Bayesian Density Estimation Framework for Model Error Quantification

Khachik Sargsyan, Xun Huan, Habib Najm

Sandia National Laboratories, Livermore, CA

SIAM AN16
Boston, MA, USA
July 11-15, 2016

Sandia National Laboratories

Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Further acknowledgements:
R. Ghanem(USC), J. Bender(LLNL), Y. Marzouk(MIT), C. Feng(MIT),
M. Eldred (SNL), C. Safta (SNL), B. Debusschere (SNL).

Bayesian Density Estimation Framework for Model Error Quantification

Khachik Sargsyan, Xun Huan, Habib Najm

Sandia National Laboratories, Livermore, CA

SIAM AN16
Boston, MA, USA
July 11-15, 2016

DOD, DARPA Enabling Quantification of Uncertainty in Physical Systems (EQUiPS) program
DOE Office of Advanced Scientific Computing Research (ASCR), Scientific Discovery through Advanced Computing (SciDAC)
DOE Office of Basic Energy Sciences, Division of Chemical Sciences, Geosciences, & Biosciences

Further acknowledgements:
R. Ghanem(USC), J. Bender(LLNL), Y. Marzouk(MIT), C. Feng(MIT),
M. Eldred (SNL), C. Safta (SNL), B. Debusschere (SNL).

Outline: model error quantification

- Motivation
- Current issues
- Method
- Toy demos
- Applications
  - Chemical reaction model
  - Atmospheric transport
  - LES computation

Main target

Model error = deviation from ‘truth’, or from a higher-fidelity model

- Represent and estimate the error associated with
  - Simplifying assumptions, parameterizations
  - Mathematical formulation, theoretical framework
  - Numerical discretization
- ...will be useful for
  - Model validation
  - Model comparison
  - Scientific discovery and model improvement
  - Reliable computational predictions
- Inverse modeling context
  - Given experimental or higher-fidelity model data, estimate the model error

Motivation

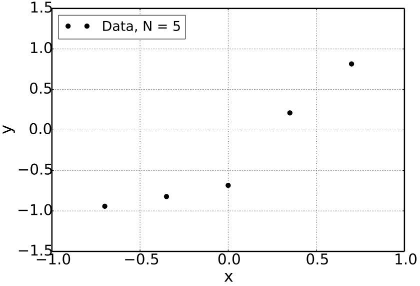

Model-data fit

- Given noisy data – Gaussian noise
- $y = g_{\text{true}}(x) + \epsilon$

Motivation

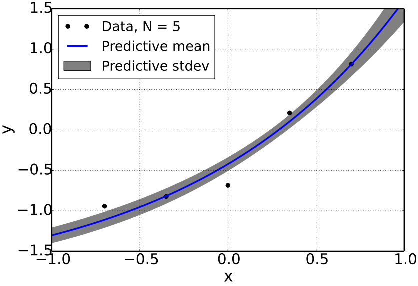

Model-data fit

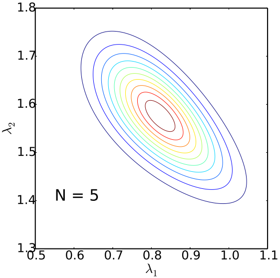

Posterior on parameters

- Employ Bayesian inference to fit an exponential model – $y_{\rm m} = f(x, \lambda)$
- Discrepancy between data and prediction presumed exclusively due to *i.i.d.* Gaussian data noise – $y = f(x, \lambda) + \epsilon_{\rm d}$
- Plotted:
    - Posterior density on the parameters
    - Preditive mean and standard deviation

Motivation

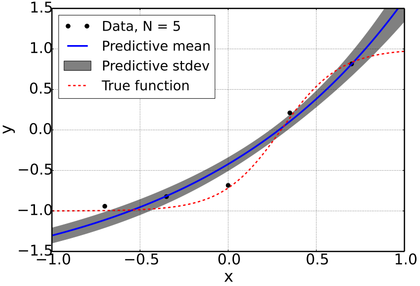

Model-data fit

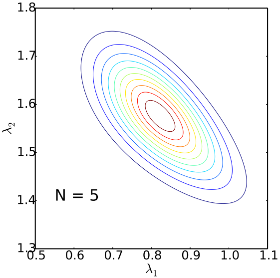

Posterior on parameters

- Employ Bayesian inference to fit an exponential model – $y_{\rm m} = f(x, \lambda)$
- Discrepancy between data and prediction presumed exclusively due to *i.i.d.* Gaussian data noise – $y = f(x, \lambda) + \epsilon_{\rm d}$
- True model $g(x)$ – dashed-red – differs from fit model $f(x, \lambda)$
- Actual discrepancy includes both data and model errors

Motivation

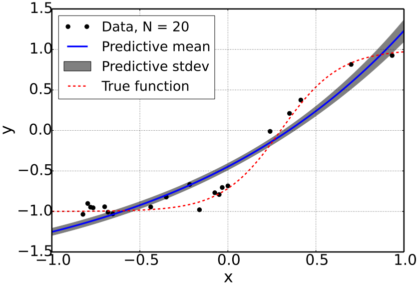

Model-data fit

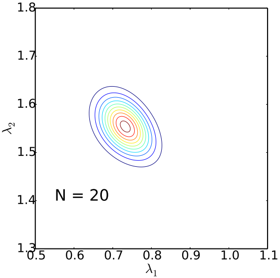

Posterior on parameters

- Increasing number of data points decreases posterior and predictive uncertainty
- We are increasingly sure about predictions based on the *wrong* model

Motivation

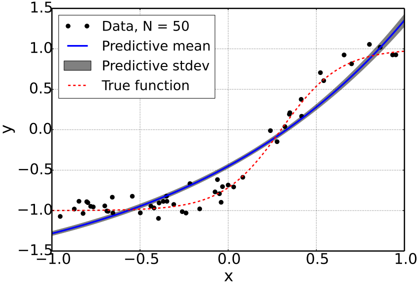

Model-data fit

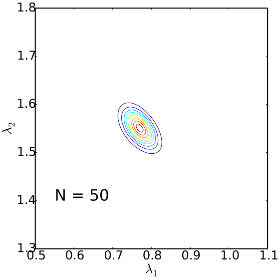

Posterior on parameters

- Increasing number of data points decreases posterior and predictive uncertainty
- We are increasingly sure about predictions based on the *wrong* model

