\boldsymbol{x}) = 0, \quad \text{on } \Gamma_D$$

$$\nabla u(\boldsymbol{x}) = 0 \text{ on } \Gamma_N = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$$

Diffusion

$$a(\boldsymbol{x}, \boldsymbol{\theta}) = \exp\left[\theta_1 \sin(x_1\pi)\sin(x_2\pi) + \theta_2 \cos(3/2x_1\pi)\cos(3/2x_2\pi)\right]$$

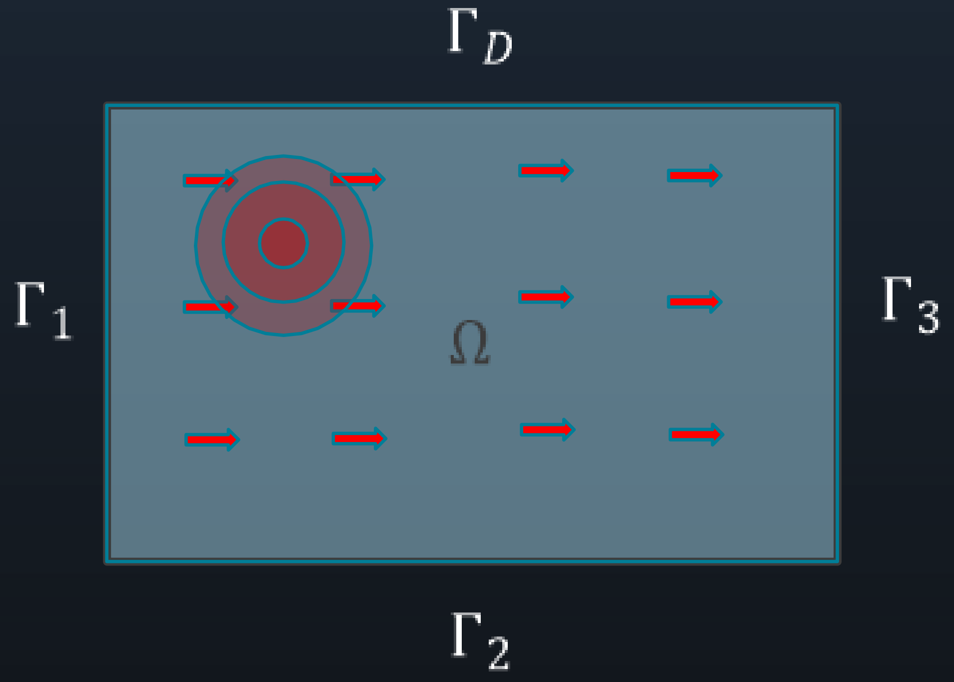

Quantify-of-interest – concentration across the domain

$$q(\boldsymbol{x}, \boldsymbol{\theta}) := u(\boldsymbol{x}, \boldsymbol{\theta}), \quad \boldsymbol{x} \in \Omega$$

The optimal experimental design problem

Determine optimal sensor locations ξ to measure the contaminant concentration $u(\boldsymbol{x}, \boldsymbol{\theta})$ to minimize uncertainty in the quantity-of-interest $q(x, \boldsymbol{\theta})$

The optimal experimental design problem

Determine optimal sensor locations ξ to measure the contaminant concentration $u(\boldsymbol{x}, \boldsymbol{\theta})$ to minimize uncertainty in the quantity-of-interest $q(x, \boldsymbol{\theta})$

Design

$$\xi = \begin{Bmatrix} \boldsymbol{x}_1, \ldots, \boldsymbol{x}_k \\ w_1, \ldots, w_k \end{Bmatrix}$$

- $\boldsymbol{x}_i \in [0,1] \cup [0,1]$ – Fixed spatial design candidates
- $w_i \in \{0,1\}$ – Binary weights
- $\sum w_i = N$ – Budget

The optimal experimental design problem

Determine optimal sensor locations ξ to measure the contaminant concentration $u(\boldsymbol{x}, \boldsymbol{\theta})$ to minimize uncertainty in the quantity-of-interest $q(x, \boldsymbol{\theta})$

Design

$$\xi = \begin{Bmatrix} \boldsymbol{x}_1, \ldots, \boldsymbol{x}_k \\ w_1, \ldots, w_k \end{Bmatrix}$$

- $\boldsymbol{x}_i \in [0,1] \cup [0,1]$ – Fixed spatial design candidates
- $w_i \in \{0,1\}$ – Binary weights
- $\sum w_i = N$ – Budget

