Reduced Basis Model Reduction for Statistical Inverse Problems with applications in Tsunami Modelling

J.H.S. de Baar [a] , **B. Harding** [a], **M. Hegland** [a] and <u>C. Oehmigara</u> [a]

[a] *Mathematical Sciences Institute, Australian National University, Australia*
Email: chaitanya.shettigara@anu.edu.au

Abstract: Tsunami modelling involves the problem of forecasting the effect that a tsunami has on a coastal area, given the recordings of the water height at several buoys. This forecast has two consecutive steps: (1) estimating the initial condition from the buoy data, and (2) propagating that initial condition to forecast the effect on the coast. In order to be effective, the forecast should be completed within reasonable time. This is challenging, as the first step involves solving a high-dimensional inverse problem, at significant computational cost. Therefore, there is a need to search for methods that accelerate the solution of the inverse problem.

One approach is to reduce the number of dimensions of the inverse problem through parameter reduction. This involves determining a reduced parameter basis, which enables us to approximate the model data to a high level at accuracy while reducing the dimensionality of the inverse problem. This approach introduces an 'offline' and an 'online' stage: during the offline stage, before a tsunami occurs, we allocate significant computational cost to determining the reduced basis. As soon as a tsunami occurs and buoy data arrives, we address the inverse problem, which is now lower-dimensional as a result of our efforts of determining the reduced basis in the offline stage.
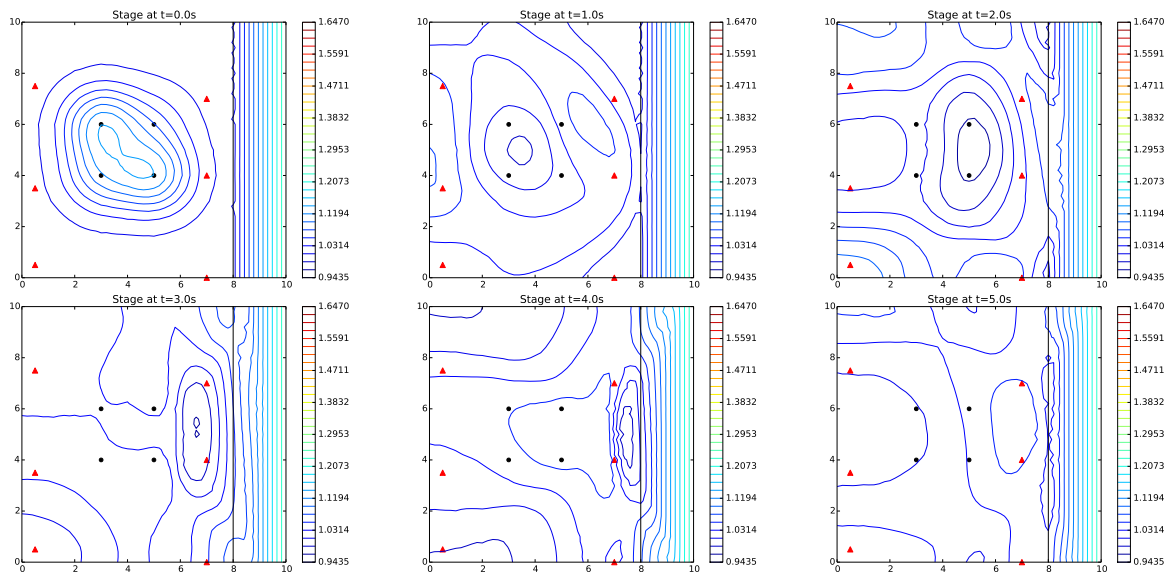
Figure 1. Time evolution of a model tsunami problem, with a coast line on the right-hand side.

