SAND2018-6585C

Simulation-based Bayesian Experimental Design for Computationally Intensive Models

Xun (Ryan) Huan

Sandia National Laboratories
Joint work with Youssef M. Marzouk (MIT)

June 20, 2018

Uncertainty in engineering and science

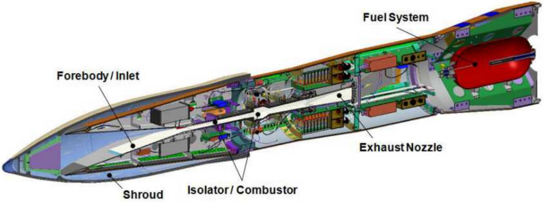

Uncertainty is everywhere

"Uncertainty is everywhere and you cannot escape from it."

— *Dennis Lindley*

Uncertainty quantification (UQ) provides a systematic approach to merging **data** with **models**, and allows us to:

- characterize
- incorporate
- propagate
- reduce

. . . uncertainty for realistic applications

UQ in engineering and science: big picture

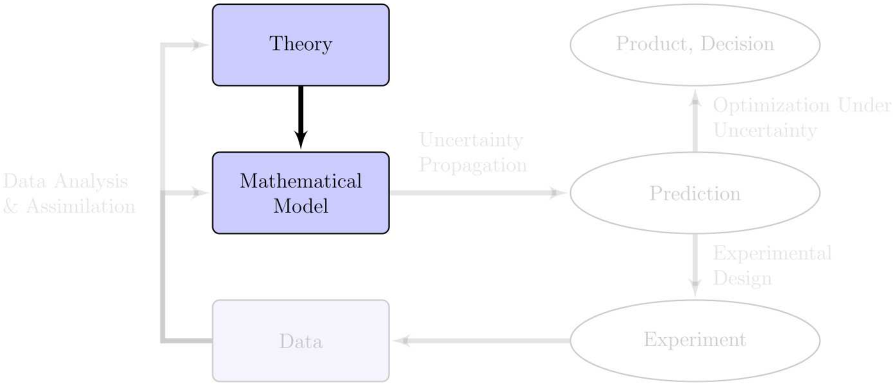

UQ in engineering and science: big picture

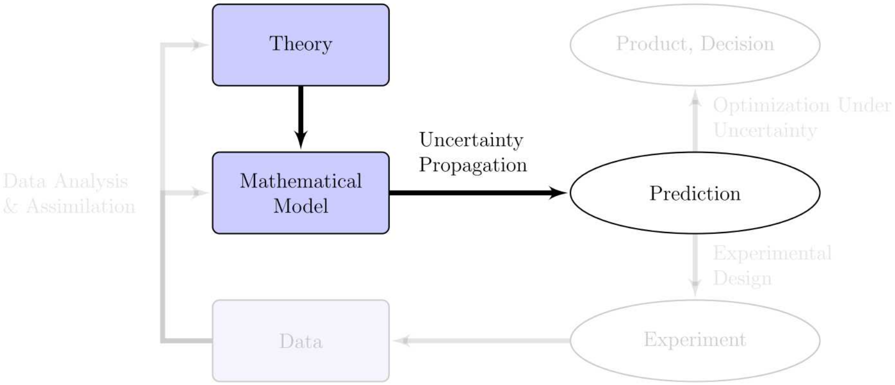

UQ in engineering and science: big picture

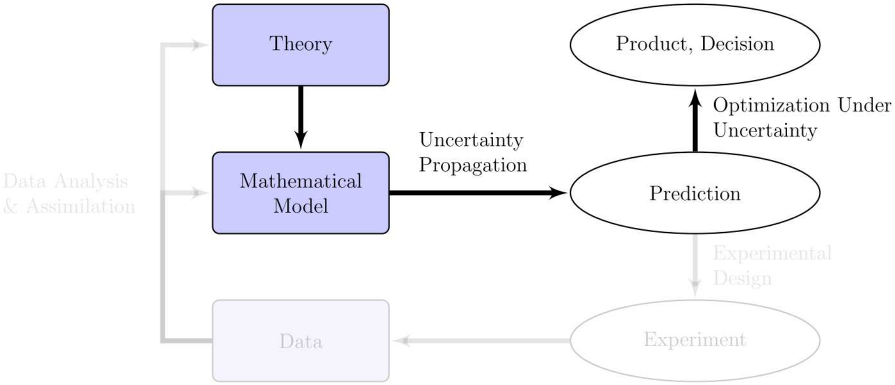

UQ in engineering and science: big picture

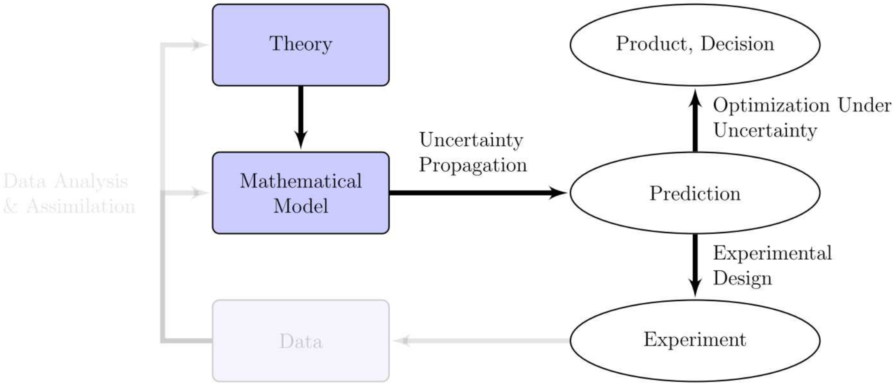

UQ in engineering and science: big picture

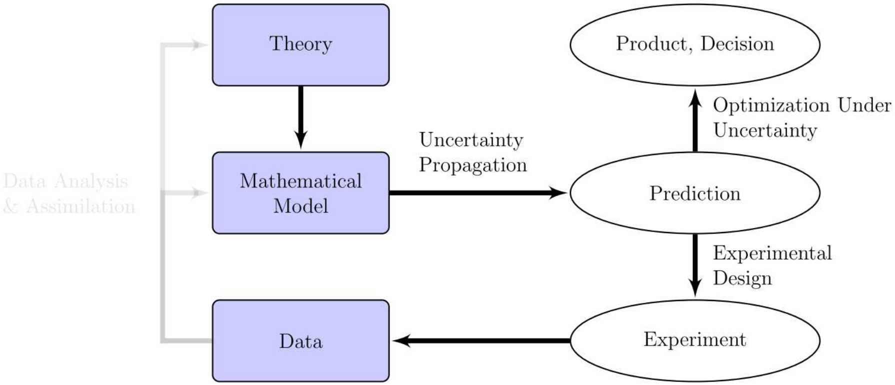

UQ in engineering and science: big picture

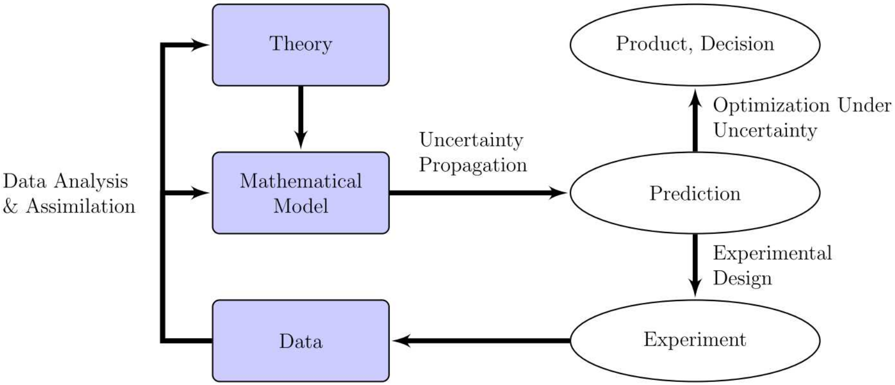

UQ in engineering and science: big picture

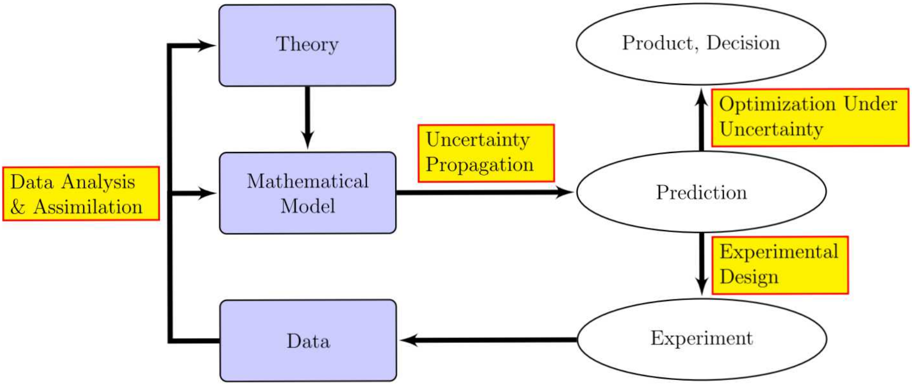

UQ in engineering and science: big picture

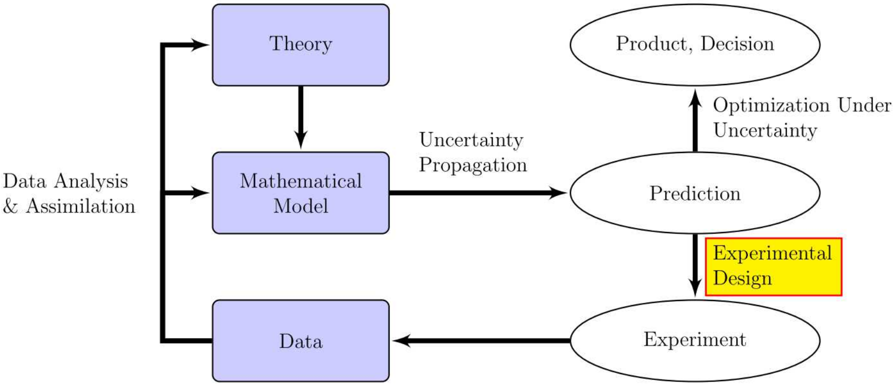

Outline

Outline

Some experiments are more useful than others

Experiments are:

- Expensive
- Time-consuming
- Delicate to perform

Experimental design helps address:

- Under what **conditions** to perform the experiment?
- **What**, **where**, and **when** to measure?
- **How many** experiments?

Batch (non-sequential) design landscape

Traditional statistical **design of experiments**: non simulation-based (e.g., space-filling, blocking) [Fisher 66, Box 87, Cox 00, Box 05]

Optimal experimental design: criterion based on a model

- Linear: information matrix (e.g., A-, D-optimal) [Fedorov 72, Atkinson 92]
- Nonlinear: intractable [Box 59, Ford 89, Chaloner 95]
- Decision-theoretic: maximize expected utility

  [Berger 85, Clyde 96, Santner 03, Müller 04, Amzal 06, Parmigiani 09]