Compare

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Vs.

$$U(\xi) = \mathrm{AVaR}_{0.95}[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Using the average-value-at-risk reduces max prediction variance

Optimal designs and corresponding prediction variances

(L) 1st optimal sensor (R) 1st & 2nd optimal sensors

(top)
$\mathcal{R} := E[\text{prediction variance}]$

(bottom)
$\mathcal{R} := \text{AVaR}_{p=0.95}[\text{prediction variance}]$

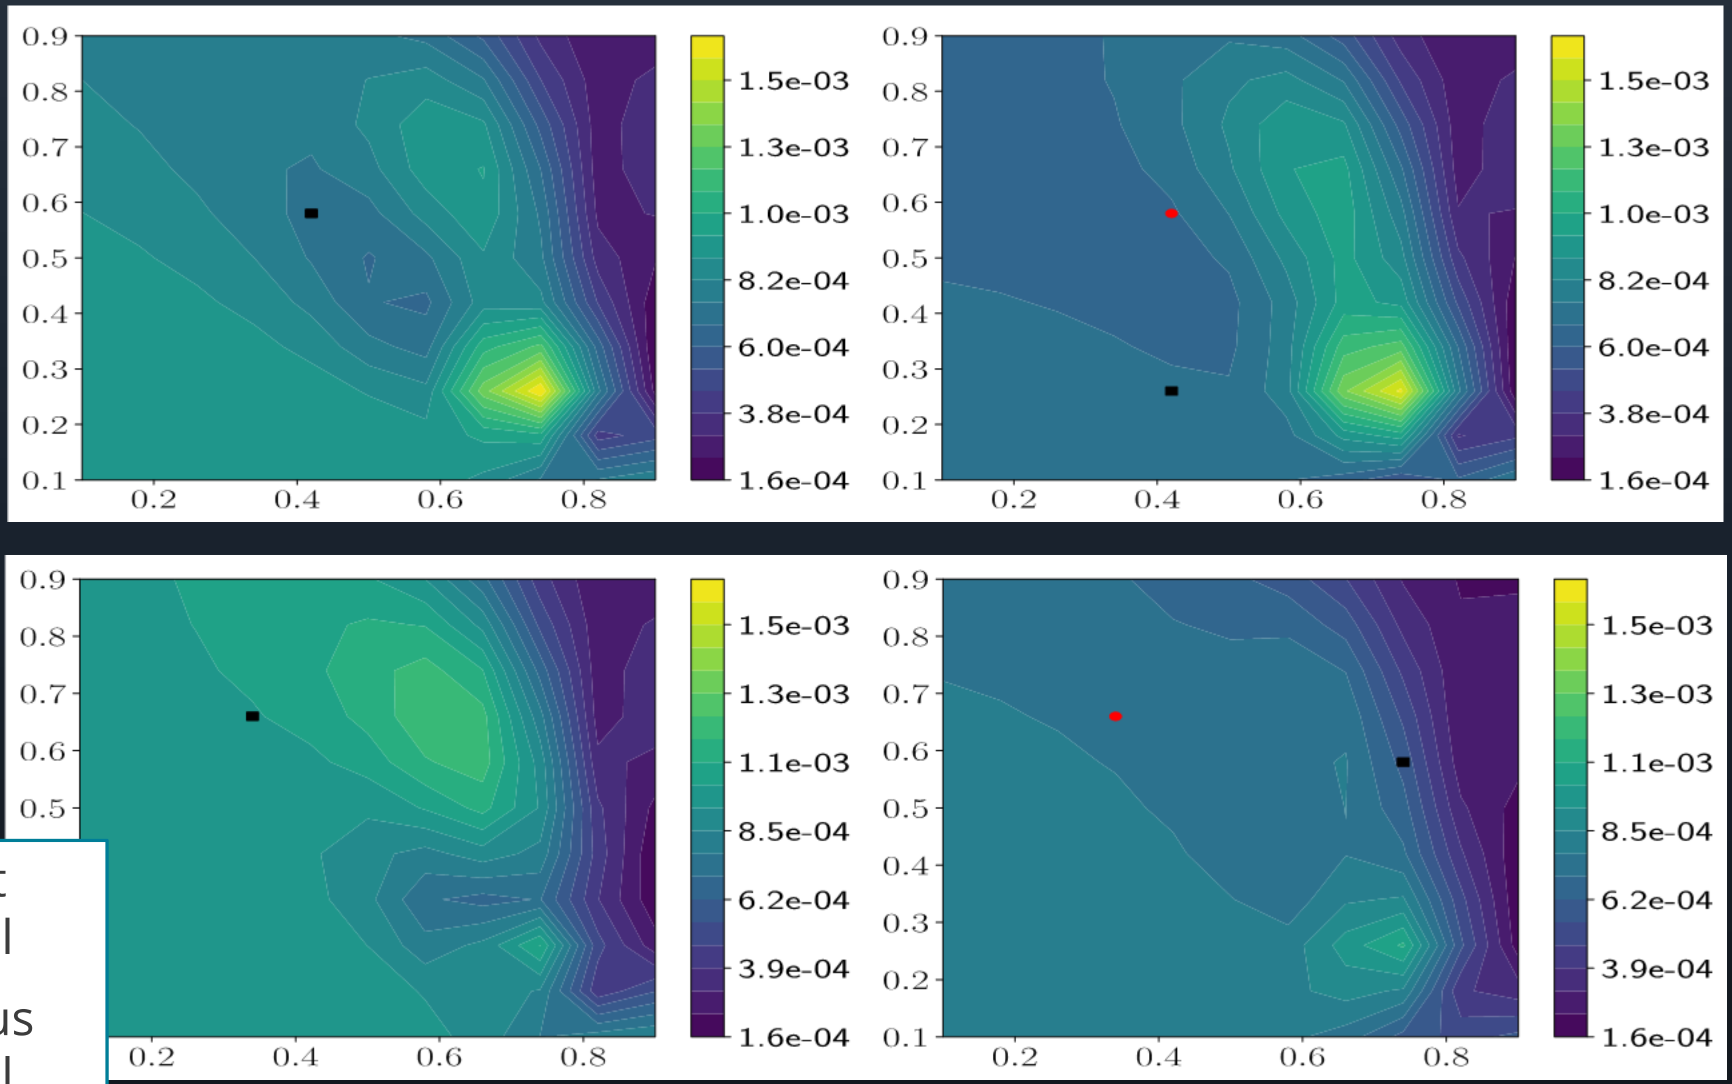

- current optimal
- previous optimal

Risk measures provided alternative statistics to compute experimental design objective functions

OED objective function

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Replace

1. average prediction uncertainty
2. Average over the likely data
3. Deviation measure

With the average-value-at-risk measures and deviations

Numerical example: Modeling a thermally activated battery

$$-\nabla \cdot \left(k(\boldsymbol{x},\theta)\, \nabla T(\boldsymbol{x},t)\right) = f(\boldsymbol{x}) \qquad in \quad \Omega = \Omega_1 \cup \Omega_2 \cup \Omega_3 \cup \Omega_4 \cup \Omega\_5$$

$$k(\boldsymbol{x},\boldsymbol{\theta}) \nabla T(\boldsymbol{x},t) \cdot n = \frac{1}{10}\left(T(\boldsymbol{x},t) - T_o\right) \qquad on \quad \partial\Omega$$

$$k(\boldsymbol{x},\boldsymbol{\theta}) = \begin{cases} 100, & \boldsymbol{x} \in \Omega_1 \\ 1, & \boldsymbol{x} \in \Omega_2 \\ \exp\left(\sum_{d=1}^{5} \lambda_d \phi_d(\boldsymbol{x}) \theta_d\right), & \boldsymbol{x} \in \Omega_3 \\ 1, & \boldsymbol{x} \in \Omega_4 \\ 100, & \boldsymbol{x} \in \Omega_5 \end{cases}$$

