

SAND2017-7756C

Sequential Optimal Experimental Design via Stochastic Control

Xun (Ryan) Huan & Youssef M. Marzouk

Sandia National Laboratories
Massachusetts Institute of Technology

July 10, 2017

Some experiments are more useful than others

Experiments are:

- Expensive
- Time-consuming
- Delicate to perform

Experimental design helps address:

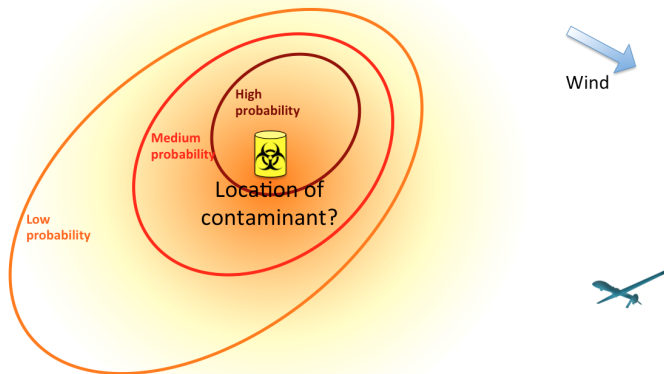
- Under what **conditions** to perform the experiment?
- **What** to measure?
- **Where** to measure?
- **When** to measure?

Experimental design is useful in many different applications

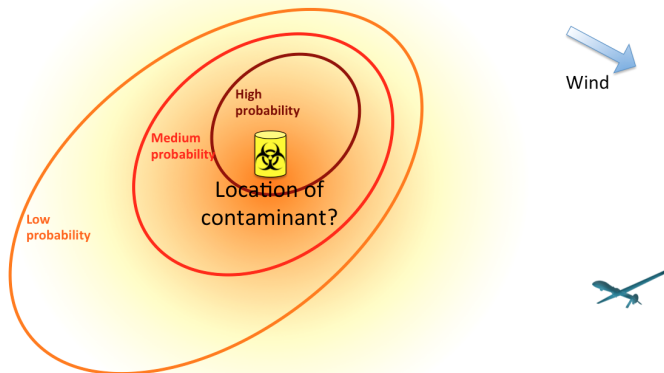


(Sources: left—Argonne National Labs; right—www.weather.com)

Planning measurements: batch (open-loop) design



Planning measurements: sequential (closed-loop) design



Sequential experimental design is relatively less developed

Batch experimental design:

- Linear: Fisher information matrix (e.g., A -, D -optimal)
 - Nonlinear: advances beyond linearization and Gaussianization
 - Information-based experimental design [Lindley 56]
-

Greedy (myopic) design:

- Repeated application of batch design [Solonen 12, Drovandi 14, Kim 14]
- But it is *not optimal*

Dynamic programming:

- Fully optimal description (has 1. feedback, 2. forward looking)
- Very difficult to tackle generally
- Thus far limited to discrete variables, special problem and solution structures, simple objectives [Carlin, Bradley 98, Brockwell 03, Berry 02]

Objectives and scope

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

Illustrate connections to *stochastic optimal control*

Scope:

- **Finite** number of experiments
- **Nonlinear** and expensive models
- **Continuous** parameter, design, and data spaces of multiple dimensions
- **Bayesian** treatment of uncertainty
- **Non-Gaussian** distributions
- **Information measure** objective (design for parameter inference)

Core components of general sequential design formulation

Experiment: $k = 0, \dots, N - 1$, 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