In this work, we consider the model tsunami problem illustrated in Figure 1, which shows how the water height evolves over time, starting from an arbitrary initial condition. The initial condition is parameterised by the input parameters $\mathbf{p}$, which result in the buoy output $\mathbf{y}(\mathbf{p})$. We focus on finding a reduced basis, such that the reduced model $\mathbf{y}(Q_k\mathbf{p})$ is an accurate estimate of the full model $\mathbf{y}(\mathbf{p})$, with Q_k an orthogonal projection corresponding to the reduced basis. We use a greedy algorithm to approximate the optimal reduced basis.

We add a novel element to the greedy algorithm by quantifying how well the reduced model approximation $\mathbf{y}(Q_k\mathbf{p})$ can recover the full model $\mathbf{y}(\mathbf{p})$ for arbitrary $\mathbf{p}$. Results for up to $N = 25$ input parameters show that we can efficiently approximate the full model by using a reduced number of parameters. We find that the relative amount of parameters that we need for an accurate approximation decreases when the number of input parameters increases.

Keywords: Statistical inverse problems, model reduction

1 INTRODUCTION

On March 11, 2011, the east coast of Japan was hit by a tsunami caused by an earthquake 70 kilometres off shore, resulting in 20,000 people dead or missing. It is of great interest to understand this destructive phenomenon, which requires reconstructing the source.

Waveforms of the tsunami were recorded by a number of observation systems in the region surrounding the source, including ocean-bottom pressure gauges and GPS buoys. Determining the initial water height from this data involves finding the solution to an inverse problem. This problem has been tackled using Time Reversal Imaging techniques by Hossen et al. (2015), and with Least Squares Inversion by Saito et al. (2011). A similar problem using Kalman filters has been addressed by Verlaan et al. (2005) and Heemink et al. (1997).

We seek to improve upon the least squares inversion method by posing the problem as a statistical inverse problem, and to improve efficiency by applying reduced basis methods. We follow the method introduced in Lieberman et al. (2010).

2 METHOD

2.1 Statistical inverse problem

Consider a parameterised model $\mathcal{M} : \mathcal{P} \to \mathcal{Y}$ of some phenomenon, where given a parameter value $\mathbf{p}$ in the parameter space $\mathcal{P} \subset \mathbb{R}^N$, the model predicts $\mathbf{y}(\mathbf{p})$ in the output space $\mathcal{Y} \subset \mathbb{R}^M$, and we observe noisy output $\mathbf{y}_d \in \mathcal{Y}$. The inverse problem is to estimate $\mathbf{p}$ given $\mathbf{y}_d$, a problem that is often ill-posed. In the deterministic setting, regularisation and optimisation techniques are used to give a single point estimate of the parameter. In contrast, a statistical formulation of the inverse problem returns a probability density function over parameter space describing the relative likelihood of observation-consistent parameters, called the 'posterior distribution', $\pi(\mathbf{p}|\mathbf{y}_d)$.

Here we infer properties of the probability distribution of model parameters $\mathbf{p} \in \mathbb{R}^N$ from observations $\mathbf{y}_d \in \mathbb{R}^m$. We adopt a Bayesian approach where a prior probability density $\pi(\mathbf{p})$ of $\mathbf{p}$ is given. We also assume that our model of the observations provides us with a likelihood function $\pi(\mathbf{y}_d \mid \mathbf{p})$ of the data $\mathbf{y}_d$. The data $\mathbf{y}_d$ is treated as a random vector $\mathbf{Y}$ with probability distribution $\pi(\mathbf{y} \mid \mathbf{p})$ and we denote the expectation as

$$\mathbf{y}(\mathbf{p}) := E(\mathbf{Y} \mid \mathbf{p}).$$

Bayes' law leads to a formula for the posterior probability density of $\mathbf{p}$ as

$$\pi(\mathbf{p} \mid \mathbf{y}_d) \propto \pi(\mathbf{p}) \, \pi(\mathbf{y}_d \mid \mathbf{p}). \tag{1}$$

This probability density could then be further explored for example, in order to find marginals or moments of the posterior or even the MAP (maximum a posteriori) estimate which is of the form

$$\mathbf{p}_{\mathrm{MAP}} = \arg\max_{\mathbf{p} \in \mathbb{R}^N} \pi(\mathbf{p} \mid \mathbf{y}_d). \tag{2}$$

In any case, the exploration of the posterior $\pi(\mathbf{p} \mid \mathbf{y}_d)$ requires its numerical evaluation typically many times. In our case each evaluation of the likelihood (and thus the posterior) requires an expensive numerical simulation.