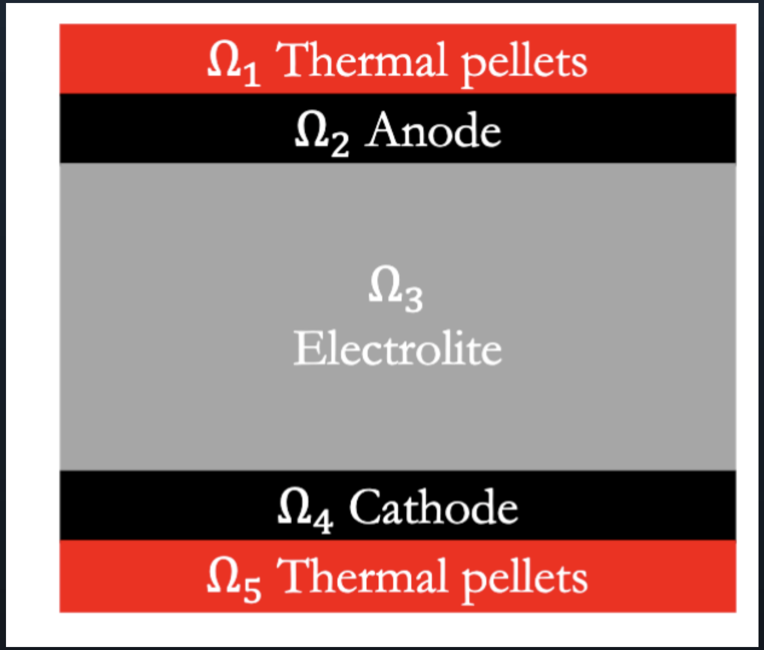

Numerical example: Modeling a thermally activated battery

$$-\nabla \cdot \big(k(\boldsymbol{x}, \theta)\, \nabla T(\boldsymbol{x}, t)\big) = f(\boldsymbol{x}) \qquad in \;\; \Omega = \Omega_1 \cup \Omega_2 \cup \Omega_3 \cup \Omega_4 \cup \Omega\_5$$

$$k(\boldsymbol{x}, \boldsymbol{\theta}) \nabla T(\boldsymbol{x}, t) \cdot n = \frac{1}{10}\big(T(\boldsymbol{x}, t) - T_o\big) \qquad on \;\; \partial\Omega$$

$$k(\boldsymbol{x}, \boldsymbol{\theta}) = \begin{cases} 100, & \boldsymbol{x} \in \Omega_1 \\ 1, & \boldsymbol{x} \in \Omega_2 \\ \exp\left(\sum_{d=1}^{5} \lambda_d \phi_d(\boldsymbol{x}) \theta_d\right), & \boldsymbol{x} \in \Omega_3 \\ 1, & \boldsymbol{x} \in \Omega_4 \\ 100, & \boldsymbol{x} \in \Omega_5 \end{cases}$$

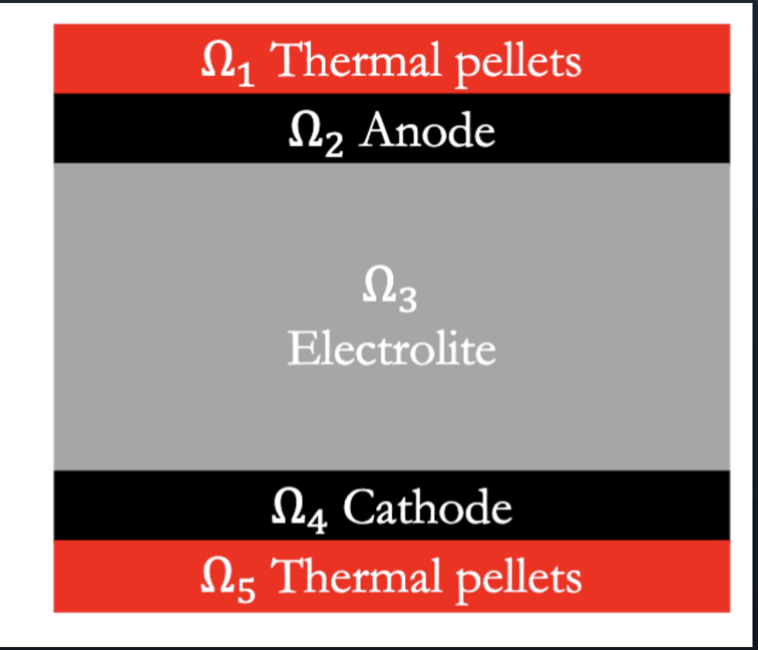

The eigenvalues and eigenfunctions – derived from Fredholm integral equations

Numerical example: Modeling a thermally activated battery

Candidate sensor locations – along the boundary of the electrolyte domain at times $t = 3.2, 5.4, 10.8$ seconds

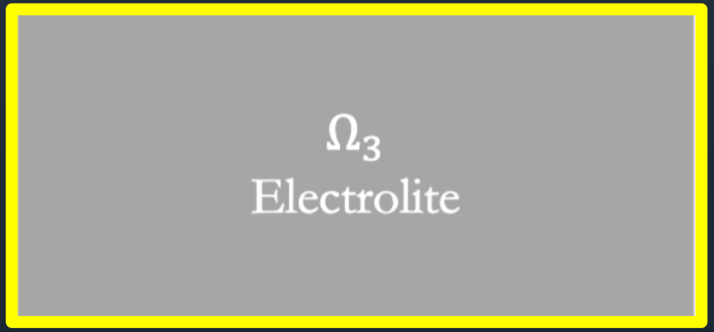

Quantify-of-interest – temperature in the electrolyte domain averaged across the three times

$$q(\boldsymbol{x}, \boldsymbol{\theta}) \coloneqq \bar{T}(\boldsymbol{x}, \boldsymbol{\theta}), \quad \boldsymbol{x} \in \Omega_3$$