Motivation

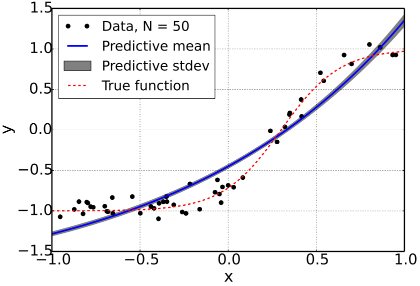

Model-data fit

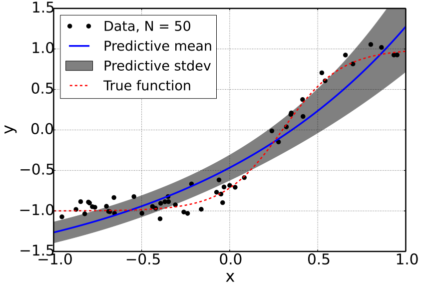

What we want

- If the model has structural uncertainty, more data leads to biased and overconfident results
- We want to quantify model-vs-truth discrepancy in a rigorous and systematic way
    - Cannot ignore model error

Explicit model discrepancy: issues for physical models

$$y_i = \underbrace{f(x_i;\lambda) + \delta(x_i)}_{\text{truth}} + \epsilon_i$$

- Explicit additive statistical model for model error $\delta(x)$ [Kennedy-O'Hagan, 2001]
- Potential violation of physical constraints
- Disambiguation of model error $\delta(x_i)$ and data error ϵ_i
- Calibration of model error on measured observable does not impact the quality of model predictions on other QoIs
- Physical scientists are unlikely to augment their model with a statistical model error term on select outputs
  - Calibrated predictive model: $f(x;\lambda) + \delta(x)$ or $f(x;\lambda)$?
- Problem is highlighted in model-to-model calibration ($\epsilon_i = 0$)
  - no a priori knowledge of the statistical structure of $\delta(x)$

Model error embedding: key idea

Ideally, modelers want predictive *errorbars*:
inserting randomness on the outputs has issues, so...

- Cast input parameters λ as a random variable Λ

  Black-box $$y_i = f(x_i; \Lambda) + \epsilon_i$$

- Generalize parameter forms,

  Random field $$y_i = f(x_i; \Lambda(x_i)) + \epsilon_i$$

- More generally, explore additional parameterizations,

  Extra 'physics' $$y_i = \tilde{f}(x_i; \lambda, \Theta) + \epsilon_i$$

Model error embedding: key idea

Cast input parameters λ as a random variable Λ

$$y_i = f(x_i;\lambda) + \delta(x_i) + \epsilon_i \longrightarrow y_i = f(x_i;\Lambda) + \epsilon_i$$

- Embed model error in specific submodel phenomenology
  - a modified transport or constitutive law
  - a modified formulation for a material property
- Allows placement of model error term in locations where key modeling assumptions and approximations are made
  - as a correction or high-order term
  - as a possible alternate phenomenology
- Naturally preserves model structure and physical constraints

Model error embedding – Bayesian density estimation

$$y_i = f(x_i; \Lambda) + \epsilon_i$$

- Parametrize embedded random variable Λ:
  - PDF form $\pi_\Lambda(\cdot;\alpha)$
  - Polynomial Chaos (PC): $\Lambda = \sum_k \alpha_k \Psi_k(\xi)$
  - Multivariate Normal (MVN):

$$\begin{cases} \Lambda_1 = \alpha_{10} + \alpha_{11}\xi_1 \\ \Lambda_2 = \alpha_{20} + \alpha_{21}\xi_1 + \alpha_{22}\xi_2 \\ \quad \vdots \\ \Lambda_d = \alpha_{d0} + \alpha_{d1}\xi_1 + \alpha_{d2}\xi_2 + \cdots + \alpha_{dd}\xi_d \end{cases}$$

- Inverse modeling context
  - Parameter estimation of $\lambda \Rightarrow$ PDF estimation of $\Lambda \Rightarrow$
    $\Rightarrow$ parameter estimation of α
  - Bayesian formulation $\underbrace{p(\alpha|y)}_{\text{Posterior}} \propto \underbrace{L_y(\alpha)}_{\text{Likelihood}} \underbrace{p(\alpha)}_{\text{Prior}}$

Model error embedding – likelihood options

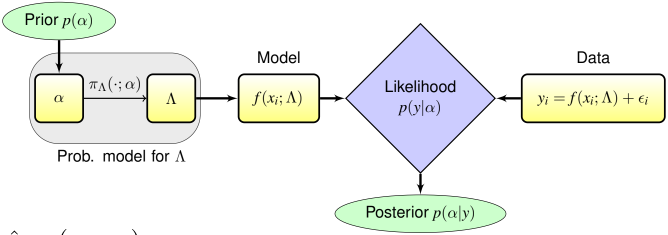

- Infer $\hat{\alpha} = (\alpha, \sigma_{\mathcal{D}})$
- Data generation model; to aid likelihood $p(y|\hat{\alpha})$ construction

$$\begin{aligned} y_i &= f(x_i, \Lambda) + \epsilon_i = \\ &= f\left(x_i, \sum_k \alpha_k \Psi_k(\xi_1, \ldots, \xi_d)\right) + \sigma_{\mathcal{D}} \xi_{d+i} = \\ \text{[NISP]} &\approx \sum_k f_{ik}(\alpha) \Psi_k(\xi_1, \ldots, \xi_d) + \sigma_{\mathcal{D}} \xi_{d+i} = h_i(\hat{\xi}; \hat{\alpha}) \end{aligned}$$