2.2 Reduced basis method

In order to decrease the computational load required we reduce the parameter space from $\mathbb{R}^N$ to $\mathbb{R}^{N_r}$ as any exploration or computation in a lower dimensional space is cheaper. For this we generate a sequence of parameter vectors $\mathbf{p}_1, \mathbf{p}_2, \dots$ in $\mathbb{R}^N$ which are defined using a greedy algorithm as in Lieberman et al. (2010)

$$\mathbf{p}_{k+1} = \arg\max_{\mathbf{p} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{y}(\mathbf{p}) - \mathbf{y}(Q_k \mathbf{p})\|^2 + \beta \log \pi(\mathbf{p}) \right), \tag{3}$$

where the $Q_k : \mathbb{R}^N \to \mathbb{R}^N$ are orthogonal projections with $\mathrm{range}(Q_k) = \mathrm{span}\{\mathbf{p}_1, \dots, \mathbf{p}_k\}$ for $k \geq 1$ and $Q_0 := 0$. This choice introduces the new $\mathbf{p}_{k+1}$ to make sure that a large number of observations $\mathbf{y}_d$ can be approximated in the range of Q_{k+1} while exploring parameter vectors which are sufficiently likely according to the prior.

The main outcome of this paper is to show that this choice of $\mathbf{p}_k$ gives good approximations for observations $\mathbf{y}_d$ generated by a tsunami model for relatively small N_r.

3   TSUNAMI MODELLING APPLICATION

Given an initial water displacement, the wave propagation at subsequent times can be modelled with the shallow water wave equations. These are commonly expressed as a non-linear system of hyperbolic conservation laws for mass and momentum (Roberts et al., 2015)

$$
\begin{aligned}
\frac{\partial h}{\partial t} + \frac{\partial (hu)}{\partial x_1} + \frac{\partial (hv)}{\partial x_2} &= 0 \\
\frac{\partial (hu)}{\partial t} + \frac{\partial}{\partial x_1}\left(u^2 h + \frac{1}{2}gh^2\right) + \frac{\partial (huv)}{\partial x_2} &= -gh\frac{\partial b}{\partial x_1} - ghB_{x_1} \\
\frac{\partial (hv)}{\partial t} + \frac{\partial (huv)}{\partial x_1} + \frac{\partial}{\partial x_2}\left(v^2 h + \frac{1}{2}gh^2\right) &= -gh\frac{\partial b}{\partial x_2} - ghB_{x_2}.
\end{aligned}
\tag{4}
$$

Here g is the acceleration due to gravity, h is the water depth, u and v are the depth-averaged velocities in the x_1 and x_2 (longitudinal and latitudinal) coordinates respectively, b is the ocean floor elevation and (B_{x_1}, B_{x_2}) is the bed friction given by

$$
B_{x_1} = \frac{u\eta^2\sqrt{u^2 + v^2}}{h^{4/3}} \text{ and } B_{x_2} = \frac{v\eta^2\sqrt{u^2 + v^2}}{h^{4/3}}
$$

where η is the Manning friction coefficient. We conducted our experiments using the open source hydrodynamic modelling software ANUGA, developed by the Australian National University (ANU) and Geoscience Australia (GA), see for example Roberts et al. (2015); Nielsen et al. (2005). ANUGA solves the shallow water wave equations using a finite-volume numerical scheme.

The computational domain is a 10 m $\times$ 10 m square with reflective boundary conditions. The ocean floor has constant Manning friction coefficient of $\eta = 0.1$ and a linearly sloping bathymetry representing a shoreline, so that the unperturbed water depth is $h = 1$ m at $x_1 = 0$ and $h = 0$ m at $x_1 = 8$.

