LA-UR-16-21255
Approved for public release; distribution is unlimited.

Title: Tractable Experiment Design via Mathematical Surrogates

Author(s): Williams, Brian J.

Intended for: Report

Issued: 2016-02-29

Tractable Experiment Design via Mathematical Surrogates

Brian Williams

Statistical Sciences Group, Los Alamos National Laboratory

Abstract

This presentation summarizes the development and implementation of quantitative design criteria motivated by targeted inference objectives for identifying new, potentially expensive computational or physical experiments. The first application is concerned with estimating features of quantities of interest arising from complex computational models, such as quantiles or failure probabilities. A sequential strategy is proposed for iterative refinement of the importance distributions used to efficiently sample the uncertain inputs to the computational model. In the second application, effective use of mathematical surrogates is investigated to help alleviate the analytical and numerical intractability often associated with Bayesian experiment design. This approach allows for the incorporation of prior information into the design process without the need for gross simplification of the design criterion. Illustrative examples of both design problems will be presented as an argument for the relevance of these research problems.

Rare Event Estimation

- Interested in rare event estimation
  - Outputs obtained from computational model
  - Uncertainties in operating conditions and physics variables
  - Physics variables calibrated wrt reference experimental data
- In particular, quantile or percentile estimation

$$Pr[\,\eta(\boldsymbol{x},\boldsymbol{\theta}) > q_\alpha\,] = \alpha$$

  - One of q_α or α is specified and the other is to be inferred
  - q_α may be random when inferring α
- Sequential importance sampling for improved inference
  - Oversample region of parameter space producing rare events of interest
  - Sequentially refine importance distributions for improved inference

Example: VR2plus Model

- Scenario: Pressurizer failure, followed by pump trip and initiation of SCRAM (insertion of control rods)
- Goal: Understand behavior of peak coolant temperature (PCT) in the reactor
- Interested in probability that PCT exceeds 700º K

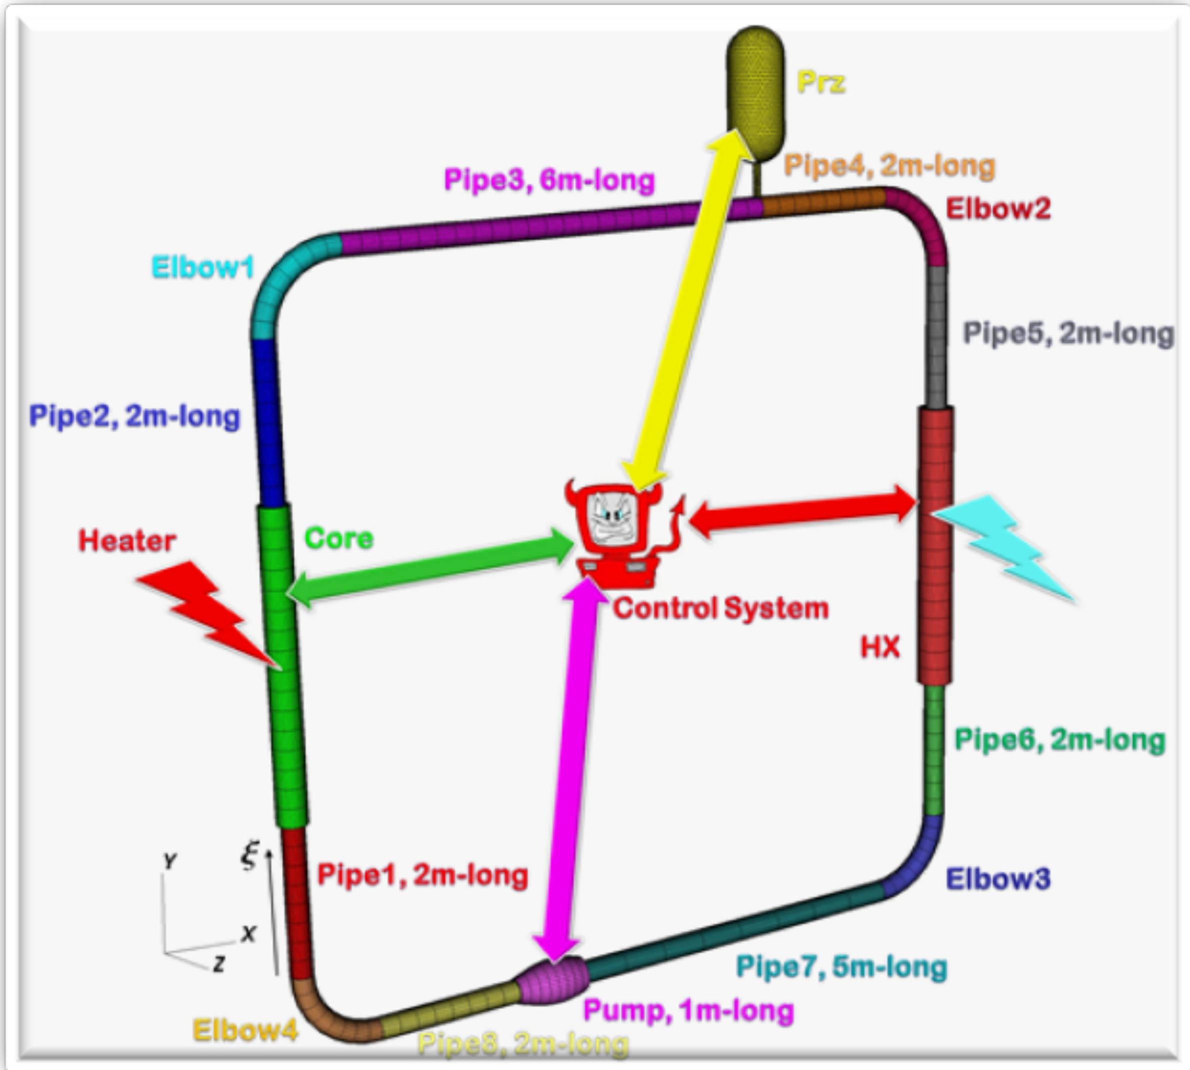

VR2plus Details

- Single thermal-hydraulics loop with 21 components
- Working coolant is water at 16MPa and 600° K, single-phase flow
- Nominal power output of this reactor is 15MW
- Calculations performed with reactor safety analysis code R7 (INL)

