University of Warwick institutional repository: http://go.warwick.ac.uk/wrap

## A Thesis Submitted for the Degree of PhD at the University of Warwick

http://go.warwick.ac.uk/wrap/3900

This thesis is made available online and is protected by original copyright.
Please scroll down to view the document itself.
Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.


# Applications of MCMC Methods on Function Spaces 

by<br>\section*{Simon L. Cotter}

## Thesis

Submitted to the University of Warwick
for the degree of
Doctor of Philosophy

## Institute of Mathematics

April 2010

## Contents

List of Figures ..... vi
Acknowledgments ..... xv
Declarations ..... xvi
Abstract ..... xvii
Chapter 1 Introduction and Mathematical Preliminaries ..... 1
1.1 Introduction ..... 1
1.1.1 Data Assimilation ..... 2
1.1.2 Filtering ..... 5
1.1.3 Generalised Polynomial Chaos ..... 7
1.1.4 Lagrangian Data Assimilation ..... 8
1.1.5 Shape Registration ..... 9
1.2 Summary of the Thesis ..... 10
1.3 Function Space Settings ..... 13
1.3.1 Function Space Setting For The Stokes Problem ..... 14
1.4 Solutions to the Stokes Problem ..... 15
1.5 The Stokes Operator ..... 18
1.6 Estimate on solutions of Stokes Flow ..... 19
1.7 Random Field Sampling ..... 21
1.8 Regularity of Gaussian Random Fields ..... 23
1.9 Absolute Continuity of Gaussian Measures ..... 26
1.10 Bayesian Statistics ..... 28
1.11 Relationship with Tikhonov Regularisation ..... 31
1.12 Markov Chain Monte Carlo Methods ..... 32
1.12.1 MCMC methods on finite space ..... 32
1.12.2 MCMC methods on function spaces ..... 36
1.13 Computer Generation of (Pseudo-) Random Numbers ..... 44
1.14 Interpolation Methods ..... 44
1.14.1 Bilinear Interpolation ..... 45
1.14.2 Bicubic Interpolation ..... 46
Chapter 2 Eulerian Data Assimilation ..... 48
2.1 Motivation ..... 48
2.2 The Model: Stokes flow ..... 50
2.3 Observational Noise Model ..... 52
2.4 Prior Distribution of $u_{0}$ ..... 53
2.5 The Posterior Distribution ..... 54
2.5.1 Bounds on $\mathcal{G}_{E}$ ..... 55
2.6 The Random Walk Metropolis Hastings Algorithm ..... 56
2.7 Explicit Solutions to Eulerian Data Assimilation ..... 57
2.7.1 Properties of the Eulerian Analytic Posterior ..... 60
2.7.2 Using the Analytical Eulerian Posterior to Assess Information Contained in Observations ..... 69
2.8 General Setup ..... 71
2.9 Limitations of the Standard RWMH method ..... 73
2.10 Validating the Random Walk Algorithm ..... 76
2.11 Inverse Crimes ..... 85
2.12 Properties of the Posterior Measure ..... 88
2.12.1 Decay of Eulerian Data in Time ..... 91
2.13 Conclusions and Future Directions ..... 96
Chapter 3 Lagrangian Data Assimilation ..... 98
3.1 Motivation ..... 98
3.2 The Model ..... 99
3.3 The Posterior Distribution ..... 101
3.3.1 Bounds on $\mathcal{G}_{L}$ ..... 101
3.4 Numerical Approximation of the Operator $\mathcal{G}_{L}$ ..... 104
3.5 General Setup ..... 106
3.6 Limitations of the Standard RWMH method ..... 106
3.7 Validating the Random Walk Algorithm ..... 108
3.8 Inverse Crimes ..... 110
3.9 Properties of the Posterior Measure ..... 111
3.10 Growth of Data in Time ..... 113
3.11 Conclusions and Future Directions ..... 118
Chapter 4 Data Assimilation of Model Error ..... 120
4.1 Motivation ..... 120
4.2 Mismatched Statistical Model and Data Environment ..... 121
4.3 The Noisy Stokes Equations ..... 132
4.4 The Prior Distribution ..... 133
4.5 Bounds on Observation Operators ..... 134
4.5.1 Bounds on $\mathcal{G}_{E N}$ ..... 134
4.5.2 Bounds on $\mathcal{G}_{L N}$ ..... 135
4.6 Sampling Model Error ..... 136
4.7 Value of Data in Assimilation in Model Error ..... 137
4.8 Properties of the Posterior Measure ..... 142
4.8.1 Eulerian Data Assimilation With Model Error ..... 142
4.8.2 Lagrangian Data Assimilation With Model Error ..... 156
4.9 Conclusions ..... 163
Chapter 5 Filtering in Data Assimilation ..... 168
5.1 Motivation ..... 168
5.2 Sequential Sampling ..... 169
5.3 Analytical Posteriors in Eulerian Sequential Sampling ..... 176
5.4 Conclusions ..... 180
Chapter 6 A Data Assimilation Problem in Shape Registration ..... 181
6.1 Motivation ..... 181
6.2 Equations of motion for the curve matching problem ..... 182
6.3 Prior Distribution on $\left(p_{0}, \nu\right)$ ..... 187
6.4 The Observational Noise Model ..... 188
6.5 The Posterior Distribution ..... 188
6.5.1 Properties of the Observation Operator ..... 189
6.6 RWMH with Deterministic Burn-In ..... 199
6.7 Numerical Approximation of $\mathcal{G}$ ..... 201
6.8 General Setup ..... 201
6.9 Numerical Results ..... 203
6.9.1 Posterior Consistency ..... 203
6.9.2 "Real Life" Data ..... 208
6.10 Conclusions ..... 215
Bibliography ..... 218

## List of Figures

1.1 Bilinear Interpolation on a mesh, from [83] ..... 45
2.1 The mean values of $\left\{a_{i}, b_{i}\right\}$ in the posterior distribution with varying. $\sigma$ ..... 60
2.2 Difference in the mean of the posterior from the least squares solution ..... 61
2.3 The diagonal dominance of the covariance matrix in the posterior distri- bution with varying $\sigma$ ..... 62
2.4 Convergence of the posterior covariance to the covariance of the prior distribution as $\sigma \rightarrow \infty$ ..... 63
2.5 Mean values of $\left\{a_{i}, b_{i}\right\}$ in the posterior as number of observations is increased ..... 64
2.6 Frobenius norm of the posterior covariance with varying number of ob- servations ..... 64
2.7 Relative difference with the actual Fourier modes of the norm of the posterior mean with varying number of observations ..... 65
2.8 Frobenius norm of the posterior covariance with varying number of ob- servations ..... 66
2.9 Relative difference with the actual Fourier modes of the norm of the posterior mean with varying number of observation times ..... 66
2.10 Frobenius norm of the posterior covariance with varying number of ob- servation times ..... 67
2.11 Relative difference with the actual Fourier modes of the norm of the expectation of the posterior mean with varying number of observations ..... 68
2.12 Frobenius norm of the expectation of the posterior covariance with vary- ing number of observation times ..... 68
2.13 Values of the Fourier coefficients in mean of the posterior, with 25 ob- servations made at time T, 2 different timescales ..... 70
2.14 Frobenius norm of $\Sigma_{(3,4)}$ with 25 observations made at time $T$ ..... 70
2.15 Relative amount of information available in $m(1)$ and $m(2)$ compared to $m(3)$ and $m(4)$ in posterior with variable observation time ..... 72
2.16 Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, SRWMH ..... 75
2.17 Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, RWMH ..... 76
2.18 Relative 2-norm error of the ergodic average of the vector $x=\left(a_{1}, b_{1}, a_{2}, b_{2}\right) 77$
2.19 Relative Frobenius-norm error of the running covariance matrix $\left(C_{i j}\right)$ of the vector $x=\left(a_{1}, b_{1}, a_{2}, b_{2}\right)$ ..... 78
2.20 Convergence of numerics to analytic distributions ..... 79
2.21 Relative error in the mean, 25 observations made at 32 times up to $\mathrm{T}=1$ ..... 80
2.22 Relative error in the covariance matrix, 25 observations made at 32 times up to $T=1$ ..... 80
2.23 Convergence of Markov chains with different initial states, with Eulerian data ..... 81
2.24 Convergence of Markov chains with different values of $\beta$ ..... 82
2.25 Marginal distributions with and without random $\beta$ ..... 83
2.26 Distribution of the accepted $\beta$ ..... 84
2.27 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, high resolution data, Eule- rian data. ..... 86
2.28 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing variance in the noise model of the algorithm, low actual noise, Eulerian case ..... 87
2.29 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Decreasing variance in the noise model of the algorithm, high variance actual noise, Eulerian case ..... 88
2.30 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, low resolution data, Eulerian case ..... 89
2.31 Increasing numbers of observations in space, Eulerian ..... 90
2.32 Increasing numbers of observations in time, Eulerian. ..... 91
2.33 PDFs For Eulerian Data, 9 Stations, Varying T ..... 92
2.34 Sketch to show decay of $u(x, t)$ into a $\delta$-ball when $f \equiv 0$ ..... 94
3.1 Spaghetti diagram of 20-day drifter trajectory segments. Colours give the mean drift direction (legend in upper-right corner). Taken from [53] ..... 99
3.2 Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, SRWMH ..... 107
3.3 Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, RWMH ..... 108
3.4 Convergence of Markov chains with different initial states, with La- grangian data ..... 109
3.5 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, high resolution data, La- grangian data. ..... 111
3.6 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing variance in the noise model of the algorithm, low actual noise, Lagrangian case ..... 112
3.7 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, low resolution data, La- grangian case ..... 113
3.8 Increasing numbers of observations in space, Lagrangian ..... 114
3.9 Increasing numbers of observations in space, Lagrangian ..... 115
3.10 PDFs For Lagrangian Data, 9 Paths, Varying T ..... 116
3.11 Sketch to show confinement of $z(t)$ into a $\delta$-ball ..... 116
3.12 PDFs Of $u_{0,1}$ For Lagrangian Data, Varying Observation Time, With Only $u_{0,1}$ And $u_{2,2}$ Present ..... 117
3.13 PDFs Of $u_{2,2}$ For Lagrangian Data, Varying Observation Time, With Only $u_{0,1}$ And $u_{2,2}$ Present ..... 118
4.1 Analytical posterior mean values using model $A$ with data from model B, low frequency modes, observation time increasing ..... 123
4.2 Analytical posterior mean values using model A with data from model $B$, high frequency modes, observation time increasing ..... 124
4.3 Analytical posterior covariance convergence to $\Sigma_{0}$ as observation time increases ..... 125
4.4 Expectation of analytical posterior mean values using model A with data from model $B$, low frequency modes ..... 126
4.5 Expectation of analytical posterior mean values using model $A$ with data from model $B$, high frequency modes ..... 127
4.6 Analytical posterior mean values using model A with data from model B, low frequency modes. Forcing dependent on time ..... 129
4.7 Analytical posterior mean values using model A with data from model $B$, high frequency modes. Forcing dependent on time ..... 130
4.8 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing number of observation times, unmatched forcing in data and model, Eulerian data ..... 131
4.9 $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing number of observation times, unmatched forcing in data and model, Lagrangian data ..... 131
4.10 $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case. ..... 143
4.11 $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case. ..... 144
4.12 $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case. ..... 145
4.13 Value of $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ in the Markov chain ..... 146
4.14 Value of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ in the Markov chain ..... 147
4.15 $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case ..... 148
$4.16\left\|\mathbb{E}(u)-u_{\text {Act }}\right\|_{L^{2}}$ : Increasing numbers of observations in space, Eulerian Model Error Case. $u_{\text {Act }}$ is actual initial condition that created the data. ..... 148
4.17 $\left\|\mathbb{E}(\eta)-\eta_{\text {Act }}\right\|_{L^{2}(0, T ; H)}$ : Increasing numbers of observations in space, Eulerian Model Error Case. $\eta_{\text {Act }}$ is the actual forcing functions that created the data. ..... 149
4.18 \|E $(F)-F_{\text {Act }} \|_{L^{2}(0, T ; H)}$ : Increasing numbers of observations in space,Eulerian Model Error Case. $F_{\text {Act }}$ is the actual forcing functions thatcreated the data.149
4.19 $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case. ..... 150
4.20 $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case ..... 150
4.21 $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case. ..... 151
4.22 $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case ..... 151
4.23 $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing ..... 152
4.24 $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing ..... 153
4.25 $\operatorname{Re}\left(\eta_{4,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing ..... 153
4.26 $\operatorname{Re}\left(\eta_{4,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing ..... 154
4.27 $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing ..... 154
4.28 $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing ..... 155
4.29 $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior ..... 156
4.30 $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior ..... 157
4.31 $\operatorname{Re}\left(\eta_{2,2}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior ..... 157
4.32 $\operatorname{Re}\left(\eta_{2,2}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior ..... 158
4.33 $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior ..... 158
4.34 $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior ..... 159
4.35 $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case ..... 161
4.36 $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case. ..... 162
4.37 $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case. ..... 163
4.38 $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case. ..... 164
4.39 $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case ..... 165
4.40 $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case. ..... 166
4.41 $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case. ..... 166
4.42 $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case. ..... 167
5.1 Results of different filtering approaches ..... 174
5.2 Results of different filtering approaches ..... 175
5.3 Results of different filtering approaches ..... 177
5.4 Convergence of $m_{n}$ to the analytical mean of the posterior using the whole data set ..... 178
5.5 Convergence of $\Sigma_{n}$ to the analytical covariance of the posterior using the whole data set ..... 179
6.1 Marginal distributions on $p_{0}(0)$ with increasing numbers of observations ..... 204
6.2 Marginal distributions on $p_{1}(0)$ with increasing numbers of observations ..... 204
6.3 Marginal distributions on $\nu_{0}$ with increasing numbers of observations ..... 205
6.4 Marginal distributions on $\nu_{1}$ with increasing numbers of observations ..... 205
6.5 Distributions of acceptance probabilities in the MCMC method with in- creasing numbers of observations, average acceptance probability tuned to $\approx 25 \%$ ..... 206
6.6 Distribution of $\Phi(\cdot)=\frac{1}{2}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$ in the MCMC method with increas- ing numbers of observations ..... 207
6.7 Distribution of $\frac{1}{N} \Phi(\cdot)=\frac{1}{2 N}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$ in the MCMC method with increasing numbers of observations ..... 207
6.8 Christmas tree shape data ..... 208
6.9 Marginal distributions on $p_{0}(0)$ with increasing numbers of observations ..... 209
6.10 Marginal distributions on $p_{1}(0)$ with increasing numbers of observations ..... 209
6.11 Marginal distributions on $\nu_{0}$ with increasing numbers of observations ..... 210
6.12 Marginal distributions on $\nu_{1}$ with increasing numbers of observations ..... 211
6.13 Representations of the mean functions ..... 211
6.14 Mesh deformation for the forward model of the mean functions. Original data included for comparison in the left hand frame ..... 212
6.15 Face shaped data ..... 213
6.16 Marginal distributions on $p_{0}(0)$ with increasing numbers of observations ..... 213
6.17 Marginal distributions on $p_{1}(0)$ with increasing numbers of observations ..... 214
6.18 Marginal distributions on $p_{0}(0)$ with increasing numbers of observations ..... 214
6.19 Marginal distributions on $p_{1}(0)$ with increasing numbers of observations ..... 215
6.20 Representations of the mean functions ..... 216
6.21 Mesh deformation for the forward model of the mean functions. Original data included for comparison in the left hand frame ..... 216

## Acknowledgments

I would first of all like to thank Andrew Stuart for all of his hard work and collaboration over the last $41 / 2$ years. He has been a constant source of encouragement, knowledge and inspiration, and I have been hugely fortunate to have him as my supervisor.

Thanks also go to Zoe Langham for her constant support and understanding, and for putting up with me.

Thanks to Colin Cotter, whose interest in my work led to a very interesting collaboration that is still ongoing.

Also to all who have graced room B2.39 over more than 3 years who have helped me with all sorts of techie problems, thanks!

Thanks go to the CSC for computing time and technical support on the IBM cluster Francesca.

Thanks go to EPSRC for funding my studies.
Finally thanks to my family and particularly my Dad.

For Mum

## Declarations

I hereby declare that the work contained herein is, unless expressly acknowledged, my own work. Chapters 2 through 5 were formulated in collaboration with Andrew Stuart. Chapter 6 was formulated in collaboration with Andrew Stuart and with Colin Cotter, who also provided code for the forward model. All other code, and all other figures were created solely by the author.

## Abstract

In the course of this thesis, several different applications of data assimilation will be looked at. In each case, a rigorous mathematical framework will be constructed, in a Bayesian context, to enable the use of various types of data to infer on various infinite dimensional parameters of the system that has been observed. After careful consideration of the forward problem, well-defined posterior distributions on function space are constructed. Using MCMC methods which are defined on these function spaces themselves, we can construct Markov chains whose invariant measures are the posterior of interest. From this point, we can implement these methods on a computer, having finally discretised the problem.

The philosophy that we adhere to throughout, is the idea that numerical methods formulated on function space are robust under discretisation, and do not suffer from the curse of dimensionality typically suffered by sampling methods formulated after disretisation.

The first few chapters (after the introductory chapter) will focus on various aspects of data assimilation of observations of Stokes flow dynamics. Chapter 2 will focus on Eulerian data where direct observations of the velocity of the fluid at various points in time and space will be made. Chapter 3 will concentrate on data assimilation of indirect observations of the field, in the form of the positions of passive tracers in the flow. In these two chapters we will assume that the forcing of the system is known and that we are merely trying to recover the initial condition of the flow field.

In chapter 4 we will consider both Eulerian and Lagrangian data assimilation, with the added complexity of trying to use the data to not only infer on the initial condition but also on the space-time dependant forcing of the system.

In chapter 5 we will try to show how these smoothing methods could be adapted into a filtering algorithm, and a simple example will be presented.

In the final chapter, 6, this Bayesian framework on function space will be applied to a shape matching problem with applications in the biomedical sciences.

## Chapter 1

## Introduction and Mathematical

## Preliminaries

### 1.1 Introduction

Applied mathematicians are constantly striving to better approximate phenomena and systems that describe the world around us. Whether it be the currents of an ocean that distribute warmer water and marine life around the globe, or the dynamics of a mixture of reacting chemicals, mathematics has a diverse range of tools available for trying to understand these systems and for attempting to reconstruct this behaviour on a computer

In many applications the systems of interest are infinite dimensional in nature. From potential functions derived from our knowledge of nuclear/atomic forces, to flow fields in the air around us, the functions that describe these systems can never be stored completely on a computer. These functions belong to infinite dimensional function spaces. A computer, however, has a finite amount of memory and computing power,
and so these functions can only be approximated finite dimensionally, usually amounting to discretizing the system onto grids in time and space, or alternatively truncating infinite sums that converge to the functions of interest.

In this thesis, we aim to demonstrate how formulating algorithms on function spaces, as opposed to already discretized spaces, can lead to better, more efficient and robust methods. We will apply this philosophy to a range of subjects in data assimilation.

### 1.1.1 Data Assimilation

Data assimilation is the act of blending a mathematical model of a system with informative observations, to give a better understanding of that system's state either during the window of observations, or even to create forecasts of what is going to happen in the future. The various methods associated with this area have many important applications, and as such there is a huge amount of literature on the subject from meteorologists, oceanographers and geophysicists to name a few.

We will focus, in chapters 2-4, on data assimilation problems in fluid dynamics. Using satellite data of the atmosphere, or GPS tracker data from ocean floats, we aim to reconstruct the underlying velocity field that was present at the time of observation. We will however use simulated data in our algorithms in a bid to assess the algorithm's ability to find the information we wish to ascertain in simple systems, so that these methods could possibly be applied to real data sets.

There is a great deal of literature on this subject, particularly from the perspective of data assimilation in weather forecasting. Since in this field the aim is to provide a forecast of what is to come, speed and efficiency in algorithms is particularly important. Data arrives on a regular basis, and using the prior knowledge of the system since the last forecast, this data is assimilated and a new forecast produced, before the whole
procedure starts again straight away. This form of data assimilation, termed filtering, dominates the meteorological community $[48,54,3,52,31]$. However, smoothing algorithms, where a whole data set is used to infer on the initial condition of the velocity field at the beginning of the assimilation window, are also commonly used offline to analyse discrepancies between the model and the dynamical system, or model error, or to estimate parameters. This is sometimes termed reanalysis or hindcasting. In [6] the basic concepts in variational data assimilation are introduced. Importantly though, this approach is then extended to the statistical perspective that we will be interested in for the purposes of this thesis, pointing out that the cost functionals can actually be considered to be the log probability densities of a probability measure on velocity fields. The paper [30] considers data assimilation in the case where data is sparse, with a view to oil recovery problems in geophysics, similarly looks to a Bayesian framework and it's relationship to Tikhonov regularisation techniques in variational frameworks.

There are many other applications of data assimilation across applications in science and engineering. Not least in the field of oil recovery[44], where sparse measurements taken from oil wells during excavation of oil are used to ascertain various properties of the porosity of the rock, for example, to give indications of where might be best to drill to maximise the output of oil from the location[59, 58, 38]. This problem is often tackled by using ensemble Kalman filters. Other sources of information, such as 4 D seismic data can also be used to better quantify the properties of the rock and the oil below[28].

The paper [34] highlights the practical demands on medium range forecasts which dictate the choice of data assimilation algorithms. In the vast majority of practical applications, variational methods are used where the problem is framed in such a way that the solution is given to be the flow field which minimises a particular functional which
includes a term to ensure that the observations are assimilated, and a regularisation term to ensure minimal amounts of regularity of that field. In [34] the 4DVAR algorithm (an example of a smoothing data assimilation algorithm) is utilised. They go on to explain the demands of medium range forecasting, where the state space is approximated by $7 \times 10^{9}$ degrees of freedom, and where data is added every 6-12 hours. In those 12 hour cycles, only two hours is given to data assimilation, so the methods need to be quick and accurate, not least because the calculations of the forward model in the algorithm are very expensive due to the huge number of degrees of freedom. In [33] the 4DVAR algorithm is again used, this time for the Lagrangian data assimilation of chemical species in the atmosphere.

With the large dimension of the state space in some applications in mind, many methods that are proposed aim to approximate the solutions as best as possible with as little computational effort as possible. For example, in [49], two different incremental 4DVAR methods are introduced. One uses the tangential linear model (TLM) where the forward model is approximated by a linearisation, and the other uses a suitable inexact linear model which if carefully chosen, can yield similar results at much less cost. It is in fact shown that the error incurred when using the inexact model can be less than if you simply reduced the resolution of the TLM to give the same level of computational cost.

In these variational methods, where we wish to find the minimisation of a functional on the velocity field's state space, gradient descent-type methods are often utilised to find the solution. In [90], examples are given where the explicit gradient of the observation operator can't be found (for example Burgers equation, or in satellite radiance data assimilation). It is shown, however, how generalisations of gradient-based methods can be derived for non-differentiable observation operators.

Variational methods are not the only area of discussion with regards to data assimilation however, and in [4, 12] a Bayesian framework is explored, along with various Monte Carlo Markov chain (MCMC) methods which allow us to sample from complex probability distributions without knowing their explicit densities. This will be the approach that is adopted in this thesis. The important thing to notice here is that the MCMC methods discussed here are framed on function space, meaning that they are robust under all discretisations. The advantages of this are highlighted in [11] where various MCMC methods are bench marked against each other for sampling diffusion bridges, which are themselves in function space. The MCMC method framed on function space is compared to another which is not, for a range of different mesh sizes. It is shown conclusively that as the mesh is refined, that the efficiency of the standard MCMC method on discretized space decreases as the mesh is refined, whereas the MCMC method posed on function space is unaffected. This is a key discovery behind the philosophy of building methods on function space on which we base this entire thesis.

### 1.1.2 Filtering

Variational smoothing techniques such as 4DVAR are important in long range forecasting, and for reanalysis of data in an attempt to better understand error in the model. However for most short range forecasting applications, filtering methods are the order of the day, as they can be continually updated with short bursts of data to reassess the weather conditions and improve forecasts. 3DVAR, a popular choice of variational filtering algorithm, is mentioned in brief in [6]. Another choice of algorithm called the maximum likelihood ensemble filter (MLEF) equations are also used in [90]. In actual fact there are a myriad of options available to anyone who wishes to conduct filtering on a given constantly concatenated data set. [54] gives a comprehensive review of recent
methods for a range of different scenarios; sparse/plentiful observations, addition of model error terms, linear and nonlinear models/dynamics to name but a few. Together with numerics for a selection of these schemes, a broad overview of the various methods' advantages and disadvantages are presented.

A great deal of filtering algorithms are based, in some part, on the Kalman filter[3, 31]. This method[46], for filtering data from linear models, requires a prior mean and error covariance matrix in each assimilation "window" (often calculated as the forecast from the previous window). Then by assimilating the data from that window, an analysis state and error covariance are calculated, via the Kalman gain matrix, which essentially applies Bayes' theorem. This analysis is then pushed forward in time through the dynamics to the end of the window to create the forecast state and error covariance, which is then in turn used as the prior for the next window. A nonlinear equivalent, the extended Kalman filter, can be used, where the nonlinear model is approximated in each window by a linearisation. This can of course cause problems if the window length is too long, or if the system is highly nonlinear.

However, these methods can still be too computationally expensive to calculate on-line as the data comes in, and further approximations must be made. One commonly used method in this regard is the ensemble Kalman filter (EnKF), where the forward model on a distribution is approximated by calculating the paths of an ensemble of particles in the system, from which a mean and covariance are approximated. A subclass of such a method, termed square root filter (SRF) are analysed in more detail in [52]. In this method, the EnKF approach is used to approximate the mean and the square root of the analysis covariance, to avoid negative diagonal entries and to ensure that the matrix is positive definite. The subsequently calculated forecast covariance is then post-multiplied by an orthogonal matrix. [52] gives conditions on this matrix to ensure
that the analysis ensemble mean will be equal to the analysis state estimate.
Other methods, such as particle filter methods, are also a subject of great interest. In these methods, an ensemble of particles is moved forwards in the dynamics, and then, based on a window of observations, those particles are given importance weights. Those weights are equivalent to an approximation of the posterior distribution at that point in time. The process is then iterated. One of the main issues with many types of particle filters, is a phenomena called sample impoverishment, where the weights on the particles will often be centred on only a small handful of particles, therefore rendering our approximation of the distribution to be very coarse. This can be addressed by periodically resampling the particles according to the current distribution approximation so that the weights are more evenly distributed once again. [48] presents one particular particle filter method, the Diffusion Kernel Filter, and compares it against another, the Bootstrap Filter, on a highly nonlinear system, the Lorenz-63 equations. There are many different versions of particle filters, and many different ways of addressing problems such as sample impoverishment. An important point of note with particle filters is that in theory, as the number of particles included in the algorithm increases, that the approximated distribution converges to the target distribution. This is in contrast to EnKF, for which no such result exists, where Gaussian approximations are made. Particle filters are a way of capturing non-linear and non-Gaussian behaviour in a distribution. They do, however, suffer from the curse of dimensionality, as shown in $[9,13,72]$.

### 1.1.3 Generalised Polynomial Chaos

An alternative approach to MCMC methods, which can be extremely expensive, can be to approximate the forward model, with random inputs, by generalised polynomial chaos (gPC)[88]. By discretizing the random space and calculating the forward model at each
of the grid points in that discretization, the forward model can then be approximated by a polynomial in the random inputs, in a process called stochastic collocation[56]. This has the advantage that once the approximating polynomial has been found, evaluating that approximation is almost trivial. Once we have this approximation, it can then be fed into any of the other methods we have been discussing, for instance the ensemble Kalman filter[51]. These types of methods have also been considered in the field of data assimilation of incompressible flows[89]. One drawback to these methods is the curse of dimensionality. For each additional dimension in the state space, many more forward models must be run to incorporate that degree of freedom into the polynomial expansion. At the time of writing of [88], only forward models with 50 degrees of freedom had been attempted, and although this figure will undoubtedly be larger now due to the continuing improvement in computing power, this method still remains completely unfeasible for the size of state spaces that the meteorological community deals with, for example.

