Laboratories

SAND2020-2778C

Calibration, Propagation, and Validation of Model Discrepancy Across Experimental Settings

PRESENTED BY

Kathryn Maupin and Laura Swiler

SIAM Conference on Uncertainty Quantification

Overview

- Motivation
- Background
- Dakota implementation
- Example
- Philosophical Issues

Motivation: Predictions Under Uncertainty

We need to make predictions that incorporate both parametric and model form uncertainties

- Predictions may be interpolatory or extrapolatory
- Central to high-consequence model and simulation activities

Here, we focus on non-intrusive methods to support black-box simulations

- Perform predictions under uncertainty with explicit discrepancy models
- Explore challenges from algorithmic and deployment perspectives

Calibration of Computer Models

Experimental data = Model output + error

$$d(x_i) = M(\boldsymbol{\theta}, x_i) + \varepsilon_i$$

- $\boldsymbol{\theta}$ = variables to be calibrated
- x = scenario or configuration variables
  - Represent different experimental settings at which data is taken (temperature, pressure, etc)
- $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ = i.i.d. measurement/observation error

The likelihood over n experiments is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^{n} \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{\left(d(x_i) - M(\boldsymbol{\theta}, x_i)\right)^2}{2\sigma^2}\right]$$

A surrogate model $\hat{M}(\boldsymbol{\theta}, x)$ may be used in place of the simulation model $M(\boldsymbol{\theta}, x)$ for computational efficiency

Calibration of Computer Models

Often, even with calibration, the agreement between the data and the model is not very close. This can be due to model form error, also called model discrepancy or structural error

$$\Rightarrow \quad d(x_i) = M(\boldsymbol{\theta}, x_i) + \delta(x_i) + \varepsilon_i$$

Goal: Make predictions in the presence of parametric and model form uncertainties

Philosophical and implementation issues:

- How do we estimate δ?
- What model form is appropriate for δ?
- How can we understand if there is significant confounding between our estimates of $\boldsymbol{\theta}$ and δ?
- How can we appropriately use δ to improve the predictive capability of the model?
- How do we capture and propagate uncertainty

Kennedy and O'Hagan, "Bayesian calibration of computer models." J R Stat Soc, B63, pg 425-464

Additional work on model discrepancy

The parameters of a model are non-identifiable when multiple combinations of θ and δ yield equally good fits to the data

- Arendt, P. D., Apley, D. W. & Chen, W. (2012)
  - Recommend modular Bayes approach to encourage identifiability
- Ling, Y., Mullins, J., Mahadevan, S. (2014)
  - Introduced a practical method to determine non-identifiability of parameters
  - Examined model discrepancy function selection
- Brynjarsdottir, J. and O'Hagan, A. (2014)
  - If model discrepancy is ignored, then predictions (both interpolations and extrapolations) and inferences about parameters are biased
  - If model discrepancy is used, it must be informed with physical information and a carefully selected prior

Additional work on model discrepancy

Discrepancy may be embedded in the model instead of explicitly added

- Sargsyan, K., Najm, H. N., Ghanem, R. (2015)
- Sargsyan, K., Huan, H., Najm, H. N. (2018)
  - Explicit discrepancy may not adhere to physical laws
  - Explicit discrepancy is confounded with measurement error
- Morrison, R. E., Oliver, T. A., Moser, R. D. (2016)
- Portone, T., McDougall, D. Moser, R. D. (2017)
  - Use of a stochastic operator as an inadequacy representation

How do we make the answers to our philosophical questions general?

DAKOTA

Explore and predict with confidence.