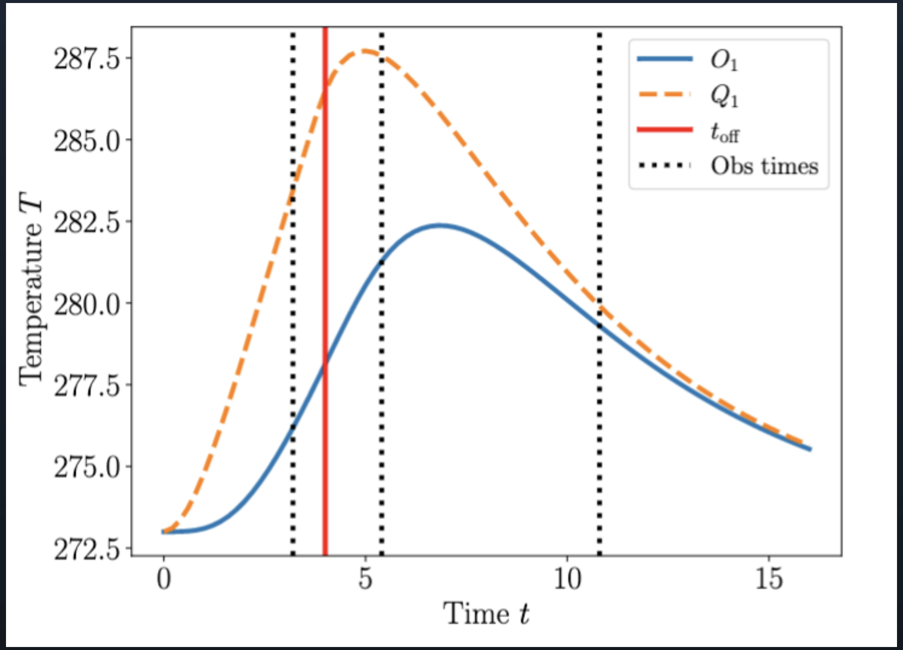

Numerical example: Modeling a thermally activated battery

Plot of the expected deviations across the electrolyte domain and the optimal sensor location

