Quantifying Uncertainty in the Laser Flash Experiment

Simon Cotter, Catherine Powell, James Rynn, Louise Wright

Manchester Numerical Analysis

National Physical Laboratory

The Laser Flash Experiment

The laser flash experiment is used to determine the thermal conductivity of a material. A sample of the material is placed in a furnace and hit with a laser pulse. The temperature at the other end of the material is then measured at certain time points using an IR sensor.

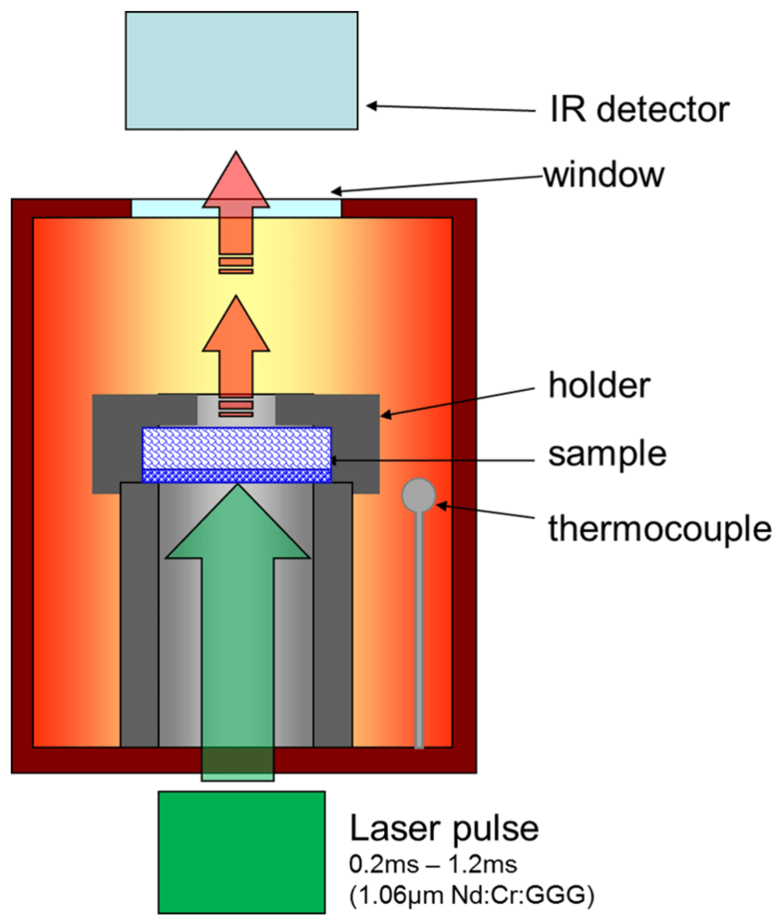

For a full description see [1], for example.

Uncertainty Quantification

In inverse uncertainty quantification, our aim is to recover a probability distribution for unknowns in a system's model using noisy readings of the system at some points in space and time.

We suppose our data is of the form

$$y = \mathcal{G}(\theta) + \eta, \quad \eta \sim \mathcal{N}(0, \Sigma), \tag{1}$$

where θ is the true value of the unknown we are trying to build a probability distribution for, $\mathcal{G}$ is the so-called *observation operator* (which maps values of the unknown into data) and η is observational noise. Here we assume the noise is Gaussian with mean zero and covariance matrix Σ.

We use a simple one-dimensional model with time dependence to model the temperature of the sample at the point x and time t, denoted $u(x,t)$. In this case, the data points y are temperature measurements at the end of the sample at the measurement times and the observation operator $\mathcal{G}$ gives the solution of the transient heat equation

$$\rho c_p \frac{\partial u}{\partial t}(x,t) = \lambda \frac{\partial^2 u}{\partial x^2}(x,t) + Q(x,t) \tag{2}$$

at these points for given values of the parameters λ (thermal conductivity), I (laser intensity) and k (convection coefficient). We are interested in finding the thermal conductivity of the material and so we must somehow invert (1) whilst filtering out the noise.

The Bayesian Approach

In the 18th century, Bayes and Laplace formulated what we call Bayes' rule,

$$\mathbb{P}(\theta|y)\mathbb{P}(y) = \mathbb{P}(y|\theta)\mathbb{P}(\theta).$$

This simple formula allows us to combine our prior knowledge about an unknown, θ, along with data, y, to produce an informed probability distribution for θ given y.

This probability distribution is called the *posterior* and we use Bayes' rule to define it as

$$\mathbb{P}(\theta|y) \propto \mathbb{P}(y|\theta)\mathbb{P}(\theta), \tag{3}$$

where the likelihood function, $\mathbb{P}(y|\theta)$, follows from (1),

$$\mathbb{P}(y|\theta) = \mathcal{N}(\mathcal{G}(\theta) - y, \Sigma). \tag{4}$$

Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) methods are a flexible class of statistical algorithms which can be used to explore probability distributions.

These algorithms allow us to generate samples from a distribution where the probability density function is known only up to a constant of proportionality. Notice that this is exactly the situation described in (3).