Automate typical parameter variation studies with advanced methods and a generic interface to your simulation

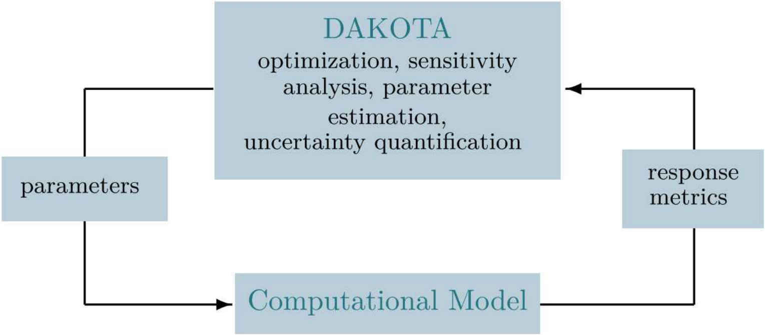

Discrepancy Formulation in Dakota

Given data, we want to calibrate model parameters $\boldsymbol{\theta}$ and calculate δ

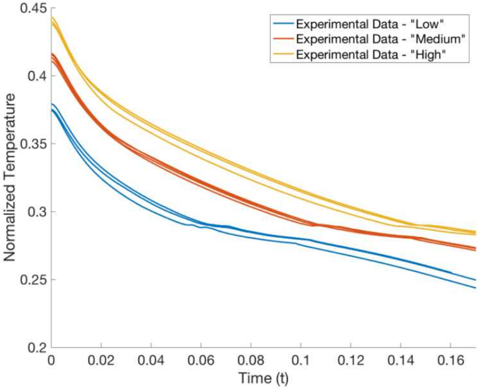

- Response = experimental value at a point in time or space
- Field = set of responses for single experiment
- Configuration = experimental setting such as temperature or pressure

Discrepancy Formulation in Dakota

Currently in Dakota

- Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $d(x)$
- Scalar responses
  - For each response function

$$\delta(x_i) = d(x_i) - M(\boldsymbol{\theta}, x_i)$$

  - $\delta = \delta(x)$ is **only** a function of the configuration variables
- Field responses
  - For each response in the field

$$\delta(t_i, x_j) = d(t_i, x_j) - M(t_i, \boldsymbol{\theta}, x_j)$$

  - $\delta = \delta(t, x)$ is a function of the configuration variables and independent field coordinates
- Prediction variance can also be computed

$$\Sigma_{total}(t, x) = \Sigma_d(t, x) + \Sigma_\delta(t, x)$$

Example: Thermal Battery Calibration

We wish to use a single model for temperature calculations for any initial condition

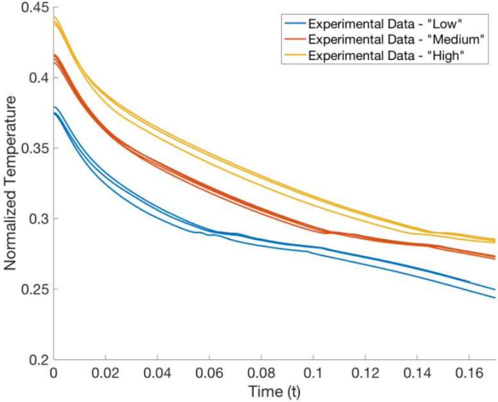

- t = time
- $\boldsymbol{\theta} = \{\theta_1, \ldots \theta_7\}$ = parameters to be calibrated
- x = configuration parameter (initial condition)

Example: Thermal Battery Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

$$\pi(\boldsymbol{\theta}|\mathbf{d}) \propto \pi(\mathbf{d}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$$

- Model is an emulator with 7 parameters
- Three cases of "leave one out" calibration
  - Calibrate to low and medium, extrapolate to high
  - Calibrate to low and high, interpolate to medium
  - Calibrate to medium and high, extrapolate to low
- Choose one experiment of each type
- $\pi(\boldsymbol{\theta}) \sim \mathcal{U}$
- $\pi(\mathbf{d}|\boldsymbol{\theta}) \sim \mathcal{N}$

Example: Thermal Battery Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

- Using $\bar{\boldsymbol{\theta}}$, the model is inadequate

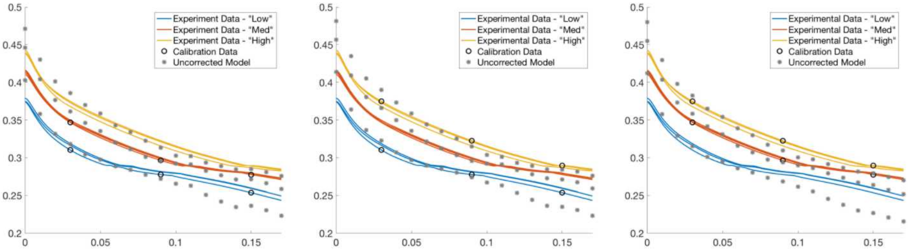

Example: Thermal Battery Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}d$ using Bayes' Rule