Input Parameter	Min	Max	Description
PumpTripPre	15.6 MPa	15.7 MPa	Min. pump pressure causing trip
PumpStopTime	10 s	100 s	Relaxation time of pump phase-out
PumpPow	0.0	0.4	Pump end power
SCRAMtemp	625° K	635° K	Max. temp. causing SCRAM
CRinject	0.025	0.24	Position of CR at end of SCRAM
CRtime	10 s	50 s	Relaxation time of CR system

Input parameters assigned independent
Uniform distributions on their ranges

Importance Sampling

- Sample from alternative distribution g(**x**,θ), evaluate computational model, and weight each output

$$\int 1_{[\eta(\boldsymbol{x},\boldsymbol{\theta})>q_\alpha]}(\boldsymbol{x},\boldsymbol{\theta})\, w(\boldsymbol{x},\boldsymbol{\theta}) g(\boldsymbol{x},\boldsymbol{\theta})\, d\boldsymbol{x}\, d\boldsymbol{\theta}$$

Importance weight $w(\boldsymbol{x},\boldsymbol{\theta}) \equiv \frac{f(\boldsymbol{x})\,\pi(\boldsymbol{\theta})}{g(\boldsymbol{x},\boldsymbol{\theta})}$ Importance density

  - Percentile estimator: $\hat{\alpha}^{(g)} = \frac{1}{N}\sum_{j=1}^{N} w(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)})\, 1_{[\eta(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)})>q_\alpha]}(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)})$
  - Quantile estimator: Smallest value of q_α for which $\hat{\alpha}^{(g)} \leq \alpha$

- Importance density g(**x**,θ) ideally chosen to increase the likelihood of observing desired rare events

$$g(\boldsymbol{x},\boldsymbol{\theta}) = \frac{1}{\alpha} 1_{[\eta(\boldsymbol{x},\boldsymbol{\theta})>q_\alpha]} f(\boldsymbol{x})\pi(\boldsymbol{\theta})$$

minimum variance importance density

Goal: Approximate minimum variance importance density to achieve substantial variance reduction

Relevant (Sensitive) Variables

- A subset of variables may be responsible for generating rare events
  - Relevant **x**$_r$, θ_r
  - Irrelevant **x**$_i$, θ_i

$$\eta(\boldsymbol{x},\boldsymbol{\theta}) = \eta(\,(\boldsymbol{x}_r,\tilde{\boldsymbol{x}}_i),(\boldsymbol{\theta}_r,\tilde{\boldsymbol{\theta}}_i)\,) \text{ for all values of } \tilde{\boldsymbol{x}}_i \text{ and } \tilde{\boldsymbol{\theta}}_i$$

- Minimum variance importance distribution

$$g_r(\boldsymbol{x}_r,\boldsymbol{\theta}_r) = \frac{1}{\alpha}1_{[\eta(\boldsymbol{x},\boldsymbol{\theta})>q_\alpha]}(\boldsymbol{x},\boldsymbol{\theta})f_r(\boldsymbol{x}_r)\pi_r(\boldsymbol{\theta}_r) \text{ for any values of } \boldsymbol{x}_i \text{ and } \boldsymbol{\theta}_i$$

$$g_i(\,(\boldsymbol{x}_i,\boldsymbol{\theta}_i)|(\boldsymbol{x}_r,\boldsymbol{\theta}_r)\,) = f_i(\boldsymbol{x}_i|\boldsymbol{x}_r)\pi_i(\boldsymbol{\theta}_i|\boldsymbol{\theta}_r)$$

- Importance distribution g(**x**,θ) becomes

$$g(\boldsymbol{x},\boldsymbol{\theta}) = g_r(\boldsymbol{x}_r,\boldsymbol{\theta}_r)f_i(\boldsymbol{x}_i|\boldsymbol{x}_r)\pi_i(\boldsymbol{\theta}_i|\boldsymbol{\theta}_r)$$

Many Applications Will Require Code Surrogates

- Many computational models run too slowly for direct use in brute force or importance sampling based rare event inference
- Use training runs to develop a statistical surrogate model for the complex code (i.e., the *emulator*) based on a Gaussian process
  - Deterministic code is interpolated with zero uncertainty

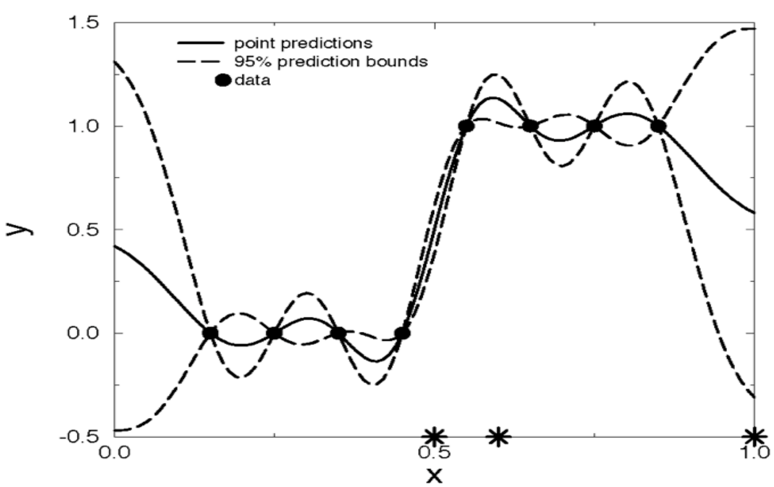

outputs evaluated at training runs $\mathbf{z}_1, \ldots, \mathbf{z}_n$

$$\tilde{\eta}(\boldsymbol{z};\boldsymbol{\beta}) = \boldsymbol{r}^T(\boldsymbol{z};\boldsymbol{\beta})\boldsymbol{R}^{-1}(\boldsymbol{\beta})\boldsymbol{\eta}_n$$

correlations between prediction site $\mathbf{z}$ and training runs $\mathbf{z}_1, \ldots, \mathbf{z}_n$

pairwise correlations between training runs $\mathbf{z}_1, \ldots, \mathbf{z}_n$

- Kriging Variance

$$\widehat{Var}(\,\tilde{\eta}(\boldsymbol{z};\boldsymbol{\beta})\,) = \sigma^2\left(1 - \boldsymbol{r}^T(\boldsymbol{z};\boldsymbol{\beta})\boldsymbol{R}^{-1}(\boldsymbol{\beta})\boldsymbol{r}(\boldsymbol{z};\boldsymbol{\beta})\right)$$

