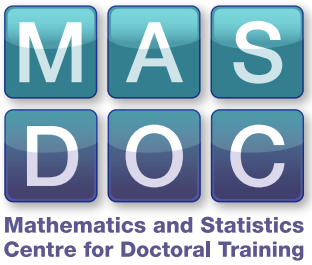

Ensemble Based Methods for Geometric Inverse Problems

by

Neil Kumar Chada

Thesis

Submitted for the degree of

Doctor of Philosophy

Mathematics Institute

The University of Warwick

July 2018

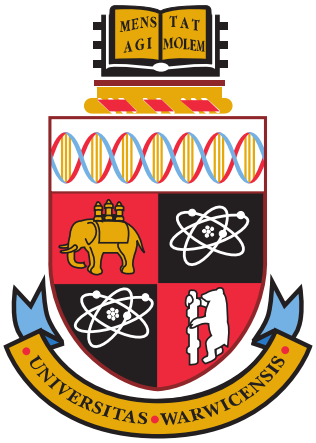

Contents

List of Tables

List of Figures

Acknowledgments

A PhD is a long journey and process where numerous people have helped in different ways. First and foremost I would like to thank my primary supervisor Prof. Andrew Stuart. It was Prof. Stuart who introduced me to the interesting areas of Bayesian inverse problems and data assimilation, and who helped with attaining an EPSRC industrial studentship. Under his supervision I developed and enhanced both my research and interpersonal skills, and for that I am deeply indebted to him. I could not have asked for a better supervisor.

I have also had the pleasure of working further with a number of individuals. Firstly to Dr. Claudia Schillings and Dr. Aretha Teckentrup, thank you for the all your help regarding the RSG and eventually my first year of my PhD. During those years I also had the pleasure to know Dr. Jonathan Carter. Despite Jonathan not interacting with my research directly he always provided advice both on a work and personal level. I also thank Prof. Charlie Elliott who acted as my internal supervisor at Warwick for my final two years. I enjoyed the reading group discussions we had and the recent project we started to work on. I have also enjoyed my trips up to Nottingham to meet Dr. Marco Iglesias, who has welcomed me. For the past two years my final thanks from collaborators goes to Dr. Lassi Roininen. Lassi and I have worked on numerous projects which have taken place both in the UK and Finland. I thank him for introducing me to Dr. Sari Lasanen, Prof. Markku Lehtinen and Dr. Petteri Piiroinen. I have reveled working with him and look forward to our future ongoing projects. My thanks also go to Dr. Björn Stinner who aided with teaching duties, and provided general advice throughout my PhD studies. I am grateful to both my examiners Prof. Chris Budd and Dr. Björn Stinner for their feedback

and comments, and also for their flexibility on the timing of the viva examination. The progression of this thesis would not be possible without industry funding from both E.ON and Premier Oil, for which I thank Dr. Jonathan Carter and Patrick Jennings.

Outside of research many people have made the PhD transition easier, including office mates and friends from before, in particular: Anil, Adam, Ati, Barri, Daniel, Dave, Jakub, Lloyd, Michael, Michael. H, Rohan and Wojciech.

Last but not least I would like to thank my family who have always supported my endeavours. To my cousins Annika and Dave thank you for providing a fun environment during the Christmas break. To my grandmother Kaushal, thank you for always encouraging me to continue the route of education. Above all this thesis would have not been possible without my parents Harjeet and Lalit. I thank them everyday for allowing me to pursue my interests and providing an escape from work when needed to. This thesis is a result of their parenting.

Declarations

This thesis has been submitted to the University of Warwick in support for my application for the degree of Doctor of Philosophy. This is a confirmation that this thesis is the sole work of the author. I hereby declare that this thesis does not include any previous material that has been submitted by a researcher for publication without acknowledgment.

A number of pieces of literature have been produced from this work, which are included in this thesis. A list of submitted papers, and ones that are in preparation, are provided below.

- N. K. Chada. Analysis of hierarchical ensemble Kalman inversion, *submitted to Journal of Inverse and Ill-Posed Problems.*
- N. K. Chada, C. M. Elliott, K. Deckelnick, A. M. Stuart and V. Styles, A Bayesian formulation of the inverse eikonal equation, *in preparation.*
- N. K. Chada, M. A. Iglesias, L. Roininen and A. M. Stuart. Parameterizations for ensemble Kalman inversion, *Inverse Problems,* **34**, 2018.

Neil K. Chada.

Abstract

Since the development of the ensemble Kalman filter, it has seen a wide application to many scientific fields ranging from signal processing to weather forecasting and reservoir simulation. One field which has recently seen a keen interest towards filtering techniques is that of inverse problems. Ensemble-based methods are a popular choice of filtering techniques as they provide a computational advantage over traditional methods whilst retaining a good level of accuracy. This thesis is concerned with developing analysis and numerics of ensemble Kalman inversion (EKI) in the context of Bayesian inverse problems. In particular we are interested in quantifying the uncertainty that can arise for problems where our unknown is defined through geometric features. In the first part of this work we are interested in developing hierarchical approaches for EKI. This motivation is taken from hierarchical computational statistics for Gaussian processes where we are interested in a number of further unknowns such as hyperparameters that define the underlying unknown for the model problem. We present numerics of these hierarchical approaches whilst understanding its long term effect through continuous-time limits. The second part of this work is aimed at improving the computational burden of the forward solver within inverse problems. This improved forward solver is based on the reduced basis method which was designed for parameterized partial differential equations. The final part of the thesis concludes with an application of EKI where we adopt a Bayesian approach of the inverse eikonal equation. Our motivation is to extend the current work to Hamilton-Jacobi equations, where there exists a rich mathematical theory. A key understanding of how to tackle the uncertainty for this equation is addressed.

Chapter 1

Background

1.1 Introduction

In numerous scientific disciplines inverse problems are ubiquitous. Inverse problems [138] are concerned with the recovery of an input to a model, or set of equations, from partial noisy measurements of an output. Some examples of common inverse problems include inferring the geological properties of an oil reservoir from production data, calculating the earths density from its gravity fields and estimating the sound speeds in the earth's subsurface from seismic data. Mathematically speaking inverse problems can be expressed as aiming to estimate $u \in \mathcal{X}$ from noisy measurements of data $y \in \mathcal{Y}$ which are in the form

$$y = \mathcal{G}(u) + \eta, \tag{1.1.1}$$

where

- $\mathcal{X}, \mathcal{Y}$ - function spaces.
- $u \in \mathcal{X}$ - input.
- $y \in \mathcal{Y}$ - output.
- $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$ - forward operator.
- η - additive noise.

Usually we assume that $(\mathcal{X}, \|\cdot\|_{\mathcal{X}}), (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ are two Banach spaces and our output, or the data, y is representative of noisy solutions to our forward model where we have some noise η within our system. A common assumption with (1.1.1) is that the additive noise is Gaussian, though other choices for the noise observational

noise are possible; however in this thesis we will focus on additive Gaussian noise for simplicity. An issue that arises when aiming to solve the inverse problem (1.1.1) is that it is ill-posed in the Hadamard sense i.e. there is no guarantee of existence, stability and uniqueness of solutions. Due to this drawback during the 20^{th} Century mathematicians aimed at producing algorithms that approximated a solution to (1.1.1). An idea that came about was that one way to approximate the solution to inverse problems was by relating it it to the least squares functional

$$\Phi_{\text{LS}}(u) = \frac{1}{2}\big|y - \mathcal{G}(u)\big|_{\mathcal{Y}}^2 \,, \tag{1.1.2}$$

with subscript LS denoting least squares. From (1.1.2) our solution u^* to the inverse problem is then given by a minimisation procedure such that

$$u^\star = \underset{u\in\mathcal{X}}{\operatorname{argmin}}\, \Phi_{\text{LS}}(u),$$

where our minimization occurs in the solution space $\mathcal{X}$. However, despite this formulation with the LS functional there still lies issues with approximating a solution. We can approximate solutions, but under certain conditions uniqueness is not guaranteed as well as the stability of solutions. One way that was proposed to overcome this issue was to modify the least square functional such that it incorporated prior information of the unknown. By doing so this leads to a slight modification of (1.1.2), where now our functional is defined as

$$\Phi_{\text{TP}}(u) = \frac{1}{2}\big|y - \mathcal{G}(u)\big|_{\mathcal{Y}}^2 + \mathcal{R}(u). \tag{1.1.3}$$

Our new inclusion in our least-square functional is a regularisation term $\mathcal{R}(u)$. One example of this is the form of $\mathcal{R}(u) = \frac{\lambda}{2}|u|_{\mathcal{Z}}^2$; this is referred to as Tikhonov-Phillips regularization [55, 107]. From our regularisation term $\lambda > 0$ denotes our regularisation constant and $\mathcal{Z}$ is an embedded subspace of $\mathcal{X}$. The motivation for adding regularisation terms is threefold: firstly to aid the inversion by reducing the amount of influence the data has on the solution, i.e. to prevent overfitting of data. Secondly the add some further prior information about our unknown. Finally adding regularization ensures that the inverse problem is continuous with respect to the data. Choosing the form of regularization can be dependent on the actual problem of interest for inversion and model [13, 54]. Now we can restate our inverse solution which is the following minimization procedure

$$u^\star = \underset{u\in\mathcal{X}}{\operatorname{argmin}}\big(\frac{1}{2}\big|y - \mathcal{G}(u)\big|_{\mathcal{Y}}^2 + \frac{\lambda}{2}|u|_{\mathcal{Z}}^2\big). \tag{1.1.4}$$

As there is a relationship between the TP functional and LS functional, optimisation methods have been traditionally been used and applied for solving inverse problems in the form of (1.1.4). Some common examples of iterative optimisation algorithms that have been used include: the Gauss-Newton method [12, 131], the Landweber method [65, 125] and the conjugate gradient method [85]. Like many optimizations methods, solving inverse problems requires information about derivatives. In the inverse problem setting this would include the Jacobian $D\mathcal{G}(u)$ and the Hessian $D^2\mathcal{G}(u)$, which are based on the forward operator. There has been extensive research into deterministic inverse problems which have included various forms of regularization. Although this aids with approximating stable solutions, still does not resolve the lack of uniqueness. Resolving this issue largely depends on the inverse problem, the setting and the form of regularization which is added. One way to alleviate this issue is to look at an alternative approach that was adopted to solve inverse problems. This approach negates the idea of characterizing a functional and aiming to minimize it for u. Instead the unknown u is set as a probabilistic distribution, where all the quantities from (1.1.1) are treated as a random variable. This approach is known as a statistical or Bayesian approach to inverse problems.

1.2 Bayesian Approach

The methodology that was discussed previously in subsection 1.1 is related to the "deterministic approach" to inverse problems. Since then there has been much development in this field which has sparked an alternative viewpoint of inverse problems. Instead of our unknown u we are now interested in characterizing a distribution of the random variable $u|y$ which is the unknown conditioned on the noisy data y. This is statistical view point of inverse problems which is commonly referred to as the "Bayesian approach" [87, 137]. Taking our inverse problem (1.1.1) the Bayesian approach allows us to treat each quantity as a random variable with a Lebesgue density. In order to characterise our new unknown we apply Bayes' Theorem which, in its general finite-dimensional form, is

$$\mathbb{P}(u|y) = \frac{\mathbb{P}(u,y)}{\mathbb{P}(y)} \tag{1.2.1}$$

$$\mathbb{P}(u|y) = \frac{\mathbb{P}(y|u) \times \mathbb{P}(u)}{\mathbb{P}(y)} \tag{1.2.2}$$

$$\propto \mathbb{P}(y|u) \times \mathbb{P}(u). \tag{1.2.3}$$

In (1.2.3) the term $\mathbb{P}(u)$ denotes our *prior* distribution, encoding initial belief about our unknown u, prior to seeing any data; on the other hand the *data-likelihood* term $\mathbb{P}(y|u)$ captures the relationship between the forward model $G(u)$ and the data y. By using Bayes' Theorem, as done in (1.2.2), we can combine both the *prior* and *data-likelihood* which results in the conditional distribution of u given y, $\mathbb{P}(u|y)$, known as the *posterior*. The form (1.2.3) is presented because the constant of proportionality $\mathbb{P}(y)$ is often hard to compute, and is not needed for many computational methods which explore the posterior. We note that in the case of our posterior form (1.2.3) we can approximate the posterior up to some level of proportionality. In many scenarios in Bayesian statistics the normalising constant $\frac{1}{\mathbb{P}(y)}$ is not known, or is to difficult to obtain. However, despite this simple form of Bayes' Theorem, the Bayesian approach has been further applied to inverse problems in a differential equation setting, specifically partial differential equations (PDEs). Due to this one has to reformulate Bayes' Theorem for PDEs, in a function space setting. Instead of a distribution we now wish to characterise our unknown as a posterior measure μ^y. This change corresponds to going from a finite-dimensional to a ∞-dimensional problem. An obvious question to ask is what advantages does the Bayesian approach have over the deterministic case? Unlike its counterpart, under certain conditions, well-posedness of the Bayesian inverse problem can be attained. This is one of the key significant advances with the Bayesian approach, which is possible due to its ∞-dimensional analysis. Secondly, arguably its main contribution is handling uncertainty within the problem. This can be tackled more effectively due to various prior forms, where a range of these are discussed in [137]. However much of the initial and existing theory on this assumed that the reference measure was of a Gaussian form. We now recall the definition of a Gaussian measure [20] and discuss some important definitions and assumptions which are required in order to present a well-posedness theorem of (1.1.1).

Definition 1.2.1. *A Borel measure μ on $\mathbb{R}$ is called a non-degenerate Gaussian measure if there exists a $m \in \mathbb{R}$ and $\sigma^2 > 0$ such that*

$$\frac{d\mu}{d\lambda}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\Big(-\frac{1}{2\sigma^2}(x-m)^2 \Big).$$

A Gaussian measure can be characterized through two quantities its mean $m \in \mathbb{R}$ and its variance (covariance) $\sigma^2 > 0$. We now extend this formal definition to Gaussian measures on Banach spaces.

Definition 1.2.2. *Let $\mathcal{X}$ be a separable Banach space and μ be a Borel measure on $\mathcal{X}$. Then μ is said to be a Gaussian measure if $\ell^{\#}\mu$ is a Gaussian measure on $\mathbb{R}$ for*

all $\ell \in \mathcal{X}^*$.

From the above definition $\ell^{\#}\mu$ is a push forward measure of the measure μ. We say that μ is centered if $\ell^{\#}\mu$ has mean 0 for all $\ell \in \mathcal{X}^*$. Our final definition on Gaussian measures is with regards to defining the covariance of a Gaussian measure on both a Banach and Hilbert space.

Definition 1.2.3. *Let μ be a centered Gaussian measure on a separable Banach space $\mathcal{X}$. Then the covariance operator $\mathcal{C} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is defined as*

$$\mathcal{C}(\ell, \ell') = \int_{\mathcal{X}} \ell(u)\ell(u')\mu(du), \tag{1.2.4}$$

is called the covariance operator of μ. If $\mathcal{X}$ is a Hilbert space, then after identification with its dual space, our covariance operator is now

$$\mathcal{C} = \int_{\mathcal{X}} (u \otimes u)\mu(du).$$

In the case of a centered Gaussian measure $\mu \sim \mathcal{N}(0, \mathcal{C})$, the measure is defined entirely through its covariance operator. In the context of Bayesian inversion in the ∞-dimensional case an important question is how to relate a measure on the prior μ_0 to the posterior measure μ^y. With Gaussian measures this relationship between both measures could be represented through a Radon-Nikodym derivative

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z}\exp(-\Phi(u; y)). \tag{1.2.5}$$

The Radon-Nikodym derivative (1.2.5) describes the change of measure from the prior to the posterior, and as seen through Bayes' Theorem (1.2.3), this is achieved through the negative log-likelihood, which in this case is $\Phi(u; y)$ up to a constant Z. To express this change of the measures, we require that μ^y to be absolutely continuous with respect to the measure μ_0. In order to provide a well-posedness theorem for the posterior measure μ^y we need a suitable metric to provide a stability result. Common metrics which are used are the total variation, the Hellinger and Kullback-Liebler metric [91].

Definition 1.2.4. *(Total variation distance) Given two measures μ and μ', and their corresponding densities $\rho(u)$ and $\rho(u')$, the total variation distance between these measures is given as*

$$d_{TV}(\mu, \mu') = \frac{1}{2}\int_{\mathbb{R}^n} \big|\rho(u) - \rho'(u)\big| du.$$

Definition 1.2.5. *(Hellinger distance) Given two measures μ and μ', and their corresponding densities $\rho(u)$ and $\rho(u')$, the Hellinger distance between these measures is given as*

$$d_{Hell}(\mu,\mu') = \Big(\frac{1}{2}\int_{\mathbb{R}^n}\Big(\sqrt{\rho(u)}-\sqrt{\rho'(u)}\Big)^2 du\Big)^{1/2}.$$

Definition 1.2.6. *(Kullback-Liebler divergence) Given two measures μ and μ', and their corresponding densities $\rho(u)$ and $\rho(u')$, the Kullback-Liebler divergence between these measures is given as*

$$d_{KL}(\mu'||\mu) = \int_{\mathbb{R}^n}\log\left(\frac{\rho'(u)}{\rho(u)}\right)\rho'(u)du.$$

An important way to compare some of these metrics is through the following relationship between the total variation and Hellinger metrics

$$0 \leq \frac{1}{\sqrt{2}}d_{TV}(\mu,\mu) \leq d_{Hell}(\mu,\mu') \leq d_{TV}(\mu,\mu')^{1/2} \leq 1.$$

A natural question one can ask regarding the above metrics is between the total variation and Hellinger metric, which one is preferred? The following lemma is crucial in understanding why the Hellinger distance is more favourable over total variation.

Lemma 1.2.1. *Let $f:\mathbb{R}^l \to \mathbb{R}^p$ be such that*

$$\big(\mathbb{E}^{\mu}|f(u)|^2 + \mathbb{E}^{\mu'}|f(u)|^2\big) < \infty.$$

Then

$$\big|\mathbb{E}^{\mu}f(u) - \mathbb{E}^{\mu'}f(u)\big| \leq 2\big(\mathbb{E}^{\mu}|f(u)|^2 + \mathbb{E}^{\mu'}|f(u)|^2\big)^{\frac{1}{2}}d_{Hell}(\mu,\mu'),$$

where as a consequence

$$\big|\mathbb{E}^{\mu}f(u) - \mathbb{E}^{\mu'}f(u)\big| \leq 2\big(\mathbb{E}^{\mu}|f(u)|^2 + \mathbb{E}^{\mu'}|f(u)|^2\big)^{\frac{1}{2}}d_{TV}(\mu,\mu')^{1/2}.$$

Lemma 1.2.1 is significant as it shows that two measures μ and μ' are $\mathcal{O}(\epsilon)$-close within the Hellinger metric. That is if the function $f(u)$ is square integrable with respect to μ and μ', then its expectations are $\mathcal{O}(\epsilon)$- close within the Hellinger metric. However this is not the case with the total variation metric as the expectations of $f(u)$ with respect to μ and μ' are $\mathcal{O}(\epsilon^{-1/2})$. In order to attain $\mathcal{O}(\epsilon)$-closeness we need stronger assumptions, which is achieved through the following lemma.

Lemma 1.2.2. *Assume that $|f|$ is finite almost surely with respect to both μ and*

μ', *where we denote the almost sure upper bound on* $|f|$ *by* $f_{\max}$, *then*

$$\left|\mathbb{E}^{\mu} f(u) - \mathbb{E}^{\mu'} f(u)\right| \leq 2 f_{\max} d_{TV}(\mu, \mu').$$

The previous two lemmas are proved in [91], and provide a basis on why we should choose the Hellinger metric over the total variation. This is due to the induced perturbations of the measure on expectations of square-integrable functions. It is for this reason why we consider the Hellinger metric as a more "appropriate" metric to use than total variation for the existence of solution to a Bayesian inverse problem. We omit using the Kullback-Liebler (KL) divergence as using the definition of a metric, it does not abide by the symmetry property of the metric. As we have discussed both necessary concepts we can present the following theorem which provides uniqueness and existence of solutions to (1.1.1), which is based on the Hellinger metric.

Theorem 1.2.1. *(Well-posedness) Assume that* μ_0 *is defined as* $\mathcal{N}(0, \mathcal{C})$, y *by* (1.1.1) *and* Φ *by* $\frac{1}{2}|y - \mathcal{G}(u)|_{\Gamma}^2$. *If* μ^y *is the regular conditional probability measure on* $u|y$, *then* $\mu^y \ll \mu_0$ *with Radon-Nikodym derivative*

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z} \exp(-\Phi(u; y)),$$

where

$$Z := \int_{\mathcal{X}} \exp(-\Phi(u; y)) \mu_0(du).$$

Furthermore μ^y *is locally Lipschitz with respect to* y *in the Hellinger distance: for all* y, y' *with* $max\{|y|_{\Gamma}, |y'|_{\Gamma}\} \leq r$, *there exists a* $C = c(r) > 0$ *such that*

$$d_{Hell}(\mu^y, \mu^{y'}) \leq C|y - y'|_{\Gamma}.$$

Remark 1.2.1. *We note that with the above theorem, we have assumed that our prior is of a Gaussian form. This theorem can be extended to include numerous other priors such as uniform and Besov, whilst the main statements of the theorem still hold. Further assumptions are required for well-posedness which can be found in [137].*

With the Bayesian setting of inverse problems extended to the ∞-dimensional case, the next natural question is to study numerical methods to explore the posterior. This can be achieved by taking, or extending, existing methods from computational statistics. We review a number of class of methods designed and optimized for Bayesian inversion, where in particular we present sampling methods and meth-

ods which arise in a deterministic optimization setting. We will discuss both the numerical method used in its original context and how the methodology can be used for solving Bayesian inverse problems. These methods will primarily be split into data assimilation methods and Markov chain Monte Carlo methods, whilst we briefly review other commonly used methods.

1.3 Markov Chain Monte Carlo Methods

In many situations arising in computational statistics one is interested in simulating from a probabilistic distribution π, or in other words calculating an expectation of a π- integrable function $f : \mathcal{X} \to \mathbb{R}$

$$\mathbb{E}_\pi(f) := \int_{\mathcal{X}} f(x)\pi(dx). \tag{1.3.1}$$

However simulating from a distribution π is not always possible due to a number of reasons that make it difficult such as:

- High dimensionality of the problem.
- Difficult to attain the normalising constant.

Traditional methods such as analytical integration and quadrature schemes do not aid significantly when aiming to solve (1.3.1). One way to overcome this issue is to construct a Markov chain $\{X_i\}_{i=1}^n$ on a measurable space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, with Borel σ-algebra that converges to π. This in a nutshell is the aim of Markov chain Monte Carlo (MCMC) methods [121], where we are interested in constructing an invariant ergodic Markov chain with respect to π. MCMC methods are used to calculate (1.3.1) based the notion of ergodic averages

$$S_n(f) := \frac{1}{n}\sum_{i=1}^{n} f(X_i), \quad \text{for } n \in \mathbb{R}^+.$$

We know from the strong law of large numbers that the ergodic average, taking a large enough n, will converge to π i.e.

$$S_n(f) \longrightarrow \mathbb{E}_\pi(f), \quad \text{when } n \to \infty.$$

From the central limit theorem (CLT) we know that the convergence is of order $\mathcal{O}(n^{-\frac{1}{2}})$, which is also the order of MCMC methods. By working in an MCMC framework, we can consider our distribution of interest π as a posterior of the form

(1.3.1), and as a result can obviate the need to calculate the normalizing constant. One of the most common type of MCMC methods that are used in practice are Metropolis-Hastings (MH) algorithms. These algorithms are based on an accept-reject scheme for proposed moves through a transition kernel $q(x, y)$, which encodes the condition that the probability of going from state x to y is the same as from y to x. It implies that π is an invariant distribution for Markov chain with kernel q. The accept-reject is designed in such a way to allow for detailed balance

$$\pi(y)q(y, x) = \pi(x)q(x, y). \tag{1.3.2}$$

This detailed balance (1.3.2) is a consequence of reversibility of the Markov chain, and it reflects the property that the probability of going from state x to y is the same as from y to x. The MH algorithm is described in more detail in Algorithm 1. We recall the properties of transition kernels through the following definition.

Definition 1.3.1. *A function $q : \mathcal{X} \times \mathcal{B}(\mathcal{X}) \to [0, 1]$ is defined as a Markov kernel if the following conditions are satisfied*

- *For each $x \in \mathcal{X}$, $q(x, \cdot)$ is a probability measure on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$.*
- *$x \mapsto q(x, A)$ is $\mathcal{B}(\mathcal{X})$-measurable for all $A \in \mathcal{B}(\mathcal{X})$.*

In the context of MH from (1.3.2) we can think of $\pi(y)$ as the target density we are interested in simulating from and $q(x, y)$ as the transition kernel, or proposal, from the current state x to y. The accept-reject scheme for MH is based on an acceptance probability which is given in the form

$$\alpha(x, y) := \min\left\{1, \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)}\right\}, \quad \text{for } x \neq y,$$

For the MH algorithm the overall kernel for the scheme can be written as,

$$p(x, y) = \alpha(x, y)q(x, y) + (1 - \alpha(x, y))p(x, y). \tag{1.3.3}$$

which includes the usual transition kernel $q(x, y)$. The intuition behind the (1.3.3) is that we have a number of choices based on proposing and accepting y. If we move from x to y and the proposed move is accepted then we have the result of $\alpha(x, y)q(x, y)$. However we also have the option of moving from state x to x. This can occur in two ways either by having the accepted proposal $\alpha(x, x)q(x, x)$ or that our proposed move from x to y is rejected.

Algorithm 1 Metropolis Hastings Algorithm

1. If current position is x propose move to y according to $q(x, y)$.
2. Compute acceptance probability

$$\alpha(x, y) = \min\left\{\frac{\pi(y)q(y,x)}{\pi(x)q(x,y)}, 1\right\}, \text{ for } x \neq y.$$

3. With acceptance probability $\alpha(x, y)$ move from from x to y otherwise stay still.
4. Go back to **1.**

For MH we can choose a number of proposals which work best under different scenarios. Three common proposal types are given by:

- Random walk Metropolis Hastings (RWMH).
- Metropolis-adjusted Langevin algorithm (MALA).
- Preconditioned Crank-Nicholson (pCN).

For RWMH this employs a Gaussian proposal that is centered at the current state. Given we are at x_n the proposal for the next step is defined as

$$y_{n+1} = x_n + \sigma\xi_{n+1}, \tag{1.3.4}$$

such that $\xi_{n+1} \sim \mathcal{N}(0, \sigma^2)$. Our initial state $x_0 \sim \mathcal{N}(0, 1)$ is also normally distributed, which implies our transition kernel is now $q(x, y) \sim \mathcal{N}(x, \sigma^2)$. Due to symmetry of the Normal distribution we can set the proposal as $q(x, y) = q(y, x)$ which leads to an acceptance probability of

$$\alpha(x, y) := \min\left\{1, \frac{\pi(y)}{\pi(x)}\right\}. \tag{1.3.5}$$

An important note to make about RWMH is that the noise ξ_{n+1} is independent of x_n, which means that the noise has no information about the current position. This is one limitation regarding this proposal. An obvious way to alleviate this problem is to consider more informative proposals. This leads to an alternative method that combines more information, while remaining in the MH accept-reject mechanism. One example of this is the MALA proposal [63, 122]. Its derivation comes from the SDE

$$dX_t = -\nabla V(X_t)dt + (2\beta^{-1})^{1/2}dW_t, \tag{1.3.6}$$

which is a common stochastic differential equation (SDE) that arises in various scientific disciplines, known as the Langevin equation. From (1.3.6) $\{W_t\}_{t\geq 0}$ is a d-dimensional standard Brownian motion, $V(x)$ is a potential with $X_t \in \mathbb{R}^d$ and $\beta > 0$ is a constant. In order to encapsulate Langevin dynamics in our MCMC we need to efficiently discretize (1.3.6). Arguably the simplest way to do so is by taking a forward Euler discretization of the SDE (1.3.6)

$$X_{n+1} \sim \mathcal{N}(X_n - h\nabla V(X_n), 2hI_d),$$

with our constant chosen as $\beta = 1$. Therefore the proposal of MALA are given as

$$Y_{n+1} := X_n - h\nabla V(X_n) + (2h)^{1/2}\xi_{n+1}, \quad \xi_{n+1} \sim \mathcal{N}(0,1). \tag{1.3.7}$$

Based on the MALA proposal given in (1.3.7) our transitional kernel is now $q(x,y) \sim \mathcal{N}(x - h\nabla V(x), 2hI_d)$. Our final proposal we discuss is a simpler variant of the RMWH which includes a slightly different proposal kernel to that of (1.3.4). The proposal itself is a redefined version of RWMH given as

$$y_{n+1} = m + (1-\beta_*^2)^{1/2}(x_n - m) + \beta_*\xi_{n+1}, \tag{1.3.8}$$

which we call the pCN method. Here again we assume that $\xi_{n+1} \sim \mathcal{N}(0, \mathcal{C})$ is Gaussian noise, but with the addition of $\beta_* \in (0,1]$ which is some tuneable parameter and m is the mean of our initial state. Equation (1.3.8) was first derived in [18] and tested numerically in [39]. The pCN proposal is specific to when our target density is a Gaussian. The intuition behind the parameter β_* is that in this context it is defined as the proposal variance and can be tuned to ensure a more reasonable acceptance probability. Combined with the initial condition x_0 it can be viewed as an improvement over the independence sampler, that is defined through the acceptance probability (1.3.5).

1.3.1 MCMC Within Inverse Problems

Statistical inverse problems in the finite-dimensional case are reviewed in the text by Kaipio and Somersalo [87], where they discuss various MCMC methods used for different problems. The text itself is oriented towards applications, including interpolation and PDE inversion. Beyond this and work of others [69], developing MCMC methods on function spaces was crucial in understanding how we can use MCMC inversion for ∞-dimensional problems. We recall that in the MCMC for Bayesian inversion we are interested in sampling from a measure μ^y as in Theorem

1.2.1.

Notably the work of Roberts and Stuart [17, 104, 118] developed much of the existing literature of various MCMC algorithms in high dimensions, and this itself posed a motivation for merging this work to the Bayesian approach. Much of the existing work has been based on the RWMH, whereas before, the key aspects of this in Bayesian inversion is through the proposal and the likelihood. Recalling that pCN proposal is defined, with slightly different notation

$$\hat{u}^{(n)} = (1-\beta_*^2)^{1/2}u^{(n)} + \beta_*\zeta^{(n)}, \quad \zeta^{(n)} \sim \mathcal{N}(0, C). \tag{1.3.9}$$

We omit defining our proposal as y as this denotes the data. As we assume our prior measure μ_0 is a centered Gaussian measure, we set $m = 0$ .We can interpret the proposal (1.3.9) as how we define our prior which is given through $\zeta^{(n)}$ with covariance C. Now for the data-likelihood, this is incorporated in the MH framework through the acceptance probability which is given by

$$\begin{aligned}\alpha(u^{(n)}, \hat{u}^{(n)}) &:= \min\left\{1, \frac{\pi(\hat{u})}{\pi(u)}\right\} \\ &= \min\left\{1, \exp\left(\Phi(u^{(n)}; y) - \Phi(\hat{u}^{(n)}; y)\right)\right\}.\end{aligned} \tag{1.3.10}$$

Combining both the proposal and the acceptance probability a simple RMWH algorithm for Bayesian inversion is presented below in Algorithm 2.

Algorithm 2 pCN Algorithm

1. Set initial state $u^{(n)} \in \mathcal{X}$ with proposal (1.3.9) with $\beta_* \in (0, 1]$.
2. Compute acceptance probability

$$\alpha(u^{(n)}, \hat{u}^{(n)}) = \min\left\{1, \exp\left(\Phi(u^{(n)}; y) - \Phi(\hat{u}^{(n)}; y)\right)\right\}.$$

3. With acceptance probability $\alpha(u^{(n)}, \hat{u}^{(n)})$ move from u to $\hat{u}$ otherwise stay still.
4. Go back to **1.**

Aside from RWMH, other MCMC algorithms have been implemented which are that of a geometric type which include Hamiltonian Monte Carlo [15], as well as MALA. We note that with these latter methods that they are not-derivative free. Much of the ongoing work on MCMC for inverse problems is not derivative-free which in numerous cases can produce better results, but at the same time adds constraints with the implementation and the cost [22, 117]. Concentrating on derivative-free methods, there have been a number of extensions such as the use of geometric priors

[80] and level set techniques [81].
In this context we can state that this particular inverse solver is a sampling method as it gains inference from a probability distribution in the form of a posterior through (1.3.3). As we will discuss later, there are other types of methods that can be used in a Bayesian setting such as methods deriving from optimization rather than sampling.

1.4 Data Assimilation Techniques

Data assimilation [57, 90, 91] is the study of state estimation of a dynamical system with the incorporation of noisy measurements. The ideas behind data assimilation can be related back to control theory and optimal control [88]. Since then there has been a growing number of applications where data assimilation techniques can be used such as weather forecasting, geosciences and meteorology [58, 59]. The two governing equations of data assimilation, in discrete time, are given as

$$u_{j+1} = \Psi(u_j) + \xi_j, \qquad \{\xi_j\}_{j\in\mathbb{Z}^+} \sim \mathcal{N}(0,\Sigma), \tag{1.4.1}$$

$$y_{j+1} = H(u_{j+1}) + \eta_{j+1}, \quad \{\eta_{j+1}\}_{j\in\mathbb{Z}^+} \sim \mathcal{N}(0,\Gamma). \tag{1.4.2}$$

Here $\{u_j\}_{j\in\mathbb{Z}^+}$ is our signal which is updated through a forward operator $\Psi : \mathbb{R}^m \to \mathbb{R}^m$, which when combined with noise, provides the update u_{j+1}. Our data is denoted as y_{j+1} which is produced by sending our updated signal through the operator $H : \mathbb{R}^m \to \mathbb{R}^n$ which is known as observational operator. Our initial conditions for the system are given as $u_0 \sim \mathcal{N}(m_0, \mathcal{C}_0)$. We also notice that there is the addition of additive Gaussian noise. In data assimilation the common goal is to use the data (1.4.2) to inform the signal dynamics governed by (1.4.1). We can think of this probabilistically where we are interested in characterising the distribution of $\mathbb{P}(u_j|y_j)$. Algorithms that are used to quantify this distribution are classified as either; (i) *smoothing*, (ii) *filtering*.
Smoothing is concerned with determining the smoothing distribution $\mathbb{P}(u|y) \propto \mathbb{P}(y|u)\mathbb{P}(u)$ where the signal is conditioned on all of the data specified in the appropriate time interval. Filtering is concerned with the determining the filtering distribution $\mathbb{P}(u_j|Y_j)$ where $Y_j = \{y_l\}_{l=1}^j$ be the accumulated data up to time j. Filtering is commonly split into two steps:

- **Prediction step:** $\mathbb{P}(u_j|Y_j) \mapsto \mathbb{P}(u_{j+1}|Y_j)$ - maps the signal into the data space.
- **Analysis/Update step:** $\mathbb{P}(u_{j+1}|Y_j) \mapsto \mathbb{P}(u_{j+1}|Y_{j+1})$ - updates the signal by

comparing it with the data.

By applying Bayes Theorem' in the analysis step we can deduce that

$$\begin{aligned}\mathbb{P}(u_{j+1}|Y_{j+1}) &= \mathbb{P}(u_{j+1}|Y_j, y_{j+1}) \\ &= \frac{\mathbb{P}(y_{j+1}|u_{j+1})\mathbb{P}(u_{j+1}|Y_j)}{\mathbb{P}(y_{j+1}|Y_j)}\end{aligned} \tag{1.4.3}$$

Unlike smoothing, filtering has the advantage of determining information from the signal at the current state. Due to this filtering algorithms are a more popular choice among practitioners. Some of the most common filtering algorithms that are used are: (i) the Kalman filter, (ii) the extended Kalman filter (ExKF), (iii) 3DVAR and (iv) the ensemble Kalman filter (EnKF). A common attribute among filtering algorithms is to update the signal based on observations through statistical quantities. As the original signals that were considered were Gaussians i.e. $u_j \sim \mathcal{N}(m_j, C_j)$, we aim to update our filtering distribution using notions of the mean and the covariance. Below we present the 3DVAR algorithm and the extended and ensemble Kalman filter.

1.4.1 3DVAR

The 3DVAR algorithm is derived from the linear Kalman filter [88], which has the assumption that the underlying signal is of a Gaussian form. It differs in that the covariance matrix is fixed such that $\hat{C}_{j+1} \equiv \hat{C}$. This leads to the equations

$$\begin{aligned}\hat{m}_{j+1} &= \Psi(m_j), \\ m_{j+1} &= (I - KH)\hat{m}_{j+1} + Ky_{j+1}, \\ K &= \hat{C}H^T S^{-1}, \\ S &= H\hat{C}H^T + \Gamma.\end{aligned}$$

3DVAR works in a variational manner and is sequentially updated at each j through the minimization procedure

$$\hat{m}_{j+1} = \operatorname*{argmin}_u I_n(u),$$

with cost function $I_n(u)$ given as

$$I_n(u) := \frac{1}{2}|y_{j+1}^{(n)} - Hu|_\Gamma^2 + \frac{1}{2}|u - \hat{m}_{j+1}^{(n)}|_{\hat{C}}^2.$$

This optimization procedure despite having a different mathematical formulation, leads to the update formulae for the ExKF. As suggestive from its name in practical application the minimisation procedure is taken over all spatial dimensions. The extension from this is 4DVAR which takes the time dimension into consideration. However by doing so our problem is now given as a smoothing algorithm. We now look to go beyond this with two methods that sequentially update our quantity of interest through both an updated mean and covariance; the extended Kalman filter and ensemble Kalman filter.

1.4.2 Extended Kalman Filter

The ExKF was developed to work with non-linear Gaussian proposals of the linear Kalman filter which propagates both the mean and the covariances. As with all filtering methods the ExKF has a prediction and analysis step where the prediction step is to define both $\hat{m}_{j+1}$ and $\hat{C}_{j+1}$.

$$\begin{aligned}\hat{m}_{j+1} &= \Psi(m_j),\\ \hat{C}_{j+1} &= D\Psi(m_j)C_j D\Psi(m_j)^T + \Sigma.\end{aligned}$$

$$\begin{aligned}S_{j+1} &= H\hat{C}_{j+1}H^T + \Gamma,\\ K_{j+1} &= \hat{C}_{j+1}H^T S_{j+1}^{-1},\\ m_{j+1} &= (I - K_{j+1}H)\hat{m}_{j+1} + K_{j+1}y_{j+1},\\ C_{j+1} &= (I - K_{j+1}H)\hat{C}_{j+1}.\end{aligned}$$

In filtering K_{j+1} is commonly referred to as the Kalman gain matrix which describes how much information we have gained through the covariance structure. From the predicted mean and covariance we have an update of our distribution in terms of m_{j+1} and C_{j+1}.

1.4.3 Ensemble Kalman Filter

Out of the filtering algorithms that were mentioned in Section 1.4, the EnKF can be thought of as a Monte Carlo approximation of the Kalman filter which has certain advantages over its counterparts. Arguably its main advantage it acquires is the reduction in cost of the algorithm for high dimensional problems. As we can see with the ExKF we constantly have to update the covariances in the analysis step through previous knowledge of information. This procedure is expensive.

With the EnKF this takes away this issue by using an ensemble of particles

to represent the predicted mean and covariances, which in turn are used for the updated moments. As before, we can split the EnKF into a prediction and analysis step:

$$\begin{aligned}
\hat{u}_{j+1}^{(n)} &= \Psi(u_j^{(n)}) + \xi_j^{(n)}, \\
\hat{m}_{j+1} &= \frac{1}{N}\sum_{n=1}^{N} u_{j+1}^{(n)}, \\
\hat{C}_{j+1} &= \frac{1}{N-1}\sum_{n=1}^{N}(u_{j+1}^{(n)} - \hat{m}_{j+1})(u_{j+1}^{(n)} - \hat{m}_{j+1})^T.
\end{aligned}$$

$$\begin{aligned}
K_{j+1} &= \hat{C}_{j+1}H^T(H\hat{C}_{j+1}H^T + \Gamma), \\
u_{j+1}^{(n)} &= (I - K_{j+1}H)\hat{u}_{j+1}^{(n)} + K_{j+1}y_{j+1}^{(n)}, \\
y_{j+1}^{(n)} &= y_{j+1} + \eta_{j+1}^{(n)}.
\end{aligned}$$

As before K_{j+1} represents the Kalman gain matrix and $\xi_j^{(n)}$ and $\eta_{j+1}^{(n)}$ are i.i.d. Gaussian noise. In the EnKF context our prediction step defines a sample mean and covariance from our signal. From this in the analysis step we define our Kalman gain through our sample covariance, which updates our signal, which is given by $u_{j+1}^{(n)}$. This is aided by aiming to minimize the discrepancy of the data $y_{j+1}^{(n)}$ and the quantity $H(u)$. To better understand this discrepancy, there is an alternative approach of looking at the EnKF is through a variational approach, where we consider the follow cost function

$$I_n(u) := \frac{1}{2}|y_{j+1}^{(n)} - H(u)|_\Gamma^2 + \frac{1}{2}|u - \hat{u}_{j+1}^{(n)}|_{\hat{C}_{j+1}}^2, \tag{1.4.4}$$

for which we aim to minimise, which is defined as the updated mean

$$\hat{m}_{j+1} = \operatorname*{argmin}_u I_n(u). \tag{1.4.5}$$

This minimization procedure relies on the updated covariance $\hat{C}_{j+1}$ which is dependent entirely on $\hat{v}^{(n)}$. As described in the prediction step and update step of filtering, a mapping is presented between distributions. As we related the distributions in the filtering setting, for each step, we can do so similarly for the EnKF, i.e.

$$\{u_j^{(n)}\}_{n=1}^N \mapsto \{u_{j+1}^{(n)}\}_{n=1}^N, \qquad \{u_{j+1}^{(n)}\}_{n=1}^N \mapsto \{\hat{u}_{j+1}^{(n)}\}_{n=1}^N.$$

1.4.4 DA Techniques within Inverse Problems

Data assimilation and Bayesian inverse problems have an important connection, arguably due to the fact that both fields have a common aim of constructing a distribution of a quantity of interest conditioned on some data. In particular one data assimilation method that has seen a significant translation to inverse problems is the EnKF. We refer this type of inversion as ensemble Kalman inversion (EKI). EKI was first considered by Reynolds and coauthors [93, 110] which was motivated by applications in subsurface flow. Since then a new method was proposed of how to effectively apply EnKF techniques to solve PDE-constrained Bayesian inverse problems. This was done through the work of Iglesias et al. [78, 79], known as the iterative EnKF method. This method proposed is very similar to that of the traditional EnKF, where in this context we are interested in updating an ensemble of particles $\{u_n^{(j)}\}_{j=1}^J$ where n is the iteration count and J is the ensemble member. This is achieved through defining the mean and the covariances

$$\bar{\mathcal{G}}_n = \frac{1}{J}\sum_{j=1}^{J}\mathcal{G}(u_n^{(j)}), \tag{1.4.6}$$

$$\bar{u}_n = \frac{1}{J}\sum_{j=1}^{J}(u_n^{(j)}), \tag{1.4.7}$$

$$C_n^{uw} = \frac{1}{J-1}\sum_{j=1}^{J}((u_n^{(j)}) - \bar{u}) \times (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}), \tag{1.4.8}$$

$$C_n^{ww} = \frac{1}{J-1}\sum_{j=1}^{J}(\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}) \times (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}), \tag{1.4.9}$$

where as discussed before $\mathcal{G}$ is the forward operator of the PDE. Then finally the update of the ensemble is achieved through the following update formula

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \Gamma)^{-1}(y_n^{(j)} - \mathcal{G}(u_n^{(j)})). \tag{1.4.10}$$

From the update formula (1.4.10) $(y^{(j)} - \mathcal{G}(u_n^{(j)}))$ can be viewed as the discrepancy between the data $y_n^{(j)}$ and the solution evaluated at the forward operator. Relating this further to the EnKF discussed in subsection 1.4.3, our Kalman gain matrix in this context is simply $C_n^{uw}(C_n^{ww} + \Gamma)^{-1}$. The iterative EnKF has two important characteristics: firstly that the method is completely derivative free: it requires no derivatives of the forward operator and its adjoint to implement. This poses computational advantages and as a result can be treated as a black-box solver. The

other is that the method can be viewed as a *semi*-Bayesian method in that it is based on an optimization framework, but can incorporate numerous Bayesian techniques, motivated from the likes of MCMC.

To derive this method we can do so in a number of ways, one is to follow a similar approach as the Kalman filter, was is shown in [88]. The other approach is to consider the functional of interest we aim to minimize as

$$I(u) = \|y - \mathcal{G}u\|_{\Gamma}^2 + \|u - \bar{u}\|_{\mathcal{C}}^2, \tag{1.4.11}$$

where u_{TP} is a solution of the functional (1.4.11). Equation (1.4.11) can be characterized, as before, as the analysis step. An important result was shown in [92] that if we have a linear operator $\mathcal{G}(u) = \mathcal{G}u$, then for all cases of $C, \mathcal{G}, \Gamma$ we have

$$u_{\text{TP}} = \bar{u} + C\mathcal{G}^*(\mathcal{G}C\mathcal{G}^* + \Gamma)^{-1}(y - \mathcal{G}\bar{u}),$$

where $\mathcal{G}^*$ is the adjoint operator. By using the estimates (1.4.6) - (1.4.9) and update formula (1.4.10), and taking the limit as $J \to \infty$, it was shown that $u \to u_{TP}$. We omit the derivation here but for the interested reader we refer them to [79]. One issue with this method, similarly with deterministic inverse solvers, is stability, where regularization is usually added. In the work of Iglesias [77] a regularized version was considered which was motivated by the work of Hanke [71] which looked at the effect of Levenburg-Marquardt regularization for nonlinear elliptic PDEs. This added an additional step to the iterative EnKF method in the form of a discrepancy principle (or stopping rule)

$$\|y - \bar{\mathcal{G}}(u)\|_{\Gamma}^2 \leq \zeta\eta, \tag{1.4.12}$$

where $\zeta > 1/\rho \in (0, 1)$ is a regularization parameter. Incorporating this regularization modifies our update equation (1.4.10) to

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \alpha_n\Gamma)^{-1}(y_n^{(j)} - \mathcal{G}(u_n^{(j)})),$$

such that $\alpha_n \equiv \alpha_n^N$ satisfies

$$\rho\|\Gamma^{-1/2}(y^{(j)} - \bar{\mathcal{G}}_n)\|_Y \leq \alpha_n^N\|\Gamma^{1/2}(C_n^{ww} + \alpha_n^N\Gamma)^{-1}(y^{(j)} - \bar{\mathcal{G}}_n)\|_Y,$$

and where α_n is chosen based on $\alpha_n^{i+1} = 2^i\alpha_n^0$. By adopting this form of regularization numerical results were substantially improved as the discrepancy principle (1.4.12) allows for a termination, before the experiment has finished. The robustness of this approach was shown for various PDEs in [81] where a numerical investigation

was considered for tuning the additional parameters mentioned above.

The use of Levenberg-Marquardt regularization has been an effect tool in inverse problems, which originally arises from its use in nonlinear optimization. For a review on both the original Levenberg-Marquardt method, and its adaption as a form of regularization in inverse problems, we refer the reader to Appendix A.

We note as well that the discrepancy principle described as in (1.4.12) is just one possible choice one can use. This discrepancy principle arises from the Levenberg-Marquardt regularization, however once can choose other principles depending on the application or model problem. A nice review on discrepancy principles can be found in the book by Hansen [72].

1.5 Other Computational Techniques

Our last two Bayesian inverse solvers that we aim to discuss differ in that one is of a Monte-Carlo form and the other is a variational method. The first one we will discuss is an extremely popular method used in computational statistics which are sequential Monte Carlo (SMC) methods.

1.5.1 Sequential Monte Carlo Methods

Gaining inference with respect to expectations from (1.3.1) can be difficult even with complex samplers. An alternative to using MCMC methods are SMC methods [42, 43]. This is a convenient approach when π can be decomposed in a sequence of distributions, where importance sampling at each step can be implemented. These methods are based on hidden Markov models (HMM).

A HMM is composed of two processes which take values in measurable spaces $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ and $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$. We assume we have a random variable $\{X_n\}_{n\geq 0}$ which take the form of a time-homogenous Markov chain with transition kernel m_θ. Our transition kernel depends on a set of parameters $\theta \in \Theta$. Our observed random variable $\{Y_n\}_{n\geq 1}$ are conditionally independent on and distributed according to

$$Y_n|X_n \sim g_\theta(\cdot|X_n),$$

where we assume we have some initial $X_0 = x_0$ which is known. With SMC we are interested in gaining inference from $\pi_{\theta,y_{1:T}}$

$$\pi_{\theta,y_{1:T}}(f) = \frac{\int_{\mathcal{X}^T} f(x_{1:T}) \prod_{n=1}^{T}(x_{n-1}, dx_n) \prod_{n=1}^{T} g_\theta(y_n|x_n)}{l(\theta; y_{1:T})},$$

for some function $f : \mathcal{X}^T \to \mathbb{R}$ where $l(\theta; y_{1:T})$ denotes the likelihood for θ such that

$$l(\theta; y_{1:T}) := \int_{\mathcal{X}^T} \prod_{n=1}^{T} m_\theta(x_{n-1} dx_n) \prod_{n=1}^{T} g_\theta(y_n | x_n).$$

Likewise with data assimilation schemes, we can characterize SMC approaches as either smoothing or filtering methods. The simplicity of SMC methods depends hugely on the choice of the kernels. In the case of a Gaussian m_θ and g_θ we can evaluate the likelihood $l(\theta; y_{1:T})$ in analytical manner recursively. However the converse of this with non-Gaussian kernels makes it more difficult where a number of algorithms have been developed. These include the bootstrap filter, the auxiliary particle filter, sequential importance resampling and more.

1.5.2 Maximum a Posteriori Estimation

An efficient and rather simple approach which is also considered in computational statistics are Maximum a posteriori (MAP) methods. The key idea behind these methods is based on quantifying the unknown target, which in this case is a posterior distribution, through its mode. The formal equation of a MAP estimator is

$$u_{\text{MAP}} := \underset{u}{\text{argmax}}\, \mathbb{P}(u|y).$$

MAP methods pose an advantage over maximum likelihood estimation (MLE)

$$u_{\text{MLE}} := \underset{u}{\text{argmax}}\, \mathbb{P}(y|u),$$

as it incorporates the prior distribution. However both MAP and MLE coincide when the prior is given as a constant function. Despite its simplistic approach MAP estimation can be viewed as a method which can perform significantly worse due a number of reasons. One of them is when the posterior is a multi-modal distribution it is not always a guarantee that the maximum value is a good estimate of the posterior. A second reason is that it is a point estimate, which is less informative than using a distribution to gain inference. For these reasons MCMC is a more favourable choice.

