Using Surrogate Models to Accelerate Bayesian Inverse Uncertainty Quantification

James Rynn

SIAM Student Chapter Conference 2017

james.rynn@manchester.ac.uk

05/05/17

Overview

Bayesian Inverse Problems

Find the unknown θ given n_{obs} observations $\mathbf{D}$, satisfying

$$\mathbf{D} = \mathcal{G}(\theta) + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \Sigma),$$

where

- $\mathbf{D} \in \mathbb{R}^{n_{\text{obs}}}$ is a given vector of **observations**,
- $\mathcal{G} \colon \mathcal{X} \to \mathbb{R}^{n_{\text{obs}}}$ is the **observation operator**,
- $\theta \in \mathcal{X}$ is the **unknown**,
- $\boldsymbol{\eta} \in \mathbb{R}^{n_{\text{obs}}}$ is a vector of **observational noise**.

We treat this as a probabilistic problem and search for a posterior distribution for θ.

Bayesian Inverse Problems

In the finite-dimensional case, from Bayes' Theorem we have

$$\begin{aligned} \pi(\theta|\mathbf{D}) &\propto L(\mathbf{D}|\theta)\ \pi_0(\theta) \\ &\propto \exp\left(-\frac{1}{2}\|\mathbf{D}-\mathcal{G}(\theta)\|_\Sigma^2\right)\ \pi_0(\theta), \end{aligned}$$

Markov Chain Monte Carlo (MCMC) Methods

- We know $\pi(\theta|\mathbf{D})$ up to a constant of proportionality,
- Use MCMC algorithm to generates samples $\theta_1, \theta_2, \ldots, \theta_N$ from the posterior distribution,
- Use these samples to construct Monte Carlo estimates of quantities of interest (means, variances and/or probabilities),
- e.g.,

$$\mathbb{E}_\pi[\phi(\Theta)] = \int_{\mathcal{X}} \phi(s)\pi(s|\mathbf{D})\mathrm{d}s \approx \frac{1}{N}\sum_{i=1}^{N} \phi(\theta_i).$$

Motivation

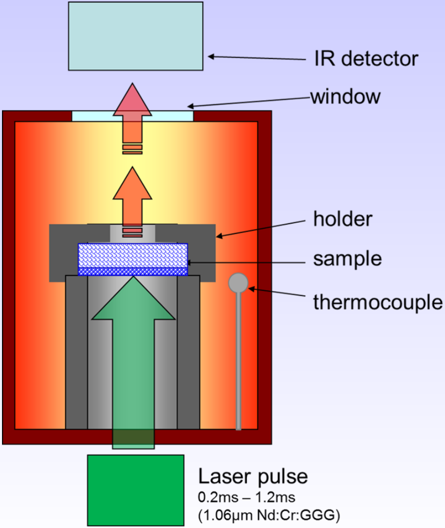
IR detector
window
holder
sample
thermocouple
Laser pulse
0.2ms – 1.2ms
(1.06μm Nd:Cr:GGG)

Motivation

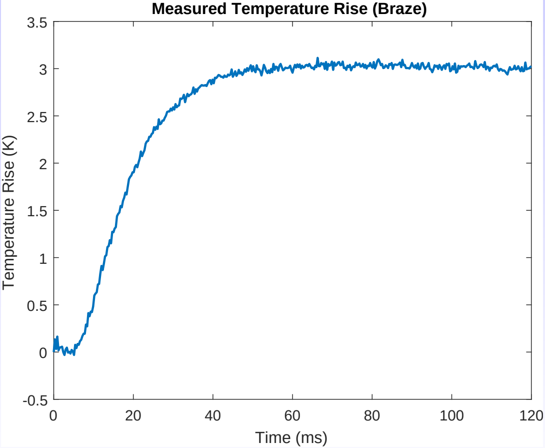
Measured Temperature Rise (Braze)
Temperature Rise (K)
3.5
3
2.5
2
1.5
1
0.5
0
-0.5
0
20
40
60
80
100
120
Time (ms)

Motivation

Unknowns:

- λ — thermal conductivity,
- I — laser intensity,
- k — boundary condition parameter,
- σ — standard deviation of measurement noise.

Example

Consider the one-dimensional steady state heat equation,

$$-\frac{\mathrm{d}}{\mathrm{d}x}\left(\lambda\frac{\mathrm{d}u}{\mathrm{d}x}(x)\right) = f(x), \quad x \in [0, H],$$

with homogeneous Dirichlet boundary conditions,

$$u(0) = u(H) = 0.$$

where $\lambda = e^{\theta}$ is the unknown thermal conductivity.