3.1   Model input

The initial sea surface displacement $u_0(x_1, x_2)$ is parameterised by a superposition of N Gaussian basis functions given by

$$
u_0(x, y) = \sum_{i=1}^{N} p_i\, \gamma_i(x_1, x_2) \text{ where } \gamma_i(x_1, x_2) = \exp\left(-\frac{|x_1 - \bar{x}_{1,i}|^2}{2L^2} - \frac{|x_2 - \bar{x}_{2,i}|^2}{2L^2}\right).
$$

The centres of the Gaussian basis functions are at uniformly spaced grid points $(\bar{x}_{1,i}, \bar{y}_{2,i})$ within the computational domain with corners at $(x_1, x_2) = (3, 4), (3, 6), (5, 4)$ and $(5, 6)$, as illustrated by the black dots in the centres of Figures 1, 2 and 3. In order to investigate the effect of the number of input parameters, we have set up the problem for different cases using $N = 4, 6, 9, 15, 25$. The input parameters $\mathbf{p} = (p_1, p_2, \ldots, p_N)$ are the coefficients of the basis functions, while the kernel size $L = 1$.

3.2   Model output

The water flow is evolved using ANUGA for 5 seconds. A typical time evolution of the water height is illustrated by Figure 1 (see extended abstract). Starting from the initial condition, a wave forms which hits the beach after some 1 to 2 s, depending on the location.

The output $\mathbf{y}(\mathbf{p})$ is given by the time series data of water height extracted at six different points, representing observation buoys, which are located at $(x_1, x_2) = (0.5, 0.5), (0.5, 3.5), (0.5, 7.5), (7, 0), (7, 4)$ and $(7, 7)$. The locations of these buoys are illustrated by the red triangles in Figure 1, 2 and 3.

3.3   Outlook: inverse problem

Because the motivation for creating a reduced basis is solving the inverse problem to infer the probability distribution of the input parameters $\pi(\boldsymbol{p}|\boldsymbol{y}_d)$ from the observed data $\boldsymbol{y}_d$, we provide a brief outlook.
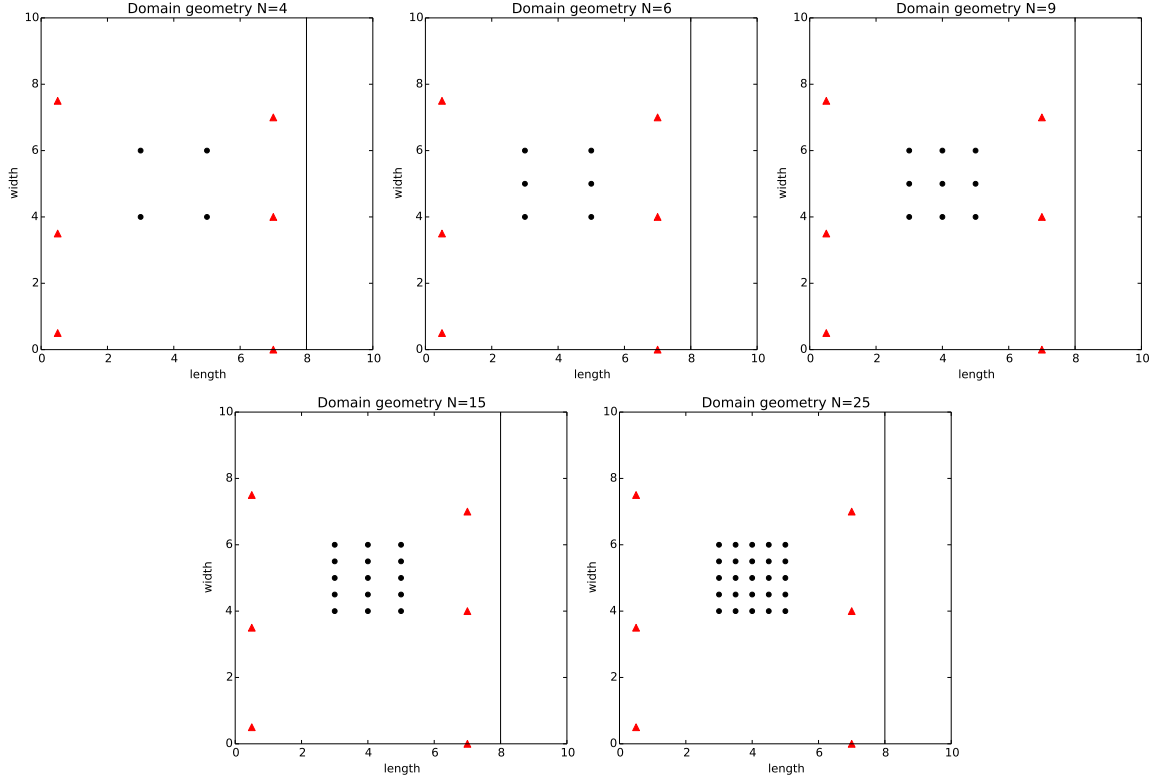
Figure 2. Geometry of the computational domain. The centres of the Gaussian basis functions are represented by circles. The observation buoys are shown as red triangles, and the shoreline by the vertical line at $x = 8$.