Surrogate-based Percentile/Quantile Estimation

- Gaussian Process model with plug-in covariance parameter estimates, e.g.

$$\frac{\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta}) - \eta(\boldsymbol{x},\boldsymbol{\theta})}{\sqrt{\widehat{Var}(\,\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta})\,)}} \sim N(0,1)$$

- Process-based inference

$$Pr\left[\eta(\boldsymbol{x},\boldsymbol{\theta}) > q | (\boldsymbol{x},\boldsymbol{\theta})\right] = \Phi\left(\frac{\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta}) - q}{\sqrt{\widehat{Var}(\,\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta})\,)}}\right)$$

$$\alpha = \int \Phi\left(\frac{\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta}) - q_\alpha}{\sqrt{\widehat{Var}(\,\tilde{\eta}(\boldsymbol{x},\boldsymbol{\theta})\,)}}\right) w(\boldsymbol{x},\boldsymbol{\theta})\, g(\boldsymbol{x},\boldsymbol{\theta})\, d\boldsymbol{x}\, d\boldsymbol{\theta}$$

$$\hat{\alpha}^{(g)} = \frac{1}{N}\sum_{j=1}^{N} \Phi\left(\frac{\tilde{\eta}(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)}) - q_\alpha}{\sqrt{\widehat{Var}(\,\tilde{\eta}(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)})\,)}}\right) w(\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)})$$

Adaptive Importance Sampling

- Choose $g(\mathbf{x},\theta)$ by iterative refinement
  1. Sample from initial importance density

$$g^{(1)}(\mathbf{x},\theta) = f(\mathbf{x})\,\pi(\theta)$$

  2. Given target probability or quantile level, determine which parameters are sensitive for producing extreme output values

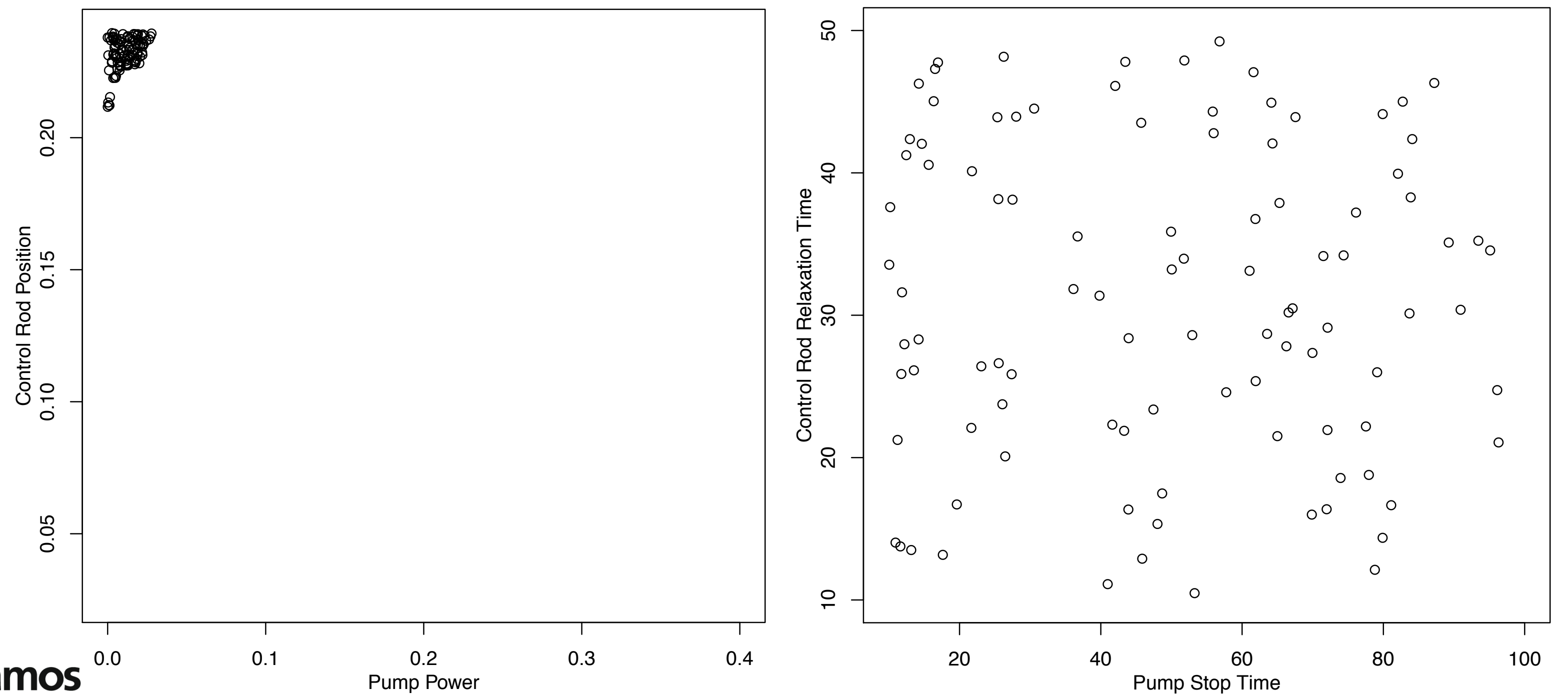

Adaptive Importance Sampling

3. For sensitive parameters, adopt a family of importance densities (e.g. Beta, (truncated) Normal, etc.). Fit the parameters of this family to the selected largest order statistics (e.g. method of moments, MLE).
4. Keep the conditional distribution for the insensitive parameters, given values of the sensitive parameters.
5. Combine the distributions of (3) and (4) to obtain $g^{(2)}(\mathbf{x},\theta)$.
6. If applicable, refine the surrogate in the input region of interest, e.g. using information provided by $g^{(2)}(\mathbf{x},\theta)$.
7. Repeat (2)-(6) until the updated importance distribution $g^{(k)}(\mathbf{x},\theta)$ "stabilizes" in its approximation of the minimum variance importance distribution.