- Full PC germ $\hat{\xi} = (\underbrace{\xi_1, \ldots, \xi_d}_{\text{Model error}}, \underbrace{\xi_{d+1}, \ldots, \xi_{d+N}}_{\text{Data noise}})$

Model error embedding – likelihood options

- Data generation model; to aid likelihood $p(y|\alpha)$ construction

$$
\begin{aligned}
y_i &= f(x_i, \Lambda) + \epsilon_i = \\
&= f\left(x_i, \sum_k \alpha_k \Psi_k(\xi_1, \ldots, \xi_d)\right) + \sigma_{\mathcal{D}} \xi_{d+i} = \\
\text{[NISP]} &\approx \sum_k f_{ik}(\alpha) \Psi_k(\xi_1, \ldots, \xi_d) + \sigma_{\mathcal{D}} \xi_{d+i}
\end{aligned}
$$

- Full Likelihood: $L(\alpha) = p(y|\alpha) = p(y_1, \ldots, y_N|\alpha) = \pi(y)$
  - Degenerate if no data noise
  - Requires multivariate KDE or high-d integration
  - Gaussian approximation:
    $L(\alpha) \propto \exp\left(-\frac{1}{2}(y - \mu(\alpha))^T \Sigma^{-1}(\alpha)(y - \mu(\alpha))\right)$
  - Non-intrusive spectral projection (NISP) relieves the expense and provides easy access to mean $\mu(\alpha)$ and covariance $\Sigma(\alpha)$

Model error embedding – likelihood options

- Data generation model; to aid likelihood $p(y|\alpha)$ construction

$$
\begin{aligned}
y_i &= f(x_i, \Lambda) + \epsilon_i = \\
&= f\left(x_i, \sum_k \alpha_k \Psi_k(\xi_1, \dots, \xi_d)\right) + \sigma_{\mathcal{D}} \xi_{d+i} = \\
\text{[NISP]} &\approx \sum_k f_{ik}(\alpha) \Psi_k(\xi_1, \dots, \xi_d) + \sigma_{\mathcal{D}} \xi_{d+i}
\end{aligned}
$$

- Marginalized Likelihood:
  $L(\alpha) = p(y|\alpha) \approx \prod_{i=1}^N p(y_i|\alpha) = \prod_{i=1}^N \pi(y_i)$
  - Requires univariate KDE
  - Neglects built-in correlations
  - Gaussian approximation:
    $L(\alpha) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^N \Sigma_{ii}^{-1}(\alpha)(y_i - \mu_i(\alpha))^2\right)$

Model error embedding – likelihood options

- Data generation model; to aid likelihood $p(y|\alpha)$ construction

$$
\begin{aligned}
y_i &= f(x_i, \Lambda) + \epsilon_i = \\
&= f\left(x_i, \sum_k \alpha_k \Psi_k(\xi_1, \dots, \xi_d)\right) + \sigma_{\mathcal{D}} \xi_{d+i} = \\
\text{[NISP]} &\approx \sum_k f_{ik}(\alpha) \Psi_k(\xi_1, \dots, \xi_d) + \sigma_{\mathcal{D}} \xi_{d+i}
\end{aligned}
$$

- Approximate Bayesian Computation (ABC): $L(\alpha) = \frac{1}{\epsilon} K\left(\frac{\rho(\mathcal{S}_{\mathcal{M}}, \mathcal{S}_{\mathcal{D}})}{\epsilon}\right)$
  - Mean of $f(x_i; \Lambda)$ is "centered" on the data
  - The width of the distribution of $f(x_i; \Lambda)$ is consistent with the spread of the data around the nominal model prediction

$$
L(\alpha) \propto \exp\left(-\frac{1}{2\epsilon^2} \sum_{i=1}^{N} \left[(\mu_i(\alpha) - y_i)^2 + \left(\sqrt{\Sigma_{ii}(\alpha)} - \gamma|\mu_i(\alpha) - y_i|\right)^2\right]\right)
$$

Model Error – Predictions

$$f(x;\Lambda) = f\left(x; \sum_k \alpha_k \Psi_k(\xi_{1:d})\right) = \sum_k f_k(x;\alpha)\Psi_k(\xi_{1:d})$$

- Non-intrusive spectral projection (NISP) will allow
  - Posterior/pushed-forward predictions
  - Easy access to first two moments:

$$\mu(x;\alpha) = f_0(x;\alpha), \qquad \sigma^2(x;\alpha) = \sum_{k>0} f_k^2(x;\alpha)||\Psi_k||^2$$

- Predictive mean

$$\mathbb{E}[y(x)] = \mathbb{E}_\alpha[\mu(x;\alpha)]$$

- Decomposition of predictive variance

$$\mathbb{V}[y(x)] = \underbrace{\mathbb{E}_\alpha[\sigma^2(x;\alpha)]}_{\text{Model error}} + \underbrace{\mathbb{V}_\alpha[\mu(x;\alpha)]}_{\text{Posterior error}}$$

Model Error – Predictions at data locations

$$f(x_i;\Lambda) = f\left(x_i; \sum_k \alpha_k \Psi_k(\xi_{1:d})\right) + \sigma_D \xi_{i+d} = \sum_k f_k(x_i;\alpha)\Psi_k(\xi_{1:d}) + \sigma_D \xi_{i+d}$$

- Non-intrusive spectral projection (NISP) will allow
  - Likelihood computation
  - Easy access to first two moments:

$$\mu(x_i;\alpha) = f_0(x_i;\alpha), \qquad \sigma^2(x_i;\alpha) = \sum_{k>0} f_k^2(x_i;\alpha)||\Psi_k||^2$$

- Predictive mean

$$\mathbb{E}[y(x_i)] = \mathbb{E}_\alpha[\mu(x_i;\alpha)]$$

- Decomposition of predictive variance

$$\mathbb{V}[y(x_i)] = \underbrace{\mathbb{E}_\alpha[\sigma^2(x_i;\alpha)]}_{\text{Model error}} + \underbrace{\mathbb{V}_\alpha[\mu(x_i;\alpha)] + \sigma_d^2}_{\text{Posterior/Data error}}$$