Information-based Bayesian design: information gain objective [Lindley 56]

Scope of this work:

- Nonlinear and expensive physical models
- Continuous parameter, design, and data spaces of multiple dimensions
- Bayesian information objective with non-Gaussian distributions

Step 1: Define experimental goals

What is a good experiment?

Depends on the experimental goals

Examples:

Goal	Quantifiable objective
learn model parameter	minimize parameter uncertainty
make accurate prediction	minimize QoI bias and variance
custom performance assessment	minimize a loss function
⋮	⋮

Step 1: Define experimental goals

What is a good experiment?

Depends on the experimental goals

Examples:

Goal	Quantifiable objective
learn model parameter	minimize parameter uncertainty
make accurate prediction	minimize QoI bias and variance
custom performance assessment	minimize a loss function
⋮	⋮

Step 2: Formulate objective function

Characterize uncertainty using **Bayes' Theorem**:

$$\frac{\overbrace{p(\theta\quad)}^{prior}}{}$$

θ — parameters $\qquad y$ — noisy data

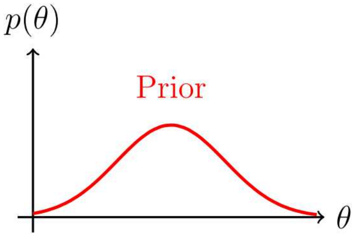

Step 2: Formulate objective function

Characterize uncertainty using **Bayes' Theorem**:

$$\frac{\overbrace{p(y|\theta\quad)}^{\textit{likelihood}}\;\overbrace{p(\theta\quad)}^{\textit{prior}}}{}$$

θ — parameters $\quad$ y — noisy data

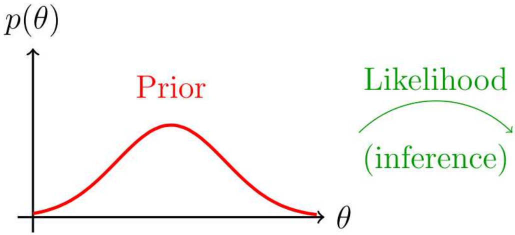

Step 2: Formulate objective function

Characterize uncertainty using **Bayes' Theorem**:

$$\overbrace{p(\theta|y)}^{\textit{posterior}} = \frac{\overbrace{p(y|\theta)}^{\textit{likelihood}}\,\overbrace{p(\theta)}^{\textit{prior}}}{\underbrace{p(y)}_{\textit{evidence}}}$$

θ — parameters $\qquad y$ — noisy data

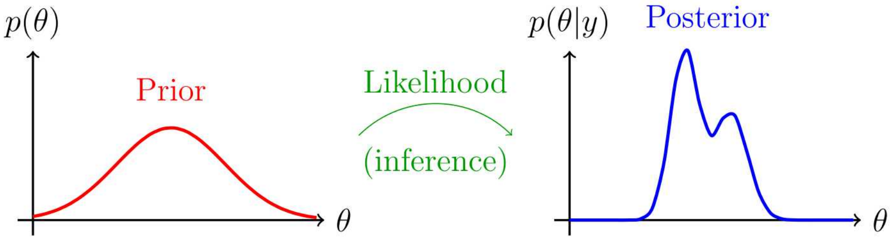

Step 2: Formulate objective function

Characterize uncertainty using **Bayes' Theorem**:

$$\overbrace{p(\theta|y,d)}^{\textit{posterior}} = \frac{\overbrace{p(y|\theta,d)}^{\textit{likelihood}}\,\overbrace{p(\theta|d)}^{\textit{prior}}}{\underbrace{p(y|d)}_{\textit{evidence}}}$$

θ — parameters $\quad$ y — noisy data $\quad$ d — design (experimental conditions)