### 1.1.4 Lagrangian Data Assimilation

In chapters 3 and 4 we will also be considering data assimilation of Lagrangian data. Lagrangian data, where the position of a tracer in a flow is observed noisily, has the added complication that even if the underlying dynamics of the flow is linear or close to linear, the dynamics of the tracer itself are highly nonlinear. In paper [43] a choice of extended Kalman filter is suggested, using a tangent linear model, where the positions of the Lagrangian tracers are included in the state vector, along with the vector field. There are many references in the literature to attempts to assimilate Lagrangian data[4, $12,5,6,68,69,33,18]$. The paper [68] details a particle filter methodology for tackling Lagrangian data assimilation, which is then implemented in numerics that appear in [69]. It also highlights the problems with this kind of data, where hyperbolic points in
the flow can lead to bimodal distributions of the tracers ("filter divergence"), which are not at all compatible with Gaussian approximations, which many methods rely on.

Much of the literature on Lagrangian data assimilation is concerned with various problems in oceanography. However this is not the only application where Lagrangian data is collected, as shown in $[14,15,87]$, where atmospheric chemical-transport models are used to approximate the movement of various pollutants in the air. This can be with a view to trying to find the source of a particular pollutant, or to recover the concentration field of that pollutant over a wide area, or even to create forecasts of levels of pollutant levels.

The paper [5] presents a Bayesian framework with which to work with Lagrangian data. It then discusses and implements different finite-dimensional MCMC methods with which to sample from well defined posterior measures, and then compares these results with approximations given by an implementation of EnKF, which is shown to break down when there are long gaps between observations. This is closer to the approach that we will take in 3, although we will be formulating our methods on function spaces.

### 1.1.5 Shape Registration

There are also plenty of other applications that data assimilation can be applied to, not least in the biomedical sciences. In chapter 6, we apply the general framework used in the previous chapters, on a shape matching problem with a view to applications in biomedical imaging from prenatal scans. The problem of finding distances between shapes in shape space is also an area with a great deal of literature, although not so much from a statistical Bayesian point of view. Other methods which find geodesic paths in shape space between two images often take on a variational format, with various choices of functional, with some based on elastic energies in deformations [29, 67, 86, 85, 2].

By adding various terms to the functional with differing properties, and giving different weightings to those terms, different geodesic paths can be calculated. Alternatively spline-based techniques can be used[62] to find these paths.

The approach that we will take in this thesis is more similar to that found in $[78,8,36]$, where the deformation of one shape to another is defined by a velocity field which flows shape $A$ at time $t=0$, to another shape $B$ at time $t=1$. The distance along any given path is then defined by the integral of the norm of that flow field over time in some Hilbert or Sobolev space. However, we will also be putting this into a Bayesian statistical setting, an approach which is not mentioned in any of the above references.

### 1.2 Summary of the Thesis

After introducing some basic concepts, theory and techniques in chapter 1, in chapter 2 we will consider the problem of data assimilation in the context of Stokes flow. We will frame this in a Bayesian context, and show that, by analysing the forward problem, enabling us to make the correct choice of prior, the posterior is absolutely continuous with respect to that prior. Using the Radon-Nikodym derivative that this absolute continuity induces, we shall show how Markov chains can be constructed on function space whose invariant measure is the posterior measure of interest. We will then implement such a method. This chapter is very much a way of benchmarking our methods, since it is in fact simple, with the observation operator being linear and the prior a Gaussian, to calculate the posterior distribution, which is also Gaussian, explicitly. We will calculate these analytic posterior distributions for some simple low-dimensional examples, and look at the effect of varying a range of parameters.

Knowing this explicit distribution is useful as we are then able to look at the
convergence of our Markov chain Monte Carlo(MCMC) methods, and then these methods can then be applied to nonlinear problems where the posterior is not Gaussian, or other problems, such as Eulerian data assimilation with model error, where the posterior distribution is not so easily analytically calculated.

Importantly, we will compare the efficiency of the standard MCMC methods, as opposed to those framed on function space. This is central to the philosophy that we adhere to throughout this entire thesis.

We will also consider some issues to do with mismatch in the system that we are observing and the model on our computer, before going on to some numerics where the amount of informative observations increases. Our hope being here that as our information increases, the posterior distribution converges to an increasingly peaked (or in the limit a Dirac delta measure) on the system state which was actually present when the observations were made

In chapter 3, an identical approach is taken for the problem of Lagrangian data assimilation in Stokes flow, where passive (inertia-less) particles are advected by a fluid current, and noisy observations are made of those particles' positions. This problem is highly nonlinear, despite the linearity of the underlying fluid motion equations. This gives us an opportunity to show the algorithms working in a scenario where the posterior cannot be calculated explicitly, due to the non linearity of the particle's dynamics. Equivalent numerical results as those presented in chapter 2 (bar the convergence of the posterior to the analytical posterior) are presented in this case also.

In chapter 4, we consider both Eulerian and Lagrangian data types, but will add some uncertainty into the model itself. We will use the data to infer not only on the initial condition of the assimilation window, but also on the space-time dependant forcing, which we now assume that we don't know. Once again we will analyse the
forward model to inform the correct choice of prior on the initial condition and on the forcing, and then invoke the Bayesian framework to enable numerical sampling from the well-defined posterior distribution. After presenting some numerics, we will then analyse what we are able to garner about the forcing inherent in the system through these two different data types. As we did in chapters 2 and 3, we will once again look at what happens to the posterior distribution as we increase the number of informative observations to be assimilated.

In chapter 5, a short example will be made of how these methods could be applied to filtering. Instead of using a gain matrix or a particle filter, the analysis for each assimilation window will be computed via the statistical methods we have used throughout chapters 2-4. This method would certainly be too computationally expensive to be practically useful at this point (in meteorological applications for instance), and the problem of moving the posterior distribution through time is still a point of concern. In this simple example, where Eulerian data assimilation without model error is considered, various Gaussian approximations of the analysis will be made, which can then be propagated through time exactly in the linear dynamics. Since this case could be analytically solved exactly, we also present a small set of results where a deterministic filter is implemented.

In chapter 6 we move away from the application of data assimilation in fluid mechanics, and turn our attentions to a shape matching problem with applications in the biomedical sciences. We once again frame this problem as a Bayesian inverse problem, and in the mode of the previous chapters, analyse the forward problem which maps one member of shape space continuously to another, in a way which minimises the energy of the deformation in a particular Hilbert space norm. Using this analysis we are able to pick a prior for which the posterior will be absolutely continuous. We
are then able to implement our methods to draw samples from these distributions, and we present results with increasing numbers of observations with the data created from a deformation from the prior distribution, and then we consider some cases with shapes of "real" objects.

In the remaining sections in this chapter we introduce mathematical concepts and algorithms which will be used frequently through the whole thesis.

### 1.3 Function Space Settings

Firstly, we will introduce some basic mathematical concepts that we will use throughout. The Sobolev space $W^{k, p}(U)$ is often used as the function space setting when dealing with PDEs. First we define weak derivatives.

Definition 1.3.1. [50] Suppose $u, v \in L_{\mathrm{loc}}^{1}(U)$, and $\alpha$ is a multiindex. We say that $v$ is the $\alpha^{\text {th }}$-weak partial derivative of $u$, written

$$
D^{\alpha} u=v
$$

provided

$$
\int_{U} u D^{\alpha} \phi d x=(-1)^{|\alpha|} \int_{U} v \phi d x
$$

for all test functions $\phi \in C_{c}^{\infty}(U)$. Note that here, $|\alpha|$ is not defined in terms of the Euclidean norm, but in fact as follows: $|\alpha|:=\sum_{i=1}^{d} \alpha_{i}$.

Definition 1.3.2. [50] $k, p, d \in \mathbb{N}$. The Sobolev space

$$
W^{k, p}(\Omega)
$$

consists of all locally summable functions $u: \Omega \rightarrow \mathbb{R}^{d}$ such that for each multiindex $\alpha \in \mathbb{N}^{d}$ with $|\alpha| \leq k, D^{\alpha} u$ exists in the weak sense and belongs to $L^{p}(\Omega)$. In the case that $p=2$, this space is a Hilbert space.

### 1.3.1 Function Space Setting For The Stokes Problem

Later in the thesis we will be considering solutions to the Stokes problem on the twodimensional torus. Solutions to this problem represent incompressible fluid motion, and at the very least the velocity field is square integrable. We introduce for this problem the space $H=\left\{u \in L^{2}(\Omega): \nabla . u=0, \int_{\Omega} u d x=0\right\}$, armed with the usual $L^{2}$ norm, $\|\cdot\|_{2}^{2}=\sum_{k}\left\langle\cdot, \phi_{k}\right\rangle^{2}$, where $\phi_{k}$ is an orthonormal basis of the space $H$. More specifically, $\left\{\psi_{k}\right\}$ is an orthonormal basis of $L^{2}\left(\mathbb{T}^{2}\right)$, where

$$
\begin{equation*}
\psi_{k}(x)=\exp (2 \pi i k \cdot x) \tag{1.1}
\end{equation*}
$$

For a function $u=\sum_{k} u_{k} \psi_{k} \in L^{2}$ to also be divergence free we require the condition that $u_{k} . k=0$. This motivates the choice of $\phi_{k}$, so that we set

$$
\begin{equation*}
\phi_{k}=\frac{k^{\perp}}{|k|} \psi_{k} \tag{1.2}
\end{equation*}
$$

This defines an orthonormal basis of $H$.
Similarly, we define $H^{s}=\left\{u \in W^{s, 2}(\Omega): \nabla . u=0 \int_{\Omega} u=0\right\}$, with norm $\|\cdot\|_{s}^{2}=\sum_{k}|k|^{2 s}\left\langle\cdot, \phi_{k}\right\rangle^{2}$. If we define the Leray projector $\mathcal{P}$ to be the projection of functions onto $H$, ie such that $\mathcal{P} f=\sum_{k}\left\langle f, \phi_{k}\right\rangle \phi_{k}$, then we can alternatively define $H^{s}$ (where $H=H^{0}$ ) to be

$$
H^{s}=\mathcal{P} W^{s, 2}(\Omega)
$$

In terms of notation, when we are dealing with a function $u \in H$, we shall assume that

$$
u=\sum_{k} u_{k} \phi_{k}
$$

where $u_{k}=\left\langle u, \phi_{k}\right\rangle$.

### 1.4 Solutions to the Stokes Problem

Stokes flow is traditionally framed as a PDE in the following way[74];

$$
\begin{array}{r}
\partial_{t} u-\nu \Delta u+\nabla p=f, \quad t \geq 0 \\
\nabla \cdot u=0, \quad t \geq 0 \\
u(x, 0)=u_{0}(x), \quad x \in \Omega \tag{1.5}
\end{array}
$$

We consider solutions of (1.3-1.5), with $\Omega=\mathbb{T}^{2}$ a unit square in two dimensions with periodic boundary conditions. We expand $u, p$ and $f$ in terms of the $\psi_{k}$ as defined in 1.1 to give

$$
\begin{aligned}
u(x, t) & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} u_{k}(t) \psi_{k}(x), u_{k}(t) \in \mathbb{C}^{2}, \\
p(x, t) & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} p_{k}(t) \psi_{k}(x), p_{k}(t) \in \mathbb{C}, \\
f(x, t) & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} f_{k}(t) \psi_{k}(x), f_{k}(t) \in \mathbb{C}^{2} .
\end{aligned}
$$

(1.3) applied to the $u_{k}$ gives

$$
\begin{equation*}
\frac{d u_{k}}{d t}+\nu 4 \pi^{2}|k|^{2} u_{k}=-2 \pi i k p_{k}+f_{k} \tag{1.6}
\end{equation*}
$$

Taking the dot product of (1.6) with $u_{k}$, while using the fact that $u_{k} \cdot k=0$ (in order to satisfy 1.4 , and rearranging for $p_{k}$ gives

$$
p_{k}=\frac{f_{k} \cdot k}{2 \pi i|k|^{2}}
$$

We can now substitute this back into our equation (1.6) to give

$$
\frac{d u_{k}}{d t}+\nu 4 \pi^{2}|k|^{2} u_{k}=\left(I-\frac{k \otimes k}{|k|^{2}}\right) f_{k}
$$

Here $a \otimes b$ denotes the outer product, defined by $(a \otimes b) c:=a(b . c)$.

Solving this ODE gives us a set of equations for the evolution of the $u_{k}$, which in the case that the $f_{k}$ are constant gives

$$
\begin{equation*}
u_{k}(t)=\frac{1}{4 \pi^{2} \nu|k|^{2}}\left(I-\frac{k \otimes k}{|k|^{2}}\right) f_{k}+\exp \left(-\nu 4 \pi^{2}|k|^{2} t\right) u_{k}(0) \tag{1.7}
\end{equation*}
$$

So given a suitable initial condition $u_{0}$ and forcing term $f$, we have an analytical solution to the Stokes equations on our periodic domain $\Omega=\mathbb{T}^{2}$.

Alternatively, we can consider a Helmholtz decomposition $[35,32]$ of $L^{2}\left(\mathbb{T}^{2}\right)$. That is, splitting our basis $\mathcal{B}=\left\{\binom{1}{0} e^{2 \pi i k \cdot x},\binom{0}{1} e^{2 \pi i k \cdot x}: k \in \mathbb{Z}^{2}\right\}$ into two sets, one being divergence-free and the other being curl-free. Since

$$
\nabla \cdot u=0 \Longleftrightarrow u_{k} \cdot k=0
$$

we can redefine our basis $\psi_{k}=e^{2 \pi i k . x}$ into two bases that form orthogonal subspaces:

$$
\begin{aligned}
\mathcal{E} & =\operatorname{Span}\left\{e_{k}=\frac{k^{\perp}}{|k|} e^{2 \pi i k \cdot x}: k \in \mathbb{Z}^{2} \backslash\{0\}\right\}, \\
\mathcal{D} & =\operatorname{Span}\left\{d_{k}=\frac{k}{|k|} e^{2 \pi i k \cdot x}: k \in \mathbb{Z}^{2} \backslash\{0\}\right\} .
\end{aligned}
$$

$\mathcal{E}$ is divergence-free, and $\mathcal{D}$ is curl-free. We will now consider a normal expansion of $u$ in terms of the $L^{2}$ Fourier basis $\mathcal{B}$,

$$
\begin{aligned}
u & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} u_{k} e^{2 \pi i k \cdot x} \\
& =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} \omega_{k} e_{k}+\tau_{k} d_{k},
\end{aligned}
$$

for $\omega_{k}, \tau_{k} \in \mathbb{C}$. So therefore $u_{k}=\frac{\omega_{k} k^{\perp}+\tau_{k} k}{|k|}$.
If we similarly expand $p$ and $f$ such that

$$
\begin{aligned}
p & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} p_{k} e_{k}+p_{k}^{\perp} d_{k}, p_{k}, p_{k}^{\perp} \in \mathbb{C} \\
f & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} f_{k} e_{k}+f_{k}^{\perp} d_{k}, f_{k}, f_{k}^{\perp} \in \mathbb{C},
\end{aligned}
$$

and substitute these expansions into (1.3), then we get

$$
\begin{aligned}
\sum_{k} \frac{d \omega_{k}}{d t} e_{k}+\frac{d \tau_{k}}{d t} d_{k}= & \nu \sum_{k}-4 \pi^{2}|k|^{2} \omega_{k} e_{k}-4 \pi^{2}|k|^{2} \tau_{k} d_{k}-\sum_{k} 2 \pi i|k| p_{k}^{\perp} d_{k} \\
& +\sum_{k} f_{k} e_{k}+f_{k}^{\perp} d_{k}
\end{aligned}
$$

Notice that the $p_{k}$ terms disappear as they are orthogonal to $k$. Once again taking the dot product of this with $k$, we are left with an equation for $p_{k}^{\perp}$.

So

$$
p_{k}^{\perp}=\frac{f_{k}^{\perp}}{2 \pi i|k|} .
$$

Effectively this means that the pressure has removed the part of $f$ which is not divergence free.

So we end up with two sets of ordinary differential equations,

$$
\begin{aligned}
& \frac{d \omega_{k}}{d t}=-4 \nu \pi^{2}|k|^{2} \omega_{k}+f_{k} \\
& \frac{d \tau_{k}}{d t}=0
\end{aligned}
$$

Since we set $u_{0}$ to be divergence free, then $\tau_{k}=0 \forall t \geq 0$. As for the $\omega_{k}$, we get

$$
\begin{aligned}
\omega_{k} & =\frac{1}{4 \nu \pi^{2}|k|^{2}} f_{k}+\left(\omega_{k}(0)-\frac{f_{k}}{4 \nu \pi^{2}|k|^{2}}\right) \exp \left(-4 \nu \pi^{2}|k|^{2} t\right) \\
& =\frac{1-\exp \left(-4 \nu \pi^{2}|k|^{2}\right)}{4 \nu \pi^{2}|k|^{2}} f_{k}-\omega_{k}(0) \exp \left(-4 \nu \pi^{2}|k|^{2} t\right)
\end{aligned}
$$

Note that this holds only if $f$ is constant in time. This is equivalent to finding weak solutions $u \in H$ which satisfy the ODE on $H$ given by;

$$
\begin{align*}
\frac{d u}{d t}+A u & =\mathcal{P} f=g, \quad t>0  \tag{1.8}\\
u(0) & =u_{0} \in H \tag{1.9}
\end{align*}
$$

In this setting, it is much simpler to solve these equations. The solution, for $u(0)=$ $\sum_{k} u_{k} \phi_{k}$ and $g=\sum_{k} g_{k} \phi_{k}$, is given by

$$
\begin{aligned}
u(t) & =\exp (-\nu A t) u(0)+\int_{0}^{t} \exp (-A(t-s)) g(s) d s \\
& =\sum_{k} \exp \left(-\nu a_{k} \omega_{k} t\right) u_{k}+\int_{0}^{t} \exp \left(-a_{k} \omega_{k}(t-s)\right) g_{k}(s) d s
\end{aligned}
$$

Note here that no assumption that $f$ is constant is made. Sums such as these can be quickly computed numerically using a Fast Fourier Transform (FFT)[60], which are implemented in packages such as FFTW (Fastest Fourier Transform in the West) [70].

### 1.5 The Stokes Operator

Throughout chapters 2-5 we will be performing various types of data assimilation on data observed from a dynamical system termed Stokes flow. Moreover, we will be sampling Gaussian random fields where the covariance operator is given as a power of the Stokes operator. While not going into details of this in this section, these are motivations for introducing here the Stokes operators and some properties which it possesses.

We define the Stokes operator $A: H \rightarrow H$ as follows. In [66] it is shown that the Stokes operator, acting on $H$, is simply the Laplacian projected onto $H$. Therefore, we can define the Stokes operator using the Leray projector $\mathcal{P}$;

$$
A=-\mathcal{P} \Delta
$$

We now look to find the eigenvalues and eigenfunctions of $A$. If we define the domain of our functions to be the unit torus $\mathbb{T}^{2}=[0,1] \times[0,1]$ with periodic boundary conditions, as we will be doing throughout chapters $2-5$, then this is a relatively simple task. Let $k \in \mathbb{Z}^{2} \backslash\{0\}$. If $k=\binom{k_{1}}{k_{2}}$, define $k^{\perp}:=\binom{k_{2}}{-k_{1}}$, which is perpendicular to $k$ and is the same size in the Euclidean norm.

We now define our eigenfunctions;

$$
\begin{equation*}
\phi_{k}(x)=\frac{e^{2 \pi i k . x}}{|k|} k^{\perp} . \tag{1.10}
\end{equation*}
$$

These satisfy the equality $\phi_{k} \cdot k=0$ which we require for them to be divergence free. The corresponding eigenvalues, with respect to the Stokes operator, are

$$
\begin{equation*}
a_{k}=4 \nu \pi^{2}|k|^{2} . \tag{1.11}
\end{equation*}
$$

These eigenfunctions form an orthonormal basis of $H$, and will be used throughout chapters 2-5.

### 1.6 Estimate on solutions of Stokes Flow

In analysing the forward problem, we will need to use various estimates on the solution to the Stokes problem. In this section we present two lemmas that we will need in doing this.

Lemma 1.6.1. Let $u \in H^{r}$ for some $r \geq 0$. Then for any $s \geq r$

$$
\left\|e^{-\mathcal{A} t} u\right\|_{s} \leq C t^{-(s-r) / 2}\|u\|_{r}, \quad \forall t>0 .
$$

Proof. Using the fact that $e^{-x} x^{\alpha}$ is bounded on $\mathbb{R}^{+}$for any $\alpha>0$ we have that

$$
\begin{aligned}
\left\|e^{-\mathcal{A} t} u\right\|_{s}^{2} & =\sum_{k} e^{-c|k|^{2} t}\left|u_{k}\right|^{2}|k|^{2 s} \\
& =\frac{1}{(c t)^{s-r}} \sum_{k} e^{-c|k|^{2} t}|k|^{2(s-r)}(c t)^{s-r}\left|u_{k}\right|^{2}|k|^{2 r} \\
& \leq \frac{C}{t^{s-r}} \sum_{k}\left|u_{k}\right|^{2}|k|^{2 r} \\
& \leq \frac{C}{t^{s-r}}\|u\|_{r}^{2}
\end{aligned}
$$

Taking the square root gives the result.

Lemma 1.6.2. Let $u_{0} \in H^{l}$ for $l \geq 0$ and $f \in C\left(0, T ; H^{r}\right)$ for some $r \geq 0$. Then for $l \leq s<r+2$, then

$$
\|u(t)\|_{s} \leq C\left(\frac{\left\|u_{0}\right\|_{l}}{t^{(s-l) / 2}}+\|f\|_{C\left(0, T ; H^{r}\right)}\right)
$$

Proof. The solution map of Stokes flow is given by the variation of constants formula[39]:

$$
u(t)=\exp (-A t) u(0)+\int_{0}^{t} \exp (-A(t-s)) f(s) d s
$$

Using lemma 1.6.1,

$$
\begin{aligned}
\|u(t)\|_{s} & \leq\left\|e^{-\mathcal{A t}} u_{0}\right\|_{s}+\left\|\int_{0}^{\tau} e^{-\mathcal{A}(t-\tau)} f(\tau) d \tau\right\|_{s} \\
& \leq \frac{C}{t^{(s-l) / 2}}\left\|u_{0}\right\|_{l}+\mathcal{I}
\end{aligned}
$$

where

$$
\begin{aligned}
\mathcal{I} & \leq \int_{0}^{t}\left\|e^{-\mathcal{A}(t-\tau)}\right\|_{\mathcal{L}\left(H^{r}, H^{s}\right)}\|f(\tau)\|_{r} d \tau \\
& \leq\|f\|_{C\left(0, T ; H^{r}\right)} \int_{0}^{t}\left\|e^{-\mathcal{A}(t-\tau)}\right\|_{\mathcal{L}\left(H^{r}, H^{s}\right)} d \tau .
\end{aligned}
$$

Using the definition of the operator norm and lemma 1.6.1 once again we get that

$$
\begin{aligned}
\left\|e^{-\mathcal{A} t}\right\|_{\mathcal{L}\left(H^{r}, H^{s}\right)}^{2} & =\sup _{u} \frac{\left\|e^{-\mathcal{A} t} u\right\|_{s}^{2}}{\|u\|_{r}^{2}} \\
& \leq \frac{C}{t^{s-r}} \sup _{u} \frac{\|u\|_{r}^{2}}{\|u\|_{r}^{2}} .
\end{aligned}
$$

Therefore

$$
\left\|e^{-\mathcal{A} t}\right\|_{\mathcal{L}\left(H^{r}, H^{s}\right)}^{2} \leq \frac{C}{t^{(s-r) / 2}}
$$

which is integrable if and only if $s-r<2$. Putting this all together we have that

$$
\|u(t)\|_{s} \leq C\left(\frac{1}{t^{(s-l) / 2}}\left\|u_{0}\right\|_{l}+\|f\|_{C\left(0, T ; H^{r}\right)}\right) .
$$

We will also need Grönwall's inequality[20]. Here we present the integral form of the lemma, taken from [84].

Lemma 1.6.3. Let $\alpha, \beta$ and $w$ be real-valued functions defined on the interval $[0, T]$. Assume that $\beta$ and $w$ are continuous and that the negative part of $\alpha$ is integrable on every closed and bounded sub interval of $[0, T]$. If

$$
w(t) \leq \alpha(t)+\int_{0}^{t} \beta(s) w(s) d s, \quad \forall t \in[0, T],
$$

then

$$
w(t) \leq \alpha(t)+\int_{0}^{t} \alpha(s) \beta(s) \exp \left(\int_{s}^{t} \beta(r) d r\right) d s, \quad \forall t \in[0, T] .
$$

If in addition $\alpha$ is a constant in time then

$$
w(t) \leq \alpha \exp \left(\int_{0}^{t} \beta(s) d s\right) \quad \forall t \in[0, T] .
$$

### 1.7 Random Field Sampling

In the chapters that follow, we will formulate several MCMC algorithms on function space (more specifically vector fields). In these algorithms, we will be required to make proposals of vector fields, which are a linear combination of the current accepted state, and a draw from the prior distribution, a measure on a space of random fields[80]. These prior distributions will be Gaussian distributions on function space. In this section we will discuss how we can go about drawing samples from such a distribution. The following definition is taken from [73].

Definition 1.7.1. A Gaussian random field $u: H \times \Omega_{0} \rightarrow \mathbb{R}^{n}$ is one where, for any integer $q \geq 1$, and any set of points $\left\{x_{k}\right\}_{k=1}^{q}$ in $H$, the random vector

$$
\left(u\left(x_{1} ; \cdot\right)^{*}, \ldots, u\left(x_{q} ; \cdot\right)^{*}\right)^{*} \in \mathbb{R}^{n q}
$$

is a Gaussian random vector. The mean function of a Gaussian random field is $m(x)=$ $\mathbb{E} u(x)$. The covariance function is $c(x, y)=\mathbb{E}(u(x) m(x))(u(y) m(y))^{*}$. For Gaussian random fields this function, together with the mean function, completely specify the joint probability distribution for $\left(u\left(x_{1} ; \cdot\right)^{*}, \ldots, u\left(x_{q} ; \cdot\right)^{*}\right)^{*} \in \mathbb{R}^{n q}$. Furthermore, if we view the Gaussian random field as a Gaussian measure on $L^{2}\left(H ; \mathbb{R}^{n}\right)$ then the covariance operator can be defined from the covariance function as follows:

$$
(C \phi)(x)=\int_{H} c(x, y) \phi(y) d y
$$

It is shown in [61] that any Gaussian distribution on a Hilbert space is defined in this way by some function $m \in H$ and covariance operator $\mathcal{C} \in \mathcal{L}(H, H)$ where $\mathcal{C}$ is trace class, non-negative and self-adjoint. Moreover, any function $m \in H$ and operator $\mathcal{C} \in \mathcal{L}(H, H)$ satisfying these conditions has a Gaussian distribution associated with it.

One way to go about drawing a sample from this is to create a Fourier series using random coefficients with specific distributions.