Step 2: Calculate discrepancies

$$\delta(t_i, x_j) = d(t_i, x_j) - M(t_i, \bar{\boldsymbol{\theta}}, x_j)$$

Example: Thermal Battery Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

Step 2: Calculate discrepancies

Step 3: Calibrate discrepancy model

- Discrepancy model corrected some areas better than others
- Experimental data is contained within the prediction intervals of the corrected model

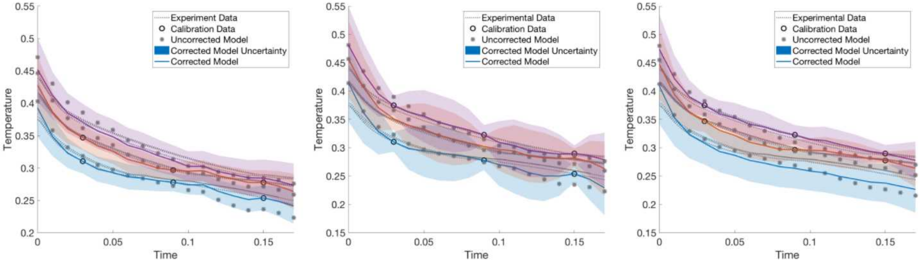

Example: Thermal Battery Calibration

Comparison to Simultaneous Calibration

The original Kennedy and O'Hagan paper proposed simultaneous estimation of $\boldsymbol{\theta}$ and δ parameters

$$\Rightarrow \pi(\boldsymbol{\theta}, \boldsymbol{\ell} | \mathbf{d}) \propto \pi(\mathbf{d} | \boldsymbol{\theta}, \boldsymbol{\ell}) \pi(\boldsymbol{\theta}, \boldsymbol{\ell})$$

- $\boldsymbol{\ell} = \{\ell_x, \ell_t\}$ = correlation lengths of δ
- $\boldsymbol{\theta} = \{\theta_1, \ldots, \theta_7\}$
- $\pi(\boldsymbol{\theta}, \boldsymbol{\ell}) \sim \mathcal{U}$
- $\pi(\mathbf{d} | \boldsymbol{\theta}, \boldsymbol{\ell}) \sim \mathcal{N}$

Example: Thermal Battery Calibration

Comparison to Simultaneous Calibration

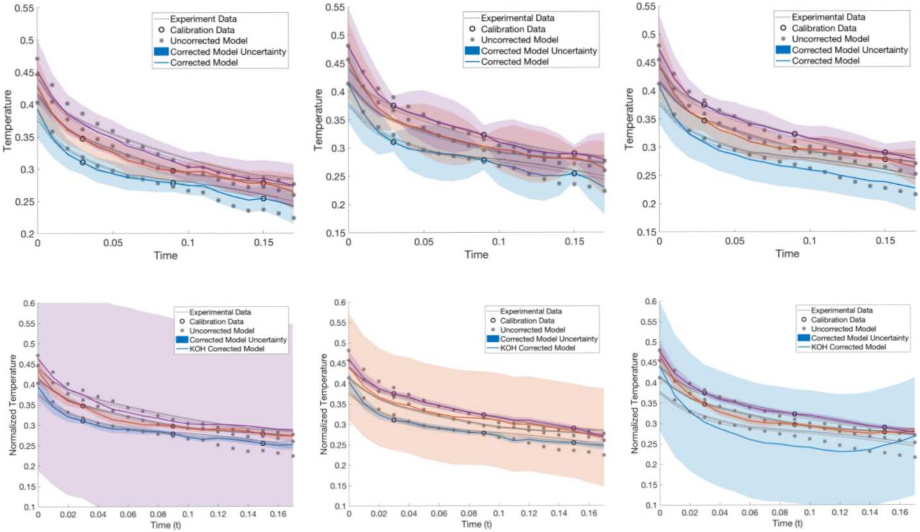

Example: Thermal Battery Calibration

Comparison to Simultaneous Calibration

	Case 1	Case 2	Case 3
Uncorrected Model	3.24%	3.88%	4.22%
Corrected Model	1.74%	2.26%	3.03%
KOH Corrected Model	1.96%	2.32%	4.47%