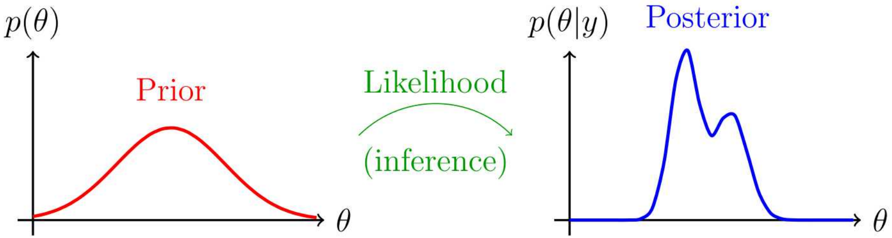

Step 2: Formulate objective function

Expected **Kullback-Leibler (KL)** divergence between posterior and prior ($\Longleftrightarrow$ **mutual information** between data and parameters)

$$\begin{aligned} U(d) &= \mathbb{E}_{y|d}\left[D_{\mathrm{KL}}(p(\theta|y,d)||p(\theta|d))\right] \\ &= \int_{\mathcal{Y}}\left[\int_{\Theta} \ln\left[\frac{p(\theta|y,d)}{p(\theta)}\right] p(\theta|y,d)\, d\theta\right] p(y|d)\, dy \\ &\approx \frac{1}{N}\sum_{i=1}^{N}\left\{\ln\left[p(y^{(i)}|\theta^{(i)},d)\right] - \ln\left[p(y^{(i)}|d)\right]\right\} \end{aligned}$$

$$p(y^{(i)}|d) \approx \frac{1}{M}\sum_{j=1}^{M} p(y^{(i)}|\theta^{(i,j)},d) \qquad \text{(nested Monte Carlo!)}$$

U — expected utility

θ — parameters of interest
y — noisy data
d — design variables

Step 3: Approximate objective function numerically

Expected **Kullback-Leibler (KL)** divergence between posterior and prior ($\Longleftrightarrow$ **mutual information** between data and parameters)

$$
\begin{aligned}
U(d) &= \mathbb{E}_{y|d}\left[D_{\mathrm{KL}}(p(\theta|y,d)||p(\theta|d))\right] \\
&= \int_{\mathcal{Y}} \left[\int_{\Theta} \ln\left[\frac{p(\theta|y,d)}{p(\theta)}\right] p(\theta|y,d)\,d\theta\right] p(y|d)\,dy \\
&\approx \frac{1}{N}\sum_{i=1}^{N}\left\{\ln\left[p(y^{(i)}|\theta^{(i)},d)\right] - \ln\left[p(y^{(i)}|d)\right]\right\}
\end{aligned}
$$

$$
p(y^{(i)}|d) \approx \frac{1}{M}\sum_{j=1}^{M} p(y^{(i)}|\theta^{(i,j)},d) \qquad \text{(nested Monte Carlo!)}
$$

U — expected utility

θ — parameters of interest
y — noisy data
d — design variables

Step 4: Perform stochastic optimization

Find optimal design: $$d^* = \underset{d \in \mathcal{D}}{\operatorname{argmax}}\ U(d)$$

$\Rightarrow$ **noisy** objective function due to Monte Carlo

Robbins-Monro stochastic approximation [Robbins 51]

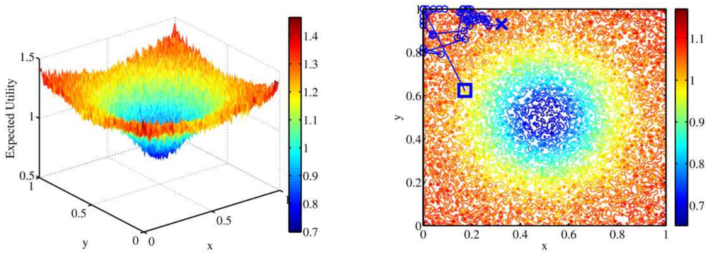

- Needs to derive gradient
- Can obtain easily from surrogate models

Accelerate using polynomial chaos (PC) surrogate models

A random variable can be expanded:

$$g(\theta(\xi)) = \sum_{\beta \in \mathcal{J}} c_\beta \Psi_\beta(\xi)$$

- c_β: coefficients (β polynomial order multi-index)
- ξ: reference random vector (e.g., uniform, Gaussian)
- Ψ_β: multivariate orthonormal polynomial (e.g., Legendre, Hermite)

Non-intrusive approach to compute coefficients

$$c_\beta = \langle G \Psi_\beta \rangle = \int G(\theta(\xi), d(\xi)) \Psi_\beta(\xi) p(\xi)\, d\xi$$

Integrate using adaptive sparse quadrature [Gerstner 98, Barthelmann 00, Gerstner 03]

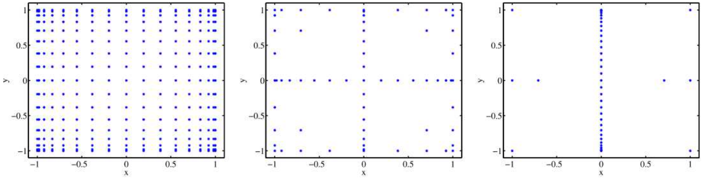

Design of shock tube experiment for combustion kinetics

H_2-O_2 homogeneous ignition, no transport, constant pressure, adiabatic

19-reaction mechanism [Yetter 91], learn two Arrhenius kinetic parameters

- $\ln A$ of $R1$: $H + O_2 \Longleftrightarrow O + OH$
- E_a of $R3$: $H_2 + OH \Longleftrightarrow H_2O + H$

Measurement: ignition delay time
Design variables: initial T and ϕ (equivalence ratio)

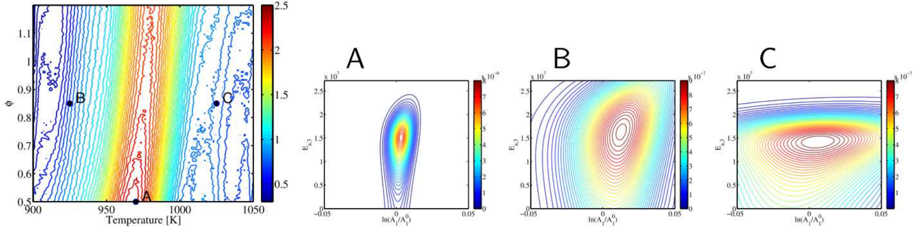

Sensor placement for source inversion

2D scalar diffusion PDE in a square domain with a time-limited source

Goal: infer uncertain source location

Measurements: concentration level at 5 equally spaced sample times
Design variables: sensor coordinates

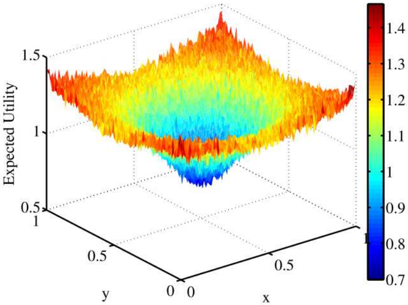

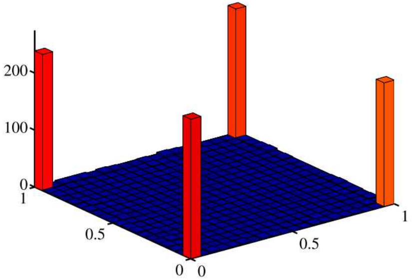

Sensor placement for source inversion: 1 sensor

Posterior PDFs: $x_{src} = (0.39, 0.52)$

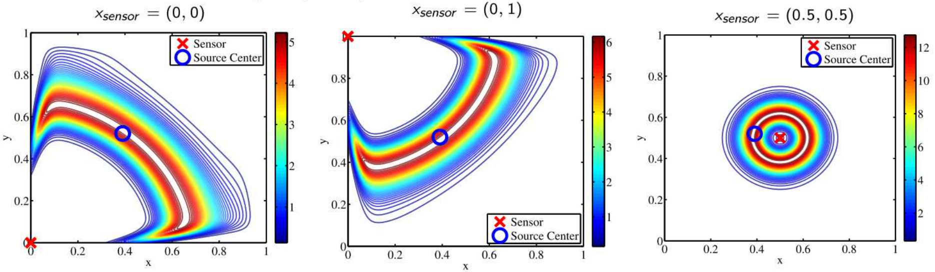

Sensor placement for source inversion: 1 sensor

Posterior PDFs: $x_{src} = (0.39, 0.52)$ (row 1) and $x_{src} = (0.09, 0.22)$ (row 2):