Convergence Assessment Through Iterative Refinement of Surrogate

- Surrogate bias may affect inference results
- Choose design augmentation that minimizes integrated mean square error with respect to the currently estimated importance distributions for sensitive parameters
  - A version of "targeted" IMSE (tIMSE)

$$IMSE(\mathcal{D}_b) = 1 - \text{trace}\left(\left[\int \boldsymbol{r}_{\{\mathcal{D}_0,\mathcal{D}_b\}}(\boldsymbol{z};\boldsymbol{\beta})\, \boldsymbol{r}^T_{\{\mathcal{D}_0,\mathcal{D}_b\}}(\boldsymbol{z};\boldsymbol{\beta}) \prod_{i=1}^{n_z} w_i(z_i)\, d\boldsymbol{z}\right] \boldsymbol{R}^{-1}_{\{\mathcal{D}_0,\mathcal{D}_b\}}(\boldsymbol{\beta})\right)$$

Estimate correlation parameters based on current design $\mathcal{D}_0$

For sensitive input *i*, select $w_i(z_i)$ to be its importance distribution

VR2plus Analysis

- Percentile inference
  - PCT > 700^o K
- Importance distribution
  - Independent Beta distributions for sensitive parameters
  - Independent Uniform distributions for insensitive parameters

PumpPow and CRinject are sensitive

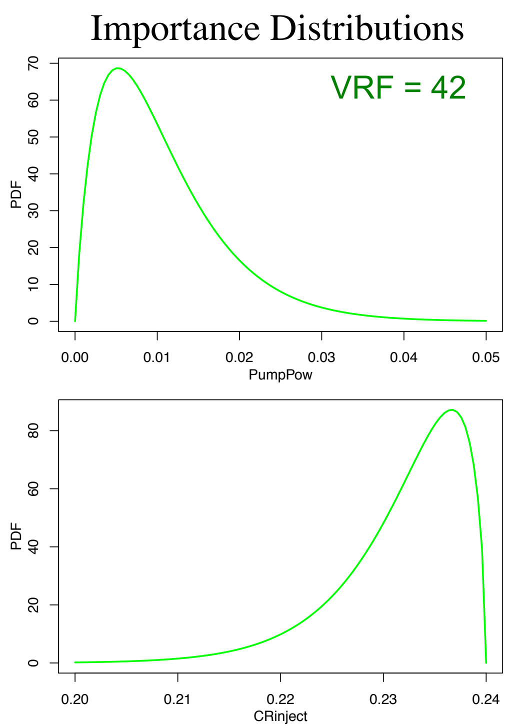

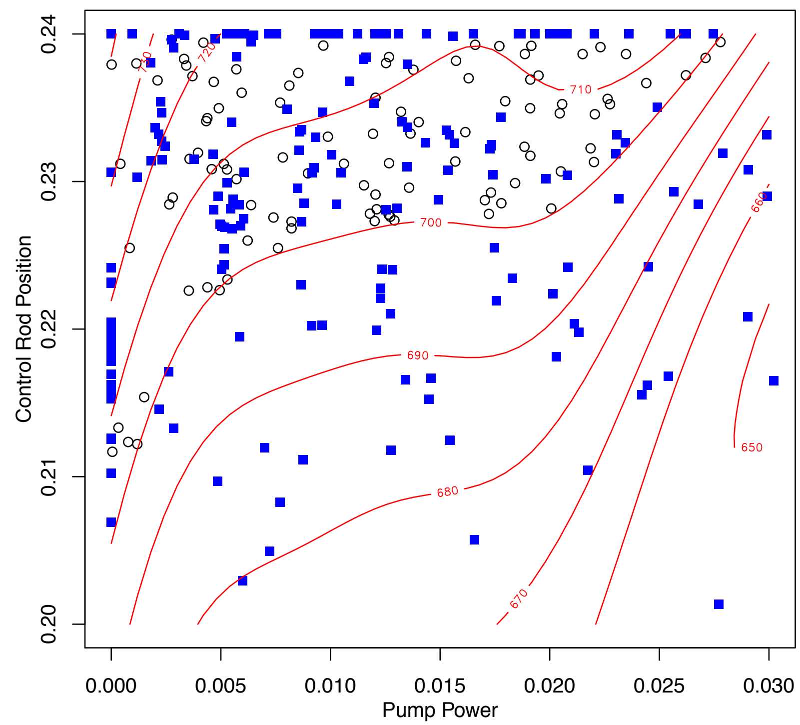

Los Alamos NATIONAL LABORATORY EST.1943

Pressurizer Failure: Percentile Inference

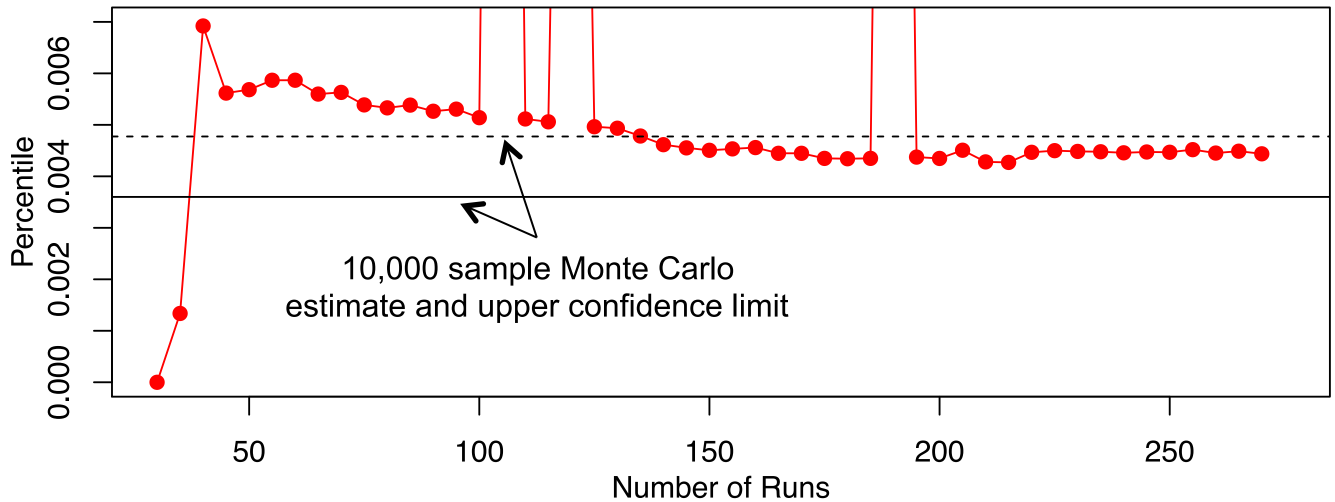

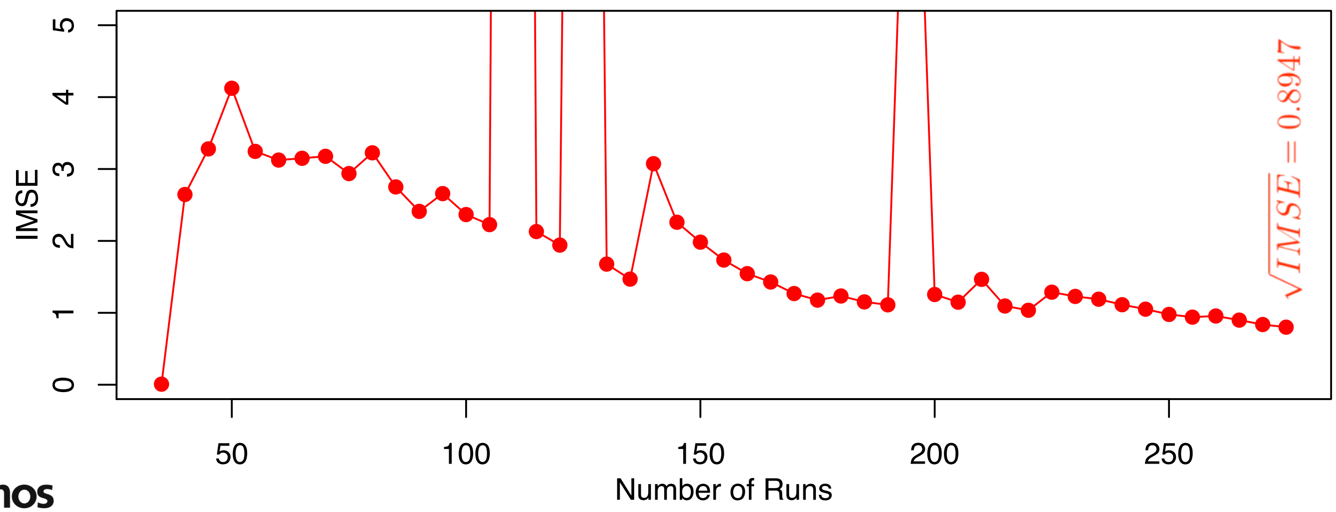

Bayesian Experimental Design

- Decision-theoretic approach to designing experiments that assimilates the objectives of the experiment with prior information to arrive at an optimal set of new runs
- Chaloner, K. and Verdinelli, I. (1995) is an excellent review article on Bayesian experimental design