Predictions account for model error

Calibrating single-exponential models
with data from a double exponential model $g(x) = e^{-0.5x} + e^{-2x}$

Linear-exponential $f(x, \lambda) = e^{\lambda_1 + \lambda_2 x}$

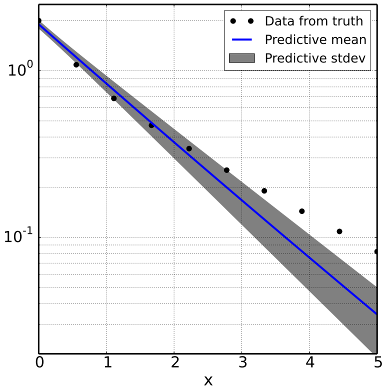

Additive Gaussian error

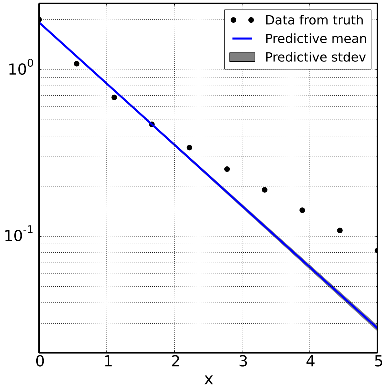

Predictions account for model error

Calibrating single-exponential models
with data from a double exponential model $g(x) = e^{-0.5x} + e^{-2x}$

Linear-exponential $f(x,\lambda) = e^{\lambda_1+\lambda_2 x}$

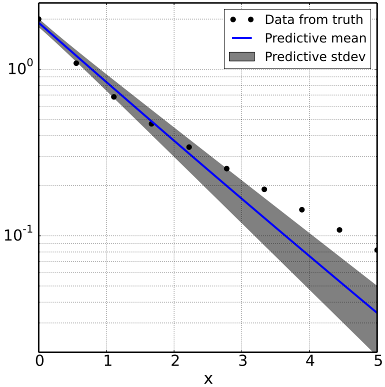

Quadratic-exponential $f_2(x,\lambda) = e^{\lambda_1+\lambda_2 x+\lambda_3 x^2}$

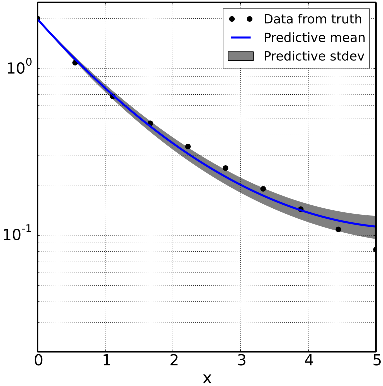

More data leads to 'leftover' model error

Calibrating a quadratic $f(x) = \lambda_0 + \lambda_1 x + \lambda_2 x^2$
w.r.t. 'truth' $g(x) = 6 + x^2 - 0.5(x+1)^{3.5}$ measured with noise $\sigma = 0.1$.

$N = 20$ $\quad$ $N = 50$ $\quad$ $N = 1000$

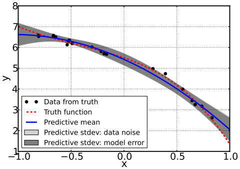

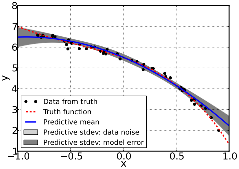

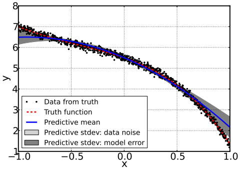

Summary of features:

- Well-defined model-to-model calibration
- Model-driven discrepancy correlations
- Respects physical constraints
- Disambiguates model and data errors
- Calibrated predictions of multiple QoIs

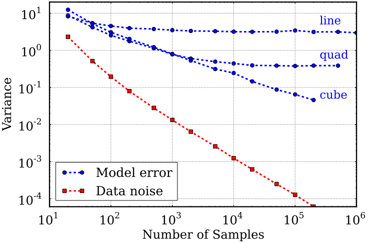

Scenarios of interest

- Model-to-model calibration
  - Chemical reaction model
- Multi-model analysis
  - Atmospheric transport
- Prediction of other QoIs
  - LES computation

Model-to-model calibration: ignition model

- Homogeneous ignition, methane-air mixture
- Single-step global reaction model calibrated against a detailed chemical kinetic model

- Data: ignition time; range of initial T & equivalence ratio
- Single-step model:

$$CH_4 + 2O_2 \rightarrow CO_2 + 2H_2O$$

$$\begin{aligned} \mathfrak{R} &= [CH_4][O_2]k_f \\ k_f &= A\exp(-E/R^oT) \end{aligned}$$

- $(\ln A, E) = \sum_k \boldsymbol{\alpha}_k \Psi_k(\xi)$

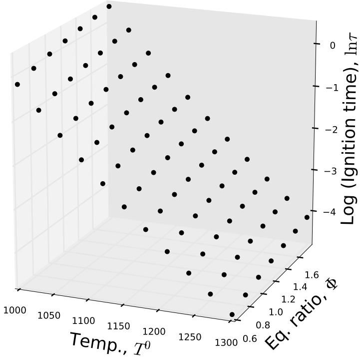

Model-to-model calibration: ignition model

Calibrated uncertain fit model is consistent with the detailed-model data.