1.5.3 SMC & MAP Within Inverse Problems

Due to the high computational burden which arises from MCMC for inverse problems, alternate approaches have been considered and developed, notably both MAP

and SMC methods. In the case of a linear inverse problem defined as

$$y = Au + \eta, \quad \eta \sim \mathcal{N}(0, C_\eta),$$

where A is linear operator and $u \sim \mathcal{N}(0, C_u)$ we can express our MAP estimate as

$$\begin{aligned}\mathbb{P}(u|y) &\propto \mathbb{P}(y|u)\mathbb{P}(u) \\ &= \exp(-\frac{1}{2}\left|C_u^{-1/2}u\right|^2 - \left|C_\eta^{-1/2}(y - Au)\right|^2),\end{aligned}$$

therefore leading to

$$\begin{aligned}u_{\text{MAP}} &= \underset{u}{\text{argmax}} \left(\exp(-\frac{1}{2}\left|C_u^{-1/2}u\right|^2 - \left|C_\eta^{-1/2}(y - Au)\right|^2)\right), \\ &\iff \underset{u}{\text{argmin}} \left(\frac{1}{2}\left|C_u^{-1/2}u\right|^2 - \left|C_\eta^{-1/2}(y - Au)\right|^2\right).\end{aligned}$$

The Gaussian MAP estimate in the finite-dimensional case has been considered quite extensively which is summarised well in [46]. Extending this to the ∞-dimensions we now think of a seeking a center of a ball with maximal probability and study the the limit of the center z as the radius δ tends to zero. This is aided through the Onsager-Machlup functional I which satisfies

$$\lim_{\delta \to 0} \frac{\mu(B^\delta(z_2))}{\mu(B^\delta(z_1))} = \exp(I(z_1) - I(z_2)),$$

such that $B^\delta(\cdot) \subset \mathcal{X}$ is an open ball with radius δ. Since there has been development of MAP methods such as considering certain geometric priors [50] and extensions and more rigorous analysis [7, 24].

In terms of SMC methods used for Bayesian inversion, there has been a number of papers developed which have analyzed this [16, 89] . However there still lacks an investigation on how these methods can be best optimized. This will of course depend on what particular SMC method is used. As the EnKF can be further seen as a HMM, the relationship between both methods could pose a nice combination for future work to be developed.

1.6 Outline of Thesis

This thesis is to consider a number of avenues of investigation concerning EKI. Furthermore, much of our focus is on geometric inversion problems, in which (part of) the unknown is a geometrical feature such as an interface. Thus far much of the

understanding of PDE-constrained Bayesian inversion have been based on MCMC. Part of the reason behind this is that there are many problems in uncertainty quantification in which MCMC has been traditionally used, due to their high-level performance and for its sampling capabilities of a probability distribution. EKI is an optimization based alternative to MCMC which, like MCMC, operates via an ensemble, but has distinct computational advantages: it typically requires far fewer forward model evaluations.

Rather than focusing on understanding one aspect of EKI, the thesis is split into numerous avenues which answer different key and important questions in the field of inverse problems, but with the common theme that we aim to do so in a setting that aligns with subsections 1.4.3 and 1.4.4. In particular the motivation behind this thesis is to consider a number of avenues in EKI which are based on solving geometric inverse problems, i.e. where our underlying unknown contains some form of geometry, such as in the geosciences and in medical imaging. Examples include piecewise constant, piecewise continuous and level set functions. Understanding inverse problems through these functions is important as numerous fields require recovering some form of geometry such as geosciences and medical imaging. Some of the questions that this thesis aims to address are:

- Can we effectively use EKI to recover some form of geometry from our unknown, which includes various discontinuities?
- Given the computational burden that can arise with the forward operator, can we reduce this cost in a practical manner and what kind of cost reduction is possible?
- Can we transfer ideas from hierarchical Bayesian inversion to the EKI methodology, and in doing so find improved inversion strategies?
- Given the current extent of the literature on Bayesian-related inversion, and the EKI in particular, can we extend this to new and potentially more challenging PDEs applications?

The layout and description of each chapter is provided as follows:

1.6.1 Chapter 2. Parameterizations of Ensemble Kalman Inversion

The prior form in Bayesian inversion is of crucial importance as depending on how it is chosen, our reconstructions can be quite varied. One way to overcome this issue is to understand the inverse problem in a hierarchical manner. This introduces further

unknowns where these additional unknowns correspond to hyperparameters which allow for parsimonious representation of the unknown, and which are learned from the data, along with the unknown itself. These priors are usually defined through three key parameters; (i) amplitude $\sigma \in \mathbb{R}$, (ii) smoothness $\alpha \in \mathbb{R}^+$ and (iii) inverse length-scale $\tau = 1/\ell \in \mathbb{R}^+$ which appear in the covariance structure $\mathcal{C}$

$$\mathcal{C} = \sigma^2(\tau^2 I - \Delta)^{-\alpha}, \tag{1.6.1}$$

where Δ is the Laplacian operator. However with EKI still lies a limitation which is known as the subspace property, that states however your prior is chosen, your reconstructed solution will be of a similar form. By placing a Gaussian initial ensemble with prescribed hyperparameters we can expect our end result to have a similar result with the same value of hyperparameters. One technique we use to break away from this property is to generate the Gaussian prior through a Stochastic partial differential approach, which in turn states that the solution u of

$$(\tau^2 I - \triangle)^{\frac{\alpha}{2}} u = \tau^{-2/d}\sqrt{\beta}\xi,$$

omits a structure similar to that of the covariance (1.6.1). By doing so we can successfully break away, and effectively learn the various unknowns. This is based on two hierarchical approaches: the centered and non-centered case. We consider this for a variety of PDEs and adopt level set techniques. A further extension we look at in this chapter is the question of whether we can extend the modelling of our hyperparameters from a scalar field to a random field, as done with our underlying unknown. For this we adopt the length-scale as both a Gaussian random field, but also a newly adopted random field in the area of Bayesian inversion, a Cauchy random field. The latter poses an advantage over the former as it performs better for edge-preserving and rougher features. We further consider analyzing some limit analysis of hierarchical EKI which can be seen as an extension for the recent work that was done on non-hierarchical EKI.

The work in this chapter appeared as a jointly co-authored paper [33]. Although the four authors, myself included, contributed in unison to the methodology developed in the paper, and to the choice of numerical experiments, the majority of the computational work is mine alone.

1.6.2 Chapter 3. Analysis of Hierarchical Ensemble Kalman Inversion

The EnKF has showcased a high level of applicability to a number of scientific disciplines, due to its ease of implementation and low computational cost. However despite this is a lack of clear foundational understanding, both for the traditional EnKF from data assimilation, and EKI for inverse problems.. Much of the work since the derivation of the work has been on improving the algorithm under different settings. Only till recently there has been some work which has aimed at understanding this. However for EKI there has only been one piece of work that has looked at developing diffusion limits aiming to understand the long-term behaviour. This chapter aims to complement Chapter 2. in that it provides a more analytical understanding of hierarchical EKI. We look to extend the results of the work done on EKI, namely understand the long-term behaviour. It was shown that in the linear case EKI could be interpreted as a gradient-flow structure. In the hierarchical case we see that this also holds, for both the centered and non-centered case. Aside from diffusion limits and gradient-flow structures, we aim to prove whether using both approaches allow us to break away from the subspace property, in the case of both a discrete and continuous case. Finally we consider hierarchical aspects of certain variants of the EnKF, which include localization and variance inflation. Both these techniques are known to break from the induced property, therefor we take a hierarchical approach and see if if we have the similar effect if not a better performance. This is highlighted through a numerical example on a 1D linear elliptic PDE.

This chapter was solely carried out by myself, and has been submitted as the solo-authored paper [31]. However the work benefitted from a number of individuals including Dr. Marco Iglesias and Prof. Claudia Schillings.

1.6.3 Chapter 4. Reduced Basis Methods for Bayesian Inverse Problems

Most inverse problem techniques involve multiple evaluations of the forward model, which can be computationally expensive. Therefore a fundamental question that arises is how to reduce the cost of the forward solver in an efficient manner which still allows for high quality inversion? Traditional PDE solvers such as the finite difference and finite element method have shown to work well under various conditions, but can be costly in high-dimensional problems. In particular this question has gained an interest from uncertainty quantification where various methods have

been developed to tackle these issues. These methods are commonly referred to as reduced order models (ROMs). One of these methods, which we consider in this chapter, are reduced basis methods (RBMs). RBMs can be interpreted as a dimension reduction method based on a Galerkin projection which is based on a greedy algorithm. The solution of the RBM is controlled via an a posteriori error bound. It is specific to parameterized PDEs, which are PDEs where the input and output are induced by some random parameter. As an example in an elliptic PDE we would have both a random diffusion coefficient $\kappa(x;\vartheta)$ and a random solution $p(x;\vartheta)$ induced by some parameter ϑ in a parameter space Γ, given by

$$\begin{aligned} -\nabla \cdot (\kappa(x;\vartheta)\nabla p(x;\vartheta)) &= f, \quad \in D, \\ p(x;\vartheta) &= 0, \quad \in \partial D. \end{aligned}$$

The RBM projects the solution p_N onto a smaller space X_N and aims to build a RB space by solving the following system given by the variational formulation

$$A(p_N, q; \vartheta) = F(q), \quad \forall q \in X_N.$$

The focus on this chapter is two fold; firstly that we consider the proposal of a new set of points to represent the parameter space Γ. Traditional choices of these points include uniform points, random points and Curtis-Clenshaw points. Our set of points we propose are Lebesgue optimal points, where our motivation arises from the work of Chen et al. [35]. By using these points we aim to see if we can get more accurate results and see how these points scale within our experiments. The second aim of this chapter is to transfer these ideas from ROMs, and apply them in to Bayesian inverse problems. Specifically we consider the inverse problem of Darcy flow, where in particular we focus on geometric inverse problems such that our random coefficient takes a piecewise constant form. We briefly consider the inverse problem of impedance tomography where we derive particular bounds related to the RBM.

The work conducted on the forward problem can be viewed as extension to the 2D case of existing work in 1D; in particular I studied the 1D case as part of a research study groups project in my MSc year Warwick with three other MSc students. This study group was led by Dr's Claudia Schillings and Aretha Teckentrup and their input was instrumental in setting directions for this work.

1.6.4 Chapter 5. A Bayesian Formulation of the Inverse Eikonal Equation

Since the development of the Bayesian approach in the infinite dimensional case, there have been a wide class of PDE-related inverse problems that have been analyzed. Most of this work has been concerned with elliptic and parabolic problems, such as those arising in groundwater flow, electrical impedance tomography and electro-magnetics (elliptic) and data assimilation in fluid, through the Navier-Stokes equation (parabolic). One class of PDEs that has a rich theory behind it, but has not being analyzed greatly in an inverse setting, are Hamilton-Jacobi (HJ) equations. We are interested in one particular HJ equation which is the eikonal equation. The eikonal equation is an important PDE which arises in numerous application such as computer vision and geosciences. The forward problem is concerned with finding T solving

$$\begin{aligned} |\nabla T(x)| &= u(x) \quad x \in \Omega \setminus \{x_0\}, \\ T(x_0) &= 0, \\ \nabla T(x) \cdot \nu(x) &\geq 0, \quad x \in \Gamma. \end{aligned}$$

Here $T(x)$ represents the shortest distance of a travel-time from a point x_0 on the boundary to a point x on the domain, from $u(x)$ known as the slowness function, characterizes the properties of the medium. The solution, rather than being expressed through a weak formulation, is given as a minimisation procedure, as done in optimal control theory, where

$$T(x) = \inf_{\zeta} \Big\{ \int_a^b u(\zeta(r))|\zeta'(r)|dr \mid \zeta \in W^{1,\infty}([0,1], \bar{\Omega}), \zeta(a) = x_0, \zeta(b) = x \Big\}.$$

The inverse eikonal equation has only been considered in the deterministic setting by Decklenick et al. [52]. Our motivation is to formulate the inverse eikonal equation in a Bayesian inverse setting. By doing so we aim to tackle uncertainty that can arise within the slowness function, such as discontinuities. Unlike the diffusion coefficient in elliptic PDEs, the slowness function poses stricter regularity conditions and can allow for only a number of discontinuities. We aim to understand this behaviour by testing a variety of various priors, including both Gaussian and geometric. We further consider deriving analytical results for the inverse eikonal equation, namely well-posedness for some of the priors considered. Hierarchical techniques developed in previous chapters will also be further applied.

The overall problem formulation and methodology in this chapter was carried

out in conjunction with Prof. Charlie Elliott, Prof. Andrew Stuart and Dr. Vanessa Styles'; however the majority of the computational work was carried out by myself, with some help from Dr. Ollie Dunbar on coding aspects relating to the forward Eikonal solver.

Chapter 2

Parameterizations for ensemble Kalman inversion

2.1 Overview

In this chapter we introduce various hierarchical approaches to understand EKI in a hierarchical manner. This includes two approaches: the centered and non-centered approach, where we derive various scaling limits and test our methodology on numerous geometric PDEs.

2.2 Introduction

2.2.1 Content

Consider finding u from y where

$$y = \mathcal{G}(u) + \eta, \tag{2.2.1}$$

$\mathcal{G}$ is a forward map taking the unknown parameter u into the data space, and η represents noise. Ensemble Kalman inversion is an attractive technique that has shown considerable success in the solution of such problems. Whilst it is derived from the application of Kalman-like thinking, with means and covariances computed from an empirical ensemble, it essentially acts as a black-box derivative-free optimizer which requires only evaluation of the forward map $\mathcal{G}(\cdot)$; in practice it can often return good solutions to inverse problems with relatively few forward map evaluations. However the choice of parameterization of the unknown is key to the success of the method. In this chapter we will demonstrate how carefully thought out parameterizations

can have substantial impact in the quality of the reconstruction.

Although our viewpoint in this chapter is to consider Ensemble Kalman inversion as an optimization method, and evaluate it from this perspective, there is considerable insight to be gained from the perspective of Bayesian inversion; this is despite the fact that the the algorithm does not, in general, recover or sample the true Bayesian posterior distribution of the inverse problem. Algorithms that can, with controllable error, approximately sample from the true posterior distribution are commonly referred to as fully Bayesian, with examples including Markov Chain Monte Carlo and sequential Monte Carlo. Ensemble Kalman inversion is not fully Bayesian but the link to Bayesian inversion remains important as we now explain. There is considerable literature available about methods to improve fully Bayesian approaches to the inverse problem through for example, geometric and hierarchical parameterizations of the unknown. The purpose of this chapter is to demonstrate how these ideas from Bayesian inversion may be used with some success to improve the capability of ensemble Kalman methods considered as optimizers. In view of the relatively low computational cost of the ensemble methods, in comparison with fully Bayesian inversion, this cross-fertilization of ideas has the potential to be quite fruitful.

2.2.2 Literature Review

The Kalman filter (KF) [88] was developed to sequentially update the probability distribution on states of partially observed linear Gaussian systems, and subsequently generalized to nonlinear problems in the form of the extended Kalman filter. However for high-dimensional systems the size of covariances makes use of these methods prohibitive. In 1994 Evensen [57, 58] proposed a Monte-Carlo based nonlinear Kalman filter which tackled this issue by using an ensemble of particles to represent the covariances and mean, resulting in what is now known as the ensemble Kalman filter (EnKF). A major success story for the EnKF has been in weather prediction models [4, 74], but it has also been deployed in numerous applications domains, including the reservoir engineering community [1] and in oceanography [59]. Variants on the idea include the randomized maximum likelihood (RML) method [97], and algorithms such as the ensemble square-root Kalman filter [140].

In this chapter we are primarily interested in use of ensemble Kalman methods to study inverse problems for parameter estimation, an approach pioneered for oil industry applications where the inverse problem is known as history matching [93, 97]; the paper [56] contains an insightful analysis of the methodology in the large ensemble limit. In this application domain such inversion methods are some-

times referred to as ensemble Kalman smoothers, although the nomenclature is not uniform. We will simply refer to *ensemble Kalman inversion* (EKI). The methodology is formulated quite generally in [79], independently of oil industry applications; and in [78] it is shown that the method performs well as an optimizer but does not capture true posterior uncertainty, in the context of oil industry applications.

The ideas introduced in this chapter concerning parameterization are independent of the particular implementation of the EKI method used; in all our numerical experiments we will use the form of iterative regularization proposed by Iglesias [77]. The general philosophy behind the method is that, as the algorithm is iterated, the solution to the inverse problem should approach the truth in the small noise limit, and hence that any regularization introduced should diminish in influence in this limit. However, because the convergence theory for ensemble Kalman inversion is in its infancy, the choice of iterative regularization is made by analogy with classical iterative methods that have been used for inverse problems [10, 12, 71], with the ensemble method using empirical covariances in place of derivatives or adjoints of the forward solver. The resulting iterative method is an ensemble version of the Levenberg-Marquardt algorithm with the inclusion of regularization as in [71].

In the form of EKI that we use the linear span of the initial ensemble is preserved by the iteration [79, 93]. The initial ensemble thus encodes prior information about the solution of the problem. This means that choice of the parameterization of the method, as well as the choice initial subspace, is key to its performance. Based on experience with (Bayesian) statistical modelling we will introduce geometric and hierarchical priors that address the issue of making good parameterizations, and we will draw from those priors to create the initial ensemble. Hierarchical models have been extensively studied in the fields of computational statistics and machine learning [94, 115, 139]. Their use in the context of Monte Carlo Markov chain (MCMC) methods for Bayesian inverse problems is overviewed in [116]; see [47, 127] for application oriented work. One important outcome of research in this area is that learning parameters such as length-scale, amplitude and regularity within Gaussian random field priors (such as Whittle-Matérn) can be of significant value [6, 98, 123]. In a series of recent papers this hierarchical modelling was extended to allow for length-scale which is itself spatially varying [102, 123]. The development of hierarchical methods within EKI, rather than fully Bayesian MCMC, has been limited to date, with the primary contribution being the work [53] where the methodology was based on building large ensembles from multiple Gaussians assigned different weights. However this work requires that the correct hierarchical parameter is in the ensemble if it is to be successful. We also note that there is some work in hier-

archical EnKF within the context of state estimation; see [143] and the references therein.

In addition to hierarchical approaches we will also study geometric parameterizations. These can be of use when the geometric object has known form, such as faults, channels and inclusions, or when it is of unknown topology the level set method may be used [128]. We will build on recent Bayesian implementation of these ideas; see [80, 81] and references therein.

2.2.3 Contribution of This Work

Our main contribution is to establish the importance of novel parameterizations which have the potential to substantially improve the performance of EKI. Although our perspective on EKI is one of optimization, the methods we introduce are all based on taking established and emerging methods from Bayesian statistics and developing them in the context of the ensemble methods. The connection to Bayesian statistics is exploited to provide insights about how to make these methods efficient. The resulting methods are illustrated by means of examples arising in both electrical impedance tomography, groundwater flow and source inversion. The contributions are:

- We develop hierarchical approaches for EKI, based on solving for the unknown function and unknown scalars which parameterize the prior.
- We generalize these hierarchical approaches to EKI to include unknown fields which parameterize the prior, rather than scalars.
- We demonstrate the key role of choosing non-centered variables when implementing hierarchical methods.
- We show the potential for geometric hierarchical priors, including the level set parameterization, for piecewise continuous reconstructions.

2.2.4 Organization

The layout of the chapter is as follows, in section 3.3 we discuss different approaches to parameterizing inverse problems. We begin by conveying the main ideas in subsection 2.3.1 in abstract. In subsection 2.3.2 we describe these ideas more concretely. In section 2.4 we describe the hierarchical version of iterative EKI as used in this chapter, and section 2.5 describes the model problems that we use to illustrate the power of the proposed parameterizations. Numerical results are presented in section 3.5, whilst in section 3.6 we make some concluding remarks.

2.2.5 Notation

Throughout the chapter we make use of common notation for Hilbert space norms and inner products, $\|\cdot\|, \langle\cdot\rangle$. We will assume that $\mathcal{X}$ and $\mathcal{Y}$ are two separable Hilbert spaces which are linked through the forward operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$. This nonlinear operator can be thought of as mapping from the space of unknown parameters $\mathcal{X}$ to the observation space $\mathcal{Y}$. Our additive noise for the inverse problems will be denoted by $\eta \sim N(0, \Gamma)$ where $\Gamma : \mathcal{Y} \to \mathcal{Y}$ is a self-adjoint positive operator. For any such operator we define $\langle\cdot,\cdot\rangle_\Gamma = \langle\Gamma^{-1/2}\cdot, \Gamma^{-1/2}\cdot\rangle$ and $\|\cdot\|_\Gamma = \|\Gamma^{-1/2}\cdot\|$, and for finite dimensions $|\cdot|_\Gamma = |\Gamma^{-1/2}\cdot|$ with $|\cdot|$the Euclidean norm. If the Gaussian measure associated to η is supported on $\mathcal{Y}$ then we will require Γ to be trace-class; however we will also consider white noise whose support is on a larger space than $\mathcal{Y}$ and for which the trace-class condition fails in the infinite dimensional setting.

2.3 Inverse Problem

2.3.1 Main Idea

Non-Hierarchical Inverse Problem

We are interested in the recovery of $u \in \mathcal{X}$ from measurements of $y \in \mathcal{Y}$ given by equation (2.2.1) in which, recall, η is additive Gaussian noise. In the Bayesian approach to inverse problems we treat each quantity within (2.2.1) as a random variable. Via an application of Bayes' Theorem [1] [137] we can characterize the conditional distribution of $u|y$ via

$$\mathbb{P}(u|y) \propto \mathbb{P}(y|u) \times \mathbb{P}(u), \tag{2.3.1}$$

where $\mathbb{P}(u)$ is the prior distribution, $\mathbb{P}(u|y)$ is the posterior distribution and $\mathbb{P}(y|u)$ is the likelihood. Although we view EKI as an optimizer in this chapter, the Bayesian formulation of (2.2.1) is important because we derive the initial ensemble from the prior distribution $\mathbb{P}(u)$. From an optimization viewpoint, our goal is to make the following least squares objective function small:

$$\Phi(u; y) = \frac{1}{2}|y - \mathcal{G}(u)|_\Gamma^2. \tag{2.3.2}$$

[1] We write all instances of Bayes' Theorem in finite dimensions for simplicity; extension to Bayes' Theorem for functions is straightforward but not central to this chapter and so we avoid the extra notation that would be needed for this.

This (upto an irrelevant additive constant) is the negative log likelihood since it is assumed that η is a mean-zero Gaussian with covariance Γ.

Centered Hierarchical Inverse Problem

In many applications it can be advantageous to add additional unknowns θ to the inversion process. In particular these may enter through the prior, as in hierarchical methods, [116], and we will refer to such parameters as hyperparameters. The inverse problem is then the recovery of (u, θ) from measurements of y given again by (2.2.1). The additional parameterization of the prior results in Bayes' Theorem in the form

$$\mathbb{P}(u, \theta|y) \propto \mathbb{P}(y|u) \times \mathbb{P}(u, \theta). \tag{2.3.3}$$

Prior samples, used to initialize the ensemble smoother, will then be of the pair (u, θ). What we will term *centered* hierarchical methods (a terminology we discuss in Section 2.4) typically involve factorization of the prior in the form $\mathbb{P}(u, \theta) = \mathbb{P}(u|\theta)\mathbb{P}(\theta)$. From an optimization point of view our goal is again to make the objective function (2.3.2) small, but now using hierarchical parameterization to construct the initial ensemble. [2]

Non-Centered Hierarchical Inverse Problem

Another variant of the inverse problem that is particularly relevant for hierarchical methods, in which θ enters only the prior, is *non-centered* reparameterization (a further terminology discussed in Section 2.4). We introduce the transformation $T : (\xi, \theta) \to u$ and note that (2.2.1) then becomes

$$y = \mathcal{G}(T(\xi, \theta)) + \eta, \tag{2.3.4}$$

and Bayes' Theorem then reads

$$\mathbb{P}(\xi, \theta|y) \propto \mathbb{P}(y|\xi, \theta) \times \mathbb{P}(\xi, \theta). \tag{2.3.5}$$

Prior samples, again used to initialize the ensemble smoother, will then be of the pair (ξ, θ). Typically the change of variables from u to ξ is introduced so that ξ and θ are independent under the prior: $\mathbb{P}(\xi, \theta) = \mathbb{P}(\xi)\mathbb{P}(\theta)$. As a result the inverse problem (2.3.4) is different to that appearing in (2.2.1), in terms of both the prior and the

[2] We note that hyperparameters θ may also enter the likelihood as well as the prior if the state variable is re-scaled in a hyperparameter-dependent fashion, as happens in the version of the Bayesian level set method advocated in [47].

likelihood. This non-centered approach is equivalent to the ancillary augmentation technique discussed in [115] which discusses the decoupling of u and θ through the variable ξ. From an optimization viewpoint, our goal is to make the following least squares objective function small:

$$\Phi(\xi, \theta; y) = \frac{1}{2}|y - \mathcal{G}(T(\xi, \theta))|^2_\Gamma. \tag{2.3.6}$$

In this chapter we consider the application of EKI for solution of the inverse problems (2.2.1), non-hierarchically and centered-hierarchically, and (2.3.4) non-centered hierarchically. Although this method has a statistical derivation, the work in [79, 90] demonstrates that the method may be thought of as a derivative-free optimizer that approximates the least squares problem (2.3.2) and (2.3.6) rather than sampling from the relevant Bayesian posterior distribution. We will show that the use of iterative EKI, as proposed by Iglesias in [77], can effectively solve a wide range of challenging inversion problems, if judiciously parameterized.

2.3.2 Details of Parameterizations

In this section we describe in detail several classes of parameterizations that we use in this chapter. The first and second are geometric parameterizations, ideal for piecewise continuous reconstructions with unknown interfaces. The third and fourth are hierarchical methods which introduce an unknown length-scale, and regularity parameter, into the inversion. For the two geometric problems we initially formulate in terms of trying to find a function $w : D \mapsto \mathbb{R}$, D a subset of $\mathbb{R}^d$, and then reparameterize w. For the hierarchical problems we initially formulate in terms of trying to find a function $u : D \mapsto \mathbb{R}$, and then append parameters θ and also rewrite in terms of $(\xi, \theta) \mapsto u$.

Geometric Approach – Finite Dimensional Parameterization

In many problems of interest the unknown function w has discontinuities, determination of which forms part of the solution of the inverse problem. To tackle such problems it may be useful to write w in the form

$$w(x) = \sum_{i=1}^{n} u_i(x)\chi_{D_i}(x). \tag{2.3.7}$$

Here the union of the disjoint sets D_i is the whole domain D. If we assume that the configuration of the D_i is determined by a finite set of scalars θ and let u denote the union of the functions u_i and the parameters θ then we may rewrite the inverse problem in the form (2.2.1). The case where the number of subdomains n is unknown would be an interesting and useful extension of this work; but we do not consider it here.

Geometric Approach – Infinite Dimensional Parameterization

If the interface boundary is not readily described by a finite number of parameters we may use the level set idea. For example if the field w takes two known values $w^{\pm}$ with unknown interfaces between them we may write

$$w(x) = w^{+}\mathbb{I}_{u>0}(x) + w^{-}\mathbb{I}_{u<0}(x), \tag{2.3.8}$$

and formulate the inverse problem in the form (2.2.1) for u. This idea may be generalized to functions which take an arbitrary number of constant values, through the introduction of level sets other than $u = 0$, or through vector level sets functions u.

Scalar-valued Hierarchical Parameterizations

To illustrate ideas we will concentrate on Gaussian priors of Whittle-Matérn type. These are characterized by a covariance function of the form

$$c(x,y) = \sigma^2 \frac{2^{1-\alpha+d/2}}{\Gamma(\alpha - d/2)} \left(\frac{|x-y|}{\ell} \right)^{\alpha-d/2} K_{\alpha-d/2}\left(\frac{|x-y|}{\ell} \right), \quad x, y \in \mathbb{R}^d, \tag{2.3.9}$$

where $K_\cdot$ is a modified Bessel function of the second kind, $\sigma^2 > 0$ is the variance and $\Gamma(\cdot)$ is a Gamma function. We will always ensure that $\ell > 0$ and $\alpha > d/2$ so that draws from the Gaussian are well-defined and continuous. On the unbounded domain $\mathbb{R}^d$, samples from this process may be generated by solving the stochastic PDE

$$(I - \ell^2 \triangle)^{\frac{\alpha}{2}} u = \ell^{d/2} \sqrt{\beta} \xi, \tag{2.3.10}$$

where $\xi \in H^{-s}(D)$, $s > \frac{d}{2}$, is a Gaussian white noise, i.e. $\xi \sim N(0, I)$, and

$$\beta = \sigma^2 \frac{2^d \pi^{d/2} \Gamma(\alpha)}{\Gamma(\alpha - \frac{d}{2})}.$$

In this chapter we will work with the scalar hierarchical parameters α and

$\tau = \ell^{-1}$. Putting $\ell = \tau^{-1}$ into (2.3.10) gives the stochastic PDE

$$\mathcal{C}_{\alpha,\tau}^{-\frac{1}{2}} u = \xi, \tag{2.3.11}$$

where the covariance operator $\mathcal{C}_{\alpha,\tau}$ has the form $\mathcal{C}_{\alpha,\tau} = \tau^{2\alpha-d}\beta(\tau^2 I - \Delta)^{-\alpha}$. Throughout this chapter we choose σ such that $\beta = \tau^{d-2\alpha}$ so that

$$\mathcal{C}_{\alpha,\tau} = (\tau^2 I - \Delta)^{-\alpha}, \tag{2.3.12}$$

and we equip the operator Δ with Dirichlet boundary conditions on D. These choices simplify the exposition, but are not an integral part of the methodology; different choices could be made.

We have thus formulated the inverse problem in the form of the centered hierarchical inverse problem (2.2.1) for (u, θ) with $\theta = (\alpha, \tau)$. The parameter θ enters only through the prior as in this particular case it does not appear in the likelihood. In this chapter we will place uniform priors on α and τ, which will be specified in subsection 2.6.1. We may also work with the variables (ξ, θ) noting that (2.3.11) defines a map $T : (\xi, \theta) \mapsto u$ and we have formulated the inverse problem in the form (2.3.4), the non-centered hierachical form. In Figures 2.1 and 2.2 we display random samples from (2.3.12) with imposed Dirichlet boundary conditions and varying values of the inverse length-scale τ and the regularity α. These samples are constructed in the domain $D = [0, 1]^2$.

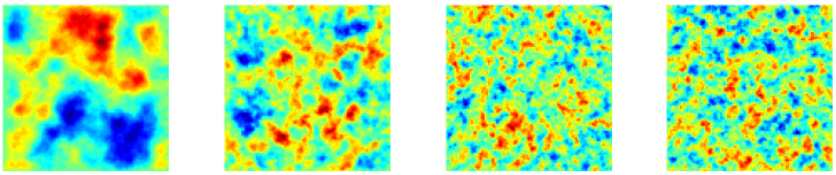

Figure 2.1: Modified inverse length-scale for $\tau = 10$, 25, 50 and 100. Here $\alpha = 1.6$.

Function-Valued Hierarchical Parameterization

In order to represent non-stationary features it is of interest to allow hierarchical parameters to themselves vary in space. To this end we will also seek to generalize (2.3.10) and work with the form

$$\left(I - \ell(x; v)^2 \Delta\right)^{\frac{\alpha}{2}} u = \ell(x; v)^{\frac{d}{2}} \xi. \tag{2.3.13}$$

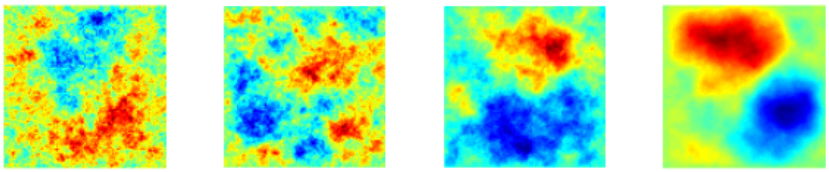

Figure 2.2: Modified regularity for $\alpha = 1.1$, 1.3, 1.5 and 1.9. Here $\tau = 15$.

where (We have set $\beta = 1$ for simplicity). In order to ensure that the length-scale is positive we will write it in the form

$$\ell(x; v) = g(v(x)), \tag{2.3.14}$$

for some positive monotonic increasing function $g(\cdot)$. We thus have a formulation as a centered hierarchical inverse problem in the form (2.2.1) noting that hyperparameter $\theta = v$ is here a function and enters only through the prior. We may also formulate inversion in terms of the variables (ξ, v) giving the inverse problem in the form (2.3.4) with $\theta = v$. We will consider two forms of prior on v. The first is based on a Gaussian random field with Whittle-Matérn covariance function (2.3.9) and we then choose $g(v) = \exp(v)$. The second, which will apply only in one dimension, is to consider a one-dimensional Cauchy process, as in [102]. In particular we will construct $\ell(x; v)$ by employing a one-dimensional Cauchy process $v(x)$, which is an α-stable Lévy motion with $\alpha = 1$ with Cauchy increments on the interval δ given by the density function f, and positivity-inducing function g, where

$$f(x) = \frac{\delta}{\pi(\delta^2 + x^2)}, \quad g(s) = \frac{a}{b + c|s|} + d, \tag{2.3.15}$$

such that $a, b, c, d > 0$ are constants. [3] Samples from these two priors on v, and hence ℓ, are shown in Figures 2.3 and 2.4.

2.4 Iterative Ensemble Kalman Inversion

In this section we describe iterative EKI as implemented in this chapter. We outline it first for inverse problems as parameterized in equations (2.2.1) and then discuss

[3] We note here that α has a different meaning from the parameter α used in the covariance function of a Gaussian prior in, for example, (2.3.12). We abuse notation in this way because the parameter α is widely used in the literature in both contexts.

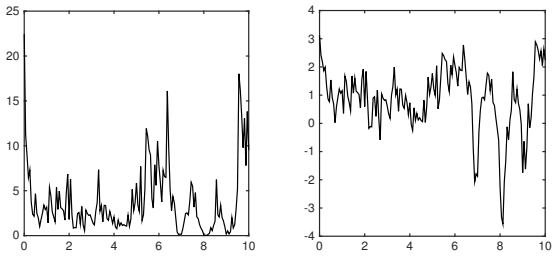

Figure 2.3: Gaussian random field. Left: Length-scale realization $\ell(x)$. Right: Realization of $v(x)$.

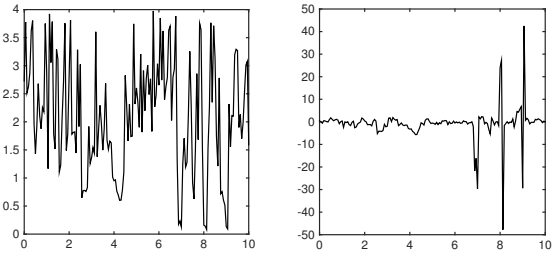

Figure 2.4: Cauchy random field. Left: Length-scale realization $\ell(x)$. Right: Realization of $v(x)$.

generalizations to centered hierarchical inversion and the non-centered hierarchical inversion (2.3.4). We briefly mention the continuous time limits of the methods as these provide insight into how EKI works, and the effect of re-parameterizing; the study of continuous limits from EKI was introduced in [129] and further details concerning their application to the problems considered here may be found in [31].

2.4.1 Formulation for (2.2.1)

The form of iterative EKI that we use is that employed in [77]. When applied to the inverse problem (2.2.1) it takes the following form, in which the subscript n denotes the iteration step, and the superscript (j) the ensemble member:

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \Upsilon_n\Gamma)^{-1}(y - \mathcal{G}(u_n^{(j)})). \tag{2.4.1}$$

The empirical covariances C_n^{uw}, C_n^{ww} are given by

$$C_n^{uw} = \frac{1}{J-1}\sum_{j=1}^{J}(u_n^{(j)} - \bar{u}_n) \otimes (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}_n) \tag{2.4.2}$$

$$C_n^{ww} = \frac{1}{J-1}\sum_{j=1}^{J}(\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}_n) \otimes (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}_n). \tag{2.4.3}$$

Here $\bar{u}_n$ denotes the average of $u_n^{(j)}$ over all ensemble members and $\bar{\mathcal{G}}_n$ denotes the average of $\mathcal{G}(u_n^{(j)})$ over all ensemble members. The parameter Υ_n is chosen to ensure that

$$\|y - \bar{\mathcal{G}}_n\|_\Gamma \leq \zeta\eta, \tag{2.4.4}$$

a form of discrepancy principle which avoids over-fitting.

If we define

$$d_n^{(j,m)} = \langle (C_n^{ww} + \Upsilon_n\Gamma)^{-1}(\mathcal{G}(u_n^{(j)}) - y), \mathcal{G}(u_n^{(m)}) - \bar{\mathcal{G}}_n\rangle,$$

then we see that

$$u_{n+1}^{(j)} = u_n^{(j)} - \frac{1}{J-1}\sum_{m=1}^{J} d_n^{(j,m)} u_n^{(m)}, \tag{2.4.5}$$

and it is apparent that the algorithm will preserve the linear span of the initial ensemble $\{u_0^{(j)}\}_{j=1}^J$. We describe the details of how Υ_n is chosen in the next subsection where we display the algorithm in full for a generalization of the setting of (2.2.1) to the hierarchical setting.

To write down the continuous-time limit of the EKI we consider the setting in which $\Upsilon_n^{-1} \equiv (J-1)h$ and view $u_n^{(j)}$ as approximating a function $u^{(j)}(t)$ at time $t = nh$. If we define

$$d^{(j,m)} = \langle \Gamma^{-1}(\mathcal{G}(u^{(j)}) - y), \mathcal{G}(u^{(m)}) - \bar{\mathcal{G}}\rangle,$$

with obvious definition of $\bar{\mathcal{G}}$, then we obtain the continuous-time limit

$$\dot{u}^{(j)} = -\sum_{m=1}^{J} d^{(j,m)} u^{(m)}, \tag{2.4.6}$$

where $\dot{u}^{(j)}$ denotes the standard time derivative of $u^{(j)}$ viewed as solving an ordinary differential equation; because the algorithm preserves the linear span of the initial ensemble [129] the dynamics take place in a finite dimensional space, even if $\mathcal{X}$ is

Algorithm 3 Hierarchical Iterative Kalman Method (Centered Version).

Let $\{u_0^{(j)}, \theta_0^{(j)}\}_{j=1}^J \subset \mathcal{X}$ be the initial ensemble with J elements.
Further let $\rho \in (0,1)$ with $\zeta > \frac{1}{\rho}$ and $\theta = (\alpha, \tau)$.

Generate $\{u_0^{(j)}, \theta_0^{(j)}\}$ i.i.d. from the prior $\mathbb{P}(u, \theta)$, with synthetic data
$y_{n+1}^{(j)} = y + \eta_{n+1}^{(j)}, \quad \eta^{(j)} \sim N(0, \Gamma)$ i.i.d.

Then for $n = 1, \ldots$

1. **Prediction step:** Evaluate the forward map $w_n^{(j)} = \mathcal{G}(u_n^{(j)})$,

   and define $\bar{w}_n = \frac{1}{J} \sum_{j=1}^J w_n^{(j)}$.

2. **Discrepancy principle:** If $\|\Gamma^{-1/2}(y - \bar{w}_n)\|_{\mathcal{Y}} \leq \zeta\eta$, stop!

   Output $\bar{u}_n = \frac{1}{J} \sum_{j=1}^J u_n^{(j)}$ and $\bar{\theta}_n = \frac{1}{J} \sum_{j=1}^J \theta_n^{(j)}$.

3. **Analysis step:** Define sample covariances:

   $C_n^{uw} = \frac{1}{J-1} \sum_{j=1}^J (u^{(j)} - \bar{u}) \otimes (\mathcal{G}(u^{(j)}) - \bar{\mathcal{G}})$,
   $C_n^{\theta w} = \frac{1}{J-1} \sum_{j=1}^J (\theta^{(j)} - \bar{\theta}) \otimes (\mathcal{G}(u^{(j)}) - \bar{\mathcal{G}})$,
   $C_n^{ww} = \frac{1}{J-1} \sum_{j=1}^J (\mathcal{G}(u^{(j)}) - \bar{\mathcal{G}}) \otimes (\mathcal{G}(u^{(j)}) - \bar{\mathcal{G}})$.

   Update each ensemble member as follows

   $u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \Upsilon_n \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)}))$,
   $\theta_{n+1}^{(j)} = \theta_n^{(j)} + C_n^{\theta w}(C_n^{ww} + \Upsilon_n \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)}))$,

   where Υ_n is chosen as $\Upsilon_n^{i+1} = 2^i \Upsilon_n^0$,

   where Υ_n^0 is an initial guess. We then define $\Upsilon_n \equiv \Upsilon_n^N$ where N is the first integer such that

   $\rho \|\Gamma^{-1/2}(y^{(j)} - \bar{w}_n)\|_{\mathcal{Y}} \leq \Upsilon_n^N \|\Gamma^{1/2}(C_n^{ww} + \Upsilon_n^N \Gamma)^{-1}(y^{(j)} - \bar{w}_n)\|_{\mathcal{Y}}$.

infinite dimensional. Note that $d^{(j,m)}$ depends on $\{u^{(k)}\}_{k=1}^J$ and that the dynamical system couples the ensemble members.

2.4.2 Generalization for Centered Hierarchical Inversion

Algorithm 3 shows the generalization of (2.4.1) to the setting of centered hierarchical inversion. The ensemble is now over both u and θ and cross covariances from the observational space to both the u and θ spaces are required. Algorithm 3 also spells out in detail how the parameter Υ_n is chosen. We now define

$$d_n^{(j,m)} = \langle (C_n^{ww} + \Upsilon_n \Gamma)^{-1} (\mathcal{G}(u_n^{(j)}) - y), \mathcal{G}(u_n^{(m)}) - \bar{\mathcal{G}}_n \rangle,$$

and we see that

$$u_{n+1}^{(j)} = u_n^{(j)} - \frac{1}{J-1} \sum_{m=1}^{J} d_n^{(j,m)} u_n^{(m)}, \tag{2.4.7}$$

and

$$\theta_{n+1}^{(j)} = \theta_n^{(j)} - \frac{1}{J-1} \sum_{m=1}^{J} d_n^{(j,m)} \theta_n^{(m)}. \tag{2.4.8}$$

It is apparent that, once again, the algorithm will preserve the linear span of the initial ensemble $\{u_0^{(j)}, \theta_0^{(j)}\}_{j=1}^J$. Furthermore we note that for the centered hierarchical method, since $\mathcal{G}$ does not depend on θ, the algorithm projected onto the u coordinate is identical to that in the preceding subsection, with the only difference being that the initial span of $\{u_0^{(j)}\}_{j=1}^J$ is constructed over a diverse set of θ, reflecting the dependency structure in $\mathbb{P}(u, \theta)$; for hierarchical priors as in subsubsections 2.3.2 and 2.3.2, the dependency structure is typically of the form $\mathbb{P}(u, \theta) = \mathbb{P}(u|\theta)\mathbb{P}(\theta)$. [4]

If we again define

$$d^{(j,m)} = \langle \Gamma^{-1} (\mathcal{G}(u^{(j)}) - y), \mathcal{G}(u^{(m)}) - \bar{\mathcal{G}} \rangle,$$

with obvious definition of $\bar{\mathcal{G}}$, then we obtain the continuous-time limit

$$\dot{u}^{(j)} = -\sum_{m=1}^{J} d^{(j,m)} u^{(m)}, \tag{2.4.9}$$

$$\dot{\theta}^{(j)} = -\sum_{m=1}^{J} d^{(j,m)} \theta^{(m)}. \tag{2.4.10}$$

[4] The centered hierarchical method, where $\mathcal{G}$ does not depend on θ, is presented for simplicity. The details of this algorithm are readily transferred to include θ dependence in the forward mapping $\mathcal{G}$ as required by the version of the Bayesian level set method advocated in [47].

Since $d^{(j,m)}$ depends only on $\{u^{(k)}\}_{k=1}^J$ and not on $\{\theta^{(k)}\}_{k=1}^J$ for the centered hierarchical method the continuous time limit for u is identical to that in the preceding subsection, with the only difference being the creation of the initial ensemble using variable θ. This severely limits the capability of the hierarchical method in the centered case, and is motivation for the non-centered approach that we now describe.

2.4.3 Generalization for Non-Centered Hierarchical Parameterization

One of the reasons for using hierarchical parameterizations is that good choices of parameters such as length-scale and regularity of u are not known a priori; this suggests that they might be learnt from the data. In this context the preservation of the linear span of the initial ensemble is problematic if the algorithm is formulated in terms of (u, θ). This is because, even though the length-scale (for example) may update as the algorithm progresses, the output for u remains in the linear span of the initial set of u, which likely does not contain a good estimate of the true length scale. Instead one can work with the variables (ξ, θ), where ξ is the forcing function in a stochastic PDE, as explained in subsubsection 2.3.2. Working with (u, θ) and with (ξ, θ) are referred to as the *centered* parameterization and the *non-centered* parameterization, respectively. The pros and cons of each method is discussed in the context of Bayesian inversion in [115, 116], where the terminology is also introduced. The provenance of the terminology has no direct relevance in our context, but we retain it to make the link with the existing literature.

The algorithm for updating (ξ, θ) is identical to that shown in subsection 2.4.1 with the identifications $u \mapsto (\xi, \theta)$ and $\mathcal{G} \mapsto \mathcal{G} \circ T$. Note that even though, in the centered case, $\mathcal{G}$ depends only on u, the mapping $\mathcal{G} \circ T$ will depend on both ξ and θ. Hence these variables are coupled through the iteration. Indeed if we now define

$$d_n^{(j,m)} = \langle (C_n^{ww} + \Upsilon_n \Gamma)^{-1} (\mathcal{G} \circ T(\xi_n^{(j)}, \theta_n^{(j)}) - y), \mathcal{G} \circ T(\xi_n^{(m)}, \theta_n^{(m)}) - \overline{\mathcal{G} \circ T}_n \rangle$$

then we see that

$$\xi_{n+1}^{(j)} = \xi_n^{(j)} - \frac{1}{J-1} \sum_{m=1}^{J} d_n^{(j,m)} \xi_n^{(m)}, \tag{2.4.11}$$

and

$$\theta_{n+1}^{(j)} = \theta_n^{(j)} - \frac{1}{J-1} \sum_{m=1}^{J} d_n^{(j,m)} \theta_n^{(m)}. \tag{2.4.12}$$

Although the algorithm will preserve the linear span of the initial ensemble $\{\xi_0^{(j)}, \theta_0^{(j)}\}_{j=1}^J$ the variable of interest $u_n^{(j)} = T(\xi_n^{(j)}, \theta_n^{(j)})$ is not in the linear span of $u_0^{(j)}$, in general. This confers a significant advantage on the non-centered parameterization in comparison with the centered approach.

As in the previous subsections we describe a continuous-time limit, now for the non-centered hierarchical approach. We define

$$d^{(j,m)} = \langle (\Gamma^{-1}(\mathcal{G} \circ T(\xi^{(j)}, \theta^{(j)}) - y), \mathcal{G} \circ T(\xi^{(m)}, \theta^{(m)}) - \overline{\mathcal{G} \circ T} \rangle,$$

again with the obvious definition of $\overline{\mathcal{G} \circ T}$. In the same setting adopted in the previous two subsections we obtain the limiting equations

$$\dot{\xi}^{(j)} = -\sum_{m=1}^{J} d^{(j,m)} \xi^{(m)},$$
$$\dot{\theta}^{(j)} = -\sum_{m=1}^{J} d^{(j,m)} \theta^{(m)}.$$

Now $d^{(j,m)}$ depends on both $\{\xi^{(k)}\}_{k=1}^J$ and on $\{\theta^{(k)}\}_{k=1}^J$ so that the dynamical system not only couples the ensemble members but in general can couple the dynamics for ξ and for θ. This is another way to understand the significant advantage of the non-centered parameterization in comparison with the centered approach.

2.5 Model Problems

In order to demonstrate the benefits of the parameterizations that we introduced here we employ a number of models on which we will base our numerical experiments. This section will be dedicated to describing the various PDEs that will be used. We will describe the forward problem, together with a basic version of the inverse problem, for each model, relevant in the non-hierarchical case. We note that the ideas such as the level set method, and hierarchical formulations from subsubsections 2.3.1 and 2.3.1, can be used to reformulate the inverse problems, and we will use these reformulations in section 3.5.

2.5.1 Model Problem 1

Our first test model is from electrical impedance tomography (EIT). This imaging method is used to learn about interior properties of a medium by injecting current, and measuring voltages, on the boundary [21, 49, 144]. We will use the complete

electrode model (CEM) introduced in [134]. The forward model is as follows: given a domain $D = \mathcal{B}(0,1)^2$ and a set of electrodes $\{e_l\}_{1=1}^{m_e}$ on the boundary ∂D with contact impedance $\{z_l\}_{l=1}^{m_e}$, and interior conductivity κ, the CEM aims to solve for the potential ν inside the domain D and the voltages $\{V_l\}_{1=1}^{m_e}$ on the boundary. The governing equations are

$$\nabla \cdot (\kappa \nabla \nu) = 0, \quad \in \ D \tag{2.5.1a}$$

$$\nu + z_l \kappa \nabla \nu \cdot n = V_l, \ \in \ e_l, \quad l = 1, \ldots, m_e \tag{2.5.1b}$$

$$\nabla \nu \cdot n = 0, \quad \in \ \partial D \backslash \cup_{l=1}^{m_e} e_l \tag{2.5.1c}$$

$$\int \kappa \nabla \nu \cdot n \ ds = I_l, \ \in \ e_l, \quad l = 1, \ldots, m_e, \tag{2.5.1d}$$

with n denoting the outward normal vector on the boundary. The linearity of the problem implies that the relationship between injected current and measured voltages can be described through an Ohm's Law of the form

$$V = R(\kappa) \times I. \tag{2.5.2}$$

In our experiments D will be a two-dimensional disc of radius 1. The inverse problem may now be stated. We write the unknown conductivity as $\kappa = \exp(u)$ and try to infer u from a set of J noisy measurements of voltage/current pairs (V_j, I_j). If we define $\mathcal{G}_j(u) = R(\kappa) \times I_j$ then the inverse problem is to find u from y given an equation of the form (2.2.1).