- Based on utility functions
  - Represent preferences: $U(x) \geq U(y)$ if and only if x is preferred to y
- Two primary experimental objectives: estimation and prediction

Expected Utility

- Design η must be chosen from some set H
- Data **y** observed from sample space Y
- Based on **y**, a decision *d* is chosen from some set D
- Unknown parameter θ defined on parameter space Θ
- General utility function U(*d*, θ, η, **y**)

$$U(\eta) = \int_Y \max_{d \in D} \int_\Theta U(d, \theta, \eta, \mathbf{y})\, p(\theta | \mathbf{y}, \eta)\, p(\mathbf{y} | \eta)\, d\theta\, d\mathbf{y}$$

$p(\mathbf{y}, \theta \mid \eta)$

- Bayesian experimental design solution η^* :

$$U(\eta^*) = \max_{\eta \in H} \int_Y \max_{d \in D} \int_\Theta U(d, \theta, \eta, \mathbf{y})\, p(\theta | \mathbf{y}, \eta)\, p(\mathbf{y} | \eta)\, d\theta\, d\mathbf{y}$$

Motivating Application: Accelerated Testing

- Items will not fail under certain conditions
- Such items are exposed to *accelerated conditions* to induce failure
- Typical application: Estimating the lifetime distribution of reliable items

LANL: Understanding the conditions that induce undesirable explosions

Accelerated testing has become commonplace

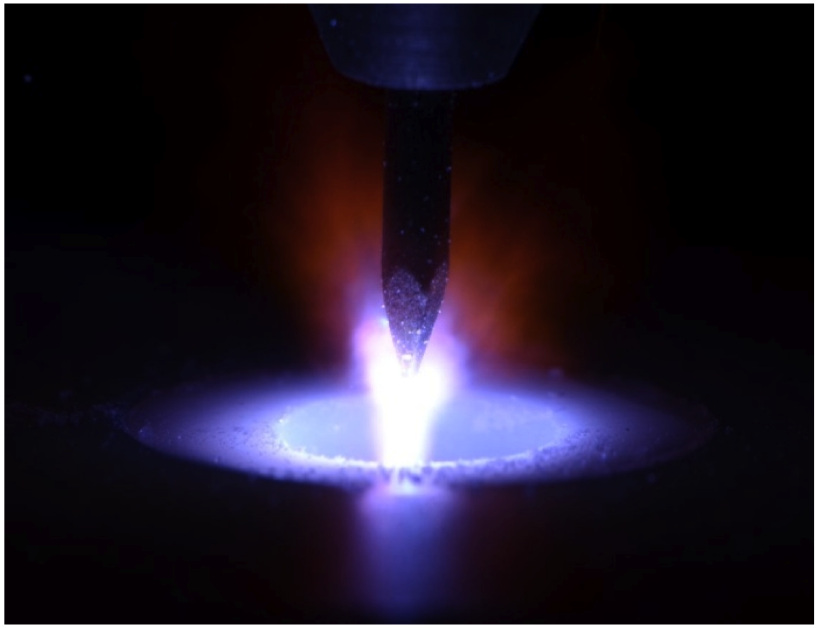

Motivating Application: Test Planning

- Key question: How can prior knowledge be incorporated into accelerated test planning?
- Prior knowledge implemented as an *informative prior distribution* on model parameters
  - Obtained from sequential testing
  - Experience with similar populations
- A *test* consists of
  - Choosing the levels of the accelerators
  - Proportion of units at each selected level