Definition 1.7.2. Let $\left\{\lambda_{k}\right\}_{k=1}^{\infty}$ and $\left\{\phi_{k}\right\}_{k=1}^{\infty}$ be the eigenvalues and eigenfunctions respectively, of the covariance operator $\mathcal{C}$, such that $m=m_{k} \phi_{k}$. Then we define the Karhunen Loeve expansion to be

$$
u=\sum_{k=1}^{\infty} \sqrt{\lambda_{k}} \xi_{k} \phi_{k}, \quad \xi_{k} \quad \text { iid, } \quad \xi_{1} \sim \mathcal{N}(0,1) .
$$

Then $u \in H$, and $u \sim \mathcal{N}(0, \mathcal{C})$. So to draw a sample from our desired distribution, note that

$$
u+m=v \sim \mathcal{N}(m, \mathcal{C}) .
$$

As long as we are able to characterise the spectrum of a given covariance operator $\mathcal{C}$, we are able to approximate samples from $\mathcal{N}(m, \mathcal{C})$ via truncation of the sum. In
practise this means picking $N \in \mathbb{N}$ as large as our computational capacity will allow, in order to keep the truncation error

$$
\left\|\sum_{k=N+1}^{\infty}\left(m_{k}+\sqrt{\lambda_{k}} \xi_{k}\right) \phi_{k}\right\|,
$$

as small as possible. We now wish to ascertain certain different regularity conditions of draws from such a distribution, dependent on the choice of covariance operator $\mathcal{C}$.

### 1.8 Regularity of Gaussian Random Fields

In this section we will look at the regularity of draws from a particular family of Gaussian random field distributions that we will be using later in the thesis. In general, excluding chapter 6, we will be considering two-dimensional domains for our random fields. However, since in chapter 6 we will be required to sample random fields in one dimension, we aim to keep this section as general as possible, and the parameter $d \in \mathbb{N}$ will denote the dimension of the domain. Consider the distribution

$$
\mathcal{N}\left(m, A^{-\alpha}\right)
$$

for some $\alpha \in \mathbb{R}$. We wish to ascertain for what values of $\alpha$ will a draw from this distribution be in certain function spaces, for instance $H$ or $H^{s}$.

Since $\phi_{k}$ is an orthogonal basis in $L^{2}(\Omega)$, we can say $\forall u \in L^{2}(\Omega), \exists u_{k}$ such that

$$
u=\sum_{k} u_{k} \phi_{k}
$$

Therefore we can define our norm to be $\|u\|_{s}^{2}=\sum_{k}\left(1+|k|^{2 s}\right) u_{k}^{2}$. This is equivalent to saying that both $u, \mathcal{K} u \in L^{2}(\Omega)$ where $\mathcal{K}$ is some $s$-fold differential operator.

Lemma 1.8.1. Suppose $v \sim \mathcal{N}\left(0, \delta A^{-\alpha}\right)$ for some $\alpha>s+\frac{d}{2}$ for some $s \geq 0$. Then $v \in H^{s}$.

Proof. Suppose $v \sim \mathcal{N}\left(0, \delta A^{-\alpha}\right)$, then by the Karhunen-Loeve expansion,

$$
v=\sum_{k} \sqrt{\lambda_{k}} \phi_{k} \xi_{k}, \quad \xi_{k} \sim \mathcal{N}(0,1) \text { i.i.d. }
$$

This implies that

$$
\begin{aligned}
\mathbb{E}\left(\|v\|_{s}^{2}\right) & =\mathbb{E}\left(\sum_{k}\left(1+|k|^{2 s}\right) \lambda_{k} \xi_{k}^{2}\right) \\
& =\sum_{k}\left(1+|k|^{2 s}\right) \lambda_{k} .
\end{aligned}
$$

Here the bigger $s$ is, the faster the $\lambda_{k}$ must decay if the sum is to converge to a finite limit.

Observe the following:

$$
\mathbb{E}\left(\|v\|_{s}^{2}\right)<\infty \Longrightarrow\|v\|_{s}<\infty \quad \text { a.s. }
$$

If is this wasn't the case, then there would be a set of non-zero measure on which $\|v\|_{s}$ was infinite. However this leads us to the conclusion that the expectation of $\|v\|_{s}^{2}$ would be infinite, leading to a contradiction. So by this argument, we have that $\mathbb{E}\left(\|v\|_{s}^{2}\right)<\infty$ is sufficient to show that $v \in H^{s}(\Omega)$ almost surely.

$$
\begin{array}{ccc}
\mathbb{E}\left(\|v\|_{s}^{2}\right) & = & \sum_{k}\left(1+|k|^{2 s}\right) \lambda_{k} \\
& \propto & \sum_{k}\left(1+|k|^{2 s}\right)|k|^{-2 \alpha} \\
\leq C \sum_{k}|k|^{2(s-\alpha)} .
\end{array}
$$

By comparison to integrals it follows that for $v$ to be in $H^{s}$ we require $\alpha>s+\frac{d}{2}$.

Now that we know for which values of $\alpha$ we are in $H^{s}$ for differing values of $s$, we can now consider the Sobolev embedding theorem, as described in [50] which is generalised for fractional Hilbert spaces in [66] as summarised by the following theorem.

Theorem 1.8.1. Sobolev Embedding Theorem: Let $\Omega$ be a bounded $C^{k}$ domain in $\mathbb{R}^{d}$, and suppose that $u \in H^{k}$.
(i) If $k<d / 2$ then $u \in L^{2 d /(d-2 k)}(\Omega)$, and there exists a constant $C$ such that

$$
\|u\|_{L^{2 d /(d-2 k)}(\Omega)} \leq C\|u\|_{H^{k}(\Omega)}
$$

(ii) If $k=d / 2$ then $u \in L^{p}(\Omega)$ for every $1 \leq p<\infty$, and for each $p$ there exists a constant $C=C(p)$ such that

$$
\|u\|_{L^{p}(\Omega)} \leq C\|u\|_{H^{k}(\Omega)} .
$$

(iii) If $k>j+(d / 2)$ then $u \in C^{j}(\bar{\Omega})$, and there exists a constant $C$ such that

$$
\|u\|_{C^{j}(\bar{\Omega})} \leq C\|u\|_{H^{k}(\Omega)}
$$

Note that since $\Omega$ is bounded, it follows trivially in part (iii) that $u \in L^{p}(\Omega)$ for every $1 \leq p \leq \infty$.

Applying these theorems, we get the following 3 cases;

Case 1: $1<\alpha \leq 2$

As shown before, $u \in H^{s}$ for $s<\alpha-1$. For this range of values of $\alpha$, we have that $0<s<1$, and therefore that $s<\frac{n}{2}=1$. We can now apply part (i) of the Sobolev embedding theorem (1.8.1).

Therefore in this case, $u \in L^{\frac{2 n}{n-2 s}}$ for $s<\alpha-1$.

Case 2: $s=\frac{n}{p}$

Next we consider the borderline case where $s p=n$. If $\alpha=2$, then we are in the case as described above, and we are in some $L^{p}$ for a finite $p$. How ever if $\alpha=2+\epsilon$ for some
$\epsilon>0$, then we can show, as we will in the next case, that $u$ is Hölder continuous. So therefore this borderline case is pretty irrelevant.

In this case $-\mathcal{L}^{\frac{1}{2}} u \in L^{2}(\Omega)$. So therefore, $u \in H_{0}^{1}(\Omega)$. We can use the Sobolev embedding theorem. Since $s p=n$ we are in case (ii) of (1.8.1), so we can say that for this value of $\alpha, u \in L^{\varphi}(\Omega)$. However as described before this does not actually occur as if we have $\alpha \leq 2$ we are in case 1 , and if $\alpha>2$ we are in case 3 .

Case 3: $\alpha>2$

For this range of values of $\alpha$, we can say something much stronger about the regularity of $u$. Since $s<\alpha-1$, we have that for this range of $\alpha$ we can guarantee that $s>\frac{p}{n}$ so that we can apply part (iii) of (1.8.1). So if $s>j+\frac{n}{2}$ for $j \in \mathbb{N} \bigcup\{0\}$, then $u$ is continuous and $j$-times differentiable.

Using this information, we are able to create Gaussian random fields which are (almost surely) in a whole range of useful function spaces. This will be crucial later in thesis.

### 1.9 Absolute Continuity of Gaussian Measures

In certain situations we will need to show absolute continuity of Gaussian measures with respect to another. Let us consider the Feldman-Hajek theorem, which is presented in [61], but here we present the version given in [73].

Theorem 1.9.1. Two Gaussian measures $\mu_{i}=\mathcal{N}\left(m_{i}, \mathcal{C}_{i}\right), i=1,2$, on a Hilbert space $\mathcal{H}$ are either singular or equivalent. They are equivalent if and only if the following three conditions hold:

- $\operatorname{Im}\left(\mathcal{C}_{1}^{\frac{1}{2}}\right)=\operatorname{Im}\left(C_{2}\right)^{\frac{1}{2}}:=E ;$
- $m_{1}-m_{2} \in E$;
- the operator $T:=\left(\mathcal{C}^{-\frac{1}{2}} \mathcal{C}_{2}^{\frac{1}{2}}\right)\left(\mathcal{C}^{-\frac{1}{2}} \mathcal{C}_{2}^{\frac{1}{2}}\right)^{*}-I$ is Hilbert-Schmidt in $\bar{E}$.

For example, we may apply this theorem to the two Gaussian measures

$$
\begin{gather*}
\mu_{1}=\mathcal{N}\left(0,(-\Delta)^{-\alpha}\right)  \tag{1.12}\\
\mu_{2}=\mathcal{N}\left(0,(\ell I-\Delta)^{-\alpha}\right) \tag{1.13}
\end{gather*}
$$

for some $\ell>0, \alpha>\frac{1}{2}$, on $\mathcal{H}=L^{2}(\Omega)$ where $\Omega=[0,1)$ with periodic boundary conditions. To show this, we first need the following lemma.

Lemma 1.9.1. For any two positive-definite, self-adjoint, bounded linear operators $\mathcal{C}_{i}$ on a Hilbert space $\mathcal{H}, i=1,2$, the condition $\operatorname{Im}\left(\mathcal{C}_{1}^{\frac{1}{2}}\right) \subset \operatorname{Im}\left(\mathcal{C}_{2}^{\frac{1}{2}}\right)$ holds if and only if there exists a constant $K>0$ such that

$$
\left\langle h, \mathcal{C}_{1} h\right\rangle \leq K\left\langle h, \mathcal{C}_{2} h\right\rangle \quad \forall h \in \mathcal{H} .
$$

Corollary 1.9.1. The Gaussian measures $\mu_{1}$ and $\mu_{2}$ as defined by (1.12-1.13) are equivalent.

Proof. We conduct the proof by showing each of the conditions in theorem 1.9.1 are met.

- The two covariance operators share the same eigenfunctions,

$$
\phi_{k}(x)=\exp (2 \pi i k \cdot x),
$$

and have the eigenvalues

$$
\lambda_{k}^{1}=\left(4 \pi^{2}|k|^{2}\right)^{-\alpha}, \lambda_{k}^{2}=\left(\ell+4 \pi^{2}|k|^{2}\right)^{-\alpha},
$$

respectively. Therefore, for any $h=\sum_{k}\left\langle h, \phi_{k}\right\rangle \phi_{k}=\sum_{k} h_{k} \phi_{k} \in \mathcal{H}$,

$$
\left(\frac{4 \pi^{2}}{\ell+4 \pi^{2}}\right)^{-\alpha} \leq \frac{\left\langle h, \mathcal{C}_{2} h\right\rangle}{\left\langle h, \mathcal{C}_{1} h\right\rangle}=\frac{\sum_{k}\left(\ell+4 \pi^{2} k^{2}\right)^{-\alpha} h_{k}^{2}}{\sum_{k}\left(4 \pi^{2} k^{2}\right)^{-\alpha} h_{k}^{2}} \leq 1
$$

Therefore, by lemma 1.9, the first condition is met.

- The second condition is met trivially.
- Note that

$$
T=\mathcal{C}_{2}^{-\frac{1}{2}} \mathcal{C}_{2} \mathcal{C}_{1}^{-\frac{1}{2}}-I
$$

is diagonalised in the same basis as the $\mathcal{C}_{i}$ and has eigenvalues

$$
\frac{1}{\left(4|k|^{2} \pi^{2}\right)^{\alpha}} .
$$

These are square summable if $\alpha>\frac{1}{2}$, and so the final condition is met.

### 1.10 Bayesian Statistics

We cannot discuss Bayesian statistics without first stating Bayes' theorem. We will first consider the case that the state space is finite dimensional.

Theorem 1.10.1 (Bayes' Theorem). Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, for $A, B \in \mathcal{F}$,

$$
\mathbb{P}(A \mid B)=\frac{\mathbb{P}(B \mid A) \mathbb{P}(A)}{\mathbb{P}(B)}
$$

Proof. Result follows from definition of conditional probabilities, namely that

$$
\mathbb{P}(A \mid B)=\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}
$$

This elementary result is pivotal in many areas of science and technology. In chapters 2-5 we will be using this theory to attack various data assimilation problems, in which we wish to use noisy observations of a system to infer on the initial state of that system $u_{0}$. Suppose that $B$ is the event that we have made observations $y$ of our dynamical system. $A$ is the event that the initial condition of the system that was observed was equal to $u_{0}$. Our aim is to use this data to statistically infer on the initial condition $u_{0}$; that is, garner information about the probability measure $\mathbb{P}\left(u_{0} \mid y\right)$. Using this, we can informally show that

$$
\begin{aligned}
\mathbb{P}\left(u_{0} \mid y\right) & =\frac{\mathbb{P}\left(y \mid u_{0}\right) \mathbb{P}\left(u_{0}\right)}{\mathbb{P}(y)} \\
& \propto \mathbb{P}\left(y \mid u_{0}\right) \mathbb{P}\left(u_{0}\right),
\end{aligned}
$$

since $\mathbb{P}(y)$ is a constant.
Therefore, given prior information about $u_{0}$, in the form of a probability distribution $\mu_{0}=\mathbb{P}\left(u_{0}\right)$, and a likelihood density that $y$ was created with any given $u_{0}$, $\mathbb{P}\left(y \mid u_{0}\right)$, we are able to analyse the measure of interest, $\mu=\mathbb{P}\left(u_{0} \mid y\right)$.

We now consider the case that the state space is a function space, and therefore infinite dimensional. In this case, we cannot calculate the densities explicitly. We can, however, calculate the Radon-Nikodym derivative, if $\mu$ (the posterior measure) is absolutely continuous with respect to $\mu_{0}$ (the prior measure). Then the infinite dimensional analogue of Bayes' theorem is given by:

$$
\begin{equation*}
\frac{d \mu}{d \mu_{0}} \propto \mathbb{P}\left(y \mid u_{0}\right) \tag{1.14}
\end{equation*}
$$

Since the right hand side is non-negative, a function $\Phi(\cdot)$ can be found such that $\mathbb{P}\left(y \mid u_{0}\right) \propto \exp \left(-\Phi\left(u_{0}\right)\right)$, meaning that the Radon Nikodym derivative is given by

$$
\begin{equation*}
\frac{d \mu}{d \mu_{0}} \propto \exp (-\Phi(\cdot)) \tag{1.15}
\end{equation*}
$$

This function is related to the likelihood function.
We wish to calculate how likely it is that a given initial condition $u_{0}$ will produce our noisy observations, $y$. This is encoded in the likelihood function of the data given $u_{0}$.

We assume that the noise on our observations is Gaussian such that

$$
y=\mathcal{G}\left(u_{0}\right)+\xi, \xi \sim \mathcal{N}(0, \Sigma),
$$

for some covariance matrix $\Sigma$, where $\mathcal{G}: H \rightarrow \mathbb{R}^{N}$ denotes our observation operator. So the likelihood that we will observe $y$ given that we have initial condition $u(0, x)=u_{0}(x)$ is simply given by

$$
\begin{aligned}
\mathbb{P}\left(y \mid u_{0}\right) & \propto \exp \left(-\frac{\|\xi\|^{2}}{2}\right) \\
& =\exp \left(-\frac{1}{2}\left\|y-\mathcal{G}\left(u_{0}\right)\right\|_{\Sigma}^{2}\right) \\
& :=\exp \left(-\frac{1}{2}\left(y-\mathcal{G}\left(u_{0}\right)\right)^{T} \Sigma^{-1}\left(y-\mathcal{G}\left(u_{0}\right)\right)\right) .
\end{aligned}
$$

Therefore, by Bayes' law (1.14), assuming $\mu$ is $\mu_{0}$-measurable, the RadonNikodym derivative is given by

$$
\begin{equation*}
\frac{d \mu\left(u_{0}\right)}{d \mu_{0}}=\exp \left(-\frac{1}{2}\left\|y-\mathcal{G}\left(u_{0}\right)\right\|_{\Sigma}^{2}\right) . \tag{1.16}
\end{equation*}
$$

The next task is to find an appropriate prior Gaussian distribution $\mu_{0}=\mathbb{P}\left(u_{0}\right)$ with enough regularity to ensure that $\mu=\mathbb{P}\left(u_{0} \mid y\right)$ is $\mu_{0}$-measurable.

Theorem 1.10.2. If $\mathcal{G}: H \rightarrow \mathbb{R}^{N}$ is $\mu_{0}$ measurable then the posterior measure $\mu(d x)=$ $\mathbb{P}(d x \mid y)$ is absolutely continuous with respect to the prior measure $\mu_{0}(d x)$ and has Radon-Nikodym derivative given by (1.16).

Proof. Proof given in Theorem 2.1 in [24].

Corollary 1.10.1. If $\mathcal{G}: H \rightarrow \mathbb{R}^{N}$ is continuous on a function space $X$ and $\mu_{0}(X)=1$, then the posterior measure $\mu(d x)=\mathbb{P}(d x \mid y)$ is absolutely continuous with respect to the prior measure $\mu_{0}(d x)$ and has Radon-Nikodym derivative given by (1.16).

Proof. Result follows since if $\mathcal{G}$ is continuous on $X$, then it is measurable on $X$, and therefore is almost surely $\mu_{0}$-measurable.

### 1.11 Relationship with Tikhonov Regularisation

One approach to data assimilation is to try to find the minimiser of the following expression:

$$
\min _{u}\|\mathcal{G}(u)-y\|_{\Sigma}^{2}
$$

Solutions to this problem can, however, be very rough and have undesirable properties, or the solution may not be unique or even exist. Therefore it may be appropriate to add a penalty term to this minimisation to ensure minimal amounts of regularity in the solution, so that we are trying to minimise $J$ where

$$
J(u)=\min _{u} \frac{1}{2}\|\mathcal{G}(u)-y\|_{\Sigma}^{2}+\frac{1}{2}\|u\|_{X}^{2}
$$

for some function space $X$. This is termed Tikhonov regularisation[45]. If $X$ is the Cameron-Martin space for some Gaussian measure $\mu_{0}$, for which the measure $\mu=\mathbb{P}(u \mid y)$ is absolutely continuous, then this problem is equivalent to trying to maximise the probability density from the Bayesian posterior we have previously defined. As such, the theorems that we will go onto describe for the various situations in chapters 2-4 are also applicable to those wishing to attempt these types of variational approach, as they give insight into good choices of penalty term to ensure existence of global minima of $J$.

### 1.12 Markov Chain Monte Carlo Methods

Markov Chain Monte Carlo (MCMC) methods are a family of methods which allow us to sample from complex distributions which are hard to characterise analytically. In essence, we create a Markov chain whose invariant measure is the measure we wish to characterise. Then, assuming the chain is ergodic, the Markov chain should converge in distribution to the desired measure. By taking statistics of a sufficiently large number of samples, we can accurately estimate the statistics of this target measure. In other words, if the $X_{n}$ are the states of our Markov chain, $f$ is an observable function of the $X_{n}$, and $\pi$ is the invariant density of the Markov chain, then assuming ergodicity, we have that

$$
\int_{\Omega} f(X) d \pi=\lim _{n \rightarrow \infty} \frac{1}{N} \sum_{n=1}^{N} f\left(X_{n}\right)
$$

### 1.12.1 MCMC methods on finite space

The following algorithm is the basis for all Metropolis Hastings based MCMC methods for sampling from a given measure $\mu$ with density $\pi$ on $\mathbb{R}^{n}[76,19] . q$ is a transition density, which can be a function of both $y$, the proposal, and $x_{i-1}$, the previously accepted state of the chain.

1: $x=x_{0}$
2: for $i=1: N$ do
3: $\quad$ Sample $y \sim q\left(x_{i-1},.\right), y \in \mathbb{R}^{n}$
4: $\quad \alpha\left(x_{i-1}, y\right)=\min \left\{1, \frac{\pi(y) q\left(y, x_{i-1}\right)}{\pi\left(x_{i-1}\right) q\left(x_{i-1}, y\right)}\right\}$
5: $\quad$ Sample $u \sim U([0,1))$
6: if $u<\alpha(x, y)$ then
7: $\quad x_{i}=y$

8: else

$$
x_{i}=x
$$

end if
end for

The function $\alpha$ is termed the acceptance probability, and this is determined completely by the target density $\pi$ and the transition density $q$.

Note here that each draw from a random distribution, at steps 3 and 5 are independent of each other, and independent for each iteration $i$.

These algorithms are defined when the target and proposal have densities $\pi(\cdot)$ and $q(x, \cdot)$ with respect to Lebesgue measure. For us it is important to generalise this setting to situations where densities are defined with respect to Gaussian measures.

## The Independence Sampler

The independence sampler is so named due to the fact that the transition density is independent of the currently accepted state. In independence samplers it is more important to carefully choose your transition density if you want a sampler that explores the state space of the target measure efficiently. The algorithm works as follows:

1: $x=x_{0}$

2: for $i=1: N$ do
3: $\quad$ Sample $y \sim q(\cdot) \in \mathbb{R}^{n}$
4: $\quad \alpha\left(x_{i-1}, y\right)=\min \left\{1, \frac{\pi(y) q\left(x_{i-1}\right)}{\pi\left(x_{i-1}\right) q(y)}\right\}$
5: $\quad$ Sample $u \sim U([0,1))$
6: if $u<\alpha(x, y)$ then
7: $\quad x_{i}=y$

```
    else
    \(x_{i}=x\)
    end if
end for
```

This algorithm only works efficiently if a good choice of the transition density $q$ is chosen. Moreover it will not converge to $\pi$ if the support of $\pi$ is not contained within the support of $q$.

## The Standard Random Walk Metropolis Hastings Algorithm

In the following two methods we are aiming to sample from a target density $\pi$ which satisfies $\pi(\cdot) \propto \exp (-\Psi(\cdot))$ for some function $\Psi$. Consider the following finite dimensional SDE:

$$
\begin{equation*}
d X=-\gamma A \nabla \Phi(X)+\sqrt{2 A} d W \tag{1.17}
\end{equation*}
$$

With $\gamma=1$ this SDE has invariant measure $\pi$. The discretization of this SDE with $\gamma=0$ gives us the proposal for the Standard Random Walk Metropolis Hastings (SRWMH) algorithm.

In the SRWMH method, commonly used in many applications, the proposal density is chosen to be a Gaussian centred on the currently accepted state, with covariance related to the prior density $\pi_{0}$ which is also a Gaussian. Let $\pi_{0} \sim \mathcal{N}(m, C)$ and without loss of generality let $m=0$. In this example we assume that the target measure is change of measure from this Gaussian, ie $\Psi(\cdot)=\Phi(\cdot)+\frac{1}{2}\left|C^{1 / 2} \cdot\right|^{2}$. Then the algorithm is as follows;

1: $x=x_{0}$
2: for $i=1: N$ do

3:

5: $\quad$ Sample $u \sim U([0,1))$
6: if $u<\alpha(x, y)$ then
$7:$

8:

9: $\quad x_{i}=x$
10 :
1: end for

This algorithm does however suffer from the curse of dimensionality. For a fixed value of $\beta$, but with increasing resolutions (ie more terms in the Fourier series are non-zero), the average acceptance probabilities decrease in such a way as to cause the complexity of the algorithm to be $\mathcal{O}(n)$ where $n$ is the number of degrees of freedom in the state space. It has been shown in $[63,65]$ that the optimal average acceptance probability for this algorithm is approximately $23.4 \%$.

## The Standard Metropolis Adjusted Langevin Algorithm

This method, once again commonly used in many applications, is similar in many ways to the random walk method described in the previous section. Similarly, we aim to draw samples from the target distribution with density $\pi \propto \exp (-\Phi(\cdot))$. However, in this algorithm gradient information about the posterior measure is included in the proposal, to allow for more efficient exploration of the state space. The proposal comes from the discretization of the $\operatorname{SDE}(1.17)$ with $\gamma=1$, with time step $\Delta t$. A preconditioner matrix $A$ can be chosen to be $I$ or any positive symmetric matrix $C$.
$x=x_{0}$
for $i=1: N$ do
Sample $y$ such that $\frac{y-x_{i-1}}{\Delta t}=-A \nabla \Phi\left(x_{i-1}\right)+\sqrt{\frac{2 A}{\Delta t}} \xi_{i}$, where $\xi_{i} \sim \mathcal{N}(0, I)$

$$
\rho\left(x_{i-1}, y\right)=\Phi\left(x_{i-1}\right)+\frac{\Delta t}{4}\left|A^{1 / 2} \Phi\left(x_{i-1}\right)\right|^{2}+\frac{1}{2}\left\langle y-x, \nabla \Phi\left(x_{i-1}\right)\right\rangle
$$

$$
\alpha\left(x_{i-1}, y\right)=\min \left\{1, \exp \left(\rho\left(x_{i-1}, y\right)-\rho\left(y, x_{i-1}\right)\right\}\right.
$$

Sample $u \sim U([0,1))$
if $u<\alpha(x, y)$ then

$$
x_{i}=y
$$

else

$$
x_{i}=x
$$

end if
end for
This algorithm does again suffer from the curse of dimensionality. However, the increased efficiency of including gradient information means that the order of complexity is reduced to $\mathcal{O}\left(n^{1 / 3}\right)$. We also have the extra condition here that the function $\Phi$ must be differential on the support of the proposal distribution. In [64] it is shown that the optimal acceptance rate for this scheme is $57.4 \%$.

### 1.12.2 MCMC methods on function spaces

We will now consider equivalent MCMC methods which instead of acting on a finite dimensional space, act on spaces of functions. These algorithms have the great advantage of not being affected by the curse of dimensionality, and so are robust under refinement of discretisations. We will show later in the thesis what effect this has in a practical sense.

## The Independence Sampler

In a function space setting we are not able to define the probability densities as these quantities are not available to us or are infinite. Instead, we define the problem through the Radon-Nikodym derivative of our target measure $\mu$ with respect to a reference measure, or a prior measure. In the case of the independence sample, we choose a proposal measure $\nu$. As in the finite dimensional case it is important that the support of the target measure is contained within the support of the proposal measure, but we also have the added condition now that the proposal density must be absolutely continuous with respect to the target density. The acceptance probabilities are then defined in terms of the Radon Nikodym which exists due to this condition, namely

$$
\frac{d \nu}{d \mu}
$$

```
\(x=x_{0}\)
for \(i=1: N\) do
    Sample \(y \sim \nu\)
    \(\alpha\left(x_{i-1}, y\right)=\min \left\{1, \frac{d \nu}{d \mu}\left(x_{i-1}\right) / \frac{d \nu}{d \mu}(y)\right\}\)
    Sample \(u \sim U([0,1))\)
    if \(u<\alpha(x, y)\) then
        \(x_{i}=y\)
    else
        \(x_{i}=x\)
    end if
    end for
```

As in the finite dimensional version, this algorithm only works efficiently if a good choice of the proposal measure $\nu$ is chosen.

## The Random Walk Metropolis Hasting Algorithm

In the following two sections we will consider target measures $\mu$ which are absolutely continuous with respect to a prior measure $\mu_{0}$, as in many Bayesian applications such as those addressed in this thesis, with Radon Nikodym derivative given by (1.15) for some function $\Phi$.