Over the range of (T^0, Φ):

- MAP predictive mean ignition-time is centered on the data
- MAP predictive stdv is consistent with the scatter of the data

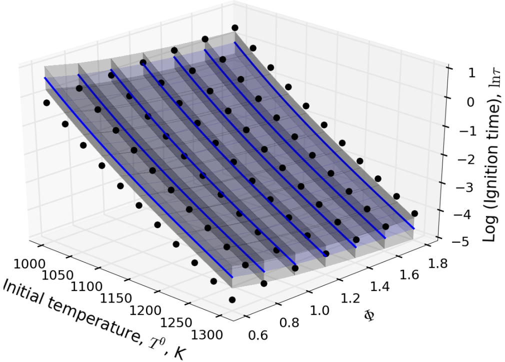

K. Sargsyan, H.N. Najm, and R. Ghanem
"On the Statistical Calibration of Physical Models"
Int. J. Chem. Kin., 47(4): 246-276, 2015

TransCom3 Experiment of CO_2 Flux Inversion

[Gurney *et al.*, Tellus B, 2003]

- Observations $\mathbf{d}$ at $N = 77$ sites around the world
- Inverse problem: find fluxes $\mathbf{s}$ at $M = 22$ locations
- Linearized 'response' model $\mathbf{R}$, such that $\mathbf{d} \approx \mathbf{R}\mathbf{s}$

$$\boxed{\mathbf{d} = \mathbf{R}\mathbf{s} + \epsilon_{\mathbf{d}}}$$

- Model $\mathbf{R}$ is never perfect thus contaminating the inversion
- The inferred values of $\mathbf{s}$ compensate for model deficiencies
- $\epsilon_{\mathbf{d}}$ is meant to capture data errors, but is 'entangled' with model errors

Consider 14 different response models $\mathbf{R}$

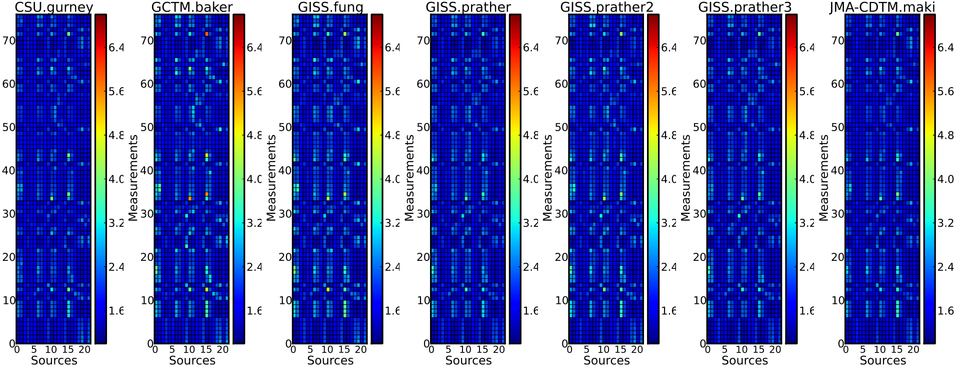

Infer fluxes $\mathbf{s}$, given measurements $\mathbf{d}$ to satisfy $\mathbf{d} \approx \mathbf{R}\mathbf{s}$

- Conventional additive Gaussian error (least-squares): $\mathbf{d} = \mathbf{R}\mathbf{s} + \xi$
- Embed probabilistic model for fluxes s: $\mathbf{d} = \mathbf{R}(\mu_{\mathbf{s}} + \mathbf{C}_{\mathbf{s}}\xi)$

Consider 14 different response models $\mathbf{R}$

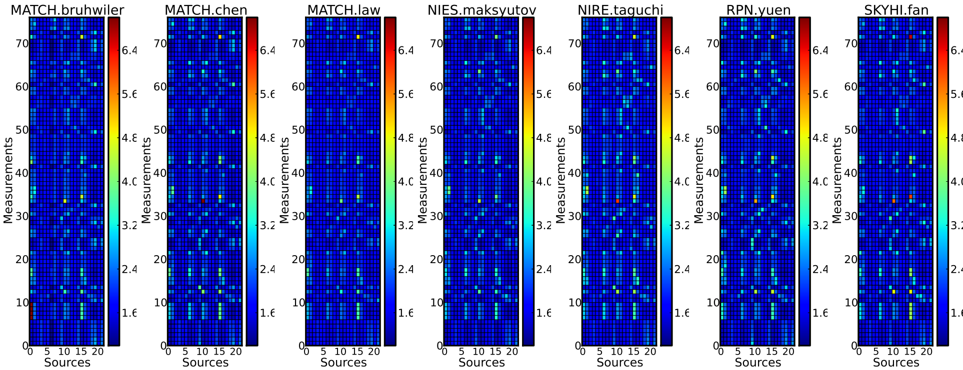

Infer fluxes $\mathbf{s}$, given measurements $\mathbf{d}$ to satisfy $\mathbf{d} \approx \mathbf{R}\mathbf{s}$

- Conventional additive Gaussian error (least-squares): $\mathbf{d} = \mathbf{R}\mathbf{s} + \xi$
- Embed probabilistic model for fluxes s: $\mathbf{d} = \mathbf{R}(\mu_{\mathbf{s}} + \mathbf{C}_{\mathbf{s}}\xi)$