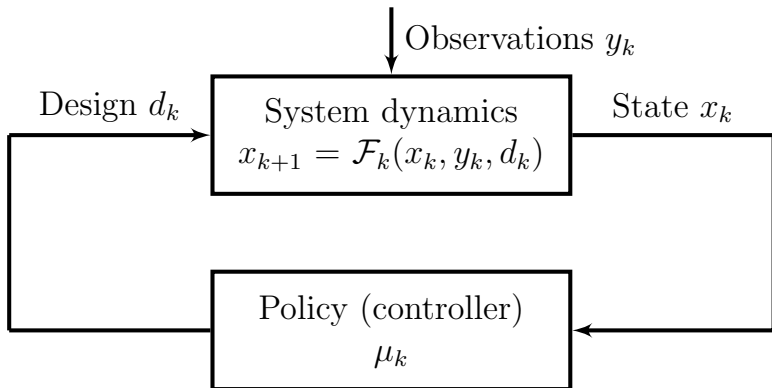
Design: $d_k = \mu_k(x_k)$

seek good *policy* $\pi \equiv \{\mu_0, \mu_1, \dots, \mu_{N-1}\}$

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

Sequential design corresp. to closed-loop feedback control



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

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

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

$$\mu_k(x_k) \in \mathcal{D}_k, \forall x_k, k$$

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])

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


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

$k = 0, \dots, N - 1$; $J_k(x_k)$ 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])

Batch (open-loop) design is a special case of the sOED problem, and thus suboptimal

- Has no feedback
- Designs all experiments concurrently as a batch
- Finds optimal *designs* (vectors) rather than a policy



$$\{d_0^*, \dots, d_{N-1}^*\} = \underset{d_0, \dots, d_{N-1}}{\operatorname{argmax}} \mathbb{E}_{y_0, \dots, y_{N-1} | d_{0:N-1}} \left[\sum_{k=0}^{N-1} g_k(x_k, y_k, d_k) + g_N(x_N) \right]$$

Greedy (myopic) design is a special case of the sOED problem (DP form), and thus suboptimal

- Uses feedback
- Considers the *next experiment* only
- Has no future effects

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

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

subject to $x_{k+1} = \mathcal{F}_k(x_k, y_k, d_k)$

Sequential Bayesian inference


For the k -th experiment:

$$\underbrace{f(\theta|y_k, d_k, I_k)}_{\text{posterior}} = \frac{\underbrace{f(y_k|\theta, d_k, I_k)}_{\text{likelihood}} \underbrace{f(\theta|I_k)}_{\text{prior}}}{\underbrace{f(y_k|d_k, I_k)}_{\text{evidence}}}$$

θ — parameters to infer

I_k — information from previous experiments, $I_k \equiv \{d_0, y_0, \dots, d_{k-1}, y_{k-1}\}$

Conceptually: belief state is posterior random variable $x_{k,b} = \theta|I_k$



Information gain objective for parameter inference

We choose to use total information gain at end of all experiments
(Kullback-Leibler (KL) divergence from final posterior to prior)

$g_k(x_k, d_k, y_k)$ = reflects experimental cost

$$g_N(x_N) = D_{KL}(f(x_{N,b}) || f(x_{0,b})) = \int_{\mathcal{H}} f(x_{N,b}) \ln \left[\frac{f(x_{N,b})}{f(x_{0,b})} \right] d\theta$$

Corresponding system dynamics:

- *Belief state*: Bayes' Theorem

$$f(x_{k+1,b}) = \frac{f(y_k | \theta, d_k, l_k) f(x_{k,b})}{f(y_k | d_k, l_k)}$$

- *Physical state*: physical process

Represent a policy using one-step lookahead form

One-step lookahead policy representation: (e.g., [Bertsekas 05])

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

Approximate value functions using **linear architecture**:

$$\tilde{J}_k(x_k) = r_k^\top \phi_k(x_k)$$

ϕ_k features (selected from heuristics), r_k weights

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 $k = N - 1, \dots, 1$

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

Belief state representation

Conceptually: belief state is posterior random variable

How to numerically represent it . . .

- for general non-Gaussian continuous random variables
- in a finite-dimensional manner
- to easily perform Bayesian inference repeatedly in evaluating


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

Traditional approaches:

- Gaussian approximation and model linearization
- Gridding or functional approximation of its PDF or CDF
- Non-parametrics (with particle filter, MCMC)

Often expensive and some do not scale well to multiple dimensions. We seek an approach that can quickly perform **many** Bayesian inferences.