Samples are generated by producing a Markov chain $X = \left(X^{(0)}, X^{(1)}, \dots, X^{(M)}\right)$ with stationary density equal to the desired (posterior) density, π.

The Metropolis-Hastings (MH) Algorithm, originally defined in the 1970s is as follows:

Table 1: The Metropolis–Hastings Algorithm.

```
set X^(0) = X_0 (initial state)
for m = 1, 2, ..., M do
    draw Y ~ q(X^(m-1), ·) (sample from proposal distribution, dependent on current state)
    set α(X^(m-1), Y) = min{1, [π(Y)/π(X^(m-1))] [q(X^(m-1), Y)/q(Y, X^(m-1))]} (compute acceptance probability)
    draw U ~ U[0,1]
    if U ≤ α(X^(m-1), Y) then
        set X^(m) = Y (accept proposal)
    else
        set X^(m) = X^(m-1) (reject proposal)
output X = (X^(0), X^(1), ..., X^(M)) (output Markov chain)
```

The Forward Problem

To generate samples using an MCMC algorithm such as the (Random Walk) Metropolis–Hastings algorithm, we must be able to evaluate the forward operator, $\mathcal{G}$. This is referred to as the *forward problem*.

We use the Crank-Nicolson (CN) implicit finite difference method to approximate the solution of (2) on a uniform space–time mesh, introducing approximation error into our solution. However, the CN method can be shown to be second order in both space and time as well as consistent [2], meaning that the resulting approximation converges to the true solution as the mesh is refined. Furthermore, it can be shown that, under certain conditions, the approximation error in our approximate posterior reduces similarly under refinement [3].

Below is a plot of the temperature of a sample of braze at the end point at the measurement times, produced using the CN method. Also plotted (in blue) is some experimental data. Optimisation was used to find the values of the unknown parameters λ, I and k which allowed the CN approximation to match the true data best in the least squares sense. These optimised values for the unknown parameters give a good starting value $X^{(0)}$ for our MCMC routine, as we expect them to lie in a region of high probability density.

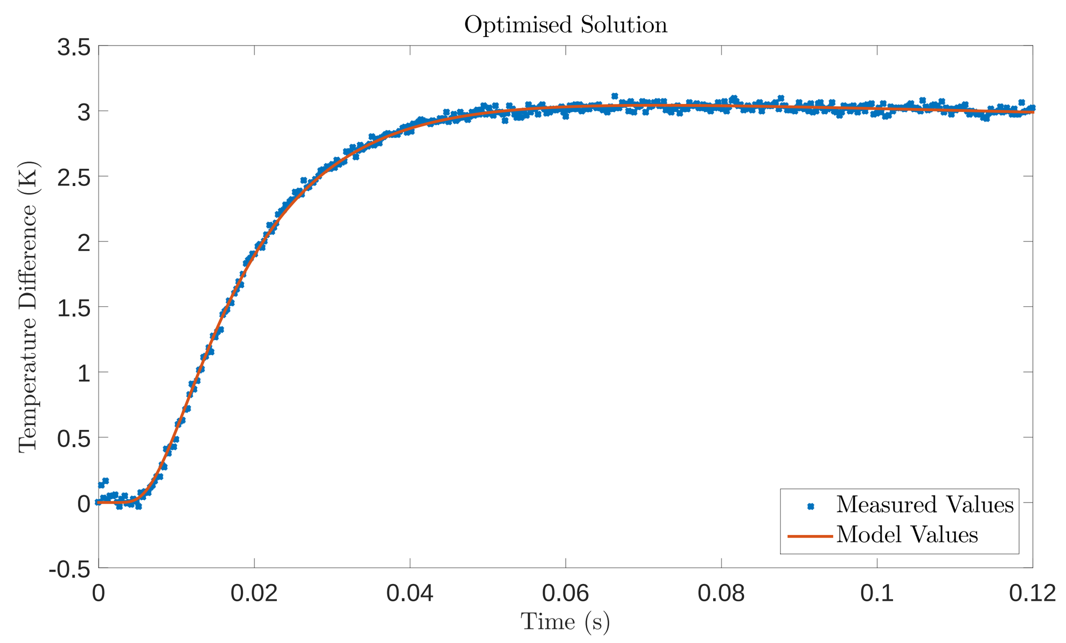

References

[1] L. Wright, X.S. Yang, C. Matthews, L. Chapman and S. Roberts, *Parameter Estimation from Laser Flash Experiment Data.*, Computational Optimization and Applications in Engineering and Industry, Springer, 2011

[2] K. Morton and D. Mayers, *Numerical Solution of Partial Differential Equations*, Cambridge University Press, New York (2005) Second Edition,

[3] S. Cotter, M. Dashti and A. Stuart, *Approximation of Bayesian Inverse Problems for PDEs.*, SIAM Journal on Numerical Analysis 48 (2010): 322–345.

Results

Firstly, we examine how the posterior mean of an MCMC sample is affected by the level of mesh refinement. This is important because the more refined the mesh, the more costly each MCMC sample is to produce. Therefore, finding the coarsest mesh such that the approximation error is negligible is computationally beneficial. The plot below shows how the posterior mean of the thermal conductivity λ varies as the temporal mesh is refined.