Consider the SDE

$$
\begin{equation*}
\frac{d u}{d s}=\mathcal{K}(\mathcal{L} u-\gamma \nabla \Phi(u))+\sqrt{2 \mathcal{K}} \frac{d b}{d s} \tag{1.18}
\end{equation*}
$$

where $b$ is a cylindrical Brownian motion, $\mathcal{K}$ is a positive, symmetric pre-conditioning operator. $\mathcal{L}=(-\mathcal{C})^{-1}, \mathcal{C}$ is the covariance operator as described in prior distribution on $u_{0}$ (this will be discussed in further detail in later chapters), and $\gamma \in\{0,1\}$.

For $\gamma=1$, this SDE has invariant measure $\mu$, the distribution that we wish to sample from. For $\gamma=0$, it has invariant density $\mu_{0}$, the prior distribution on $u$.

If our proposal distribution is closer to that of our target distribution, then we are able to make larger changes in that proposal and still be able to accept samples, so that we explore the state space quicker and more thoroughly. Therefore if we make proposals based on a discretization of this SDE we will have a more efficient MCMC algorithm.

If we discretize (1.18) using the theta method with $\theta=\frac{1}{2}$, and with time increment $\Delta t$, we get the following.

$$
\begin{equation*}
\frac{u_{n+1}-u_{n}}{\Delta s}=\frac{1}{2} \mathcal{K}\left(\mathcal{L} u_{n}+\mathcal{L} u_{n+1}\right)-\gamma \mathcal{K} \nabla \Phi\left(u_{n}\right)+\sqrt{\frac{2 \mathcal{K}}{\Delta s}} \xi_{0} \tag{1.19}
\end{equation*}
$$

where $\xi_{0} \sim \mathcal{N}(0, I)$.
This gives us a construct in which we can define our transition probabilities on the state space. In essence, given that our current state is $u$, we make the proposal for our next state to be $v$ where

$$
\begin{equation*}
\frac{v-u}{\Delta s}=\frac{1}{2} \mathcal{K}(\mathcal{L} u+\mathcal{L} v)-\gamma \mathcal{K} \nabla \Phi(u)+\sqrt{\frac{2 \mathcal{K}}{\Delta s}} \xi_{0} \tag{1.20}
\end{equation*}
$$

The parameter $\Delta s$ can be tuned to give a sensible acceptance rate. For example, if we were to have too large a value for $\Delta t$, the proposal would be quite far away from the currently accepted state, and may be in the tails of the target distribution. If this is the case, the acceptance probability would be very small, and the proposal would be rejected the majority of the time, and we would not explore the state space efficiently.

Conversely, if we choose $\Delta t$ too small, the proposal would be very close to the current state. As such the acceptance probability would be very high, but as we have not moved very far, we once again will not explore the state space efficiently. An optimal scheme with well chosen $\Delta t$ will have an acceptance rate neither close to 1 nor 0 . The optimal acceptance probability changes depending on our choice of proposal.

The algorithm, when $\gamma=1$, is termed the MCMC-adapted Langevin algorithm. This algorithm requires $\Phi$ to have a derivative which is measurable with respect to the prior measure on $u$, and as such we will not be using this version of the algorithm in this thesis. This algorithm is useful if you can show this though, as including gradient information in the proposal allows the state space to be explored more efficiently.

Instead, we will use this algorithm with $\gamma=0$, which is termed the Random Walk Metropolis Hasting (RWMH) algorithm. We will be using the preconditioned version of this algorithm with $\mathcal{K}=\mathcal{C}$, giving us the proposal;

$$
\frac{v-u}{\Delta s}=-\frac{1}{2}(u+v)+\sqrt{\frac{2 \mathcal{C}}{\Delta s}} \xi_{0}
$$

This can be rearranged, with a substitution of parameters, to read as

$$
\begin{equation*}
v=\left(1-\beta^{2}\right)^{1 / 2} u+\beta w \tag{1.21}
\end{equation*}
$$

where $\beta \in[0,1]$ scales the size of steps made in the proposal, and $w \sim \mathcal{N}(0, \mathcal{C})$. The acceptance probability of moving from the current state $u$ to a proposed state from the transition kernel defined by 1.21 is given by

$$
\alpha(u, v)=\exp (\Phi(u)-\Phi(v))
$$

To demonstrate why this is the case, consider the following informal calculation.
Proposal state: $v=\left(1-\beta^{2}\right)^{\frac{1}{2}}+\beta \mathcal{C}^{\frac{1}{2}} \xi, \xi \sim \mathcal{N}(0,1), \beta \in(0,1]$.
Target measure density:

$$
\begin{aligned}
\pi(u) & =\exp (-\Phi(u)) \exp \left(\frac{1}{2}\langle u, \mathcal{L} u\rangle\right) \\
& =\exp (-\Phi(u)) \exp \left(-\frac{1}{2}\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}\right)
\end{aligned}
$$

Transition kernel: $q(u, v)=\exp \left(-\frac{1}{2 \beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}}\left(v-\left(1-\beta^{2}\right)^{\frac{1}{2}} u\right)\right|^{2}\right)$.
All we need to show is that if we set $\pi(u) q(u, v)=\exp (-\Phi(u)) \times A(u, v)$, and show that $A(u, v)$ is symmetric in $u$ and $v$, then we have shown that $\alpha(u, v)=$ $\exp \left(\int_{\Omega} \Psi(u)-\Psi(v) d x\right)$.

$$
\begin{aligned}
-2 \log (A) & =\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}+\frac{1}{\beta^{2}} \left\lvert\, \mathcal{C}^{-\frac{1}{2}}\left(v-\left.\left(1-\beta^{2}\right)^{\frac{1}{2}} u\right|^{2}\right.\right. \\
& =\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}+\frac{1}{\beta^{2}} \left\lvert\, \mathcal{C}^{-\frac{1}{2}}\left(v-\left.\left(1-\beta^{2}\right)^{\frac{1}{2}} u\right|^{2}\right.\right. \\
& =\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}+\frac{1}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} v\right|^{2}+\frac{1-\beta^{2}}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}-\frac{2\left(1-\beta^{2}\right)^{\frac{1}{2}}}{\beta^{2}}\left\langle\mathcal{C}^{-\frac{1}{2}} u, \mathcal{C}^{-\frac{1}{2}} v\right\rangle \\
& =\frac{1}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} v\right|^{2}+\frac{1}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}-\frac{2\left(1-\beta^{2}\right)^{\frac{1}{2}}}{\beta^{2}}\left\langle\mathcal{C}^{-\frac{1}{2}} u, \mathcal{C}^{-\frac{1}{2}} v\right\rangle
\end{aligned}
$$

which is symmetric, since $\langle u, v\rangle=\langle v, u\rangle \forall u, v \in H$.
Unlike the SRWMH algorithm, this acceptance probability is completely independent of the number of degrees of freedom that you allow to be non-zero in the Fourier truncation. This means that the method is robust under refinements of mesh. This will be demonstrated in later chapters.

To clarify, the RWMH algorithm is defined as follows.

```
\(x=x_{0}\)
for \(i=1: N\) do
    Sample \(y=\left(1-\beta^{2}\right)^{1 / 2} x_{i-1}+\beta w\), where \(w \sim \pi_{0}\)
    \(\alpha\left(x_{i-1}, y\right)=\min \left\{1, \exp \left(\Phi\left(x_{i-1}\right)-\Phi(y)\right\}\right.\)
    Sample \(u \sim U([0,1))\)
    if \(u<\alpha(x, y)\) then
        \(x_{i}=y\)
    else
        \(x_{i}=x\)
    end if
end for
```

In the numerics that follow later, we allow a 'burn-in' period where we run the Markov chain for a period of time, and wait for it hopefully to converge in distribution to its invariant measure, $\pi$. During this period we can tune the parameter $\Delta t$ to give the Markov chain an average acceptance rate that is close to the optimal value. In [63] it is shown that for high dimensional SRWMH algorithms, the optimal acceptance probability is $23.4 \%$. This method however is different. However it is still pretty safe to expect the optimal average acceptance probability for this algorithm to be somewhere between 20-60\%.

Before considering a result concerned with the acceptance probabilities, let us define a set of assumptions on $\Phi$.

Assumptions 1. The function $\Phi: X \times \mathbb{R}^{m} \rightarrow \mathbb{R}$ satisfies the following:

1. there exists $p>0$ and for every $r>0$ a $K=K(r)>0$ such that, for all $u \in X$ and all $y:|y|<r$,

$$
0 \leq \Phi(u ; y) \leq K(r)\left(1+\|u\|_{\mathrm{X}}^{p}\right) ;
$$

2. for every $r>0$ there is $K(r)>0$ such that, for all $u, v \in X$ and $y \in \mathbb{R}^{m}$ with $\max \left\{|y|,\|u\|_{\mathrm{X}},\|v\|_{\mathrm{X}}\right\}<r$,

$$
|\Phi(u ; y)-\Phi(v ; y)| \leq K(r)\|u-v\|_{\mathrm{X}} ;
$$

3. there is $q>0$ and for every $r>0$ a $K=K(r)>0$ such that, for all $u \in X$ and all $y_{1}, y_{2} \in \mathbb{R}^{m}$ with $\max \left\{\left|y_{1}\right|,\left|y_{2}\right|\right\}<r$,

$$
\left|\Phi\left(u ; y_{1}\right)-\Phi\left(u ; y_{2}\right)\right| \leq K(r)\left(1+\|u\|_{\mathrm{X}}^{q}\right)\left|y_{1}-y_{2}\right| .
$$

Consider the following theorem regarding acceptance probabilities, adapted from a result that appears in [23].

Theorem 1.12.1. Let $\mu_{0}$ be a Gaussian measure on a Hilbert space $\left(X,\|\cdot\|_{\mathrm{x}}\right)$ with $\mu_{0}(X)=1$ and let $\mu$ be a second measure on $X$ given by the Radon-Nikodym derivative (1.16), satisfying assumptions 1 ;

Then the preconditioned random walk algorithm with fixed $\beta$ are defined on $X$ and, furthermore, the acceptance probability satisfies

$$
\lim _{\beta \rightarrow 0} \mathbb{E}^{\eta} a(u, v)=1 .
$$

It is important to point out here that whilst it is not expected that an acceptance probability close to 1 is optimal for MCMC methods, it is still desirable to be able to tune the acceptance probability arbitrarily in $[0,1]$, as can be done in finite dimensions.

## The Metropolis Adjusted Langevin Algorithm

Similarly, there is a function space equivalent to the MALA method. This method uses a proposal derived from a discretization of (1.18) with time step $\Delta t$, partially using the theta method with $\theta=\frac{1}{2}$. We present here the preconditioned version with $A=C$, the covariance of the prior distribution $\mu_{0} \sim \mathcal{N}(0, C)$.

```
\(x=x_{0}\)
    for \(i=1: N\) do
    Sample \((2+\Delta t) y=(2-\Delta t) x_{i-1}-2 \Delta C \nabla \Phi\left(x_{i-1}\right)+\sqrt{8 C \Delta t} \xi\), where \(\xi \sim \mathcal{N}(0, I)\)
    \(\rho\left(x_{i-1}, y\right)=\Phi\left(x_{i-1}\right)+\frac{1}{2}\left\langle y-x_{i-1}, \nabla \Phi\left(x_{i-1}\right)\right\rangle+\frac{\Delta t}{4}\left\langle x_{i-1}+y, \nabla \Phi\left(x_{i-1}\right)\right\rangle+\)
        \(\frac{\Delta t}{4}\left\|C^{-1 / 2} \nabla \Phi\left(x_{i-1}\right)\right\|^{2}\)
    \(\alpha\left(x_{i-1}, y\right)=\min \left\{1, \exp \left(\rho\left(x_{i-1}, y\right)-\rho\left(y, x_{i-1}\right)\right\}\right.\)
    Sample \(u \sim U([0,1))\)
    if \(u<\alpha(x, y)\) then
        \(x_{i}=y\)
        else
        \(x_{i}=x\)
        end if
    end for
```

As in the finite-dimensional case, this algorithm requires fewer samples to give an accurate approximation of the target distribution than the random walk method. However, $\Phi$ must be shown to be differentiable on the support of $\mu_{0}$. The calculation of $\nabla \Phi$ is also far from trivial, requiring computation of the adjoint problem alongside that of the forward problem for each proposed state $y$, so that the computational cost per sample is increased.

For the purposes of this thesis, we will deal exclusively with the RWMH algorithm, with a view to attempting to implement the Langevin algorithms at some point in the future.

### 1.13 Computer Generation of (Pseudo-) Random Numbers

When dealing with the generation of very large numbers of random numbers computationally, as we will be doing throughout the whole thesis, it is very important to think carefully about what algorithms we are using to create these numbers. The built-in random number generators in many languages, including $C$, are often thought to be inadequate for applications such as MCMC methods. Therefore, we have chosen a "Mersenne Twister" algorithm from the GNU GSL library to produce uniformly distributed numbers in all computations throughout this thesis, which has a period of $2^{19937}-1$, plenty big enough for our needs [1]. These are then converted into normally distributed random numbers via the Box-Müller[75] algorithm. These two algorithms used together produce high-quality random numbers which have passed standardised statistical tests such as "DIE HARD" [55].

### 1.14 Interpolation Methods

In section 1.4 it was shown how the Stokes problem can be solved analytically in Fourier space. In practise, we cannot compute this value for infinite dimensional functions, but instead we aim to compute this as accurately as we can on a truncated set of Fourier modes, that we then evaluate in real space through an FFT calculation. This returns an approximation of the value of the function on a set of grid points. However, we may wish to know the value of the function in between these grid points. We can approximate
these values through interpolating the values of the nearest neighbour grid points.

### 1.14.1 Bilinear Interpolation

Consider Figure 1.1 where we wish to approximate the value of $u(P)$, where $P$ is a point inside a square between points $Q_{1,1}, Q_{1,2}, Q_{2,1}$, and $Q_{2,2}$.


Figure 1.1: Bilinear Interpolation on a mesh, from [83]

Bilinear interpolation suggests that we take the approximation by first taking the linear interpolation in the $x$-direction. This yields

$$
\begin{array}{lll}
u_{i}\left(R_{1}\right) \approx \frac{x_{2}-x}{x_{2}-x_{1}} u_{i}\left(Q_{11}\right)+\frac{x-x_{1}}{x_{2}-x_{1}} u_{i}\left(Q_{21}\right) \quad \text { where } \quad R_{1}=\left(x, y_{1}\right) \\
u_{i}\left(R_{2}\right) \approx \frac{x_{2}-x}{x_{2}-x_{1}} u_{i}\left(Q_{12}\right)+\frac{x-x_{1}}{x_{2}-x_{1}} u_{i}\left(Q_{22}\right) \quad \text { where } \quad R_{2}=\left(x, y_{2}\right)
\end{array}
$$

We proceed by interpolating in the y-direction.

$$
u_{i}(x, y) \approx \frac{y_{2}-y}{y_{2}-y_{1}} u\left(R_{1}\right)+\frac{y-y_{1}}{y_{2}-y_{1}} u\left(R_{2}\right)
$$

If we re-label $Q_{1,1}=(0,0)$ and $Q_{2,2}=(\Delta, \Delta)$, then we can consider this simply
to be the bilinear form:

$$
u_{i}((x, y), t) \approx\left(\begin{array}{ll}
1-x & x
\end{array}\right)\left(\begin{array}{ll}
u((0,0), t) & u((0, \Delta), t) \\
u((\Delta, 0), t) & u((\Delta, \Delta), t)
\end{array}\right)\binom{1-y}{y}
$$

To calculate an approximation of the vector $u(x, y)$, we simply calculate this for $i \in\{1,2\}$. We then have an approximation for the value of our vector field within our whole state space $\mathbb{T}^{2}$, so we can successfully numerically approximate the solution of the Stokes problem with initial condition $u_{0}$ and forcing $g=\mathcal{P} f$. Therefore, we are also able to approximate the observation operator $\mathcal{G}_{E}$, and in turn, calculate relative densities of any two initial conditions with respect to well-defined posterior distributions satisfying corollary 2.5.3.

### 1.14.2 Bicubic Interpolation

We could alternatively use higher order interpolation methods which would give more accurate results, but at a higher computational cost. For example, we could use a bicubic interpolation[17] which approximates the function in the $x$ and $y$ directions by two cubic functions. This requires extra computational effort as the value of 3 different partial derivatives are required at each of the four corners of the cell. These could be calculated using finite difference methods - but these will add significant error on a course grid. Alternatively, these derivatives can be found easily in Fourier space, and the extra computational effort goes into computing 3 extra FFTs, essentially multiplying the complexity of the algorithm by a factor of 4 . For further details of the algorithm, see [82].

There are of course many other higher order methods which we omit here, but their computational cost eventually outweighs the improvement in accuracy of the algorithm.

For the vast majority of the numerics that will follow later in the thesis, we will employ the bilinear method, but we will also briefly explore the impact that a higher order interpolation method can have on the results.

## Chapter 2

## Eulerian Data Assimilation

### 2.1 Motivation

Collecting and analysing Eulerian data is crucially important for the weather forecasting community. Very often, this is satellite data comprising of observations of the flow of the atmosphere at fixed points in space. There are huge amounts of such data collected every day, with the intention of using it to find the current state of the atmosphere, in an effort to make more accurate predictions about the state of the atmosphere in the future.

These observations are of course not perfect, so a big part of this problem is also to ascertain a reasonable noise scenario for the model. We will, for the purposes of this thesis, assume that all observation errors are mean-zero Gaussian with known covariance, although the techniques which are described herein are perfectly adaptable to being used in other scenarios.

This type of inverse problem, where we have noisy observations of a system of which we wish to ascertain it's state, is termed data assimilation in meteorological[47]
and oceanographic[10] circles. These problems, when undertaken in an optimisation context, are often ill posed. That is, if we simply find the vector field which minimises the covariance-weighted distance between our observations and the observation operator applied to the vector field, then our solutions may be rough, undifferentiable and even discontinuous. This problem lies in the damping out of high frequency behaviour in dissipative fluid models. Therefore, the inverse map, which takes vector fields backwards through time, causes these frequencies to blow up. With the addition of noise, the least squares solution may cause many of these frequencies to have non-trivial values, causing us a problem if we try to regress this solution backwards in time.

The approach that is often taken in the optimisation community, is to enforce some kind of regularisation condition on the family of possible solutions. This can be done by using Tikhonov regularisation, or truncated iterative methods, or many others, as discussed in [45]. This is just one way that we can make this problem well posed.

In the following chapter, we will show how this ill posed optimisation problem can be translated into a well posed Bayesian inverse problem. We will analyse the forward problem in detail, so that we may determine prior measures with respect to which our observation operator is measurable. As discussed in section 1.11, this also gives insight into good choices of penalty term in a Tikhonov regularisation to ensure existence of global minima of the functional which we would wish to minimise.

We will then describe MCMC methods which are formulated on function space, that allow us to sample from well defined posterior measures. We will then present numerical results showing that these methods are valid, and then use these methods to probe the types of posterior distribution that we are interested in.

Let us note here that since this is a linear problem, and we intend to apply a Gaussian prior to the initial condition of the velocity field, $u_{0}$, that the posterior is

Gaussian and can be calculated explicitly. This means that in fact it is not necessary to estimate this distribution through implementation of MCMC methods. However, it is an excellent test of our methods, and is used here as a simple test-bed. Once we have shown that these methods work on this simple case, we will be able to extend their use to non-linear and harder problems such as those presented in chapters 3 and 4 .

### 2.2 The Model: Stokes flow

We consider a much simplified version of the real world problem. In reality the flow of geophysical fluids is highly non-linear, and can be best described by a solution $u \in$ $\left\{L^{\infty}(0, T ; H) \cap L^{2}(0, T ; H) \cap \nabla \cdot u=0\right\}$ to the Navier-Stokes equations (where $\gamma=1$ below):

$$
\begin{array}{r}
\partial_{t} u-\nu \delta u+\gamma(u \cdot \nabla) u+\nabla p=f, \quad t \geq 0 \\
\nabla \cdot u=0, \quad t \geq 0 \\
u(x, 0)=u_{0}(x), \quad x \in \Omega, \tag{2.3}
\end{array}
$$

The domain of solution in applications is also often complex, whether it be the atmosphere we are modelling where we need to consider the shape of mountain ranges etc, or in ocean dynamics where we need to consider the shape of coastlines and the varying depth of the water.

Solving these types of problems also gives us severe theoretical problems, not least because a theory of global solutions to the Navier-Stokes equations does not exist in three dimensions. Since we wish to build up a sound theoretical framework on which to do data assimilation, and we want to be able to compute numerics quickly and efficiently to demonstrate how this frameworks and can be implemented, we will need to simplify these equations and the domain $\Omega$.

If $\gamma=0$, then the equations (2.1)-(2.3) are termed the Stokes equations. In [24], Eulerian data assimilation problems concerning the Navier-Stokes equations are tackled in detail, from a theoretical standpoint. In this thesis, however, we will consider only the Stokes equations. We do this simply because by removing the nonlinear term from the equations, they become far easier to solve numerically. A fully spectral numerical method can be invoked easily when dealing with a linear problem, meaning that solving the equations is as simple as running a Fast Fourier Transform (FFT) each time you wish to know the value of the solution at any point in time. Spectral methods have been widely used in fluid mechanics as an alternative method to finite difference methods and finite element methods[27, 42].

We also simplify the geometry of our domain significantly, picking it to be $\Omega=$ $(0,1)^{2} \subset \mathbb{R}^{2}$ with periodic boundary conditions. That is, we are dealing with the unit torus $\mathbb{T}^{2}$. These simplifications make analysing the forward model a great deal more straightforward, yet there is still plenty of complexity in formulating the problem.

For the purposes of this chapter, we will also be assuming that $f$ is known. We will consider the case where $f$ is unknown in chapter 4 . In this chapter, we will discuss how we can formulate an MCMC method to sample from a well defined posterior measure which marries the observed data with our prior knowledge and the model to infer upon only the initial condition $u_{0} \in H$.

Since we are making Eulerian observations, given an initial condition $u_{0}$ we merely need to evaluate the weak solution to the equations (2.1-2.3) (with $\gamma=0$ ) at a given discrete set of points over a discrete set of times. Given sufficient regularity of $u_{0}$ and $f$, we can find the unique strong solution for $u$ for $t>0$.

As shown in section 1.4, this problem is actually equivalent to solving the fol-
lowing ODE on $H$;

$$
\begin{align*}
\frac{d u}{d t}+A u & =\mathcal{P} f=g, \quad t>0  \tag{2.4}\\
u(0) & =u_{0} \in H, \tag{2.5}
\end{align*}
$$

where $\mathcal{P}$ is the Leray projector onto the divergence-free space $H$.

### 2.3 Observational Noise Model

We make the assumption that our Eulerian observations are given by

$$
y_{j, k}=u\left(x_{j}, t_{k}\right)+\xi_{j, k}, \quad k \in\{1 \ldots N\} \quad \xi_{j, k} \sim \mathcal{N}(0, \Sigma), \quad \text { iid }
$$

That is, Eulerian observations are simply point-wise evaluations of the solution to the Stokes equations. We have assumed here that our observations are noisy, and that the noise is additive and mean zero Gaussian, with known covariance. We will adopt a vector notation,

$$
y=\mathcal{G}_{E}\left(u_{0}\right)+\xi \in \mathbb{R}^{2 J K}
$$

where $\mathcal{G}_{E}: H \rightarrow \mathbb{R}^{2 J K}$ is the Eulerian observation operator. $\mathcal{G}_{E}$ takes as its argument the initial condition $u_{0} \in H$ and returns the values of the weak solution of the Stokes equations at the observation points in space and time that are specified.

We wish to use these observations to gain information about the initial state of the vector field $u(x, 0)=u_{0}$. We will do this by constructing a probability distribution on function space for $u_{0}$, conditioned on the Eulerian observations. We can then sample from this by using Monte Carlo methods.

### 2.4 Prior Distribution of $u_{0}$

The prior distribution represents our prior knowledge of the velocity field, and our level of uncertainty in that knowledge. In the case where we have no specific knowledge about the initial condition prior to the data assimilation process, then the equilibrium distribution for the dynamical system is often used

We will employ a Gaussian prior on $u_{0}$, with covariance specified in terms of the stokes operator $A$. This choice is merely illustrative, but allows us to easily control regularity of the prior, which is necessary for analysis later on. To do this we must discover what the eigenfunctions and eigenvalues of the Stokes operator are, with a view to constructing samples from the prior using the Karhunen-Loeve expansion.

We wish to sample $u_{0}$ from a Gaussian measure $\mu_{0}$ on which we have defined the covariance operator $\mathcal{C}$. In this case we choose

$$
\mathcal{C}=\delta A^{-\alpha},
$$

where $\delta, \alpha>0$. Since $A$ is a smoothing operator, the larger the value of $\alpha$, the smoother the distribution $\mathcal{N}(0, \delta \mathcal{C})$ is. This is the reason for choosing such a prior distribution, as we now have complete control of the regularity of the samples that are drawn from it. This is important, as we will show later, for ensuring measurability of the likelihood with respect to the prior.

The eigenfunctions of $\mathcal{C}$ are $\phi_{k}$ given by (1.10), and the eigenvalues are $\lambda_{k}=a_{k}^{-\alpha}$ where the $a_{k}$ are given by (1.11). For further details on these and other matters relating to the Stokes' operator, see section 1.5. If we wish to sample $u_{0} \sim \mathcal{N}(0, \delta \mathcal{C})$, we use
the Karhunen-Loeve expansion, so that

$$
\begin{aligned}
u_{0} & =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} \sqrt{\delta \lambda_{k}} \phi_{k} \xi_{k}, \quad \xi_{k} \in \mathbb{C}, \quad \xi_{k} \sim \mathcal{N}(0,1) \\
& =\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} \frac{\sqrt{\delta}}{\left(4 \pi^{2}|k|^{2}\right)^{\alpha / 2}} \frac{k^{\perp}}{|k|} e^{2 \pi i k . x} \xi_{k}
\end{aligned}
$$

Definition 2.4.1. Consider a complex Gaussian random variable $\xi=X+i Y$, where $X$ and $Y$ are real and independent Gaussian variables with equal variances $\sigma^{2}$. The density of the joint variables is then

$$
\frac{1}{2 \pi \sigma^{2}} e^{-\left(x^{2}+y^{2}\right) /\left(2 \sigma^{2}\right)} .
$$

Then as $\operatorname{Var}(Z):=\mathbb{E}\left(|Z|^{2}\right)=2 \sigma^{2}$, where $|$.$| is the complex modulus, if we wish to$ have $Z \sim \mathcal{N}(0, I)$, we let $\sigma=\frac{1}{\sqrt{2}}$.

So now we are able to draw samples from $\mu_{0}$ for our initial condition on the vector field, $u_{0}$. Samples from $\mu_{0}$ can be made as rough or smooth as we like, by changing the parameter $\alpha$. In practise, when implementing this on a computer, it is also necessary to truncate the number of Fourier modes.

### 2.5 The Posterior Distribution

Before we can choose appropriate parameters for the prior distribution, we must analyse the forward problem, and more specifically the observation operator, to ascertain what regularity the prior distribution must have to ensure that the posterior measure $\mu$ is $\mu_{0}$ measurable so that theorem 1.10 .2 holds. One way to ensure that $\mathcal{G}_{E}$ is $\mu_{0}$-measurable is to show that it is continuous on a set $X$ which has full measure with respect to $\mu_{0}$, as shown in corollary 1.10.1. We can do this by bounding the operator in certain ways, and then applying Sobolev embedding theorems.