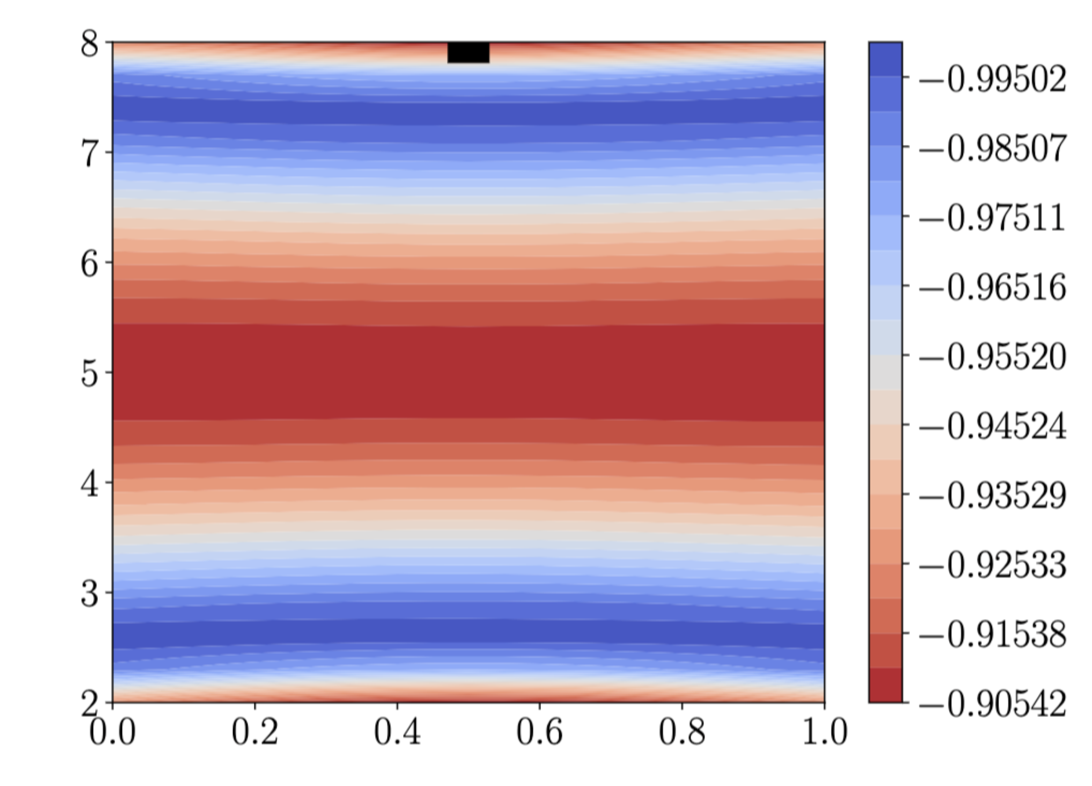

Risk neutral

$$U(\xi) = E[E_y[\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

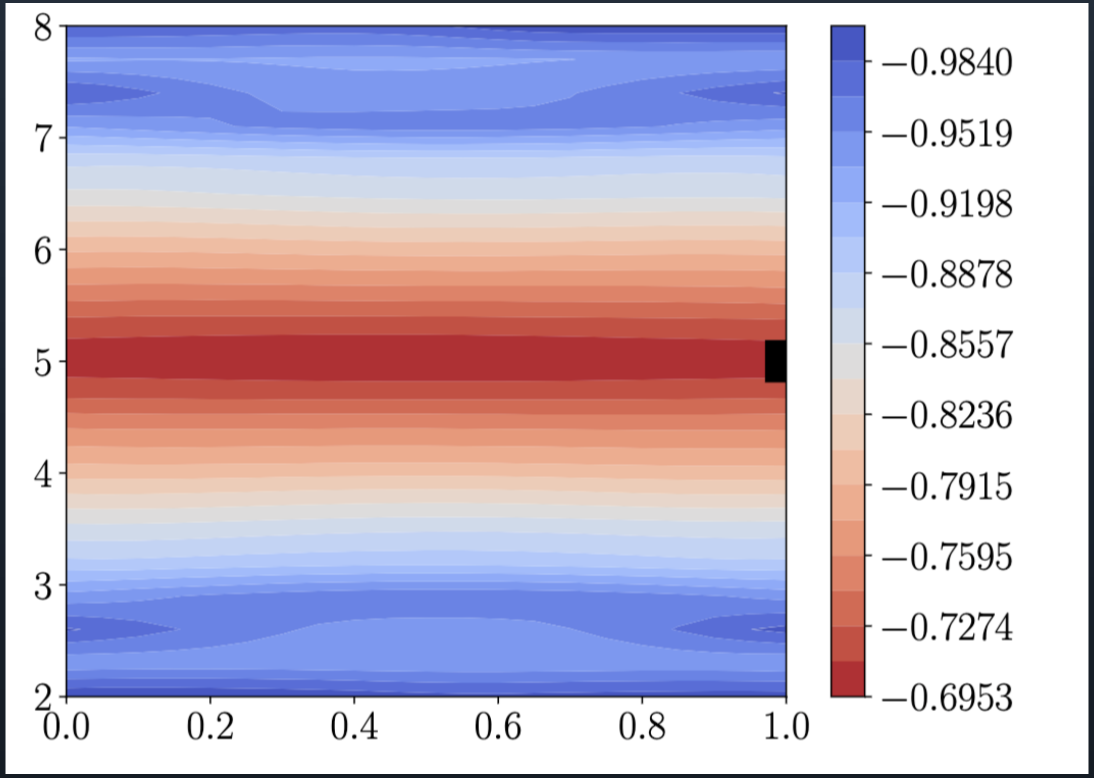

Risk averse

$$U(\xi) = R[R_y[R_\sigma[q(\boldsymbol{x}, \boldsymbol{\theta})]]]$$

Conclusions and future work

- Risk measures provide a more flexible experimental design framework
- Accounting for risk preferences changes the optimal experimental design
- Goal-oriented approaches are beneficial

Future research

- Efficient computation for large-scale, nonlinear problems

Publication

R. White, J. Jakeman, A. Alexanderian, D. Kouri, and B. van Bloemen Waanders. A Bayesian approach to risk-averse optimal experimental design. *In-progress*