2.5.2 Model Problem 2

Our second model problem arises in hydrology: the single-phase Darcy flow equations. The concrete instance of the forward problem is as follows: given the domain $D = [0,6]^2$ and real-valued permeability function κ defined on D, the forward model is to the determine real-valued pressure (or hydraulic head) function p on D from

$$-\nabla \cdot (\kappa \nabla p) = f, \quad x \ \in \ D, \tag{2.5.3}$$

with mixed boundary conditions

$$p(x_1, 0) = 100, \qquad \frac{\partial p}{\partial x_1}(6, x_2) = 0, \quad -\kappa \frac{\partial p}{\partial x_1}(0, x_2) = 500, \qquad \frac{\partial p}{\partial x_2}(x_1, 6) = 0, \tag{2.5.4}$$

and the source term f defined as

$$f(x_1, x_2) = \begin{cases} 0, & \text{if } 0 \leq x_2 \leq 4, \\ 137, & \text{if } 4 \leq x_2 \leq 5, \\ 274, & \text{if } 5 \leq x_2 \leq 6. \end{cases}$$

The inverse problem concerned with (2.5.3) is as follows: write $\kappa = \exp(u)$ and determine u from J linear functionals of the pressure $\mathcal{G}_j(u) = l_j(p)$. This may thus be cast in the form (2.2.1). We take the linear functionals as mollified pointwise observations on a regular grid. This specific set-up of the PDE model is that tested by Hanke [54] in his consideration of the regularized Levenberg-Marquardt algorithm. More information on this setting can be found by Carera et al. [29].

2.5.3 Model Problem 3

Our final model is a simple linear inverse problem which we can describe directly. The aim is to reconstruct a function u from noisy observation of J linear functionals $\mathcal{G}_j(u) = l_j(p)$, $j = 1, \cdots, J$, where p solves the equation

$$\frac{d^2p}{dx^2} + p = u, \ \in D, \tag{2.5.5a}$$

$$p = 0, \ \in \partial D. \tag{2.5.5b}$$

This may also be cast in the form (2.2.1). We use equally spaced pointwise evaluations as our linear functionals. We will assume our domain is chosen such that $D = [0, 10]$.

2.6 Numerical Examples

To assess the performance of each parameterization we present a range of numerical experiments on each of the three model problems described in the previous section. Our experiments will be presented in a consistent fashion, between the different models and the different algorithms. Each model problem will be tested using each of the non-hierarchical and hierarchical approaches, although we will not use the centered approach for Model Problem 3. Within each of these approaches we will show the progression of the inverse solver from the first to the last iteration. This will include five images ordered by iteration number, with the first figure displaying the first iteration and the last displaying the final iteration. These figures will be accompanied with figures demonstrating the learning of the hyperparameters, as

the iteration progresses, for Model Problems 1 and 2, but not for Model Problem 3 (where the hyperparameter is a field).

In order to illustrate the effect of the initial ensemble we will show output of the non-centered approach for ten different initializations, for each model problem. We will plot the final iteration reconstruction arising from four of those initializations. We observe the variation across the initializations through the relative errors in the unknown field u_{EKI}, with respect to the truth $u^\dagger$, and in the data misfit, as the iteration progresses:

$$\frac{\|u_{\text{EKI}} - u^\dagger\|_{L^2(D)}}{\|u^\dagger\|_{L^2(D)}},$$
$$\|y - \bar{\mathcal{G}}(u_{\text{EKI}})\|_\Gamma.$$

2.6.1 Level Set Parameterization

Level set methods are a computationally effective way to represent piecewise constant functions, and there has been considerable development and application to inverse problems [23, 25, 77], starting from the paper [128], in which interfaces are part of the unknown. We apply level set techniques, combined with hierarchical parameter estimation, in the context of ensemble inversion; we are motivated by the recent Bayesian level set method developed by Lu et al. in [81], and its hierarchical extensions introduced in [34, 47].

When applying the level set technique to inverse problems of the form (2.2.1), we modify our forward operator to

$$\mathcal{G} = \mathcal{O} \circ G \circ F, \tag{2.6.1}$$

where $G : X \mapsto Y$ maps the coefficient of the PDE to its solution, $\mathcal{O} : Y \to \mathcal{Y}$ is our observational operator, and $F : \mathcal{X} \to X$ is the level set map described by

$$(Fu)(x) \to \kappa(x) = \sum_{i=1}^{n} \kappa_i 1_{D_i}(x). \tag{2.6.2}$$

The sets $\{D_i\}_{i=1}^n$ are n disjoint subdomains with union D and whose boundaries define the interfaces. The boundaries are assumed to be defined through a continuous real-valued function u on D via its level sets. In order to model the level set method hierarchically we will base our reconstructions on the approaches taken in subsection 2.3.2. In general it can be helpful to re-scale the level values as the

hierarchical parameter τ is learned [47]; however if the unknown is binary and the level set taken as zero, as used in our numerical experiments here, then this is not a consideration.

We apply level set inversion to Model Problem 1 (EIT) from subsection 2.5.1. We reconstruct a binary field and the variable u is, rather than the logarithm of the conductivity κ, the level set function defining (2.6.2): specifically the level set formulation is achieved through representing the conductivity as

$$\kappa(x) = (Fu)(x) = \kappa_- \chi_{u \leq 0} + \kappa_+ \chi_{u>0}, \tag{2.6.3}$$

where χ_A denotes a characteristic function of A with κ_- and κ_+ being known positive constants that help define low and high levels of our diffusion coefficient.

We place 16 equidistant electrodes on the boundary of the unit disc D in order to define our observations. All experiments are conducted using the MATLAB package EIDORS [2]. The contact impedances $\{z_l\}_{l=1}^{m_e}$ are chosen with value 0.05 and all electrodes chosen subjected to an input current of 0.1. This provides a matrix of stimulation patterns $I = \{I^{(j)}\}_{j=1}^{15} \in \mathbb{R}^{16\times 15}$ given as

$$I = 0.1 \times \begin{pmatrix} +1 & 0 & \dots & 0 \\ -1 & +1 & \dots & 0 \\ 0 & -1 & \ddots & 0 \\ \vdots & \vdots & \ddots & +1 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

For our iterative method we choose $J = 200$ ensemble members with regularization parameter $\rho = 0.8$. The covariance of our noise η is chosen such that $\Gamma = 10^{-4} \times I$. Our truth for the EIT problem will take the form given in Figure 2.5 where we have high levels of conductivity within the two inclusions. This is constructed by thresholding a Whittle-Matérn Gaussian random field defined by (2.3.11), (2.3.12); true values for the hierarchical parameters used are shown in Table 2.1.

Remark 2.6.1. *We do not display the underlying Gaussian random field u which is thresholded to obtain the true conductivity in Figure 5 as this Gaussian random field cannot be expected to be reconstructed accurately, in general. Furthermore it is important to appreciate that in general a true conductivity will not be constructed by such thresholding; the field u is simply an algorithmic construct. We do however show u, and its evolution, in the algorithm, because this information highlights the roles of the length-scale and regularity parameters.*

When performing inversion we sample initial ensembles using the prior distributions shown in Table 2.2. We have set our prior distributions in such a way that the true value for each hyperparameter lies within the range specified.

Hyperparameter	Value
$\alpha^\dagger$	3
$\tau^\dagger$	10

Table 2.1: Model problem 1. True values for each hyperparameter.

Hyperparameter	Prior
α	$\mathcal{U}[1.3, 4]$
τ	$\mathcal{U}[5, 30]$

Table 2.2: Model problem 1. Prior distribution for each hyperparameter.

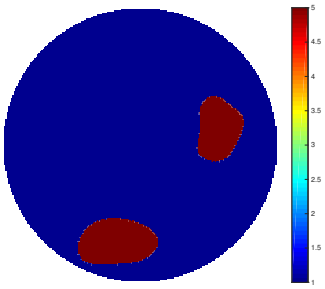

Figure 2.5: Model problem 1: true log-conductivity.

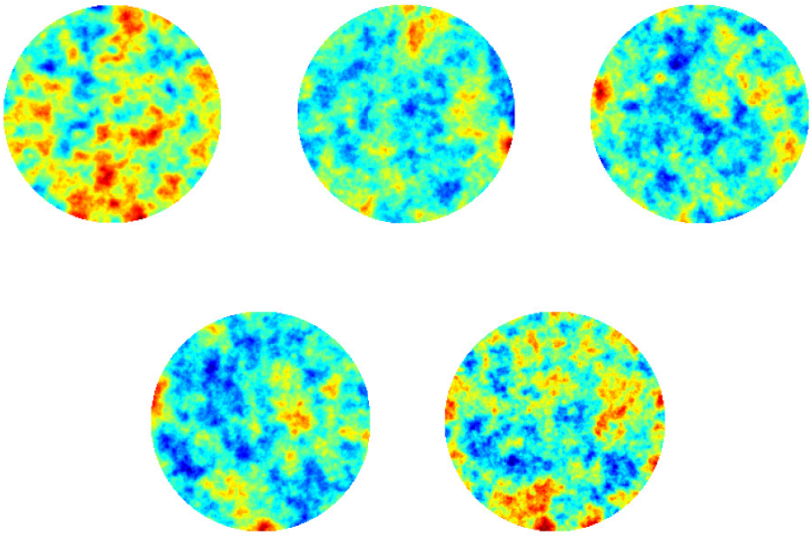

Figure 2.6: Model problem 1. Progression through iterations of non-hierarchical method.

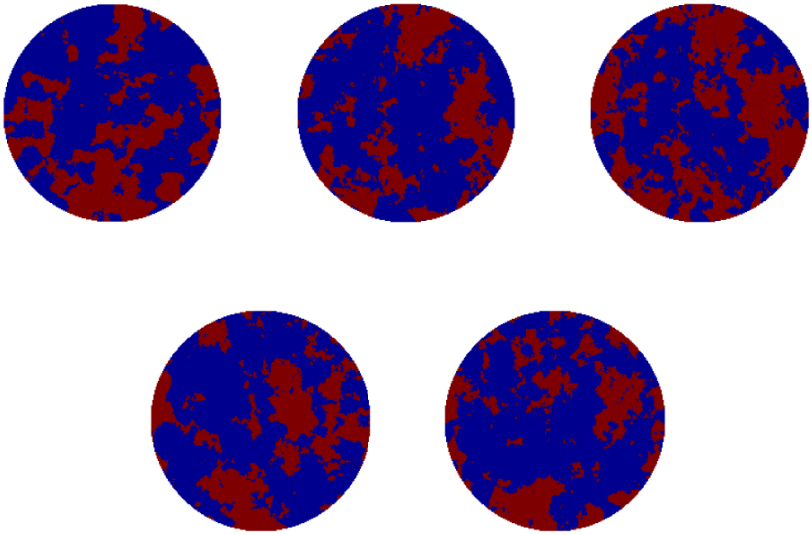

Figure 2.7: Model problem 1. Progression through iterations of non-hierarchical method with level set.

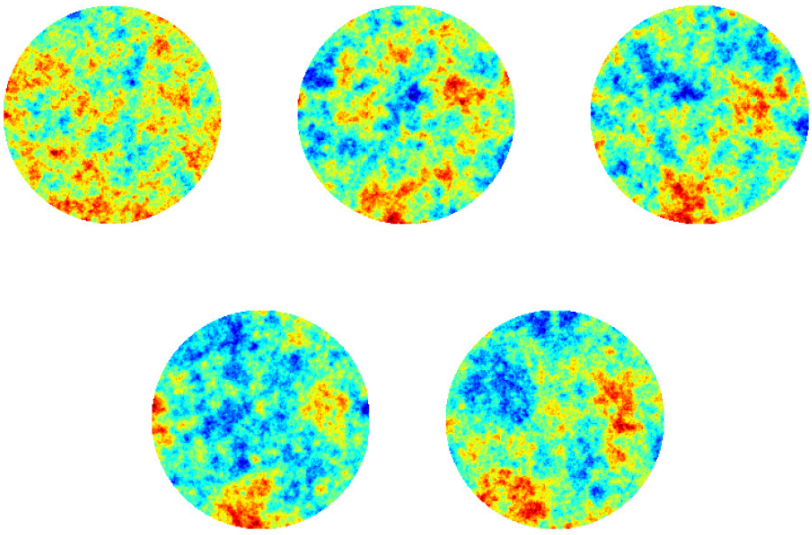

Figure 2.8: Model problem 1. Progression through iterations of centered hierarchical method.

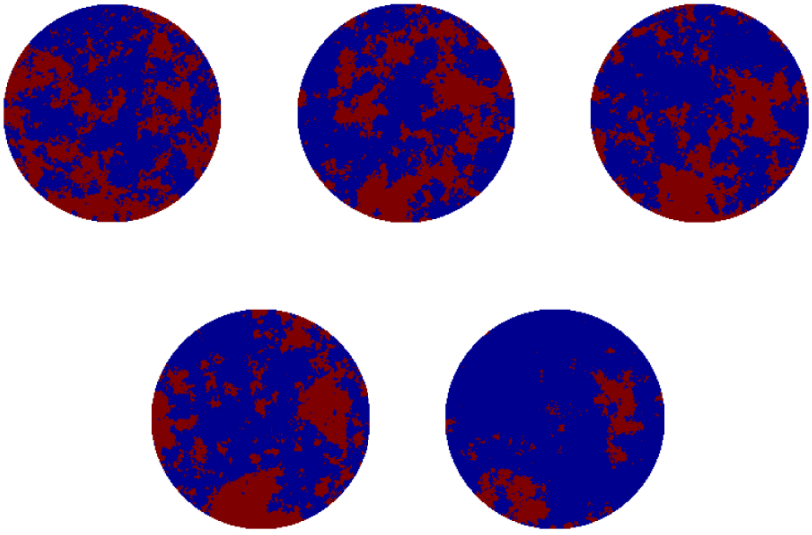

Figure 2.9: Model problem 1. Progression through iterations of centered hierarchical method with level set.

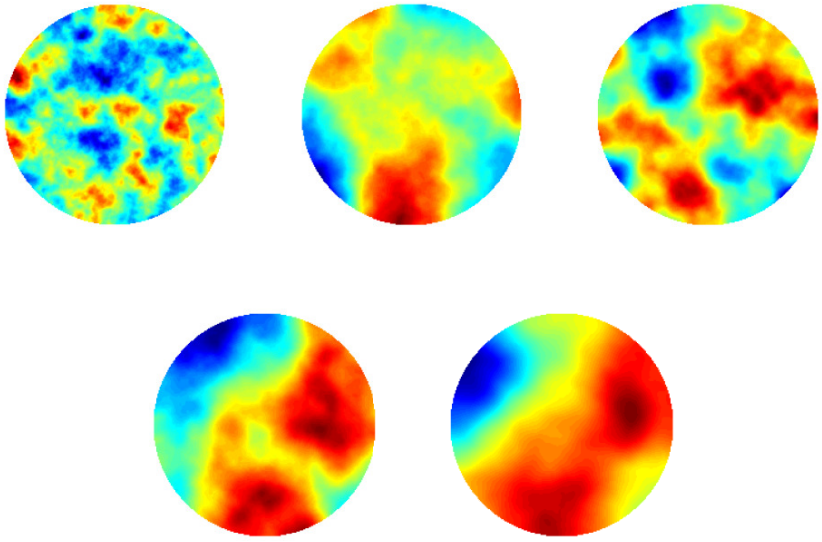

Figure 2.10: Model problem 1. Progression through iterations of non-centered hierarchical method.

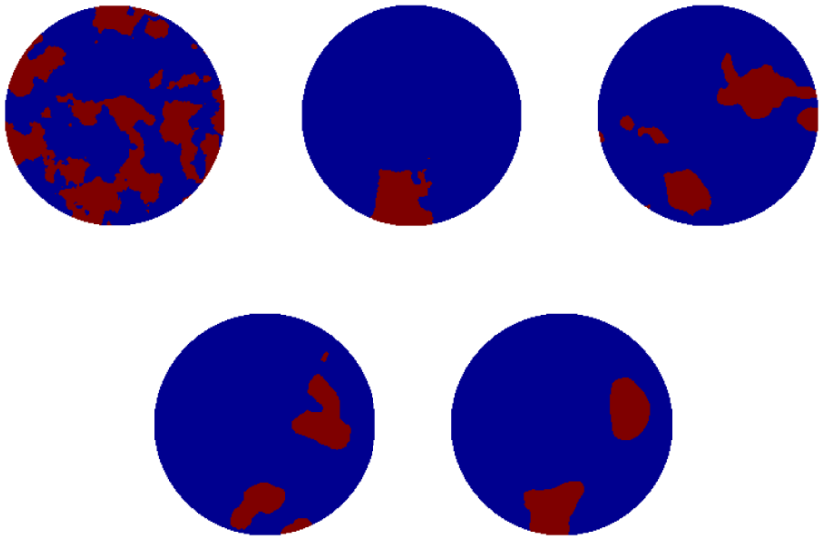

Figure 2.11: Model problem 1. Progression through iterations of non-centered method with level set.

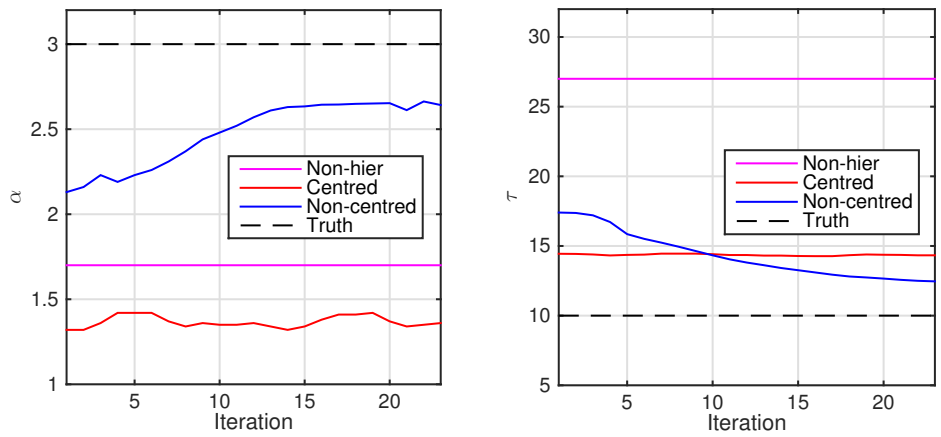

Figure 2.12: Model problem 1. Progression of average value for α and τ.

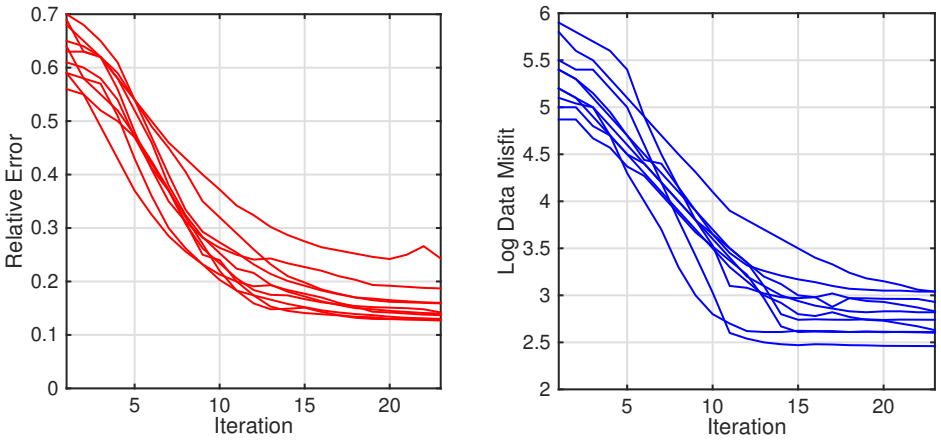

Figure 2.13: Model problem 1. Left: relative error. Right: log-data misfit.

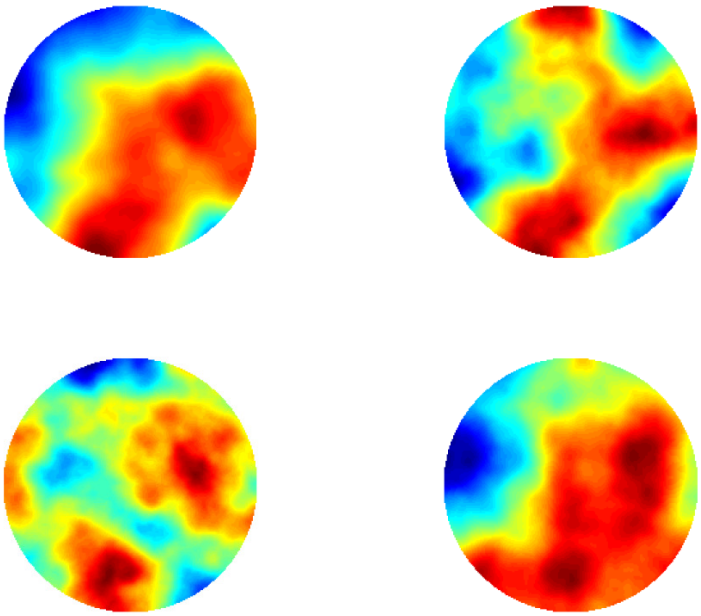

Figure 2.14: Model problem 1. EKI for the final iteration for the non-centered approach with Whittle-Matérn from four different initializations.

Figures 2.6 - 2.11 show the progression of both the level-set function u, and the permeability κ, through five iterations of the method. Non-hierarchical, centered and non-centered methods are considered in turn. By comparing the reconstructions with the true conductivity, these figures clearly demonstrate two facts: (a) that being hierarchical is necessary to obtain a good reconstruction; (b) that implementing the hierarchical method using a non-centered parameterization has significant benefits when compared to the centered method. These points are further demonstrated in Figures 2.12 which shows the learning of the hyperparameters, in comparison with the truth, for the three methods.

Concentrating solely on the non-centered approach, we run ten different initializations of the EKI. We display the resulting data-misfit and error, as a function of iteration, for all ten in Figure 2.13. We display the last iteration of four of these ten in Figures 2.14 - 2.15.

In summary, Figures 2.10 - 2.12 clearly show the superiority of the non-centered hierarchical method. For all the initializations shown, the EKI produces conductivities which concentrate near to the true conductivity and have length-scale similar to those appearing in the truth; see Figures 2.13 - 2.15. The centered hierar-

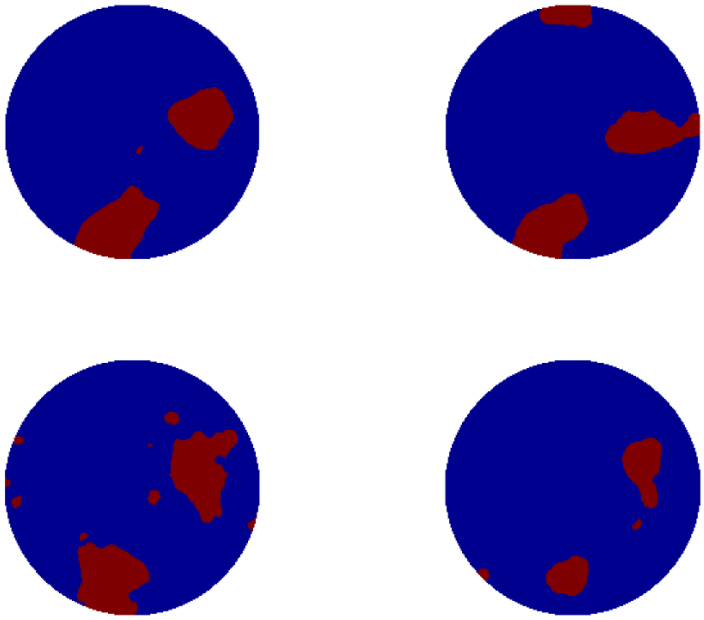

Figure 2.15: Model problem 1. EKI for the final iteration for the non-centered method with level set from four different initializations.

chical and non-hierarchical methods fail to do this; see Figures 2.6 - 2.9. However it is important to note the non-centered method does produce substantial variation in the predicted solution, depending on which initialization is used as shown in Figures 2.13 - 2.15.

2.6.2 Geometric Parameterization

In this subsection we employ Model Problem 2 from subsection 2.5.2. We consider reconstruction of a piecewise continuous channel which is defined through two heterogeneous Gaussian random fields, scalar geometric parameters specifying the geometry and scalar hierarchical parameters characterizing the length-scale and regularity of the two fields.

The truth $u^\dagger$ is shown in Figure 2.16. It is drawn from a prior distribution which we now describe; details may be found in [80]. The channel is described by five parameters: d_1 – amplitude; d_2 – frequency; d_3 – angle; d_4 – initial point; and d_5 – width. We generate two Gaussian random fields $\{\kappa_i\}_{i=1}^2$, both defined on the whole of the domain D but entering the permeability κ only inside and outside

(respectively) the channel. The unknown u thus comprises the five scalars $\{d_i\}_{i=1}^5$ and the two fields $\{\kappa_i\}_{i=1}^2$. We do not explicitly spell out the mapping from u to the coefficient κ appearing in the Darcy flow, but leave this to the reader. The $\{\kappa_i\}_{i=1}^2$ are specified as log-normal random fields and the underlying Gaussians are of Whittle-Matérn type, defined by (2.3.11), (2.3.12); different uniform distributions on α and τ are used for the two fields $\{\kappa_i\}_{i=1}^2$ The parameters $\{d_i\}_{i=1}^5$ are also given uniform distributions. The entire specification of the prior is given in Table 2.4. For the truth the true hierarchical parameters are provided in Table 2.3.

In our inversion we employ 64 mollified pointwise observations $\{l_i(p)\}_{i=1}^{64}$ given by, for some $\sigma > 0$,

$$l_t(p) = \int_D \frac{1}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}(x-x_t)^2} p(x)dx, \tag{2.6.4}$$

where the x_i are uniformly distributed points on D. We discretize the forward model using a second order centered finite difference method with mesh spacing 10^{-2}. For our EKI method we use the same values for our parameters as in subsection 2.6.1.

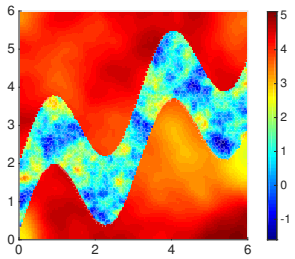

Figure 2.16: Model problem 2. True log-permeability.

Parameter	Value
$\alpha_1^\dagger$	2
$\alpha_2^\dagger$	2.8
$\tau_1^\dagger$	30
$\tau_2^\dagger$	10

Table 2.3: Model problem 2. Parameter selection of the truth.

Parameter	Prior
d_1	$\mathcal{U}[0,1]$
d_2	$\mathcal{U}[2,13]$
d_3	$\mathcal{U}[0.4,1]$
d_4	$\mathcal{U}[0,1]$
d_5	$\mathcal{U}[0.1,0.3]$
κ_1	$N(1,(I-\tau_1^2\Delta)^{-\alpha_1})$
κ_2	$N(4,(I-\tau_2^2\Delta)^{-\alpha_2})$
α_1	$\mathcal{U}[1.3,3]$
τ_1	$\mathcal{U}[8,30]$
α_2	$\mathcal{U}[1.3,3]$
τ_2	$\mathcal{U}[8,30]$

Table 2.4: Model problem 2. Prior associated with channelised flow.

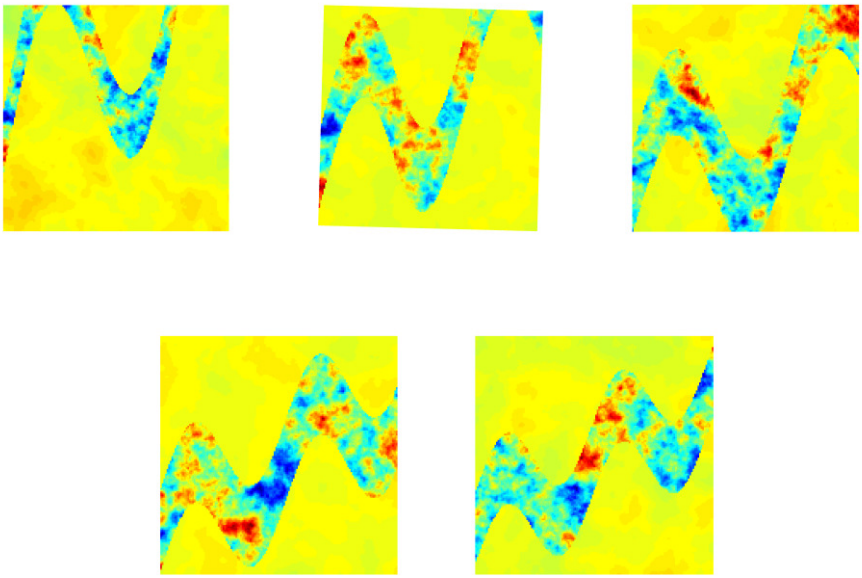

Figure 2.17: Model problem 2. Progression through iterations of non-hierarchical method.

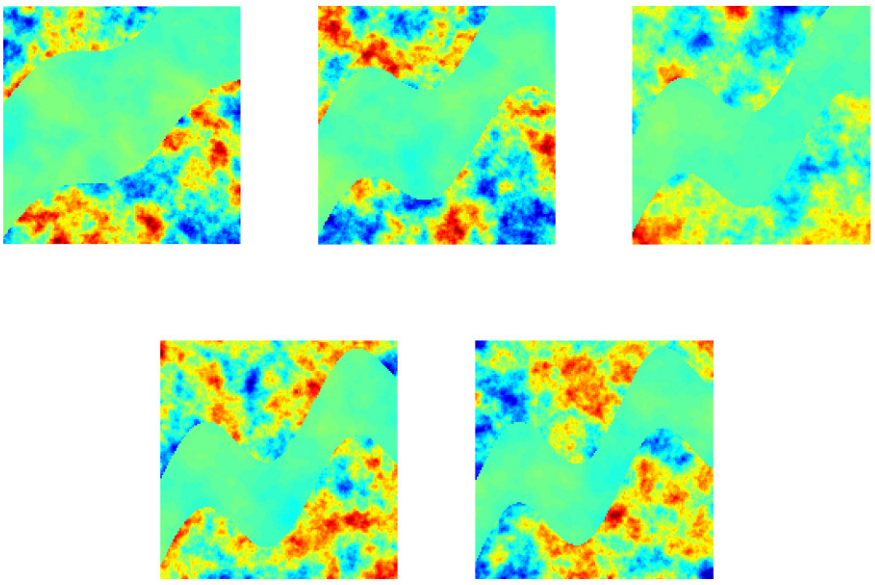

Figure 2.18: Model problem 2. Progression through iterations of centered hierarchical method.

Figures 2.17 - 2.19 are consistent with the previous subsection in that we notice that hierarchical methods are needed and that non-centring is necessary to make hierarchical methods perform well. We qualify this by noting that the geometry is well-learned in all cases, but that the reconstructions of the random fields u_1 and u_2 inside and outside the geometry are sensitive to needing non-centered hierarchical representation. Even then the reconstruction is only accurate in terms of amplitude and length-scales and not pointwise.

Figures 2.20 and 2.21 give further insight into this, showing how the smoothness and length-scale parameters are learned differently in the hierarchical, centered and non-centered methods. Again the conclusions are consistent with the previous subsection. Figures 2.22 and 2.23 concentrate on the application of the non-centered approach, using ten different initializations. The data misfit and relative error in the field are shown for all ten cases in Figure 2.22; four solution estimates are displayed in Figure 2.23. The results are similar to those in the previous subsection. However there is more variability across initializations. This might be ameliorated by use of a larger ensemble size.

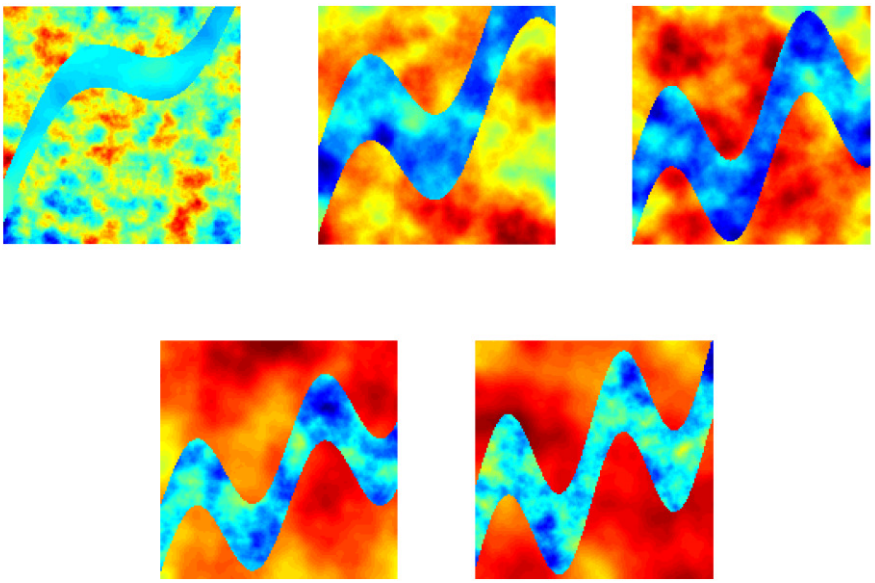

Figure 2.19: Model problem 2. Progression through iterations of non-centered hierarchical method.

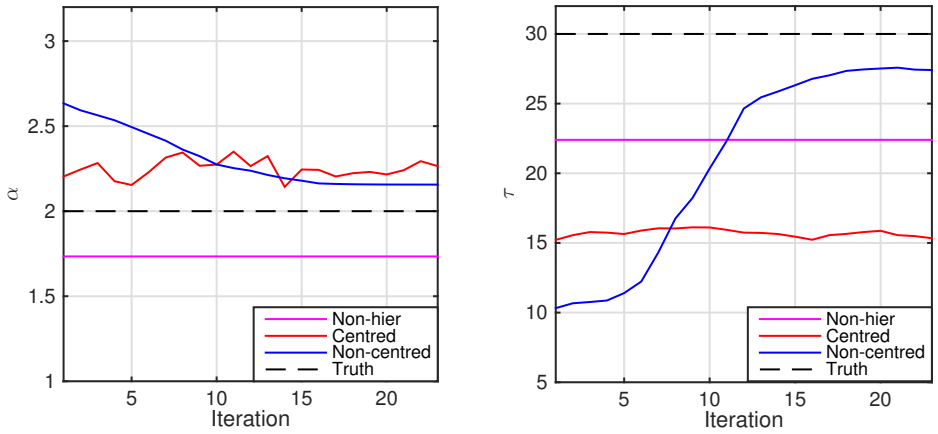

Figure 2.20: Model problem 2. Progression of average value for α_1 and τ_1.

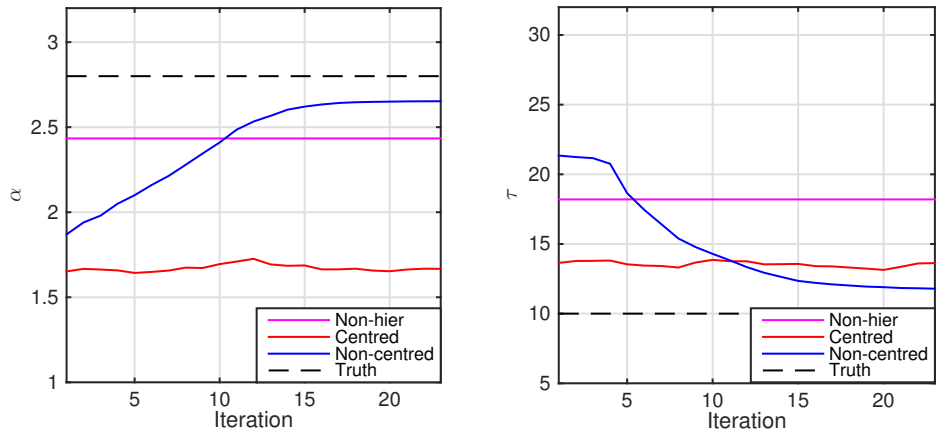

Figure 2.21: Model problem 2. Progression of average value for α_2 and τ_2.

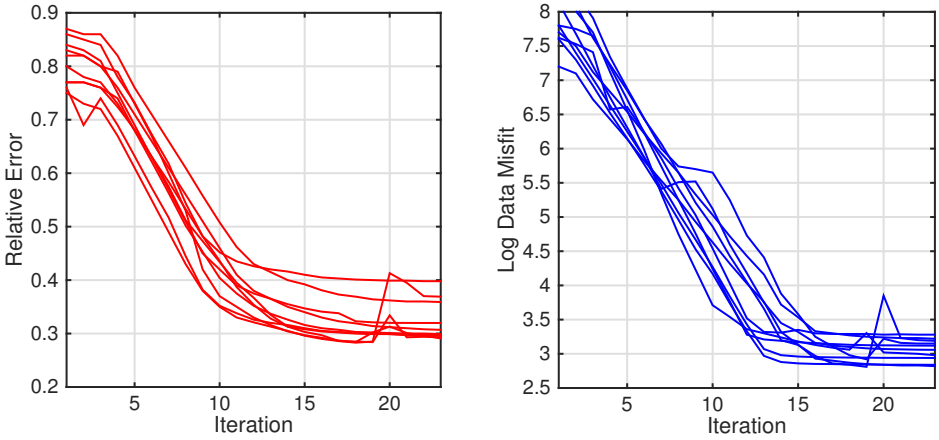

Figure 2.22: Model problem 2. Left: relative error. Right: log-data misfit.

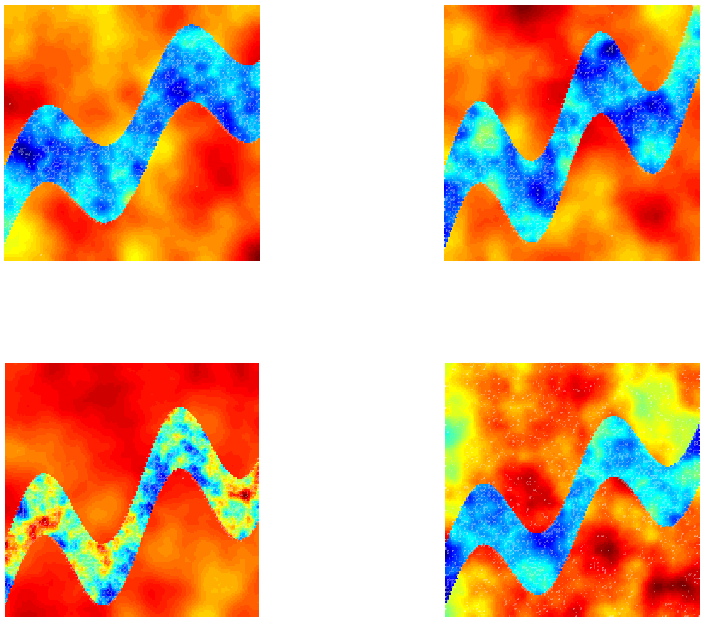

Figure 2.23: Model problem 2. EKI for the final iteration for the non-centered approach from four different initializations.

2.6.3 Function-valued Hierarchical Parameterization

Our final set of experiments will be based on hierarchical inversion of non-stationary random fields, using Model Problem 3. We consider reconstruction of truths that are not drawn from the prior; the prior will be a hierarchical Gaussian model with spatially varying inverse length-scale as hyperparameter. Examples of such truths are ones which contain both rough and smooth features. We discretize the forward model using a piecewise-linear finite element method (FEM), with a mesh of $h = 1/100$. The truth we aim to recover is given by

$$u^\dagger(x) = \begin{cases} \exp\left(4 - \frac{25}{x(5-x)}\right), & x \in (0,5) \\ 1, & x \in [7,8] \\ -1, & x \in (8,9] \\ 0, & \text{otherwise}, \end{cases} \tag{2.6.5}$$

which incorporates both rough and smooth features. Our parameters for the iterative methods are identical to those used in subsection 2.6.1. Because the results of subsection 2.6.1 and 2.6.2 clearly demonstrate the need for non-centring in hierarchical methods, we do not include results for the centered hierarchical approach here; we compare non-hierarchical methods with the use of non-centered hierarchical methods with both Cauchy and Gaussian random fields as priors on the hyperparameter v.

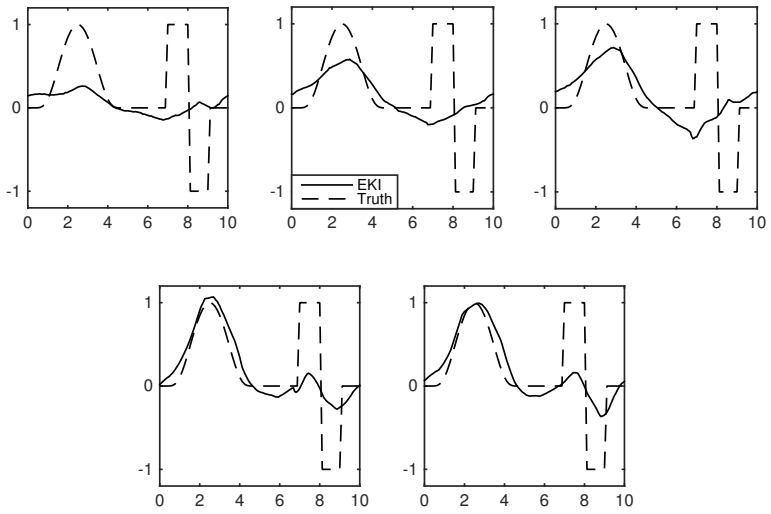

Figure 2.24: Model problem 3. Progression through iterations of non-hierarchical method.

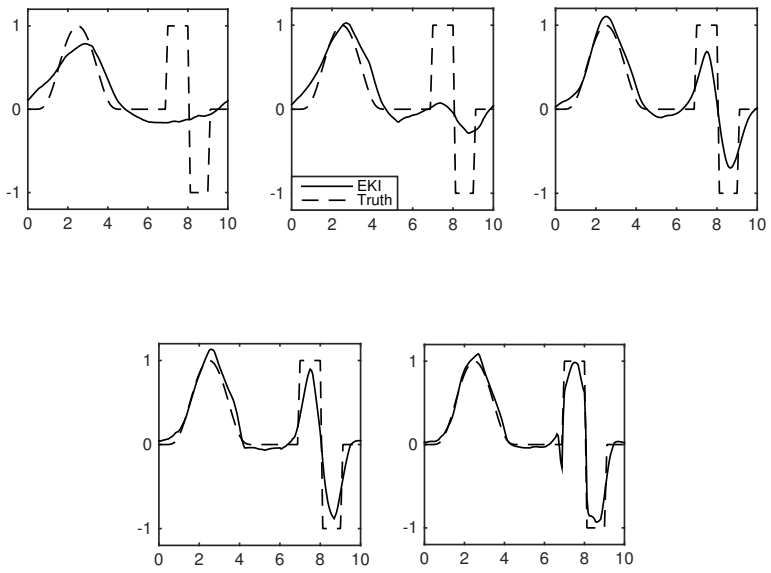

Figure 2.25: Model problem 3. Progression through iterations with hierarchical Gaussian random field.

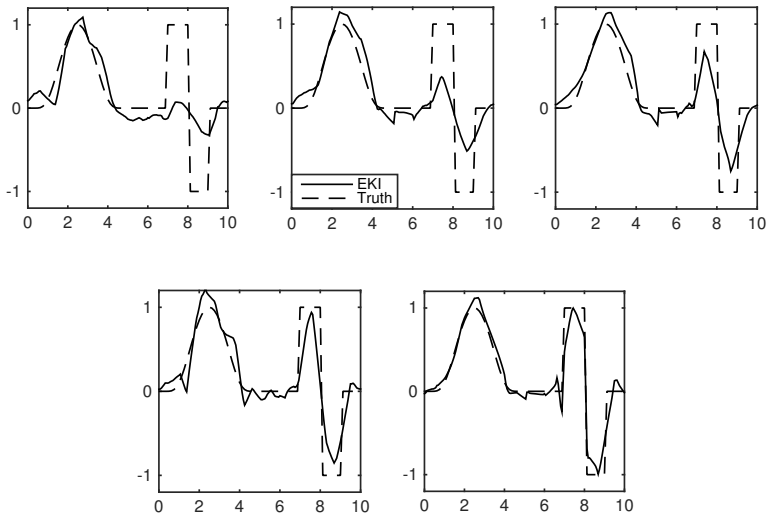

Figure 2.26: Model problem 3. Progression through iterations with hierarchical Cauchy random field.

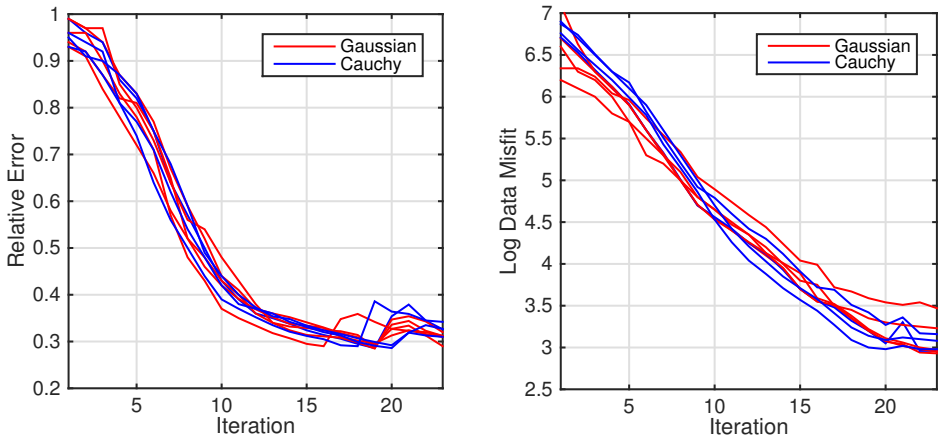

Figure 2.27: Model problem 3. Left: relative error. Right: log-data misfit.

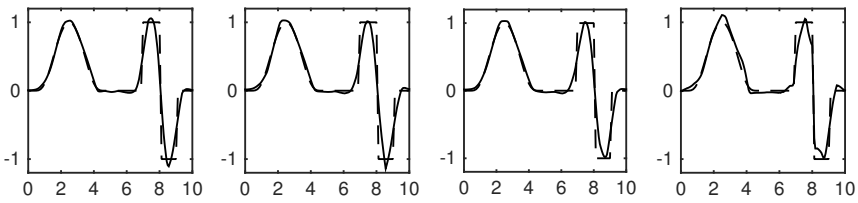

Figure 2.28: Model problem 3. EKI for the final iteration for the Gaussian hierarchical method from four different initializations.

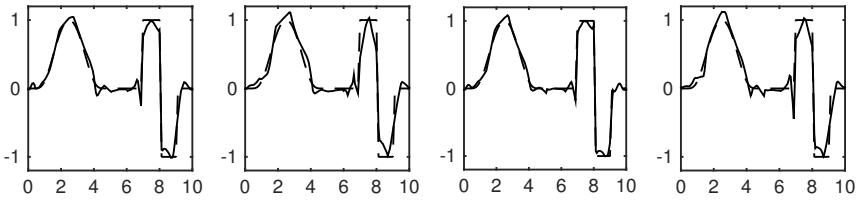

Figure 2.29: Model problem 3. EKI for the final iteration for the Cauchy hierarchical method from four different initializations.

For our Cauchy density function (2.3.15), we set $a = 4$ and $b = d = 0$. Our length-scale for the non-hierarchical method will be based on a Gaussian random field (2.3.14), similarly with the non-centered Gaussian approach. Our comparison of each approach is provided in Figures 2.24 - 2.26. We notice that the non-hierarchical method struggles to reconstruct the truth, and in particular the piecewise constant part of it with discontinuities. In contrast both non-centered approaches perform well. The effectiveness of the non-centered approaches are highlighted in Figures 2.27 - 2.29. The first shows the data-misfit and relative error over ten realizations, and

the second shows four reconstructions chosen from these ten at random. Substantial robustness to the choice of realization is clear. Note however that for a multi-dimensional problem inversion the variability with respect to initialization may be more significant, as in the previous two subsections.

2.7 Conclusion & Discussion

In this chapter we have considered several forms of parameterizations for EKI. In particular our main contribution has been to highlight the potential for the use of hierarchical techniques from computational statistics, and the use of geometric parameterizations, such as the level set method. Our perspective on EKI is that it forms a derivative free optimizer and we do not evaluate it from the perspective of uncertainty quantification. However the hierarchical and level set ideas are motivated by Bayesian formulations of inverse problems. We have shown that our parameterizations do indeed lead to better reconstructions of the truth on a variety of model problems including groundwater flow, EIT and source inversion. There is very little analysis of EKI, especially in the fixed, small, ensemble size setting where it is most powerful. Existing work in this direction may be found in [19, **?**]; it would be of interest to extend these analyses to the parameterizations introduced here. Furthermore, from a practical perspective, it would be interesting to extend the deployment of the methods introduced here to the study of further applications.

Chapter 3

Analysis of hierarchical ensemble Kalman inversion

3.1 Overview

The motivation in this chapter is to generalize the scaling results derived in Chapter 2. We do so by considering both the noisy and noise-free limits while giving a better overview of both approaches. As much of this chapter will follow from its predeceasing one, we state a major difference in this chapter, which is to introduce certain variants of the EnKF hierarchically and to define their continuous-time limit. A numerical example will highlight the importance of this through a discretization of the limit on a 1D elliptic PDE.

3.2 Introduction

The ensemble Kalman filter (EnKF) [57, 59] was proposed by Evensen in 1994 as a Monte-Carlo approximation of the Kalman filter (KF). Its motivation was based on mitigating the computational challenges associated with the KF, replacing the updated mean and covariances with an ensemble of particles. Since then the EnKF has been widely applied in numerous fields such as weather prediction and oceanography [4, 58, 97]. Given its robustness and Bayesian formulation paradigm, the EnKF has been further applied to inverse problems. Inverse problems are concerned with the recovery of some quantity of interest $u \in \mathcal{X}$ from noisy measurements $y \in \mathcal{Y}$ given by

$$y = \mathcal{G}(u) + \eta, \quad \eta \sim \mathcal{N}(0, \Gamma). \tag{3.2.1}$$

By allowing for a Bayesian approach one is interested in constructing, via an application of an infinite-dimensional Bayes' Theorem [137], a posterior measure of the random variable $u|y$

$$\mu(du) = \frac{1}{Z}\exp(-\Phi(u;y))\mu_0(du),$$

with normalizing constant

$$Z := \int_{\mathcal{X}} \exp(-\Phi(u;y))\mu_0(du),$$

such that our data-likelihood is in the form of a potential

$$\Phi(u;y) = \frac{1}{2}\|\Gamma^{-1/2}(y - \mathcal{G}(u))\|^2,$$

with the addition of a prior measure μ_0. This has been recently studied where there have been advancements in both computational and theoretical understanding [19, 77, 78, 129]. From the computational aspect the EnKF was derived as a derivative-free inverse solver, which can be thought of as an optimizer which uses techniques from the Levenberg-Marquardt (LM) scheme [71] combined with elements of the EnKF. It has been shown that applying these regularization techniques from LM [77, 71] can improve the performance of the method. Regarding the theory of the EnKF for inverse problems, there has been progress on gaining analytical insight such as approximating continuous-time limits [19, 129] within the context of inverse problems. A new direction in this field which has emerged is the incorporation of hierarchical approaches for inverse problems [6, 47, 98, 123, 124]. In hierarchical inverse problems we are interested in recovering our unknown and a corresponding hyperparameter $\theta \in \mathbb{R}^+$ that defines the unknown i.e. we wish to recover an unknown $(u, \theta) \in \mathcal{X} \times \mathbb{R}^+$ from noisy measurements y where

$$y = \mathcal{G}(u, \theta) + \eta.$$

This allows for richer reconstructions as more information about the underlying unknown is available. An important feature of the EnKF applied to inverse problems is that it produces an ensemble of particles which lies within the linear span of the initial ensemble. This effect is known as the "subspace property". By incorporating various hierarchical approaches we look to break this subspace property. This allows the solution to learn from information which may not be given within the span, but instead the data. Specifically for EnKF inversion a hierarchical methodology was proposed in [33] which demonstrated improvements over its non-hierarchical

counterpart. The newly proposed method provides a way to effectively learn both the unknown and its hyperparameters that define it. This work used ideas from hierarchical computational statistics and applied it in an inverse problem setting [115, 116].