In chapter 4 we will be considering an Eulerian observation operator $\mathcal{G}_{E N}$ which is a function of not only the initial condition but also the time dependent forcing $\eta$ which is assumed to be unknown. Since the estimates require the same calculation, we define here the observation operator $G_{E}(\cdot, \cdot)$ which is a function of the initial condition and the time dependent forcing $f$ which drives the Stokes equations. We will then, in this chapter, set $\mathcal{G}_{E}\left(u_{0}\right)=G_{E}\left(u_{0}, f=\right.$ known $)$.

### 2.5.1 Bounds on $\mathcal{G}_{E}$

Lemma 2.5.1. Assume that $u_{0} \in H$ and that $f \in L^{2}\left(0, T, H^{r}\right)$ for some $r>0$. There exists a constant $C$ independent of $u_{0}$ and $f$ such that,

$$
\left|G_{E}\left(u_{0}, f\right)\right| \leq C\left(\left\|u_{0}\right\|+\|f\|_{L^{2}\left(0, T ; H^{r}\right)}\right) .
$$

Proof. Result follows from lemma 1.6.2.

Lemma 2.5.2. . Suppose $u_{0}, v_{0} \in H$, and $f, g \in L^{2}(0, T ; H)$. Then

$$
\left|G_{E}\left(u_{0}, f\right)-G_{E}\left(v_{0}, g\right)\right| \leq\left\|u_{0}-v_{0}\right\|+C\|f-g\|_{L^{2}\left(0, T ; H^{r}\right)}
$$

Proof. Proof follows from linearity of $G_{E}$ and application of lemma 2.5.1.

$$
\begin{aligned}
\left|G_{E}\left(u_{0}, f\right)-G_{E}\left(v_{0}, f\right)\right| & =\left|G_{E}\left(u_{0}-v_{0}, 0\right)\right| \\
& \leq\left\|u_{0}-v_{0}\right\|
\end{aligned}
$$

We can now use these bounds on $G_{E}$ to calculate equivalent bounds of $\mathcal{G}_{E}$, where $f$ is known.

Corollary 2.5.1. Assume that $u_{0} \in H$ and that $f \in L^{2}\left(0, T, H^{r}\right)$ for some $r>0$.
There exists a constant $C$ independent of $u_{0}$ and $f$ such that,

$$
\left|\mathcal{G}_{E}\left(u_{0}\right)\right| \leq C\left\|u_{0}\right\| .
$$

Proof. Follows from 2.5.1 with $f$ known.

Corollary 2.5.2. . Suppose $u_{0}, v_{0} \in H$, and $f \in L^{2}(0, T ; H)$. Then

$$
\left|\mathcal{G}_{E}\left(u_{0}\right)-\mathcal{G}_{E}\left(v_{0}\right)\right| \leq\left\|u_{0}-v_{0}\right\| .
$$

Proof. Follows from 2.5 .2 with $f=g$ known.

Thus, we have a space $X=H$ on which $\mathcal{G}_{E}$ is continuous, and therefore measurable with respect to any measure on which $X$ has full measure. We can now combine this with lemma 1.8 .1 to give the following result.

Corollary 2.5.3. Let $\mu_{0}=\mathcal{N}\left(0, \delta A^{-\alpha}\right)$ for $\alpha>1$. Then $\mathcal{G}_{E}$ is measurable with respect to $\mu_{0}$, and the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by (1.16).

Proof. Result follows by corollary 1.10.1, lemma 1.8.1, and corollaries 2.5.1 and 2.5.2.

### 2.6 The Random Walk Metropolis Hastings Algorithm

For the numerics that follow in this chapter, we will consider only the preconditioned RWMH algorithm, which we described in subsection 1.12.2. That is, given a currently accepted state in the chain $u_{n}$, we propose a new state $v$ where

$$
\begin{equation*}
v=\left(1-\beta^{2}\right)^{1 / 2} u_{n}+\beta w, \quad w \sim \mu_{0} \tag{2.6}
\end{equation*}
$$

where $\mu_{0}$ is the prior measure that we have prescribed on the initial condition $u_{0}$. Recall that the acceptance probability in this algorithm is given by

$$
\alpha(u, v)=\exp (\Phi(u)-\Phi(v))
$$

where $\Phi(\cdot)=\frac{1}{2}\left\|\mathcal{G}_{E}(\cdot)-y\right\|_{\Sigma}^{2}$.

### 2.7 Explicit Solutions to Eulerian Data Assimilation

We now have all the tools we need to begin sampling from approximations of welldefined posterior distributions, in the case of Eulerian data assimilation. However, since the problem is a linear one, and the prior we have chosen is Gaussian, our posterior distribution is itself a Gaussian, with mean and covariance that are calculable explicitly. In this section we will calculate these, and using this knowledge, analyse certain properties of these distributions.

Let us truncate our function space so that it is spanned by only two basis functions, for example $\phi_{0,1}$ and $\phi_{2,2}$. This calculation can also be carried out for any number of basis functions, but for computational ease we shall keep this number small for now. The solution $u$ to (2.4-2.5) with $f \equiv 0$ is given by

$$
u(x, t)=a_{1} e^{-\lambda_{1,0} t} \phi_{0,1}+a_{2} e^{-\lambda_{2,2} t} \phi_{2,2} .
$$

Let us define some notation for this special case. Let $x=\binom{a_{1}}{a_{2}}$. Let $B \in \mathbb{R}^{2 N \times 2}$, such that $B x=\left(\begin{array}{c}u\left(x_{1}, t^{*}\right) \\ \vdots \\ u\left(x_{N}, t^{*}\right)\end{array}\right)$, where the $x_{i}$ are the positions at which we are making the observations, and $t^{*}$ is the observation time.

If we let the prior distribution have density $\pi_{0}(u)=\mathcal{N}\left(m_{0}, \Sigma_{0}\right)$, our posterior distribution is given by

$$
\pi(x) \propto \exp \left(-\frac{1}{2 \sigma^{2}}|y-B x|^{2}-\frac{1}{2}\left|\Sigma_{0}^{-1 / 2}\left(x-m_{0}\right)\right|^{2}\right)
$$

This distribution is also Gaussian, so we should be able to find $m \in \mathbb{R}^{2}$ and $\Sigma \in \mathbb{R}^{2 \times 2}$ such that $\pi \sim \mathcal{N}(m, \Sigma)$. Let us consider the following.

$$
\begin{aligned}
-2 \log (\pi)= & \frac{1}{\sigma^{2}}|y-B x|^{2}+\left|\Sigma_{0}^{-1 / 2}\left(x-m_{0}\right)\right|^{2} \\
= & \frac{1}{\sigma^{2}}\langle y-B x, y-B x\rangle+\left\langle\Sigma_{0}^{-1 / 2}\left(x-m_{0}\right), \Sigma_{0}^{-1 / 2}\left(x-m_{0}\right)\right\rangle \\
= & \frac{1}{\sigma^{2}}\langle y, y\rangle+\frac{1}{\sigma^{2}}\langle B x, B x\rangle-\frac{2}{\sigma^{2}}\langle y, B x\rangle+\left\langle\Sigma_{0}^{-1 / 2} x, \Sigma_{0}^{-1 / 2} x\right\rangle \\
& +\left\langle\Sigma_{0}^{-1 / 2} m_{0}, \Sigma_{0}^{-1 / 2} m_{0}\right\rangle-2\left\langle\Sigma_{0}^{-1 / 2} x, \Sigma_{0}^{-1 / 2} m_{0}\right\rangle \\
= & \left\langle x,\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right) x\right\rangle-2\left\langle x, \Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}} y\right\rangle+\rho
\end{aligned}
$$

where $\rho$ is independent of $x$. Similarly,

$$
\begin{aligned}
-2 \log (\pi) & =\left|\Sigma^{-1 / 2}(x-m)\right|^{2} \\
& =\left\langle x, \Sigma^{-1} x\right\rangle-2\left\langle x, \Sigma^{-1} m\right\rangle+\left\langle m, \Sigma^{-1} m\right\rangle
\end{aligned}
$$

Therefore, by comparing coefficients, we can conclude that

$$
\begin{aligned}
\Sigma^{-1} & =\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1} \\
m & =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}} y\right)
\end{aligned}
$$

Suppose we wish to calculate the Kalman "gain matrix". That is the matrix $K$ such that

$$
m=m_{0}+K\left(y-B m_{0}\right)
$$

To calculate this we will need a particular identity. First note that

$$
\frac{B^{*}}{\sigma^{2}}\left(\sigma^{2}+B \Sigma_{0} B^{*}\right)=\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right) \Sigma_{0} B^{*}
$$

since $\Sigma_{0}$ is positive definite. Since $\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{1}$ and $\sigma^{2}+B \Sigma_{0} B^{*}$ are both positive definite as well, we can conclude that

$$
\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1} \frac{B^{*}}{\sigma^{2}}=\Sigma_{0} B^{*}\left(\sigma^{2}+B \Sigma_{0} B^{*}\right)^{-1}
$$

Consider the following,

$$
\begin{aligned}
m & =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}} y\right) \\
& =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*} B}{\sigma^{2}} m_{0}+\frac{B^{*}}{\sigma^{2}} y-\frac{B^{*} B}{\sigma^{2}} m_{0}\right) \\
& =m_{0}+\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1} \frac{B^{*}}{\sigma^{2}}\left(y-B m_{0}\right) \\
& =m_{0}+\Sigma_{0} B^{*}\left(\sigma^{2}+B \Sigma_{0} B^{*}\right)^{-1}\left(y-B m_{0}\right)
\end{aligned}
$$

Therefore the Kalman gain matrix is given by

$$
K=\Sigma_{0} B^{*}\left(\sigma^{2}+B \Sigma_{0} B^{*}\right)^{-1}
$$

Similarly, for the covariance matrix,

$$
\begin{aligned}
\Sigma & =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1} \\
& =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\frac{B^{*} B}{\sigma^{2}} \Sigma_{0}+\Sigma_{0}^{-1} \Sigma_{0}-\frac{B^{*} B}{\sigma^{2}} \Sigma_{0}\right) \\
& =\Sigma_{0}-\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1} \frac{B^{*} B}{\sigma^{2}} \Sigma_{0} \\
& =\Sigma_{0}-\Sigma_{0} B^{*}\left(\sigma^{2}+B \Sigma_{0} B^{*}\right)^{-1} B \Sigma_{0} \\
& =\Sigma_{0}-K B \Sigma_{0}
\end{aligned}
$$

This Kalman gain matrix $K$ could be used in a Kalman filter. We will briefly look at filtering problems in chapter 5 .

### 2.7.1 Properties of the Eulerian Analytic Posterior

Now that we have a way of exploring Eulerian posteriors exactly, or at least very simple ones, we can change certain parameters, and the quantity of observations to see how this affects the posterior distribution.

For example, let us consider what happens in the two limits, $\sigma \rightarrow 0$ and $\sigma \rightarrow$ $\infty$. As an example we choose 5 randomly placed observation stations, with a single observation being made at each station at time $T=0.1$, with only $u_{0,1}$ and $u_{2,2}$ non-zero. Specifically we set $u_{0,1}=\sqrt{2}\left(a_{1}-i b_{1}\right)$ and $u_{2,2}=\sqrt{2}\left(a_{2}-i b_{2}\right)$, with $a_{1}=1, b_{1}=-0.7, a_{2}=-0.2$, and $b_{2}=0.1$. We then look at the analytical posterior distribution with varying $\sigma$.


Figure 2.1: The mean values of $\left\{a_{i}, b_{i}\right\}$ in the posterior distribution with varying. $\sigma$

Figure 2.1 shows that as $\sigma \rightarrow 0$, the mean values of the $\left\{a_{i}, b_{i}\right\}$ converge to a value close but not exactly the same as the actual values that created the field that the observations were taken from. They are converging, in fact to the least squares solution
of the functional $J$ where

$$
J(u)=\frac{1}{2}\|\mathcal{G}(u)-y\|_{\Sigma}^{2}
$$

It also shows that as $\sigma \rightarrow \infty$, the mean values of the $\left\{a_{i}, b_{i}\right\}$ converge to their mean values in the prior distribution, which in this case are all zeros.


Figure 2.2: Difference in the mean of the posterior from the least squares solution

Figure 2.2 shows the difference in the mean of the posterior from the least squares solution. That is, specifically, the solution of

$$
\min _{x}\left\{\left\|\mathcal{G}_{E}(x)-y\right\|^{2}\right\} .
$$

As $\sigma$ tends to infinity, as shown in Figure 2.1, the mean of the posterior tends to the zero vector. Hence in 2.2 , as $\sigma$ gets large, the relative difference with the least squares solution tends to 1 . As $\sigma$ tends to zero however, the mean converges to the least squares solution.


Figure 2.3: The diagonal dominance of the covariance matrix in the posterior distribution with varying $\sigma$

Figure 2.3 shows how the diagonal dominance of the covariance matrix of the posterior distribution changes as we vary $\sigma$. We quantify this by comparing $\Sigma$ with $D=\operatorname{diag}(\Sigma)$, where

$$
\operatorname{diag}(\Sigma)_{i, j}= \begin{cases}\Sigma_{i, j} & \text { if } i=j \\ 0 & \text { if } i \neq j\end{cases}
$$

Then we can find out how diagonally dominant the matrix is relatively by calculating

$$
\frac{\|\Sigma-D\|_{F}}{\|D\|_{F}}
$$

where $\|A\|_{F}=\sqrt{\operatorname{Tr}\left(\mathrm{A}^{\mathrm{T}} \mathrm{A}\right)}$ is the Frobenius norm.
Figure 2.4 shows how the posterior covariance converges to that of the prior as $\sigma \rightarrow \infty$.

Now let us consider what happens as we increase the number of observations, and let $\sigma$ be fixed. The observations are made at points on a grid of varying refinement.


Figure 2.4: Convergence of the posterior covariance to the covariance of the prior distribution as $\sigma \rightarrow \infty$

Figure 2.5 shows the mean values of the $\left\{a_{i}, b_{i}\right\}$ in the posterior distribution with increasing number of observations, where the dashed lines represent the actual values that created the field from which we made the observations. As the number of observations increases, the mean values become closer to the actual values of the coefficients present in the field.

Figure 2.6 shows how the Frobenius norm of the posterior distribution's covariance matrix decreases exponentially with the number of observations made. Figures 2.5 and 2.6 together show how as we increase the number of observations, our posterior measure is converging to a Dirac measure on the field of which we made our observations.

Further to this, let us consider a set of observations made at a single time, made at a random set of points that are increasing in number. Figure 2.7 shows that as the number of observations increase, our mean is converging to the actual Fourier modes


Figure 2.5: Mean values of $\left\{a_{i}, b_{i}\right\}$ in the posterior as number of observations is increased


Figure 2.6: Frobenius norm of the posterior covariance with varying number of observations
present in the field at time $t=0$. Figure 2.8 shows that as the number of observations increases, the posterior is converging to a Dirac measure.


Figure 2.7: Relative difference with the actual Fourier modes of the norm of the posterior mean with varying number of observations

Now we fix the observation on a 5 by 5 grid, and look at what happens at posterior as we increase the number of observations times. Similarly, Figures 2.9 and 2.10 show that as we increase the number of observation times, our posterior measure is converging to a Dirac measure on the actual underlying vector field that we wish to understand.

It is worth noticing that the graphs in Figures 2.7 and 2.9 are not linear and seem to be affected by the noise in their observations. We can eliminate this by considering the analytic form of $m$, given that $y=B X+\xi$, where $X$ denotes the true Fourier components in the initial condition of the vector field.


Figure 2.8: Frobenius norm of the posterior covariance with varying number of observations


Figure 2.9: Relative difference with the actual Fourier modes of the norm of the posterior mean with varying number of observation times


Figure 2.10: Frobenius norm of the posterior covariance with varying number of observation times

$$
\begin{align*}
m & =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*} B}{\sigma^{2}} X+\frac{B^{*} \xi}{\sigma^{2}}\right)  \tag{2.7}\\
\Rightarrow \mathbb{E}(m) & =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*} B}{\sigma^{2}} X\right) . \tag{2.8}
\end{align*}
$$

The expectation of the mean of the posterior distribution, as given above, is no longer dependant on the observations themselves, and should give us a cleaner result. Figures 2.11 and 2.12 show how the expectation of the posterior mean converges to the vector $X$ as we increase the number of observations in time and space respectively.

Note that in Figures 2.10 and 2.12 there is a nonlinear effect for small numbers of observation times. This is simply due to the fact that as we are splitting the time interval $[0,1]$ into $n$ equal parts, at first all of the observation times are greater than $t=0.1$. At times greater than this, we do not gain much information at all about the higher frequency Fourier modes (we shall demonstrate this later in the chapter), and so therefore our error in the mean is higher, as is the uncertainty, which is reflected in a


Figure 2.11: Relative difference with the actual Fourier modes of the norm of the expectation of the posterior mean with varying number of observations


Figure 2.12: Frobenius norm of the expectation of the posterior covariance with varying number of observation times
larger Frobenius norm of the covariance matrix.
As the number of observation times increases, we have many more observation times less than $t=0.1$, so we observe the higher frequency modes much better and so we settle into a steadier exponential decay.

### 2.7.2 Using the Analytical Eulerian Posterior to Assess Information Contained in Observations

We are now in a position to investigate further into what information is actually contained in Eulerian data. We will continue to use the initial vector field as described before, with $x=\left(\begin{array}{llll}1 & -0.7 & -0.2 & 0.1\end{array}\right)^{\prime}$.

We consider 25 observations, made on a grid, at a variable time $t=T$. We can then look at the analytic posterior distributions for each time, and compare how much data we have about the Fourier coefficients of the initial condition of the vector field. Figure 2.13(a) shows the values of the Fourier coefficients in the mean of the posterior, with 25 observations made at varying time $T$. This graph shows that after quite a short time, the information relating to the higher frequency Fourier modes quickly disappears and reverts to the prior. Figure 2.13(b) shows that on a larger time scale the information about the low frequencies also decays. These are things that we will explicitly proved later in the thesis.

Let $\Sigma_{(i, j)}$ be the truncated matrix, containing only the entries such that $\Sigma_{(i, j)}=$ $\left(\begin{array}{cc}\Sigma_{i, i} & \Sigma_{i, j} \\ \Sigma_{j, i} & \Sigma_{j, j}\end{array}\right)$. The Frobenius norm of $\Sigma_{(i, j)}$ is an indicator of how much information we have about the values of the $i^{\text {th }}$ and $j^{\text {th }}$ Fourier coefficients. The smaller this norm, the closer our distribution is to a Dirac on these Fourier coefficients, and therefore the more information we have about them.


Figure 2.13: Values of the Fourier coefficients in mean of the posterior, with 25 observations made at time T, 2 different timescales


Figure 2.14: Frobenius norm of $\Sigma_{(3,4)}$ with 25 observations made at time $T$

Figure 2.14 shows how the information about the higher frequency Fourier modes decays as the observation time increase. Notice that the value of the norm hits the ceiling of the value of $\left(\Sigma_{0}\right)_{(3,4)}$ as the information vanishes and $\Sigma_{(3,4)}$ has reverted to the prior distribution. An identical situation is present for $\Sigma_{(1,2)}$, but on a larger timescale, but I have omitted it here.

The value of $\frac{\left\|\Sigma_{(3,4)}\right\|_{F}}{\left\|\Sigma_{(1,2)}\right\|_{F}}$ is an indicator of how much information is available in the posterior about $m(1)$ and $m(2)$ in comparison with $m(3)$ and $m(4)$. Figure 2.15 presents an interesting phenomena. Since the information in Eulerian data is encoded in how different Fourier modes decay at different rates, observations at small times tell you more about the value of the high frequency modes (which are decaying quickly) than the low frequency modes. As these high frequencies are damped out and the effect of them on the solution becomes negligible, the relative amount of information about the low frequency Fourier modes increases, until the damping effect damps these modes out to negligible levels as well.

### 2.8 General Setup

We now consider numerical results in situations where the true posterior is not so easily available. We describe here the general setup for all of the numerics that follow. We will engage in simulation studies in which the data is itself produced by employing the numerical simulation of a forward PDE model. However it is important in what follows to make a distinction between the algorithm that created the data which will be used by the data assimilation algorithm (MCMC method), and the model that is used within the assimilation algorithm itself. In particular we do not necessarily assume that our model is perfect; there may be some mismatch between the forward model used for data generation and for data assimilation. These issues will be covered in subsection 2.11.


Figure 2.15: Relative amount of information available in $m(1)$ and $m(2)$ compared to $m(3)$ and $m(4)$ in posterior with variable observation time

However, in all the numerics described in subsection 2.10, the data was created using the same resolution and noise environment as the numerical model used by the data assimilation algorithm itself.

For now we describe the set-up used in subsection 2.10 and for the experiments in section 2.12; slight variations apply in subsection 2.11. In each example, the same initial velocity field is used for creation of the data, which is itself a draw from the prior distribution $\mathcal{N}\left(0, \delta A^{-\alpha}\right)$, with $\delta=400$ and $\alpha=2$. Note also that this choice ensures that the conditions on the prior for the initial condition required in Theorem 2.5.3 hold. It is assumed that the covariance matrix for the observational noise is diagonal, and equal to $\sigma^{2} I$, with $\sigma=0.01$.

The approximation of the vector fields are truncated to 100 Fourier modes ${ }^{1}$ in

[^0]both the creation of the data, and for the assimilation algorithm. For the purposes of this chapter the forcing in the system is assumed to be $f \equiv 0$ in the data, and in the assimilation model. For problems with model error, where $f$ is unknown, see chapter 4.

In each experiment, the positions of the observation stations are given on a grid. Observation times are evenly spaced, with the final observation time given by $T=1$.

### 2.9 Limitations of the Standard RWMH method

In this section we aim to show the advantages of the RWMH posed on function space as opposed to the SRWMH method which works on a discretized space.

These methods were outlined in sections 1.12 .1 and 1.12.2. The key thing to notice here is the differences in their respective acceptance probabilities. The method framed on function space has proposal given by

$$
v=\left(1-\beta^{2}\right)^{1 / 2} u+\beta w, \quad w \sim \mathcal{N}(0, C)
$$

where $u$ is the currently accepted state in the chain. The acceptance probability is then given by

$$
a(u, v)=\min \{1, \exp (\Phi(u)=\Phi(v))\}
$$

where $\Phi=\frac{1}{2}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$. This acceptance probability does not contain any dependence on the dimension of the state space, which in practical terms will have to be finite

On the other hand, the SRWMH has the simpler proposal given by

$$
v=u+\beta w, \quad w \sim \mathcal{N}(0, C)
$$

The acceptance probability is given by

$$
\alpha(u, v)=\min \left\{1, \exp \left(\Phi(u)-\Phi(v)+\frac{1}{2}\left|C^{1 / 2} u\right|^{2}-\frac{1}{2}\left|C^{1 / 2} v\right|^{2}\right\}\right.
$$

the number of grid points used in approximation of the velocity field
which in function space would be infinite. The reasoning behind this is simple. The state $u$ is a linear combination of draws from the prior. Any draw from the prior is given by

$$
w=C^{-1 / 2} \xi, \quad \xi \sim \mathcal{N}(0, I) .
$$

## Therefore

$$
\begin{aligned}
\mathbb{E} \frac{1}{2}\left|C^{1 / 2} w\right|^{2} & =\frac{1}{2} \mathbb{E}|\xi|^{2} \\
& =\frac{1}{2} \sum_{k} \mathbb{E}\left|\xi_{k}\right|^{2}
\end{aligned}
$$

Since each $\mathbb{E}\left|\xi_{k}\right|^{2}=1$ by definition, this sum is infinite. In practical terms, this is not the case since we truncate the Fourier expansion to some number $N$. However, this does mean that as the mesh is refined and more non-zero terms are added to the Fourier expansion, this term in the acceptance probability increases, and is in fact $\mathcal{O}(N)$. In turn, this reduces the average acceptance probabilities for a given fixed step size $\beta$.

In the following graphs we aim to show this happening in practise. In both cases, the same data set was taken which consists of 9 Eulerian observations stations on a grid with one observation time $T=0.1$. A draw from the prior was taken as the initial condition in the data creation process. Markov chains were started with a range of values of $\beta$ for both the RWMH and SRWMH method. The average acceptance probabilities were taken until they had converged sufficiently. This process was repeated for a range different grid sizes, with $\Delta x \in\left\{\frac{1}{10}, \frac{1}{20}, \frac{1}{50}, \frac{1}{100}, \frac{1}{200}, \frac{1}{500}\right\}$. The results are presented in Figures 2.16 and 2.17.

Figure 2.16 shows clearly the degeneration of the algorithm as the grid is refined. As the Fourier truncation has more non-zero terms, the average acceptance probability for a given value of $\beta$ reduces. This means that as we increase the number of Fourier modes, we are unable to make such large steps in state space. This means that we


Figure 2.16: Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, SRWMH
explore the state space slower and the method becomes increasingly inefficient.
In contrast, Figure 2.17 shows the same plots for the RWMH method framed on function space, whose acceptance probabilities have no dependence on the dimension of the discretisation. This independence is borne out in the results which clearly show that as the mesh is refined, the average acceptance probabilities stay exactly the same for a given step size $\beta$, allowing us to make larger steps in state space in comparison to the SRMH method.

These results are an indictment of the SRWMH algorithm, and show the advantages of formulating methods on function space before discretising. As a result, the RWMH posed on function space is the algorithm that we will use throughout this chapter, and indeed through later chapters in this thesis.


Figure 2.17: Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, RWMH

### 2.10 Validating the Random Walk Algorithm

In this section we aim to verify numerically that the MCMC algorithm that we have described (i) converges to an exact analytical solution; (ii) converges from different initial states to the same solution; (iii) converges for different $\beta$ to the same solution. We can use the fact that we know exactly the form of the posterior distribution in the Eulerian case to verify that our MCMC algorithm is drawing samples correctly.

Consider the case where we have 25 observation points on a 5 by 5 grid, and where we have once again truncated the possible velocity fields to those with two nonzero complex Fourier modes, with frequencies $k_{1}=(0,1)$ and $k_{2}=(2,2)$. We let the observation points be on the grid points of our discretized domain so that we can eliminate any error from interpolation. We observe the vector field at all of these points at the time $t=0.1$, with initial velocity field equal to

$$
\begin{aligned}
u(x, 0) & =\sum_{k} a_{k} \frac{k^{\perp}}{|k|} \sqrt{2} \sin (2 \pi i k \cdot x)+b_{k} \frac{k^{\perp}}{|k|} \sqrt{2} \cos (2 \pi i k \cdot x) \\
& =\sum_{k} \sqrt{2}\left(b_{k}-i a_{k}\right) \frac{k^{\perp}}{|k|} e^{2 \pi i k \cdot x}
\end{aligned}
$$

with $a_{0,1}=1, b_{0,1}=-0.7, a_{2,2}=-0.2, b_{2,2}=0.1$ and all other values equal to zero.
Figure 2.18 shows the evolution of the error of the ergodic average in the Markov chain of the Fourier coefficients with respect to that of our analytic solution. Similarly 2.19 shows the evolution of the error of the covariance matrix of the Fourier coefficients.


Figure 2.18: Relative 2-norm error of the ergodic average of the vector $x=$ $\left(a_{1}, b_{1}, a_{2}, b_{2}\right)$

These graphs show us that our Markov chain seems to be converging in distribution to our analytic solution as we take more samples, in this situation. Let us now consider what happens to our numerics as we increase the number of observations that we make at one time, $T=0.1$, on a grid. Figures $2.20(\mathrm{a})-2.20(\mathrm{f})$ show a selection of


Figure 2.19: Relative Frobenius-norm error of the running covariance matrix $\left(C_{i j}\right)$ of the vector $x=\left(a_{1}, b_{1}, a_{2}, b_{2}\right)$
the results showing the numerics converging to our analytic distributions.
We should next check that if we introduce more observation times, that our code is still drawing samples from the corresponding analytic distribution. Figures 2.21 and 2.22 show an example with 32 observation times, which is converging to the correct distribution. All of these tests show that the code appears to be drawing samples from the correct distribution.