In the ideal case, the prior would be constructed by carefully analysing all historical tsunami data for the region of interest. In those cases where such data is not available, one might choose a uniform prior. However, since we assume some smoothness, we define the prior

$$\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbf{S}), \tag{5}$$

with prior covariance matrix

$$S_{ij} \propto \exp \left\{ \frac{-\|\overrightarrow{x}_i - \overrightarrow{x}_j\|^2}{2b^2} \right\}$$

using smoothness coefficient b. We assume the likelihood

$$\mathbf{y}_d \sim \mathcal{N}\left(\mathbf{y}(\mathbf{p}), \sigma^2 \mathbf{I}\right), \tag{6}$$

of observing the experimental data $\mathbf{y}_d$ for a given input $\mathbf{p}$, with σ the measurement uncertainty. From the prior (5) and the likelihood (6) we can compute the desired posterior (1).
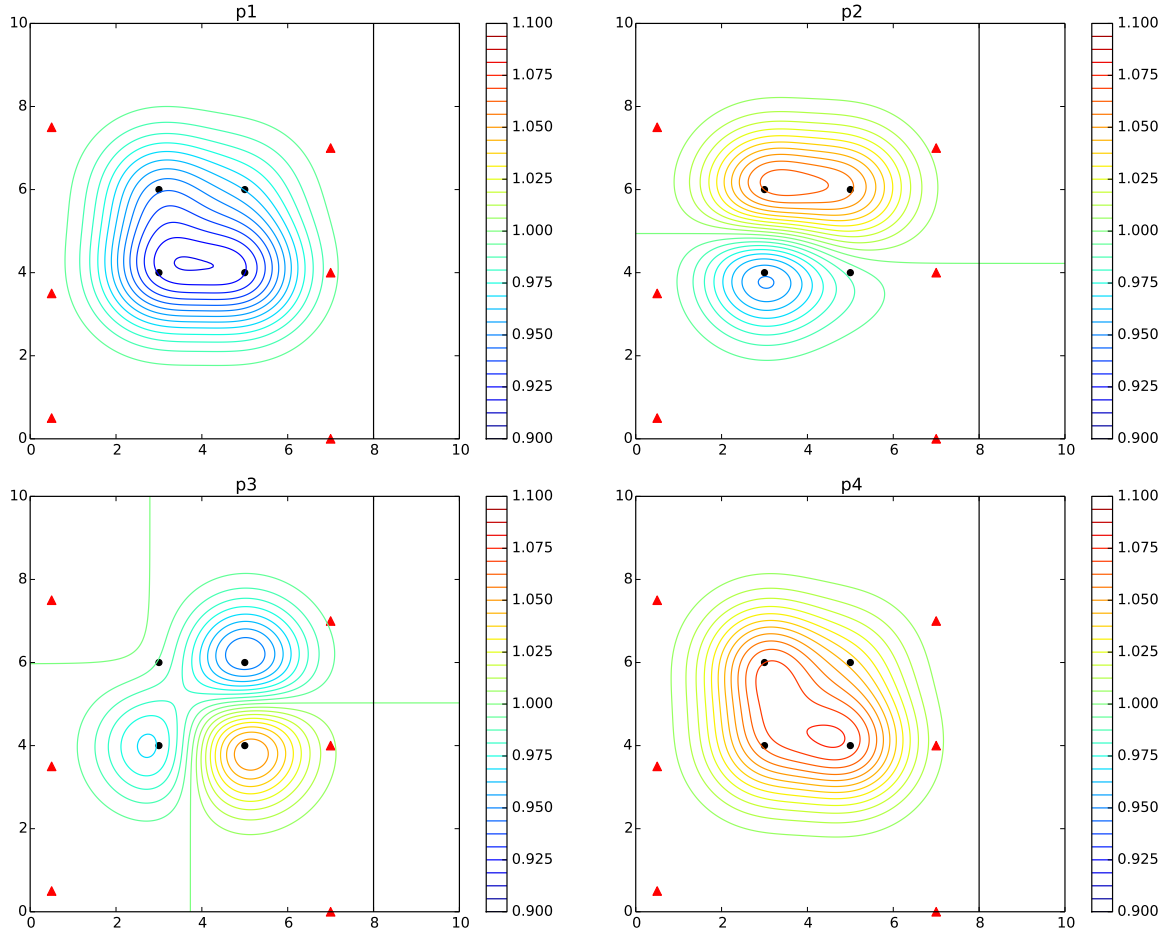
4 RESULTS