However regarding analytical results there has been no development in understanding these hierarchical approaches for the EnKF. This can be related to the lack of analysis on the EnKF. As of yet there has been work done on estimating non-hierarchical continuous-time limits [64]. The purpose of this chapter is to build some analytical insight for hierarchical approaches that were used in [33] for Bayesian inverse problems. It was shown in the linear noise-free case that one can attain a preconditioned gradient flow structure. Much of this will be based on extending the current theory in a hierarchical manner to the nonlinear noisy case, while providing an overview of the limit results attained in [33]. We also aim to understand these approaches with modified versions of the EnKF, namely localization [70] and covariance inflation [5]. Both these techniques were developed to improve errors based on a small ensemble size. Similarly with some of the hierarchical approaches, localization and covariance inflation have the ability to break the subspace property. As a result it would be of interest to understand the limiting behaviour of these techniques. This includes conducting numerical experiments to verify hierarchical results obtained. We emphasize that with this chapter, rather than deriving new results for the EnKF, we aim to shed some light on hierarchical EnKF approaches for inverse problems and their respective continuous-time limits.

3.2.1 Structure

The layout of this chapter is as follows; in Section 3.3 we provide an overview of the EnKF applied to inverse problems. This will lead onto the formal derivation of the continuous-time limits applied to inverse problems. In Section 3.4 we give a brief introduction for hierarchical approaches to EnKF inversion, while in Section 3.5 we derive and present continuous-time limits for a list of variants on the EnKF. We verify these results through means of numerics in Section 3.6. Finally in Section 3.7 we summarize our results and provide a brief mention on future work to consider.

3.2.2 Notation

We assume that $(\mathcal{X}, \|\cdot\|, \langle\cdot\rangle)$ and $(\mathcal{Y}, \|\cdot\|, \langle\cdot\rangle)$ are two separable Hilbert spaces which are linked through the forward operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$. The operator can be thought of as mapping from the space of parameters $\mathcal{X}$ to the observation space $\mathcal{Y}$. We denote

the space of our hyperparameters as $\theta = (\sigma, \alpha, \ell) \in \mathbb{H}$ where $\mathbb{H} := \mathbb{R} \times \mathbb{R}^+ \times \mathbb{R}^+$. For any such operator we define $\langle \cdot, \cdot \rangle_\Gamma = \langle \Gamma^{-1/2} \cdot, \Gamma^{-1/2} \cdot \rangle$ and $\| \cdot \|_\Gamma = \| \Gamma^{-1/2} \cdot \|$, while for finite dimensions $| \cdot |_\Gamma = | \Gamma^{-1/2} \cdot |$ with $| \cdot |$ denoting Euclidean norm. $u_n^{(j)}$ will denote an ensemble of particles where n is the iteration count and $j \in \{1, \ldots, J\}$ is the j^{th} ensemble member.

3.3 EnKF for Inverse Problems

The iterative EnKF method was first proposed in [79] to tackle Bayesian inverse problems in a partial differential equation (PDE)-constrained framework. The method can be derived as a sequential Monte-Carlo (SMC) approximation, where our probability measures of interest μ_n are defined by, for $h = N^{-1}$,

$$\mu_n(du) \propto \exp(-nh\Phi(u; y))\mu_0(du),$$

thus leading to

$$\mu_{n+1}(du) = \frac{1}{Z_n} \exp(-h\Phi(u; y))\mu_n(du),$$

where

$$Z_n := \int_{\mathcal{X}} \exp(-h\Phi(u; y))\mu_n(du).$$

We can construct our update for our probability measures μ_{n+1} through the operation

$$\mu_{n+1} = L_n \mu_n, \tag{3.3.1}$$

where L_n can be treated as a non-linear operator from μ_n to μ_{n+1} via an application of Bayes' Theorem. The idea behind the formulation of (3.3.1) is that it can be viewed as an artificial discrete-time dynamical system mapping the prior measure μ_0 to the posterior measure μ_n. Recall that with SMC methods one is interested in approximating a sequence of particles and weights which take the form

$$\mu_n \simeq \sum_{j=1}^{J} w_n^{(j)} \delta_{u_n^{(j)}}, \quad j \in \{1, \ldots, J\},$$

with $\delta_{u_n^{(j)}}$ denoting the delta-Dirac mass at $u_n^{(j)}$. The weights associated with our sequence of particles satisfy the condition

$$\sum_{j=1}^{J} w_n^{(j)} = 1.$$

The SMC approach poses computational advantages over other Monte-Carlo methods, but still has limitations within it. These arise when the weights $\{w_n^{(j)}\}_{j=1}^J$ become degenerate i.e. that one of the weights becomes close to one where the rest are negligible [8]. The EnKF poses an improvement on this as its approximation has the form

$$\mu_n \simeq \frac{1}{J}\sum_{j=1}^{J}\delta_{u_n}^{(j)},$$

which excludes the weights. The EnKF for inverse problems, similarly to the EnKF, can be into two steps: a *prediction step* and an *update step*. The prediction step can be interpreted as mapping an ensemble of particles $u_n^{(j)}$ into the data space where we define our sample means for J ensemble members

$$\bar{u} = \frac{1}{J}\sum_{j=1}^{J} u_n^{(j)},$$
$$\bar{\mathcal{G}} = \frac{1}{J}\sum_{j=1}^{J} \mathcal{G}(u_n^{(j)}),$$

and our empirical covariances

$$C_n^{uw} = \frac{1}{J-1}\sum_{k=1}^{J}(u^{(k)} - \bar{u}) \otimes (\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}) \tag{3.3.2}$$

$$C_n^{ww} = \frac{1}{J-1}\sum_{k=1}^{J}(\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}) \otimes (\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}). \tag{3.3.3}$$

The update step matches the mapped ensemble of particles to the data $y_{n+1}^{(j)}$ by using the calculated mean and covariances through the update formula

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}\big(C_n^{ww} + h\Gamma\big)^{-1}\big(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})\big), \tag{3.3.4}$$

where

$$y_{n+1}^{(j)} = y + \iota_{n+1}^{(j)}, \quad \iota_{n+1}^{(j)} \sim \mathcal{N}(0, h^{-1}\Gamma). \tag{3.3.5}$$

The EnKF for inverse problems possesses an important characteristic known as the *subspace property* [78, 93]. The property was first discussed [93] which states that the updated ensemble of particles $u_{n+1}^{(j)}$ is preserved by the linear span of the initial ensemble $\mathcal{A} := \text{span}\{u_0^{(j)}\}$ for $j \in \{0, \dots, J\}$. In the context of Gaussian priors, in the discrete case, this was proved in the following theorem.

Theorem 3.3.1. *For every* $(n, j) \in \mathbb{N} \times \{1, \dots, J\}$ *we have* $u_{n+1}^{(j)} \in \mathcal{A}$ *and hence* $u_{n+1} \in \mathcal{A}$.

Proof. The proof can be found in [79] by Iglesias et al.. □

The property can be interpreted as given an initial ensemble, with particular set features depending how it is chosen, our solution to the inverse problem (3.2.1) will remain in the form that it is chosen initially. This can be advantageous if we know that the underlying unknown u is of a similar form to the initial ensemble, where the converse of this is that it poses a limitation if they differ significantly.

3.3.1 Continuous-Time Limit

Nonlinear Noisy Case

The continuous-time limit of the EnKF applied to inverse problems was considered in the work of Schillings et al. [129]. We briefly recall the limit analysis here, firstly by considering the nonlinear noisy case. The limit here arises by taking the parameter $h \to 0$. We define $u_n = \{u_n^{(j)}\}_{j=1}^J$ and assume that $u_n \approx u(nh)$. Our update step (3.3.4) can now be written in the form

$$\begin{aligned} u_{n+1}^{(j)} &= u_n^{(j)} + hC_n^{uw}(u_n)\big(hC_n^{ww}(u_n) + \Gamma\big)^{-1}\big(y - \mathcal{G}(u_n^{(j)})\big) \\ &\quad + hC_n^{uw}(u_n)\big(hC_n^{ww}(u_n) + \Gamma\big)^{-1}\iota_{n+1}^{(j)} \\ &= u_n^{(j)} + hC_n^{uw}(u_n)\big(hC_n^{ww}(u_n) + \Gamma\big)^{-1}\big(y - \mathcal{G}(u_n^{(j)})\big) \\ &\quad + h^{\frac{1}{2}}C_n^{uw}(u_n)\big(hC_n^{ww}(u_n) + \Gamma\big)^{-1}\sqrt{\Gamma}\zeta_{n+1}^{(j)}, \end{aligned}$$

where $\zeta_{n+1}^{(j)} \sim \mathcal{N}(0, I)$. By taking the limit $h \to 0$, our limit can be viewed as a tamed Euler-Maruyama type discretization of the stochastic differential equations (SDEs)

$$\frac{du^{(j)}}{dt} = C^{uw}(u)\Gamma^{-1}\big(y - \mathcal{G}(u^{(j)})\big) + C^{uw}(u)\sqrt{\Gamma^{-1}}\frac{dW^{(j)}}{dt},$$

with $W^{(j)}$ denoting independent cylindrical Brownian motions. By substituting the form of the covariance operator (3.3.2) we see

$$\frac{du^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle \mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}, y - \mathcal{G}(u^{(j)}) + \sqrt{\Gamma}\frac{dW^{(j)}}{dt}\Big\rangle_{\Gamma}(u^{(k)} - \bar{u}). \tag{3.3.6}$$

This derivation of the limit satisfies a generalization of the subspace property in continuous-time provided there is a solution, as the vector field is in the linear span of the ensemble. As we have just analyzed the limit in the noisy-case we will now turn our attention towards the linear noise-free case.

Linear noise-free case

For this we take our forward operator $\mathcal{G}(\cdot) = A\cdot$ to be bounded and linear. Using this notion and by substituting our linear operator A in (3.3.6) we have the following diffusion limit

$$\frac{du^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\left\langle A(u^{(k)} - \bar{u}), y - Au^{(j)}\right\rangle_{\Gamma}(u^{(k)} - \bar{u}). \tag{3.3.7}$$

By defining the empirical covariance operator

$$C(u) = \frac{1}{J-1}\sum_{k=1}^{J}(u^{(k)} - \bar{u}) \otimes (u^{(k)} - \bar{u}),$$

and taking $\Gamma = 0$ we can express (3.3.7) as

$$\frac{du^{(j)}}{dt} = -C(u)D_u\Phi(u^{(j)}; y), \tag{3.3.8}$$

with

$$\Phi(u; y) = \frac{1}{2}\|\Gamma^{-1/2}(y - Au)\|^2.$$

Thus we note that each particle performs a preconditioned gradient descent for $\Phi(\cdot; y)$ where all the gradient descents are preconditioned through the covariance $C(u)$. Since our covariance operator $C(u)$ is semi-positive definite we have that

$$\frac{d}{dt}\Phi(u(t); y) = \frac{d}{dt}\frac{1}{2}\|\Gamma^{-1/2}(y - Au)\|^2 \leq 0,$$

which provides a bound on $\|Au(t)\|_{\Gamma}$. In this case it was shown, through Theorem 2. in [129], that the gradient flow structure provides the existence of a solution satisfying the subspace property.

3.4 Hierarchical Ensemble Kalman Inversion

In order to derive continuous-time limits we first recall a few properties of the hierarchical ensemble Kalman inversion (EKI). This will include newly defined update

equations where we consider both the centred and non-centred approaches towards generating our prior measure μ_0. Our prior $\mu_0 \sim \mathcal{N}(0, \mathcal{C})$ will be assumed to be of a Gaussian form with a Whittle-Matérn covariance function

$$c(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{|x - x'|}{\ell} \right)^{\nu} K_\nu \left(\frac{|x - x'|}{\ell} \right), \qquad x, x' \in \mathbb{R}^d, \tag{3.4.1}$$

where K_ν denotes a modified Bessel function of the second kind and $\Gamma(\nu)$ is a Gamma function. From (3.4.1) we also have the inclusion of three hyperparameters; the amplitude $\sigma \in \mathbb{R}$, the regularity $\nu = \alpha + d/2 \in \mathbb{R}^+$ and the length-scale $\ell \in \mathbb{R}^+$. We can explicitly represent this covariance function through the following stochastic partial differential equation (SPDE), which is derived in [124],

$$(I - \ell^2 \triangle)^{\frac{\alpha}{2}} u = \ell^{d/2} \sqrt{\beta} \xi, \tag{3.4.2}$$

where $\xi \in H^{-s}(D)$, $s > \frac{d}{2}$, for $D \subset \mathbb{R}^d$ is Gaussian white noise and

$$\beta = \sigma^2 \frac{2^d \pi^{d/2} \Gamma(\alpha)}{\Gamma(\alpha - \frac{d}{2})}.$$

Taking the SPDE defined above with $\beta \equiv 1$ we can rewrite (5.4.8) as

$$\mathcal{C}_\theta^{-\frac{1}{2}} u = \xi, \tag{3.4.3}$$

where $\theta = (\sigma, \alpha, \ell) \in \mathbb{H}$ denotes the collection of hyperparameters. We specifically choose the value of $\beta \equiv 1$ for simplicity. The SPDE (3.4.2) is a common way of representing and expressing Gaussian random fields. This approach introduced by Lindgren et al. [96] was motivated to act as alternative to the Karhunen-Loève expansion which posed significant computational benefits. They showed that the solution to the SPDE (3.4.2) omitted a covariance structure of the form (3.4.1). Hierarchical modelling in statistics [116] has become quite crucial for better understanding of estimating the underlying unknown.

This can be translated to inverse problems where we are not only interested in the field u but its hyperparameters associated with it. Within hierarchical modelling there are commonly two approaches one can take: the *centred* approach and the *non-centred* approach. These approaches were derived by Papaspiliopoulos et al. in [115, 116] in the context of Gaussian processes for computational statistics. Translating this to our inverse setting, the non-centred approach can be viewed as the parameterization under which we aim to solve $(\xi, \theta) \in H^{-s}(D) \times \mathbb{H}$ from (3.4.3). While the centred approach differs as under its parameterization we aim to solve for

$(u, \theta) \in \mathcal{X} \times \mathbb{H}$ from (3.4.3). In terms of how the quantities (u, θ) and (ξ, θ) differ, their respective prior forms will be different as for the non-centred approach ξ and θ are independent. Before discussing each approach in more detail we present an important proposition which states both approaches are equivalent when generating samples from (3.4.2).

Proposition 3.4.1. *Given a Gaussian random field u with covariance operator $\mathcal{C}_\theta$, the centred and non-centred approaches to generate u are equivalent.*

Proof. Let $T : (\xi, \theta) \to u$ be a mapping where we choose $\mathcal{C}_\theta := \ell^d \beta (I - \ell^2 \Delta)^{-\alpha}$ for equation (3.4.3). We can express u through the Karhunen-Loève expansion

$$u = \sum_k \sqrt{\lambda_k} \hat{\xi_k} \phi_k, \qquad \hat{\xi_k} \sim \mathcal{N}(0, 1),$$

where (λ_k^2, ϕ_k) are the eigenpairs of $\mathcal{C}_\theta$ for $k = 1, 2$. Using the fact that both

$$u = \sum_k \hat{u_k} \phi_k, \tag{3.4.4}$$

$$\xi = \sum_k \hat{\xi_k} \phi_k, \tag{3.4.5}$$

we see after substituting (3.4.4) and (3.4.5) into (3.4.3), where $k = \begin{pmatrix} k_1 \\ k_2 \end{pmatrix}$, that

$$\frac{1}{\ell^{d/2}\sqrt{\beta}} (I - \ell^2 |k|^2)^{\frac{\alpha}{2}} \sum_k \hat{u_k} \phi_k = \sum_k \hat{\xi_k} \phi_k,$$
$$\frac{1}{\ell^{d/2}\sqrt{\beta}} (I - \ell^2 |k|^2)^{\frac{\alpha}{2}} \hat{u_k} = \hat{\xi_k}.$$

This implies

$$\hat{u_k} = \ell^{d/2} \sqrt{\beta} (I - \ell^2 |k|^2)^{-\frac{\alpha}{2}} \hat{\xi_k},$$

which is equivalent to $\lambda_k^2 := (I - \ell^2 |k|^2)^{-\alpha}$. □

3.4.1 Centred Formulation

We now characterize our inverse problem through the centre formulation. For this approach our prior will have the form

$$\mathbb{P}(u, \theta) = \mathbb{P}(u|\theta)\mathbb{P}(\theta), \tag{3.4.6}$$

via the definition of conditional probability. We are interested in the recovery of our unknown $u \in \mathcal{X}$ from noisy measurements of our data y where

$$y = \mathcal{G}(u) + \eta, \qquad \eta \sim \mathcal{N}(0, \Gamma). \tag{3.4.7}$$

We can further define a potential for our inverse problem $\Phi(u; y) : \mathcal{X} \to \mathbb{R}$ where

$$\Phi(u; y) = \frac{1}{2}|y - \mathcal{G}(u)|_{\Gamma}^2. \tag{3.4.8}$$

From the potential given in (3.4.8) we can define our data-likelihood as

$$\mathbb{P}(y|u) = \exp\big(-\Phi(u; y)\big). \tag{3.4.9}$$

Combing both our prior (3.4.6) and data-likelihood (3.4.9), via Bayes' Theorem, we can construct our posterior probability

$$\begin{aligned}\mathbb{P}(u, \theta|y) &\propto \mathbb{P}(y|u)\mathbb{P}(u, \theta) \\ &= \exp\big(-\Phi(u; y)\big)\mathbb{P}(u|\theta)\mathbb{P}(\theta).\end{aligned}$$

Remark 3.4.1. *We note that the inverse problem associated with the centred approach* (3.4.7) *is the exact same as the non-hierarchical inverse problem* (3.2.1) *as the data does not depend on the updated hyperparameters. Thus in deriving continuous-time limits, the limit for our updated random field* $u_n^{(j)}$ *should be equivalent.*

As with the non-hierarchical method, we are interested in analyzing the hierarchical approaches influence on the subspace property, specifically whether they can break away from this property. With the centred approach we know that the data is only conditioned on the field u and not its hyperparameters. Due to this we expect that with the centred approach, (u, θ) to lie within the span of the initial ensemble $\mathcal{A}$. The following theorem verifies this in the discrete case.

Theorem 3.4.1. *For every* $(n, j) \in \mathbb{N} \times \{1, \dots, J\}$ *we have* $u_{n+1}^{(j)}$, $\theta_{n+1}^{(j)} \in \mathcal{A}$ *and hence* u_{n+1}, $\theta_{n+1} \in \mathcal{A}$.

Proof. The proof follows similarly to that in [79] which is based on simple induction, but with the key difference of the inclusion of our hyperparameters $\theta_n^{(j)}$. We define our Kalman gain matrices as

$$K_n^u = \begin{pmatrix} C_n^{uw}\left(C_n^{ww} + \Gamma\right)^{-1} \\ C_n^{uu}\left(C_n^{ww} + \Gamma\right)^{-1} \end{pmatrix},$$

$$K_n^\theta = \begin{pmatrix} C_n^{\theta w}\left(C_n^{ww} + \Gamma\right)^{-1} \\ C_n^{\theta\theta}\left(C_n^{ww} + \Gamma\right)^{-1} \end{pmatrix},$$

with empirical covariances $C_n^{uw}, C_n^{uu}, C_n^{\theta w}$. Recalling that the update equations are given as

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})), \tag{3.4.10}$$

$$\theta_{n+1}^{(j)} = \theta_n^{(j)} + C_n^{\theta w}(C_n^{ww} + \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})). \tag{3.4.11}$$

By defining

$$d_{n+1}^{(j)} = (C_n^{ww} + \Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})),$$

Then the update formulas (3.4.10) and (3.4.11) can be defined as

$$\begin{aligned} u_{n+1}^{(j)} &= u_n^{(j)} + \frac{1}{J}\sum_{j=1}^{J}\langle \bar{\mathcal{G}}_{n+1}, d_{n+1}^{(j)}\rangle u_{n+1}^{(j)} \\ &= u_n^{(j)} + \frac{1}{J}\sum_{j=1}^{J}\langle \bar{\mathcal{G}}_{n+1}, d_{n+1}^{(j)}\rangle u_n^{(j)}, \\ \theta_{n+1}^{(j)} &= \theta_n^{(j)} + \frac{1}{J}\sum_{j=1}^{J}\langle \bar{\mathcal{G}}_{n+1}, d_{n+1}^{(j)}\rangle \theta_{n+1}^{(j)} \\ &= \theta_n^{(j)} + \frac{1}{J}\sum_{j=1}^{J}\langle \bar{\mathcal{G}}_{n+1}, d_{n+1}^{(j)}\rangle \theta_n^{(j)}. \end{aligned}$$

At step size n this shows that $u_{n+1}^{(j)}, \theta_{n+1}^{(j)} \in \mathcal{A}$ for $j \in \{1, \dots, J\}$. Hence since our

outputs u_{n+1}, θ_{n+1} at the end are defined as

$$u_{n+1} = \frac{1}{J}\sum_{j=1}^{J} u_{n+1}^{(j)},$$
$$\theta_{n+1} = \frac{1}{J}\sum_{j=1}^{J} \theta_{n+1}^{(j)},$$

it follows that both $u_{n+1}, \theta_{n+1} \in \mathcal{A}$. □

3.4.2 Non-Centred Formulation

As done previously in subsection 3.4.1 we characterize our inverse problem but now for the non-centred formulation. For this approach our prior will have the form

$$\mathbb{P}(\xi, \theta) = \mathbb{P}(\xi)\mathbb{P}(\theta), \tag{3.4.12}$$

via the definition of the non-centred approach in [115]. We are interested in the recovery of our unknown $(u, \theta) \in \mathcal{X} \times \mathbb{H}$ from noisy measurements of our data y where

$$y = \mathcal{G}(T(\xi, \theta)) + \eta, \qquad \eta \sim \mathcal{N}(0, \Gamma), \tag{3.4.13}$$

where $T : (\xi, \theta) \to u$ is an operator such that $u = T(\xi, \theta)$. This modified formulation of our unknown arises from the SPDE (3.4.2). As before we can further define a potential for our inverse problem $\Phi_{\text{NC}}(\xi, \theta; y) : \mathcal{X} \times \mathbb{H} \to \mathbb{R}$ where

$$\Phi_{\text{NC}}(\xi, \theta; y) = \frac{1}{2}|y - \mathcal{G}(T(\xi, \theta))|_{\Gamma}^2. \tag{3.4.14}$$

With NC denoting non-centred. From the potential given in (3.4.14) we can define our data-likelihood as

$$\mathbb{P}(y|\xi, \theta) = \exp\big(-\Phi_{\text{NC}}(\xi, \theta; y)\big). \tag{3.4.15}$$

Combing both our prior (3.4.12) and data-likelihood (3.4.15), via Bayes' Theorem, we can construct our posterior probability

$$\begin{aligned}\mathbb{P}(\xi, \theta|y) &\propto \mathbb{P}(y|\xi, \theta)\mathbb{P}(\xi, \theta) \\ &= \exp\big(-\Phi_{\text{NC}}(\xi, \theta; y)\big)\mathbb{P}(\xi)\mathbb{P}(\theta).\end{aligned}$$

Remark 3.4.2. *Unlike the centred approach, the non-centred formulation also differs as shown in the inverse problem* (3.4.13)*, namely that the data it is dependent*

on both the field u and the set of hyperparameters θ which is based on the transformation T. This would suggest the continuous-time limits would be different to the centred approach.

The difference in the prior form between both approaches is important in understanding why the non-centred approach is advantageous. Given that we are using ξ and that it is independent on the initialization of θ in the prior form, and under the transformation T, this allows a much less restriction induced by the subspace property. As a result both ξ and θ mix and update well, showcasing improvements over the centred approach. Also numerics in [33] demonstrated this for a range of non-linear PDE based inverse problems. The following theorem highlights this key difference related to the subspace property.

Theorem 3.4.2. *For every $(n, j) \in \mathbb{N} \times \{1, \dots, J\}$ we have $\xi_{n+1}^{(j)}$, $\theta_{n+1}^{(j)} \in \mathcal{A}$ and ξ_{n+1}, $\theta_{n+1} \in \mathcal{A}$ hence $u_{n+1} \notin T\mathcal{A}$, where $T\mathcal{A}$ is the space containing the transformed ensemble of particles.*

Proof. The proof follows very similarly to Theorem 3.4.1 but with the difference of the transformation $T(\xi, \theta) = u$ which abides by a difference space than the one of the initial ensemble $\mathcal{A}$. Therefore $u_{n+1} \notin T\mathcal{A}$. □

	Centred approach	**Non-centred approach**
Inverse problem	$y = \mathcal{G}(u) + \eta$	$y = \mathcal{G}(T(\xi, \theta)) + \eta$
Prior	$\mu_0 \equiv \mathbb{P}(u, \theta)$ $\mu_0 \equiv \mathbb{P}(u\|\theta) \times \mathbb{P}(\theta)$	$\mu_0 \equiv \mathbb{P}(\xi, \theta)$ $\mu_0 \equiv \mathbb{P}(\xi) \times \mathbb{P}(\theta)$
Likelihood	$\Phi(u; y) = \frac{1}{2}\|y - \mathcal{G}(u)\|_\Gamma^2$ $\mathbb{P}(y\|u) = e^{-\Phi(u;y)}$	$\Phi_{\text{NC}}(\xi, \theta; y) = \frac{1}{2}\|y - \mathcal{G}(T(\xi, \theta))\|_\Gamma^2$ $\mathbb{P}(y\|\xi, \theta) = e^{-\Phi_{\text{NC}}(\xi, \theta; y)}$
Posterior	$\mathbb{P}(u, \theta\|y) \propto \mathbb{P}(y\|u) \times \mathbb{P}(u, \theta)$ $\mathbb{P}(u, \theta\|y) \propto e^{-\Phi(u;y)}\mathbb{P}(u\|\theta)\mathbb{P}(\theta)$	$\mathbb{P}(\xi, \theta\|y) \propto \mathbb{P}(y\|\xi, \theta) \times \mathbb{P}(\xi, \theta)$ $\mathbb{P}(\xi, \theta\|y) \propto e^{-\Phi_{\text{NC}}(\xi, \theta; y)}\mathbb{P}(\xi)\mathbb{P}(\theta)$

Table 3.1: Comparison of both hierarchical approaches.

3.5 Hierarchical Continuous-Time Limits

3.5.1 Centred Approach

3.5.2 Nonlinear Noisy Case

We begin our derivation of a continuous-limit for the hierarchical iterative EnKF method by considering firstly the centred approach. As we are interested now in

$(u, \theta) \in \mathcal{X} \times \mathbb{H}$ we can construct a general posterior measure for $(u, \theta|y)$

$$\mu(du, d\theta) = \frac{1}{Z} \exp(-\Phi(u; y))\mu_0(du, d\theta),$$

with

$$Z := \int_{\mathcal{X}\times\mathbb{H}} \exp(-\Phi(u; y))\mu_0(du, d\theta).$$

Similarly with the non-hierarchical EnKF, we can derive an approximation of the posterior measure through introducing an artificial dynamical system $\mu_{n+1} = L_n\mu_n$ where

$$\mu_{n+1}(du, d\theta) = \frac{1}{Z_n} \exp(-h\Phi(u; y))\mu_n(du, d\theta),$$

and

$$Z_n := \int_{\mathcal{X}\times\mathbb{H}} \exp(-h\Phi(u; y))\mu_n(du, d\theta).$$

To construct our continuous-time limit we recall that the updates equations with the hierarchical iterative EnKF for

$$\begin{aligned} u_{n+1}^{(j)} &= u_n^{(j)} + C_n^{uw}(C_n^{ww} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})) \\ \theta_{n+1}^{(j)} &= \theta_n^{(j)} + C_n^{\theta w}(C_n^{ww} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})). \end{aligned}$$

Our update equations contain empirical covariance operators

$$\begin{aligned} C_n^{uw} &= \frac{1}{J-1}\sum_{k=1}^{J}(u^{(k)} - \bar{u}) \otimes (\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}) \\ C_n^{\theta w} &= \frac{1}{J-1}\sum_{k=1}^{J}(\theta^{(k)} - \bar{\theta}) \otimes (\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}) \\ C_n^{ww} &= \frac{1}{J-1}\sum_{k=1}^{J}(\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}) \otimes (\mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}), \end{aligned}$$

where, as before,

$$\bar{\theta} = \frac{1}{J}\sum_{k=1}^{J}\theta_n^{(k)}, \quad \bar{u} = \frac{1}{J}\sum_{k=1}^{J}u_n^{(k)}, \quad \bar{\mathcal{G}} = \frac{1}{J}\sum_{k=1}^{J}\mathcal{G}(u_n^{(k)}),$$

for $j = 1, \ldots, J$. By taking the limit of our update equations as $h \to 0$ this leads to an Euler-Maruyama (EM) discretization of the form

$$\frac{du^{(j)}}{dt} = C_n^{uw}(u)\Gamma^{-1}\big(y - \mathcal{G}(u^{(j)})\big) + C_n^{uw}(u)\sqrt{\Gamma^{-1}}\frac{dW^{(j)}}{dt} \tag{3.5.1}$$

$$\frac{d\theta^{(j)}}{dt} = C_n^{\theta w}(u)\Gamma^{-1}\big(y - \mathcal{G}(u^{(j)})\big) + C_n^{\theta w}(u)\sqrt{\Gamma^{-1}}\frac{dW^{(j)}}{dt}, \tag{3.5.2}$$

such that $W^{(j)}$ are cylindrical Brownian motions. By substituting the covariance operators $C_n^{uw}, C_n^{\theta w}$ in (3.5.1) and (3.5.2) this leads to

$$\frac{du^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle \mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}, y - \mathcal{G}(u^{(j)}) + \sqrt{\Gamma}\frac{dW^{(j)}}{dt}\Big\rangle_{\Gamma}(u^{(k)} - \bar{u}) \tag{3.5.3}$$

$$\frac{d\theta^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle \mathcal{G}(u^{(k)}) - \bar{\mathcal{G}}, y - \mathcal{G}(u^{(j)}) + \sqrt{\Gamma}\frac{dW^{(j)}}{dt}\Big\rangle_{\Gamma}(\theta^{(k)} - \bar{\theta}). \tag{3.5.4}$$

In the hierarchical case the key distinguishment we see is firstly that our formulation of our measure differs as we take more than one underlying unknown, but also, when taking the limit $h \to 0$ we see we have coupled systems of SDEs. Using the same arguments in the non-hierarchical case given there is a solution to both (3.5.3) and (3.5.4)

3.5.3 Linear Noise-Free Case

which after further substitution of the linear operator $A \in \mathcal{L}((\mathcal{X}\times\mathbb{H}), \mathcal{Y})$ our coupled SDEs read

$$\frac{du^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle A(u^{(k)} - \bar{u}), y - Au^{(j)}\Big\rangle_{\Gamma}(u^{(k)} - \bar{u}),$$

$$\frac{d\theta^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle A(\theta^{(k)} - \bar{\theta}), y - Au^{(j)}\Big\rangle_{\Gamma}(\theta^{(k)} - \bar{\theta}).$$

Given our covariance operators for the centred approach

$$C(u) = \frac{1}{J-1}\sum_{k=1}^{J}(u^{(k)} - \bar{u}) \otimes (u^{(k)} - \bar{u}), \tag{3.5.5}$$

$$C(\theta) = \frac{1}{J-1}\sum_{k=1}^{J}(\theta^{(k)} - \bar{\theta}) \otimes (\theta^{(k)} - \bar{\theta}), \tag{3.5.6}$$

and $\Gamma = 0$, we can express (3.5.5) and (3.5.6) as

$$\frac{du^{(j)}}{dt} = -C(u)D_u\Phi(u^{(j)}; y), \tag{3.5.7}$$

where our potential is defined as

$$\Phi(u; y) = \frac{1}{2}\|\Gamma^{-1/2}(y - Au)\|^2.$$

As before we can interpret (3.5.7) as each particle $\{u^{(j)}\}_{j=1}^J$ performing a gradient descent for $\Phi(\cdot; y)$. This is the exact same limit and gradient flow structure that we have in the non-hierarchical case (3.3.8).

3.5.4 Non-Centred Approach

3.5.5 Nonlinear Noisy Case

Our construction of our posterior measure differs with the non-centred approach as we have a modified potential (3.4.14). Using this potential our posterior measure for $(\xi, \theta|y)$ now reads

$$\mu(d\xi, d\theta) = \frac{1}{Z}\exp(-\Phi_{\text{NC}}((\xi, \theta); y))\mu_0(d\xi, d\theta),$$

with

$$Z := \int_{H^{-s}(D)\times\mathbb{H}} \exp(-\Phi_{\text{NC}}((\xi, \theta); y))\mu_0(d\xi, d\theta).$$

As similarly done for the centred approach we can derive an approximation by an artificial dynamical system $\mu_{n+1,\text{NC}} = L_{n,\text{NC}}\mu_{n,\text{NC}}$ where

$$\mu_{n+1}(d\xi, d\theta) = \frac{1}{Z_n}\exp(-h\Phi_{\text{NC}}((\xi, \theta); y))\mu_{n,\text{NC}}(d\xi, d\theta),$$

and

$$Z_n := \int_{H^{-s}(D)\times\mathbb{H}} \exp(-h\Phi_{\text{NC}}((\xi, \theta); y))\mu_n(d\xi, d\theta).$$

The prediction step of the non-centred approach is a mirror to that of the centred approach but with the difference of updating ξ instead of u, and we evaluate both(ξ, θ) in the forward evaluation. By defining $\mathcal{G}^T = \mathcal{G} \circ T$ our update equations

for the non-centred approach are

$$\xi_{n+1}^{(j)} = \xi_n^{(j)} + C_n^{\xi p}(C_n^{ww} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}^T(\xi_n^{(j)}, \theta_n^{(j)}))$$
$$\theta_{n+1}^{(j)} = \theta_n^{(j)} + C_n^{\theta w}(C_n^{ww} + h^{-1}\Gamma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}^T(\xi_n^{(j)}, \theta_n^{(j)})),$$

where we again assume that $\iota_{n+1} \sim \mathcal{N}(0, h^{-1}\Gamma)$ such that $y_{n+1}^{(j)} = y + \iota_{n+1}$, and that our empirical covariances are defined as

$$C_n^{\xi p} = \frac{1}{J-1}\sum_{k=1}^{J}(\xi^{(k)} - \bar{\xi}) \otimes (\mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) - \overline{\mathcal{G}^T}),$$
$$C_n^{\theta w} = \frac{1}{J-1}\sum_{k=1}^{J}(\theta^{(k)} - \bar{\theta}) \otimes (\mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) - \overline{\mathcal{G}^T}),$$
$$C_n^{ww} = \frac{1}{J-1}\sum_{k=1}^{J}(\mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) - \bar{\mathcal{G}}) \otimes (\mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) - \overline{\mathcal{G}^T}).$$

We see that with the covariances defined above we have the addition of the hyper-parameter included in the evaluation of the forward operator which coincides with the inverse problem formulation (3.4.13) where

$$\overline{\mathcal{G}^T} = \frac{1}{J}\sum_{j=1}^{J}\mathcal{G}^T(\xi_n^{(j)}, \theta_n^{(j)}), \quad j = 1, \dots, J.$$

Therefore by taking the limit of our update equations as $h \to 0$, we have the coupled SDEs

$$\frac{d\xi^{(j)}}{dt} = C_n^{\xi p}(\cdot)\Gamma^{-1}\big(y - \mathcal{G}^T(\xi^{(j)}, \theta^{(j)})\big) + C_n^{\xi p}(\cdot)\sqrt{\Gamma^{-1}}\frac{dW^{(j)}}{dt} \tag{3.5.8}$$
$$\frac{d\theta^{(j)}}{dt} = C_n^{\theta w}(\cdot)\Gamma^{-1}\big(y - \mathcal{G}^T(\xi^{(j)}, \theta^{(j)})\big) + C_n^{\theta w}(\cdot)\sqrt{\Gamma^{-1}}\frac{dW^{(j)}}{dt}, \tag{3.5.9}$$

such that $W^{(j)}$ are cylindrical Brownian motions. Using the formula for the covariances from (3.5.8) and (3.5.9)

$$\frac{d\xi^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle \mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) - \overline{\mathcal{G}^T}, y - \mathcal{G}^T(\xi^{(j)}, \theta^{(j)}) + \sqrt{\Gamma}\frac{dW^{(j)}}{dt}\Big\rangle_\Gamma(\xi^{(k)} - \bar{\xi})$$
$$\frac{d\theta^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\Big\langle \mathcal{G}^T(\xi^{(j)}, \theta^{(k)}) - \overline{\mathcal{G}^T}, y - \mathcal{G}^T(\xi^{(j)}, \theta^{(j)}) + \sqrt{\Gamma}\frac{dW^{(j)}}{dt}\Big\rangle_\Gamma(\theta^{(k)} - \bar{\theta}).$$

3.5.6 Linear Noise-Free Case

As before we work in a linear setting where we define $\mathcal{G}^T(\cdot) = A\cdot$. Substituting $\mathcal{G}^T(\xi^{(k)}, \theta^{(k)}) = Au^{(k)}$, for $k = 1, \ldots, J$, yields

$$\frac{d\xi^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\left\langle A(u^{(k)} - \bar{u}), y - Au^{(k)}\right\rangle_{\Gamma}(\xi^{(k)} - \bar{\xi}) \tag{3.5.10}$$

$$\frac{d\theta^{(j)}}{dt} = \frac{1}{J}\sum_{k=1}^{J}\left\langle A(u^{(k)} - \bar{u}), y - Au^{(k)}\right\rangle_{\Gamma}(\theta^{(k)} - \bar{\theta}). \tag{3.5.11}$$

We notice with the SDEs the inclusion of the hyperparameter θ highlights one of the differences for the non-centred approach. Given our covariance operators for the non-centred approach

$$C(\xi) = \frac{1}{J-1}\sum_{k=1}^{J}(\xi^{(k)} - \bar{\xi}) \otimes (\xi^{(k)} - \bar{\xi})$$

$$C(\theta) = \frac{1}{J-1}\sum_{k=1}^{J}(\theta^{(k)} - \bar{\theta}) \otimes (\theta^{(k)} - \bar{\theta}),$$

which we can express (3.5.5) and (3.5.6), where $\Gamma = 0$, as

$$\frac{d\xi^{(j)}}{dt} = -C(\xi)D_u\Phi_{\mathrm{NC}}(u^{(j)}; y) \tag{3.5.12}$$

$$\frac{d\theta^{(j)}}{dt} = -C(\theta)D_u\Phi_{\mathrm{NC}}(u^{(j)}; y), \tag{3.5.13}$$

with potential

$$\Phi_{\mathrm{NC}}(\xi, \theta; y) = \frac{1}{2}\|\Gamma^{-1/2}(y - Au)\|^2.$$

For the non-centred approach we have derived a coupled gradient flow system for both the underlying unknown (3.5.12) and the hyperparameters (3.5.13) that differs from its centred counterpart.

3.5.7 Hierarchical Covariance Inflation

With the developments of the EnKF there has been considerable advancements which have looked at alternative approaches that provide improvements. An issue that can arise with the EnKF is rank deficiency. This problem occurs from the empirical covariances when the number of ensemble particles J in the data space $\mathcal{Y}$

is less than that of the input space $\mathcal{X}$. One way to counteract this issue is through the technique of covariance inflation [57]. We now aim to derive continuous-time limits of hierarchical covariance inflation, for EnkF inversion. We will do so specifically for the non-centred case, given its advantages we have discussed and shown in [33]. This allows for a modification of our covariances $C(\xi), C(\theta)$ given by

$$C(\xi) \to \gamma C_0 + C(\xi) \tag{3.5.14}$$
$$C(\theta) \to \gamma \theta_0 + C(\theta), \tag{3.5.15}$$

with $\gamma \in \mathbb{R}^+$. Substituting (3.5.14) and (3.5.15) in our gradient flow system leads to, for $j = 1, \dots, J$,

$$\begin{aligned} \frac{d\xi^{(j)}}{dt} &= -(\gamma C_0 + C(\xi)) D_u \Phi_{\mathrm{NC}}(u^{(j)}; y) \\ \frac{d\theta^{(j)}}{dt} &= -(\gamma \ell_0 + C(\theta)) D_u \Phi_{\mathrm{NC}}(u^{(j)}; y). \end{aligned}$$

By taking the inner product with $D_u \Phi_{\mathrm{NC}}(u^{(j)}; y)$ we have

$$\begin{aligned} \frac{d\Phi_{\mathrm{NC}}(u^{(j)}; y)}{dt} &\leq -\gamma \| C_0^{1/2} D_u \Phi_{\mathrm{NC}}(u^{(j)}; y) \|^2 \\ \frac{d\Phi_{\mathrm{NC}}(u^{(j)}; y)}{dt} &\leq -\gamma \| \theta_0^{1/2} D_u \Phi_{\mathrm{NC}}(u^{(j)}; y) \|^2. \end{aligned}$$

which indicates that all limits are contained in the critical points of both potentials.

3.5.8 Hierarchical Localization

A further issue with the EnKF can arise from the correlation between the empirical covariances. If the correlation distance is long this can cause problems with updating our unknowns. Localization [58] is a method that aids by cutting off these long distances which helps improve the update of the estimate. It is usually achieved through the aid of convolution kernels that reduce distances of distant regions. The convolution kernels $\rho : D \times D \to \mathbb{R}$ are usually of the form

$$\rho(x, y) = \exp\big(- (x - y)^T \big),$$

given $D \subset \mathbb{R}^d$ for $d \in \mathbb{N}$, thereby allowing us to define continuous-time limits

$$\frac{d\xi^{(j)}}{dt} = C^{\text{loc}}(\xi) D_u \Phi_{\text{NC}}(u^{(j)}; y)$$
$$\frac{d\theta^{(j)}}{dt} = C^{\text{loc}}(\theta) D_u \Phi_{\text{NC}}(u^{(j)}; y),$$

where

$$C^{\text{loc}}(\xi)\Phi(x) = \int_D \phi(y) k(x,y) \rho(x,y) dy$$
$$C^{\text{loc}}(\theta)\Phi(x) = \int_D \phi(y) k(x,y) \rho(x,y) dy,$$

given that $k(x,y)$ corresponds to the kernel of the covariances and $\phi \in \mathcal{X}$.

3.6 Numerical Experiments

We now wish to add some numerics to the theory discussed regarding the variants of localization and covariance inflation. We have seen through numerical investigation in [33] that the theory discussed here matches with the results attained for various non-linear and linear inverse problems. In the context of this work we will only test for linear inverse problems, specifically a 1D elliptic PDE. Our numerics will consist of learning rates of hyperparameters and the reconstruction of the truth for both hierarchical localization and covariance inflation. Given a domain $D \subset \mathbb{R}^d$, for $d = 1$, with boundary ∂D, our forward model is concerned with solving for $p \in H_0^1(D)$ from

$$\frac{d^2 p}{dx^2} + p = u \quad x \in D, \tag{3.6.1}$$
$$p = 0 \quad x \in \partial D. \tag{3.6.2}$$

Here we assume a domain of $D = (0, \pi)$ with prescribed zero Dirichlet boundary conditions (3.6.2). The inverse problem associated with the forward problem (3.6.1) is the recovery of noisy measurements from the right hand side u where

$$y_j = l_j(p) + \eta_j, \tag{3.6.3}$$

such that $l_j \in V^*$ where V^* is the dual space of $H_0^1(D)$. By defining $\mathcal{G}_j(T(\xi,\theta)) = l_j(p)$, where we take our unknown function $T(\xi,\theta) = u$, we can rewrite (3.6.3) as

$$y = \mathcal{G}(T(\xi,\theta)) + \eta. \tag{3.6.4}$$

Our inverse solver for our numerics will be the iterative ensemble Kalman method [79], where we aim to reconstruct a Gaussian random field. We will use the exact continuous-time limits and take a discretization of the ODEs, using the MATLAB solver ODE45. The time-stepping we use for our ODE solver will be chosen as $h = 0.01$.

Remark 3.6.1. *Note the choice of the time-stepping can be crucial. Our results we will present are for the case of $h = 0.01$. However we have tested other values for h for which the experiments differed marginally.*

Initially we set our initial ensemble based on a prior distribution. Our initial field will be set such that $\xi_0^{(j)} \sim \mathcal{N}(0, \mathcal{C}_\theta)$ where $\mathcal{C}_\theta$ takes the form (3.4.3). To generate our initial ensemble with covariance structure of (3.4.1) we first discretize our SPDE (3.4.2) for u using a 1D centred finite difference method

$$u_i - \ell^2 \frac{u_{i+1} - 2u_i + u_{i-1}}{h_*^2} = \xi_i, \quad \xi_i \sim \mathcal{N}(0, \alpha\ell/h_*),$$

which in matrix form is given as

$$\begin{pmatrix} 1+2\frac{\ell^2}{h_*^2} & -\frac{\ell^2}{h_*^2} & 0 & \dots & 0 \\ -\frac{\ell^2}{h_*^2} & 1+2\frac{\ell^2}{h_*^2} & -\frac{\ell^2}{h_*^2} & \ddots & \vdots \\ 0 & -\frac{\ell^2}{h_*^2} & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -\frac{\ell^2}{h_*^2} \\ 0 & \dots & 0 & -\frac{\ell^2}{h_*^2} & 1+2\frac{\ell^2}{h_*^2} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_I \end{pmatrix} = \begin{pmatrix} \xi_1 \\ \xi_2 \\ \vdots \\ \xi_I \end{pmatrix}.$$

After generating u we take our linear mapping $T : \mathcal{X} \to \mathcal{X}$ to generate samples of ξ. Our mesh size for our discretization is given as $h_* = 1/50$ where $I = 50$. From θ we will only treat the parameter of the length-scale ℓ hierarchically. Our reason for this is that in a 1D numerical example the length-scale has a more notable effect on how the input is generated. We keep $\sigma = 1$ and $\alpha = 0.8$ while setting a prior now on the inverse length-scale $\tau = 1/\ell$

$$\tau \sim \mathcal{U}[10, 40], \tag{3.6.5}$$

This modification is for numerical purposes to notice the learning at a better rate. We generate our prior form $\mathbb{P}(\xi, \theta)$ by solving the SPDE (3.4.2) using a piecewise linear finite element method. Our truths will be chosen such that $\xi^\dagger \sim \mathcal{N}(0, \mathcal{C}_\theta^\dagger)$, similar to the initial ensemble, where $\theta^\dagger = (\sigma^\dagger, \alpha^\dagger, \tau^\dagger) = (1, 0.8, 37)$. For our iterative method we set an ensemble size of $J = 50$ and an iteration count of $n = 15$, with covariance noise $\Gamma = 0.01^2 I$. We discretize our PDE model (3.6.1) with a different mesh size of $h^* = 1/50$ using a centred finite difference method. We make inference of our unknown through 16 chosen observations which lie on the true value of the unknown. For implementing covariance inflation we set the parameter as $\gamma = 0.1$. For our prior covariance θ_0, as we consider it as uniformly distributed, we modify this based on a Gaussian where we have $\theta_0 \sim \mathcal{N}(25, \sqrt{7}^2)$. This is the case for the inverse length-scale.

In Figure 3.1 we analyze the performance of hierarchical localization by comparing it with non-hierarchical localization and the standard EnKF. We see that in the left subfigure the standard EnKF and localization perform similarly emulating a smooth function. However for hierarchical localization we see an improved reconstruction which is more closely related to the truth, which incorporates its sharper features. This can be attributed to changes in the length-scale which are verified in the right sub figure, where we see that by adopting a hierarchical approach we can effectively learn the true value of the length-scale which is $\tau^\dagger = 37$. The learning of the length-scale remains consistent with the results of [33] where the hyperparameters learn the true value quickly and reach a limit before the learning stops prior to the termination of the experiment.

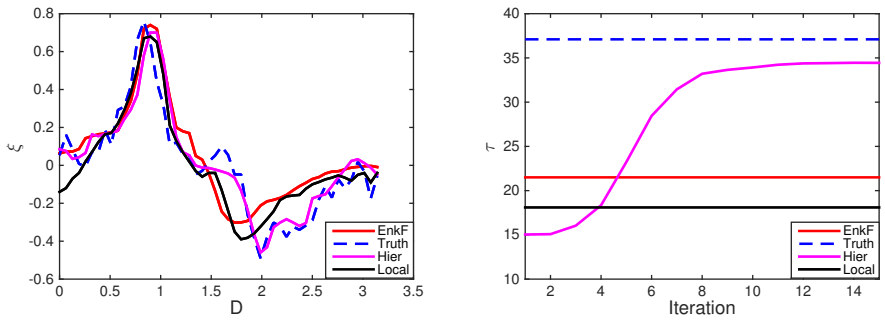

Figure 3.1: Performance of hierarchical localization. Left: reconstruction of the truth. Right: learning rate of the length-scale.