Motivating Application: Device A Data

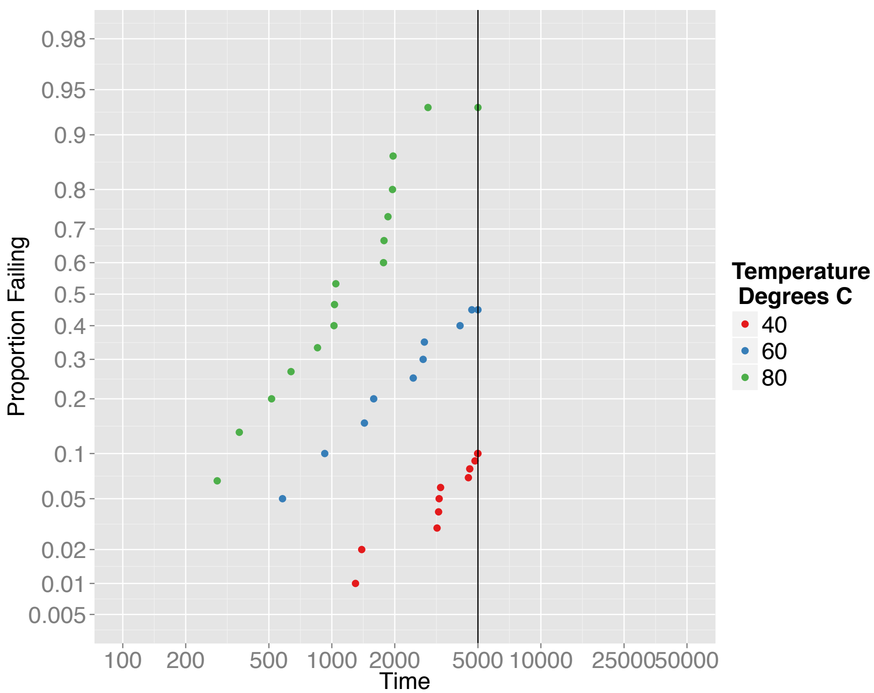

- Data are from Meeker and Escobar (1998)
- 165 items are exposed to accelerated temperatures
  - 30 units at 10°C with 0 failures
  - 100 units at 40°C with 10 failures
  - 20 units at 60°C with 9 failures
  - 15 units at 80°C with 14 failures
- Meeker and Escobar (1998) assume a lognormal distribution with Arrhenius relation for failure time T
  - $\log(T) \sim N(\mu(x), \sigma^2)$
  - $\mu(x) = \alpha + \beta x$
  - $x = 11605/(273.15 + \text{Temp}(^{o}C))$
  - $\theta = (\alpha, \beta, \sigma)$

Los Alamos NATIONAL LABORATORY EST. 1943

Motivating Application: New Data

- Suppose
  - Different yet similar items to Device A are to be tested next
  - Expect a similar failure mechanism as Device A
  - Posterior distribution of θ resulting from Device A analysis is relevant
    - Becomes the prior distribution p(θ) for planning and analyzing the next test

- Notation
  - η is an accelerated test
    - Temperatures to test and corresponding proportions
  - **t** is a vector of failure and censoring times corresponding to η
  - $t_{0.1}$ is the 0.1 quantile of the failure-time distribution
  - var($t_{0.1}$ | **t**, η, 10^{o}C) is the posterior variance of $t_{0.1}$ at 10^{o}C

Motivating Application: Test Constraints

- Experimentalists specify a temperature range of [10^oC, 80^oC] for new tests
  - 10^oC represents the use condition
  - Testing above 80^oC will introduce a new failure mechanism not seen at the use condition
- All tests end at a censoring time of $t_c = 5000$ hours, independent of test η
- Interest in estimating $t_{0.1}$ at 10^oC with high precision
  - $t_{0.1}(x) = \exp\{ \alpha + \beta x + \sigma\, \Phi^{-1}(0.1) \}$, Φ denotes standard normal CDF
  - $x = 11605/(273.15 + \text{Temp}(^oC)) = 11605/(273.15 + 10) \approx 41$
  - Therefore, $t_{0.1}$ at 10^oC is a function $t_{0.1}(\theta)$ of parameters θ only

Exact Bayesian Experimental Design

- Quadratic loss: $U(d,\theta,\eta,\mathbf{t}) = -(t_{0.1}(\theta) - d)^2$

$$\max_{d \in D} \int_\Theta U(d,\theta,\eta,\mathbf{t})\, p(\theta|\mathbf{t},\eta)\, d\theta = -\mathrm{var}(t_{0.1}|\mathbf{t},\eta,10^\circ C)$$

- Expected utility:

$$\begin{aligned} U(\eta) &= -\int_T \mathrm{var}(t_{0.1}|\mathbf{t},\eta,10^\circ C)\, p(\mathbf{t}|\eta)\, d\mathbf{t} \\ &= -\int_\Theta \int_T \mathrm{var}(t_{0.1}|\mathbf{t},\eta,10^\circ C)\, p(\mathbf{t}|\theta,\eta)\, d\mathbf{t}\, p(\theta)\, d\theta \end{aligned}$$

- Choose η to minimize

$$\Lambda(\eta) = \int_\Theta \int_T \mathrm{var}(t_{0.1}|\mathbf{t},\eta,10^\circ C)\, p(\mathbf{t}|\theta,\eta)\, d\mathbf{t}\, p(\theta)\, d\theta$$