We wish to find a posterior distribution for λ (equivalently θ), given observations of $u(x)$ at $x_1, x_2, \ldots, x_{n_{obs}} \in [0, H]$.

Example

Here, our observation operator $\mathcal{G}$ is of the form

$$\mathcal{G}(\theta) = (u(x_1;\theta), u(x_2;\theta), \dots, u(x_{n_{obs}};\theta))^T,$$

and approximated by $\mathcal{G}_h$ given by

$$\mathcal{G}_h(\theta) = (u_h(x_1;\theta), u_h(x_2;\theta), \dots, u_h(x_{n_{obs}};\theta))^T,$$

where u_h is the finite element solution to the ODE on a mesh of width h.

Note: For each value of θ, to evaluate $\mathcal{G}_h$ we are required to compute a FEM solve.

Random Walk Metropolis Hastings Algorithm (FEM)

Algorithm 1: RWMH Algorithm

set initial state $X^{(0)} = \theta_0$
for $m = 1, 2, \ldots, N$ **do**
  draw proposal
  evaluate likelihood by **computing** $\mathcal{G}_h$ **(expensive!)**
  compute acceptance probability α
  accept proposal with probability α
output chain $X = (\theta_0, \theta_1, \ldots, \theta_N)$

Here $N \gg 10^5$.

Results

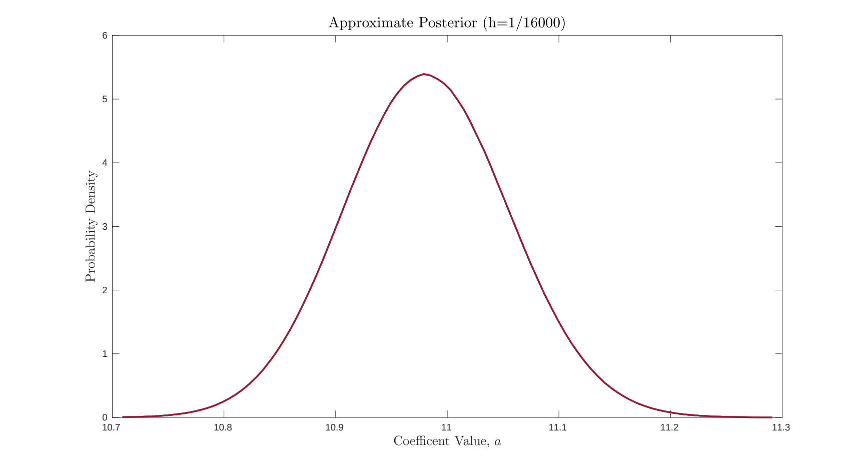
Approximate Posterior (h=1/16000)
Probability Density
Coefficent Value, a
0
1
2
3
4
5
6
10.7
10.8
10.9
11
11.1
11.2
11.3

Stochastic Galerkin Finite Element Method

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and consider the problem

$$-\frac{\mathrm{d}}{\mathrm{d}x}\left(e^{\theta(\omega)}\frac{\mathrm{d}u}{\mathrm{d}x}(x,\omega)\right) = f(x), \quad x \in [0,H], \quad \omega \in \Omega,$$

with homogeneous Dirichlet boundary conditions,

$$u(0,\omega) = u(H,\omega) = 0, \quad \omega \in \Omega.$$

Assuming θ is of the form

$$\theta(\omega) = \theta(\xi(\omega)),$$

we can transform this into a parametric equation on $[0,H] \times \xi(\Omega)$.

Stochastic Galerkin Finite Element Method

Parametric form:

$$-\frac{\mathrm{d}}{\mathrm{d}x}\left(e^{\theta(y)}\frac{\mathrm{d}u}{\mathrm{d}x}(x,y)\right) = f(x), \quad x \in [0,H], \quad y \in \Gamma := \xi(\Omega),$$

with homogeneous Dirichlet boundary conditions,

$$u(0,y) = u(H,y) = 0, \quad y \in \Gamma.$$

Construct a stochastic Galerkin FEM solution u_{hP} on a finite dimensional subspace of $L^2(\Gamma, \mathcal{H}_0^1(D))$ of size $(P+1) \times \frac{1}{h}$.

Random Walk Metropolis Hastings Algorithm Algorithm (SGFEM)

Algorithm 2: RWMH Algorithm with SGFEM Surrogate

compute SGFEM solution u_{hP}
set initial state $X^{(0)} = \theta_0$
for $m = 1, 2, \ldots, N$ **do**
    draw proposal
    evaluate likelihood by **evaluating** $\mathcal{G}_{hP}$ **(cheap!)**
    compute acceptance probability α
    accept proposal with probability α
output chain $X = (\theta_0, \theta_1, \ldots, \theta_N)$

Here $N \gg 10^5$.