Inferred fluxes show less variability across models

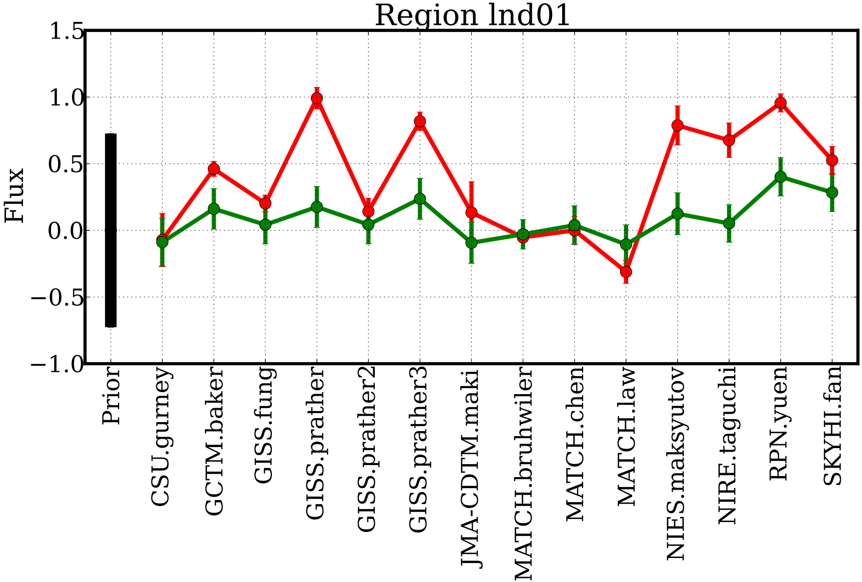

Inferred fluxes show less variability across models

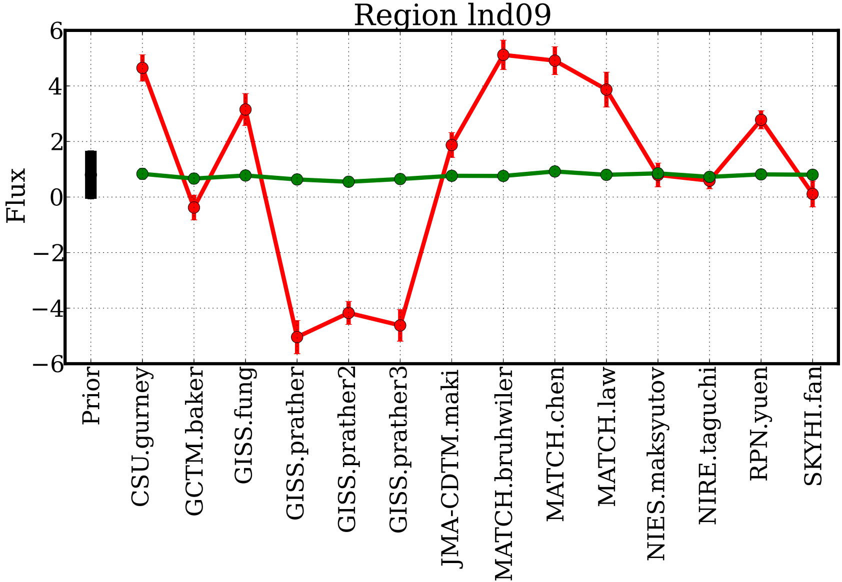

Inferred fluxes show less variability across models

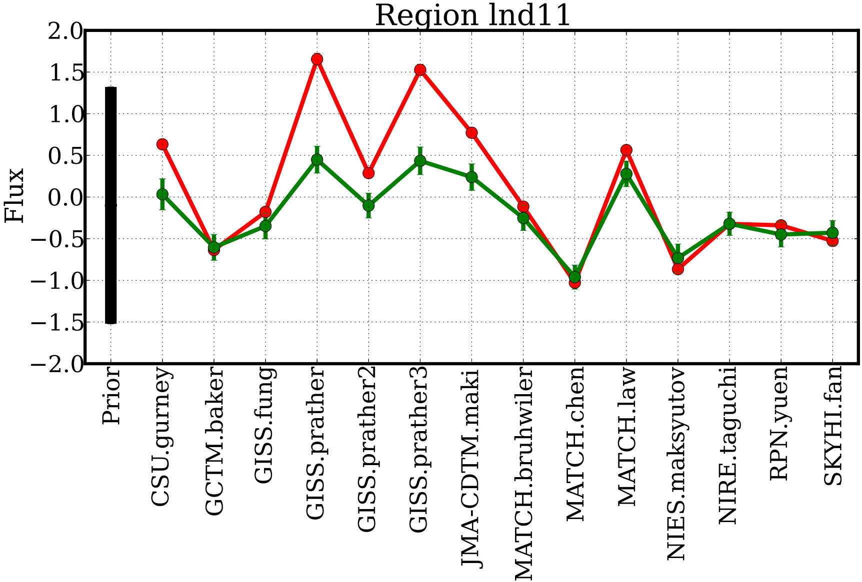

Inferred fluxes show less variability across models

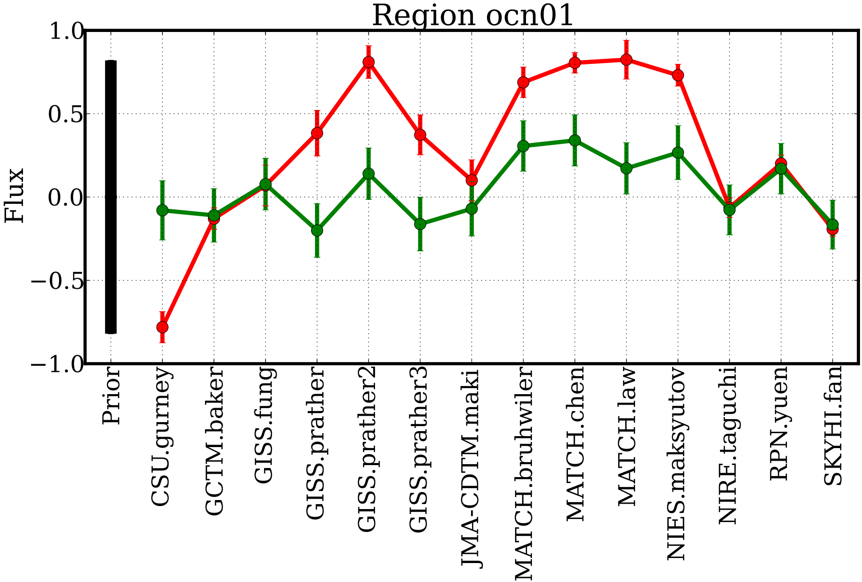

Inferred fluxes show less variability across models

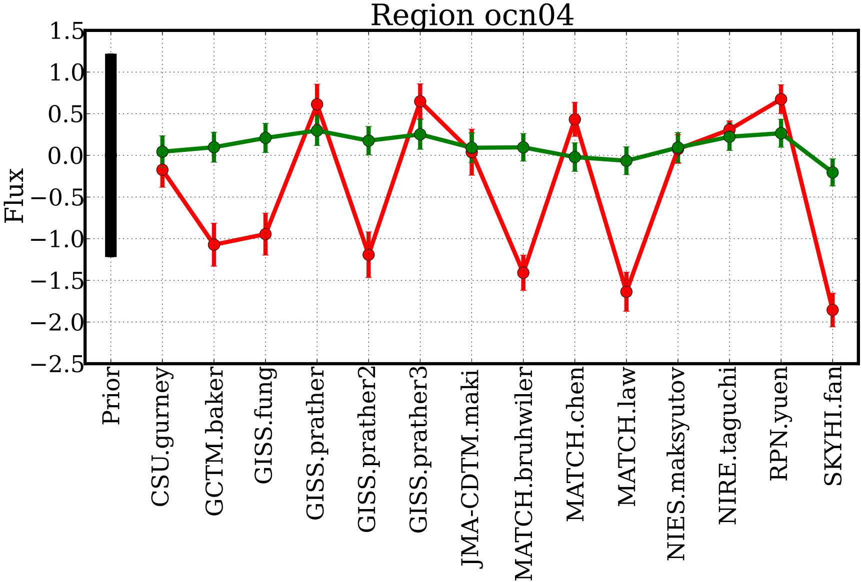

Inferred fluxes show less variability across models