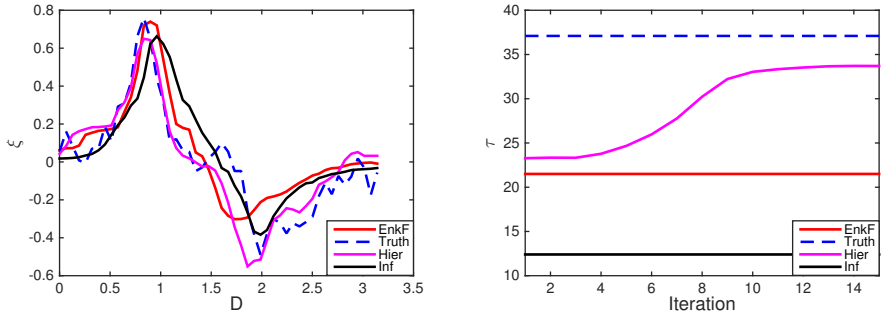

Figure 3.2: Performance of hierarchical covariance inflation. Left: reconstruction of the truth. Right: learning rate of the length-scale.

We see similar results when analyzing hierarchical covariance inflation, where learning the length-scale improves on the overall reconstruction of the truth as shown in Figure 3.2.

3.7 Conclusion

The objective of this work was to introduce analysis regarding the recent hierarchical approaches that were applied to EKI [33]. We have given a detailed description and comparison of both the centred and non-centred approaches. For each case we have shown how they relate to the subspace property where we further derived continuous-time limits in in both the noisy and noise-free case. Our analysis clarifies that by taking a non-centred approach one can significantly improve the performance of EKI. This is verified through the transformation which allows the ensemble of particles to leave the span of the initial ensemble. We introduced certain variants of the EnKF to show that hierarchically this can be achieved too, which was verified through a numerical experiment.

One avenue of interest is to consider, as done in [129], the behaviour of the gradient flow structure defined for the non-centred approach (3.5.12) and analyze the relationship with the subspace property. This is beyond the scope of this chapter, but analyzing the behaviour could potentially result in improved convergence results over the non-hierarchical case. A further direction is to extend this work by using certain SDE discretizations of EKI. This was analyzed in [19], the natural extension of this would be to translate this in a hierarchical manner.

Chapter 4

Reduced basis methods for Bayesian inverse problems

4.1 Overview

This chapter is concerned with the deployment of reduced order models to reduce the cost of the forward model. This will be focused on one reduced order model which is the reduced basis method. As the aim of this chapter is on reducing computational cost, this is distinguishable compared to the other chapters. However we note this chapter contains a brief review of Bayesian inverse problems and EKI, which have already been mentioned in previous chapters.

4.2 Introduction

Uncertainty quantification (UQ) has risen as a topic of interest for both academics and practitioners in recent years. It is concerned with quantifying the complexity and uncertainty that can arise within a model or a set of equations. Due to its popularity it has sparked the development of various numerical methods [9, 42, 38, 67, 148] which aid by attempting to quantify this uncertainty. One form of how uncertainty arises is through quantities within the model such as the diffusion coefficient, boundary conditions and the source terms. In order to model these quantities better, while taking into account the randomness which can arise, we pose the model as a parametric system. Many methods have been developed to solve parametric systems of differential equations. One group of numerical methods which have shown promise of recent is that of reduced order models (ROMs) [119]. Some popular examples include

- Proper orthogonal decomposition
- Reduced basis method
- Krylov subspaces

The general philosophy of these methods can be thought of as a means to reduce the complexity within a system by reducing the dimensionality. Due to this one can think of these ROMs as projection-based methods, which aim to approximate the random solution on a lower-dimensional subspace through a Galerkin projection. With ROMs they seek a reduced subspace. This work will specifically be looking at one of the ROMs stated above, the reduced basis method (RBM). This particular method is based on a greedy optimisation algorithm which builds the reduced subspace. The RBM is specific for parametric partial differential equations (PDEs). This subspace consists of approximate solutions to the variational formulation with respect to the parameters. These reduced approximation solutions in ROMs are commonly referred to as snapshots. Comparing ROMs, a key difference and advantage of the RBM is that it uses a posteriori error estimates to control the stability and error of the solution, while other methods such as the proper orthogonal decomposition method (POD) are based on a priori estimates. Recent work [141] has indicated that within lower dimensional parameter spaces both methods perform similarly but once the dimension size is increased the RBM performs considerably more effectively.

As well as developments in handling randomness within models, inverse problems have seen an increase in popularity within UQ. Traditionally inverse problems were numerically solved through optimisation techniques in the classic approach but an alternative approach was formulated which guarantees well-posedness and tackles uncertainty, which is known as the Bayesian approach [87, 137]. Since then a wide array of methods have been used to solve Bayesian inverse problems, most notably Markov chain Monte Carlo (MCMC) [39, 80] and ensemble-based methods [77, 79]. A comparison of these types of methods can be found by Iglesias et al. [78].

Our aims in this chapter is to combine the ideas behind ROMs and inverse problems where we look to implement the RBM as the forward solver. The motivation behind this is that in inverse problems a common issue is the computational burden of the forward solver. Thus by working with the RBM we hope to get a cheaper approximation of the inverse solution. Our focus will be on two PDE-based inverse problems, that of groundwater flow (GWF) and electrical impedance tomography (EIT). For the GWF problem we aim to test our newly proposed method both in terms of accuracy and cost. While for EIT our primary focus will be on

deriving a new a posteriori error bound, where we omit numerics. One important aspect of RBMs is the training set which is a collection of the points represented as the parameters. To complement the work discussed above, we further consider the question of how to effectively choose the parameters. Various collections of points (or training sets) have already been tested, however we propose a new training set which is motivated from ideas of stochastic collocation and Lagrangian interpolation. This new training set will be tested in the context of both forward and inverse problems.

The outline of this chapter is as follows: at the beginning in Section 4.3 we provide an overview of random PDE theory and finite element methods before discussing the RBM in detail in Section 4.4. We then provide a discussion of the various training sets in Section 4.5 which includes a number of numerical investigations. We then present an introduction into Bayesian inverse problems 4.6 in which a review will be done, whilst introducing the inverse solver which is the iterative Kalman method. This will lead onto Section 4.7 where we present various numerical examples of the coupled inverse solver. We conclude by deriving a posteriori error bound for the RBM applied to EIT in Section 4.8, and mentioning a number of final remarks on further directions of research.

Literature Review

The last decade has seen further development and applications [26, 108] of reduced basis methods in the context of solving high dimensional problems. Since its application towards differential equations it has been extensively used in UQ related problems which has sparked much popularity within the UQ community. Knowing that it has considerable improvements on computational cost over traditional methods it seems a natural choice to implement into inverse solvers. Despite this idea there has not been much work which has exploited this. Recently there has been some work done which has looked at doing this but in the context of analysing inverse problems in the classical sense. What was proposed was a reduced basis Landweber method by Garmatter et al. [65]. As of yet this has been one of few examples where inverse problems have exploited the benefits of the RBM. In the Bayesian setting of inverse problems this still has to be explored.

Ensemble based methods [77, 79] have proven to be a suitable choice for solving inverse problems which aim to use ensemble approximations for the derivative operators. One of these methods that looks to improve on these perhaps limitation is the iterative ensemble Kalman method which serves as a derivative-free optimization

method. The method combines ensemble techniques from the Kalman filter with that of the regularized Levenburg-Marqardt method [71]. Comparing to MCMC methods one key difference is that the number of forward evaluations required is significantly less as shown in [78]. Given this advantage of computational power, this could speed up the inverse problem considerably further.

Currently there has been recent literature on employing RBM techniques into the EnKF but in a different context to what we are aiming to do. [65] has discussed the role of the RBM within inverse problems but from a classical non-Bayesian approach. There has also been work specifically implementing it in a Bayesian setting but from a data-assimilation perspective by Manzoni at al. [114], thus so far no work has been done on the iterative Kalman method.

As of yet there has been no advancements on applying the RBM in conjunction with EIT. One of the reasons behind this could be the difficulty of attaining rigorous RB a posteriori error bounds for forward problem associated with EIT, or even with the numerical implementation.

4.3 Background Material

4.3.1 Random PDE Theory

Given some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where Ω is our sample space, $\mathcal{F}$ is a sigma-algebra and $\mathbb{P} : \mathcal{F} \to [0,1]$ is a probability measure where we define samples from our sample space as $\vartheta \in \Omega$. We assume we have a domain $D \subset \mathbb{R}^d$ for $d < \infty$ where D is a Lipschitz domain with boundary ∂D. Given a random field $\kappa(x;\vartheta)$ we are interested in finding the solution $p(x;\vartheta)$ to the random PDE

$$-\nabla \cdot (\kappa(x;\vartheta)\nabla p(x;\vartheta)) = f, \quad \in D, \tag{4.3.1a}$$

$$p(x;\vartheta) = 0, \quad \in \partial D, \tag{4.3.1b}$$

where we have imposed Dirichlet boundary conditions. Before discussing a well-posedness theorem of (4.3.1) we first review a number of assumptions regarding both the source term f and the random coefficient $\kappa(x;\vartheta)$.

Assumption 4.3.1. *There exist constants* $0 < \kappa_{min} < \kappa_{max} < \infty$ *such that*

$$\mathbb{P}(\kappa_{min} \leq \kappa(x;\vartheta) \leq \kappa_{max}, \ \forall x \in D) = 1. \tag{4.3.2}$$

where $\kappa(x;\vartheta) \in L^\infty(D;\mathbb{R})$ *and* $\operatorname{esssup}\kappa(x;\vartheta) = \kappa_{max}(\vartheta) > 0, \mathbb{P} - a.s.$

We also assume that $\kappa(x;\vartheta)$ has a very particular structure.

Assumption 4.3.2. *Let $\Gamma \subset \mathbb{R}^K$ be called the parameter space. Let $\kappa_0, \kappa_1, \ldots, \kappa_K \in L^\infty(D)$, and let $u_1, \ldots, u_K$ be independent random variables from Ω taking values in $\mathbb{R}$ such that $u = (u_1, u_2, \ldots, u_K) \in \Gamma$. Then our random coefficient can be be expressed through the following series*

$$\kappa(x;\vartheta) = \kappa_0(x) + \sum_{k=1}^{K} \kappa_k(x) u_k(\vartheta). \tag{4.3.3}$$

That is, $\kappa(x;\vartheta)$ has a linear dependence on finitely many random variables. For ease of computation and analysis, we will assume that the u_k's are chosen uniformly at random from the uniform distribution $\mathcal{U}[-1,1]$. Given these two assumptions we present the uniqueness and existence theorem of random PDEs which is an application of the Lax-Millgram Theorem:

Theorem 4.3.1. *Assume that $f \in V^*, \mathbb{P}-a.s$, and*

$$\operatorname{essinf} \kappa(x;\vartheta) = \kappa_{min}(\vartheta) > 0, \quad \mathbb{P}-a.s.$$

Then $\mathbb{P}-a.s$ (4.3.1) has a unique solution which satisfies

$$\|p(\cdot;\vartheta)\|_V \leq \frac{\|f\|_{V^*}}{\kappa_{min}(\vartheta)}.$$

Assume further that either:
(i) $f \in V^*$ is deterministic and κ is distributed according to a uniform prior.
(ii) $\kappa = e^u,\ u \in L^\infty(D;\mathbb{R})$ deterministic and $f \in L^2_{\mathbb{P}}(\Omega; V^*)$.

Then (4.3.1) has a solution $\mathbb{P}-a.s$ and in $L^2_{\mathbb{P}}(\Omega; V)$.

Proof. The proof can be found in [99] which is based on the Lax-Milgram Lemma. □

Remark 4.3.1. *We note that the Lax-Milgram Theorem in Theorem 4.3.1 holds also for a random source term $f(x;\vartheta)$, but for the purposes of this work we keep f deterministic as stated above.*

4.3.2 Finite Element Method

To solve numerically a realisation of (4.3.1), we use a finite element method (FEM), which is based on the Galerkin projection. In particular, let $D = [0,1]$ and $h \in (0,1)$ and let $x_0 = 0 < x_1^h < \ldots < x_{N_h} < x_{N_h+1}^h = 1$ be a partition of D such that

$x_i^h - x_{i-1}^h = h$, for all $i \in \{1, \dots, N_h\}$. Define the basis functions $\{\phi_j\}_{j=1}^{N_h}$ such that $\phi_j(x_i^h) = \delta_{ij}$, for all $i \in \{0, 1, \dots, N_h\}$, and interpolate linearly between any two points of the partition. Then $\{\phi_j\}_{j=1}^{N_h}$ is a basis of a finite dimensional subspace V_h of $H_0^1(D)$ which contains all the functions $q \in C^0(D)$ such that $q|_{[x_{i-1}^h, x_i^h]}$ is a linear polynomial, for all $i \in \{1, \dots, N_h\}$.

We now consider the family $\{V_h\}_{h \in (0,1)}$ of the finite dimensional subspaces of $H_0^1(D)$, generated by the discretisation parameters $h \in (0, 1)$, and the finite dimensional equations

$$A(p_h(\vartheta), q_h; \vartheta) = l(q_h), \quad \forall q_h \in V_h. \tag{4.3.4}$$

We know that there exists a unique solution $p_h(\vartheta)$ of (4.3.4) and it is called the Galerkin projection of the solution $p(\vartheta)$ onto V_h. As $V_h = \text{span}\{\phi_1, \dots, \phi_{N_h}\}$, where $N_h = \dim V_h$ and $\{\phi_i\}_{i=1}^{N_h}$ is the basis of V_h, we can express the solution $p_h(\vartheta)$ in terms of the basis functions $\{\phi_i\}_{i=1}^{N_h}$

$$(p_h(\vartheta))(x) = \sum_{i=1}^{N_h} P_i(\vartheta)\phi_i(x), \tag{4.3.5}$$

where $\{P_i(\vartheta)\}_{i=1}^{N_h}$ are real numbers still to be calculated. In fact such a calculation can be done by solving the following system of linear equations

$$\sum_{i=1}^{N_h} A(\phi_i, \phi_j; \vartheta) P_i(y) = l(\phi_j), \quad j = 1, \dots, N_h. \tag{4.3.6}$$

It is more convenient to think of the above linear system as

$$A_h^\vartheta \cdot P_h^\vartheta = l_h, \tag{4.3.7}$$

where $A_h^\vartheta \in \mathbb{R}^{N_h \times N_h}$ is called the stiffness matrix given by $A_h^\vartheta(i, j) = A(\phi_i, \phi_j; \vartheta)$, for all $i, j \in \{1, \dots, N_h\}$, $P_h^y = (P_1(\vartheta), \dots, P_{N_h}(\vartheta))$ and $l_h = (l(\phi_1), \dots, l(\phi_{N_h}))$. Note that the stiffness matrix is tri-diagonal because each basis function only overlaps with the two neighbouring basis functions. This implies that for large enough N_h the matrix is sparse.

A good approximation result of the Galerkin projections $\{p_h(\vartheta)\}_{h \in (0,1)}$ is given by Céas Lemma which in general states the following.

Lemma 4.3.1. *Let H be a Hilbert space, A be a bilinear form on H, which is coercive with constant $\kappa_{min} > 0$ and bounded with constant $\kappa_{max} > 0$, and l a linear*

functional on H. Let $p \in H$ such that $A(p,q) = l(q)$, for all $q \in H$, and consider a finite dimensional subspace V of H and p_V such that $A(p_V, q) = l(v)$, for all $q \in V$. Then

$$\|p - p_V\|_H \leq \frac{\kappa_{max}}{\kappa_{min}} \inf_{q \in V} \|p - q\|_H. \tag{4.3.8}$$

In the case of the FEM the lemma above takes the following form

$$\|p(\vartheta) - p_h(\vartheta)\|_{H_0^1(D)} \leq C \inf_{q_h \in V} \|p(\vartheta) - q_h\|_{H_0^1(D)}, \tag{4.3.9}$$

for $C = \frac{\kappa_{max}}{\kappa_{min}} > 0$.

Finally, we have the following convergence result in terms of h, provided that $p(\vartheta) \in H^2(D)$, the coefficient $\kappa(\cdot, \vartheta) \in C^1(D)$ and D is a convex, bounded, Lipschitz boundary

$$\|p(\vartheta) - p_h(\vartheta)\|_{H_0^1(D)} \leq Ch\|p(\vartheta)\|_{H^2(D)}. \tag{4.3.10}$$

Note that the classical theory of PDEs implies that $p(\vartheta) \in H^2(D)$ since $f \in L^2(D)$.

4.4 Reduced Basis Method

In this section we provide an overview of the RBM beginning with a general understanding before discussing the main components of the numerical method. This will lead to discussing our new proposed set of points which can be used as the training set for our experiments. Note for ease of notation we label $p = p_h$ noting that now p denotes our true solution, which is a finite dimensional approximate solution of (4.3.1). Specifically for this chapter we take it be of the form (4.3.5).

The RBM is a projection based method which is to serve as a cheaper evaluation of a forward problem. The idea behind this method is to replace the high order (N_h) FEM basis by a lower one consisting of solutions induced by parameters $\{p(\vartheta^n)\}_{n=1}^N$. It uses a greedy sampling algorithm to aid with parameter selection $\vartheta^1, \ldots, \vartheta^N \in \Gamma$ which in turn builds a reduced basis space $X_N := \text{span}\{p(\vartheta^1), \ldots, p(\vartheta^N)\}$, by solving a Galerkin projection problem

$$A(p_N, q; \vartheta) = F(q), \ \forall q \in X_N, \tag{4.4.1}$$

for $N = 2, \ldots, N_{\max}$, such that $N_{\max} \ll N_h$. Instead of viewing the RB solution as

an approximation of (4.3.1), it can be viewed and expressed in the following way

$$p_N(\vartheta) = \sum_{m=1}^{N} p_{Nm}(\vartheta)\zeta_m, \tag{4.4.2}$$

where $\{\zeta_m\}_{m=1}^N$ are the orthonormalized solutions. This procedure is done for computational purposes where now $X_N = \text{span}\{\zeta_1, \ldots, \zeta_N\}$. The solutions which correspond to the reduced basis space X_N are referred to as *snapshots*, and Γ is the space where the parameters lie. Usually for further computational purposes instead of using the whole parameter space we take a training sample $\Xi_{\text{train}} \in \Gamma$, hoping it is a good representative of the parameter space. As briefly mentioned, an important question to ask is how are the parameters chosen in an optimal way?

Initially we look to build the reduced basis space X_N in a hierarchical manner which is initially done by sampling the first parameter

$$\vartheta^{N+1} = \text{argmax}_{\vartheta\in\Gamma}\|p(\vartheta) - p_N(\vartheta)\|_X,$$

However this can be expensive to compute so commonly the first parameter is usually chosen randomly within Γ. Once the first parameter is attained we now have our reduced space as

$$X_1 = \text{span}\{p(\vartheta^1)\}.$$

For the rest of the parameters we seek them at which the error between the reduced basis solution and the true solution attains its maximum, however as before this optimisation procedure is quite costly. Instead we replace the true error $p(\vartheta) - p_N(\vartheta)$ by an error estimator $\Delta_N(\vartheta)$. We treat the error estimator as an upper bound for the reduced solution error. Now our parameter section is based on the following optimisation procedure for $N = 2, ..., N_{\max}$

$$\vartheta^N = \text{argmax}_{\vartheta\in\Xi_{\text{train}}}\Delta_{N-1}(\vartheta),$$

where now we seek our parameters within a subset $\Xi_{\text{train}} \in \Gamma$ of the parameter space known as the training set. The motivation behind using a training set, as mentioned, is namely improve on computational efficiency. Ideally the training set should be a good representative of the parameter space and be also cheap to evaluate. A more thorough discussion on training sets and various forms it can take are presented in Section 4.5. We present below in Algorithm 4 the RBM method.

Algorithm 4 Greedy-Reduced Basis Algorithm

For $N = 1$,

1. Choose $\vartheta^1 \in \Xi_{\text{train}}$ and set tolerance,

2. Compute $\Delta_1(\vartheta)$ for each $\vartheta \in \Xi_{\text{train}}$.

For $N = 2, \ldots, N_{\max}$,

3. Set $\vartheta^N = \text{argmax}_{\vartheta \in \Xi_{\text{train}}} \Delta_{N-1}(\vartheta)$.

4. Compute $p(\vartheta^N)$ via (4.4.1).

5. Build the reduced basis space $X_N = X_{N-1} \oplus \text{span}\{p(\vartheta^N)\}$.

A posteriori error bound

The efficiency of the reduced basis approximation is an important factor where one way to control this is via an a posteriori error bound Δ_N. We derive an error bound as follows:

let $R(q;\vartheta) \in X'$ be the residual in the dual space of X, which is defined as

$$R(q;\vartheta) := F(q) - A(p_N, q; \vartheta),\ \forall q \in X,$$

where $A(\cdot,\cdot;\vartheta)$ is the associated bilinear form and $F(\cdot)$ defining our right hand side. Through the Riesz-representation theorem, there exists a unique $\hat{e}(\vartheta) \in X$ such that $(\hat{e}(\vartheta), q)_X = R(q;\vartheta)$ and $\|\hat{e}(\vartheta)\|_X = \|R(\cdot;\vartheta)\|_{X'}$. We define the error between our true solution and our approximated reduced basis solution as

$$e(\vartheta) := p(\vartheta) - p_N(\vartheta),$$

which leads to the equation $A(e(\vartheta), q; \vartheta) = R(q;\vartheta),\ \forall q \in X$. Letting $q = e(\vartheta)$ and using Cauchy-Schwarz inequality we obtain

$$\begin{aligned}\kappa_{LB}(\vartheta)\|e(\vartheta)\|_X^2 &\leq A(e(\vartheta), e(\vartheta);\vartheta) = R(e(\vartheta);\vartheta)\\ &\leq \|R(\cdot;\vartheta)\|_{X'}\|e(\vartheta)\|_X\\ &= \|\hat{e}(\vartheta)\|_X\|e(\vartheta)\|_X,\end{aligned}$$

where $\kappa_{LB}(\vartheta)$ is the lower bound of the coercivity constant κ_{min}. From this we get

our bound as

$$\Delta_N := \frac{\|\hat{e}(\vartheta)\|_X}{\kappa_{LB}(\vartheta)}. \tag{4.4.3}$$

From the bound given in (4.4.3), both quantities $\|\hat{e}(\vartheta)\|_X$ and $\kappa_{LB}(\vartheta)$ are of interest to compute. Commonly for the lower bound an optimisation procedure is taken to calculate it, known as a successive constrained linear optimisation method [76]. For the estimate $\|\hat{e}(\vartheta)\|_X$, we compute this by an offline-online decomposition procedure.

Offline-online decomposition

As stated in the previous subsection, the offline-online procedure helps with estimating the a posteriori error bound by calculating the value $\|\hat{e}(\vartheta)\|_X$. We initiate the discussion of the method by recalling the reduced basis solution which is

$$p_N(\vartheta) = \sum_{m=1}^{N} p_{Nm}(\vartheta)\zeta_m.$$

Then by substituting this into (4.4.1) and choosing $\zeta_n = q$, for $1 \leq n \leq N$, we obtain the equation

$$\sum_{m=1}^{N} \left(A_0(\zeta_m, \zeta_n) + \sum_{k=1}^{K} \vartheta_k A_k(\zeta_m, \zeta_n) \right) p_{Nm}(\vartheta) = F(\zeta_n). \tag{4.4.4}$$

where we aim to solve for $p_{Nm}(\vartheta)$. From the above expression, we note that $F(\zeta_n)$, $A_0(\zeta_m, \zeta_n)$ and $A_k(\zeta_m, \zeta_n)$ are all independent of ϑ which we can precompute and store. This is known as the offline procedure. After the offline procedure we assemble the stiffness matrix of (4.4.4) and solve, this is known as the online procedure. However now with these stored components we still have to compute $\|\hat{e}(\vartheta)\|_X$ for each corresponding ϑ, therefore we expand the residual as

$$R(q;\vartheta) = F(q) - A(p_N, q; \vartheta) = F(q) - \sum_{n=1}^{N} p_{Nn} \left(\sum_{k=0}^{K} A_k(\zeta_n, q) \right).$$

By the Riesz representation theorem we can find Φ and $\Psi_n^k \in X$ where

$$\begin{aligned} (\Phi, q)_X &= F(q) \\ (\Psi_n^k, q)_X &= -A_k(\zeta_n, q), \end{aligned}$$

$\forall q \in X_N, 1 \leq n \leq N, 0 \leq k \leq K$. Substituting this into (4.4.4) and noting that $(\hat{e}(\vartheta), q)_X = R(q;\vartheta)$ leads to the following expression

$$\|\hat{e}(\vartheta)\|_X^2 = (\Phi,\Phi)_X + \sum_{k=0}^{K}\sum_{n=1}^{N} \vartheta_k p_{Nn}(\vartheta)\left(2(\Phi,\Psi_n^k)_X + \sum_{k'=0}^{K}\sum_{n'=1}^{N} \vartheta_{k'} p_{Nn'}(\vartheta)(\Psi_n^k,\Psi_{n'}^{k'})_X\right).$$

4.5 Training Set

As mentioned it is of importance to pick a sensible choice of a training set $\Xi_{train} \in \Gamma$ where ideally it should be a good representative of Γ and one that minimises the error. For the purpose of these experiments we decide to propose a new training set which are the Lebesgue optimal points (LOPs) whose motivation arises from a result presented in [35]. We wish to compare this training set to the Clenshaw-Curtis sparse grid points, which is a commonly used training set, in order to see the effectiveness of the LOPs. Before discussing some results we present an overview of both training sets and the motivation behind our proposed training set.

4.5.1 Clenshaw-Curtis Points

The Clenshaw-Curtis points arise as part of a quadrature rule. The goal of a quadrature rule is to approximate the integral $I[f]$ of some say continuous function f on $[-1,1]$. We are taking $[-1,1]$ for concreteness and since this is the interval (and its K-dimensional products) that we will be investigating in our numerical experiments.

Consider the N-th approximation of $I[f]$ given by $I_N[f] = \sum_{k=0}^{N} w_k f(x_k)$. The weights w_k here are chosen so that the expansion is exact for polynomials up to degree N. The Clenshaw-Curtis points are a particular choice for the evaluation points x_k. They are defined as follows:

Consider the N-th degree Chebyshev polynomial $T_N(x)$ defined by $T_N(\cos(\theta)) = \cos(N\theta)$. The Clenshaw-Curtis points (of degree N) are the extrema in $[-1,1]$ of $T_N(x)$ and they are given by the explicit formula:

$$x_n = \cos\left(\frac{n\pi}{N}\right) \quad n = 0,\ldots,N \tag{4.5.1}$$

A property of the Clenshaw-Curtis points that can be easily deduced from the explicit formula above is the asymptotic (for N large) clustering of points near the boundary (i.e. near -1 and 1). The following result gives error stability for the Clenshaw-Curtis quadrature:

If I_N is constructed using the Clenshaw-Curtis points then for any $f \in C[-1,1]$ we have the estimate:

$$|I[f] - I_N[f]| \leq 4\|f - p_N^*\|_\infty$$

where p_N^* is the best approximation for f in terms of degree n polynomials. Since $\|f - p_N^*\|_\infty \to 0$ as $N \to \infty$ the Clenshaw-Curtis quadrature does indeed converge to the actual value of the integral: $I_N[f] \to I[f]$.

Of course there are a number of other quadrature rules, for example Gauss quadrature. The relevant points here are the Gauss points which are the roots of the normalised Legendre polynomials. This quadrature rule (and others) satisfy similar (and even better in the case of Gauss quadrature) error stability estimates; so the reader might wonder what is the merit of using the Clenshaw-Curtis points which were introduced much later in the 1960s. The reason is the following computational one: using the Fast Fourier Transform one needs $\mathcal{O}(N\log(N))$ operations to calculate the Clenshaw-Curtis weights while Gauss quadrature needs $\mathcal{O}(N^2)$ operations. For a further discussion of these facts and more on Clenshaw-Curtis and Gauss quadrature see [37].

4.5.2 Sparse Grid

In uncertainty quantification one common problem that arises is the "curse of dimensionality" which says as the dimension size of the associated problem gets bigger so does the cost of the evaluation. This causes problems for numerical schemes when trying to work in a high dimensional setting. One way to overcome this issue is through the use of sparse grids [133].

Suppose we start off with one dimensional points $\{x_1, \ldots x_N\}$, we would like a way to build a collection $\mathcal{H}_{K,q}^s$ of K-dimensional points i.e. our sparse grid. Let $i = (i_1, \ldots, i_K)$ be a multi-index and consider the norm $|i|_{l_1} = \sum_{k=1}^K i_k$ and let $q \geq K$. Take a nested family of collections of points $\{\Theta^j\}$ so that $\Theta^j \subset \{x_1, \ldots x_N\}$, $\forall j$, $\Theta^j \subset \Theta^l$ if $l > j$ and $|\Theta^j| = 2^{j-1} + 1$. Now define the sparse grid in K dimensions (note that this depends on how we picked the nested family of points $\{\Theta^j\}$) as follows,

$$\mathcal{H}_{K,q}^s = \bigcup_{q-K+1 \leq |i|_{l_1} \leq q} \left(\Theta^{i_1} \times \cdots \times \Theta^{i_K} \right).$$

Due to this constrained building of our grid (the l_1 constraint on the multi-index does not allow us to choose too many high index/high cardinality families Θ^j) the

number of points is much less. It is important and of interest to note with sparse grids that despite the actual points in the sparse grid depend on how we chose $\{\Theta^j\}$ the cardinality of the sparse grid is independent of this choice.

In order to motivate out final proposed training set, we present the following theorem from the work of Chen et al. [35] which combines the the UQ techniques of the reduced basis method and the stochastic collocation method.

Theorem 4.5.1. *Provided that the training set Ξ_{train} for the reduced basis method is taken as the set of collocation points Θ, then*

$$\|p - p_N\|_{L^\infty(\Gamma;H_0^1(D))} \leq C\|p - L_M(p)\|_{L^\infty(\Gamma;H_0^1(D))}, \quad \forall N \leq M, \tag{4.5.2}$$

where $C = 3\kappa_{max}/\kappa_{min}$ is independent of N, $L_M(p)$ is the Lagrangian operator and κ_{max} and κ_{min} are the upper and lower ellipticity constants.

Proof. The proof can be found in [35]. □

What Theorem 4.5.1 is stating is that provided the training set used is a set of collocation points Θ, the RBM will perform just as good if not better as the stochastic collocation method. Of course one has to propose a set of optimal points than can be described as a set of collocation points. Our points we aim to use in this context are the Lebesgue optimal points. In order to describe them we review both the stochastic collocation method and the Lagrangian interpolation problem which in turn discuss the Lagranian operator $L_M(p)$, which appears in Theorem 4.5.1.

4.5.3 Stochastic Collocation Method

We now introduce the univariate case of the stochastic collocation Method (SCM). Further details can be found in [9]. For $K = 1$, the 1-dimensional stochastic collocation method approximates a solution $p(\vartheta) \in H_0^1(D)$, $\vartheta \in \Gamma$, using the Lagrangian interpolant. Given any set of collocation points $\Theta = \{\vartheta^0 < \vartheta^1 < \ldots < \vartheta^M\} \subset \Gamma$ and the corresponding snapshots $\{p(\vartheta^n)\}_{n=0}^M$, define the Lagrangian operator,

$$\left(L_M(p)(\vartheta)\right)(x) = \sum_{n=0}^{M} \left(p(\vartheta^n)\right)(x) l^n(\vartheta), \quad \vartheta \in \Gamma, \quad x \in D, \tag{4.5.3}$$

where $\{l^n\}_{n=0}^M$ are the Lagrangian characteristic polynomials, $l^n(\vartheta^k) = \delta_{nk}$, $0 \leq n, k \leq M$, given by the formula,

$$l^n(\vartheta) = \prod_{m \neq n} \frac{\vartheta - \vartheta^m}{\vartheta^n - \vartheta^m}, \quad 0 \leq n \leq M. \tag{4.5.4}$$

If $K > 1$, then by using tensor products or sparse grids we can extend the definition of the Lagrangian operator in K dimensions. Note that this gives us interpolation by tensor product polynomials.

4.5.4 Lebesgue Optimal Points

We next describe the Lagrange interpolation problem and its relation to the stochastic collocation method. Let $\Pi_\beta^K = \{\chi : \mathbb{R}^K \to \mathbb{R} | \chi \text{ polynomial with } deg(\chi) \leq \beta\}$, the set of polynomials on $\mathbb{R}^K$ with total degree at most β and $q = dim(\Pi_\beta^K) = \binom{\beta+K}{K}$. Let $X_{\beta,K} = \{\xi_j\}_{j=1}^q \subset \Gamma$, where Γ is a compact subset of $\mathbb{R}^K$. Then, the Lagrange Interpolation problem is the following:

For $f \in C^0(\Gamma)$, find a polynomial $\chi_f \in \Pi_\beta^K$, such that

$$\chi_f(\xi_j) = f(\xi_j), \quad j = 1, 2, \ldots, q.$$

The set $X_{\beta,K}$ is said to be unisolvent if such a unique polynomial exists for each $f \in C^0(\Gamma)$ and we call χ_f the Lagrangian interpolant of f. In general, not every set $X_{\beta,K}$ is unisolvent but if the Vandermonde determinant formed by the points $\{\xi_j\}_{j=1}^q$ is non-zero the existence of the Lagrangian interpolant is guaranteed. More specifically, if the set $X_{\beta,K}$ is unisolvent the explicit formula for the Lagrangian interpolant of f is given by

$$L_q(f(x)) = \sum_{n=0}^{q} f(\xi_n) l^n(x), \tag{4.5.5}$$

where $\{l^n\}_{n=1}^q$ are the Lagrangian Characteristic polynomials $l^n(\xi_j) = \delta_{nj}$, $n, j \in \{1, \ldots, q\}$, which form a basis for Π_β^K and it can be written as the Vandermonde determinants

$$l^n(x) = \frac{\det[V(\xi_1, \ldots, \xi_{n-1}, x, \xi_{n+1}, \ldots, \xi_q)]}{\det[V(\xi_1, \ldots, \xi_q)]}, \quad n = 1, \ldots, q,$$

where $V \in \mathbb{R}^{q \times q}$ such that $(V(\lambda_1, \lambda_2, \ldots, \lambda_q))_{jn} = l^n(\lambda_j)$ (In general, instead of l^n, we can use any other basis of Π_β^K). Notice that in the 1-dimensional case the expression above is actually (4.5.4).

We would like now to find a set $X_{\beta,K}$ which gives the lowest upper bound on the interpolation error. To find such a set we first have to find an upper bound of the error which applies to all $f \in C^0(\Gamma)$. If we denote by χ_f^* the best approximation of f in Π_β^K in the $\|\cdot\|_\infty$ then $L_q(\chi_f^*) = \chi_f^*$, thus

$$\|f - L_q(f)\|_\infty \leq \|f - \chi_f^*\|_\infty + \|L_q(\chi_f^*) - L_q(f)\|_\infty \leq (1 + \lambda_{q,K})\|f - \chi_f^*\|_\infty,$$

where

$$\lambda_{q,K} := \|L_q\|_\infty = \max_{x\in\Gamma} \sum_{n=1}^{q} |l^n(x)|.$$

The upper bound above is called the Lebesgue constant of the set $X_{\beta,K}$ and we easily see that the optimal choice $X_{\beta,K}^*$ of such a set is one that solves the following optimisation problem:

$$X_{\beta,K}^* = \operatorname*{argmin}_{\{\xi_j\}_{j=1}^q} \max_{x\in\Gamma} \sum_{n=1}^{q} |l^n(x)|. \tag{4.5.6}$$

We know that such an optimal set exists but is not unique in general. However, in [6] there is an extended description of the structure of the optimal set $X_{\beta,K}^*$, especially in the case of the cube $[-1,1]^K$, which is the state space we are concerned with.

Remark 4.5.1. *We have seen that stochastic collocation method constructs the Lagrangian polynomials using any set of collocation points and then approximates the solution as a linear combination of these polynomials based on the set of snapshots. On the other hand, the solution to the Lagrange interpolation problem provides us with a set of collocation points which are "universally" the best choice of reducing the error to such an approximation. Thus, a priori, the solution of (4.5.6), seems to be a promising choice of a collocation set.*

However, there is a slight difference between the two problems, which lies on the space where the Lagrangian operator acts. In (4.5.3) L_N acts on $L^\infty(\Gamma, H_0^1(D))$ while in (4.5.5) L_q on $L^\infty(\Gamma)$.

To convince ourselves that this is not as terrible as it seems, we only have to observe two important facts. First, the minimisation problem (4.5.6) minimises the Lebesgue constant $\lambda_{q,K}$, which is the norm of the operator L_q, depending only on the set Γ. Second, the operator norm of L_N is also a Lebesgue constant. More specifically, we have the following theorem:

Theorem 4.5.2. *Let $L_N : L^\infty(\Gamma; H_0^1(D)) \to L^\infty(\Gamma; H_0^1(D))$, where Γ is any com-*

pact subset of $\mathbb{R}^K$*, be the Lagrangian operator. Then,*

$$\|L_N\| = \max_{\vartheta\in\Gamma} \sum_{n=0}^{N} |l^n(\vartheta)|.$$

Proof. Let $p \in L^\infty(\Gamma, H_0^1(D))$ such that $\|p\|_{L^\infty(\Gamma,H_0^1(D))} \leq 1$. Then for $\vartheta \in \Gamma$

$$\begin{aligned}
\|L_N(p)(\vartheta)\| &= \|\sum_{n=0}^{N} l^n(\vartheta)p(\vartheta^n)\|_{H_0^1(D)} \\
&\leq \|\sum_{n=0}^{N} |l^n(\vartheta)|p(\vartheta^n)\|_{H_0^1(D)} \\
&\leq |l^n(\vartheta)|,
\end{aligned}$$

since $\|p(\vartheta^n)\|_{H_0^1(D)} \leq \|p\|_{L^\infty(\Gamma,H_0^1(D))}$, for all n. Thus, taking the maximum over all $\vartheta \in \Gamma$ we have

$$\|L_N\| \leq \max_{\vartheta\in\Gamma} \sum_{n=0}^{N} |l^n(\vartheta)|.$$

For the opposite direction, first we notice that the function $g : \Gamma \to \mathbb{R}$ is continuous so that is $\vartheta_0 \in \Gamma$ such that $g(\vartheta_0) = \max_{\vartheta\in\Gamma} \sum_{n=0}^{N} |l^n(\vartheta)|$, where we use the fact that Γ is compact. Consider a function $\phi \in C_c^\infty(D)$ where $\|\phi\|_{H_0^1(D)} = 1$. Define $p \in L^\infty(\Gamma; H_0^1(D))$ such that $p(\vartheta^n) = sign(l^n(\vartheta_0)) \cdot \phi$ for all $n \in \{0, \ldots, N\}$ and $p(\vartheta) = 0 \in H_0^1(D)$ for all $\vartheta \in \Gamma \setminus \{\vartheta^n\}_{n=0}^N$. From this we have

$$\begin{aligned}
\|L_N(p)\vartheta_0\|_{H_0^1(D)} &= \|\sum_{n=0}^{N} |l^n(\vartheta_0)| \cdot \phi\|_{H_0^1(D)} \\
&= \sum_{n=0}^{N} |l^n(\vartheta_0)| \cdot \|\phi\|_{H_0^1(D)} \\
&= \max_{\vartheta\in\Gamma} \sum_{n=0}^{N} |l^n(\vartheta)|
\end{aligned}$$

the proof is now complete, since

$$\|L_N(p)\|_{H_0^1(D)} \geq \|L_N(p)\vartheta_0\|_{H_0^1(D)}.$$

□

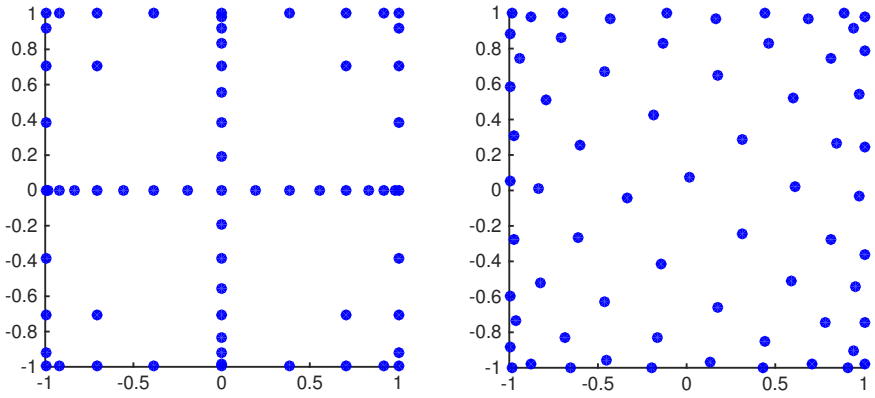

Figure 4.1: Left: spare grid training set in $[-1,1]^2$. Right: LOPs training set in $[-1,1]^2$

4.5.5 RBM Numerics

We now present a 2D example of using the RBM to solve a partial differential equation, specifically to start we will test the method on the elliptic partial differential equation with Dirichlet boundary conditions.

$$-\nabla \cdot (\kappa(x;\vartheta)\nabla p(x;\vartheta)) = f, \quad \in D, \tag{4.5.7a}$$

$$p(x;\vartheta) = 0, \quad \in \partial D, \tag{4.5.7b}$$

Our domain for all of these experiments unless stated otherwise will be $D = [0,1]^2$. For our random coefficient $\kappa(x;\vartheta)$ we represent this by an affine representation given as

$$\kappa(x;\vartheta) = \kappa_0(x) + \sum_{k=1}^{K} \kappa_k(x) u_k(\vartheta), \tag{4.5.8}$$

where we choose $\kappa_0(x) = 2$ and $\kappa_k(x) = \frac{1}{2^k} \cdot \sin(2\pi k x)$ as our basis functions. From equation (4.5.8) $\{u_k\}_{k=1}^K$ are independent random variables on Ω taking values in $[-1,1]$. As a result our lower and upper bounds for our random coefficient are such that $\kappa_{min} = 1$ and $\kappa_{max} = 3$. High differences in this values can lead to certain instabilities hence our choice for the values. For our true solution $p(\vartheta)$ to (4.5.7) we use a piecewise linear finite element method (FEM) method with mesh size $h = 1/40$. For our right hand side we set $f = 1$. Our training set Ξ_{train} for this experiment will be the LOPs in 1D and 9D and sparse grid points in 1D and 9D

with a tolerance level $TOL = 10^{-9}$. A visual comparison of these training sets can be found in Figure 4.1 where we take $\Gamma = [-1, 1]^K$. For $K = 1$ we set the number of points to be 11 for both training sets. Our motivation for this is to understand better the performance of the LOPs, which we will showcase through further visual representations. For the case of $K = 9$ we set the number of points higher to 50 points, aiming to test better the efficiency of the both training sets. Instead of using the the successive constrained linear optimisation method as described in [76], for simplicity we explicitly calculate the lower bound of our function which is a sine function.

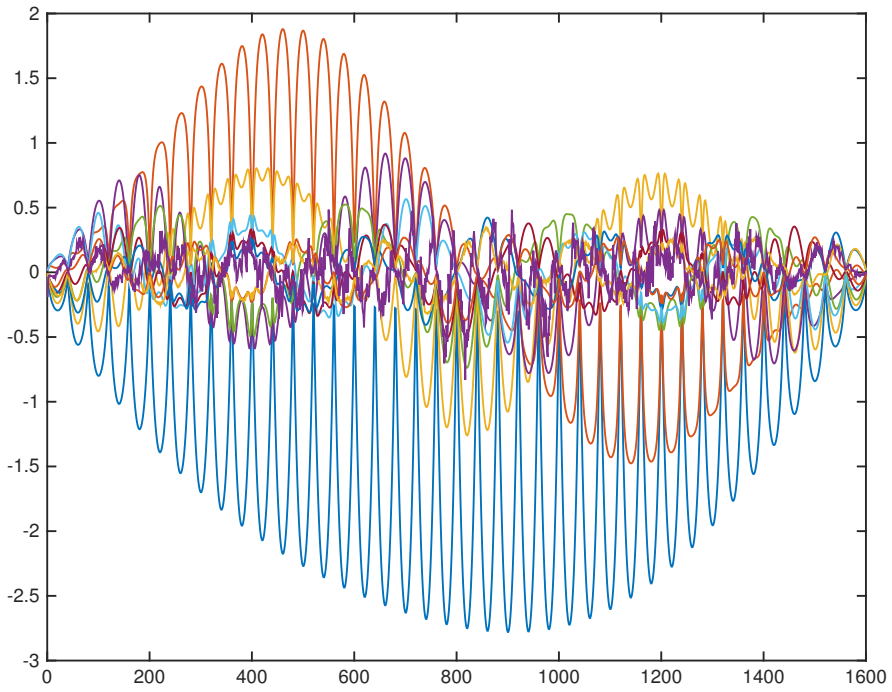

Figure 4.2: RBM numerics. Orthonormalized snapshots for LOPs in $[-1, 1]^1$.

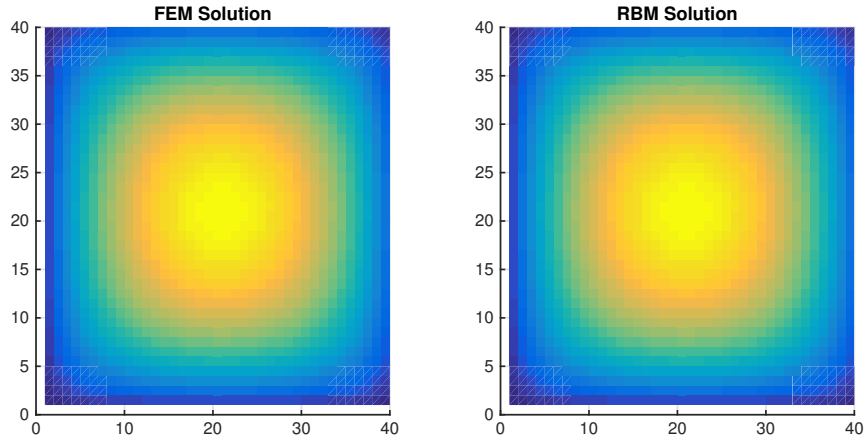

Figure 4.3: RBM numerics. Left: FEM solution. Right: RBM solution.

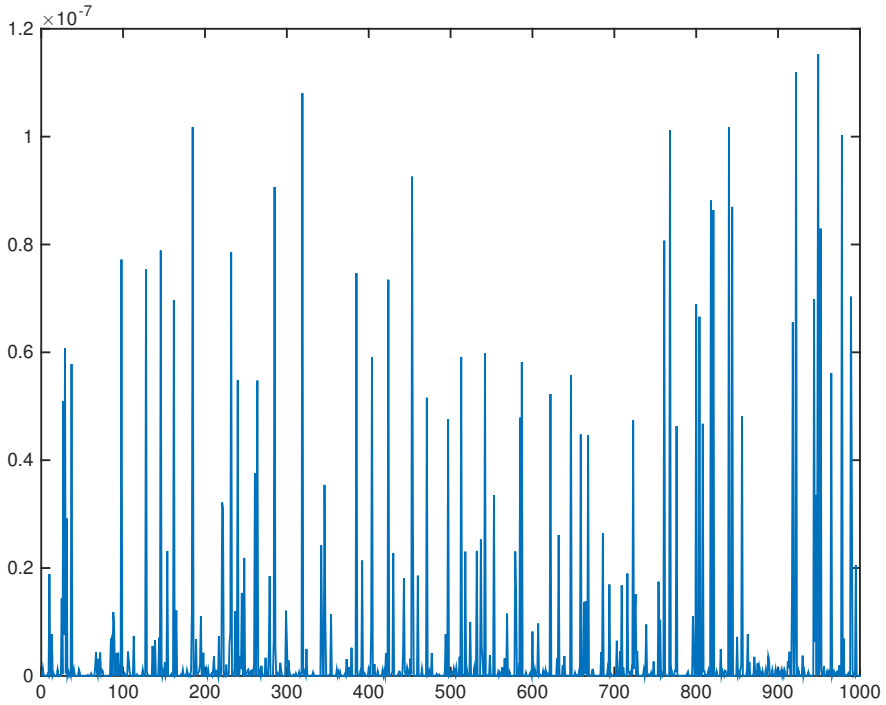

Figure 4.4: RBM numerics. Associated errors for Lebesgue optimal points in $[-1,1]^1$.

These experiments were conducted in Matlab with a processor of 2.6 GHz Intel Core i5. We first observe that from Figure 4.3 that both the RBM and FEM solutions corresponding to the elliptic PDE are approximately the same. For this

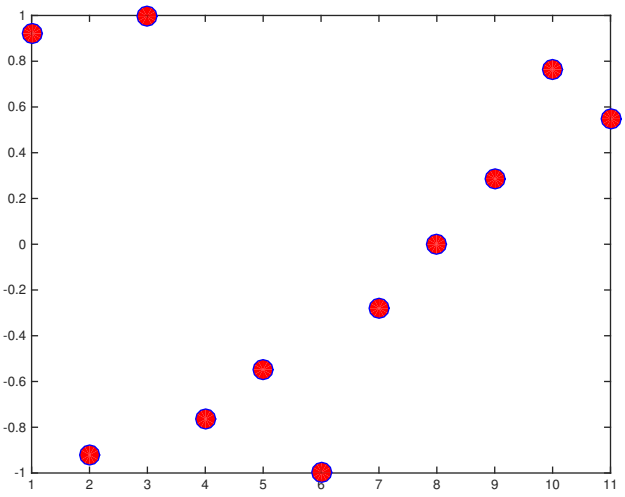

Figure 4.5: RBM numerics. Parameter selection for Lebesgue optimal points in $[-1,1]^1$.

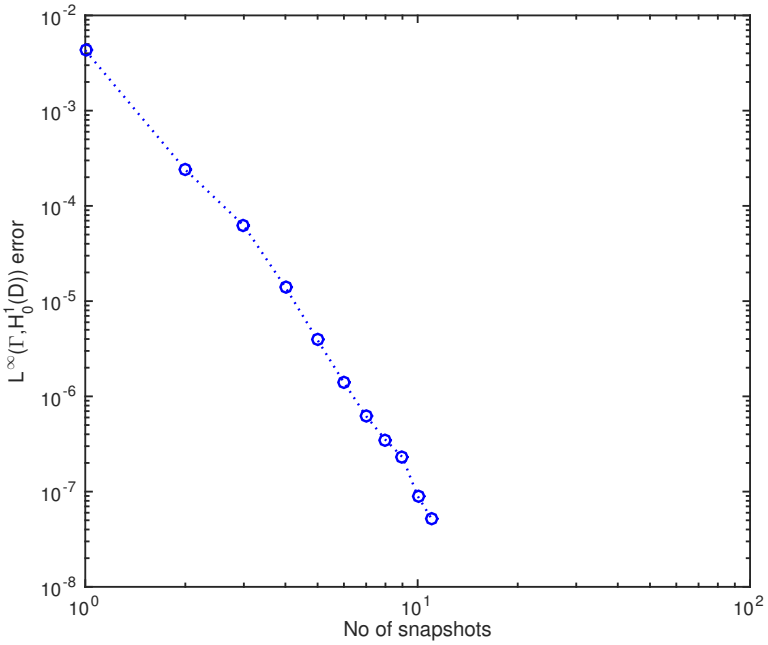

Figure 4.6: RBM numerics. Errors with respect to our snapshots in $[-1,1]^1$.

plot the 1D LOPs were used as the training set. The errors associated with the LOPs are shown in Figure 4.4. More noticeably for the final RB error, we have the error corresponding to each snapshot provided in Figure 4.6 and the corresponding

Training set	Dimension	Time (seconds)	$L^\infty(\Gamma, H_0^1(D))$
Lebesgue optimal points	1	757	1.142×10^{-7}
Lebesgue optimal points	9	1193	1.104×10^{-9}
Sparse grid	1	892	1.234×10^{-7}
Sparse grid	9	1716	2.428×10^{-8}

Table 4.1: Performance of different training sets.

snapshots in Figure 4.5. Simultaneously we also ran the RBM for the 1D sparse grid points. An analysis of the performance is presented in Table 4.3.