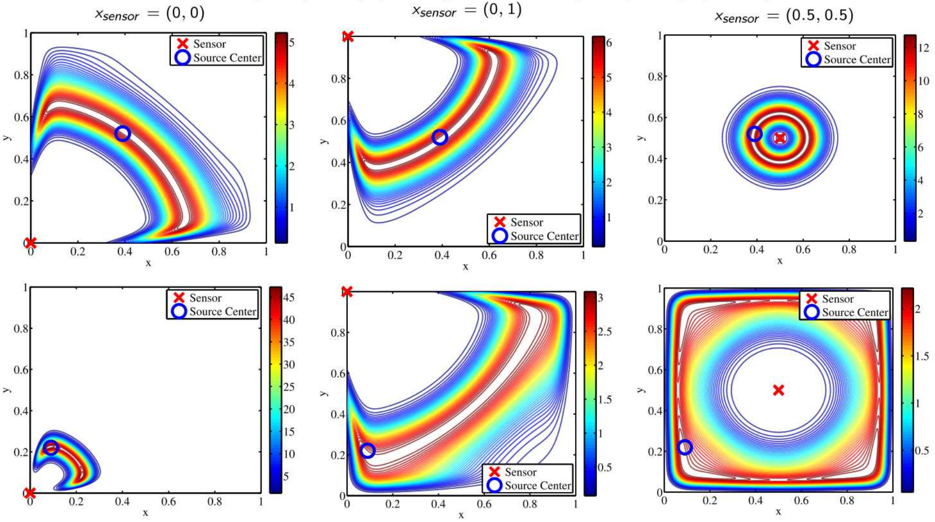

Sensor placement for source inversion: 2 sensors

Posterior PDFs: $x_{src} = (0.39, 0.52)$ (row 1) and $x_{src} = (0.09, 0.22)$ (row 2):

$x_{sensors} = (0,0), (0,0)$ $x_{sensors} = (0,0), (0,1)$ $x_{sensor} = (0.5, 0.5), (0.5, 0.5)$

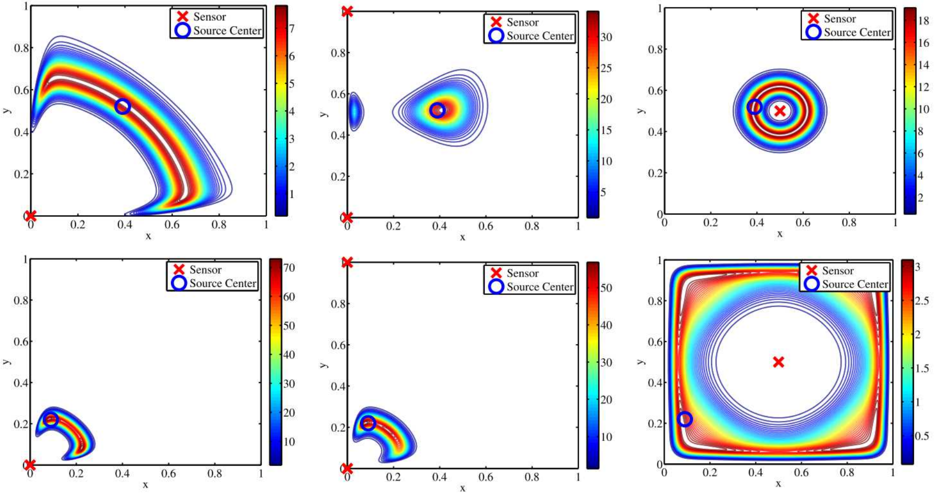

Sensor placement for source inversion: 2 sensors

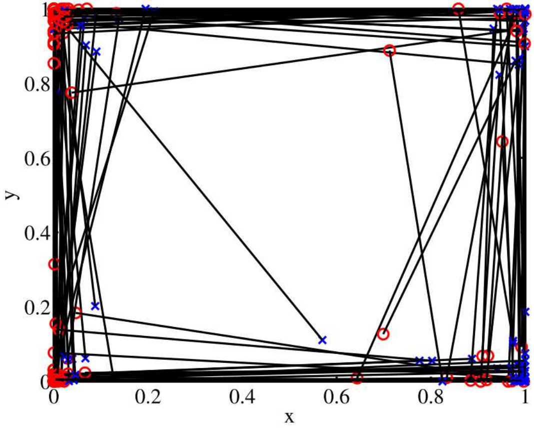

- 100 independent optimization runs
- Each pair of circle-cross represents the two sensor locations

Outline

Planning measurements: batch (non-sequential) design

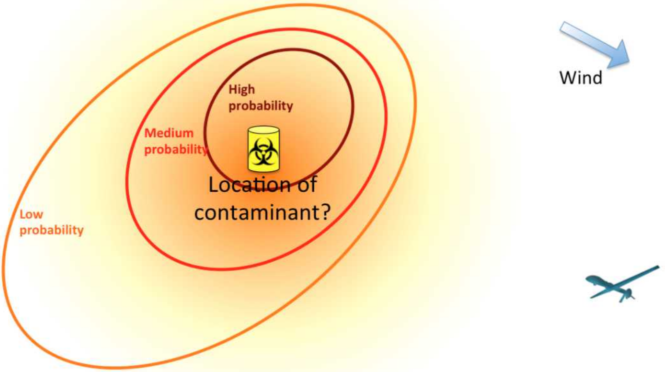

Planning measurements: sequential design

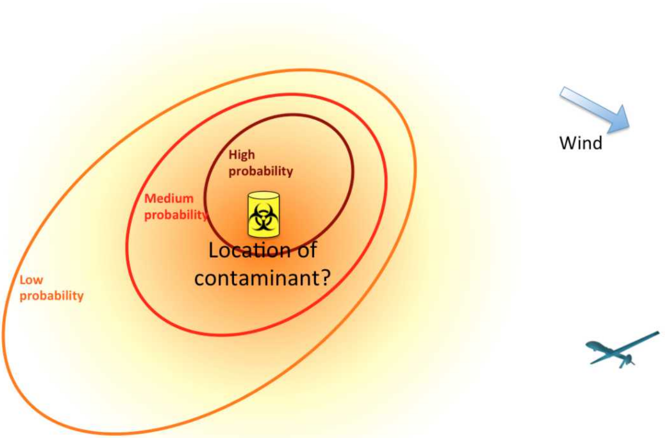

Sequential design is less developed than batch design

Greedy (myopic) design:

- Repeated application of batch design [Solonen 12, Drovandi 14, Kim 14]
- *Not optimal*

Dynamic programming:

- Fully optimal description (has 1. feedback, and 2. lookahead) [Müller 07, Von Toussaint 11]
- Thus far limited to discrete variables, special problem and solution structures [Carlin, Bradley 98, Berry 02, Brockwell 03, Christen 03, Murphy 03, Wathen 06]
- Only simple objectives

Aim of this work:

Develop a mathematical framework and numerical tools to find optimal sequential experimental designs in a computationally feasible manner

Core components of general sequential design formulation

Experiment: $k = 0, \ldots, N-1$, total N experiments; $N < \infty$

State: $x_k = [x_{k,b}, x_{k,p}]$ all information needed for optimal future designs

- *Belief state*: $x_{k,b}$ current state of uncertainty
- *Physical state*: $x_{k,p}$ deterministic design-relevant variables

Design: $d_k = \mu_k(x_k)$
seek good *policy* $\pi \equiv \{\mu_0, \mu_1, \ldots, \mu_{N-1}\}$

Noisy data: y_k distributed according to likelihood $f(y_k|\theta, d_k)$
(e.g., $y_k = G(\theta, d_k) + \epsilon$, with ϵ Gaussian)

System dynamics: $x_{k+1} = \mathcal{F}(x_k, y_k, d_k)$ state evolution

Sequential design exhibits a closed-loop behavior

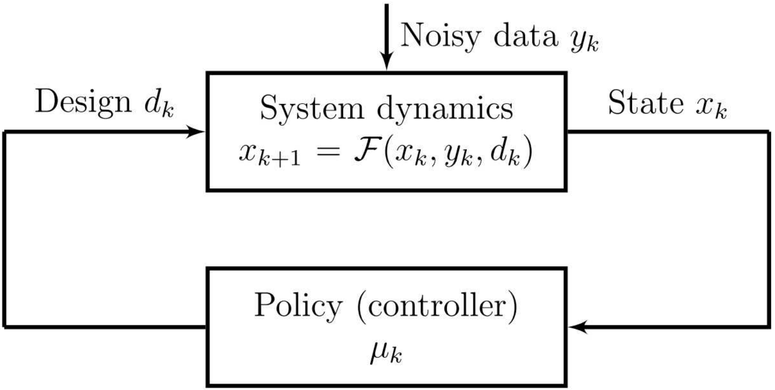

The sOED problem: find optimal policy that maximizes the expected total reward

Stage reward: $g_k(x_k, y_k, d_k)$ **Terminal reward:** $g_N(x_N)$

The sequential optimal experimental design (sOED) problem:

Find π^* where

$$\pi^* = \underset{\pi=\{\mu_0,\dots,\mu_{N-1}\}}{\text{argmax}} \ \mathbb{E}_{y_0,\dots,y_{N-1}|\pi} \left[\sum_{k=0}^{N-1} g_k\left(x_k, y_k, \mu_k(x_k)\right) + g_N\left(x_N\right) \right]$$

$$\text{s.t.} \quad x_{k+1} = \mathcal{F}(x_k, y_k, d_k), \forall k$$
$$\mu_k(x_k) \in \mathcal{D}_k, \forall x_k, k$$

Difficult to solve directly, involves optimization of a functional

The sOED problem in dynamic programming (DP) form

Re-express using Bellman's Principle of Optimality [Bellman 53]

Dynamic programming form (Bellman equations): (e.g., [Bertsekas 05])

$$J_k(x_k) = \max_{d_k \in \mathcal{D}_k} \mathbb{E}_{y_k|x_k,d_k}\left[g_k(x_k, d_k, y_k) + J_{k+1}\left(\mathcal{F}(x_k, d_k, y_k)\right)\right]$$

$$J_N(x_N) = g_N(x_N)$$

$k = 0, \ldots, N-1$; $J_k(x_k)$ are value functions

- A set of smaller tail subproblems
- Optimal policy functions implicitly in argmax: $d_k^* = \mu_k^*(x_k)$
- "Curse of dimensionality": exponential scenario growth from recursion
- Large body of approximate methods: *approximate dynamic programming*

  (e.g., [Bertsekas 96, Kaelbling 96, Sutton 98, Powell 11])

Approximate value iteration with adaptive exploitation

Approximate value iteration (backward induction with regression):

$$J_{k+1}\left(\mathcal{F}_k(x_k, d_k, y_k)\right)$$

Start with $\tilde{J}_N(x_N) \equiv g_N(x_N)$, and proceed backwards $k = N-1, \ldots, 1$

Approximate value iteration with adaptive exploitation

Approximate value iteration (backward induction with regression):

$$J_k(x_k) = \max_{d_k \in \mathcal{D}_k} \mathbb{E}_{y_k|x_k,d_k}[g_k(x_k, d_k, y_k) + J_{k+1}(\mathcal{F}_k(x_k, d_k, y_k))]$$

Start with $\tilde{J}_N(x_N) \equiv g_N(x_N)$, and proceed backwards $k = N-1, \ldots, 1$

Approximate value iteration with adaptive exploitation

Approximate value iteration (backward induction with regression):

$$\tilde{J}_k(x_k) = \mathcal{P}\left\{ \max_{d_k \in \mathcal{D}_k} \mathbb{E}_{y_k|x_k,d_k}\left[g_k(x_k, d_k, y_k) + J_{k+1}\left(\mathcal{F}_k(x_k, d_k, y_k)\right)\right] \right\}$$

Start with $\tilde{J}_N(x_N) \equiv g_N(x_N)$, and proceed backwards $k = N-1, \ldots, 1$

$\mathcal{P}$: regression operator

Approximate value iteration with adaptive exploitation

Approximate value iteration (backward induction with regression):

$$\tilde{J}_k(x_k) = \mathcal{P}\left\{\max_{d_k \in \mathcal{D}_k} \mathbb{E}_{y_k|x_k,d_k}[g_k(x_k, d_k, y_k) + \tilde{J}_{k+1}(\mathcal{F}_k(x_k, d_k, y_k))]\right\}$$

Start with $\tilde{J}_N(x_N) \equiv g_N(x_N)$, and proceed backwards $k = N-1, \ldots, 1$

$\mathcal{P}$: regression operator

Approximate value iteration with adaptive exploitation

Approximate value iteration (backward induction with regression):

$$\tilde{J}_k(x_k) = \mathcal{P}\left\{ \max_{d_k \in \mathcal{D}_k} \mathbb{E}_{y_k|x_k,d_k}\left[g_k(x_k, d_k, y_k) + \tilde{J}_{k+1}\left(\mathcal{F}_k(x_k, d_k, y_k)\right)\right] \right\}$$

Start with $\tilde{J}_N(x_N) \equiv g_N(x_N)$, and proceed backwards $k = N-1, \ldots, 1$

$\mathcal{P}$: regression operator, using **exploration** and **exploitation** samples

Introduce **adaptive update** to the exploitation policy:

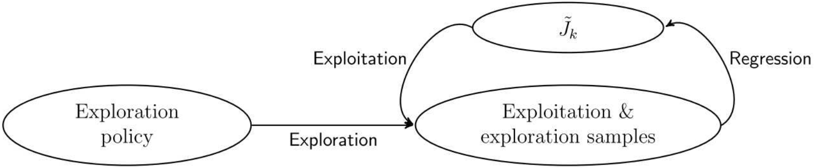

Final algorithm for overall sequential OED

1. **Set parameters**
2. **Initial exploration**
3. **Precompute objects needed for repeated Bayesian inference**
4. Iterate to refine . . .
    a. **Exploration**
    b. **Exploitation**
    c. **Approximate value iteration**
5. **Extract final policy**