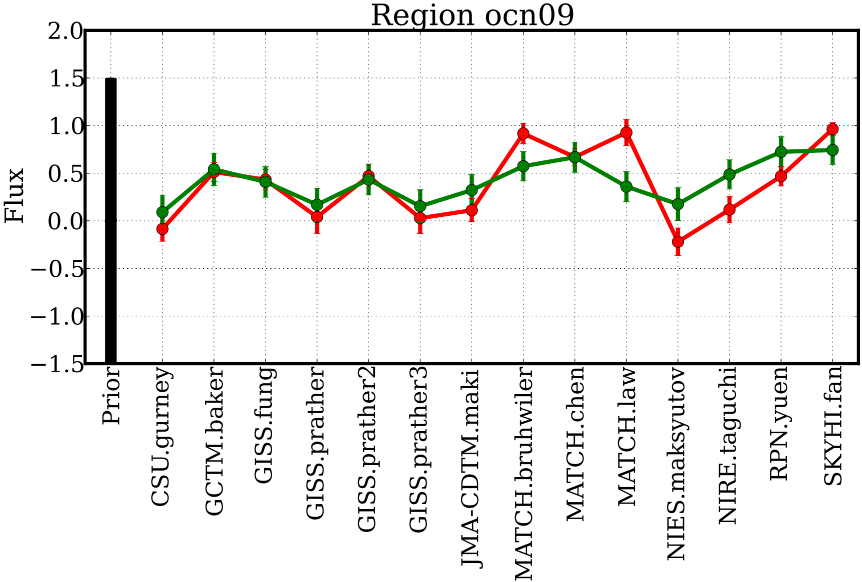

Inferred fluxes show less variability across models

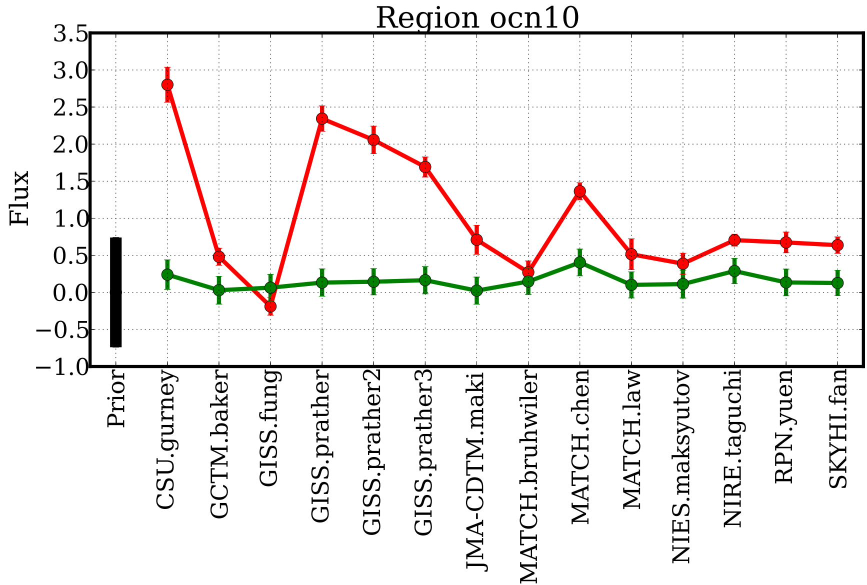

LES computation in Scramjet engine: static-vs-dynamic SGS model calibration

Calibrate with Turbulent Kinetic Energy (TKE) data, predict both TKE and Pressure

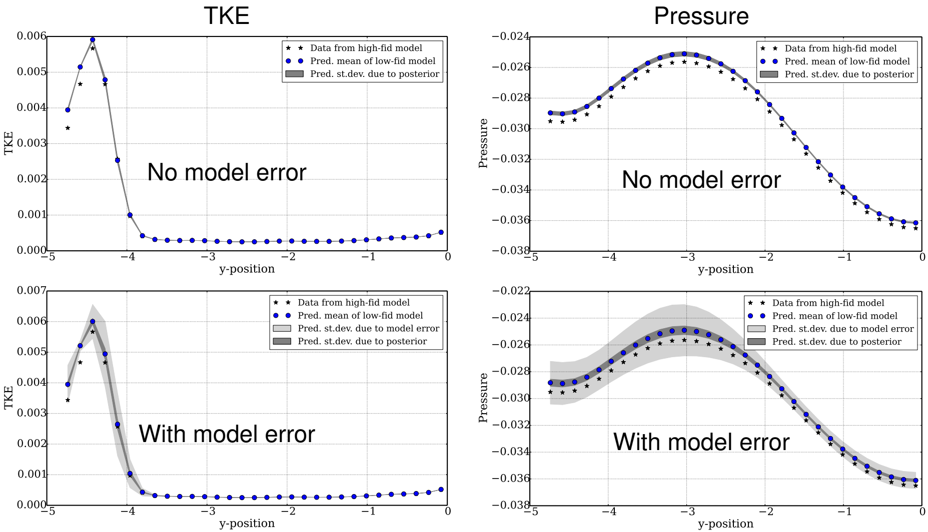

Challenges and Mitigation

- Density estimation is more challenging than parameter estimation
  - Inverse problem is ill-posed or intractable
  - $\Rightarrow$ Employ approximate or empirical likelihoods