We see in the 1D case that the $L^\infty(\Gamma, H_0^1(D))$ errors are different by a factor of 10%, despite the time taken to be similar. Going from dimension 1 to 9 with respect to the parameter space we see a considerable difference for the errors. However more noticeably we see a substantial increase in time for the sparse grid points, but not such a dramatic increase for the optimal points. This was tested with 50 points in the parameter space. It can be said that from both training sets in dimension $K = 1$ they perform almost identically, but for the case of $K = 9$ we see a slight improvement in terms of both accuracy and computational efficiency.

4.6 Bayesian Inverse Problems

As we have discussed the potential of using the RBM over finite element methods, we now in this section aim to couple our new forward solver within inverse problems. Inverse problems require repetitive evaluations of the forward problem, therefore by hopefully implementing the RBM we can acquire a cheaper evaluation whilst maintaining similar accuracy. Our inverse solver for this chapter will be based on the iterative Kalman method. Before we present the algorithm we need to define our inverse problem in a Bayesian setting which will take a similar form as done in [137] but now with a dependence on our parameters ϑ. Following on from this we then given an overview of the iterative Kalman method and some of it's key properties before defining the method. We now begin this section with a quick review of the Bayesian approach towards inverse problems.

Given two separable Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$ and a forward operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$ we are interested in the recovery of the quantity of interest $u \in \mathcal{X}$ from noisy observations $y \in \mathcal{Y}$ which are given by

$$y = \mathcal{G}(u) + \eta, \tag{4.6.1}$$

where $\eta \sim \mathcal{N}(0, \Sigma)$ with Σ denoting a positive self-adjoint operator. Trying to invert (4.6.1) can cause difficulty as there is no guarantee with well-posedness through the classical approach. One way to alleviate this is through the Bayesian approach where now we are interested in the posterior distribution μ^y on the random variable $u|y$ which can be evaluated through Bayes' Theorem. From this we can characterise the posterior as

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z}\exp\big(-\Phi(u;y)\big), \tag{4.6.2}$$

with Z being

$$Z = \int_{\mathcal{X}} \exp\big(-\Phi(u;y)\big)\mu_0(du), \tag{4.6.3}$$

where our misfit functional is given as $\Phi(u;y) = \frac{1}{2}\big|y - \mathcal{G}(u)\big|^2_{\Sigma}$.

4.6.1 Iterative Kalman Method

Our inverse solver for this chapter will be the iterative Kalman method, which was proposed by Iglesias et al. [79] as an optimisation based technique to produce stable solutions to constrained PDE problems. The method can be derived from the least-squares formulation, which takes motivation from data assimilation. The main idea behind the method is to represent a noise controlled system i.e. $u \to u^\dagger$ as $\eta \to 0$. Assume we have an ensemble of J members $\big\{u_0^{(j)}\big\}_{j=1}^J \subset \mathcal{X}$. We label this as our initial ensemble at iteration level 0 which is said to be a linear space of our solution space $\mathcal{X}$. We wish to build upon our ensemble $\big\{u_n^{(j)}\big\}_{j=1}^J$ which at each iteration level n is updated through by combining the artificial dynamics with artificial data y_n, resulting in a new ensemble $\big\{u_{n+1}^{(j)}\big\}_{j=1}^J$ which is achieved by using the ensemble mean

$$\bar{u}_n = \frac{1}{J}\sum_{j=1}^{J} u_n^{(j)}, \tag{4.6.4}$$

to approximate the solution of the inverse problem.The iterative Kalman method can be split into two parts, a prediction step and an analysis step similar to the ensemble Kalman filter [57]. The purpose of the prediction step is to map the ensemble of particles into the observational space $\mathcal{Y}$ implying information is introduced into the forward model.

The analysis step takes the mapped ensemble in the data space and compares it with the data where the the ensemble is modified to better match the data. As stated previously the scheme attains regularisation properties, this is achieved through the discrepancy principle. In order to define the principle a regularisation parameter $\tau > \frac{1}{\rho}$ is usually introduced where $\rho \in (0,1)$ and from this the discrepancy

Algorithm 5 Regularized Iterative Kalman Method

Let $\{u_0^{(j)}\}_{j=1}^J \subset \mathcal{X}$ be the initial ensemble with J elements. Let $\rho \in (0,1)$ with $\tau > \frac{1}{\rho}$

We wish to generate,

$u_0^{(j)} \sim \mu_0, \quad y^{(j)} = y + \eta^{(j)}, \quad \eta^{(j)} \sim N(0, \Sigma).$

Then for $n = 1, \ldots$

Prediction Step

1. Evaluate the forward map,

$w_n^{(j)} = \mathcal{G}(u_n^{(j)})$ for $j \in \{1, \ldots, J\}$,

and define $\bar{w}_n = \frac{1}{J} \sum_{j=1}^J w_n^{(j)}$.

Discrepancy Principle

2. If $\|\Sigma^{-1/2}(y - \bar{w}_n)\|_{\mathcal{Y}} \leq \tau\eta$, stop!

Output $\bar{u}_n = \frac{1}{J} \sum_{j=1}^J u_n^{(j)}$.

Analysis Step

3. Define sample covariances:

$C_n^{ww} = \frac{1}{J-1} \sum_{j=1}^J (\mathcal{G}(u_n^{(j)}) - \bar{w}_n) \langle \mathcal{G}(u_n^{(j)}) - \bar{w}_n, \cdot \rangle_Y.$

$C_n^{uw} = \frac{1}{J-1} \sum_{j=1}^J (u_n^{(j)} - \bar{u}_n) \langle \mathcal{G}(u_n^{(j)} - \bar{w}_n, \cdot \rangle_Y.$

Update each ensemble member as follows

$u_n^{(j+1)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \alpha_n \Sigma)^{-1}(y^{(j)} - w_n^{(j)}),$

where $\alpha_n \equiv \alpha_n^N$ satisfies

$\rho \|\Sigma^{-1/2}(y^{(j)} - \bar{w}_n)\|_Y \leq \alpha_n^N \|\Sigma^{1/2}(C_n^{ww} + \alpha_n^N \Sigma)^{-1}(y^{(j)} - \bar{w}_n)\|_Y,$

and where α_n is chosen based on

$\alpha_n^{i+1} = 2^i \alpha_n^0.$

principle is given as

$$\|\Sigma^{-1}(y - \bar{w}_n)\|_{\mathcal{Y}} \leq \tau\eta, \tag{4.6.5}$$

where $\bar{w}_n = \frac{1}{J}\sum_{j=1}^{J} \mathcal{G}(u_n^{(j)})$. Usually with iterative inverse solvers it is common to add some regularisation. We note that for all the experiments we will work with the regularized iterative version which is given by Algorithm 5. EKI uses regularisation properties taken from the Levenburg-Marqardt scheme [71]. We note also that this discrepancy principle is slightly different to typical one as it is applied to the average of the output $\bar{w}_n$. We will now refer to inversion as ensemble Kalman inversion (EKI).

4.6.2 RB-EKI

As we have discussed both the forward and inverse solver in detail we will now present the coupled scheme. We begin the development of the algorithm by reformulating the inverse problem with the RBM.

As we have defined our parametric PDE of interest, we can use this to formulate a Bayesian inverse problem. Recall that our solution $p(\cdot;\vartheta) \in X := H_0^1(D)$ and $l_j \in X^*$ are continuous linear functionals. Then we can define our observed data

$$y_j = l_j(p(\cdot;\vartheta)) + \eta_j, \quad j = 1, \ldots, J, \tag{4.6.6}$$

where $\{\eta\}_{j=1}^{J} \sim N(0, \Sigma)$ is Gaussian additive noise. From this we can further define our forward operator $\mathcal{G} : \mathbb{R}^k \to \mathbb{R}^J$ where

$$\mathcal{G}_j(u) = l_j(p(\cdot;\vartheta)) = \begin{cases} p(x_j) & \in D = [0,1] \\ \displaystyle\int_D p(x_j) g_j(x) dx & \in D = [0,1]^2, \end{cases}$$

where $g_j(x)$ is a covariance kernel for our inverse problem. This allows us to rewrite (4.6.6) as the inverse problem

$$y = \mathcal{G}(u) + \eta.$$

Note here we define our domains to be $D = [0,1]^d$ for $d = 1, 2$ as these are the specific domains we will be using for the proceeding numerics.

4.7 Numerical Results

4.7.1 Uniform Prior

From our experiments in subsection 4.5.5 our random coefficient was based on the expansion (4.5.8). For the RB-EKI our random coefficient will be based on a prior which is that of a geometric type. These priors will have the general form of

$$u = \phi_0 + \sum_{j=1}^{\infty} u_j \phi_j, \tag{4.7.1}$$

where $u \equiv \kappa$. (4.7.1) consists of random functions $u_j = \gamma_j \xi_j$ for $j = 1, \ldots, \infty$, where $\xi = \{\xi_i\}_{j=1}^{\infty}$ is an i.i.d. sequence with $\xi_1 \sim \mathcal{U}[-1, 1]$ and $\{\gamma\}_{j=1}^{\infty} \in \ell^1$. We note that $\{\phi_j\}_{j=1}^{\infty} \in L^{\infty}(D)$ is an infinite sequence. Note the form of the random coefficient in (4.7.1) is a general case, for our experiments we will use a truncated expansion.

4.7.2 Single Phase 2D Prior

Recalling the RB experiments conducted for the forward problem had a random coefficient that was based on a form given in (4.7.1). Now that we are working towards implementing the RBM within the inverse solver we need to define the form of the random coefficient as a prior μ_0. In order to do so we have to remain consistent with the assumptions of the random coefficient; that it has an affine form which is independent of the parameters $\mu = (\mu_1, \ldots, \mu_k)$.

In the RBM literature it is common to choose the random coefficient either uniformly or log-normally. Modifying the coefficient based on this can change the setup of the RBM, due to this we continue to use to assume a uniformly distributed random coefficient. In the context of groundwater flow recent priors that have been developed by Iglesias et al. [80] have showcased to perform well which are based on channelized flow. This is based on some non-linearities. For our prior will we use a modification of the channel flow prior defined in [80] neglecting the non-linearities.

In order to define our channelized flow we have two equations which govern the channel which are defined as

$$t_1 = \frac{x}{d_3} + d_1, \tag{4.7.2}$$

$$t_2 = \frac{x}{d_3} + d_1 + d_2. \tag{4.7.3}$$

From equations (4.7.2) and (4.7.3) we have three main parameters within our prior. The first being $d_3 \in \mathbb{R}^+$, which can be though of as determining the steepness of the

channel, $d_1 \in \mathbb{R}^+$ which defines the initial point and $d_2 \in \mathbb{R}^+$ defining the height of the channel. As well as the channel parameters we also have the the values of κ_1, κ_2 which are the values of the permeability in and out the channel. All the parameters within the model are distributed accordingly to a uniform which is provided in Table 4.2. By relating this to the assumptions on the random coefficient, we design our prior such that $\kappa_{min} = 1$ ad κ_{max} =5.5.

Parameter	Prior distribution
d_1	$\mathcal{U}[0, 0.5]$
d_2	$\mathcal{U}[0, 1]$
d_3	$\mathcal{U}[1, 20]$
κ_1	$\mathcal{U}[1, 1.5]$
κ_2	$\mathcal{U}[5, 5.5]$

Table 4.2: Prior associated with single phase flow.

The unknown parameter for this model is $u = (d_1, \ldots, d_3, \kappa_1, \kappa_2) \in \mathbb{R}^5$.

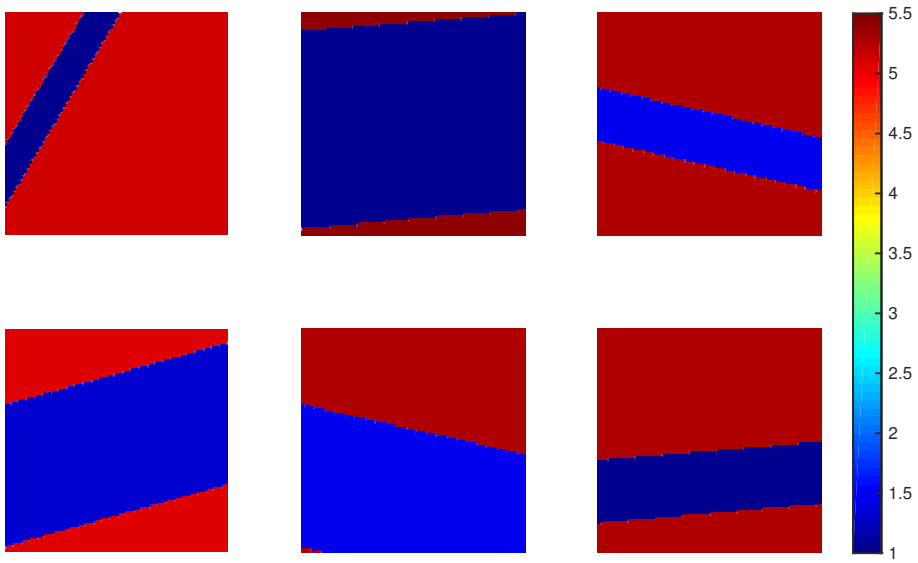

Figure 4.7: Random draws from single phase prior.

4.7.3 RB-EKI Numerics

We now look to test the RB-EKI method for the model problem (4.5.7). For the inclusion of the inverse solver we need to define our truth $u^\dagger$ which we stated would

be piecewise constant. Our truth is given below in Figure 4.8.

Our aim for choosing this particular prior form is in hope for a good reconstruction, hence why it is of a similar form to the truth. We are interested in an underdetermined system where we have 64 observational points uniformly spread. For the regularisation we choose $\rho = 0.8$ and $\tau \approx 1.25$. The number of ensemble members chosen is $J = 150$, with an iteration count of $n = 20$. Our noise $\eta \sim \mathcal{N}(0, \gamma^2 I)$ will be chosen such that $\gamma = 0.4$.

For the forward solver we initialise the numerics similarly as before where we have a mesh size of $h = 1/40$ where our training set Ξ_{train} will be based on the Lebesgue optimal points for $\Gamma = [-1, 1]^9$ with a tolerance level of $TOL = 10^{-9}$. As before we generate 50 points in our parameter space. Apart from the mean of the output u_{EKI} we are also interested in two other quantities of interest:

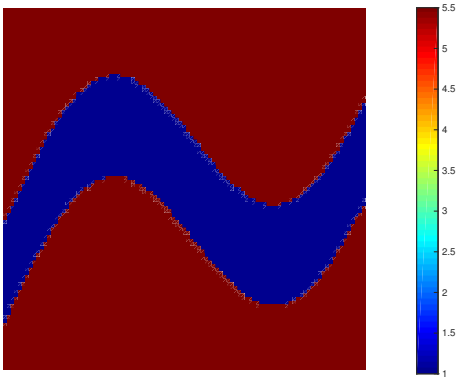

Figure 4.8: Channelized geometric truth $u^\dagger$.

- Relative error - $\frac{\|u^\dagger - u_{\text{EKI}}\|_{L^2(D)}}{\|u^\dagger\|_{L^2(D)}}$.
- Data misfit - $\left\|y - \bar{\mathcal{G}}(u_{\text{EKI}})\right\|_\Gamma^2$.

Method	Dimension	Time (seconds)
FEM-EKI	9	3773
RB-EKI	9	2287

Table 4.3: 2D RB-EKI numerics. Performance of the different iterative methods.

From the numerics conducted we gain an indication of the performance of RB-EKI in terms of both the mean reconstruction of the truth i.e. u_{EKI}, the difference in computation time and the two quantities of interest: the log data misfit

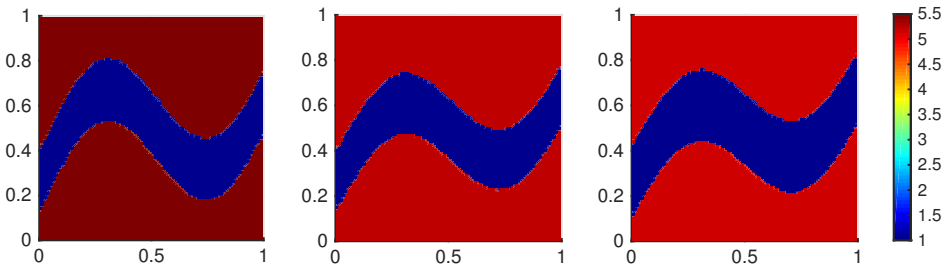

Figure 4.9: 2D RB-EKI numerics. Left: True permeability. Centre: Finite element reconstruction. Right: Reduced basis reconstruction.

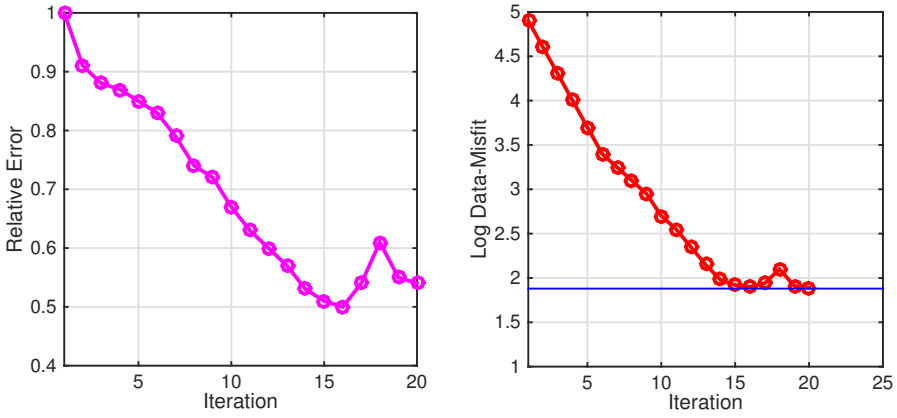

Figure 4.10: 2D RB-EKI numerics. Left: RB-relative error. Right: RB-log-data misfit.

and the relative error. Firstly regarding Figure 4.9 we see the performance of the RB-EKI which shows a relatively good recovery of the true permeability shown on the left hand side. We also see the reconstruction of the iterative method with a FEM where we see a similar performance to that of the RB-EKI. Despite the permeability levels being slightly off as well as the width of the channel the overall structure is recovered. This is aided by Figure 4.10 which demonstrates the effectiveness of the regularized properties within the iterative method. As the number of iterations increase we see a decline in both the data misfit and the relative error which terminates after the 20th iteration.

In terms of the error and the computational time when comparing both methods, we see an decrease in computational time needed with the inclusion of the RBM while showing similar errors with the FEM. This speed is up is significantly better than the results we were obtaining within the 1D elliptic problem.

4.8 Electrical Impedance Tomography

Another inverse problem that we are interested is that of impedance tomography. electrical impedance tomography, often referred to as EIT [21], is a common imaging technique concerned with measuring changes in the conductivity distribution inside a body. This is achieved by applying electrical currents to the electrodes which lie on the boundary, and applying Ohms' Law to record voltages.

The inverse problem is concerned with the recovery of the conductivity distribution from the voltage measurements of the electrodes. This inverse problem is both non-linear and ill-posed. The most common approach to model the forward model associated with EIT is referred to as the complete electrode model (CEM). The equations associated with the random-CEM are

$$\nabla \cdot (\kappa(\cdot;\vartheta)\nabla\nu(\cdot;\vartheta)) = 0, \quad \in \ D, \tag{4.8.1a}$$

$$\nu + z_l\kappa(\cdot;\vartheta)\nabla\nu(\cdot;\vartheta)\cdot n = V_l, \ \in \ e_l, \quad l = 1,\dots,m_e, \tag{4.8.1b}$$

$$\nabla\nu(\cdot;\vartheta)\cdot n = 0, \quad \in \ \partial D \backslash \cup_{l=1}^{m_e} e_l, \tag{4.8.1c}$$

$$\int \kappa(\cdot;\vartheta)\nabla\nu(\cdot;\vartheta)\cdot n \ ds = I_l, \ \in \ e_l, \quad l = 1,\dots,m_e. \tag{4.8.1d}$$

The log-conductivity distribution of the electrodes is denoted by κ, while ν is the electric potential from the electrodes $\{e_l\}_{l=1}^{m_e}$. From the model given above $\{I_l\}_{l=1}^{m_e}$ and $\{V_l\}_{l=1}^{m_e}$ are the currents and voltages associated with the electrodes while $\{z_l\}_{l=1}^{m_e}$ are the contact impedances. Further details on the derivation of the EIT model can be found in [21]. The forward solution associated with the CEM (4.8.1) is to find $(\nu, V) \in \mathbb{H}$ where $\mathbb{H} := H^1(D) \oplus \mathbb{R}^{m_e}$ which involves the m_e surface voltages and the electric potential in the interior. The CEM is based on Ohms' law which is used to construct the voltage i.e.

$$V(\kappa) = I(\kappa) \times R,$$

where R denotes the resistivity. In order to ensure well-posedness of the CEM we require conservation of charge i.e.

$$\sum_{l=1}^{m_e} I_l = 0.$$

As of yet there has been no extensive literature which has applied the RBM to solve the EIT problem. We emphasize again that the purpose of this section is to consider deriving the a posteriori bound for EIT where numerics are omitted from this work.

4.8.1 A Posteriori Bound

In order to test the RB-EKI for the EIT problem we need to derive an a posteriori error bound in a similar fashion to the elliptic PDE as done in Section 4.4. To achieve this we begin with defining the weak formulation of the CEM. From this we then look to use the definition of the residual and the elliptic structure of the PDE to obtain our bound. Given the CEM (4.8.1) we can define its weak formulation. For simplicity we stick to the non-random weak formulation, which can easily be extended to the random case. For the weak formulation we say that $(\nu, V) \in \mathbb{H}$ is a weak solution is for any $(q, Q) \in \mathbb{H}$ such that

$$\int_D \kappa \nabla \nu \cdot \nabla q \; dx + \sum_{l=1}^{m_e} z_l^{-1} \int_{e_l} (\nu - V_l)(q - Q_l) ds = \sum_{l=1}^{m_e} I_l Q_l, \tag{4.8.2}$$

where our bilinear form $A : \mathbb{H} \times \mathbb{H} \to \mathbb{R}$ and right hand side $F : \mathbb{H} \to \mathbb{R}$ are defined as

$$A(\nu_N, q; \vartheta) = \int_D \kappa \nabla \nu \cdot \nabla q \; dx + \sum_{l=1}^{m_e} z_l^{-1} \int_{e_l} (\nu - V_l)(q - Q_l) ds, \tag{4.8.3}$$

and

$$F(q) = \sum_{l=1}^{m_e} I_l Q_l.$$

We can prove existence and uniqueness of solutions to (4.8.2) via the Lax-Milgram lemma. However this poses difficulty as the bilinear form (4.8.3) is not coercive. From the expression

$$A((\nu, V), \nu, V)) = \int_D \kappa |\nabla \nu|^2 dx + \sum_{l=1}^{m_e} z_l^{-1} \int_{e_l} |\nu - V|^2 ds,$$

we see the difficulty in $A((\nu, V), \nu, V)) = 0$ implying $(\nu, V) = 0$, in other words solutions can only be defined up to the addition of a constant. Instead now we seek our solution $(\nu, V) \in \dot{\mathbb{H}}$ where $\dot{\mathbb{H}} = \mathbb{H}/\mathbb{R}$ is an alternative space with the induced norm

$$\|(\nu, V)\|_{\dot{\mathbb{H}}} = \inf_{c \in \mathbb{R}} \left(\|\nu - c\|_{H^1(D)}^2 + \|V - c\|_{\mathbb{R}^{m_e}}^2 \right)^{1/2}.$$

Now using the weak formulation defined by equation (4.8.2) we can define our RB Galerkin projection method as

$$A(\nu_N, q; \vartheta) = F(q), \; \forall q \in X_N. \tag{4.8.4}$$

Now that we have defined our bilinear form we can express this in an affine representation, as done with the general elliptic PDE, where

$$A(\nu, q; \vartheta) = A_0(\nu, q) + \sum_{k=1}^{K} A_k(\nu, q)\vartheta_k, \quad \vartheta_k \sim \mathcal{U}[-1, 1], \tag{4.8.5}$$

$$A_k(\nu, q) = \int_D \kappa_k \nabla\nu \cdot \nabla q \; dx + \sum_{l=1}^{m_e} z_l^{-1} \int_{e_l} (\nu - V_l)(q - Q_l) ds. \tag{4.8.6}$$

As we are working with the CEM our solution space is now defined as $X := \dot{\mathbb{H}} = \frac{H^1(D)\oplus\mathbb{R}^{m_e}}{\mathbb{R}}$. From this we can define the RB weak formulation for (4.8.1) as

$$A(\nu_N, q; \vartheta) = F(q), \; \forall q \in X_N. \tag{4.8.7}$$

Our residual corresponding to (4.8.7) is given as

$$r(q; \vartheta) = A(\nu_N, q; \vartheta) - F(q) \in X^{'}. \tag{4.8.8}$$

As similarly done in section 4.4 we can define our error between our true solution and our approximated reduced basis solution as $e(\vartheta) := \nu(\vartheta) - \nu_N(\vartheta) \in X_N$ which satisfies

$$A(e(\vartheta), q; \vartheta) = r(q; \vartheta), \; \forall q \in X_N. \tag{4.8.9}$$

As before through the Riesz representation Theorem there exists a $\hat{e}(\vartheta) \in X$ such that

$$(\hat{e}(\vartheta), q)_X := r(q; \vartheta), \; \forall q \in X_N. \tag{4.8.10}$$

Using equation (4.8.10) we can also write the error residual (4.8.9) as

$$A(e(\vartheta), q; \vartheta) = (\hat{e}(\vartheta); q)_X, \; \forall q \in X_N. \tag{4.8.11}$$

Using the same argument as in Section 4.4 it leads to our a posteriori error bound for the CEM as

$$\|e(\vartheta)\|_X := \|\nu(\vartheta) - \nu_N(\vartheta)\|_X \leq \Delta_N(\vartheta) = \frac{\|\hat{e}(\vartheta)\|_X}{\kappa_{LB}(\vartheta)} := \frac{\|r(\cdot; \vartheta)\|_{X'}}{\kappa_{LB}(\vartheta)}. \tag{4.8.12}$$

Note: with the new a posteriori error bound defined by (4.8.12), it looks identical to the one used for the groundwater flow problem, but with the key differences of the space X changing as well as the form of the residual (4.8.8).

4.8.2 Offline-Online Decomposition

With the newly defined error bound defined by (4.8.12) we now seek to evaluate $\|\hat{e}(\vartheta)\|_X$. In order to do so we begin, as before, by expanding the form of $\nu_N(\vartheta)$ as

$$\nu_N(\vartheta) = \sum_{m=1}^{N} \zeta_m \nu_{Nm}(\vartheta), \tag{4.8.13}$$

where $\{\zeta_m\}_{m=1}^N$ is the RB matrix. If we substitute the above expression into (4.8.4) and choosing $\zeta_n = q,\ 1 \leq n \leq N$, we obtain an equation where we aim to solve for $\nu_{N_m}(\vartheta)$

$$\sum_{m=1}^{N} \left(A_0(\zeta_m, \zeta_n) + \sum_{k=1}^{K} \vartheta_k A_k(\zeta_m, \zeta_n) \right) \nu_{Nm}(\vartheta) = F(\zeta_n). \tag{4.8.14}$$

From the expression (4.8.14) we can evaluate the quantities independently $A_k(\zeta_m, \zeta_n)$, $F(\zeta_n)$ of ϑ. In the RBM setting we treat this as the online procedure where we precompute the quantities and store them. In order to evaluate the a posteriori bound (4.8.12) we use the residual form defined as:

$$r(q;\vartheta) = F(q) - \sum_{m=1}^{N} \nu_{N_m}(\vartheta) \left(\sum_{k=0}^{K} \vartheta_k A_k(\zeta_m, q) \right), \quad \vartheta_0 = 1. \tag{4.8.15}$$

As done before to aid with the computation of $\|e(\vartheta\|_X^2$ we can use the Riesz representation theorem i.e. we can define Φ and Ψ_n^k as

$$(\Phi, q)_X := F(q), \qquad \forall q \in X_N \tag{4.8.16}$$

$$(\Psi_n^k, q)_X := -A_k(\zeta_n, q), \quad \forall q \in X_N, \quad 1 \leq n \leq N. \tag{4.8.17}$$

Recalling that $(\hat{e}(\vartheta), q)_X := r(q;\vartheta)$ and by substituting (4.8.16), (4.8.17) into (4.8.14) we get an expression for $\|\hat{e}(\vartheta)\|_X^2$ that reads

$$\|\hat{e}(\vartheta)\|_X^2 := (\Phi, \Phi)_X + \sum_{k=0}^{K} \sum_{n=1}^{N} \vartheta_k \nu_{Nn}(\vartheta) \left(2(\Phi, \Psi_n^k)_X + \sum_{k'=0}^{K} \sum_{n'=1}^{N} \vartheta_{k'} \nu_{Nn'}(\vartheta) (\Psi_n^k, \Psi_{n'}^{k'})_X \right).$$

4.8.3 Inverse Problem

As stated the inverse problem of the CEM is known as EIT which is interested in recovering the conductivity distribution from voltages $(V_l)_{l=1}^{M_e}$ on the boundary ∂D. We know our forward solution can be constructed via Ohm's Law therefore assuming

we have J linear independent current patterns $I^{(j)} \in \mathbb{R}^{m_e}$ and noisy measurements from $V(\kappa) = I(\kappa) \times R$ in the form:

$$y_j = V^{(j)} + \eta_j, \quad \eta_j \sim \mathcal{N}(0, \Gamma).$$

By the relationship between the operator $\mathcal{G}$ and $V^{(j)}$ we can rewrite our inverse problem as

$$y_j = \mathcal{G}_j(\kappa) + \eta_j,$$

where

$$\mathcal{G}_j(\kappa) = I^{(j)} \times R(\kappa). \tag{4.8.18}$$

Finally the inverse problem (4.8.18) can be expressed as

$$y = \mathcal{G}(\kappa) + \eta, \quad \eta \sim \mathcal{N}(0, \Gamma). \tag{4.8.19}$$

Via an application of Bayes' Theorem we can propose an existence and well-posedness theorem for the Bayesian approach to EIT. The proof of this and further information regarding the construction of the posterior can be found in [49] by Dunlop et al. By incorporating randomness our inverse problem is now reformulated to solving

$$y_j = \mathcal{G}_j(\kappa; \vartheta) + \eta_j,$$

where now

$$\mathcal{G}_j(\kappa; \vartheta) = I^{(j)} \times R(\kappa; \vartheta),$$

which can be rewritten in the same way as (4.8.19) is defined.

4.9 Conclusion

RBMs are a class of powerful forward solvers aimed at improving efficiency for parametric systems of equations. The goal of this chapter was to use ideas from RBMs and implement them within Bayesian inverse problems to improve on the computational burden of traditional solvers. In particular we proposed a new training set which were the LOPs, which were motivated from [35] which stated that a set of collocation points were an effective choice. By doing so we saw an improvement to our experiments while retaining a good level of accuracy. We transferred these ideas to solve the a geometric inverse problem for a 2D elliptic PDE. As before results were consistent and showcased a bigger reduction in time taken. At the end of the chapter we briefly considered the analysis of the RBM applied to the CEM, where we derived an a posteriori bound similar to that of groundwater flow. Despite the improvement of using the LOPs as a training set, there is still the limitation in that it can only be expressed up to a dimension size of 10 for the parameter space. Commonly the parameter space for the sorts of problems that were discussed are much bigger, which can incorporate more points.

There are many further investigations one can take from this, mainly conducting experiments for the RB-EKI for EIT. One reason why this was omitted was that Theorem 4.5.1 is specific for the function space $H_0^1(D)$ which is the not the solution space of (4.8.1). Another reason being that the task of implementing the RBM with the EIT is of course more complicated and issues could arise. However it is of interest to see from Section 4.8 that the posteriori bound (4.8.12) has an almost identical form to that of (4.4.3). Given how the numerics conducted were not best optimized, this poses an interesting question on how best to approach this. There is also the issue that a reduction in the error of the forward solver does not imply a reduction in the inverse solver. One possibly way of sorting this out could be a discrepancy principle which incorporates both errors for the forward and inverse solver. This has been done in the deterministic inverse problem framework [65]. As a final remark, we explicitly calculated the lower bound of the coercivity constant as we had our basis functions as a sine function. It would be interest to test this against the the optimization method in [76] which is commonly used to compute the lower bound.

Chapter 5

A Bayesian formulation of the inverse eikonal equation

5.1 Overview

The basis of this chapter is to apply the current literature on Bayesian inversion and UQ to new classes of PDEs. Our PDE we consider in this chapter is the eikonal equation. This chapter will differ from previous chapters as it is application based, however there are certain similarities such as using the hierarchical methodologies developed in Chapter 2 and 3. As before in previous chapters we present a review and brief introduction into Bayesian inversion and EKI.

5.2 Introduction

The eikonal equation [11, 28, 44, 95, 111] is a fundamental partial differential equation (PDE) that arises in numerous fields such as geosciences [109] and imaging [86, 136]. The eikonal equation can be related via electromagnetic potential, that describes the relationship between electrical potential and electrical strength. The interpretation of this for an electromagnetic example is that any charge in a region is pushed to move at right angles to the lines of constant potential that is determined through the the electric strength. Mathematically the eikonal equation is a well-studied model, with a mathematical interpretation of determining travel-time. More specifically it is concerned with calculating the time required $T(x)$ to travel from the boundary $x_0 \in \partial D$ to a point in the interior of the domain $x \in D$ given a

prescribed speed, known as the slowness function $u(x)$ where

$$\begin{cases} |\nabla T(x)| &= u(x) \quad x \in \Omega \setminus \{x_0\}, \\ T(x_0) &= 0. \end{cases}$$

The forward problem has been extensively studied both numerically and analytically. Well-posedness of the eikonal equation follows from the theory of HJ equations, which is achieved by formulating an optimal control problem. Numerically there has been significant research on this where a number of powerful and efficient computational methods have been developed, namely the fast marching method and the fast sweeping method. [130]. Now if we consider the inverse problem associated with the eikonal equation, i.e. recovering our unknown u from noisy measurements of y

$$y = \mathcal{G}(u) + \eta, \quad \eta \sim \mathcal{N}(0, \Gamma), \tag{5.2.1}$$

our quantity of interest is the slowness function. This problem can be attributed as essentially a gradient recovery problem, which has been looked at in a general sense [113, 149, 150]. Unlike the forward model, the inverse eikonal equation has not been as well studied. This was recently tackled in the work of Deckelnick et al. [52] where they showed well-posedness of the inverse solution through the aid of gamma-convergence. The inverse problem was analyzed in a deterministic framework, where they showcased numerics for a number of different piecewise constant truths. Beyond this further work was done in [142] where the inverse problem of travel-time tomography was analyzed. This chapter primarily looked at implementing a modified fast marching algorithm where they used the eikonal equation as the forward model. However it still remains of interest to see how the slowness function can be recovered under various scenarios, due to such constraints as positivity and continuity. This body of work is concerned with understanding the behaviour of the slowness function and tackling the level of uncertainty that can arise. In order to do we propose an alternative viewpoint which is to adopt a Bayesian formulation of the inverse eikonal equation.

By adopting a Bayesian approach for our inverse problem (5.2.1) we now seek a probabilistic distribution of the random variable $u|y$ known as the posterior. The Bayesian formulation of inverse problems has seen a significant body of recent work since the formulation of Bayesian inverse problems in infinite dimensions [137]. Due to this instead of a distribution we can construct a posterior measure μ^y in the

form of a Gaussian measure with Radon-Nikodym derivative given as

$$\frac{d\mu^y}{d\mu_0} = \frac{1}{Z}\exp(-\Phi(u;y)), \tag{5.2.2}$$

such that

$$\begin{aligned} Z &:= \int_{\mathcal{X}} \exp(-\Phi(u;y))\mu_0(du), \\ \Phi(u;y) &= \frac{1}{2}|y - \mathcal{G}(u)|^2_\Gamma, \end{aligned} \tag{5.2.3}$$

where denote the normalizing constant and misfit functional. As a result the Bayesian approach has been readily applied for numerous PDEs. Specifically there has been a wide development of computational methods based on Markov chain Monte Carlo (MCMC) theory [15, 80, 81] and ensemble based methods [33, 77, 78]. Much of the work done has been looking at developing priors and techniques which handle geometric features of the unknown. This includes both geometric priors and level set techniques. More recently there has been an interest in understanding inverse problems in a hierarchical manner [6, 33, 47]. These approaches are concerned with understanding the unknown and its hyperparameters that define it.

Our motivation for analysing the inverse eikonal equation is to exploit the recent advancements made in the Bayesian approach for inverse problems, looking to extend from the current literature [52, 66]. Specifically our aim is to understand the uncertainty that can arise by considering a wide array of priors and to establish well-posedness of the inverse problem. The level set method [25] is a popular technique at solving interface based problems which are usually concerned with models governed by some geometry. Recently there has been a method which has applied the level set method to Bayesian inverse problems, the Bayesian level set method [81]. This method poses advantages over traditional level set methods for deterministic problems. By allowing for a Bayesian setting we can further pose geometric constraints as prior information. This has shown to be successful at reconstructing geometric truths within inverse problems [80]. In this chapter we will consider a number of various prior ranging from piecewise deterministic priors, to level set priors and finally priors based on Whittle-Matérn Gaussian random fields [124]. Our inverse solver for this problem will be an ensemble-based method, the iterative ensemble Kalman method [79]. The core advantage of using these over its MCMC counterparts is computational cost, where they pose a significantly lower computational burden. This has shown to be promising with applications in PDEs such as electrical impedance tomography and Darcy flow. We emphasize that our

work will comprise mainly of a numerical investigation, with regards to introducing a Bayesian approach to the inverse problem. Our contributions to this piece of work are as follows:

- To the best of our knowledge this is the first piece of work to consider a Bayesian formulation of the inverse eikonal equation.
- We propose a number of prior forms to represent the slowness function ranging from a continuous Whittle-Matérn Gaussian prior to discontinuous priors which have some form of geometry attributed to them such as fixed shape priors and level set priors.
- We exploit the recent hierarchical developments made for ensemble Kalman inversion as done in [33].

5.2.1 Outline

The layout of this work is as follows; in Section 5.3 we present and formulate the forward model of the eikonal equation reviewing existing theory of well-posedness. This will lead onto Section 5.4 where we formulate the inverse problem in a Bayesian setting while describing our inverse solver, the iterative Kalman method. This will include an overview of the various prior forms used. Finally in Section 5.5 we present numerical examples of the inverse problem combining our priors and inverse solver, while concluding in Section 5.6 with some final remarks on further areas of research.

5.2.2 Notation

Throughout this work we make use of common notation for Hilbert space norms and inner products, $\|\cdot\|, \langle\cdot\rangle$. We will assume that $\mathcal{X}$ and $\mathcal{Y}$ are two separable Hilbert spaces. These spaces are related through the forward operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$. This nonlinear operator can be thought of as mapping from the space of parameters $\mathcal{X}$ to the observation space $\mathcal{Y}$. Our additive noise for the inverse problems will be denoted by $\eta \sim \mathcal{N}(0, \Sigma)$ where $\Sigma : \mathcal{Y} \to \mathcal{Y}$ is a self-adjoint positive operator. For any such operator we define $\langle\cdot,\cdot\rangle_\Sigma = \langle\Sigma^{-1/2}\cdot, \Sigma^{-1/2}\cdot\rangle$ and $\|\cdot\|_\Sigma = \|\Sigma^{-1/2}\cdot\|$, and for finite dimensions $|\cdot|_\Sigma = |\Sigma^{-1/2}\cdot|$ with $|\cdot|$ the Euclidean norm. The superscript † will denote the true value of the particular quantity. We denote $\mathcal{U}[a, b]$ as a uniform distribution with parameters a and b.

5.3 The Forward Model

We begin by defining and reviewing the forward model of the eikonal equation.

Assume an open bounded domain $\Omega \subset \mathbb{R}^d$ for $d = 2, 3$ with a Lipschitz boundary Γ. Let $x_0 \in \Omega$ be fixed and $u : \bar{\Omega} \to \mathbb{R}$ be a positive continuous function with $x \in \bar{\Omega}$. The eikonal equation is concerned with the first arrival time of a signal. It can be interpreted as finding T associated with the smallest path ζ needed to travel from x_0 to x. There are referred to as *arrival times*. A formal definition of the first arrival time is provided in the following definition.

Definition 5.3.1. *Given a slowness function $u(x)$ our first arrival time is characterized by*

$$T(x) = \inf_{\zeta} \Big\{ \int_0^1 u(\zeta(r))|\zeta'(r)|dr \mid \zeta \in W^{1,\infty}([0,1], \bar{\Omega}), \zeta(0) = x_0, \zeta(1) = x \Big\}, \tag{5.3.1}$$

where $W^{1,\infty}([0,1], \bar{\Omega})$ is the space of paths ζ.

We denote by $u(x) = \frac{1}{c(x)}$ as the slowness function where $c(x)$ is the speed of the signal in the medium. A representation of this can be seen in Figure 5.1. We say $T(x)$ is a solution of the eikonal equation if it solves the following partial differential equation

$$|\nabla T(x)| = u(x) \quad x \in \Omega \setminus \{x_0\}, \tag{5.3.2}$$

$$T(x_0) = 0, \tag{5.3.3}$$

$$\nabla T(x) \cdot \nu(x) \geq 0, \quad x \in \Gamma. \tag{5.3.4}$$

The condition (5.3.4) is a Soner boundary condition where ν denotes a unit outer normal vector to the boundary Γ. The eikonal equation was first introduced in 1827 by Hamilton as an approximation to equations which model light behaviour travelling through various materials. The eikonal equation is characterized as a Hamilton-Jacobi (HJ) equation. As with HJ equations a common approach to show uniqueness and existence of solutions is through the theory of viscosity solutions, which is based on optimal control theory.

Viscosity solutions correspond to equations of a HJ form which satisfy the vanishing viscosity form. When it it is difficult to show well-posedness some form of regularization is added such as $\varepsilon\Delta$ for some $\varepsilon > 0$. From this by taking $\varepsilon \to 0$ we hope that T^ε will converge to a weak solution. We refer the interested reader to the book by Evans [60] and Lions [28] for a richer understanding on viscosity

solutions to PDEs. Now applying the theory of viscous solutions to our model problem we have the following definitions which are required to show well-posedness for (5.3.2)-(5.3.3).

Definition 5.3.2. *A function $T \in C^0(\bar{\Omega})$ is called a viscosity subsolution of* (5.3.2) *in $\Omega \setminus x_0$ if for each $\zeta \in C^\infty(\Omega)$: if $T - \zeta$ has local maximum at a point $x \in \Omega \setminus x_0$ then*

$$|\nabla \zeta(x)| \leq u(x). \tag{5.3.5}$$

Definition 5.3.3. *A function $T \in C^0(\bar{\Omega})$ is called a viscosity supersolution of* (5.3.2) *in $\bar{\Omega} \setminus x_0$ if for each $\zeta \in C^\infty(\mathbb{R}^n)$: if $T - \zeta$ has local minimum at a point $x \in \bar{\Omega} \setminus x_0$, relative to $\bar{\Omega}$,then*

$$|\nabla \zeta(x)| \geq u(x). \tag{5.3.6}$$

Then from Definition 5.3.2 and 5.3.3 we can generalise a viscosity solution of (5.3.2) - (5.3.4) with the following definition.

Definition 5.3.4. *A viscosity solution of* (5.3.2) - (5.3.4) *is then a function $T \in C^0(\Omega)$ which is a viscosity subsolution in $\bar{\Omega} \setminus \{x_0\}$, a viscosity supersolution in $\bar{\Omega} \setminus \{x_0\}$, and which which satisfies $T(x_0) = 0$.*

Using Definitions 5.3.2 - 5.3.4 the following theorem guarantees uniqueness and existence of a viscous solution to (5.3.2).

Theorem 5.3.1. *Suppose that $u \in C^0(\bar{\Omega})$ is non-negative, then there exists a unique viscosity solution $T \in C^0(\bar{\Omega})$ of* (5.3.2) - (5.3.4) *given as*

$$T(x) = \inf_{\zeta} \Big\{ \int_0^1 u(\zeta(r))|\zeta'(r)|dr \mid \zeta \in W^{1,\infty}([0,1], \bar{\Omega}), \zeta(0) = x_0, \zeta(1) = x \Big\}. \tag{5.3.7}$$

Furthermore there exists a a constant $C = C(\Omega)$ where T is Lipschitz continuous in Ω with an upper bound on C such that

$$lip(T) \leq C \max_{\bar{\Omega}} u \tag{5.3.8}$$

Proof. The proof can be found in [28] and [135]. □

In order to attain well-posedness for (5.3.2) - (5.3.3) we need a number of assumptions regarding both the slowness function u and the solution T. We will discuss the assumptions required for u in the proceeding section, as they will be discussed in Section 5.4 where we introduce various prior forms.

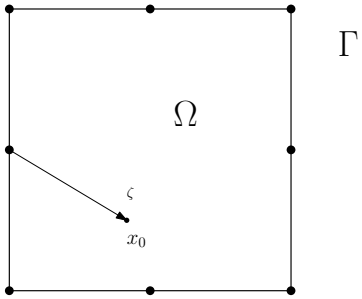

Figure 5.1: Representation of the eikonal equation.

5.3.1 Forward Finite Difference Solver

The eikonal equation also poses an important question on how to effectively numerically solve it. Since the late 20th Century there has been a significant body of work on answering this question. Namely the work by Sethian [130] helped answer this for boundary value problems of the eikonal equation where he developed a method entitled the fast marching method (FMM). Before explaining and defining the FMM we first require to discretize (5.3.2) - (5.3.4).

Assume our domain $\Omega \subset \mathbb{R}^d$ has a boundary Γ which is piecewise C^2. We will discretize our domain Ω, where $h > 0$, given the following grid

$$\mathbb{Z}_h^2 := \{x_\varphi = (h\varphi_1, h\varphi_2) |\ \varphi_i \in \mathbb{Z},\ \text{for } i = 1, 2\}.$$

Now suppose that x_0 is a grid point where $x_0 = x_{\varphi_0}$ for some $\varphi_0 \in \mathbb{Z}^2$. We define a set of inner grid points as $\Omega_h = \Omega \cap \mathbb{Z}_h^2$. If for some $x_\varphi \in \Omega_h$ there are $\iota \in \{-1, 1\}, k \in \{1, 2\}$ with $x_{\varphi+\iota e_k} \notin \Omega$, then there exists $s \in (0, 1]$ such that $x_\varphi + s\iota h e_k \in \Gamma$ and we set $\psi := \varphi + s\iota h e_k$ as well as $x_\psi := x_\varphi + s\iota e_k$. We further denote $\Gamma_h \subset \Gamma$ to be the set of all points obtained in this manner. By defining $Q_h := \Omega_h \cup \Gamma_h$ for a point $x_\varphi \in Q_h$

$$\mathcal{N}_\varphi := \begin{cases} \{x_\psi \in Q_h | x_\psi \text{ is a neighbour of } x_\varphi\}, x_\varphi \in \Omega_h, \\ \{x_\psi \in \Omega_h | x_\psi \text{ is a neighbour of } x_\varphi\}, x_\varphi \in \Gamma_h. \end{cases} \tag{5.3.9}$$

We note for points $x_\varphi \in \Gamma_h$ the region $\mathcal{N}_\varphi$ contains only the interior points. As we have provided our discretized domain and mesh with (5.3.9) our discretized system of equations representing (5.3.2) - (5.3.3) is

$$\sum_{x_\psi \in \mathcal{N}_\varphi} \Big[\big(\frac{T(x_\varphi) - T(x_\psi))}{h_{\varphi\psi}} \big)^+ \Big]^2 = u(x_\varphi)^2, \quad x_\varphi \in Q_h \setminus x_{\varphi_0}, \tag{5.3.10}$$

$$T_{\varphi_0} = 0, \tag{5.3.11}$$

where $h_{\varphi\psi} := |x_\varphi - x_\psi|$.

Lemma 5.3.1. *Suppose that $u \in C^0(\overline{\Omega})$ then* (5.3.10) - (5.3.11) *has a unique solution $T : Q_h \to \mathbb{R}$ and*

1. $T_\varphi > 0$, $x_\varphi \in Q_h$.
2. $\big|T_\varphi - T_\psi\big| \leq C\Big(\max_{\overline{\Omega}} u \Big) \big|x_\varphi - x_\psi\big|$, $x_\varphi, x_\psi \in Q_h$.

Proof. A proof can be found by Deckelnick et al. in [52]. □

Now that we have formulated our set of discretized equations (5.3.10) - (5.3.11), we are now in a position to discuss our forward solver. The method is based on a discretization of the mesh and from this the grid points needed are calculated using neighbouring grid points. Here we denote our unique solution V_φ at grid point x_φ, which as mentioned depends, only on the neighbouring values V_ψ such that $0 \leq V_\psi \leq V_\varphi$. An example of an FMM method for our problem is provided in Algorithm 6. Here we label x_{φ_0} as *known* points and any other grid points that are one point away from this *known* point as a *trial* point. Other various methods that were proposed for the numerical analysis of the eikonal equation include the fast sweeping method.

Algorithm 6 Fast Marching Method

1. Compute a trial value of $\tilde{V}_\varphi$ for every $\varphi \in$ *trial* according to (2.4), assuming it is smaller or equal to its *trial* neighbours.

2. Let x_ψ be any *trial* point where every *trial* point satisfies $\tilde{V}_\mu \leq \tilde{V}_\varphi$ $\forall x_\varphi \in$ *trial*.