Transport map transforms *distributions* (e.g., [Villani 08])



- $\xi \sim$ reference distribution, $z \sim$ target distribution
- Equivalence in distribution $\xi \stackrel{i.d.}{=} T(z)$
- Knothe-Rosenblatt (KR) map: defined by conditional distributions, is triangular and monotone, exists and is unique [Rosenblatt 52, Carlier 10]
- Easy to construct from samples: convex optimization problem
- Target *joint* distribution for fast approx Bayesian inference

Final algorithm

- 1 **Set parameters**
- 2 **Initial exploration**
- 3 **Make joint map**
- 4 Iterate to refine . . .
 - (a) **Exploration**
 - (b) **Exploitation**
 - (c) **Approximate value iteration**
- 5 **Extract final policy parameterization**

1D source inversion problem: problem settings

$$y_k = \frac{s}{\sqrt{2\pi}2\sqrt{(0.3 + Dt)}} \exp\left(-\frac{\|\theta + d_w(t) - z_{k+1}\|^2}{2(4)(0.3 + Dt)}\right) + \epsilon_k$$


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

1D source inversion problem: case 1


advantages of sOED over greedy design

Greedy (myopic) design does not account for future wind conditions

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



greedy design



sOED


1D source inversion problem: case 2

advantages of sOED over batch design


A more precise instrument available only if prior variance < 3

Batch (open-loop) design does not have feedback

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



batch design




sOED

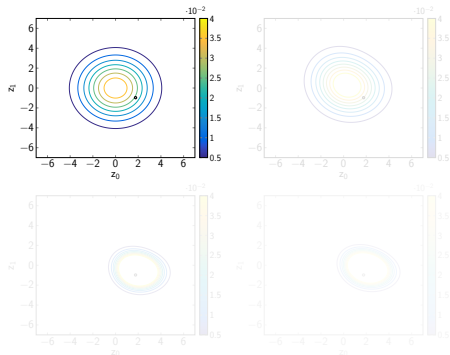
2D source inversion problem: trajectory example #1

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




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




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




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



Conclusions

- Formulated the sequential optimal experimental design (sOED) problem rigorously (has 1. feedback, 2. forward looking)
 - common principles with stochastic optimal control
 - best achievable in theory
 - difficult to solve in practice(future work: tradeoffs)
- Developed numerical methods to find approximate solutions to the sOED problem in a computationally-feasible manner, using
 - approximate dynamic programming
 - transport maps representing joint distributions
- Demonstrated computational effectiveness on realistic applications

Acknowledgment

- Sandia National Laboratories¹
- Massachusetts Institute of Technology
- Defense Advanced Research Projects Agency (DARPA)
- Air Force Office of Scientific Research (AFOSR)
Computational Mathematics Program

¹Sandia 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 Energys National Nuclear Security Administration under contract DE-NA-0003525.

References I



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.



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.



Dimitri P. Bertsekas.

Dynamic Programming and Optimal Control, Vol. 1.

Athena Scientific, Belmont, MA, 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.



Guillaume Carlier, Alfred Galichon & Filippo Santambrogio.

From Knothe's Transport to Brenier's Map and a Continuation Method for Optimal Transport.

SIAM Journal on Mathematical Analysis (Society for Industrial and Applied Mathematics), vol. 41, no. 6, pages 2554–2576, 2010.

References II



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.



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.



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.



Warren B. Powell.

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



Murray Rosenblatt.

Remarks on a Multivariate Transformation.
The Annals of Mathematical Statistics, vol. 23, no. 3, pages 470–472, 1952.

References III



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.



Cédric Villani.

Optimal Transport: Old and New.

Springer-Verlag Berlin Heidelberg, Berlin, Germany, 2008.