- Potentially a high-dimensional Bayesian problem
  - Full posterior may be inaccessible...
  - $\Rightarrow$ Resort to optimization algorithms in no-noise case
  - ... or hard to sample from
  - $\Rightarrow$ Adaptive MCMC algorithms, Likelihood-informed subspaces

- Sparse data or expensive high-fidelity simulations
  - With low information content, calibration may struggle
  - $\Rightarrow$ More informative priors/regularization

Summary and Future Thank You

- Represent, quantify and propagate physical model errors
- Parameter estimation $\Rightarrow$ density estimation
- Bayesian machinery to find parameters of the PDFs
- Approximate/empirical likelihoods impose constraints of interest
- Differentiates from data noise; allows model-to-model calibration
- Implemented in UQTk (www.sandia.gov/UQToolkit)
- K. Sargsyan, H. Najm, and R. Ghanem. “On the Statistical Calibration of Physical Models”. *International Journal for Chemical Kinetics*, 47(4): 246-276, 2015.

- Optimal design for maximum information
- Bayesian problem still hard; MCMC, priors, ...
- Hierarchical Bayesian viewpoint
- More intrusive embedding; problem specific

Additional Material

Likelihood construction – variants

- Full Likelihood

$$L(\alpha) = p(D|\alpha) = p(y_{\text{data},1}, \ldots, y_{\text{data},N}|\alpha)$$

- Marginalized Likelihood

$$L(\alpha) = p(D|\alpha) = \prod_{i=1}^{N} p(y_{\text{data},i}|\alpha)$$

- Approximate Bayesian Computation (ABC)
  – seek to satisfy the constraints:
  - $p(y|D)$ is "centered" on the data
  - The width of the distribution $p(y|D)$ is "consistent" with the spread of the data around the nominal model prediction

Full Likelihood

$$L(\alpha) = p(D|\alpha) = \pi_f(y_{\text{data},1}, \ldots, y_{\text{data},N}|\alpha)$$

where:
$\pi_f(\cdot, \alpha)$: N-variate density of the random variable $(f_1, \ldots, f_N)$
with $f_i = f(x_i, \lambda(\alpha))$

Problem: $\pi_f(\cdot)$ is degenerate in general when $N > M$

Consider a case with $M = 1$, $\lambda \sim N(\mu, \sigma^2)$, and $f = \lambda x$
Let $N = 2$, hence $(f_1, f_2) = (\lambda x_1, \lambda x_2)$ for any λ sample
With $f_1/x_1 = f_2/x_2 = \lambda$, (f_1, f_2) are dependent and
$\pi_f(\cdot|\mu, \sigma)$ is non-zero only along the line $f_2 = (x_2/x_1)f_1$

hence $\pi_f(y_{\text{data},1}, y_{\text{data},2}|\mu, \sigma)$ is non-zero only along the line
$y_{\text{data},2}/x_2 = y_{\text{data},1}/x_1$

Marginalized Likelihood

$$L(\alpha) = p(D|\alpha) = \prod_{i=1}^{N} \pi_{f_i}(y_{\text{data},i}|\alpha)$$

where $\pi_{f_i}(\cdot, \alpha)$ is the univariate density of the RV $f_i = f(x_i, \lambda(\alpha))$

Problem: the likelihood has multiple singularities corresponding to α values leading to vanishing marginal variances at each x_i

Gaussian example: Let $f_i \sim \mathrm{N}(\mu_i(\alpha), \sigma_i(\alpha)^2)$, then

$$L(\alpha) = \frac{1}{(2\pi)^{N/2}} \prod_{i=1}^{N} \frac{1}{\sigma_i(\alpha)} \exp\left(\frac{(\mu_i(\alpha) - y_{\text{data},i})^2}{2\sigma_i(\alpha)^2} \right)$$

Multiple singularities, $\sigma_i(\alpha) = 0$, $i = 1, \ldots, N$

Posterior maximization always finds one of these singularities, fitting one point perfectly, while misfitting the rest (⇒ priors)

Approximate Bayesian Computation (ABC)

Employ a kernel density as a pseudo-likelihood to enforce select constraints:

- Uncertain prediction $p(y|D)$ is centered on the data
  - With $\mu_i(\alpha) = \mathrm{E}_\xi[f(x_i, \lambda(\xi; \alpha))]$: minimize $\|\mu_i(\alpha) - y_{\mathrm{data},i}\|_2^2$
- The width of the distribution $p(y|D)$ is consistent with the spread of the data around the nominal model prediction
  - With $\sigma_i(\alpha)^2 = \mathrm{V}_\xi[f(x_i, \lambda(\xi, \alpha))]$:
    minimize $\|(\sigma_i(\alpha) - \gamma|\mu_i(\alpha) - y_{\mathrm{data},i}|)\|_2^2$
  - γ is a factor that specifies the desired match between σ_i and the discrepancy $|\mu_i(\alpha) - y_{\mathrm{data},i}|$, on average

ABC Likelihood

With $\rho(\mathcal{S})$ being a metric of the statistic $\mathcal{S}$, use the kernel function as an ABC likelihood:

$$L_{\text{ABC}}(\alpha) = \frac{1}{\epsilon} K\left(\frac{\rho(\mathcal{S})}{\epsilon}\right)$$

where ϵ controls the severity of the consistency control

Propose the Gaussian kernel density:

$$L_{\epsilon}(\alpha) = \frac{1}{\epsilon\sqrt{2\pi}} \prod_{i=1}^{N} \exp\left(-\frac{(\mu_i(\alpha) - y_{\text{d},i})^2 + (\sigma_i(\alpha) - \gamma|\mu_i(\alpha) - y_{\text{d},i}|)^2}{2\epsilon^2}\right)$$