In this section we investigate how well we are able to approximate the data $\mathbf{y}(\mathbf{p})$ by reduced models $\mathbf{y}(Q_k \mathbf{p})$ for the specific example described in the previous section and for 5 different initial conditions using $N = 4, 6, 9, 15$ and 25 parameters respectively. Here we assume that the parameter space is $\mathcal{P} = \{\mathbf{p} \mid \|\mathbf{p}\| = 0.1\}$.

We define approximations to the parameter vectors $\mathbf{p}_k$ defined in equation (3) by

$$\mathbf{p}_{k+1} = \arg \max_{\mathbf{p} \in \mathcal{M}_k} \frac{1}{2} \|\mathbf{y}(\mathbf{p}) - \mathbf{y}(Q_k \mathbf{p})\|^2. \tag{7}$$

where each $\mathcal{M}_k$ contains m independent samples of the uniform distribution over $\mathcal{P}$ and the the elements of $\mathcal{M}_k$ and $\mathcal{M}_j$ are independent for $j \neq k$. In the experiment we chose $m = 1000$. Note that this allows us to

Figure 3. Initial water height given by the four reduced basis vectors calculated for $N = 4$.

drop the $\beta \log \pi(\mathbf{p})$ term that appears in equation 3. We denote the approximations of $\mathbf{p}_k$ by the same symbol $\mathbf{p}_k$ for simplicity.

Figure 3 illustrates the resulting basis vectors for $N = 4$ input parameters. Note that the first basis vector results in a more global water height, while the second and third basis vectors provide more refinement. The fourth basis vector only serves to complete the basis.

We are now interested in how well our reduced model approximation $\mathbf{y}(Q_k \mathbf{p})$ can recover the full model $\mathbf{y}(\mathbf{p})$. To get a qualitative impression of the accuracy of the approximation, Figure 4 shows the time series data $\mathbf{y}$ at the different buoys. For an arbitrary initial condition, the red data is created using a full basis of $N = 25$ input parameters, while the blue data is created using a reduced basis of $N_r = 13$ parameters. The blue lines provide a good approximation of the red lines for each of the buoys.