Design Criterion Estimation

$$\Lambda(\eta) = \int_{\Theta} \int_{T} \text{var}(t_{0.1}|\mathbf{t}, \eta, 10^{\circ}C)\, p(\mathbf{t}|\theta, \eta)\, d\mathbf{t}\, p(\theta)\, d\theta$$

- Given η, $\Lambda(\eta)$ is estimated as follows:
  1. Draw θ* from p(θ)
  2. Simulate lifetimes **t*** from their lognormal sampling distribution given θ* and η
  3. Censor the elements of **t*** exceeding t_c
  4. Using importance sampling or MCMC to generate a posterior sample of θ given **t*** and η, estimate var($t_{0.1}$(θ) | **t***, η, 10°C) and denote the resulting sample variance by V*(η)
  5. Repeat steps 1-4 to obtain $V_1^*(\eta), \ldots, V_J^*(\eta)$ for *J* large (e.g. *J* = 10,000)
  6. Compute the Monte Carlo estimate $\hat{\Lambda}(\eta) = \frac{1}{J}\sum_{j=1}^{J} V_j^*(\eta)$

Using GPs to Estimate Optimal Designs

1. Select a set of initial designs $\eta_1, \eta_2, \ldots, \eta_k$ and compute design criterion estimates $\hat{\Lambda}(\eta_1), \hat{\Lambda}(\eta_2), \ldots, \hat{\Lambda}(\eta_k)$
   – Initial designs are generated, for example, via random or Latin hypercube sampling
2. Fit a GP to $\hat{\Lambda}_k = [\hat{\Lambda}(\eta_1), \hat{\Lambda}(\eta_2), \ldots, \hat{\Lambda}(\eta_k)]$
3. Apply the Expected Quantile Improvement (EQI) algorithm with this GP to select the next design η_{k+1}.
   – EQI selects a design that minimizes an upper quantile of $\Lambda(\eta)$.
4. Increment k to k+1
5. Repeat steps 2-4 until run budget expires or EQI stopping criteria are satisfied.

Motivating Application: EQI Implementation

- Design η has the following form

$$\eta = \begin{bmatrix} \text{Temp}_1 & \pi_1 \\ \text{Temp}_2 & \pi_2 \\ \vdots & \vdots \\ \text{Temp}_B & \pi_B \end{bmatrix} \text{ where } \pi_i \geq 0 \text{ for } i = 1, \ldots, B \text{ and } \sum_{i=1}^{B} \pi_i = 1$$

- Zhang, Y. and Meeker, W.Q. (2006) suggests a near-optimal test will have $B = 2$
  - Our search restricted to this case, so that GP is a function of the 3 inputs (Temp_1, Temp_2, π_1):

$$\eta = \begin{bmatrix} \text{Temp}_1 & \pi_1 \\ \text{Temp}_2 & 1 - \pi_1 \end{bmatrix}$$