3. Set $V_\psi = \tilde{V}_\psi$ for all such x_ψ and add x_μ to *known* and remove from *trial*.

4. Tag all neighbours of *known* as *trial* if they are not *known*.

5. If *trial* $= \{\emptyset\}$ then STOP.

6. Return to **1.**

5.4 Inverse Problem

In this section we discuss the inverse problem of reconstructing the slowness function from the eikonal equation. We begin with a discussion on the various priors we wish to use and how they relate to the slowness function. We then describe both approaches to inverse problems with an emphasis on Bayesian formulation of (5.3.2). This will in turn lead to the motivation for our chosen inverse solver which is the iterative ensemble Kalman method.

5.4.1 Prior

One of our motivations behind this work is to test the inverse eikonal equation on a number of priors. The prior form is crucial within inverse problems as depending hows similar it is to the truth, reconstructions can differ substantially between different priors. Before discussing the four prior forms we review the slowness function and what certain conditions we need to ensure the formulation of an inverse problem.

Assumption 5.4.1. *$u : \Omega \to \mathbb{R}$ is Borel measurable and there exists $0 < m \leq M$ such that*

$$m \leq u(x) \leq M, \quad \forall x \in \Omega. \tag{5.4.1}$$

Assumption 5.4.2. *For every $x_0 \in \Omega$ there exists $\epsilon_{x_0} > 0$ and a direction $n_{x_0} \in S^{d-1}$ so that for all $x \in \Omega$, $r > 0, d \in S^{n-1}$ with $|d - n_{x_0}| < \epsilon_{x_0}$ and $x + rd \in \Omega$ we have*

$$u(x + dr) - u(x) \leq (\omega|x_0 - x| + r). \tag{5.4.2}$$

where $\omega : [0, \infty) \to [0, \infty)$ is a continuous function which is nondecreasing where $\omega(0) = 0$. A stronger result of (5.4.2) is given below.

Assumption 5.4.3. *There exists $\epsilon > 0$ and $K \geq 0$ such that for all $x_0 \in \Omega$ there is a direction $n = n_{x_0} \in S^{n-1}$ with*

$$u(x + dr) - u(x) \leq Kr, \tag{5.4.3}$$

for all $y \in \Omega, r > 0$ and $d \in S^{n-1}$ with $|x_0 - x| \leq \epsilon, |d - n| < \epsilon$ and $y + rd \in \Omega$.

Both Assumption 5.4.1 and 5.4.2 are important in characterizing a viscosity solution to (5.3.2). Assumption 5.4.1 implies that our continuous slowness function $u(x)$ needs to be strictly. Both Assumption 5.4.2 and 5.4.3 states that the slowness function for all x is continuous, specifically at least Lipschitz continuous. However

Assumption 5.4.2 is unique in that it allows for certain discontinuities of f. An example of this can be found in [44], where the authors state that if the geometry of the discontinuities are known, then well-posedness can be shown.

Remark 5.4.1. *As a note we emphasise that with this work the analytical solution of the forward problem* (5.3.2) *may not be unique. This will of course depend on the actual prior form of the slowness function, and the assumptions stated above. Numerically all our priors are of a discrete form where the forward problem admits a unique discrete solution.*

Whittle-Matérn Priors

Our first prior form will be based on a Gaussian random field i.e. $u \sim \mathcal{N}(0, \mathcal{C})$, specifically a Whittle-Matérn random field. The function representing the prior is chosen such that $u \in C^0(\overline{\Omega})$. These are a family of Gaussian random fields which are isotropic and stationary with covariance function

$$c(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{|x - x'|}{\ell} \right)^{\nu} K_\nu \left(\frac{|x - x'|}{\ell} \right), \qquad x, x' \in \mathbb{R}^d. \tag{5.4.4}$$

where K_ν denotes a modified Bessel function of the second kind and $\Gamma(\nu)$ is a Gamma function. From (5.4.4) we also have the inclusion of three hyperparameters; the amplitude $\sigma \in \mathbb{R}$, the regularity $\nu = \alpha - d/2 \in \mathbb{R}^+$ and the length-scale $\ell \in \mathbb{R}^+$. The covariance function has a corresponding covariance operator which is given as

$$\mathcal{C} := \sigma^2 (I - \ell^2 \Delta)^{-\alpha}, \tag{5.4.5}$$

such that $\Delta u = \sum_{i=1}^{d} \partial_{x_i}^2 u$ is the Laplacian operator. The connection between the covariance structures of (5.4.4) and (5.4.4) can be found in Chapter 2.. There is a vast amount of literature which discusses how to generate random fields of the covariance form (5.4.4), which are provided in the text [99]. We will simulate our Gaussian prior based on the Karhunen-Loève (KL) expansion, which is given in the form of

$$\log u = \sum_{j=1}^{\mathcal{J}} \sqrt{\lambda_j} \xi_j \phi_j. \tag{5.4.6}$$

(λ_j, ϕ_j) are given as the eigensystem which satisfies $\mathcal{C}\phi_j = \lambda_j \phi_j$ and $\xi_j \sim \mathcal{N}(0, 1)$ is Gaussian white noise. Depending on the boundary conditions of the PDE, the eigenfunctions ϕ_j are taken as an inverse discrete Fourier transform. For the case of the eikonal equation we would impose Neumann conditions. We usually assume the series (5.4.6) is truncated, for computational purposes and that the eigenvalues

are defined as $\lambda_j := \sigma^2(I - \ell^2\Delta)^{-\alpha}$. By defining

$$c_{\mathcal{J}} = -\max_j(|\sqrt{\lambda_j}|\|\phi_j\|_\infty)\mathcal{J},$$

We can explicitly bound the variation of the values our unknown can take i.e. $\log u \in [-c_{\mathcal{J}}, c_{\mathcal{J}}]$. By taking the the exponential of $\log u$ we can bound u by

$$0 < m_{\mathcal{J}} \leq u \leq M_{\mathcal{J}}$$

given that $m_{\mathcal{J}} := \exp(-c_{\mathcal{J}})$ and $M_{\mathcal{J}} := \exp(c_{\mathcal{J}})$. Combining everything we present our first prior through the following definition.

Definition 5.4.1. *Given a set of fixed hyperparameters $\theta = (\sigma, \alpha, \ell) \in \mathbb{H}$ defined in subsection 5.4.1 corresponding to the covariance operator $\mathcal{C}$ with eigensystem (λ_j, ϕ_j), and white noise ξ, our Whittle-Matérn prior is defined such that*

$$\mathcal{P}^{\mathrm{WM}}_{(\sigma,\alpha,\ell)} = \{u \in C^0(\bar{\Omega}) \mid \log u = \sum_{j=1}^{\mathcal{J}} \sqrt{\lambda_j}\xi_j\phi_j\}, \tag{5.4.7}$$

where WM *stands for Whittle-Matérn.*

Our prior defined in (5.4.7) satisfies both assumptions 5.4.1 and 5.4.2. The positivity of the prior is indicated through the log where we take the exponential of (5.4.6). The continuity of the prior comes through the condition that $\alpha > \frac{d}{2}$ which ensures continuous sample paths, and from the continuity of the eigenfunctions ϕ_j. This is discussed in detail in [99]. Random draws from the KL expansion (5.4.6) with covariance (5.4.5) are provided in Figures 5.2 and 5.3.

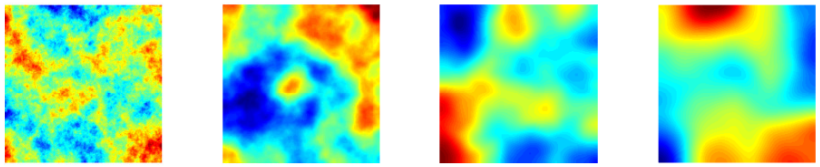

Figure 5.2: Random draws from the KL expansion with $\ell = 0.1$ and $\alpha = 1.5, 2.5, 3.5$ and 4.5 with a fixed $\sigma = 1$.

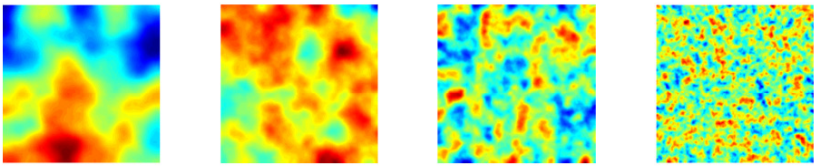

Figure 5.3: Random draws from the KL expansion with $\alpha = 3$ and $\ell = 0.1, 0.05, 0.02$ and 0.01 with a fixed $\sigma = 1$.

Hierarchical Whittle-Matérn Priors

We now wish to consider the extension of the first prior in the previous subsection to a hierarchical version. These priors consider our problem in a hierarchical manner with respect to the Whittle-Matérn covariance function 5.4.4. By doing so our unknown is not only u but $\theta = (\sigma, \alpha, \ell) \in \mathbb{H}$ which are the hyperparameters that define u. To keep our work consistent with [33] we keep the amplitude constant, and consider the recovery of only $(\alpha, \ell) \in \mathcal{H} := \mathbb{R}^+ \times \mathbb{R}^+$.

In order to do so we use an alternative approach to the KL expansion (5.4.6) known as the stochastic partial differential (SPDE) approach. This approach looks to solve the following SPDE for u

$$(I - \ell^2 \Delta)^{\frac{\alpha}{2}} \log u = \sqrt{\beta} \ell^{2/d} \xi, \tag{5.4.8}$$

where $\xi \in H^{-s}(D), s \geq \frac{d}{2}$ is Gaussian white noise and

$$\beta = \sigma^2 \frac{2^d \pi^{d/2} \Gamma(\alpha)}{\Gamma(\alpha - \frac{d}{2})}. \tag{5.4.9}$$

This approach is different to the KL expansion as rather than being a spectral method, our prior is generated by solving the SPDE (5.4.8). Numerically this can be achieved using finite element methods, where further details on the discretization can be found in Chapter 3.. As before to ensure continuity of the prior we assume $\alpha > \frac{d}{2}$, similarly with s. By solving (5.4.8) our solution has a representation of the form (5.4.4). Further information on the derivation of (5.4.8) from (5.4.4) can found in [124] but Roininen et al.. Now if we wish to work in hierarchical manner we can do so in two ways: firstly to consider the unknown as (u, θ) or as (ξ, θ). The former is referred to as the centered approach, while the later is known as the non-centered approach [115, 116], for which one reparamaterize (5.4.8) in terms of ξ. A comparison of both methods was conducted in the context of inverse problems by

Chada et al. [33], where it was shown that the non-centered approach outperformed its counterpart. The reason for this arises from the difference of both methods which is through the prior. Our prior distributions for both approaches are constructed as

$$\mathbb{P}(u,\theta) = \mathbb{P}(u|\theta)\mathbb{P}(\theta), \tag{5.4.10}$$

$$\mathbb{P}(\xi,\theta) = \mathbb{P}(\xi)\mathbb{P}(\theta). \tag{5.4.11}$$

As shown in the non-centered case (5.4.11) both the hyperparameters and the unknown are independent under the prior. Due to this loss of dependence it allows for richer reconstructions. In order to hierarchically learn our Gaussian random field, we place uniform bounds on our hyperparameters provided in Table 5.1.

Parameter	Prior distribution
α	$\mathcal{U}[1.1, 4.1]$
ℓ	$\mathcal{U}[0.02, 0.1]$

Table 5.1: Prior distributions of the hyperparameters.

Definition 5.4.2. *Given a covariance operator $\mathcal{C}$ defined in subsection 5.4.1, with constant $\beta > 0$, fixed amplitude σ and white noise ξ, our hierarchical Whittle-Matérn prior is defined such that*

$$\mathcal{P}_\sigma^{\mathrm{HWM}} = \left\{(\alpha,\ell) \in \mathcal{H}, u \in C^0(\bar{\Omega}) \mid (I-\ell^2\Delta)^{\frac{\alpha}{2}} \log u = \sqrt{\beta}\ell^{2/d}\xi\right\}, \tag{5.4.12}$$

where HWM *denotes hierarchical Whittle-Matérn.*

Level Set Priors

Level set methods [128, 130] have been important when analyzing problems regarding interfaces between domains, such as when the unknown is piecewise constant. These methods work by thresholding a continuous function $\tilde{u}$ which results in a level set function u. We can then formulate the inverse problem in the form (5.2.1) for u. Our level set function is defined as

$$F\tilde{u}(x) \to u(x) = Thr(\tilde{u}(x)),$$

where $F : C^0(\overline{\Omega};\mathbb{R}) \to \mathcal{D}(\overline{\Omega};\mathbb{R})$ is our level set mapping. Here $\mathcal{D}(\overline{\Omega};\mathbb{R})$ is a space of discontinuous slowness functions. Level set methods have seen wide applications to inverse problems [23, 25]. Recently there have been extensions of this to Bayesian

inverse problems where the Bayesian level set method [47, 81] was derived. For our level set prior we assume $\tilde{u}$ is a Gaussian random field generated through (5.4.8). for which thresholding results in $n = 3$ subdomains. We have fixed constants $[c_1, c_2]$ which define the thresholding levels, given as

$$c_1 = 0.1 \times \ell^{-(\frac{d}{2}-\alpha)}, \tag{5.4.13}$$

$$c_2 = -0.1 \times \ell^{-(\frac{d}{2}-\alpha)}, \tag{5.4.14}$$

which are scaled through the regularity α and length scale ℓ. Given these threshold levels, we define multiple subdomains $\{\Omega\}_{i=0}^{n-1}$ defined as

$$\begin{aligned}
\Omega_1 &= \{x \in \Omega \mid \tilde{u}(x) < c_1\}, \\
\Omega_2 &= \{x \in \Omega \mid c_1 \leq \tilde{u}(x) \leq c_2\}, \\
\Omega_3 &= \{x \in \Omega \mid \tilde{u}(x) > c_2\},
\end{aligned}$$

such that $\Omega_i \cap \Omega_j = \emptyset$ for $i \neq j$. Therefore our level set function is defined as

$$Thr(\tilde{u}(x)) = \begin{cases} u_0, & \tilde{u}(x) < c_1, \\ u_1, & c_1 \leq \tilde{u}(x) \leq c_2, \\ u_2, & \tilde{u}(x) > c_2. \end{cases}$$

The amount of subdomains and thresholding is not restricted to the setting we have chosen. As our threshold levels are determined through the parameters in Table 5.1, we work in a hierarchical manner. We are now interested in recovering both the regions u_i and the threshold levels $[c_1, c_2]$.

Definition 5.4.3. *Given a Gaussian random field $\tilde{u}$ defined in subsection 5.4.1, and thresholding constants $[c_1, c_2]$ given by* (5.4.13)-(5.4.14)*, our level set prior is given as*

$$\mathcal{P}^{\text{LS}}_{(\sigma,c)} = \big\{(\alpha, \ell) \in \mathcal{H}, \tilde{u} \in C^0(\bar{\Omega}), u \in \mathcal{D}(\bar{\Omega}) \mid u(x) = Thr(\tilde{u}(x))\big\}, \tag{5.4.15}$$

where LS *denotes level set.*

Unlike the previous priors defined in subsections 5.4.1 and 5.4.1, our level set prior is not continuous, but abides by Assumption 5.4.1. An example of the level set prior is given in Figure 5.4 which includes $n = 3$ defined interfaces through thresholding.

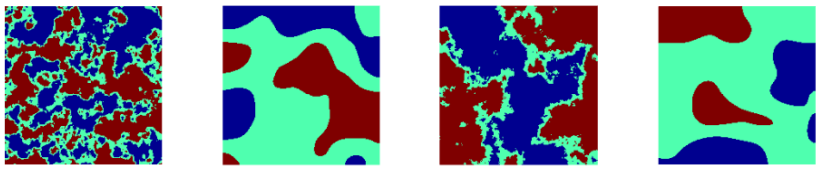

Figure 5.4: Random draws from level set thresholding of three interfaces from Gaussian draws with varying regularity and length scale with three interfaces.

Vector-Valued Level Set Priors

One disadvantage that arises with the level set formulation in subsubsection 5.4.1 is that it orders each of different layers or interfaces. This implies, in the case of three interfaces, that each regions can not be directly connected. One way to alleviate this issue is to consider a different level set formulation. Given a level set function $S : C^0(\overline{\Omega}; \mathbb{R}^n) \to \mathcal{D}(\overline{\Omega}; \mathbb{R}^n)$ where S is defined in such a way that

$$S\tilde{u}(x) \to u(x) = e_{r(x;\tilde{u})}, \tag{5.4.16}$$

where

$$r(x; \tilde{u}) = \underset{i \in \{1,\dots,n\}}{\operatorname{argmax}} \tilde{u}_i(x). \tag{5.4.17}$$

Similarily with the level set prior (5.4.15), $\tilde{u}$ denotes a Gaussian random field generated through the SPDE (5.4.8). In this formulation our level set map $S\tilde{u} = u$ is represented through $\{e_r\}_{i=1}^n$ which is the standard orthonormal basis on $\mathbb{R}^n$ determined through the operation (5.4.17). The number of interfaces is denoted by n. Not only does this alternative formulation allow for all interfaces to coincide but it allows for unknowns that can not be generated from the level set formulation as described through (5.4.15). This motivation is taken by the work of Bertozzi et al. [75], where they consider it for classification problems. The alternative level set method is defined through the following.

Definition 5.4.4. *Given a set of hyperparameters $\theta = (\alpha, \ell) \in \mathcal{H}$ defined in subsection 5.4.1, a collection of n Gaussian random fields through* (5.4.12) *and a thresholding level r, our level set prior is given as*

$$\mathcal{P}^{\text{VLS}}_{(\sigma,r)} = \left\{(\alpha, \ell) \in \mathcal{H}, \{\tilde{u}\}_{i=1}^n \in C^0(\bar{\Omega}), u \in \mathcal{D}(\bar{\Omega}) \mid u(x) = e_{r(x;\tilde{u})}\right\}, \tag{5.4.18}$$

where VLS *denotes vector level set. As an example we show draws from the prior form* (5.4.18) *for the case of $n = 3$ in Figure 5.5.*

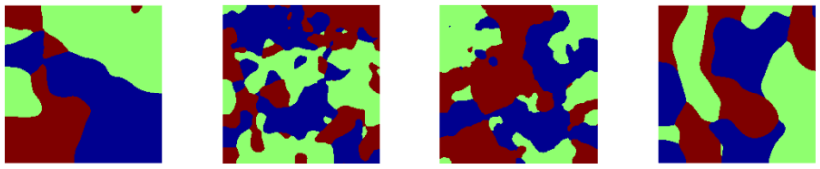

Figure 5.5: Random draws from vector level set thresholding with varying regularity and length scale with three interfaces.

Fixed Shape Priors

Our last prior we consider are priors that can be defined through piecewise constant functions where shape is fixed. These priors differ in the case of the level set priors as they are not based on thresholding a Gaussian random field. Our fixed shape prior will include circular subdomains $B_R(a_i)$, for $i = 1, \ldots, n$ where n denotes the number of balls. The geometric parameters associated with each circle $\{u_i\}_{i=1}^n$ is a common radius R and center $\{a_i\}_{i=1}^n = (a_{i,x}, a_{i,y})$. Our function defining each circle are given as

$$u_i = (x_1 - a_{i,x})^2 + (x_2 - a_{i,y})^2 - R^2.$$

Therefore our fixed shape circular prior will take the general form

$$u(x) = \sum_{i=1}^{n} u_i \chi_{B_R(a_i)}(x), \tag{5.4.19}$$

where $\chi_{B_R(a_i)}(x)$ is a characteristic function of the form

$$\chi_{B_R(a_i)}(x) = \begin{cases} 1, & x \in B_R(a_i) \\ 0, & x \notin B_R(a_i) \end{cases} \tag{5.4.20}$$

such that $B_R(a_i) \cup B_R(a_j) = \emptyset$ for $i \neq j$. These prior were first tested in the Bayesian formulation by Iglesias et al. [80] where they considered geometric priors for subsurface flow. Our prior $u \in BV(\Omega)$ belongs the space of functions of bounded variation where $BV(\Omega) = \{u \in L^1(\Omega) : \int_\Omega |Du| < \infty\}$. Specifically for this work (5.4.19) is now chosen were we have a two circular inclusion prior. Therefore we introduce an unknown $\hat{u} = (a_1, a_2, R) \in \mathcal{F} := \mathbb{R}^5$ with parameters

- $a_1 = (a_{1,x}, a_{1,y}) \in \mathbb{R}^2$ - first circular center.
- $a_2 = (a_{2,x}, a_{2,y}) \in \mathbb{R}^2$ - second circular center.
- $R \in \mathbb{R}^+$ - circular radius.

In order to ensure that the circles do not coincide with each other, we ideally would like a radius, such that for $\epsilon > 0$

$$|a_i - a_j| \geq R + \epsilon.$$

For this prior we impose conditions which prevent this type of phenomenon. This is highlighted in Table 5.2 below. The following definition presents the fixed shape prior.

Definition 5.4.5. *Given a set of geometric parameters $(a_1, a_2, R) \in \mathcal{F}$ defined in subsection 5.4.1, our fixed shape prior is constructed as*

$$\mathcal{P}^{\text{FS}} = \big\{(a_1, a_2, R) \in \mathcal{F}, u \in BV(\Omega) \ | u(x) = \sum_{i=1}^{n} u_i \chi_{B_R(a_i)}(x)\big\}, \tag{5.4.21}$$

where FS *denotes fixed shape.*

Our fixed shape prior (5.4.21) will abide by Assumption 5.4.1, as similarly with the level set prior (5.4.15) to ensure positivity we assign positivity to the balls and the region outside the balls. However for continuity Assumption 5.4.3 holds for some discontinuities, such as numerous curves in the domain which are highlighted in an example of Definition 2.1 in [44].

Parameter	Prior distribution
R	$\mathcal{U}[0, 0.2]$
a_1	$\mathcal{U}[0.10, 0.45]^2$
a_2	$\mathcal{U}[0.55, 0.85]^2$

Table 5.2: Distributions of the geometric parameters.

From Table 5.2, four random draws generated through (5.4.21) are shown in Figure 5.18.

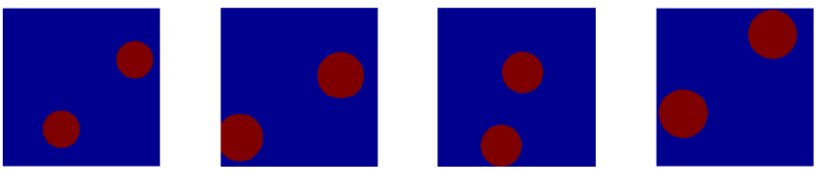

Figure 5.6: Random draws from circular fixed shape prior with varying positions and radii.

Remark 5.4.2. *We emphasize that for our fixed shape prior we keep the various levels of the slowness function constant, however it can be of interest to model this prior further with the inclusion of random parameter for the values of* $u_i(x)$.

To understand the effect of each prior on the solution $T(x)$ we see visually the differences in $T(x)$ as we consider each prior individually. This is seen through Figure 5.7 where we place 3 source points situated at $(0.4, 0.5), (0.3, 0.35)$ and $(0.75, 0.85)$ in the domain $D = [0, 1]^2$.

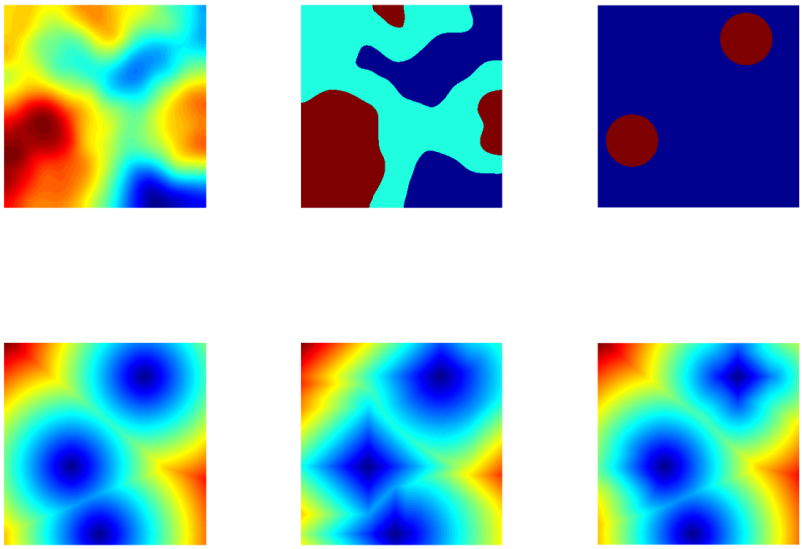

Figure 5.7: Top row: various slowness functions of priors. Bottom row: corresponding forward solution.

5.4.2 Deterministic Approach

The inverse problem associated with the eikonal equation was initially considered in a deterministic setting [52]. It was presented in an optimization framework which uses a least squares formulation: given n number of measured arrival times denoted as $T_{\text{obs}} : \Gamma \to \mathbb{R}^+$ we aim to minimise the misfit functional, i.e. compute

$$\underset{u \in \mathcal{X}}{\operatorname{argmin}} J(u),$$

such that our functional is expressed as

$$J(u) = \frac{1}{2}\int_{\Gamma} |T_u(x) - T_{\text{obs}}(x)|^2 do_x, \tag{5.4.22}$$

where $T_u(x)$ is our forward solution. Now by considering a number of source points x_0^i, for $i = 1, \ldots, I$ we can rewrite the functional (5.4.22) as

$$J(u) = \frac{1}{2}\int_{\Gamma} |T_u^i(x) - T_{\text{obs}}^i(x)|^2 do_x. \tag{5.4.23}$$

An important characteristic in the deterministic framework is regularization, which adds some information of the underlying unknown to improve reconstruction of the unknown. One of the simplest and most popular forms of regularization is Tikhonov regularization which was considered in [52]. When applied, modifies the functional (5.4.22) to

$$J(u) = \frac{1}{2}\int_{\Gamma} |T_u(x) - T_{\text{obs}}(x)|^2 do_x + \frac{\lambda}{2}\int_{\Omega} |\nabla u|^2, \tag{5.4.24}$$

for $\lambda > 0$. Here the form of regularization comes through gradient of the unknown u. Well-posedness for (5.3.2) with prescribed Sonar boundary conditions was shown in [52], which was achieved through showing a subsequence converging to the solution as the mesh size went to zero. This was achieved through using techniques from Γ - convergence.

5.4.3 Bayesian Approach

We now consider a statistical approach to model the inverse problem. The Bayesian approach to inverse problems [137] seeks to solve (5.2.1) for $u(x)$ by treating each quantity as a random variable. Through an application of finite-dimensional Bayes' Theorem we can construct a posterior distribution

$$\begin{aligned}\mathbb{P}(u|y) &= \frac{\mathbb{P}(y|u)\mathbb{P}(u)}{\mathbb{P}(y)} \\ &\propto \mathbb{P}(y|u)\mathbb{P}(u),\end{aligned}$$

where $\mathbb{P}(u)$ denotes the prior distribution under our unknown u and $\mathbb{P}(y|u)$ represents the data-likelihood. By relating (5.2.1) we can construct our inverse problem in a Bayesian setting. We are now in a position to relate the Bayesian approach to our model of interest (5.3.2). Recall that with the eikonal equation our space of slowness functions is defined as either $u(x) \in \mathcal{X} := C^0(\overline{\Omega})$ or $u(x) \in \mathcal{X} := \mathcal{D}(\overline{\Omega})$, and

similarly with our solution space $T(x) \in V := C^0(\overline{\Omega})$. As before, assume a forward operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$ where our forward operator is defined as $\mathcal{G} = \mathcal{O} \circ G : \mathcal{X} \to \mathbb{R}^m$, composing of two mappings. Here $\mathcal{O} : V \to \mathbb{R}^m =: \mathcal{Y}$ is defined as our observational operator and $G : \mathcal{X} \to V$ is our parameter to unobserved operator. We are interested in the undetermined problem of recovering u_j from noisy measured functionals

$$y_j = (\mathcal{O}_j \circ G_j)(u) + \eta_j.$$

We concatenate these observations to get data $y \in \mathbb{R}$ such that

$$y = \mathcal{G}(u) + \eta. \tag{5.4.25}$$

We note that (5.4.25) is a general expression of the inverse problem. Though in order to account for the numerous priors we have mentioned, we have to reformulate the inverse problem.

Level set prior

Given our level set prior $\mathcal{P}^{\mathrm{LS}}_{(\sigma,c)}$ defined as (5.4.15), we now have the inclusion of an additional mapping $F : \mathcal{X} \to X$ which is our level set map. With the inclusion of this map our forward operator is now given as $\mathcal{G}_{\#} = \mathcal{O} \circ G \circ F : \mathcal{X} \to \mathbb{R}^m$. Using the same arguments as in Subsection 5.4.3 we have

$$y = \mathcal{G}_{\#}(u) + \eta. \tag{5.4.26}$$

Remark 5.4.3. *We can alter* (5.4.26) *to consider the vector level set formulation using the prior* $\mathcal{P}^{VLS}_{(\sigma,r)}$ *defined in* (5.4.18)*. The only modification would be to change in the level set map* S*.*

Hierarchical Whittle-Matérn prior

Extending our Gaussian prior to a hierarchical prior requires the SPDE formulation (5.4.8). In particular now our unknown is (ξ, θ) such that $\xi \in H^{-s}(\Omega)$ is white noise with $\theta = (\alpha, \ell) \in \mathcal{H}$ denoting the collection of hyperparameters. In order to account for ξ, which we refer to as the non-centered approach, our forward mapping is such that $\mathcal{G}^{\mathcal{T}} = \mathcal{G} \circ \mathcal{T} = \mathcal{O} \circ G \circ \mathcal{T} : \mathcal{H}^* \to \mathbb{R}^m$, where $\mathcal{H}^* := \mathcal{H} \times H^{-s}(D)$. The mapping $\mathcal{T} : (\xi, \theta) \to u$ is associated with the non-centered approach arising from (5.4.8). Again using similar arguments we can express our inverse problem as

$$y = \mathcal{G}(\mathcal{T}(\xi, \theta)) + \eta. \tag{5.4.27}$$

We omit the formulation of the inverse problem for the fixed shape prior (5.4.21) as it is the same to (5.4.25).

5.4.4 Likelihood and Posterior

In order to characterize a solution of our inverse problems defined we need three components as discussed from Bayes' Theorem (1.2.2). All that remains is the likelihood for which we can represent our solution as a posterior distribution. We aim to do so for all the inverse problems that were discussed. We firstly recall the inverse problem of recovering u from noisy measurements y where

$$y = \mathcal{G}(u) + \eta, \quad \eta \sim \mathcal{N}(0, \Sigma). \tag{5.4.28}$$

We can further define a potential for our inverse problem $\Phi(u; y) : \mathcal{X} \to \mathbb{R}$ where

$$\Phi(u; y) = \frac{1}{2}|y - \mathcal{G}(u)|^2_\Sigma. \tag{5.4.29}$$

From the potential given in (5.4.29) we can define our data-likelihood as

$$\mathbb{P}(y|u) = \exp\big(-\Phi(u; y)\big). \tag{5.4.30}$$

Combing both our prior and data-likelihood (5.4.30), via Bayes' Theorem, we can construct our posterior probability

$$\mathbb{P}(u|y) \propto \mathbb{P}(y|u) \times \mathbb{P}(u) \tag{5.4.31}$$

$$= \exp\big(-\Phi(u; y)\big) \times \mathbb{P}(u). \tag{5.4.32}$$

Remark 5.4.4. *As our non-hierarchical Gaussian and fixed shape priors associated inverse problems can be expressed through* (5.4.28)*, we further omit specifying their data-likelihoods and posteriors which are defined as* (5.4.30) *and* (5.4.32)*.*

Level Set Posterior

For our level set inverse problem, we emphasis two differences. Firstly that our prior is of the form (5.4.15), and that our forward operator includes a level set mapping. Therefore our inverse problem is to recover u from noisy measurements y where

$$y = \mathcal{G}_\#(u) + \eta, \quad \eta \sim \mathcal{N}(0, \Sigma),$$

The potential for our inverse problem $\Phi_{\#}(u;y) : \mathcal{X} \to \mathbb{R}$ is defined such that

$$\Phi_{\#}(u;y) = \frac{1}{2}|y - \mathcal{G}_{\#}(u)|^2_{\Sigma}. \tag{5.4.33}$$

From (5.4.33) we can define our data-likelihood as

$$\mathbb{P}(y|u) = \exp\big(-\Phi_{\#}(u;y)\big). \tag{5.4.34}$$

Via Bayes' Theorem, we can construct our posterior probability through our prior ad data-likelihood (5.4.34)

$$\begin{aligned}\mathbb{P}(u|y) &\propto \mathbb{P}(y|u) \times \mathbb{P}(u) \\ &= \exp\big(-\Phi_{\#}(u;y)\big) \times \mathbb{P}(u).\end{aligned}$$

Hierarchical Posterior

Recall that with our hierarchical Whittle-Matérn prior (5.4.12) we are interested in recovering (ξ, θ) from y, where θ is the collection of hyperparameters and ξ is the forcing term of the SPDE (5.4.8). An important distinguishment with this prior is that both ξ and θ are independent, implying $\mathbb{P}(\xi, \theta) = \mathbb{P}(\xi)\mathbb{P}(\theta)$. Therefore our inverse problem is such that

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

with potential $\Phi_{\mathrm{NC}}(\xi, \theta; y) : \mathbb{H} \to \mathbb{R}$ given as

$$\Phi_{\mathrm{NC}}(\xi, \theta; y) = \frac{1}{2}|y - \mathcal{G}(\mathcal{T}(\xi, \theta))|^2_{\Sigma}, \tag{5.4.35}$$

where NC denotes non-centered. From the potential given in (5.4.35) we can define our data-likelihood as

$$\mathbb{P}(y|\xi, \theta) = \exp\big(-\Phi_{\mathrm{NC}}(\xi, \theta; y)\big). \tag{5.4.36}$$

Combing both our prior and data-likelihood (5.4.36), via Bayes' Theorem, we can construct our posterior probability

$$\begin{aligned}\mathbb{P}(\xi, \theta|y) &\propto \mathbb{P}(y|\xi, \theta) \times \mathbb{P}(\xi, \theta) \\ &= \exp\big(-\Phi_{\mathrm{NC}}(\xi, \theta; y)\big) \times \mathbb{P}(\xi)\mathbb{P}(\theta).\end{aligned}$$

Remark 5.4.5. *For this work we omit characterizing well-posedness of the inverse*

problems. Showing this will depend both on the prior forms, but also the forward problem used. Specific to the case of our inverse solver, as there is no posterior to analytically infer we do not require a well-posedness theorem. However we note this is a useful direction to pursue, for which we leave as future work to consider which could enhance this work.

5.4.5 Iterative Ensemble Kalman Method

The field of Bayesian inverse problems has seen developments of numerous computational algorithms. In particular a large majority of them have been for Monte Carlo methods, such as MCMC and sequential Monte Carlo, which have been favourable for extracting information well from the posterior distribution. An alternative to these methods are taken from data assimilation which are ensemble methods. These methods use an ensemble of particles to update the quantity of interest. Rather than actually sampling from the posterior these methods can be thought of as optimisers which are based on update equations. In particular common examples ensemble-based methods include the randomized maximum likelihood and the ensemble Kalman filter (EnKF) [57]. Our inverse solver that we are interesting in using is an extension of the EnKF which is the iterative ensemble Kalman method [77, 78].

The iterative ensemble Kalman method was originally developed by Iglesias et al. [79] to solve PDE-constrained optimization problems, specifically for inverse problems. The key idea behind the method is to update an ensemble of particles $\{u_n^{(j)}\}_{j=1}^J \subset \mathcal{X}$ where n denotes the iteration count and J are the number of ensemble particles. The method can be thought of as a black-box which is derivate free. It combines the usual procedure of the EnKF where there is a *prediction* and *update step* and, while taking techniques from deterministic inverse problems such as the Levenberg-Marquardt (LM) method [71]. Like the EnKF the *prediction* step defines

a sample mean and empirical covariances based on the ensemble of particles

$$\bar{u} = \frac{1}{J}\sum_{j=1}^{J} u_n^{(j)}, \tag{5.4.37}$$

$$\bar{\mathcal{G}} = \frac{1}{J}\sum_{j=1}^{J} \mathcal{G}(u_n^{(j)}), \tag{5.4.38}$$

$$C_n^{ww} = \frac{1}{J-1}\sum_{j=1}^{J} (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}) \otimes (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}), \tag{5.4.39}$$

$$C_n^{uw} = \frac{1}{J-1}\sum_{j=1}^{J} ((u_n^{(j)}) - \bar{u}) \otimes (\mathcal{G}(u_n^{(j)}) - \bar{\mathcal{G}}), \tag{5.4.40}$$

and maps these from the parameter space $\mathcal{X}$ to the data space $\mathcal{Y}$. The update compares the mapped ensemble with the data and matches it with it based on the update equation

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \Sigma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})), \tag{5.4.41}$$

which uses the sample mean and empirical covariances. The derivation and motivation behind this method arising by relating it to the least squares formulation. In the deterministic framework, suppose our functional is of the form $I(u) = \|y - \mathcal{G}(u)\|_\Gamma^2 + \|u - \bar{u}\|_\mathcal{C}$, our inverse solution is then defined as

$$u_{\text{TP}} = \operatorname{argmin} I(u),$$

with TP denoting Tikhonov-Phillips regularization. the work of Lehtinen et al. [92] showed that for a linear inverse problem, i.e. $\mathcal{G}(u) = \mathcal{G}u$, that for any $C, \mathcal{G}, \Gamma$ the inverse solution has the form

$$u_{\text{TP}} = \bar{u} + C\mathcal{G}^*(\mathcal{G}C\mathcal{G}^* + \Gamma)^{-1}(y - \mathcal{G}\bar{u}), \tag{5.4.42}$$

with $\mathcal{G}^*$ expressed as the adjoint forward operator. Now if we consider (5.4.37) - (5.4.40), then by taking the ensemble size $J \to \infty$ then our update equation (5.4.41) remains consistent with (5.4.42). This is justified in the following theorem.

Theorem 5.4.1. *Assume a linear operator $\mathcal{G}(u) = \mathcal{G}u$ with update equation* (5.4.41) *and least squares solution* (5.4.42) *to $I(u)$, then*

$$u \to u_{\text{TP}}, \qquad \text{as } J \to \infty. \tag{5.4.43}$$

Proof. The proof can be found in the work of Iglesias et al. [79]. □

We emphasis with this method, unlike the least squares approach, that it remains derivative free. As stated the iterative ensemble Kalman method is also based on the LM method, specifically it uses regularization properties that were first used in [71] to prevent the over-fitting of data. This form of regularization comes in the form a discrepancy principle, where if the following condition is satisfied

$$\|\Sigma^{-1/2}(y - \bar{w}_n)\|_{\mathcal{Y}} \leq \delta\eta, \tag{5.4.44}$$

then we terminate the experiment. From (5.4.44) $\delta > 1/\rho$ where $\rho \in (0, 1)$ denotes a regularization term. With the inclusion of the discrepancy principle this modifies the update equation (5.4.41) thus leading to a new update equation

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \beta_n\Sigma)^{-1}(y_{n+1}^{(j)} - \mathcal{G}(u_n^{(j)})). \tag{5.4.45}$$

A full numerical investigation on the effect of tuneable parameters for the regularization are provided in [77]. The full method of our inverse solver is provided by Algorithm 5 in Chapter 4.

5.5 Numerical Experiments

This section will be devoted to applying our priors, which were defined in Section 5.4, to the inverse eikonal equation. We seek to test each priors performance within the inverse solver, where we will quantify this through various ways:

$$\|y - \bar{\mathcal{G}}(u_{\text{EKI}})\|_{\Sigma}^2 \quad \text{(Data misfit)}, \tag{5.5.1}$$

$$\frac{\|u_{\text{EKI}} - u^{\dagger}\|_{L^2(\Omega)}}{\|u^{\dagger}\|_{L^2(\Omega)}} \quad \text{(Relative error)}, \tag{5.5.2}$$

$$\frac{1}{J}\sum_{j=1}^{J} u_{\text{EKI}}^{(j)} \quad \text{(Mean output at iteration)}, \tag{5.5.3}$$

where $u_{\text{EKI}}^{(j)}$ denotes the current iterate solution. As well as considering the data misfit (5.5.1) and relative error (5.5.2), our mean output will be on five different values throughout the iteration count which we will set $n = 23$. Each subfigure will be assigned an iteration number i.e. $(1^{st}, 2^{nd}, 3^{rd}, 4^{th}, 5^{th}) = (1, 5, 11, 17, 23)$. We denote $u^{\dagger}$ as our truth we are interested in reconstructing and our reconstructed

solution thus far in the iterative solver as u_{EKI} where EKI stands for ensemble Kalman inversion. As the inversion can largely depend on the initial ensemble, we repeat each experiment for each corresponding prior eight times. We showcase four final reconstructions from the eight experiments, where the best experiment will be shown in terms of the iteration number mentioned above. Our forward solver for (5.3.2) will be based on the FMM which is described in Algorithm 1. We will choose for our discretization $h = 1/100$. We specify 12 source points on the boundary with 3 on each corresponding side of the boundary. Our inverse solver, as mentioned, will be based on EKI with an ensemble size of $J = 200$, a regularization parameter of $\rho = 0.8$, the variance of the noise as $\Sigma = \gamma^2 I$ such that $\gamma^2 = 0.01$. Our motivation for choosing these prescribed values for the iterative solver comes from [33, 77] where Iglesias tested optimal values for various parameters from the inverse scheme. We further assume 64 pointwise observations $\{l_t(T)\}_{t=1}^{64}$ given as

$$l_t(T) = \int_\Omega \frac{1}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}(x-x_t)^2} T(x)dx,$$

that are defined on the boundary. From our measurements synthetic data y are generated through

$$y = (l_1(h^\dagger), \ldots, l_{64}(h^\dagger)) + \eta, \quad \eta \sim \mathcal{N}(0, \Sigma). \tag{5.5.4}$$

All of the numerics are run through MATLAB@2014 on a computer with 8GB RAM DDR3 and a processor of Intel core 3.2 GHz.

Remark 5.5.1. *Our noisy measurements y in the form of our forward solution can be taken in various ways. Usually the measurements are taken either in the domain Ω or placed on the boundary Γ. An another alternative to this is a combination where the measurements are over both the boundary and domain. For the purposes of this work we keep the measurements on the domain, which remain consistent with the work of [52].*

5.5.1 Hierarchical Whittle-Matérn Prior

Our first set of numerical experiments will involve the Gaussian Whittle-Matérn prior defined in Section 5.4.1 through the SPDE formulation. We consider this inverse problem in a hierarchical manner. As discussed in subsection 5.4.1, acting hierarchically allows us to further consider the reconstruction of certain hyperparameters. For these experiments we will only consider $\theta = (\alpha, \ell) \in \mathbb{R}^+ \times \mathbb{R}^+$ where we exclude the scaling constant $\sigma > 0$. Part of the reason for this is to remain consistent

with the work carried out in [33], and that differences in the tested hyperparameters are considerably more noticeable. Our true values for our hyperparameters are provided in Table 5.3.

Parameter	True value
α	3.8
ℓ	0.08

Table 5.3: True value of the hyperparameters.

For our initial ensemble our priors for θ will be chosen based on Table 5.1. Our truth constructed through Table 5.3 is shown in Figure 5.9.

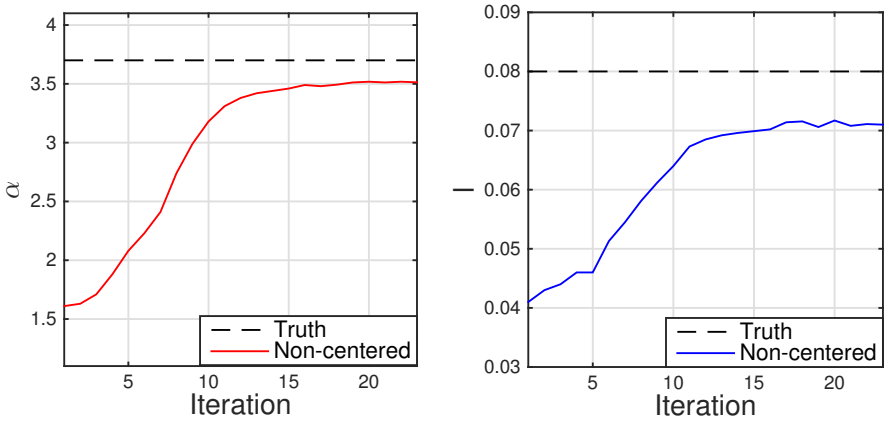

Figure 5.8: Learning rates of hyperparameters. Left: α. Right: ℓ.

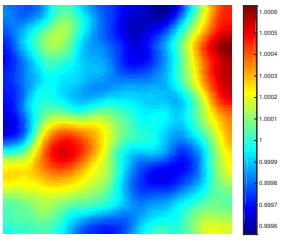

Figure 5.9: Hierarchical truth.

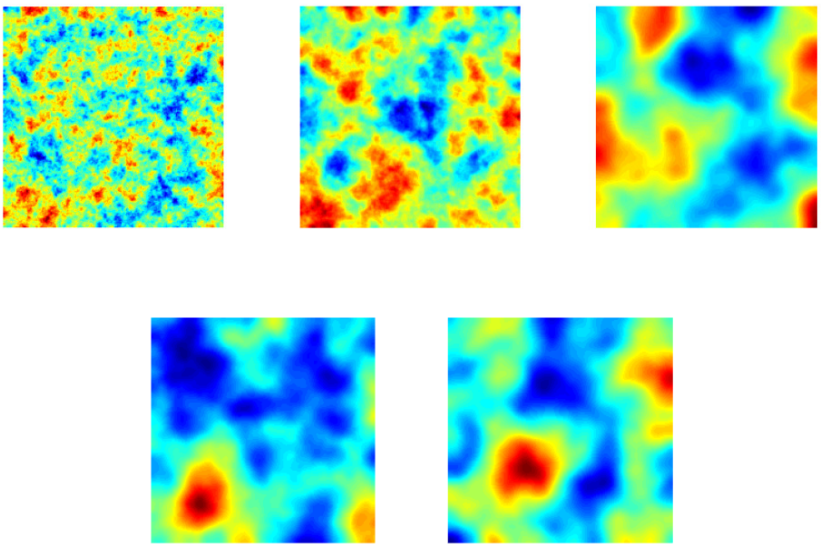

Figure 5.10: Reconstruction of hierarchical prior.

Our first set of numerical experiments were conducted with the hierarchical prior. If we first analyze Figure 5.10 we see the progression of the iterative method. We notice initially that we see an improvement in terms of both the hyperparameters where we see an increase in both the length-scale and regularity. As the iterative method terminates, our solution at the end is consistent with the truth, both in terms of the true values of the hyperparameters and the overall unknown, as seen in Figure 5.9. This is verified through Figure 5.15 which shows the learning rates of both the hyperparameters.

The performance of the iterative scheme in terms of the data misfit and relative error are given in Figure 5.12. As mentioned an important note to make on the scheme is how the initial ensemble is chosen. Figure 5.11 shows solutions to

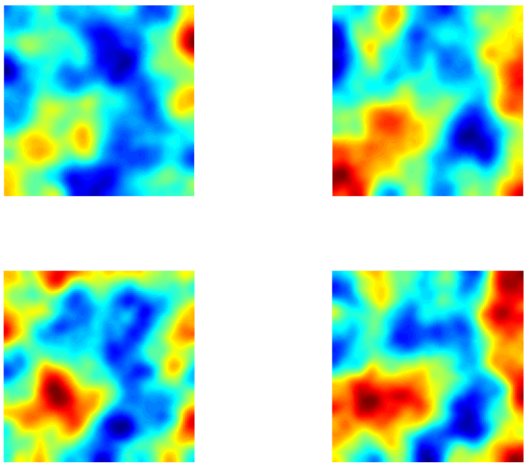

Figure 5.11: Final iteration from four different initializations of hierarchical prior.

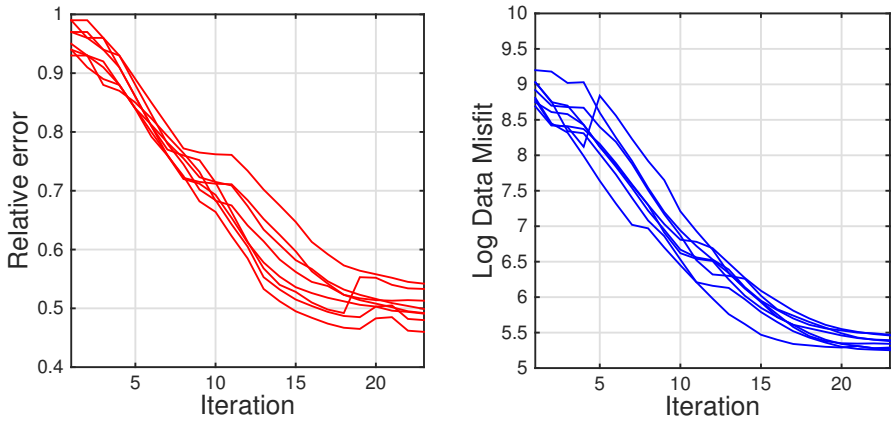

Figure 5.12: Hierarchical prior. Left: Relative error. Right: Log data misfit.

four of the different initializations, which gives an indication that regardless of how the initial ensemble is chosen, based on Table 5.1, the iterative method is able to reconstruct the truth.

5.5.2 Vector Level Set Prior

As discussed in subsection 5.4.1 we have introduced two level set priors, one which is based on a layering system. The other method negates this issue which is based on a vector level set approach. Due to this we will test only the vector level method, which we aim to test hierarchically. For our level set experiments we will consider the case of $k = 3$ classes. As a result our truth is constructed through 3 Gaussian random fields which is presented in Figure 5.13.

The true values of the hyperparameters are chosen based on Table 5.5, where as the prior on both hyperparameters are chosen in the same manner as before, namely Table 5.1.

Parameter	True value
α	4.0
ℓ	0.06

Table 5.4: True value of the hyperparameters.

Figures 5.15 and 5.14 provide a representation of how the algorithm works and how the learning of the hyperparameters proceed. For Figure 5.14 we see the progression of the iterative method, as shown similarly in subsection 5.5.1. The overall learning of the algorithm seems to work well, however in the case of vector level set method it takes slightly longer to reconstruct the unknown. However despite this, in the the learning rate of the hyperparameters we see a similar phenomenon where after roughly the 15^{th} iteration, the rate remains roughly constant.