These figures verify numerically that as we increase the number of MCMC samples the empirical measure generated by the preconditioned random walk method with proposal (6.9) converges to the true Gaussian posterior distribution. Figure 2.23 shows how the value of one particular Fourier mode of the initial condition in the Markov chains with different starting states all converge to the same distribution, in the case of Eulerian data. A large amount of data is used so that we see how the chains quickly

(a) Relative 2-norm error of $m_{n}, 1$ obser- (b) Relative Frobenius-norm error of $\Sigma_{n}$, vation at $T=0.1$

1 observation at $T=0.1$


(c) Relative 2-norm error of $m_{n}, 25$ ob- (d)
(d) Relative Frobenius-norm error of $\Sigma_{n}$, servations at $T=0.1$

25 observations at $T=0.1$


(e) Relative 2-norm error of $m_{n}, 2500$ ob- (f) Relative Frobenius-norm error of $\Sigma_{n}$, servations at $T=0.1$

2500 observations at $T=0.1$
Figure 2.20: Convergence of numerics to analytic distributions


Figure 2.21: Relative error in the mean, 25 observations made at 32 times up to $T=1$


Figure 2.22: Relative error in the covariance matrix, 25 observations made at 32 times up to $\mathrm{T}=1$
converge to the area of high probability in the state space. The Markov chains appear to converge to a single value, but actually the posterior is just very concentrated around this value due to the high volume of data, leading to low levels of uncertainty.


Figure 2.23: Convergence of Markov chains with different initial states, with Eulerian data

Similarly, it has been verified numerically that Markov chains with different values of $\beta$ in the random walk proposal converge to the same distribution, although they do converge at different rates. Figure 2.24 shows the different rates of convergence of the algorithm with different values of $\beta . \beta_{\mathrm{opt}}$ here is the value of $\beta$ which gives an acceptance rate of approximately $50 \%$. With $\beta$ too small, the algorithm accepts proposed states often, but these changes in state are too small so the algorithm does not explore the state space efficiently. This can be seen in the diagram in the trace which is approximately a straight line with a small fixed negative gradient.

In contrast, with $\beta$ too big, larger jumps are possible. However, having large proposals often means that the algorithm is attempting to jump to a state that has


Figure 2.24: Convergence of Markov chains with different values of $\beta$
much smaller probability density, and so the algorithm will often reject proposals, and therefore will not explore the state space efficiently. This can be seen in the diagram in the trace which looks like a step function.

Our aim is to pick a $\beta$ which makes big enough steps to allow us to move quickly enough away from the present state, but also one which is small enough so that we don't reject proposals all of the time [65]. Such a choice can be seen in the diagram in the trace which moves quickly to a region around 0.6 , and then explores the measure which is centred around this value.

We may also wish to consider the possibility of having a random $\beta$, with a distribution which is chosen in such a way that we still have a decent rate of convergence. The theoretical justification for this is given in [23]. This has the potential advantage of including the possibility of large and small steps in the proposal. This would allow us to


Figure 2.25: Marginal distributions with and without random $\beta$
make occasional large jumps to other parts of the state space, and smaller moves that allow us to explore the immediate neighbourhood. In the following example, a value of $\beta$ was first found that yields approximately a $50 \%$ acceptance probability in the sampler run with a given set of data, which we will denote as $\beta_{\mathrm{opt}}$.

Two instances of the sampler were run with the same data, one with a static value of $\beta=\beta_{\mathrm{opt}}$, and one with $\beta \sim U\left(\left[0.1 \times \beta_{\mathrm{OPT}}, 1.9 \times \beta_{\mathrm{OPT}}\right]\right)$. The marginal distributions for both Markov chains are shown in Figure 2.25. The two computed distributions are very close indeed. In this case, where the posterior distribution is a very Gaussian-looking mono-modal distribution both of these methods converged in approximately the same amount of time. However, it is conceivable that if the posterior distribution had more than one mode, that the variable- $\beta$ method might be better suited and converge quicker, allowing transitions between modes


Figure 2.26: Distribution of the accepted $\beta$

Figure 2.26 shows the distribution of the $\beta$ for which the proposed state was accepted. As you would expect, the initial uniform distribution is skewed as proposals with smaller jumps are more likely to be accepted.

The numerical experiments in this section demonstrate that the implementation of the preconditioned random walk algorithm that we have used is successfully sampling from the posterior distribution for the data assimilation problems of interest. The amount of Monte Carlo steps that are required to give results of this quality varies depending on the form of the observation operator $\mathcal{G}$. For small numbers of observations with large amounts of uncertainty, where the posterior distribution is relatively close to the prior distribution, the chain will converge very quickly in a matter of $\mathcal{O}\left(10^{4}-10^{5}\right)$ steps. However, if there are a large amount of observations, sometimes up to $\mathcal{O}\left(10^{7}-10^{8}\right)$ iterations are required before the samples are representative of the
distribution we are interested in.

### 2.11 Inverse Crimes

When testing algorithms to tackle inverse problems, we must always be certain that we test it in a fair and objective way. Namely, we must not assume the model in our algorithm is a perfect match for the dynamics of the system that we are observing. It is potentially problematic to study only experiments where the data and assimilation algorithm use the same model and parameters. This is sometimes referred to as an inverse crime [45]. Here we further test the algorithms by using data created with different parameters, resolutions, observation noise models, and time step sizes, to ensure that we are not committing such a crime.

It is important to consider the effect of differing resolutions in our data creation algorithm and in the sampling algorithm. In reality, of course, the true dynamical system is infinite dimensional, so we consider the case where our data is created using a much higher resolution than we use for the sampling method. In all of the following figures (apart from those that specify a lower resolution) the data was created using a grid for the vector field of $1000 \times 1000$ points (or $5 \times 10^{5}$ complex Fourier modes). The data assimilation algorithm was run with varying numbers of grid points for the velocity field approximation, running through 16, 100, 196 and finally 400 points. Figure 2.27 shows that as the number of grid points used in the velocity field approximation for the algorithm is increased, the marginal distribution for this particular Fourier mode appears to converge to a limit. The noise levels and quantity of data are such that the posterior is not a peaked distribution on the true Fourier mode which was present in the initial condition that created the data. For more details on the convergence of the posterior as the mesh used in the forward problem is refined, see [25].


Figure 2.27: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, high resolution data, Eulerian data.

An in depth discussion of what can happen when there is a mismatch in the forcing functions for the dynamical system between the data environment and the algorithm for both Eulerian and Lagrangian data can be found in chapter 4, so we shall skip over this particular area in this section.

We now consider several different cases of a mismatch in the actual noise in our observations, and the noise model that we use in the likelihood function in the algorithm. First we consider the case of low variance noise in our observations, with $\Sigma=0.001 I$, where we use a larger variance of varying size in the likelihood function.

Figure 2.28 shows that as we increase the variance of $\Sigma$, the influence of the observations decreases, until we are sampling from a distribution that is very close to the marginal prior distribution for this Fourier mode. This makes perfect sense, as if we assume more uncertainty in our data, then we can draw less information from it, and have to rely on our prior beliefs.


Figure 2.28: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing variance in the noise model of the algorithm, low actual noise, Eulerian case

We may also consider the case where the observational noise in our data has much larger variance than the variance we use in our likelihood. Figure 2.29 shows what happens in this case. This scenario can cause a great deal of problems, as we are trying to infer accurately using poor data with a high signal to noise ratio (SNR). Essentially, this shows that the algorithm still works in this case, but the results may be poor.

We may also consider the case in which we use a very low resolution model for creating our data. This is not realistic in terms of data that is collected in the field, but gives us another opportunity to show how the algorithm copes with a mismatch between model and data.

Figure 2.30 shows the marginal distributions of one Fourier mode with Eulerian data created with a $4 \times 4$ grid, with a varying number of grid points used in the assimilation model. The resolution in the algorithm which gets closest to the correct answer is actually the lowest resolution, with 16 grid points on which the velocity field is


Figure 2.29: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Decreasing variance in the noise model of the algorithm, high variance actual noise, Eulerian case
approximated. This is not surprising given that this is the resolution at which the data was created.

So in conclusion, we have demonstrated numerically that if we increase the resolution of the approximation of the velocity field in our algorithm, that the posterior converges in distribution. We have also presented how various mismatches in data creation and model used in the algorithm can affect the results.

### 2.12 Properties of the Posterior Measure

Now that we have a successful implementation of an MCMC method that allows us to draw samples from well defined posterior distributions, we can analyse exactly what kind of information is present in different data scenarios, and what happens to the posterior in the limit of large amounts of informative data. In the following figures we display


Figure 2.30: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, low resolution data, Eulerian case
only the marginal distribution of the Fourier mode $\operatorname{Re}\left(u_{0,1}\right)$ for the initial condition of the vector field. Other low wave-number Fourier modes behave similarly to this mode. (However high wave number modes are not greatly influenced by the data, and remain close to their prior distribution). For simplicity, we also assume that there is no forcing in the system, setting $f=0$.

In each of the following Figures we demonstrate posterior consistency of the low wave-number modes: we show that increasing the amount of data in sensible ways leads to posterior measures which approach a Dirac measure on the true value of the Fourier mode that was present in the initial condition of the vector field that created the data.

In our first example, the observations are made at 100 evenly spaced times, on an evenly spaced grid with an increasing number of points. Figure 2.31 shows how the posterior distribution on $\operatorname{Re}\left(u_{0,1}\right)$ changes as we increase the number of points in space at which we make Eulerian observations of the velocity field. This plot shows that the


Figure 2.31: Increasing numbers of observations in space, Eulerian.
chains have converged in their marginal distribution to approximate Gaussian curves on $\operatorname{Re}\left(u_{0,1}\right)$. We also see that we have posterior consistency, since as the number of spatial observations increases, the posterior distribution on $\operatorname{Re}\left(u_{0,1}\right)$ appears to be converging to an increasingly peaked distribution on the value that was present in the true initial condition, denoted by the solid line.

Now, if we instead have a fixed number of 25 observations in space on a grid, made at an increasing number of times evenly spaced on the unit interval, we arrive at Figure 2.32. This plot shows posterior consistency for an increase in temporal observations, giving us a sequence of increasingly peaked approximate Gaussian curves centred closely to the value that we wish to recover.


Figure 2.32: Increasing numbers of observations in time, Eulerian.

### 2.12.1 Decay of Eulerian Data in Time

We are also interested in how much information about the entire state of the vector field is actually contained within the observation that we are making. The following experiments were made with observations from 9 paths/stations, with one observation made a time $T$. We would like to know how the information we get about the initial condition varies as we take Eulerian observations further into the future from the starting time. Figure 2.33 shows histogram plots for Eulerian Markov chains with varying observation time $T$.

Figure 2.33 shows that as the time at which we make our observations increases, the information we have about the initial condition of the velocity field eventually decays, after initially improving, mirroring the effect that we saw in Figure 2.15 in section 2.7.2. Before we can prove a theorem showing the (eventual) decay, we need the following lemma.


Figure 2.33: PDFs For Eulerian Data, 9 Stations, Varying T

Lemma 2.12.1. The function $f(x)=e^{-x} x^{s}$ with $s \geq 0$, is bounded above on $\mathbb{R}^{+}$, since

$$
f(x) \leq e^{-s} s^{s} \quad \forall x \geq 0 .
$$

Proof.

$$
\frac{d f}{d x}=e^{-x} x^{s-1}(s-x) .
$$

Therefore we have the following two cases.
(i) If $s \leq 1$, we have 1 critical point at $x=s$.
(ii) If $s>1$, we have 2 critical points, at $x=s, 0$.

Since $f(x) \geq 0$ on $\mathbb{R}^{+}$, is continuous, and $f(0)=f(\infty)=0$, the critical point $x=s$ must be a maximum. Therefore

$$
f(x) \leq e^{-s} s^{s} \quad \forall x \geq 0 .
$$

We can use this result to prove the following theorem.

Theorem 2.12.1. Given a weak solution $u \in H^{\alpha}$ to the Stokes equations (2.4-2.5), with $\alpha>1, f=0$, then $u$ is decaying exponentially.

Proof. Since $u \in H, \exists\left\{u_{k}\right\}$ s.t.

$$
u(x, t)=\sum_{k \in \mathbb{Z}^{2} \backslash\{0\}} \phi_{k}(x) u_{k} e^{-\lambda_{k} t}, \quad \lambda_{k}=4 \nu \pi^{2}|k|^{2} .
$$

By using part (iii) of theorem 1.8.1, and lemma 2.12.1, we have that

$$
\begin{aligned}
\|u\|_{L^{\infty}}^{2} & \leq C\|u\|_{H^{\alpha}}^{2} \\
& =C \sum_{k}\left|u_{k}\right|^{2} e^{-2 \lambda_{k} t}|k|^{2 \alpha} \\
& =C \sum_{k}\left|u_{k}\right|^{2} e^{-\lambda_{k} t} \frac{e^{-\left.4 \nu \pi^{2}|k|\right|^{2} t}\left(4 \nu \pi^{2}|k|^{2}\right)^{\alpha}}{\left(4 \nu \pi^{2} t\right)^{\alpha}} \\
& \leq C \sum_{k}\left|u_{k}\right|^{2} \frac{e^{-\lambda_{k} t}}{\left(4 \nu \pi^{2} t\right)^{\alpha}} \\
& <C \sum^{\left|u_{k}\right|^{2}} e^{-\lambda_{k} t} \quad \text { for } \quad t>\frac{1}{4 \nu \pi^{2}}, \\
& \leq C e^{-\lambda_{0,1} t} \sum_{k}\left|u_{k}\right|^{2} \quad \text { since } \quad \lambda_{0,1} \leq \lambda_{k} \forall k \in \mathbb{Z}^{2} \backslash\{0\} .
\end{aligned}
$$

Since we have an exponential bound on the maximum velocity in the field, we can see that the velocity field tends to zero everywhere. Therefore in finite time, the noise on our Eulerian observation will dominate the actual velocity field that is present, rendering the observation useless. So therefore the information in Eulerian observations eventually decays as the time of observation increases.

Suppose now we introduce a time-independent forcing to the system, with $f=$ $\sum_{k} f_{k} \phi_{k}$.


Figure 2.34: Sketch to show decay of $u(x, t)$ into a $\delta$-ball when $f \equiv 0$
Theorem 2.12.2. Given a weak solution $u \in H^{\alpha}$ to the Stokes equations (2.4-2.5), with $\alpha>1, f=\sum_{k} f_{k} \phi_{k} \in H^{\alpha-2}$, where $f_{k} \in \mathbb{R}$ are constants, then $\|u\|_{L^{\infty}}$ is bounded. In particular, for any $\delta>0, \exists T_{\delta}$ s.t. $\forall t>T_{\delta}$,

$$
\|u(x, t)\|_{L^{\infty}} \leq C\left[\left(4 \nu \pi^{2}\right)^{-2}\|f\|_{H^{\alpha-1}}+\delta\right] .
$$

Proof. The weak solution of this system is given by

$$
u(x, t)=\sum_{k}\left(\frac{1-e^{-\lambda_{k} t}}{\lambda_{k}} f_{k}+e^{-\lambda_{k} t} u_{k}(0)\right) \phi_{k} .
$$

Therefore, for $t>T_{\delta}=\frac{1}{4 \nu \pi^{2}} \max \left\{1, \ln \left(\frac{C}{\delta}\right)\right\}$,

$$
\begin{aligned}
\|u\|_{L^{\infty}} & \leq C\|u\|_{H^{\alpha}} \\
& =C \sum_{k}\left|\frac{1-e^{-\lambda_{k} t}}{\lambda_{k}} f_{k}+e^{-\lambda_{k} t} u_{k}(0)\right|^{2}|k|^{2 \alpha} \\
& \leq C\left(\sum_{k}\left|\frac{1}{\lambda_{k}} f_{k}\right|^{2}|k|^{2 \alpha}+\sum_{k} e^{-2 \lambda_{k} t}\left|u_{k}(0)\right|^{2}|k|^{2 \alpha}\right) \\
& =C\left(S_{1}+S_{2}\right) .
\end{aligned}
$$

By the proof of theorem 2.12.1, $S_{2}<C e^{-\lambda_{0,1} t}<\delta$.

$$
\begin{aligned}
S_{1} & =\left(4 \nu \pi^{2}\right)^{-2} \sum_{k}\left|f_{k}\right|^{2} \frac{|k|^{2 \alpha}}{|k|^{4}} \\
& =\left(4 \nu \pi^{2}\right)^{-2} \sum_{k}\left|f_{k}\right|^{2}|k|^{2(\alpha-2)} \\
& =\left(4 \nu \pi^{2}\right)^{-2}\|f\|_{H^{\alpha-2}} \in \mathbb{R}^{+} .
\end{aligned}
$$

We can generalise to systems which have differing initial conditions but which have the same (non-constant) forcing.

Theorem 2.12.3. Given two weak solutions $u$ and $v$ to the Stokes' flow problem on the torus with initial conditions $u_{0}, v_{0} \in H$ respectively, both forced by the same function $f=\sum f_{k}(t) \phi_{k} \in L^{2}([0, T], H)$, then

$$
\|u-v\|_{L^{\infty}} \rightarrow 0
$$

as $T \rightarrow \infty$.

Proof. The weak solution of these systems are given by

$$
\begin{aligned}
& u(x, t)=\sum_{k}\left(\frac{1-e^{-\lambda_{k} t}}{\lambda_{k}} f_{k}+e^{-\lambda_{k} t} u_{k}(0)\right) \phi_{k} \\
& v(x, t)=\sum_{k}\left(\frac{1-e^{-\lambda_{k} t}}{\lambda_{k}} f_{k}+e^{-\lambda_{k} t} v_{k}(0)\right) \phi_{k}
\end{aligned}
$$

Both of these solutions converge to the same solution as $T \rightarrow \infty$ since

$$
\begin{aligned}
\|u-v\|_{L^{\infty}}^{2} & =\left\|\sum_{k} e^{-\lambda_{k} t}\left(u_{k}(0)-v_{k}(0)\right) \phi_{k}\right\|_{L^{\infty}}^{2} \\
& \leq e^{-2 \lambda_{0,1} t} \sum_{k}\left|u_{k}(0)-v_{k}(0)\right|^{2} \\
& =e^{-2 \lambda_{0,1} t}\left\|u_{0}-v_{0}\right\|^{2} .
\end{aligned}
$$

Since $u_{0}, v_{0} \in H$, the result follows.

Since after a period of time, the two solutions are almost indistinguishable despite having different initial conditions, we can conclude that Eulerian information about the initial condition decays in time.

### 2.13 Conclusions and Future Directions

By careful analysis of the forward problem, we have been able to formulate a well-posed Bayesian inverse problem regarding Eulerian data of the Stokes' flow dynamical system. We have shown that the likelihood function is continuous with respect to a space that has full measure with respect to a specified choice of Gaussian prior measure. Using this, we have shown how to draw samples from well defined posterior distributions on function space using the RWMH MCMC sampler. We have then implemented this algorithm in C, and gained insight into what kind of information is available in Eulerian data. We have also shown how the standard RWMH method loses efficiency as the grid is refined, and that the method framed on function space does not. We have also verified the algorithm by checking against explicit posterior distributions that are calculable in certain situations, as well as showing that we are mindful of committing inverse crimes.

This chapter demonstrates a method of tackling inverse problems that we will be using in several different scenarios throughout the thesis. The core to the philosophy underlying these methods is the belief that formulating numerical methods on infinite dimensional spaces and only discretizing once we choose to implement such a method gives us better algorithms that are robust under different discretisations and refinements.

This subject could be extended in many directions, some of which will be addressed in later chapters. We could also look to implement this problem for the full Navier-Stokes equations, which would then capture the non-linear behaviour that is hugely important in real life applications of data assimilation in fluid mechanics. The
necessary theoretical results concerning framing data assimilation of the Navier-Stokes equations has already been addressed in [24]. Implementing this problem for real life applications would also involve formulating the problem on approximations of real domains, as opposed to the torus domain that we have picked for ease in this simple example.

The sampling method can also be improved. If an adjoint to the forward model were to be implemented, we would be able to calculate the gradient of the observation operator. This would allow us to include gradient information in the proposal distribution, as in, for example, the Metropolis adjusted Langevin algorithm (MALA) presented in section 1.12.2. Moreover, we could also replace the burn-in with a deterministic method to find an approximation to the state of highest probability density, via an equivalent Tikhonov regularisation to the smoothing that the prior exerts (see section 1.11). Many of the deterministic methods used for these variational problems also use the gradient to find the solution, for example the gradient descent or conjugate gradient methods.

Another alternative, which could reduce the cost of the problem, would be to replace the forward model with a gPC approximation, as discussed in section 1.1. This would in all probability work quite well for the size of problem we have considered in the majority of the numerics in this chapter, where the state space has 100 degrees of freedom. However, if we were to increase this dimension for practical applications, the costs would soon become unfeasible, as unlike the MCMC method we have presented here, gPC approximations suffer from the curse of dimensionality.

In the next chapter, we will consider a very similar scenario in data assimilation of data observed from a Stokes' flow system. This time, however, the observations will be Lagrangian in nature.

## Chapter 3

## Lagrangian Data Assimilation

### 3.1 Motivation

We now consider a similar problem with a different data type informing us about the state of the velocity field. Hundreds of GPS floaters (which float on the surface of the ocean) and drifters (which submerge to a given depth at which they are transported by the flow of the water) are currently distributed throughout the planet's oceans in an attempt to better understand these complex dynamical systems. Periodically transmitting their position, as well as other data concerning salinity and temperature of the water amongst other things, these Lagrangian tracers create huge banks of data.

Figure 3.1 shows an example of such a data set, in the form of a spaghetti diagram. We wish to infer, from this data, the entire state of the vector field in question, just as we did in the Eulerian case. The main difference between these two forms of data, is that in the Eulerian case, the observation operator is linear, as we are making direct observations of a linear system. In the case of Lagrangian dynamics, the observation operator is highly nonlinear. Moreover, the dynamics of the tracers themselves can


Figure 3.1: Spaghetti diagram of 20-day drifter trajectory segments. Colours give the mean drift direction (legend in upper-right corner). Taken from [53]
display chaotic behaviour[7]. Herein we will explore the differences between these two observation operators, and then using the same approach as in the previous chapter, we will describe Monte Carlo methods with which we can sample from well defined posterior distributions which give us information about the flow of the dynamical system.

### 3.2 The Model

The evolution of the vector field in the Lagrangian case mirrors that of the Eulerian case. The only difference is the way in which we observe this vector field, this time indirectly through the positions of a set of tracers in the flow. Once we have a weak solution $u$ to the Stokes problem, we can follow the paths of $J$ particles $\left\{z_{j}\right\}_{j=1}^{J}$ with initial positions at time $t=0$ given by $\left\{x_{j}\right\}_{j=1}^{J}$, by solving the following set of ordinary differential equations,

$$
\begin{align*}
\frac{d z_{j}(t)}{d t} & =u\left(z_{j}(t), t\right)  \tag{3.1}\\
z_{j}(0) & =x_{j} . \tag{3.2}
\end{align*}
$$

We then observe the particles at $K$ times, $\left\{t_{k}\right\}_{k=1}^{K}$. We then define the Lagrangian observation operator $\mathcal{G}_{L}$ to be defined as follows;

$$
\mathcal{G}_{L}\left(u_{0}, f\right)=\left\{z_{j}\left(t_{k}\right)\right\}_{j, k=1}^{J, K} .
$$

For $\mathcal{G}_{L}$ to be well defined we need the paths of these particles to exist and be unique. We refer to the main result of [26].

Theorem 3.2.1. Let $\Omega$ be an open bounded subset of $\mathbb{R}^{d}, d=2,3$, with a sufficiently smooth boundary. Consider $u \in L^{\infty}\left(0, T ; H^{(d / 2)-1}(\Omega)\right) \bigcap L^{2}\left(0, T ; H^{d / 2}(\Omega)\right)$, a unique solution of the Navier-Stokes equations with $u_{0} \in H^{(d / 2)-1}(\Omega)$ and $f \in$ $L^{2}\left(0, T ; H^{(d / 2)-1}(\Omega)\right)$. Then the ordinary differential system

$$
X(t)=\int_{0}^{t} u(X(t), t) d t+X(0), \quad X(0)=a,
$$

has a unique solution $X \in C\left([0, t], \mathbb{R}^{d}\right.$.

This theorem can trivially be adapted to uniqueness of paths in Stokes flow, and shows that our operator $\mathcal{G}_{L}$ is indeed well defined, providing we have sufficient regularity of $u_{0}$ and $f$. We can now use this to define a likelihood function on $u_{0}$, conditional on the noisy observations $y$, equivalent to Eulerian example in the previous chapter. Once again, we assume that our observations are noisy, and that $f=0$ is known, with

$$
y=\mathcal{G}_{L}\left(u_{0}\right)+\xi, \quad \xi \sim \mathcal{N}(0, \Sigma),
$$

for some known covariance matrix $\Sigma$. Then the likelihood function is given by

$$
\mathbb{P}\left(y \mid u_{0}, f\right) \propto \exp \left(-\frac{1}{2}\left\|\mathcal{G}_{L}\left(u_{0}\right)-y\right\|_{\Sigma}^{2} .\right.
$$

### 3.3 The Posterior Distribution

As in the Eulerian case, for the desired posterior distribution to be well defined, we require the observation operator to be continuous on a set with full measure with respect to the prior measure. To choose an appropriate prior measure, we must first consider results concerning bounds on this operator, given in the next section. Our aim is to find an appropriate prior measure $\mu_{0}$ such that the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by

$$
\begin{equation*}
\frac{d \mu}{d \mu_{0}}=\exp \left(-\frac{1}{2}\left\|\mathcal{G}_{L}\left(u_{0}\right)-y\right\|_{\Sigma}^{2}\right) \tag{3.3}
\end{equation*}
$$

Analogously to the previous chapter on Eulerian data assimilation, we now invoke Bayes' theorem, giving us the posterior distribution we desire to explore. However, we must once again place certain conditions on our choice of prior distribution to ensure that the posterior distribution is absolutely continuous with respect to this choice of prior. As we did before, we look to the properties of the forward problem, namely of the observation operator $\mathcal{G}_{L}$.

In chapter 4 we will be considering an Lagrangian observation operator $\mathcal{G}_{L N}$ which is a function of not only the initial condition but also the time dependent forcing $\eta$ which is assumed to be unknown. Since the estimates require the same calculation, we define here the observation operator $G_{L}(\cdot, \cdot)$ which is a function of the initial condition and the time dependent forcing $f$ which drives the Stokes equations. We will then, in this chapter, set $\mathcal{G}_{L}\left(u_{0}\right)=G_{L}\left(u_{0}, f=\right.$ known $)$.

### 3.3.1 Bounds on $\mathcal{G}_{L}$

Using the results from section 1.6 we can now consider estimates on the observation operator itself.

Lemma 3.3.1. Assume that $u_{0} \in H$ and that $f \in C\left(0, T, H^{r}\right)$ for any $r>0$. Then there exists a constant $C$ independent of $u_{0}$ and $f$ such that

$$
\left|G_{L}\left(u_{0}\right)\right| \leq|z(0)|+C\left(\left\|u_{0}\right\|+\|f\|_{C\left(0, T ; H^{r}\right)}\right) .
$$

Proof. Using the Sobolev embedding theorem and lemma 1.6.2, for $1<s<2$

$$
\begin{aligned}
\left|G_{L}\left(u_{0}\right)\right| & =|z(t)| \\
& \leq|z(0)|+\int_{0}^{t}\|u(\tau)\|_{L^{\infty}} d \tau \\
& \leq|z(0)|+C \int_{0}^{t}\|u(\tau)\|_{s} d \tau \\
& \leq|z(0)|+C\left(\frac{1}{t^{s / 2}}\left\|u_{0}\right\|+\|f\|_{C\left(0, T ; H^{r}\right)}\right) .
\end{aligned}
$$

Now we consider the following lemma, which we require to prove that the operator $G_{L}$ is Lipschitz.