For the simultaneous approach:

- Smaller variance along calibration temperatures
- Some validation data falls outside of the 95% confidence intervals
- Larger variance along prediction temperatures
  - Larger ℓ_x, likely due to difficulty exploring full parameter space during MCMC
- Calibration times increased by 49%, 39%, and 41% for Cases 1, 2, and 3, respectively
  - Rejection rate much higher ($> 30\%$)
  - Need to build new GP for each sample

Example: Material Failure Calibration

We wish to use a single phenomenological model for stress calculations for any temperature

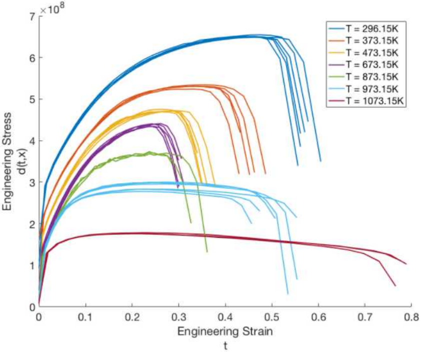

$$m(t, \theta, x) = \theta_1 \left[\frac{\log(100t+1)}{x^{0.5}} - \frac{1}{x^{0.2}(100t - 1.05(\frac{x}{100} - 6.65)^2 \theta_2)^2} \right]$$

- t = strain
- $\boldsymbol{\theta} = \{\theta_1, \theta_2\}$ = parameters to be calibrated
- x = configuration parameter (Temperature)

Example: Material Failure Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

- Calibrate against data from $x = 373K, 673K, 973K$
  - Calculate mean and variance for each
  - $\pi(\boldsymbol{\theta}) \sim \mathcal{U}$
  - $\pi(\mathbf{d}|\boldsymbol{\theta}) \sim \mathcal{N}$
- Using $\bar{\boldsymbol{\theta}}$, the model is inadequate

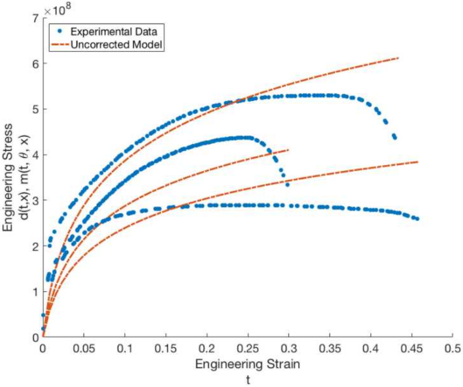

Example: Material Failure Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

Step 2: Calculate discrepancies

$$\delta(t_i, x_j) = d(t_i, x_j) - M(t_i, \bar{\boldsymbol{\theta}}, x_j)$$

Example: Material Failure Calibration

Step 1: Parameters $\boldsymbol{\theta}$ are calibrated to experimental data $\mathbf{d}$ using Bayes' Rule

Step 2: Calculate discrepancies

Step 3: Calibrate discrepancy model

- δ is able to correct the model for the calibratio configurations
- How well does the corrected model perforn for the prediction configurations?