We repeated this experiment numerous times where the four best overall reconstructions are found in Figure 5.17. Unlike with the hierarchical prior $\mathcal{P}_{\sigma}^{\text{HWM}}$ we see a greater variation between the solutions, which could be argued due to the intensity and difficult of working with the vector level set prior. The relative errors and the data misfits of all repeated experiments, are presented in Figure 5.16.

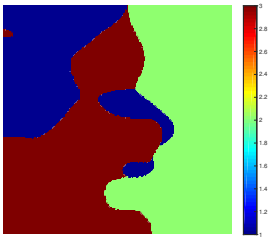

Figure 5.13: Vector level set truth.

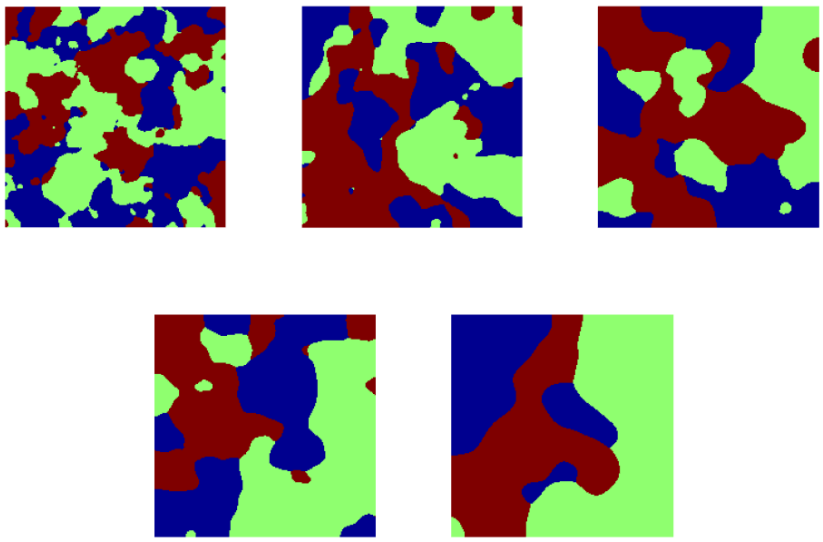

Figure 5.14: Reconstruction of vector level set prior.

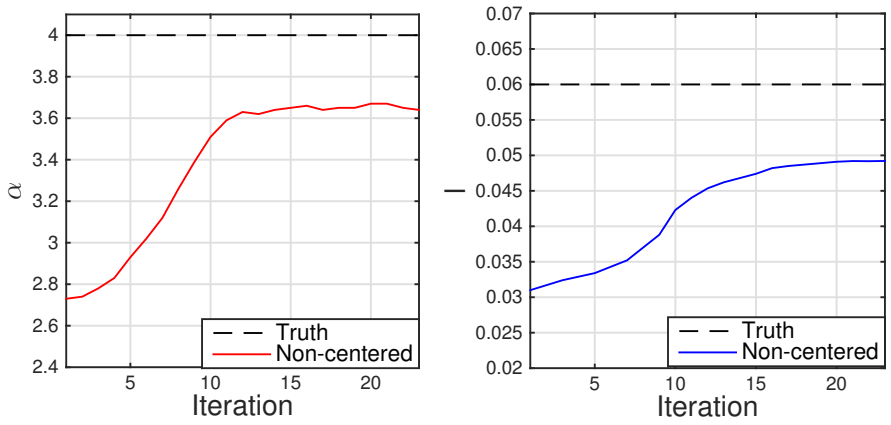

Figure 5.15: Learning rates of hyperparameters. Left: α. Right: ℓ.

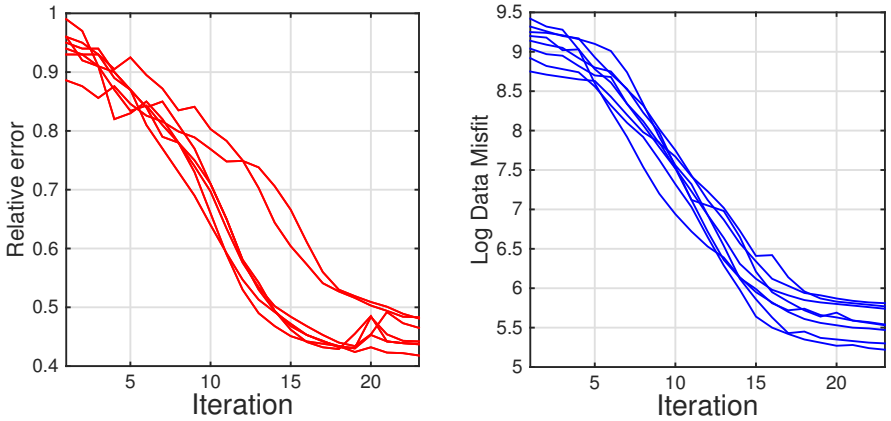

Figure 5.16: Vector level set prior. Left: Relative error. Right: Log data misfit.

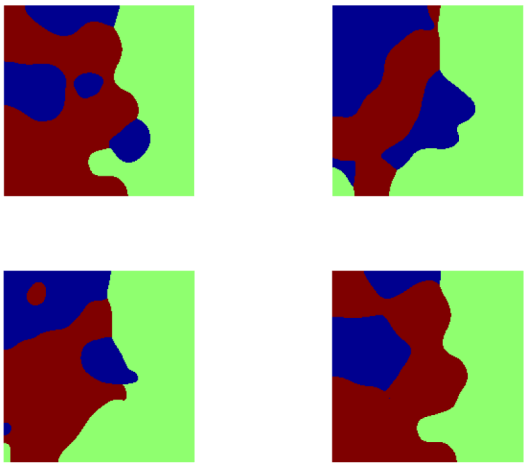

Figure 5.17: Final iteration from four different initializations of vector level set prior.

5.5.3 Fixed Shape Prior

The final prior we wish to test is our fixed shape prior with circular inclusions. The prior is defined in the form (5.4.19) where we set two circular inclusions in our domain, and a common radius between the circles. We specify a truth which is provided in Figure 5.18, which is aided through the geometric parameters chosen as shown in Table 5.5. As before we keep the setting the same for both the inverse and forward solver, where we are interested in assessing the performance through (5.5.1) - (5.5.3).

We emphasis that our focus of this experiment is to recover the geometry of the slowness function.

Parameter	True value
$(a_{1,x}, a_{1,y})$	(0.375,0.375)
$(b_{2,x}, b_{2,y})$	(0.625,0.625)
R	0.12

Table 5.5: True value of the geometric parameters.

The numerical experiments conducted through the fixed shape prior showcase a strong performance within the iterative method. Figure 5.19 gives an indication of this where we see that the learning occurs fairly quickly, as seen in the third subfigure which is the 11^{th} iteration. Considering this with the data misfit and

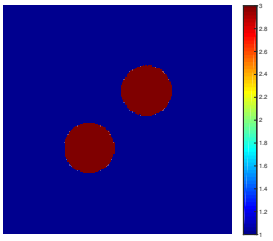

Figure 5.18: Fixed shape truth.

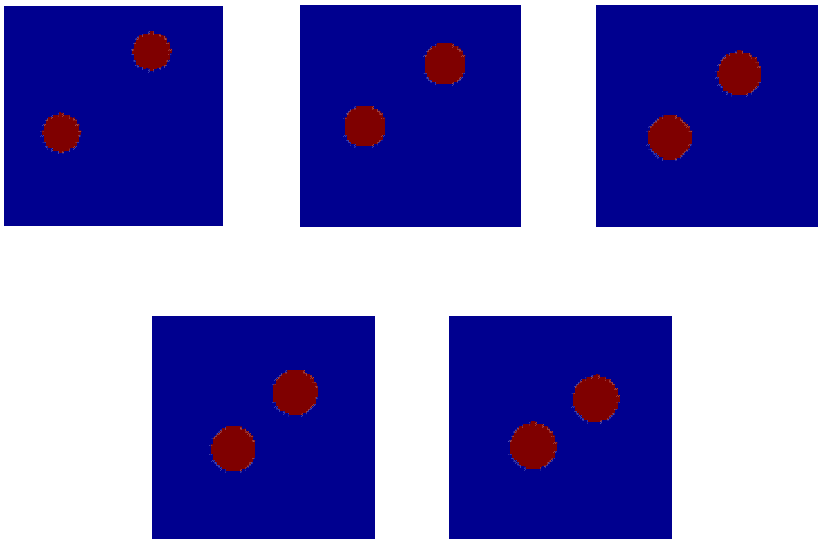

Figure 5.19: Reconstruction of fixed shape prior.

relative error in Figure 5.20, we see a strong decline in both these quantities with a final relative error of just under 0.1. Finally we repeat the experiment where four of the best initializations are presented in Figure 5.21, which is based on Table 5.2. Each final solution of the initializations produces a strong reconstruction with respect to the truth. This, with the performances discussed above, can be attributed to the restrictions given on the initial conditions of the prior presented in Table 5.2.

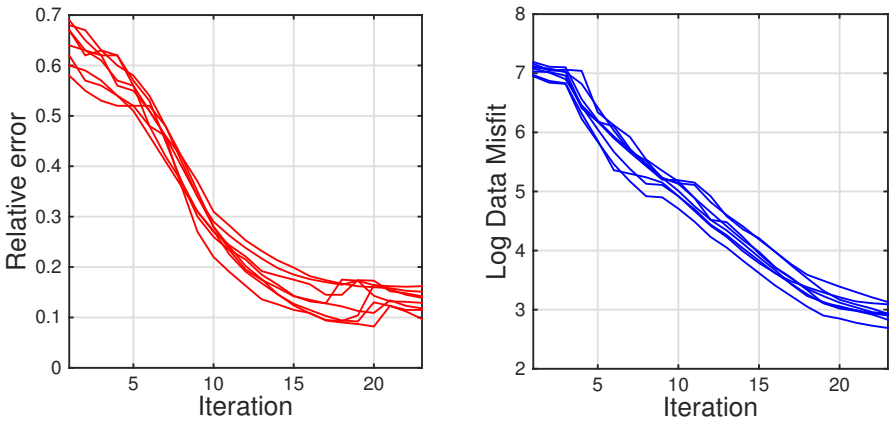

Figure 5.20: Fixed shape prior. Left: Relative error. Right: Log data misfit.

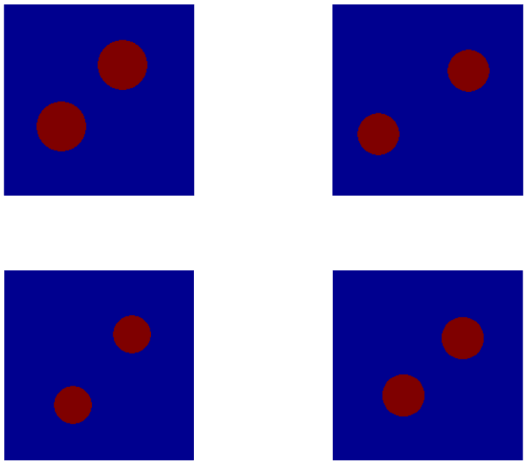

Figure 5.21: Final iteration from four different initializations of fixed shape prior.

5.6 Conclusion

The Bayesian approach to inverse problems is an attractive one, largely due to incorporating methodology from other fields. One of those fields is uncertainty quantification (UQ) which can be aided by both Bayesian techniques and applied mathematics. In this chapter we used some of those UQ techniques to tackle the problem of the inverse eikonal equation. This was achieved by introducing various prior forms that the slowness function could take, ranging from both heteregenous and piecewise constant fields. Under the various priors we showed, despite stricter assumptions on the eikonal equation, that attainable numerical results were achieved

for recovering the slowness function. It is still to be shown analytically whether we can attain well-posedeness both for the inverse and forward problem. Analyzing the inverse eikonal equation leads to a number of interesting further directions of research, namely

- The area of pedestrian dynamics has sparked recent attention. The models used in pedestrian dynamics are commonly referred to as either macroscopic or microscopic models. One model which is have been analyzed is the Hughes model which is given as

$$\frac{\partial \rho}{\partial t} - \nabla \cdot (\rho f(\rho)^2 \nabla \phi) = 0$$
$$\|\nabla \phi\| = \frac{1}{f(\rho)}.$$

  The model is concerned with how quickly a group of pedestrians can reach an exit in a particular region. From the above equation we notice that we have a coupled system with includes the eikonal equation. Given the work that has been done, applying those tools to the inverse Hughes model would be a nice extension. Other insights which would be useful in this context would be to characterize a mean field limit.

- In the context of the eikonal equation, it is common to model the slowness function as a piecewise constant function. Due to this instead of using level set techniques, it may be worth pursing a phase field approach. This offers a much richer understanding mathematically of the problem, but would be more beneficial to consider this in a deterministic setting. This would mea considering the solution as a minimizer of the functional

$$J_\epsilon(u) := \frac{1}{2}\int_\Gamma |T(x) - T_{\text{obs}}(x)|^2 ds_x + \sigma \int_\Omega \Big(\frac{\epsilon}{2}|\nabla u|^2 + \frac{1}{\epsilon}\phi(u)\Big) dx,$$

  which can be interpreted as the function (5.4.22) with an additional phase-field regularization term.

Chapter 6

Conclusion & discussion

The motivation behind this thesis was to understand and develop methodology for ensemble Kalman inversion, where we addressed various avenues of research. Our results and findings for each chapter are summarized below:

- **Chapter 2.** - Our interest for this thesis chapter was to develop hierarchical approaches for ensemble Kalman inversion. Taking motivation from Roberts et al. [115, 116], we characterized two approaches; the centred and non-centred approach. We saw that with these approaches, specifically the non-centred one, that we were able to learn not only the underlying field but the hyperparameters that define it. This was tested for a range of inverse problems. The key point from this work is that when understanding ensemble based methods hierarchically, that we are able to break away from the subspace property which can pose a restriction on the inversion through the initial ensemble.

- **Chapter 3.** - This chapter was concerned with developing theory of these hierarchical approaches discussed in Chapter 2. Namely we were interested in deriving diffusion limits (continuous-time) in the linear case, and providing a better formulation and definition of these approaches. This relates to the generation of the Gaussian random fields, where we further analyzed the approaches regarding the subspace property. Lastly we considered a number of variants on the ensemble Kalman filter such as variance inflation and localization. Diffusion limits were further derived here, where we showed through the means of a simple numeric experiment how these techniques act hierarchically.

- **Chapter 4.** - This chapters aim was to improve on a challenge within uncertainty quantification and inverse problems, which is to improve the efficiency of the forward solver. This was done through the reduced basis method, a method which is based on a Galerkin projection. A question we considered for this was the effect of the training set used in the reduced basis method, where we tested this for an elliptic partial differential equation. We then used this to further test this training set in a setting for geometric inverse problems. The improved forward solver showcased an improvement on computation time, without losing much accuracy through the inversion.

- **Chapter 5.** - The final body of work was dedicated to understanding, from a Bayesian approach, the inverse eikonal equation. The eikonal equation has been well studied, however the inverse problem associated with it has seen a lack of literature. We sought to extend the work of [52] which considered a deterministic setting. In particular we tested a number of various priors, such as geometric and log-Gaussian, on the equation and derived a Bayesian formulation of the equation. We showed that with these priors we could get good reconstructions of the slowness function. The motivation here was to understand how inverse HJ equations act under the Bayesian formulation.

Overall the thesis looked to shed light on a number of interesting areas being considered in the inverse problem community. As mentioned one of these is the issue with the forward solver where, for more computational intensive partial differential equations, alternatives are being proposed to traditional FEMs. We considered the RBM which has been well studied and showed improvements on the computation of experiments. However we still think that in a Bayesian setting, that the RBM could be studied further to gain an insight, as most of work has been done in a deterministic fashion. It would be of interest consider other ROM methods, such as Gaussian process regression, in a more computational demanding manner.

Hierarchical learning has been a recent trend as well within the inverse problem community, where the aim is to produce good enough priors to negate the influence of the data-likelihood. In order to do this, hierarchical inverse problems usually involve a number of unknowns to reconstruct. In the context of EKI we say that we could do this well, for a number of hyperparameters, whilst not sacrificing any efficiency in the experiments. As of now other hierarchical methods are being considered and developed for various solvers such as MAP estimation and Gaussian process regression. We note again that each hierarchical approach taken in can be interpreted in a different manner. This includes the analysis for each corresponding

hierarchical solver.

As a side topic, with the huge improvements on the developments of methods for Bayesian inverse problems, it would be of interest to test and clarify an understanding of further mathematical models. Thus far elliptic PDEs have been (and are still) being mainly considered, but moving to non-variational formulation cases poses many interesting questions. This was one of the reasons for initiation on the work of the eikonal equation. We combined methodologies of both EKI and various priors and noticed similar performances with other models.

Finally we summarize our findings and points discussed above by answering the questions from Chapter 1 namely:

- Can we effectively use EKI to recover some form of geometry from our unknown, which includes various discontinuities?
  Yes, this is shown through the various prior forms which we introduce such as level set, piecewise constant and continuous priors. Furthermore this has been shown on numerous PDEs where the underlying unknown can be represented through a geometric structure, such as EIT and groundwater flow. This was evident from Chapters 2, 4 and 5.

- Given the computational burden that can arise with the forward operator, can we reduce this cost in a practical manner and what kind of cost reduction is possible?
  To some degree yes this was tackled well where we introduced the RBM in Chapter 4. The motivation was to exactly reduce this computational burden where we saw a reduction in cost. However the limitations on this are that this was not compared with other reduced order models, and that the RBM is PDE specific.

- Can we transfer ideas from hierarchical Bayesian inversion to the EKI methodology, and in doing so find improved inversion strategies?
  Yes, this was shown for the work conducted in both Chapter 2 and 3. Much of the methodology was taken from computational statistics, where both assumed that the prior was of a Gaussian form. As a result certain hyperparameters were successfully recovered. This was loosely extended to the case where the hyperparameters were modelled as a random field and not a constant value.

- Given the current extent of the literature on Bayesian-related inversion, and the EKI in particular, can we extend this to new and potentially more challenging PDEs applications?
  Yes, this was done solely in Chapter 5 where we considered the extension to an inverse eikonal equation. The eikonal equation appears in numerous scientific areas which has further mathematical constraints compared to other PDEs tested. Our aim was to test how various prior forms could tackle the uncertainty within the model, through a numerical investigation, which was successfully achieved.

To conclude we highlight a number of future research directions which can be taken regarding inverse problems.

6.1 Further Areas of Research

- **Machine learning** - Big data and machine learning have seen a recent trend, due to the large amounts of data that is now available. With this comes the question of how to efficiently process it and gain inference. Many techniques which are used in Bayesian inverse problems seem natural to apply to these fields. Thus far the topics of mathematical classification and clustering has been addressed through inverse problems, both analytically and computationally [14, 48, 61, 62]. One way this has been done is the incorporation of MCMC methods in machine learning. The issue with this, and especially in machine learning problems, is the computational burden. Developing more efficient UQ methods for machine learning is a challenging task currently, but one with much potential.

- **Optimal transport** - A key issue that still lies within MCMC methods is the cost. Despite this issue it still very favourable due to the ability to reconstruct images with various prior forms. This problem is nullified with ensemble based methods, however its disadvantage lies with the accuracy and large ensemble limit. A new class of methods which are being developed are to negate any major assumptions on the prior and posterior. These methods are based on an optimal transport problem which considers a linear programming problem. There has been significant research into optimal transport maps for inverse problems and data assimilation [51, 103, 105]. A potential step forward could be to combine the ideas of EKI with optimal transport.

- **Pedestrian dynamics** - Pedestrian dynamics based on PDE models have sparked an interest in the applied mathematics community. In the context of this field developing and analzying inverse problems is a new direction which can be taken. The eikonal equation which was the focus in Chapter 6. appears in pedestrian models, notably the Hughes model [27, 30]. This model can be be interpreted as a coupled system between the eikonal equation and a reaction diffusion equation. So through the work done in Chapter 5., this would seem a natural extension to which it opens up with the additional theory of pedestrian dynamics. As other models are based on an SDE formulation, we could apply techniques used as well.

- **Non-Gaussian priors** - A key issue that is still apparent in Bayesian inverse problems is how to recover rough features and edges. Examples which have been used include the Besov prior and the Laplace prior. However in the context of numerically testing these, it is not necessarily clear how the sampling done. Given how the α-stable priors [102] have shown to work well for edge-preserving [33, 123], a direction to take with this is to build both theory and a numerical investigation of these priors. In terms of the theory given how these processes incorporate a variety of different priors, it motivates questions in statistical theory such as posterior convergence, consistency and contraction rates.

Appendices

Appendix A

Levenberg-Marquart algorithm

For non-linear optimization, the Levenberg-Marquart algorithm (LMA) seeks to minimize a function $F(x)$, which is the sum of squares of non-linear functions

$$F(x) = \frac{1}{2}\sum_{j=1}^{m} f_j(x)^2. \tag{A.0.1}$$

By denoting the Jacobian matrix as $J_i(x)$ of $f_i(x)$, the LMA uses a search direction, based on derivative information, in the direction of the solution p which satisfies

$$(J_k^T J_k + \lambda_k I)p_k = -J_k^T f_k, \tag{A.0.2}$$

where λ_k are a set of non-negative scalars, commonly referred to as the damping factor. The LMA can be viewed as combination between traditional gradient desent and the Gauss-Newton method. Like many optimization methods it works by finding a local minimizer as opposed to a global minimizer.

The LMA applied to inverse problems can be thought of as a regularization scheme, where the solution is of a similar form to (A.0.2), but with the addition of a discrepancy principle. This was first discussed by Hanke [71] where the solution of the inverse problem

$$y = \mathcal{G}(u) + \eta,$$

can be characterized as

$$u_{j+1} = u_j - \left(\alpha_j I + \mathcal{G}'(u_j)^T \mathcal{G}(u_j)\right)^{-1} \mathcal{G}'(u_j)^T (\mathcal{G}(u_j) - y), \tag{A.0.3}$$

with discrepancy principle

$$\|\mathcal{G}(u_j) - y\| \leq \tau\eta \leq \|\mathcal{G}(u_{j-1}) - y\|. \tag{A.0.4}$$

From (A.0.3), α_k is a sequence of posititve numbers determined by certain strategies. This inclusion is for the incorporation of the regularization within the iterative formula, which is linked to the discrepancy principle (A.0.4). The discrepancy principle is to ensure stable solutions which is controlled through the noise level and $\tau > 1$, which is a tuneable parameter.

The LMA is of particular interest to our setting as the update formula is reminiscent to that of the EKI, with the addition of the positive numbers α_k. As a result it is a natural choice of regularization for the EnKF.

Appendix B

Eikonal equation

In this appendix we aim to derive the solution of the eikonal equation

$$|\nabla T(x)| = u(x), \tag{B.0.1}$$

as a minimization procedure through the theory of optimal control. This is achieved as the solution to Hamilton-Jacobi equations can be characterized through Poincaré-Cartan integrals, or path intergrals. However in order to do so we will make use of the method of characteristics. Firstly given (B.0.1) is a Hamilton-Jacobi equation, we can rewrite the equation using its corresponding Hamiltonian, i.e.

$$H(x, \nabla T(x)) = u(x),$$

where $H : \mathbb{R}^n \to \mathbb{R}$ is defined as the Hamiltonian. A number of assumptions are required on the Hamiltonian, to ensure we can characterize a solution to (B.0.1), such as

- H is strictly convex.
- $H : \mathbb{R}^n_p \to \mathbb{R}$ is a function of p.
- $\lim_{|p|\to\infty} \frac{H(p)}{p} = 0$.

Under these assumptions we can define a Lagrangian by

$$L(u) := \sup_{u\in\mathbb{R}^n} \{p \cdot u - H(p)\}, \tag{B.0.2}$$

This arises from the Hamilton which can be written through the Lagrangian in (B.0.2), so

$$H(p) := \sup_{p\in\mathbb{R}^n} \{p \cdot u - L(u)\},$$

Through the properties of a Legendre transform, we can express the eikonal equation as

$$\max_{p\in\mathbb{R}^n}\{p\cdot\nabla T(x)-L(u)\}=u(x). \tag{B.0.3}$$

Now given how our usual source term/slowness function is $f(x)=u(x)$, we can express a general optimization problem for $T(x)$, from (B.0.3),

$$T(x)=\min_u\left\{\int_0^t L(u(s))ds\right\}=\min_\zeta\left\{\int_0^t L(\dot{\zeta}(s))ds\right\},$$

where the minimum is taken over all paths $\zeta(s)$ such that $\zeta(0)=x_0$ and $\zeta(t)=x$. From the theory of Lagrangian mechanics, as the Lagrangian depends only on u and not x, the minimum can be interpreted as a straight line path.

Now through a slight change of notation where we place $r=s$ and set our interval to be $[0,1]^d$ for $d\geq 1$, our solution to the eikonal equation (B.0.1) can be expressed as

$$T(x)=\inf_\zeta\left\{\int_0^1 u(\zeta(r))|\zeta'(r)|dr \mid \zeta\in W^{1,\infty}([0,1],\bar{\Omega}),\zeta(0)=x_0,\zeta(1)=x\right\}.$$

Bibliography

[1] S. I. Aanonsen, G. Naevdal, D. S. Oliver, A. C. Reynolds and B. Valles. The Ensemble Kalman Filter in Reservoir Engineering-a Review, *SPE J.* **14**, 393-412, 2009.

[2] A. Adler and W. R. B. Lionheart. Uses and abuses of EIDORS: An extensible software base for EIT, *Physiol Meas* **27**, S25-S42, 2006.

[3] M. A. Anastasio, J. Zhang, D. Modgil and P. J. La Rivière. Application of inverse source concepts to photoacoustic tomography, *Inverse Problems*, **(23)**(6), 2007.

[4] J. L. Anderson. An ensemble adjustment Kalman filter for data assimilation, *Monthly Weather Review*, **129**, 28-84, 2002.

[5] J. L. Anderson and S. L. Anderson. A Monte Carlo implementation of the non-linear filtering problem to produce ensemble assimilations and forecasts, *Monthly Weather Review*, **127**, p2741-2758, 1999.

[6] S. Agapiou, J. M. Bardsley, O. Papaspiliopoulos and A. M. Stuart. Analysis of the Gibbs sampler for hierarchical inverse problems, *SIAM Journal on Uncertainty Quantification*, **1**, p511-544, 2014.

[7] S. Agapiou, M. Burger, M. Dashti and T. Helin. Sparsity-promoting and edge-preserving maximum a posteriori estimators in non-parametric Bayesian inverse problems, arXiv:1705.03286, 2017.

[8] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso and A. M. Stuart. Importance sampling: computational complexity and intrinsic dimension,, *To appear in Statistical Science*, 2017.

[9] I. Babuska, F. Nobile and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM Journal on Numerical Analysis*, **45** (3), 1005-1034, 2007.

[10] A.B. Bakushinsedky and M. Y. Kokurin Iterative Methods for Approximate Solution of Inverse Problems, *Mathematics and its applications*, Springer, 2004.

[11] M. Bardi and I. Capuzzo-Dolcetta. Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations, *Modern Birkhauser Classics. Birkhauser, Boston Basel Berlin*, 2008.

[12] F. Bauer, T. Hohage and A. Munk. Iteratively Regularized Gauss-Newton Method for Nonlinear Inverse Problems with Random Noise, *SIAM J. Numer. Anal* **47**, p1827-1846, 2009

[13] M. Benning and M. Burger. Modern regularization methods for inverse problems, *Acta Numerica*, **27**, 2018.

[14] A .L. Bertozzi, X. Luo, A. M. Stuart and K. C. Zygalakis. Uncertainty Quantification in the Classification of High Dimensional Data, arXiv:1703.08816, 2017.

[15] A. Beskos, M. Girolami, S. Lan, P. E. Farrell and A. M. Stuart. Geometric MCMC for infinite-dimensional inverse problems, *Journal of Computational Physics,* Vol **335**, p327-351, 2017.

[16] A. Beskos, A. Jasra, E. Muzaffer, A. M. Stuart. Sequential Monte Carlo Methods for Bayesian Elliptic Inverse Problems,*Statistics and Computing*, Volume 25, Issue 4, p727-737, 2014.

[17] A. Beskos, G. Roberts and A. M. Stuart. Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions,*Ann.Appl. Probab.*, **19**, p863-898, 2009.

[18] A. Beskos, G. Roberts, A. M. Stuart and J. Voss. MCMC methods for diffusion bridges, *Stochastics and Dynamics*, pp. 319-350, 2008.

[19] D. Blomker, C. Schillings and P. Wacker, A strongly convergent numerical scheme from EnKF continuum analysis, arXiv:1703.06767, 2017.

[20] V. I. Bogachev. *Gaussian measures*, Vol 62. American Mathematical Society Providence, 1998

[21] l. Borcea. Electrical Impedance Tomography, *Inverse Problems* **18**, 2002.

[22] T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler. A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems Part I: The Linearized Case, with Application to Global Seismic Inversion, *SIAM J. Sci. Comput.*, **35** (6), A2494-A2523, 2013.

[23] M. Burger. A framework for the construction of level set methods for shape optimization and reconstruction. *Interfaces and Free Boundaries*, Vol **5**, p301-329, 2002.

[24] M. Burger and T. Helin. Maximum a posteriori probability estimates in infinite-dimensional Bayesian inverse problems, *Inverse problems* **31** (8), 2015.

[25] M. Burger and S. Osher. A survey on level set methods for inverse problems and optimal design. *Eur. J. Appl. Math*, Vol **16**, p263-301, 2005.

[26] O. Burkovska, B. Haasdonk, J. Salomon, and B. Wohlmuth. Reduced Basis Methods for Pricing Options with the Black-Scholes and Heston Models, *SIAM J. Finan. Math.,* **6** (1), 685-712, newblock 2015.

[27] F. Camilli, A. Festa and S. Tozza. A discrete Hughes' model for pedestrian flow on graphs,*Networks and Heterogeneous Media*, Vol **12**, Issue 1, Pages: 93 - 112, 2017.

[28] I. Capuzzo-Dolcetta and P. L. Lions. Hamilton-Jacobi equations with state constraints, *Trans. Amer. Math. Soc.*, 318:643-683, 1990.

[29] J. Carera and S. P. Neuman. Estimation of aquifer parameters under transient and steady state conditions: application to synthetic and field data, *Water Resources Research*, **22**(2):228-242, 1986.

[30] J. A. Carillo, S. Martin, M. Wolfram. *A local version of the Hughes model for pedestrian flow*,arXiv:1501.07054, 2015.

[31] N. K. Chada. Analysis of hierarchical ensemble Kalman inversion, arXiv:1801.00847, 2018.

[32] N. K. Chada, K. Deckelnick, C. M. Elliott, A. M. Stuart and V. Styles. Bayesian formulation of the inverse eikonal equation, *in preperation.*

[33] N. K. Chada, M. A. Iglesias, L. Roininen and A. M. Stuart. Parameterizations for ensemble Kalman inversion, *To appear in Inverse Problems.*

[34] V. Chen , M. M. Dunlop O. Papaspiliopoulos and A. M. Stuart. Robust MCMC Sampling with Non-Gaussian and Hierarchical Priors in High Dimensions, *in preparation*, 2017.

[35] P Chen, A Quarteroni and G Rozza. Comparison between reduced basis and stochastic collocation methods for elliptic problems, *Journal of Scientific Computing,* **59** (1), 187-216, 2014.

[36] Y. Cheng and S. Reich. Assimilating data into scientific models: An optimal coupling perspective, In: Nonlinear Data Assimilation. Frontiers in Applied Dynamical Systems: Reviews and Tutorials, Vol 2. Springer, 2015.

[37] C. W. Clenshaw and A. R. Curtis. A Method for Numerical Integration on an Automatic Computer, *Numer. Math*, **2**, p197-205, 1960.

[38] B. Cockburn, G. E. Karniadakis and C. W. Shu. Discontinuous Galerkin methods. Theory, computation and applications, *Lecture Notes in Computational Science and Engineering*, 11. Springer-Verlag, Berlin, 2000.

[39] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. MCMC Methods for Functions: Modifying Old Algorithms to Make Them Faster. Statistical Science, Volume **28**, Number 3, 424-446, 2013.

[40] A. Cohen and R. DeVore. Approximation of high-dimensional parametric PDE, *Acta Numerica,* **24**:1-159, 2015.

[41] K. Deckelnick, C. M. Elliott and V. Styles. Double obstacle phase field approach to an inverse problem for a discontinuous diffusion coefficient. *Inverse problems*, **32**(4), 2016.

[42] A. J. Pierre Del Moral and A. Doucet. Sequential Monte Carlo samplers, *Journal of the Royal Statistical Society*, Series B 411-436, 2006.

[43] A. Doucet, N. de Freitas and N. J. Gordon. Sequential Monte Carlo Methods in Practice, *Statistics for Engineering and Information Science,* New York: Springer-Verlag, 2001.

[44] K. Deckelnick and C. M. Elliott. Uniqueness and error analysis for Hamilton-Jacobi equations with discontinuities. *Interfaces and Free Boundaries*, **6**, 329-349, 2004.

[45] K. Deckelnick, C. M. Elliott and V. Styles. Double obstacle phase field approach to an inverse problem for a discontinuous diffusion coefficient, *Inverse problems*, **32**(4), 2016.

[46] M. Dashti, K. J. H. Law, A. M. Stuart and J. Voss. MAP estimators and their consistency in Bayesian nonsparametric inverse problems, *Inverse problems*, **29** (9), 2013.

[47] M. M. Dunlop, M. A. Iglesias and A. M. Stuart. Hierarchical Bayesian Level Set Inversion, *Statistics and Computing*, 2016.

[48] M. M. Dunlop, D Slepcev, AM Stuart and M Thorpe. Large data and zero noise limits of graph based semi-supervised learning algorithms, *In preparation.*

[49] M. M. Dunlop and A. M. Stuart. Bayesian formulation of EIT: analsis and algorithms, *Inverse problems*, 2016.

[50] M. M. Dunlop and A. M. Stuart. MAP estimators for piecewise continuous inversion, Inverse Problems, **32**, 2016

[51] T. A. El Moselhy and Y. M. Marzouk. Bayesian Inference with Optimal Maps, *Journal of Computational Physics*, Volume **231**(23), Pages 7815-7850, 2012.

[52] C. M. Elliott, K. Deckelnick and V. Styles. Numerical analysis of an inverse problem for the eikonal equation. *Numerische Mathematik*, 2011.

[53] A. A. Emerick. Towards a hierarchical parametrisation to address prior uncertainty in ensemble-based data assimilation, *Computational Geosciences*, **20**:35-47, 2016.

[54] H.W. Engl, K. Hanke and A. Neubauer. Regularization of inverse problems, *Mathematics and its Applications*, Volume **375**, Kluwer Academic Publishers Group, Dordrecht, 1996.

[55] H.W. Engl, K. Kunisch and A. Neubauer. Convergence rates for Tikhonov regularisation of non-linear ill-posed problems, *Inverse problems*, **5**, 523 1989.

[56] O. Ernst, and B. Sprungk and H. Starkloff. Analysis of the ensemble and polynomial chaos Kalman filters in Bayesian inverse problems, *SIAM-ASA JUQ.,* **3** 823-851, 2015.

[57] G. Evensen. Data Assimilation: The Ensemble Kalman Filter. *Springer*, 2009.

[58] G. Evensen Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geophysical Research: Oceans*, **99**, p10143-10162, 1994

[59] G. Evensen and P. J. Van Leeuwen. Assimilation of geosataltimeter data for the agulhas currentusing the ensemble Kalman filter with a quasi-geostrophic model, *Monthly Weather Review*, **128**, p85-86, 1996.

[60] L. C. Evans. Partial Differential Equations. *American Mathematical Society*, 1998.

[61] N. Garcia Trillos, Z. Kaplan, T. Samakhona and D. Sanz-Alonso. Graph-based Bayesian Learning for Semi-supervised Regression and Inverse Problems, *in preperation.*

[62] N. Garcia Trillos and D. Sanz-Alonso. Continuum Limit of Posteriors in Graph Bayesian Inverse Problems, arXiv:1706.07193, 2017.

[63] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods,*Journal of the Royal Statistical Society, Series B (Statistical Methodology).* **73** (2): 123-214, 2011.

[64] F. Le Gland, V. Monbet and V. D. Tran. Large sample asymptotics for the ensemble Kalman filter, *The Oxford Handbook of Nonlinear Filtering*, Oxford University Press, p598-631, 2011.

[65] D. Garmatter, B. Haasdonk and B. Harrach. A Reduced Basis Landweber method for nonlinear inverse problems, *Inverse Problems*, **32** (3), 2016.

[66] G. Gaullier, P. Charbonnier, F. Heitz and P. Cote. Introducing shape constraints into object- based traveltime tomography. *Inverse problems*, **32**(9), 2016.

[67] M. B. Giles. Multilevel Monte Carlo methods, *Acta Numerica*, **24**, 259, 2015.

[68] M. Gunzburger and A. L. Teckentrup. Optimal Point Sets for Total Degree Polynomial Interpolation in Moderate Dimensions, arXiv:1407.3291, 2014.

[69] H. Haario, M. Laine, M. Lehtinen, J. Tamminen. MCMC methods for high dimensional inversion in remote sensing,*Journal of the Royal Statistical Society Series B (Statistical Methodology)* **66** (3): p591-607, 2004.

[70] T. M. Hamill and J. S. Whitaker. Accounting for error due to unresolved scales in ensemble data assimilation: a comparison of different approaches, *Monthly Weather Review*, **133**, p3132-3147, 2004.

[71] M. Hanke. A regularizing Levenberg-Marquardt scheme, with applications to inverse groundwater filtration problems, *Inverse Problems,* **13**, p79-95, 1997.

[72] C. Hansen. *Discrete Inverse Problems: Insight and Algorithms: Insight and Algorithms*, SIAM: Fundamentals of Algorithms, 2010.

[73] V. H. Hoang. Bayesian inverse problems for Burgers and Hamilton-Jacobi equations with white-noise forcing, *Inverse Problems*, **(28)**2, 2012.

[74] P. L. Houtekamer and H. L. Mitchell. A sequential ensemble Kalman filter for atmospheric data assimilation, *Monthly Weather Review*, **129**, 123-137, 2001.

[75] H. Hu, J. Sunu and A. Bertozzi. Multi-class graph Mumford-Shah model for plume detection using the MBO scheme, *Lecture Notes in Computer Science*, 2015; 8932.

[76] D.B.P. Huynh, G. Rozza, S. Sen, A.T. A. Patera. Successive Constraint Linear Optimization Method for Lower Bounds of Parametric Coercivity and inf-sup Stability Constants. *Comptes Rendus Mathematique, Analyse Numerique* **345**(8), pp 473-478, 2007.

[77] M. A. Iglesias. A regularising iterative ensemble Kalman method for PDE-constrained inverse problems, *Inverse Problems*, **32**, 2016.

[78] M. A. Iglesias, K. J. H. Law and A. M. Stuart. Evaluation of Gaussian approximations for data assimilation in reservoir models, *Computational Geosciences*, Volume **17**(5) , pp 851-885, 2013.

[79] M. A. Iglesias, K. J. H. Law and A. M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, **29** 2013.

[80] M. A. Iglesias, K. Lin and A. M. Stuart. Well-posed Bayesian geometric inverse problems arising in subsurface flow, *Inverse problems*, **30**, 2014.

[81] M. A. Iglesias, Y. Lu and A. M. Stuart. A Bayesian Level Set Method for Geometric Inverse Problems . *Interfaces and Free Boundary Problems*, 2015.

[82] H. Ishii. A simple, direct proof of uniqueness for solutions of the Hamilton-Jacobi equations of eikonal type, *Proc. Amer. Math. Soc.,* **100**:247-251, 1987.

[83] H. Ishii. Hamilton-Jacobi equations with discontinuous Hamiltonians an arbitrary open sets, *Bull. Fac. Sci. Engrg. Chuo. Univ.,* 28, pp. 33-77, 1985.

[84] J. Jia, S. Yu, J. Peng and J. Gao. Infinite-dimensional Bayesian approach for inverse scattering problems of a fractional Helmholtz equation, arXiv:1603.04036, 2016.

[85] B. Jin. Conjugate gradient method for the Robin inverse problem associated with the Laplace equation, *Interational Journal of Numerical Methods in Engineering*, 2007.

[86] S. Jin, P. Markowich, and C. Sparber. Mathematical and computational methods for semi- classical schrodinger equations, *Acta Numerica*, **20** pp. 121?209, 2011.

[87] J. Kaipio and E. Somersalo. *Statistical and Computational Inverse problems.* Springer Verlag, New York, 2004.

[88] R. E. Kalman. A new approach to linear filtering and prediction problems, *Trans ASME (J. Basic Engineering)*, **82**, p35-45, 1960.

[89] N. Kantas, A. Beskos and A. Jasra. Sequential Monte Carlo Methods for High-Dimensional Inverse Problems: A case study for the Navier-Stokes equations, *SIAM/ASA J. Uncertain. Quantif,* **2**, 464-489, 2014.

[90] K. J. H Law and A.M. Stuart. Evaluating Data Assimilation Algorithms, *Mon. Weather Rev*, **140**, p37-57 2012.

[91] K. J. H. Law, A. M. Stuart and K. Zygalakis. Data Assimilation: A Mathematical Introduction, *Springer*, 2015.

[92] M.S. Lehtinen, L. Paivarinta and E. Somersalo. Linear inverse problems for generalised random variables, *Inverse Problems*, **5**(4), p599-612, 1989.

[93] G. Li and A. C. Reynolds. Iterative ensemble Kalman filters for data assimilation, *SPE J* Vol. **14**, p496-505, 2009.

[94] P. Liang, M. I. Jordan and D. Klein. Learning Programs: A Hierarchical Bayesian Approach, *Conference: Proceedings of the 27th International Conference on Machine Learning.* 2010.

[95] P.L. Lions. *Generalized solutions of Hamilton-Jacobi equations*, Pitman, Boston, 1982.

[96] F. Lindgren, H. Rue and J. Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach, Volume **73**, p423-498, 2011.

[97] N. Liu and D. S. Oliver. Ensemble Kalman filter for automatic history matching of geologic facies, *Journal of Petroleum Science and Engineering*, **47**, p147-161, 2005.

[98] W. Liu, J. Li and Y. M. Marzouk. An approximate empirical Bayesian method for large-scale linear-Gaussian inverse problems, arXiv:1705.07646, 2017.

[99] G. Lord, C.E. Powell and T. Shardlow. *An Introduction to Computational Stochastic PDEs,* Cambridge Texts in Applied Mathematics, 2014.

[100] J. Mandel, L. Cobb, and J. D. Beezley. On the convergence of the ensemble Kalman filter, *Applications of Mathematics*, **56**, p 533-541, 2011.

[101] A. Majda and X. Wang. *Non-linear Dynamics and Statistical Theories for Basic Geophysical Flows,* Cambridge University Press, 2006.

[102] M. Markkanen, L. Roininen, J. M. J. Huttunen and S. Lasanen Cauchy difference priors for edge-preserving Bayesian inversion with an application to X-ray tomography, arXiv:1603.06135, 2015.

[103] Y. Marzouk, T. Moselhy, M. Parno, A. Spantini. An introduction to sampling via measure transport, *Handbook of Uncertainty Quantification*, Springer, 2016.

[104] J. C. Mattingly, N. S. Pillai and A. M. Stuart. Diffusion limits of the random walk Metropolis algorithm in high dimensions,*Ann. Appl. Probab.* **22**, 881-930, 2012.

[105] A. Myers, T. Bui-Thanh, K. Wang and A. Thiéry. A Novel discrete optimal transport problem for bayesian inverse problems, *in preparation.*

[106] I. Myserth and H. Omre. Hierarchical Ensemble Kalman Filter, *SPE Journal*, vol **15**, issue 02, 2010.

[107] A. Neubauer Tikhonov regularisation for non-linear ill-posed problems: optimal convergence rates and finite-dimensional approximation, *Inverse problem*, **5**, 541 1989

[108] F. Negri, G. Rozza, A. Manzoni and A. Quarteroni. Reduced Basis Method for Parametrized Elliptic Optimal Control Problems, *SIAM J. Sci. Comput.,* 35(5), A2316-A2340, 2013.

[109] R. L. Nowack. Wavefronts and solutions of the eikonal equation, *Geophysical Journal International*, Volume 110, Issue 1, p55?62, 1992.

[110] D. Oliver, A. C. Reynolds and N. Liu. *Inverse Theory for Petroleum Reservoir Characterization and History Matching*, Cambridge University Press, 1st edn, 2008.

[111] Daniel N. Ostrov. Extending Viscosity Solutions to Eikonal Equations with Discontinuous Spatial Dependence, *Nonlinear Analysis*, Vol. 42, no. 4, 2000.

[112] Daniel N. Ostrov. Solutions of Hamilton-Jacobi Equations and Scalar Conservation Laws with Discontinuous Space-time Dependence, *Journal of Differential Equations*, 182, no. 1, pp. 51-77, 2002.

[113] J. S. Ovall. Function, Gradient, and Hessian Recovery Using Quadratic Edge-Bump Functions, *SIAM Journal on Numerical Analysis*, Vol. **45**(3), 1064-1080, 2007.

[114] S. Pagani, A. Manzoni and A. Quarteroni. A reduced basis ensemble Kalman filter for state/parameter identification in large-scale nonlinear dynamical systems, 2017.

[115] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. Non-centered parameterisations for hierarchical models and data augmentation, *In Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, Oxford University Press, USA, 2003.

[116] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. A general framework for the parametrization of hierarchical models, *Statistical Science,* pages 59-73, 2007.

[117] N. Petra, J. Martin, G. Stadler, O. Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems. Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems, *SIAM J. Sci. Comput.*, 36(4), A1525-A1555, 2014.

[118] N. Pillai, A. M. Stuart and A. H. Thiéry. Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions, *Ann. Appl. Probab.*, **22**, 2320-2358, 2012.

[119] A. Quarteroni and G. Rozza. *Reduced Order Methods for Modeling and Computational Reduction*, Springer, Volume 9, 2014.

[120] S. Reich. A non-parametric ensemble transform method for Bayesian inference, *SIAM J. Sci. Comput.,* **35**(4), A2013-A2024, 2013.

[121] C. Robert and G. Casella. *Monte Carlo statistical methods,* Springer Science & Business Media, 2013.

[122] G. O. Roberts and J. S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions, *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, **60** (1): 255-268, 1998.

[123] L. Roininen, M. Girolami, S. Lasanen and M. Markkanen. Hyperpriors for Matérn fields with applications in Bayesian inversion, arXiv:1612.02989, 2016.

[124] L. Roininen, J. M. J. Huttunen and S. Lasanen. Whittle-Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography, *Inverse problems and Imaging*, **8**, 2014.

[125] R. Ramlau. A modified landweber method for inverse problems, *Numerical Functional Analysis and Optimization*, **20**:1-2, p79-98, 2007.

[126] G. Rozza, D.B.P. Huynh and A. T. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Archives of Computational Methods in Engineering*, 2008.

[127] F. Ruggeri, Z. Sawlan, M. Scavino and R. Tempone. A Hierarchical Bayesian Setting for an Inverse Problem in Linear Parabolic PDEs with Noisy Boundary Conditions, *Bayesian Anal*, 2016.

[128] F. Santosa. A level-set approach for inverse problems involving obstacles. *The European Series in Applied and Industrial Mathematics: Control, Optimization and Calculus of Variations*, Vol **1**, p17-33, 1996.

[129] C. Schillings and A. M. Stuart. Analysis of the ensemble Kalman filter for inverse problems, *SIAM J. Num. Anal.*, 2017.

[130] J. A. Sethian. Level set methods and fast marching methods. *Cambridge Monographs on Applied and Computational Mathematics*, Cambridge University Press, 1999.

[131] M. Schweiger, S. Arridge and I. Nissila. Gauss?Newton method for image reconstruction in diffuse optical tomography, *Inverse Problems,* **50**(10), 2005.

[132] H. Silvennoinen. *3D Structure of the crust and upper mantle below northern part of the fennoscandian shield,* PhD Dissertation, 2015.

[133] S.A. Smolyak. Quadrature and interpolation formulas for Tensor Products of Certain Classes of Functions, *Dokl. Akad. Nauk SSSR*, **148**: p1042-1043, 1963.

[134] E. Somersalo, M. Cheney and D. Isaacson. Existence and Uniqueness for Electrode Models for Electric Current Computed Tomography, *SIAM J. Appl. Math.,* **52**, p1023-1040, 1992.

[135] H. M. Soner. Optimal control with state-space constraint, *SIAM Journal on Control and Optimization,* Vol. **24**, p552-561, 1986.

[136] A. Spira and R. Kimmel. An efficient solution to the eikonal equation on parametric manifolds, *Interfaces and Free Boundaries*, **6**, 315?327, 2004.

[137] A. M. Stuart. Inverse problems: A Bayesian perspective. *Acta Numerica*, Vol. **19**, p451-559, 2010.

[138] A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*, Elsevier, 1987.

[139] Y. W. Teh, M. I. Jordan, M. J. Beal, and M. D. Blei. Hierarchical Dirichlet processes, *Journal of the American Statistical Association*, **101**, 1566-1581, 2016.

[140] M. K. Tippett. Ensemble Square Root Filters , *Monthly Weather Review*, **131**, 2002.

[141] T. Tonn, K. Urban and S Volkwein. Comparison of the Reduced-Basis and POD a-Posteriori Error Estimators for an Elliptic Linear-Quadratic Optimal Control Problem, *Mathematical and Computer Modelling of Dynamical Systems*, 2010.

[142] E. Treister and E. Haber. A fast marching algorithm for the factored eikonal equation. *Journal of Computational Physics*, Vol **324**, p210-225, 2016.

[143] M. Tsyrulnikov and A. Rakitko. Hierarchical Bayes Ensemble Kalman Filter, *Physica D: Nonlinear Phenomena*, 2016.

[144] G. Uhlmann. Electrical impedance tomography and Calderón?s problem, *inverse Problems*, 2006.

[145] C. Villani Topics in optimal transportation, *American Mathematical Soc.*, **58**, 2003.

[146] C. Villani. Optimal transport: old and new, *Springer Science & Business Media*, Vol. **338** 2008.

[147] C. R. Vogel. *Computational Methods for Inverse Problems*, Frontiers in Applied Mathematics, 2002.

[148] D. Xiu. *Numerical methods for stochastic computations: a spectral method approach*, Princeton University Press, 5th version, 2010.

[149] Z. Zhang. Polynomial Preserving Gradient Recovery and A Posteriori Estimate for Bilinear Element on Irregular Quadrilaterals, *SIAM J. Sci. Comput.*, **1**(1), 1-24, 2004.

[150] Z. Zhang and A. Naga. A New Finite Element Gradient Recovery Method: Superconvergence Property, *SIAM J. Sci. Comput.*, **26**(4), 1192-1213, 2006.