Lemma 3.3.2. Assume that $u_{0}, v_{0} \in H^{l}$ where $l \in(0, r+2)$, and that $f \in C\left((0, T), H^{r}\right)$ for $r>0$. Then there exists $C$ such that

$$
\left|G_{L}\left(u_{0}\right)-G_{L}\left(v_{0}\right)\right| \leq C\left(\|u\|_{l},\|f\|_{C\left(0, T ; H^{r}\right)}\right)\left(\left\|u_{0}-v_{0}\right\|_{l}+\|f-g\|_{C\left(0, T ; H^{1 / 2}\right)}\right) .
$$

Proof. Let $v(t)$ be a solution of the Stokes equations (2.4-2.5) with initial condition $v_{0} \in H^{l}$ and driving force $g$. Let $y(t)$ be the trajectory of the particle initially at $a \in \Omega$ and moving under the velocity field $v(t)$ according to the following ODE:

$$
\frac{d y}{d t}=v(y, t), \quad y(0)=a .
$$

$z$ is as we have previously defined, with respect to $u_{0}, f$ and $z(0)$. We assume without
loss of generality that $z(0)=y(0)$. Then we have that for $1<s<1+l$

$$
\begin{aligned}
\left|G_{L}\left(u_{0}\right)-G_{L}\left(v_{0}\right)\right|= & |z(t)-y(t)| \\
\leq & |z(0)-y(0)|+\int_{0}^{t}|u(z(\tau), \tau)-v(y(\tau), \tau)| d \tau \\
\leq & \int_{0}^{t}\|D u(\cdot, \tau)\|_{L^{\infty}}|z(\tau)-y(\tau)| d \tau \\
& +\int_{0}^{t}\|u(\cdot, \tau)-v(\cdot, \tau)\|_{L^{\infty}} d \tau \\
\leq & \int_{0}^{t}\|u(\cdot, \tau)\|_{1+s}|z(\tau)-y(\tau)| d \tau+\int_{0}^{t}\|u(\cdot, \tau)-v(\cdot, \tau)\|_{s} d \tau \\
\leq & \int_{0}^{t} C\left(\frac{1}{\tau^{(1+s-l) / 2}}\left\|u_{0}\right\|_{l}+\|f\|_{C\left([0, T] ; H^{r}\right)}\right)|z(\tau)-y(\tau)| d \tau \\
& +\int_{0}^{t} C\left(\frac{1}{\tau^{(s-l) / 2}}\left\|u_{0}-v_{0}\right\|_{l}+\|f-g\|_{C\left([0, T] ; H^{r}\right)}\right) d \tau .
\end{aligned}
$$

The singularities here are integrable. Furthermore, using Grönwall's lemma (lemma 1.6.3) with $w(t)=|z(t)-y(t)|, \beta=C\left(\left\|u_{0}\right\|_{l}+\|f\|_{C\left(0, T ; H^{1 / 2}\right)}\right)$ and $\alpha=C\left(\| u_{0}-\right.$ $\left.v_{0}\left\|_{l}+\right\| f-g \|_{C\left(0, T ; H^{1 / 2}\right)}\right)$ gives us that

$$
\begin{aligned}
|z(t)-y(t)| & \leq \alpha \leq \alpha \exp \left(\int_{0}^{t} \beta(s) d s\right) \\
& \leq C\left(\|u\|_{l},\|f\|_{C\left(0, T ; H^{r}\right)}\right)\left(\left\|u_{0}-v_{0}\right\|_{l}+\|f-g\|_{C\left(0, T ; H^{1 / 2}\right)}\right)
\end{aligned}
$$

We can now use these bounds on $G_{L}$ to calculate equivalent bounds of $\mathcal{G}_{E}$, where $f$ is known.

Corollary 3.3.1. Assume that $u_{0} \in H$ and that $f \in C\left(0, T, H^{r}\right)$ for any $r>0$. Then there exists a constant $C$ independent of $u_{0}$ and $f$ such that

$$
\left|\mathcal{G}_{L}\left(u_{0}\right)\right| \leq|z(0)|+C\left(\left\|u_{0}\right\|+1\right)
$$

Proof. Result from Lemma 3.3.1 with $f$ known.
Corollary 3.3.2. Assume that $u_{0}, v_{0} \in H^{l}$ where $l \in(0, r+2)$, and that $f \in$ $C\left((0, T), H^{r}\right)$ for $r>0$. Then there exists $C$ such that

$$
\left|\mathcal{G}_{L}\left(u_{0}\right)-\mathcal{G}_{L}\left(v_{0}\right)\right| \leq C\left(\|u\|_{l}\right)\left\|u_{0}-v_{0}\right\|_{l} .
$$

Proof. Result from Lemma 3.3.2 with $f$ known.
Using these results, we can make the following assertions.
Corollary 3.3.3. Let $\mu_{0}=\mathcal{N}\left(0, \delta \mathcal{A}^{-\alpha}\right)$ for $\alpha>1$. Then $\mathcal{G}_{L}$ is measurable with respect to $\mu_{0}$, and the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by (3.3).

Proof. Result follows by corollary 1.10.1, lemma 1.8.1, and corollaries 3.3.1 and 3.3.2.

### 3.4 Numerical Approximation of the Operator $\mathcal{G}_{L}$

Suppose we now have a weak solution of (2.4-2.5) and we wish to calculate $\mathcal{G}_{L}\left(u_{0}\right)$, the Lagrangian observation operator. We assume that we are observing $J$ passive tracers $\left\{z_{j}(t)\right\}$ in the velocity field, with initial positions $\left\{x_{j}\right\}$, and observing each of these tracers noisily at $K$ times $\left\{t_{k}\right\}$. We now wish to numerically approximate the solutions to the set of ODEs given by

$$
\begin{align*}
\frac{d z_{j}}{d t} & =u\left(z_{j}(t), t\right) \forall t>0  \tag{3.4}\\
z_{j}(0) & =x_{j} . \tag{3.5}
\end{align*}
$$

We can approximate this by discretizing in time. We choose $\Delta t \ll \sup \{u(x, t), t \in$ $\mathbb{R}, x \in \Omega\}$. Suppose we are calculating the path of a tracer whose initial condition is
$z_{0}$, at some time $T$. We calculate iteratively for $1 \leq k \leq \frac{T}{\Delta t}$

$$
z((k+1) \Delta t)=z(k \Delta t)+\Delta t u(z(k \Delta t), k \Delta t) .
$$

This approximation has $\mathcal{O}(\Delta t)$ error. Moreover, since we only have the value of the field $u$ on a grid on the space, we also incur more error when interpolating these values to points in the space between grid points. This makes up a significant part of the error in calculating $\mathcal{G}_{L}$, as we shall show later.

There is also another source of error in approximating the trajectories of Lagrangian tracers in this manner. In order to be able to work with this problem on a computer, we are required to approximate the infinite dimensional function $u$ by a finite dimensional vector. Since we only know the value of the velocity field on a grid, we incur further error through the interpolation of these values to any point in the domain.

As discussed in section 1.14, there are various methods for interpolating the values of the velocity field on the grid points to any point in the domain. For the numerics that follow, we choose to forgo the improvement in accuracy that the bicubic interpolation algorithm provides in favour of the simpler bilinear interpolation, due to the factor of four increase on the computational cost that it requires.

Now that we can approximate the value of the observation operator $\mathcal{G}$, we can calculate the likelihood function,

$$
\mathbb{P}\left(y \mid u_{0}\right) \propto \exp \left(-\frac{1}{2}\left\|\mathcal{G}\left(u_{0}\right)-y\right\|_{\Sigma}^{2}\right) .
$$

We are now ready to use MCMC methods to sample from the well defined posterior distribution, combining prior knowledge, mathematical model, and observations.

### 3.5 General Setup

We once again implement a RWMH sampler as we did in the previous chapter. To recap, given currently accepted state (initial condition of the vector field) $u_{n}$, we propose $v$, where

$$
v=\left(1-\beta^{2}\right)^{1 / 2} u_{n}+\beta w,
$$

and $w \sim \mu_{0}=\mathcal{N}\left(0, \delta \mathcal{A}^{-\alpha}\right)$. With probability

$$
a(u, v)=\exp \left(\Phi\left(u_{n}\right)-\Phi(v)\right),
$$

where $\Phi(\cdot)=\frac{1}{2}\left\|\mathcal{G}_{L}(\cdot)-y\right\|_{\Sigma}^{2}, u_{n+1}=v$. Otherwise $u_{n+1}=u_{n}$. This algorithm gives us a Markov chain whose invariant distribution is the posterior distribution that we are interested in. For all of the numerics that follow, the general setup described in subsection 2.8 still holds.

Moreover, since we have shown Lipschitz continuity of the observation operator as shown in lemma 3.3.2, theorem 1.12 .1 holds, so we can scale $\beta$ appropriately to give larger average acceptance probabilities. Furthermore, a time step size of $\Delta t=0.01$ will be used throughout all the numerics that follow in this chapter.

### 3.6 Limitations of the Standard RWMH method

As we did in section 2.9, in this section we aim to show the advantages of the RWMH posed on function space as opposed to the standard RWMH method which works on discretized space.

In the following graphs we aim to show this happening in practise. In both cases, the same data set was taken which consists of 9 Lagrangian particles whose initial positions are placed on a grid, each of which follow the flow of the velocity field
and are observed at one time $T=0.1$. A draw from the prior was taken as the initial condition in the data creation process. Markov chains were started with a range of values of $\beta$ for both the RWMH and SRWMH method. The average acceptance probabilities were taken until they had converged sufficiently. This process was repeated for a range different grid sizes, with $\Delta u \in\left\{\frac{1}{10}, \frac{1}{20}, \frac{1}{50}, \frac{1}{100}, \frac{1}{200}, \frac{1}{500}\right\}$. The results are presented in Figures 3.2 and 3.3.


Figure 3.2: Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, SRWMH

Figure 3.2 shows clearly the degeneration of the SRWMH algorithm as the grid is refined. As the Fourier truncation has more non-zero terms, the average acceptance probability for a given value of $\beta$ reduces. This means that as we increase the number of Fourier modes, we are unable to make such large steps in state space. This means that we explore the state space slower and the method becomes increasingly inefficient.

In contrast, Figure 3.3 shows the same plots for the RWMH method framed on


Figure 3.3: Average acceptance probabilities for a range of different step sizes $\beta$ and grid sizes, RWMH
function space, whose acceptance probabilities have no dependence on the dimension of the discretization. This independence is borne out in the results which clearly show that as the mesh is refined, the average acceptance probabilities stay exactly the same for a given step size $\beta$, allowing us to make large steps in state space even if we have a very refined grid.

These results are consistent with those in section 2.9 and show exactly why the MCMC method posed on function space is superior to that which is posed on a problem which is already discretized.

### 3.7 Validating the Random Walk Algorithm

In the Eulerian case, we were able to calculate explicitly without sampling error the distribution that we were trying to characterise due to the linearity of the observation
operator. In the Lagrangian case we can no longer do this, however we can still check that the algorithm behaves as we would expect.

Equivalently to Figure 2.23 in subsection 2.10, Figure 3.4 shows how the value of one particular Fourier mode of the initial condition in the Markov chains with different starting states all converge to the same distribution, in the case of Lagrangian data. A large amount of data is used so that we see how the chains quickly converge to the area of high probability in the state space. The Markov chains appear to converge to a single value, but actually the posterior is just very concentrated around this value due to the high volume of data, leading to low levels of uncertainty.


Figure 3.4: Convergence of Markov chains with different initial states, with Lagrangian data

As in the Eulerian case, the value of $\beta$ has no bearing on the posterior distribution, but merely on the speed of convergence of the Markov chain. We omit numerical examples of this here, but equivalent results can be seen in section 2.10, in Figures 2.24 and 2.25 .

### 3.8 Inverse Crimes

As we did in section 2.11, we further test the algorithms by using data created with different parameters, resolutions, observation noise models, and time step sizes, to ensure that we are not committing inverse crimes.

In all of the following figures (apart from those that specify a lower resolution) the data was created using a grid for the vector field of $1000 \times 1000$ points (or $5 \times 10^{5}$ complex Fourier modes). The data assimilation algorithm was run with varying numbers of grid points for the velocity field approximation, running through 16, 100, 196 and finally 400 points. Figure 3.5 shows that as the number of grid points used in the velocity field approximation for the algorithm is increased, the marginal distribution for this particular Fourier mode appears to converge to a limit. (The noise levels and quantity of data are such that the posterior is not a peaked distribution on the true Fourier mode).

An in depth discussion of what can happen when there is a mismatch in the forcing functions for the dynamical system between the data environment and the algorithm for both Eulerian and Lagrangian data can be found in chapter 4, so we shall skip over this particular area in this section.

Figure 3.6 shows how the posterior changes if we have low noise the data, and gradually increase the variance of the noise in the likelihood function in the model. As we increase our uncertainty in the data, the posterior measure converges to the prior distribution.

Figure 3.7 shows the marginal distributions of one Fourier mode with Lagrangian data created with a $4 \times 4$ grid, with a varying number of grid points used in the assimilation model. The resolution in the algorithm which gets closest to the correct answer is actually the lowest resolution, with 16 grid points on which the velocity field is


Figure 3.5: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, high resolution data, Lagrangian data.
approximated. This is not surprising given that this is the resolution at which the data was created.

So in conclusion, we have demonstrated numerically that if we increase the resolution of the approximation of the velocity field in our algorithm, that the posterior converges in distribution. We have also presented how various mismatches in data creation and model used in the algorithm can affect the results.

### 3.9 Properties of the Posterior Measure

We now consider analogous experiments in the Lagrangian case to those presented in section 2.12 with Eulerian data. Consider an example where we have 100 evenly spaced times at which we observe a set of passive tracers, whose initial positions at time $t=0$ are evenly spaced on a grid. Figure 3.8 shows that the chains have converged to


Figure 3.6: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing variance in the noise model of the algorithm, low actual noise, Lagrangian case
approximate Gaussian distributions. Moreover, we once again have posterior consistency as we increase the number of tracers that we are observing. Unlike the case for Eulerian data, we cannot verify that these distributions are correct as the non-linearity of the observation operator causes the posterior to be non-Gaussian. However, the fact that the chains have converged, and exhibit posterior consistency, strongly suggest that the algorithm is working well.

Next, we consider an example where we have a fixed number of 25 tracers, whose initial positions are evenly spaced on a grid. We observe each tracer at an increasing number of times evenly spaced on the unit interval. Figure 3.9 shows that the marginal distributions for $\operatorname{Re}\left(u_{0,1}\right)$ have converged to approximate Gaussian distributions, and exhibit posterior consistency as the number of temporal observations increases.


Figure 3.7: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing resolution in model, low resolution data, Lagrangian case

### 3.10 Growth of Data in Time

In the Lagrangian case where we set $f \equiv 0$, since the velocity field is decaying, we would expect the passive tracers to slow down to arbitrarily slow speeds in finite time. This leads us to a theorem regarding passive tracers in Stokes's flow.

Theorem 3.10.1. Given a weak solution $u \in H$ to the Stokes equations (2.4-2.5) with $f=0$, and $\delta>0, \exists T_{\delta}$ s.t for any passive tracer $z(t)$ satisfying the $O D E$

$$
\begin{aligned}
\frac{d z}{d t} & =u(z(t), t), \quad t>0 \\
z(0) & =x \in \Omega
\end{aligned}
$$

for any $x \in \Omega$, then

$$
\left|z\left(T_{\delta}\right)-z(t)\right|<\delta \quad \forall t>T_{\delta}
$$



Figure 3.8: Increasing numbers of observations in space, Lagrangian.

Proof. From theorem 2.12.1, we have an exponential bound on

$$
\|u(x, t)\|_{L^{\infty}}<C e^{-a t}
$$

for some $a>0$, as long as $t>\frac{1}{4 \nu \pi^{2}}$. So therefore, for $T_{\delta}>\max \left\{\frac{1}{4 \nu \pi^{2}}, \frac{1}{a} \ln \left(\frac{C}{a \delta}\right)\right\}$

$$
\begin{aligned}
\left|z\left(T_{\delta}\right)-z(t)\right| & =\left|\int_{T_{\delta}}^{t} u(z(s), s) d s\right| \\
& \leq \int_{T_{\delta}}^{t}|u(z(s), s)| d s \\
& <C \int_{T_{\delta}}^{t} e^{-a s} d s \\
& =-\frac{C}{a}\left[e^{-a s}\right]_{T_{\delta}}^{t} \\
& =\frac{C}{a}\left(e^{-a T_{\delta}}-e^{-a t}\right) \\
& <\frac{C}{a} e^{-a T_{\delta}} \\
& <\delta
\end{aligned}
$$



Figure 3.9: Increasing numbers of observations in space, Lagrangian.

In Figure 3.10, it shows that for Lagrangian observations, the information improves as you take observations further into the future, up to a point, after which the information stays pretty much the same. This improvement is due to the fact that if we observe the tracer further in the future, it has had longer to be influenced by (particularly the low frequency) Fourier modes of the initial condition of the velocity field, allowing it to accrue more information about them. The information no longer improves once the position of the tracer is confined within a $\delta$-ball that has a significantly smaller radius than the variance of the additive noise, which then dominates, or in the case where forcing is present, where the value the initial condition has only negligible effect on the current value of the velocity field, as shown in theorem 2.12.3.

Let us consider an initial condition of the velocity field with $u_{0,1}=1, u_{2,2}=$


Figure 3.10: PDFs For Lagrangian Data, 9 Paths, Varying $T$


Figure 3.11: Sketch to show confinement of $z(t)$ into a $\delta$-ball
-0.2 , with all other Fourier modes zero. Running the sampler with observations being made on a grid at one varying time $T$ results in Figures 3.12 and 3.13.


Figure 3.12: PDFs Of $u_{0,1}$ For Lagrangian Data, Varying Observation Time, With Only $u_{0,1}$ And $u_{2,2}$ Present

We see in both of these figures that the information for both the high and low frequency Fourier modes improves as the observation time increases. In the case of Figure 3.13, this improvement hits a maximum as the influence on the path of the particle of this Fourier mode becomes insignificant. After around $t=0.3$, as it's value has reduced to

$$
-0.2 \times e^{-0.3 * 4 * \nu * \pi^{2}\left(2^{2}+2^{2}\right)} \approx-0.00175
$$

the effect of this Fourier mode on the trajectory of the tracer is very small indeed. However, the information gained from making observations of the tracers about this mode does not decay, it simply remains at this level, as it has already influenced the trajectory of the tracer.


Figure 3.13: PDFs Of $u_{2,2}$ For Lagrangian Data, Varying Observation Time, With Only $u_{0,1}$ And $u_{2,2}$ Present

The position of the particle remains sensitive to the initial conditions of the velocity field for all time. This is in stark comparison to Eulerian data.

### 3.11 Conclusions and Future Directions

By careful analysis of the forward problem, we have been able to formulate a wellposed Bayesian inverse problem regarding Lagrangian data of the Stokes' flow dynamical system. We have shown that the likelihood function is continuous with respect to a space that has full measure with respect to a specified choice of Gaussian prior measure. Using this, we have shown how to draw samples from well defined posterior distributions on function space using the RWMH MCMC sampler. We have then implemented this algorithm in C, and gained insight into what kind of information is available in Lagrangian data. We have also shown how the standard RWMH method loses efficiency as the grid
is refined, and that the method framed on function space does not. We have also verified the algorithm behaves as we would expect in a range of situations and have shown that we are mindful of committing inverse crimes.

This chapter demonstrates again our philosophy underlying these methods; the belief that formulating numerical methods on infinite dimensional spaces and only discretizing once we choose to implement such a method gives us better algorithms that are robust under different discretisations and refinements. This time however, we have demonstrated that the algorithm works in a non-linear data environment where the solution we are looking for cannot be found via analytical techniques.

This subject could be extended in many directions, some of which will be addressed in later chapters. We could (similarly to the future work suggestions in the previous chapter) implement the full Navier-Stokes equations, the MALA algorithm and deterministic burn-in (requiring the implementation of an adjoint to the forward model for calculation of the gradient of the observation operator), and gPC approximations of the forward model.

In the next chapter, we shall consider an extension to chapters 2 and 3 where we add uncertainty in the model itself, in the form of model error. We will attempt to recover not only the initial condition of the velocity field but also the space-time dependant forcing that was present during the entire assimilation window, from both Eulerian and Lagrangian observations.

## Chapter 4

## Data Assimilation of Model Error

### 4.1 Motivation

The models inherent in oceanography and meteorology are incredibly complex, and are constantly being refined and updated. As such, reanalysis of data to try to quantify the differences between computational models and the actual dynamical system are of great interest to both communities. In thinking about this problem in this chapter we will consider both data scenarios that we have discussed in the previous two chapters; Eulerian and Lagrangian.

The method which we set out herein could be used in reanalysis of data sets to try to learn more about these discrepancies, or model error as they are often termed. This problem is equivalent, assuming that your model is a reasonable description of the real world dynamics, to finding an external forcing function to the system.

In the following section, we will analyse what can happen if your model is not representative of the environment from which your observations are made. More specifically we look at a low dimensional Eulerian problem enabling us to calculate the exact
posterior distribution as we did in subsection 2.7. We will then go on to show some results from the algorithms in chapters 2 and 3 with mismatched forcing in algorithmic model and the actual data environment.

### 4.2 Mismatched Statistical Model and Data Environment

Suppose that we model our velocity field using model A, with the following equations,

$$
\begin{gather*}
\frac{d u}{d t}+A u=f(t), \quad \forall t>0  \tag{4.1}\\
u(0)=u_{0} \tag{4.2}
\end{gather*}
$$

but that in fact the dynamical system we are looking at, model $B$, is governed by the equations

$$
\begin{gather*}
\frac{d u}{d t}+A u=g(t), \quad \forall t>0  \tag{4.3}\\
u(0)=u_{0} \tag{4.4}
\end{gather*}
$$

It is important to understand how a discrepancy in the model such as this could affect the results of the data assimilation algorithm described in chapter 2. For example, suppose we set $f \equiv 0$, and $g=\sum_{k} g_{k} \phi_{k}$ a constant function in time. Let $u_{A}$ be the solution to (4.1-4.2), and $u_{B}$ be the solution to (4.3-4.4) with $u_{A}(0)=u_{B}(0)=u_{0}=\sum_{k} u_{k} \phi_{k}$. Then they are given by

$$
\begin{aligned}
u_{A}(t) & =\sum_{k} u_{k} \exp \left(-4 \nu \pi^{2}|k|^{2} t\right) \phi_{k} \\
u_{B}(t) & =\sum_{k}\left(\frac{g_{k}}{4 \nu \pi^{2}|k|^{2}}\left(1-\exp \left(-4 \nu \pi^{2}|k|^{2} t\right)\right)+u_{k} \exp \left(-4 \nu \pi^{2}|k|^{2} t\right)\right) \phi_{k}
\end{aligned}
$$

The difference between these two functions is then given by

$$
u_{B}(t)-u_{A}(t)=\sum_{k}\left(\frac{g_{k}}{4 \nu \pi^{2}|k|^{2}}\left(1-\exp \left(-4 \nu \pi^{2}|k|^{2} t\right)\right) \phi_{k}\right)
$$

This will naturally drastically alter both Eulerian and Lagrangian data, depending on the choice of the $\left\{g_{k}\right\}$. Let us consider a case with only a finite number of Fourier modes being non-zero so that we may once again calculate the analytical posterior distribution as we did in subsection 2.7. Using the same notation as subsection 2.7, we define our initial condition by letting $a_{1}=1, b_{1}=-0.7, a_{2}=-0.2$, and $b_{2}=0.1$. Similarly, to define the forcing $g$, we let $g_{k_{1}}=0.1, h_{k_{1}}=-0.1, g_{k_{2}}=0.05, h_{k_{2}}=0.05$, and all other $g_{k}, h_{k}=0$, where

$$
g=\sum_{k} g_{k} \sqrt{2} \sin (2 \pi(k \cdot x))+h_{k} \sqrt{2} \cos (2 \pi(k . x)) .
$$

We now look at the analytical posterior distribution if we take observations from this system and naively use them with the model using equations (4.1-4.2).

Since the field in (4.1-4.2) converges to the steady state

$$
u_{\text {steady }}=\sum_{k} \frac{g_{k}}{4 \nu \pi^{2}|k|^{2}} \phi_{k},
$$

but our chosen model expects the velocity field to decay exponentially in each Fourier mode, we would expect to see a big mismatch in this situation. Let us look at the posterior distributions when 100 observations are made on a grid at a time $t$. Figure 4.1 shows how the mean values for the low frequency modes change as we vary the observation time $t$. This shows how, due to the forcing, we overestimate the initial condition of the velocity field. However, as time increase, the prior becomes more and more dominant and overrides this effect. Figure 4.2 shows the same thing for the higher frequency modes. The prior becomes dominant sooner in time for the higher frequency modes. Figure 4.3 shows how the posterior covariance also converges to the prior as $t \rightarrow \infty$.

This dominance of the prior as $t \rightarrow \infty$ is easily explained by looking at the form of the analytical mean and covariance of the posterior. First we note the behaviour of $B$


Figure 4.1: Analytical posterior mean values using model $A$ with data from model $B$, low frequency modes, observation time increasing
(the observation matrix) as the observation time increases. Let us denote $k=(0,1)^{\mathrm{T}}$, and $k^{\prime}=(2,2)^{\mathrm{T}}$, and the positions of the observation stations $\left\{x_{i}\right\}_{1}^{N}$. Let us also define $c_{k}(x)=\sqrt{2} \cos (2 \pi(k . x))$ and $s_{k}(x)=\sqrt{2} \sin (2 \pi(k . x))$. Since

$$
B=\left(\begin{array}{cccc}
e^{-\lambda_{k} t} s_{k}\left(x_{1}\right) \frac{k_{2}}{|k|} & e^{-\lambda_{k} t} c_{k}\left(x_{1}\right) \frac{k_{2}}{|k|} & e^{-\lambda_{k^{\prime}} t} s_{k^{\prime}}\left(x_{1}\right) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & e^{-\lambda_{k^{\prime}} t} c_{k^{\prime}}\left(x_{1}\right) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\
e^{-\lambda_{k} t} s_{k}\left(x_{1}\right) \frac{-k_{1}}{|k|} & e^{-\lambda_{k} t} c_{k}\left(x_{1}\right) \frac{-k_{1}}{|k|} & e^{-\lambda_{k^{\prime}} t} s_{k^{\prime}}\left(x_{1}\right) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & e^{-\lambda_{k^{\prime}} t} c_{k^{\prime}}\left(x_{1}\right) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} \\
e^{-\lambda_{k} t} s_{k}\left(x_{2}\right) \frac{k_{2}}{|k|} & e^{-\lambda_{k} t} c_{k}\left(x_{2}\right) \frac{k_{2}}{|k|} & e^{-\lambda_{k^{\prime}} t} s_{k^{\prime}}\left(x_{2}\right) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & e^{-\lambda_{k^{\prime}} t} c_{k^{\prime}}\left(x_{2}\right) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\
e^{-\lambda_{k} t} s_{k}\left(x_{2}\right) \frac{-k_{1}}{|k|} & e^{-\lambda_{k} t} c_{k}\left(x_{2}\right) \frac{-k_{1}}{|k|} & e^{-\lambda_{k^{\prime}} t} s_{k^{\prime}}\left(x_{2}\right) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & e^{-\lambda_{k^{\prime}} t} c_{k^{\prime}}\left(x_{2}\right) \frac{-\frac{k_{1}^{\prime}}{\left|k^{\prime}\right|}}{\vdots} \quad \vdots
\end{array}\right),
$$

each term is bounded by $e^{-\lambda_{k^{\prime}} t}$, since $\lambda_{k}<\lambda_{k^{\prime}}$. Therefore, $B$ converges exponentially to the zero matrix as $t \rightarrow \infty$. With this is in mind, we can see that

$$
\Sigma^{-1}=\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1} \longrightarrow \Sigma_{0}^{-1}
$$



Figure 4.2: Analytical posterior mean values using model $A$ with data from model $B$, high frequency modes, observation time increasing
and similarly

$$
m=\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}} y\right) \longrightarrow \Sigma_{0}\left(\Sigma_{0}^{-1}\right) m_{0}=m_{0}
$$

as $t \rightarrow \infty$.
A consequence of this is that if we assume that there is no forcing in our system, no matter how far our observations are from zero, if they are made at a large time $t$, our model will return the prior distribution.