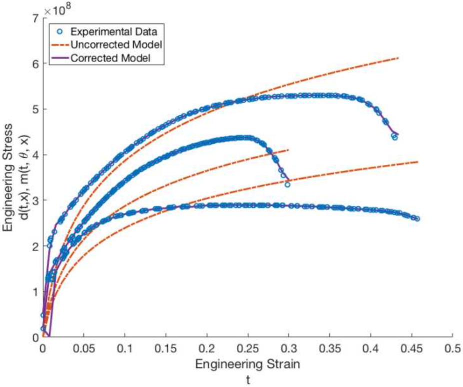

Example: Material Failure Calibration

Interpolation of δ

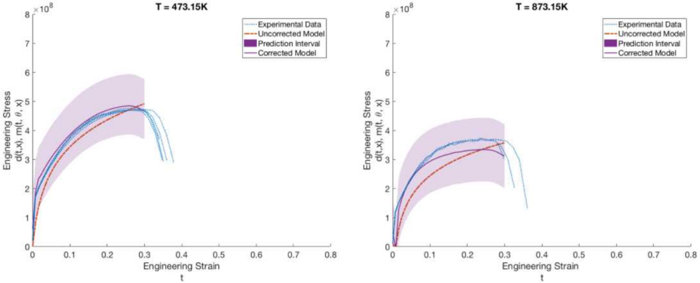

Extrapolation of δ

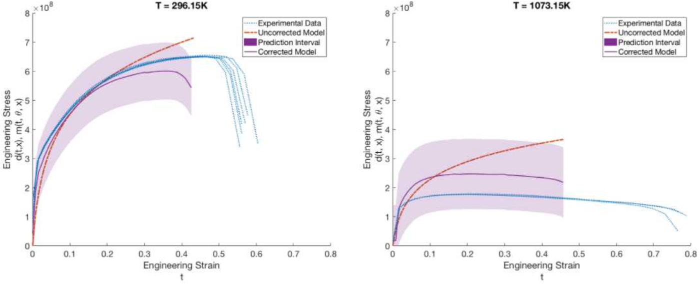

Example: Material Failure Calibration

Comparison to Simultaneous Calibration

As before,

$$\pi(\boldsymbol{\theta}, \boldsymbol{\ell} | \mathbf{d}) \propto \pi(\mathbf{d} | \boldsymbol{\theta}, \boldsymbol{\ell}) \pi(\boldsymbol{\theta}, \boldsymbol{\ell})$$

- $\boldsymbol{\ell} = \{\ell_x, \ell_t\}$ = correlation lengths of δ
- $\boldsymbol{\theta} = \{\theta_1, \theta_2\}$
- $\pi(\boldsymbol{\theta}, \boldsymbol{\ell}) \sim \mathcal{U}$
- $\pi(\mathbf{d} | \boldsymbol{\theta}, \boldsymbol{\ell}) \sim \mathcal{N}$

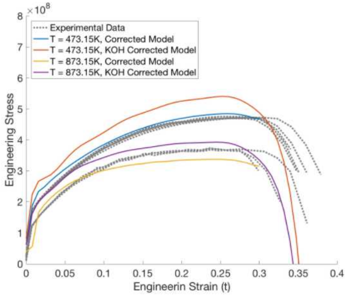

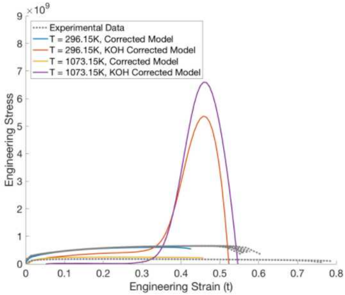

Example: Material Failure Calibration

For each temperature, the original model is inadequate

Sequential calibration approach:

- Experimental data is contained within the prediction intervals of the corrected model
- Corrected model captures general shape of experimental data
- Point of failure is difficult to predict

Simultaneous calibration approach:

- Larger variance along prediction temperatures
- Extrapolation predictions yield unphysical shapes
- Point of failure is difficult to predict

Philosophical Issues

How do we know when a discrepancy model is appropriate?

- Zero mean residuals?
- Autocorrelation of errors?

Simultaneous optimization of discrepancy parameters and calibration parameters

- How do we know whether $\boldsymbol{\theta}$ and δ are identifiable?

How/can we properly use δ in prediction?

- Especially from the sub-system level to larger, more complex scenarios

How do we propagate and capture extrapolation uncertainty?

- In the model form
- In the parameters
- In the discrepancy

Summary

Developed a capability to calculate model discrepancy with field data

Addresses problems with data under different experimental conditions (configurations)

- Example: Calibration of material models using stress-strain data at different temperatures

This capability allows us to investigate tradeoffs between amount of data, number of parameters, and identifiability to better assess future R&D needs

- How do discrepancy predictions perform with less data and more parameters?
- Plan to add diagnostics, such as test for residuals and test for parameter identifiability

Questions?