To quantify how well the reduced model approximation $\mathbf{y}(Q_k \mathbf{p})$ can recover the full model $\mathbf{y}(\mathbf{p})$ for arbitrary $\mathbf{p}$, we use the relative L_2 norm

$$\frac{\sqrt{\int_{\mathcal{P}} \|\mathbf{y}(\mathbf{p}) - \mathbf{y}(Q_k \mathbf{p})\|^2 \, d\omega}}{\sqrt{\int_{\mathcal{P}} \|\mathbf{y}(\mathbf{p})\|^2 \, d\omega}},$$

where ω is a probability measure over $\mathcal{P}$. We approximate this relative root mean squared error by Monte Carlo quadrature as

$$\rho_k = \frac{\sqrt{\frac{1}{m} \sum_{\mathbf{p} \in \mathcal{M}_{k+1}} \|\mathbf{y}(\mathbf{p}) - \mathbf{y}(Q_k \mathbf{p})\|^2}}{\sqrt{\frac{1}{m} \sum_{\mathbf{p} \in \mathcal{M}_{k+1}} \|\mathbf{y}(\mathbf{p})\|^2}}.$$

We now compute these ρ_k for five examples with $N = 4, 6, 9, 15, 25$ respectively. The results are plotted in

Figure 5. This figure shows that for more complex problems (higher N) one requires more basis functions to get the same error, but also that the relative number of basis functions required decreases, as is illustrated in Figure 6. In the case of $N = 25$, a reduced model error below $\rho_k = 0.1$ is achieved with $N_r = 13$ basis vectors, a model reduction of 48%. The output buoy data for both the full and reduced models are shown in Figure 4.
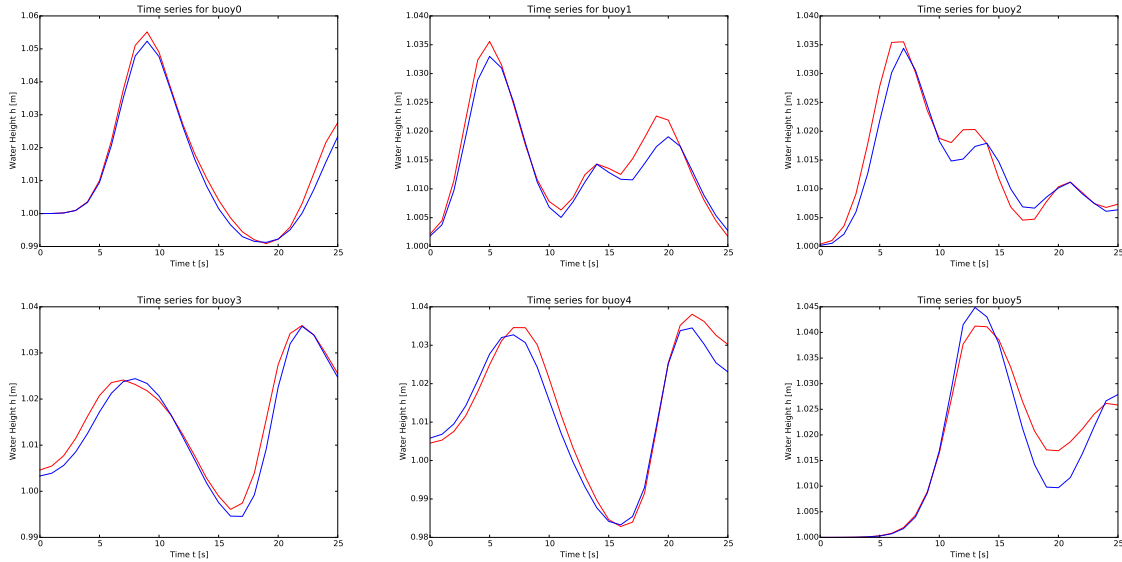
Figure 4. Comparison of time series data from the six observation buoys for full model parameter dimension $N = 25$. The full model data is shown in red and reduced model with $N_k = 13$ shown in blue. This reduced basis has root mean square error $\rho_{13} = 8.97 \times 10^{-2}$
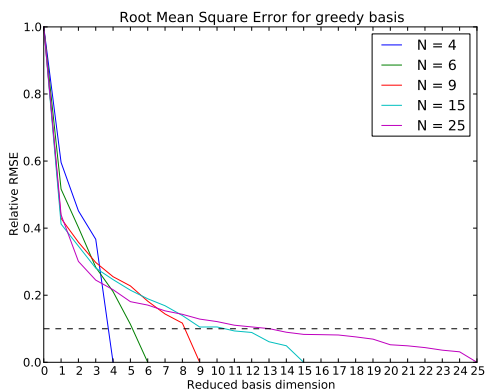
Figure 5. Relative root mean squared error ρ_k vs. reduced basis dimension k for $N = 4, 6, 9, 15, 25$. The horizontal line at $\rho_k = 0.1$ illustrates that the number of reduced basis vectors required to achieve a given accuracy increases with N, however the relative number required decreases.
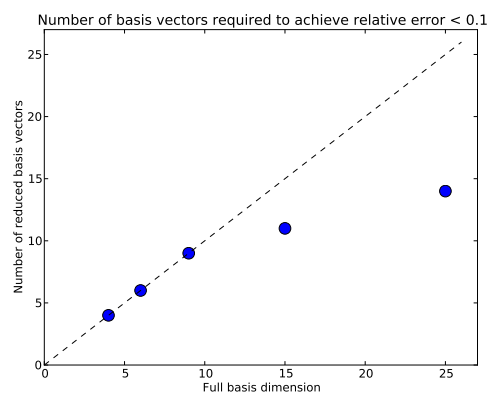
Figure 6. Number of reduced basis vectors required to achieve accuracy $\rho_k < 0.1$ vs full basis dimension. This indicates the relative number of reduced basis vectors required to achieve a given accuracy decreases as the full basis dimension N increases.