Posterior Convergence in N (Number of Samples)

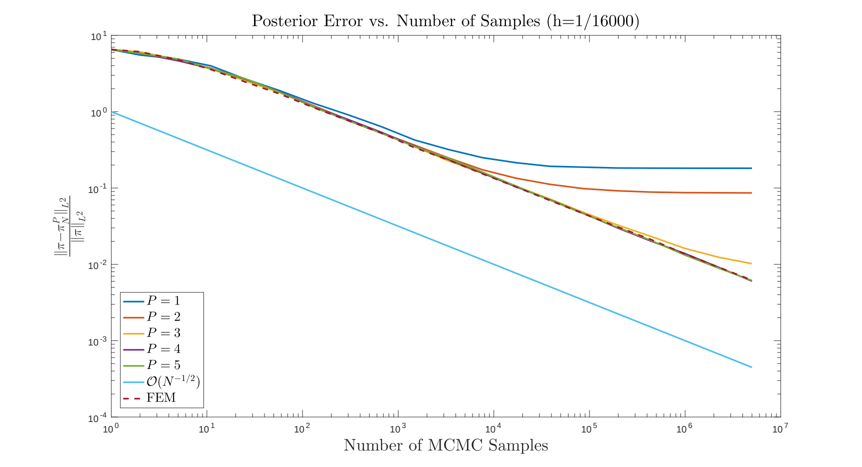

Posterior Convergence in P (Polynomial Degree)

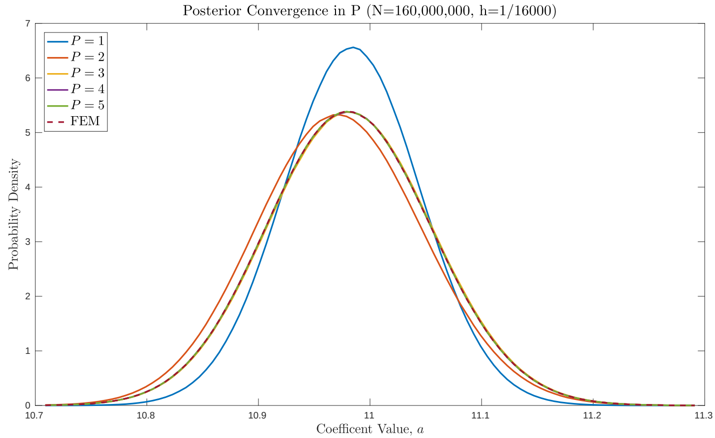

Time vs Error

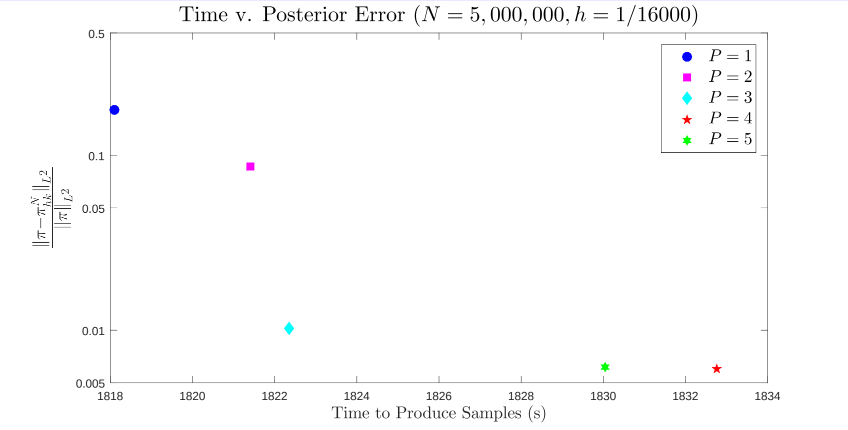

Future Work

- More realistic forward problem:
  - time-dependent PDE,
  - second spatial dimension,
  - multiple random variables,
- Error analysis,
- More sophisticated MCMC algorithm.

References

I. Babuška, R. Tempone, and G. Zouraris, *Galerkin Finite Element Approximations of Stochastic Elliptic Partial Differential Equations*, SIAM Journal on Numerical Analysis, 42 (2004), pp. 800–825.

S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo*, Chapman and Hall/CRC, 2011.

A. Stuart, *Inverse Problems: A Bayesian Perspective*, Acta Numer., 19 (2010), pp. 451–559.

L. Wright, L. Chapman, and D. Partridge, *Laser Flash Experiment on Layered Materials: Parameter Estimation and Uncertainty Evaluation*.
National Physical Laboratory (NPL), Internal Communication.