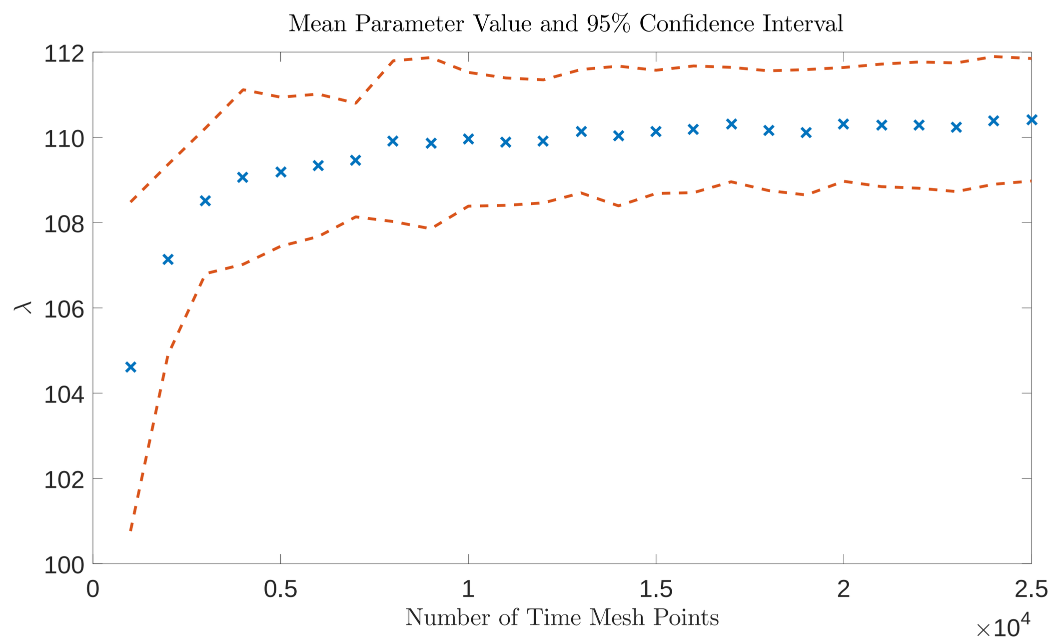

Secondly, normalised histograms of the states visited by a (long) Markov chain are plotted to give an approximation to the posterior density. An approximation of the posterior density for the thermal conductivity λ given the data is plotted below, along with the corresponding prior density.

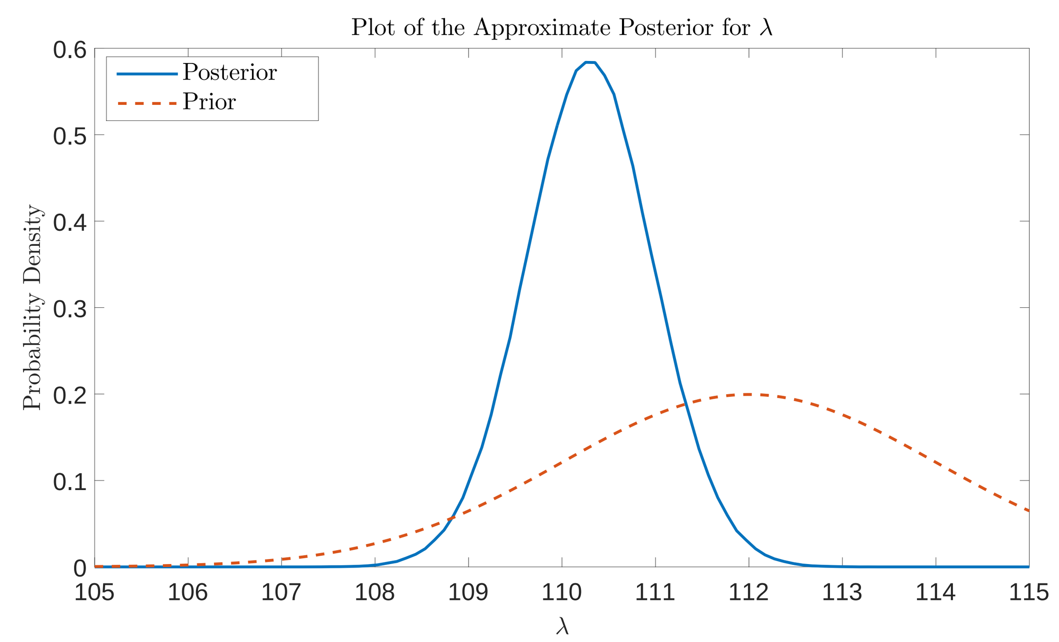

We can see from the above plot that by incorporating the data into our estimate for the parameter we have reduced the variance compared to the prior and thus the uncertainty regarding the value of that parameter.

Future Work

In order to approximate the posterior well, we must generate a large number M of samples. Reducing the time taken to generate these samples therefore is a key issue. Two possible candidates for reducing the cost of the forward solve are reduced basis or surrogate models.