5 CONCLUSION

The results presented show that accurate shallow water wave data predictions can be made using a reduced model method. This is promising for the solution of the inverse problem of recovering initial water heights from data, as the search space for the optimisation problem (2) can be as much as halved.

Future work will further investigate the relative basis reduction for full parameter space dimension N, efficiency gains and error sensitivity for solving the inverse problem using reduced basis methods, and extending this analysis to a more realistic bathymetry.

REFERENCES

Heemink, A., K. Bolding, and M. Verlaan (1997). Storm surge forecasting using kalman filtering. *Journal of the Meteorological Society of Japan Series 2 75*, 195–208.

Hossen, M., P. Cummins, S. Roberts, and S. Allgeyer (2015). Time reversal imaging of the tsunami source. *Pure and Applied Geophysics 172*(3-4), 969–984.

Lieberman, C., K. Willcox, and O. Ghattas (2010, August). Parameter and state model reduction for large-scale statistical inverse problems. *SIAM J. Sci. Comput. 32*(5), 2523–2542.

Nielsen, O., S. Roberts, D. Gray, A. McPherson, and A. Hitchman (2005). Hydrodynamic modelling of coastal inundation. *MODSIM 2005 International Congress on Modelling and Simulation*, 518–523.

Roberts, S., O. Nielsen, D. Gray, J. Sexton, and G. Davies (2015). *ANUGA User Manual*. Geoscience Australia.

Saito, T., Y. Ito, D. Inazu, and R. Hino (2011). Tsunami source of the 2011 tohoku-oki earthquake, japan: Inversion analysis based on dispersive tsunami simulations. *Geophysical Research Letters 38*(7), n/a–n/a. L00G19.

Verlaan, M., A. Zijderveld, H. de Vries, and J. Kroos (2005). Operational storm surge forecasting in the Netherlands: developments in the last decade. *Philosophical Transactions A 363*, 14411453.