Motivating Application: Results

Initial 30 Designs

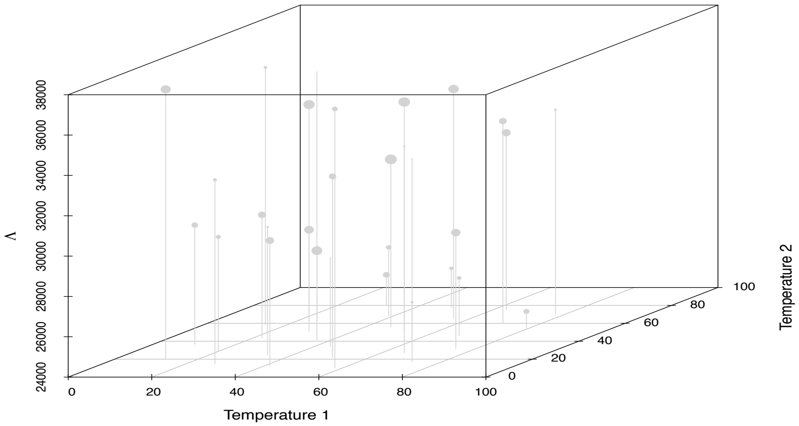

EQI Designs 1-10

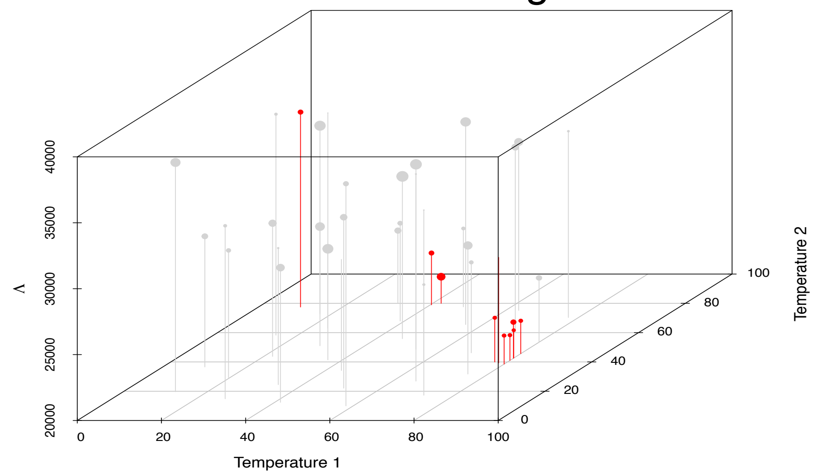

EQI Designs 1-20

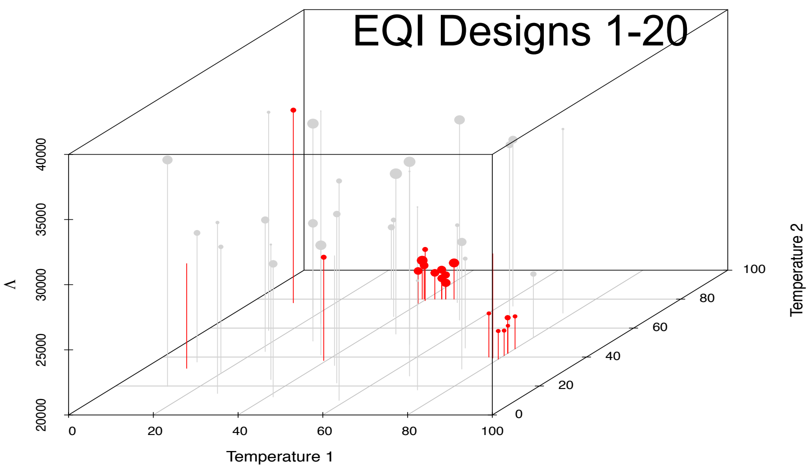

EQI Designs 1-30

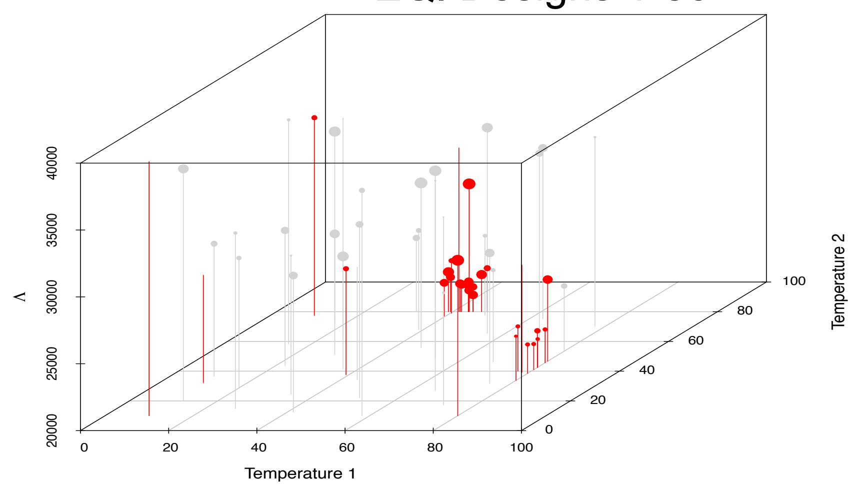

Motivating Application: Comparisons

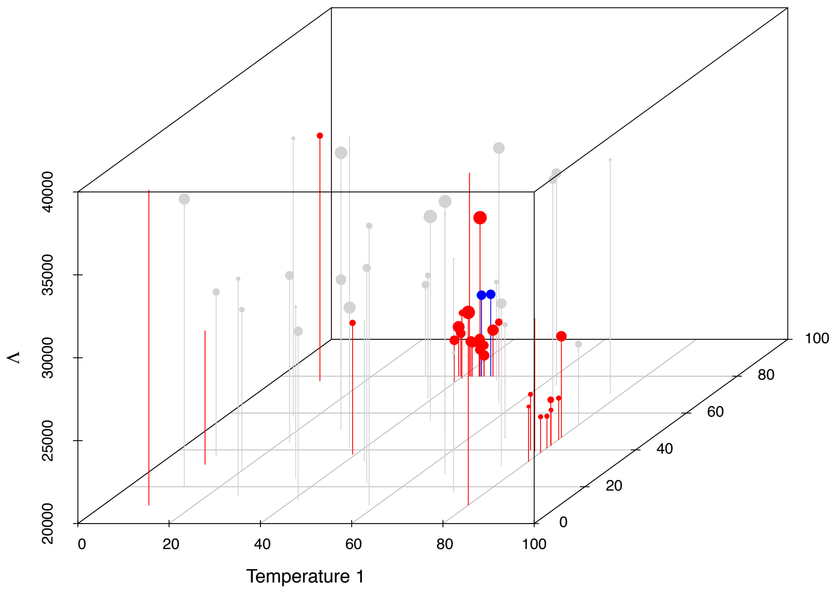

Exact Bayes: EQI Iterates

Approximate Bayes: MLE &

Device A posterior for p(θ)

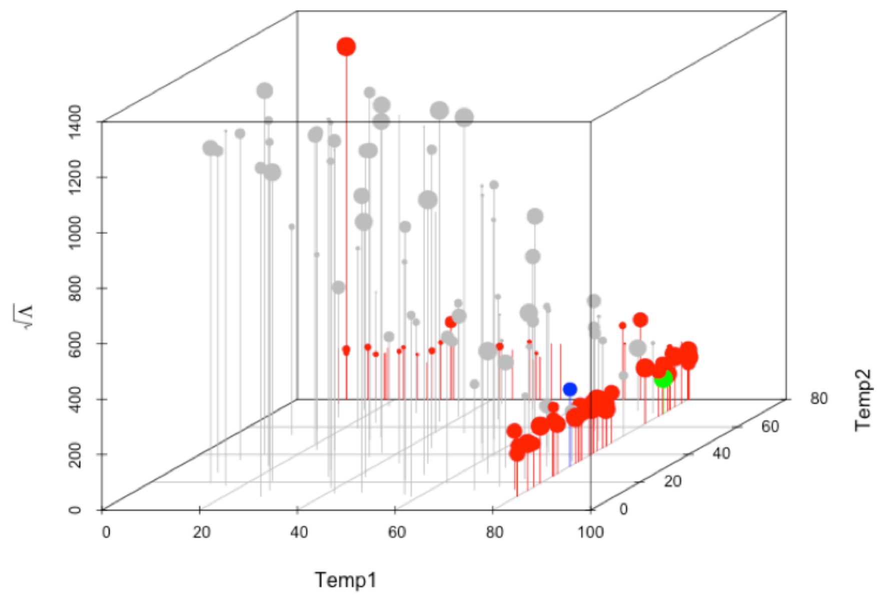

Reduce censoring time so that posterior is not approximately Gaussian:

Exact Bayes: EQI Iterates

EQI Optimal Design

Approximate Bayes

Conclusions

- Surrogates enable optimization of computationally intensive experimental design criteria
  - Slow running codes
  - Analytically intractable criterion functions
- Percentile and quantile estimation
  - Importance sampling improves UQ of percentile and quantile estimates relative to brute force approach
  - Sequential design improves surrogate quality in region of parameter space indicated by importance distributions
- Bayesian experimental design
  - Powerful framework for incorporating prior information into optimal designs
  - Estimation of analytically intractable design criteria can be mitigated through the application of surrogate-based optimization techniques