1D source inversion problem: problem setup

- 2 experiments
- Source location $\theta \sim \mathcal{N}(0, 2^2)$ Sensor starting location: 5.5
- Strong wind blows to the right after first experiment
- Quadratic movement penalty

An example scenario:

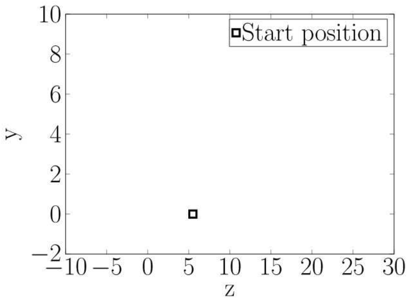

physical state and plume

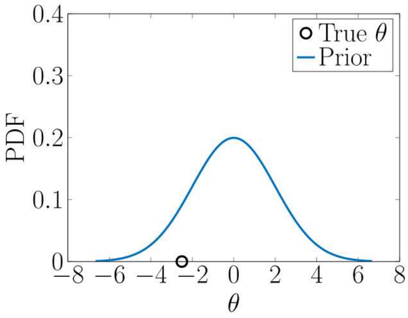

belief state

1D source inversion problem: problem setup

- 2 experiments
- Source location $\theta \sim \mathcal{N}(0, 2^2)$ Sensor starting location: 5.5
- Strong wind blows to the right after first experiment
- Quadratic movement penalty

An example scenario:

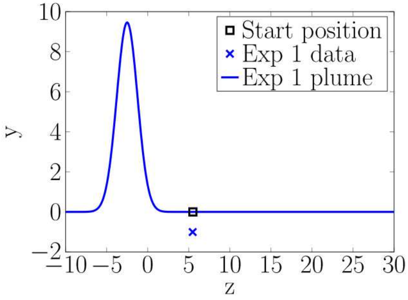

physical state and plume

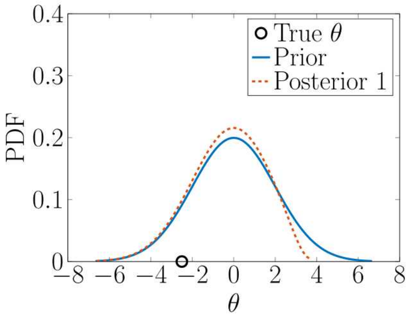

belief state

1D source inversion problem: problem setup

- 2 experiments
- Source location $\theta \sim \mathcal{N}(0, 2^2)$ Sensor starting location: 5.5
- Strong wind blows to the right after first experiment
- Quadratic movement penalty

An example scenario:

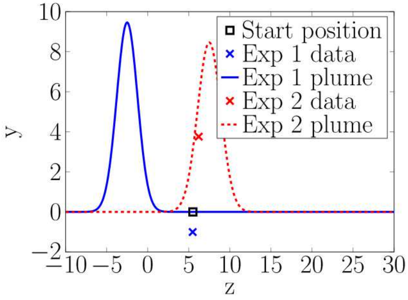

physical state and plume

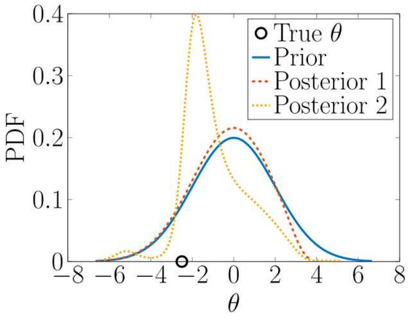

belief state

1D source inversion problem: greedy design vs. sOED

Greedy design does not account for future wind conditions

Expected reward: greedy (0.07), sOED (0.15)

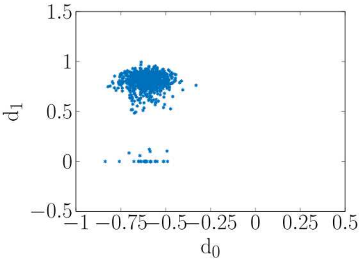

greedy design

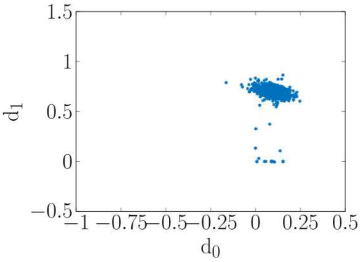

sOED

1D source inversion problem: batch design vs. sOED

A more precise instrument available only if prior variance < 3

Batch design does not have feedback

Expected reward: batch (0.15), sOED (0.26)

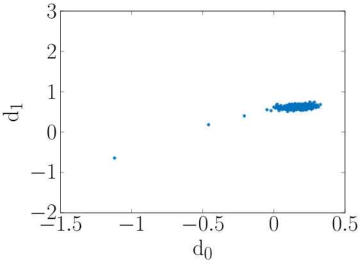

batch design

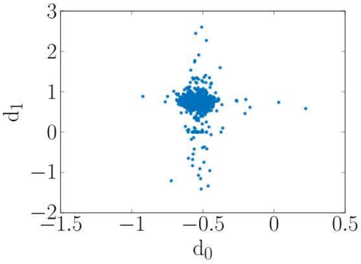

sOED

2D source inversion problem: trajectory example #1

An example scenario:

physical state and plume

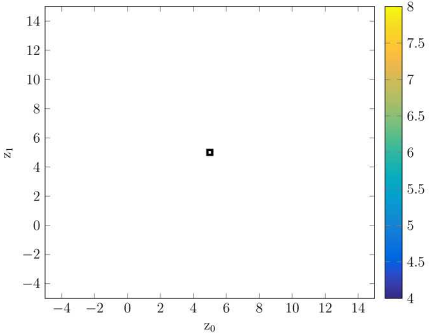

belief states

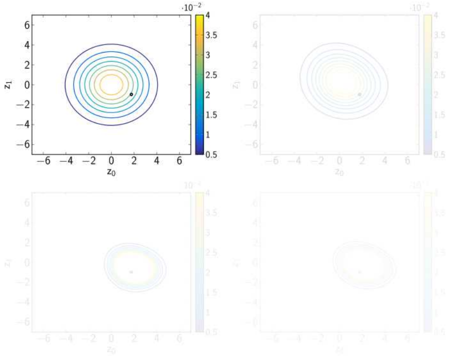

2D source inversion problem: trajectory example #1

An example scenario:

physical state and plume

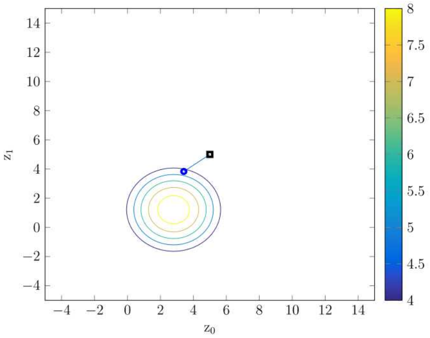

belief states

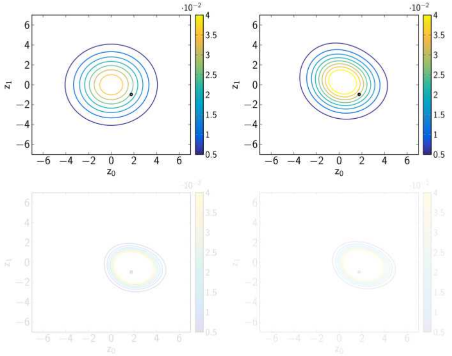

An example scenario:

physical state and plume

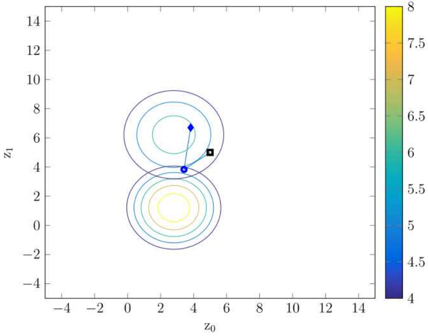

belief states

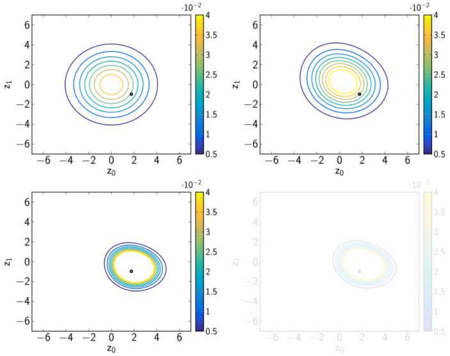

An example scenario:

physical state and plume

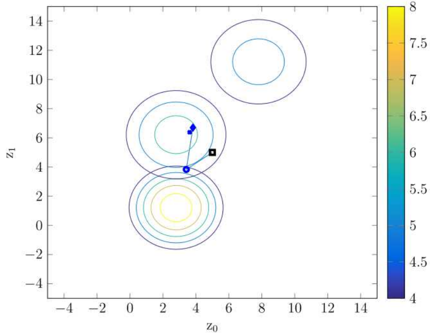

belief states

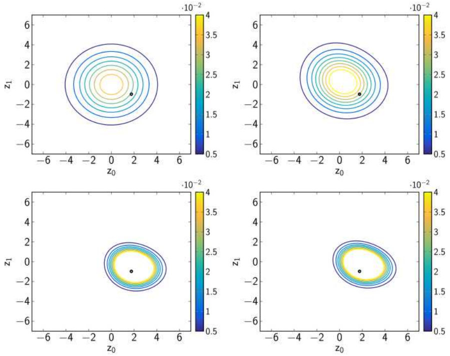

Another example scenario:

physical state and plume

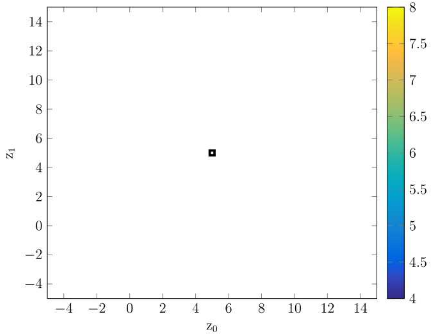

belief states

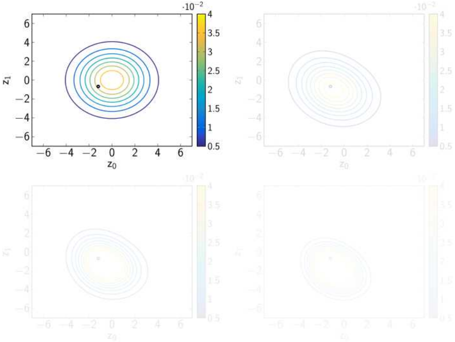

Another example scenario:

physical state and plume

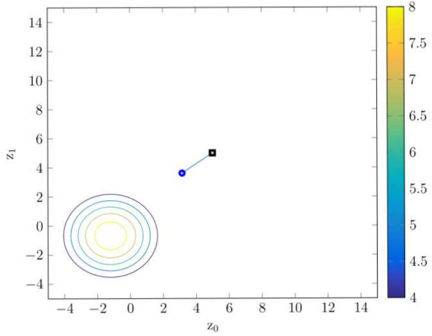

belief states

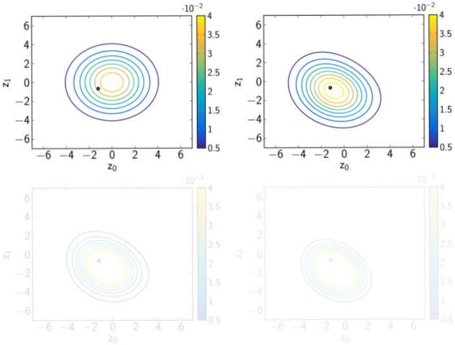

2D source inversion problem: trajectory example #2

Another example scenario:

physical state and plume

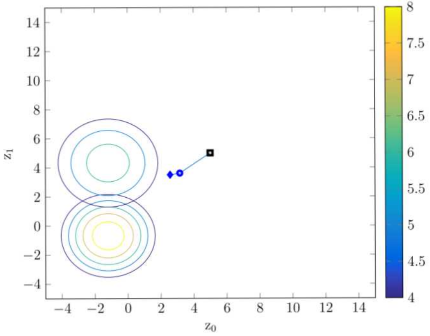

belief states

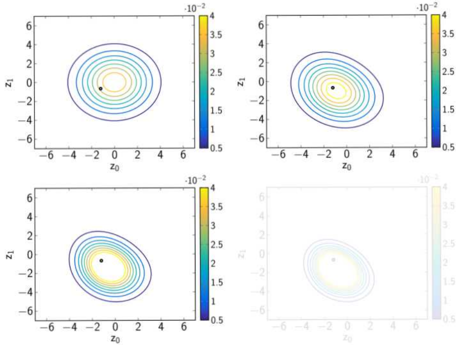

2D source inversion problem: trajectory example #2

Another example scenario:

physical state and plume

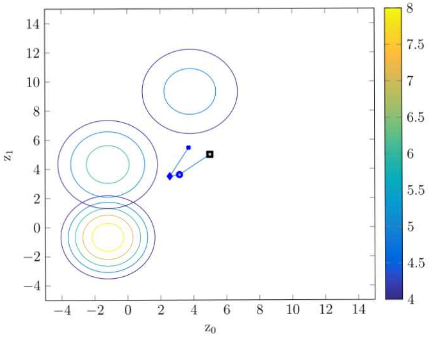

belief states

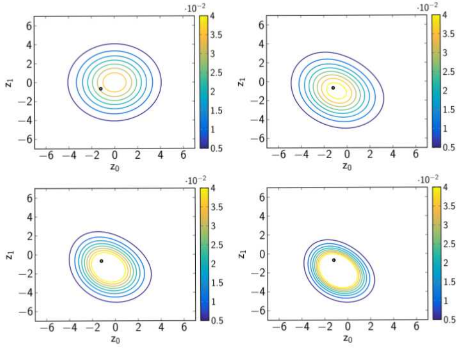

Outline

Summary

- Enabled tractable computation of optimal experimental designs targeting nonlinear and computationally intensive models
- Formulated the optimal sequential experimental design rigorously (with feedback and lookahead) using dynamic programming
- Integrated and advanced new computational approaches to optimal experimental design:
  - stochastic optimization
  - polynomial chaos surrogate models
  - adaptive sparse quadrature
- Developed new approximate dynamic programming technique: adaptive regression for approximate value iteration
- Demonstrated design of combustion experiments for learning rate parameters, and sensor placement (both batch and sequential) for source inversion

Acknowledgment

- Sandia National Laboratories[1]
- Massachusetts Institute of Technology
- Defense Advanced Research Projects Agency (DARPA)
- Air Force Office of Scientific Research (AFOSR) Computational Mathematics Program

[1] Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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-NA-0003525. The views expressed in the presentation do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References I

Billy Amzal, Frédéric Y. Bois, Eric Parent & Christian P. Robert.
Bayesian-Optimal Design via Interacting Particle Systems.
Journal of the American Statistical Association, vol. 101, no. 474, pages 773–785, 2006.

Anthony C. Atkinson & Alexander N. Donev.
Optimum Experimental Designs.
Oxford University Press, New York, NY, 1992.

Volker Barthelmann, Erich Novak & Klaus Ritter.
High dimensional polynomial interpolation on sparse grids.
Advances in Computational Mathematics, vol. 12, pages 273–288, 2000.

Richard Bellman.
Bottleneck Problems and Dynamic Programming.
Proceedings of the National Academy of Sciences of the United States of America, vol. 39, no. 9, pages 947–951, 1953.

James O. Berger.
Statistical Decision Theory and Bayesian Analysis.
Springer New York, New York, NY, 1985.

Donald A. Berry, Peter Müller, Andy P. Grieve, Michael Smith, Tom Parke, Richard Blazek, Neil Mitchard & Michael Krams.
Adaptive Bayesian Designs for Dose-Ranging Drug Trials.
In Case studies in Bayesian statistics, pages 99–181. Springer New York, New York, NY, 2002.

Dimitri P. Bertsekas & John N. Tsitsiklis.
Neuro-Dynamic Programming.
Athena Scientific, Belmont, MA, 1996.

References II

Dimitri P. Bertsekas.
Dynamic Programming and Optimal Control, Vol. 1.
Athena Scientific, Belmont, MA, 2005.

George Edward Pelham Box & H. L. Lucas.
Design of Experiments in Non-Linear Situations.
Biometrika, vol. 46, no. 1-2, pages 77–90, 1959.

George Edward Pelham Box & Norman Richard Draper.
Empirical Model-Building and Response Surfaces.
John Wiley & Sons, Inc., Hoboken, NJ, 1987.

George Edward Pelham Box, J. Stuart Hunter & William G. Hunter.
Statistics for Experimenters: Design, Innovation and Discovery.
John Wiley & Sons, Inc., Hoboken, NJ, 2nd edition, 2005.

Anthony E. Brockwell & Joseph B. Kadane.
A Gridding Method for Bayesian Sequential Decision Problems.
Journal of Computational and Graphical Statistics, vol. 12, no. 3, pages 566–584, 2003.

P. Carlin, Bradley, Joseph B. Kadane & Alan E. Gelfand.
Approaches for Optimal Sequential Decision Analysis in Clinical Trials.
Biometrics, vol. 54, no. 3, pages 964–975, 1998.

Kathryn Chaloner & Isabella Verdinelli.
Bayesian Experimental Design: A Review.
Statistical Science, vol. 10, no. 3, pages 273–304, 1995.

References III

J. Andrés Christen & Miguel Nakamura.
Sequential Stopping Rules for Species Accumulation.
Journal of Agricultural, Biological & Environmental Statistics, vol. 8, no. 2, pages 184–195, 2003.

Merlise A. Clyde, Peter Müller & Giovanni Parmigiani.
Exploring Expected Utility Surfaces by Markov Chains.
Rapport technique, Duke University, 1996.

D. R. Cox & N. Reid.
The Theory of the Design of Experiments.
Chapman & Hall/CRC, Boca Raton, FL, 2000.

Christopher C. Drovandi, James M. McGree & Anthony N. Pettitt.
A Sequential Monte Carlo Algorithm to Incorporate Model Uncertainty in Bayesian Sequential Design.
Journal of Computational and Graphical Statistics, vol. 23, no. 1, pages 3–24, 2014.

V. V. Fedorov.
Theory of Optimal Experiments.
Academic Press, New York, NY, 1972.

Ronald Aylmer Fisher.
The Design of Experiments.
Oliver & Boyd, Edinburgh, United Kingdom, 8th edition, 1966.

Ian Ford, D. M. Titterington & Christos P. Kitsos.
Recent Advances in Nonlinear Experimental Design.
Technometrics, vol. 31, no. 1, pages 49–60, 1989.

References IV

Thomas Gerstner & Michael Griebel.
Numerical integration using sparse grids.
Numerical Algorithms, vol. 18, pages 209–232, 1998.

T. Gerstner & M. Griebel.
Dimension-Adaptive Tensor-Product Quadrature.
Computing, vol. 71, no. 1, pages 65–87, 2003.

Leslie Pack Kaelbling, Michael L. Littman & Andrew W. Moore.
Reinforcement Learning: A Survey.
Journal of Artificial Intelligence Research, vol. 4, pages 237–285, 1996.

Woojae Kim, Mark A. Pitt, Zhong-Lin Lu, Mark Steyvers & Jay I. Myung.
A Hierarchical Adaptive Approach to Optimal Experimental Design.
Neural Computation, vol. 26, pages 2565–2492, 2014.

Dennis V. Lindley.
On a Measure of the Information Provided by an Experiment.
The Annals of Mathematical Statistics, vol. 27, no. 4, pages 986–1005, 1956.

Peter Müller, Bruno Sansó & Maria De Iorio.
Optimal Bayesian Design by Inhomogeneous Markov Chain Simulation.
Journal of the American Statistical Association, vol. 99, no. 467, pages 788–798, 2004.

Peter Müller, Donald A. Berry, Andy P. Grieve, Michael Smith & Michael Krams.
Simulation-based sequential Bayesian design.
Journal of Statistical Planning and Inference, vol. 137, no. 10, pages 3140–3150, 2007.

References V

S. A. Murphy.
Optimal dynamic treatment regimes.
Journal of the Royal Statistical Society: Series B (Statistical Methodology), vol. 65, no. 2, pages 331–355, 2003.

Giovanni Parmigiani & Lurdes Y. T. Inoue.
Decision Theory: Principles and Approaches.
John Wiley & Sons, Inc., West Sussex, United Kingdom, 2009.

Warren B. Powell.
Approximate Dynamic Programming: Solving the Curses of Dimensionality.
John Wiley & Sons, Inc., Hoboken, NJ, 2nd edition, 2011.

Herbert Robbins & Sutton Monro.
A Stochastic Approximation Method.
The Annals of Mathematical Statistics, vol. 22, no. 3, pages 400–407, 1951.

Thomas J. Santner, Brian J. Williams & William I. Notz.
The Design and Analysis of Computer Experiments.
Springer New York, New York, NY, 2003.

Antti Solonen, Heikki Haario & Marko Laine.
Simulation-Based Optimal Design Using a Response Variance Criterion.
Journal of Computational and Graphical Statistics, vol. 21, no. 1, pages 234–252, 2012.

Richard S. Sutton & Andrew G. Barto.
Reinforcement Learning: An Introduction.
The MIT Press, Cambridge, MA, 1998.

References VI

Udo Von Toussaint.
Bayesian inference in physics.
Reviews of Modern Physics, vol. 83, pages 943–999, 2011.

J. Kyle Wathen & J. Andrés Christen.
Implementation of Backward Induction for Sequentially Adaptive Clinical Trials.
Journal of Computational and Graphical Statistics, vol. 15, no. 2, pages 398–413, 2006.

R. A. Yetter, F. L. Dryer & H. Rabitz.
A Comprehensive Reaction Mechanism For Carbon Monoxide/Hydrogen/Oxygen Kinetics.
Combustion Science and Technology, vol. 79, no. 1-3, pages 97–128, 1991.