The effects of an overestimation of the initial condition in the least squares solution, and the eventual domination of the prior distribution as the observation time $t$ increases leaves us the effect seen in Figure 4.1.

Notice that the plots in this graph are still affected by the noise in the observations. We can remedy this as we did before, by calculating the expectation of the mean


Figure 4.3: Analytical posterior covariance convergence to $\Sigma_{0}$ as observation time increases
of posterior. First note that

$$
y=B x+B_{F} f+\eta
$$

where $f=(0.1,-0.1,0.05,-0.05)^{\mathrm{T}}$ and

$$
\begin{aligned}
& B_{F}= \\
& \left(\begin{array}{cccc}
s_{k}\left(x_{1}\right) v \frac{k_{2}}{|k|} & c_{k}\left(x_{1}\right) v \frac{k_{2}}{|k|} & s_{k^{\prime}}\left(x_{1}\right) v^{\prime} \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{1}\right) v^{\prime} \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\
s_{k}\left(x_{1}\right) v \frac{-k_{1}}{|k|} & c_{k}\left(x_{1}\right) v \frac{-k_{1}}{|k|} & s_{k^{\prime}}\left(x_{1}\right) v^{\prime} \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{1}\right) v^{\prime} \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} \\
s_{k}\left(x_{2}\right) v \frac{k_{2} \mid}{|k|} & c_{k}\left(x_{2}\right) v \frac{k_{2}}{|k|} & s_{k^{\prime}}\left(x_{2}\right) v^{\prime} \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{2}\right) v^{\prime} \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\
s_{k}\left(x_{2}\right) v \frac{-k_{1}}{|k|} & c_{k}\left(x_{2}\right) v \frac{-k_{1}}{|k|} & s_{k^{\prime}}\left(x_{2}\right) v^{\prime} \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{2}\right) v^{\prime} \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} \\
\vdots & \vdots & \vdots & \vdots
\end{array}\right)
\end{aligned}
$$

where $v=\left(1-e^{-\lambda_{k} t}\right)$ and $v^{\prime}=\left(1-e^{-\lambda_{k}^{\prime} t}\right)$. Therefore,

$$
\begin{aligned}
\mathbb{E}(m) & =\mathbb{E}\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}} y\right) \\
& =\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}}\left(B x+B_{F} f\right)\right)
\end{aligned}
$$

Using this we can produce graphs which show much more clearly what is happening. Figure 4.4 shows the same behaviour as in Figure 4.1, with the noise removed.


Figure 4.4: Expectation of analytical posterior mean values using model $A$ with data from model B, low frequency modes

Using the expectation of the mean of posterior distribution gives us a much better chance to see what is happening in the high frequencies. Figure 4.5 shows the evolution of $m_{3}(t)$ and $m_{4}(t)$ on a faster timescale. Bearing in mind that in our example, the forcing is of the same sign as the initial condition for the low frequency modes and of the opposite sign for this higher frequency, we can see how both $m_{3}$ and $m_{4}$ swap signs very quickly, reaching a peak, and then decaying as the prior distribution begins to dominate.

This reinforces the fact that the forcing is dominating our results in comparison to the initial condition after a relatively short amount of time. Therefore, if our model and our data are mismatched in terms of their underlying dynamical systems, then our data is only useful for inferring on the condition of the vector field at times close to the time of the observation.


Figure 4.5: Expectation of analytical posterior mean values using model $A$ with data from model $B$, high frequency modes

Let us next consider a forcing which is no longer simply a constant and is dependant on time. For example, consider the case when

$$
f=\cos (2 \pi t)\left(\begin{array}{c}
0.1 \\
-0.1 \\
0.05 \\
-0.05
\end{array}\right)
$$

First off we must consider the solution of the following ODE to find the rate of change of each of the Fourier modes with respect to time,

$$
\frac{d a_{k}}{d t}=-4 \nu \pi^{2}|k|^{2} a_{k}+f_{k}(t)
$$

Using the integration factor $e^{4 \nu \pi^{2}|k|^{2} t}$, we obtain

$$
\frac{d}{d t}\left(a_{k} e^{4 \nu \pi^{2}|k|^{2} t}\right)=e^{4 \nu \pi^{2}|k|^{2} t} f_{k}(t)=e^{4 \nu \pi^{2}|k|^{2} t} \cos (2 \pi t) f_{k}(0)
$$

Integrating both sides gives us

$$
a_{k}(t)=a_{k}(0) e^{-4 \nu \pi^{2}|k|^{2} t}+g_{k}(t) e^{-4 \nu \pi^{2}|k|^{2} t}
$$

where, if we set $\lambda_{k}=4 \nu \pi^{2}|k|^{2}$,

$$
\begin{aligned}
g_{k}(t) & =\int_{0}^{t} e^{\lambda_{k} s} f_{k}(s) d s \\
& =-\int_{0}^{t} \frac{\lambda_{k}}{2 \pi} e^{\lambda_{k} s} \sin (2 \pi s) d s+\frac{1}{2 \pi} \sin (2 \pi t) e^{\lambda_{k} t} \\
& =-\int_{0}^{t} \frac{\lambda_{k}^{2}}{(2 \pi)^{2}} e^{\lambda_{k} s} \cos (2 \pi s) d s+\frac{1}{2 \pi} \sin (2 \pi t) e^{\lambda_{k} t}-\frac{\lambda_{k} e^{\lambda_{k} t}}{(2 \pi)^{2}} \cos (2 \pi t) \\
& =\left(1+\frac{\lambda_{k}^{2}}{(2 \pi)^{2}}\right)^{-1}\left(\frac{1}{2 \pi} \sin (2 \pi t) e^{\lambda_{k} t}-\frac{\lambda_{k} e^{\lambda_{k} t}}{(2 \pi)^{2}} \cos (2 \pi t)\right)
\end{aligned}
$$

Therefore, if $\mathbb{E}(m)=\left(\frac{B^{*} B}{\sigma^{2}}+\Sigma_{0}^{-1}\right)^{-1}\left(\Sigma_{0}^{-1} m_{0}+\frac{B^{*}}{\sigma^{2}}\left(B x+B_{F} f\right)\right)$, where

$$
f=(0.1,-0.1,0.05,-0.05)^{\mathrm{T}}
$$

then
$B_{F}=\left(\begin{array}{cccc}s_{k}\left(x_{1}\right) g_{k}(t) e^{-\lambda_{k} t} \frac{k_{2}}{|k|} & c_{k}\left(x_{1}\right) g_{k}(t) \frac{k_{2}}{|k|} & s_{k^{\prime}}\left(x_{1}\right) g_{k^{\prime}}(t) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{1}\right) g_{k^{\prime}}(t) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\ s_{k}\left(x_{1}\right) g_{k}(t) \frac{-k_{1}}{|k|} & c_{k}\left(x_{1}\right) g_{k}(t) \frac{-k_{1}}{|k|} & s_{k^{\prime}}\left(x_{1}\right) g_{k^{\prime}}(t) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{1}\right) g_{k^{\prime}}(t) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} \\ s_{k}\left(x_{2}\right) g_{k}(t) \frac{k_{2}}{|k|} & c_{k}\left(x_{2}\right) g_{k}(t) \frac{k_{2}}{|k|} & s_{k^{\prime}}\left(x_{2}\right) g_{k^{\prime}}(t) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{2}\right) g_{k^{\prime}}(t) \frac{k_{2}^{\prime}}{\left|k^{\prime}\right|} \\ s_{k}\left(x_{2}\right) g_{k}(t) \frac{-k_{1}}{|k|} & c_{k}\left(x_{2}\right) g_{k}(t) \frac{-k_{1}}{|k|} & s_{k^{\prime}}\left(x_{2}\right) g_{k^{\prime}}(t) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} & c_{k^{\prime}}\left(x_{2}\right) g_{k^{\prime}}(t) \frac{-k_{1}^{\prime}}{\left|k^{\prime}\right|} \\ \vdots & \vdots & \vdots & \vdots\end{array}\right)$.


Figure 4.6: Analytical posterior mean values using model $A$ with data from model $B$, low frequency modes. Forcing dependent on time

Figures 4.6 and 4.7 show how the mean of the analytical posterior varies as a function of time with data from a system that is forced in this way. Once again, as the oscillatory nature of the posterior mean shows, the forcing is dominant in the posterior mean, and we are not inferring accurately on the initial condition of the velocity field. Notice that this oscillatory nature is not so evident in the high frequency modes, purely because the prior becomes dominant before the effect can be seen.

We will now consider some results produced using data from a Stokes' system with non-zero constant forcing, but with the algorithmic model assuming that there is no forcing, so that $f \equiv 0$. We attempt to explain the data arising from a forced model through the initial condition for an unforced model. Figure 4.8 and Figure 4.9 show the marginal distributions of one Fourier mode in such a situation, in the case of Eulerian and Lagrangian data respectively, and where we steadily increase the number of observation


Figure 4.7: Analytical posterior mean values using model A with data from model B, high frequency modes. Forcing dependent on time
times. Two things are noteworthy: (i) the posterior tends towards a peaked distribution as the amount of data increases; (ii) this peak is not located at the true initial condition (marked with a black line). This incorrect estimate of the initial condition is, of course, because of the mismatch between model used for the assimilation and for the data generation. In particular the energy in the posterior on the initial condition is increased in an attempt to compensate for the model error in the forcing.

What we can conclude from all of this is that if we do not accurately represent the true dynamics of the system into our model then we run into very serious problems. Data assimilation is a marriage of data and model, and without high enough quality of either, our conclusions will be useless. We look to solve this problem by inferring from our data not only on the initial condition, but also on the force that is driving the system.


Figure 4.8: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing number of observation times, unmatched forcing in data and model, Eulerian data


Figure 4.9: $\operatorname{Re}\left(u_{0,1}(t)\right)$ : Increasing number of observation times, unmatched forcing in data and model, Lagrangian data

### 4.3 The Noisy Stokes Equations

We may wish to make sense of the Stokes equations driven by a noise process $\eta$ :

$$
\begin{align*}
\partial_{t} u-\nu \Delta u+\nabla p=\eta, & \forall(x, t) \in \mathbb{T}^{2} \times(0, \infty)  \tag{4.5}\\
\nabla \cdot u=0, & \forall t \in(0, \infty)  \tag{4.6}\\
u(x, 0)=u_{0}(x), & x \in \mathbb{T}^{2} . \tag{4.7}
\end{align*}
$$

Alternatively, this can be formulated as an ODE on the Hilbert space $H$ as follows;

$$
\begin{align*}
\frac{d u}{d t}+A u & =\mathcal{P} \eta, \quad t>0  \tag{4.8}\\
u(0) & =u_{0} \in H \tag{4.9}
\end{align*}
$$

From herein we shall abuse notation and denote $\eta=\mathcal{P} \eta$. We let

$$
\begin{align*}
|\eta|^{2} & =\int_{0}^{T}\|\eta(t)\|^{2} d t  \tag{4.10}\\
|\eta|_{l}^{2} & =\int_{0}^{T}\|\eta(t)\|_{l}^{2} d t \tag{4.11}
\end{align*}
$$

denote the norms on $\mathcal{H}:=L^{2}(0, T, H)$ and $\mathcal{H}^{l}:=L^{2}\left(0, T, H^{l}\right)$ respectively.
We now define two operators, $\mathcal{G}_{E N}$ and $\mathcal{G}_{L N}$, the Eulerian and Lagrangian observation operators for the equations (4.5-4.7).

Definition 4.3.1. Suppose that we have initial condition for the velocity field $u_{0} \in H$ and forcing function $\eta \in \mathcal{H}$. Let $u$ denote the solution with these inputs of the noisy Stokes equations. Then, given observation station positions $\left\{x_{j}\right\}_{j=1}^{J}$ and observation times $\left\{t_{k}\right\}_{k=1}^{K}$, the observation operator $\mathcal{G}_{E N}$ is defined to be the vector

$$
\mathcal{G}_{E N}\left(u_{0}, \eta\right)=\left\{u\left(x_{j}, t_{k}\right)\right\}_{j, k=1}^{J, K}
$$

Definition 4.3.2. Suppose that we have initial condition for the velocity field $u_{0} \in H$ and forcing function $\eta \in \mathcal{H}$. Let $u$ denote the solution with these inputs of the noisy Stokes equations. Then, given initial Lagrangian tracer positions $\left\{x_{j}\right\}_{j=1}^{J}$, consider the solutions to the set of ODEs given below;

$$
\begin{aligned}
\frac{d z_{j}}{d t} & =u\left(z_{j}(t), t\right) \\
z_{j}(0) & =x_{j} .
\end{aligned}
$$

Then, given observation times $\left\{t_{k}\right\}_{k=1}^{K}$, the observation operator $\mathcal{G}_{L N}$ is defined to be the vector

$$
\mathcal{G}_{L N}\left(u_{0}, \eta\right)=\left\{z_{j}\left(t_{k}\right)\right\}_{j, k=1}^{J, K} .
$$

### 4.4 The Prior Distribution

As in the previous two chapters, we will adopt a Gaussian prior on the initial condition, namely $\mathcal{N}\left(0, \delta A^{-\alpha}\right)$ for some $\alpha, \delta \in \mathbb{R}$. We wish to choose a prior on the model error $\eta=\sum_{k} \eta_{k}(t) \phi_{k}$ which allows us to similarly control this function's regularity. Consider the following infinite dimensional Ornstein-Uhlenbeck (OU) process.

$$
\begin{equation*}
\frac{d \eta}{d t}=-\Gamma \eta+\sqrt{\Lambda} \frac{d \beta}{d t}, \quad t \in[0, T] \tag{4.12}
\end{equation*}
$$

where $\beta$ is a canonical Brownian motion. $\eta(t)$, for any $t \geq 0$, is then a mean-zero Gaussian random variable, with variance

$$
\mathbb{E} \eta(t)^{2}=\frac{\Lambda}{2 \Gamma}
$$

If we enforce that $\frac{\Lambda}{2 \Gamma}=\delta A^{-\alpha}$, and set $\eta(0) \sim \mathcal{N}\left(0, \delta A^{-\alpha}\right)$, then we have ensured that for all times $t \geq 0, \eta(t) \sim \mathcal{N}\left(0, \delta A^{-\alpha}\right)$. This is desirable since we already have results about the regularity of this distribution. Therefore, we can pick the invariant measure of this SDE as the prior measure on the model error function $\eta$.

Theorem 4.4.1. Let $\mu_{0}$ denote the product measure $\mathcal{N}\left(0, \delta A^{-\alpha}\right) \times \chi$ where $\chi$ is the invariant measure on time valued functions in $H$ of the SDE (4.12) such that for all times $t \geq 0, \eta(t) \sim \mathcal{N}\left(0, \delta A^{-\alpha}\right)$. Suppose $\alpha>s+1$, then almost surely for $\left(u_{0}, \eta\right) \sim \mu_{0}$, $u_{0} \in H^{s}$ and $\eta \in C\left([0, T], H^{s}\right)$.

Proof. Proof mainly comes from lemma 1.8.1, which gives us that $u_{0}, \eta(t) \in H^{s}$. Since the solution to the SDE (4.12) has a continuous solution, and we have shown that $\eta(t) \in H^{s}$ for all $t \geq 0$, we are done.

Note that you could require different amounts of regularity in $u_{0}$ and $\eta$, and this is perfectly possible.

### 4.5 Bounds on Observation Operators

As we did in the two previous chapters, we once again wish to consider bounds on $\mathcal{G}_{E N}$ and $\mathcal{G}_{L N}$ to allow us to show that the posterior with Radon-Nikodym derivative

$$
\begin{equation*}
\frac{d \mu}{d \mu_{0}} \propto \exp \left(-\frac{1}{2}\|\mathcal{G}(u, \eta)-y\|_{\Sigma}^{2}\right) \tag{4.13}
\end{equation*}
$$

is well defined.

### 4.5.1 Bounds on $\mathcal{G}_{E N}$

Lemma 4.5.1. Consider the noisy Stokes equations (4.5-4.7). Assume that $u_{0} \in H$ and that $\eta \in L^{2}\left(0, T, H^{r}\right)$ for some $r>0$. Then there exists a constant $C=C\left(t_{0}\right)$ independent of $u_{0}$ and $\eta$ such that for any $s \in(1, r+1]$

$$
\left|\mathcal{G}_{E N}\left(u_{0}, \eta\right)\right| \leq C\left(\left\|u_{0}\right\|+\|\eta\|_{L^{2}\left(0, T ; H^{r}\right)}\right)
$$

provided that $\min _{k} t_{k}>t_{0}>0$.

Proof. The result follows from Lemma 2.5.1, with $f=\eta$.

We also later require Lipschitz continuity of this operator to ensure that theorem 1.12.1 holds.

Lemma 4.5.2. Suppose $u_{0}, v_{0} \in H$, and $\eta, \xi \in L^{2}(0, T ; H)$. Then

$$
\left|\mathcal{G}_{E N}\left(u_{0}, \eta\right)-\mathcal{G}_{E N}\left(v_{0}, \xi\right)\right| \leq C\left(\left\|u_{0}-v_{0}\right\|+\|\eta-\xi\|_{L^{2}(0, T ; H)}\right) .
$$

Proof. By linearity of $\mathcal{G}_{E N}$, and by lemma 4.5.1,

$$
\begin{aligned}
\left|\mathcal{G}_{E N}\left(u_{0}, \eta\right)-\mathcal{G}_{E N}\left(v_{0}, \xi\right)\right| & =\left|\mathcal{G}_{E N}\left(u_{0}-v_{0}, \eta-\xi\right)\right| \\
& \leq C\left(\left\|u_{0}-v_{0}\right\|+\|\eta-\xi\|_{L^{2}(0, T ; H)}\right) .
\end{aligned}
$$

Corollary 4.5.1. Let $\mu_{0}$ be the prior on $\left(u_{0}, \eta\right)$ as described in section 4.4 for $\alpha>1$. Then $\mathcal{G}_{E N}$ is measurable with respect to $\mu_{0}$, and the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by (4.13).

Proof. Result follows by corollary 1.10.1, lemmas 4.4.1, 4.5.1 and 4.5.2.

### 4.5.2 Bounds on $\mathcal{G}_{L N}$

Lemma 4.5.3. Assume that $u_{0} \in H$ and that $\eta \in C\left(0, T, H^{r}\right)$ for any $r>0$. Then there exists a constant $C$ independent of $u_{0}$ and $f$ such that

$$
\left|\mathcal{G}_{L N}\left(u_{0}\right)\right| \leq|z(0)|+C\left(\left\|u_{0}\right\|+\|\eta\|_{C\left(0, T ; H^{r}\right)}\right) .
$$

Proof. Follows from lemma 3.3.1 with $f=\eta$.

We also later require Lipschitz continuity of this operator to ensure that theorem 1.12.1 holds.

Lemma 4.5.4. Assume that $u_{0}, v_{0} \in H^{l}$ where $l \in(0, r+2)$, and that $\eta, \xi \in$ $C\left((0, T), H^{r}\right)$. Then there exists $C$ such that

$$
\left|\mathcal{G}_{L N}\left(u_{0}\right)-\mathcal{G}_{L N}\left(v_{0}\right)\right| \leq C\left(\|u\|_{l},\|f\|_{C\left(0, T ; H^{r}\right)}\right)\left(\left\|u_{0}-v_{0}\right\|_{l}+\|\eta-\xi\|_{C\left(0, T ; H^{1 / 2}\right)}\right)
$$

Proof. The result follows from lemma 3.3.2.

Corollary 4.5.2. Let $\mu_{0}$ be the prior on $\left(u_{0}, \eta\right)$ as described in section 4.4 for $\alpha>1$. Then $\mathcal{G}_{L N}$ is measurable with respect to $\mu_{0}$, and the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by (4.13).

Proof. Result follows by corollary 1.10.1, lemmas 4.4.1, 4.5.3 and 4.5.4.

### 4.6 Sampling Model Error

To draw a sample from our prior distribution for the model error function $\eta(x, t)=$ $\sum_{k} \eta_{k}(t) \phi_{k}(x) \in \mathcal{H}=L^{2}([0, T], H)$, we must calculate an approximation of a solution to the set of ODEs on $H$ given by

$$
\frac{d \eta_{k}}{d t}=-\gamma_{k} \eta_{k}+\sqrt{\lambda_{k}} \frac{d \beta_{k}}{d t}
$$

where the $\left\{\beta_{k}\right\}$ are iid Brownian motions. We can approximate solutions to this by discretizing in time,

$$
\begin{aligned}
\eta_{k}(t+\Delta t) & =e^{-\gamma_{k} \Delta t} \eta_{k}(t)+\sqrt{\lambda_{k}} \int_{t}^{t+\Delta t} e^{-\gamma_{k}(s-t)} d w(s) \\
& =e^{-\gamma_{k} \Delta t} \eta_{k}(t)+\xi
\end{aligned}
$$

where $\xi$ is a mean-zero Gaussian random variable, with variance

$$
\begin{aligned}
\mathbb{E} \xi^{2} & =\lambda \int_{t}^{t+\Delta t} e^{-2 \gamma(s-t)} d s \\
& =\frac{\lambda}{2 \gamma}\left(1-e^{-2 \gamma \Delta t}\right)
\end{aligned}
$$

Given a particular instance of $\eta$, defined on a discrete set of times $\{0, \Delta t, \ldots, T-$ $\Delta t, T\}$, we can approximate the vector field $u$ that satisfies (4.8-4.9), using a trapezoidal approximation of $\eta$. That is, if we assume

$$
\eta(s)=(\eta(t)-\eta(t+\Delta t))\left(\frac{t+\Delta t-s}{\Delta t}\right)+\eta(t+\Delta t) \quad s \in[t, t+\Delta t]
$$

then each $u_{k}$ is given by the iteration

$$
u_{k}(t+\Delta t)=e^{-a_{k} \Delta t} u_{k}(t)+\int_{t}^{t+\Delta t} e^{-a_{k}(t+\Delta t-s)} \eta_{k}(s) d s
$$

where $a_{k}=4 \nu \pi^{2}|k|^{2}$. This integral can be calculated using integration by parts, so that

$$
\begin{aligned}
\int_{t}^{t+\Delta t} e^{-a_{k}(t+\Delta t-s)} \eta_{k}(s) d s= & \left(\eta_{k}^{n+1}+\frac{t+\Delta t}{\Delta t}\left(\eta_{k}^{n}-\eta_{k}^{n+1}\right)\right) \int_{t}^{t+\Delta t} e^{-a_{k}(t+\Delta t-s)} d s \\
& +\frac{\eta_{k}^{n+1}-\eta_{k}^{n}}{\Delta t} \int_{t}^{t+\Delta t} s e^{-a_{k}(t+\Delta t-s)} d s \\
= & \left(\eta_{k}^{n+1}+\frac{t+\Delta t}{\Delta t}\left(\eta_{k}^{n}-\eta_{k}^{n+1}\right)\right) I_{1}+\frac{\eta_{k}^{n+1}-\eta_{k}^{n}}{\Delta t} I_{2},
\end{aligned}
$$

where

$$
\begin{aligned}
& I_{1}=\frac{1}{a_{k}}\left(1-e^{-a_{k} \Delta t}\right), \\
& I_{2}=\frac{t\left(1-e^{-a_{k} \Delta t}\right)+\Delta t-I_{1}}{a_{k}} .
\end{aligned}
$$

We now have the tools we require to draw samples from the prior distribution, which the posterior is absolutely continuous with respect to. We can also approximate both the observation operators $\mathcal{G}_{E N}$ and $\mathcal{G}_{L N}$, which allows us to calculate likelihoods of those samples.

### 4.7 Value of Data in Assimilation in Model Error

After some initial numerics, it became obvious that Eulerian data was not particularly informative about the value of the model error function $\eta$ at any given time. This is
mainly due to the form of the solution map for Stokes flow. The following theorem formalises this.

Theorem 4.7.1. Suppose our observation operator returns observations at a sequence of times $0<t_{1} \leq t_{2} \leq \ldots \leq t_{N} \leq T$. Then given an initial condition and forcing $\left(u_{0}, \eta\right) \in$ $H \times L^{2}((0, T), H)$, there exists an infinite number of alternative $\eta^{\prime} \in \times L^{2}((0, T), H)$ with $\eta \neq \eta^{\prime}$ almost everywhere, such that

$$
\mathcal{G}_{E}\left(u_{0}, \eta^{\prime}\right)=\mathcal{G}_{E}\left(u_{0}, \eta\right)
$$

Proof. We assume that we have found a $\psi \in L^{2}\left(\left(0, t_{l-1}\right), H\right)$ which is piecewise linear on each interval $\left(t_{i_{1}}, t_{i}\right)$ for $i \in\{1, \ldots, l-1\}$, with

$$
\psi(t)=\frac{t_{i}-t}{t_{i}-t_{i-1}} b^{i-1}+\frac{t-t_{i-1}}{t_{i}-t_{i-1}} b^{i} \quad \forall t \in\left[t_{i-1}, t_{i}\right]
$$

We assume that $\left\{b^{0}, b^{1}, \ldots, b^{l-1}\right\} \subset H$ with $b^{0}$ chosen arbitrarily, $t_{0}=0$, and that

$$
\begin{aligned}
\int_{0}^{t_{1}} e^{-A\left(t_{1}-s\right)} \psi(s) d s & =0 \\
\int_{0}^{t_{2}} e^{-A\left(t_{2}-s\right)} \psi(s) d s & =0 \\
\int_{0}^{t_{3}} e^{-A\left(t_{3}-s\right)} \psi(s) d s & =0 \\
& \vdots \\
\int_{0}^{t_{l-1}} e^{-A\left(t_{l-1}-s\right)} \psi(s) d s & =0
\end{aligned}
$$

We aim to show that $\exists b^{l} \in H$ s.t if $\psi$ is extended so that it is linear on $\left[t_{l-1}, t_{l}\right)$ with $\psi\left(t_{l}\right)=b^{l}$, then

$$
\int_{0}^{t_{l}} e^{-A\left(t_{l}-s\right)} \psi(s) d s=0
$$

Therefore we require that

$$
\begin{aligned}
\int_{0}^{t_{l}} e^{-A\left(t_{l}-s\right)} \psi(s) d s= & \int_{0}^{t_{l-1}} e^{-A\left(t_{l}-s\right)} \psi(s) d s+\int_{t_{l-1}}^{t_{l}} e^{-A\left(t_{l}-s\right)}\left(\frac{t_{l}-s}{t_{l}-t_{l-1}} b^{l-1}\right) d s \\
& +\int_{t_{l-1}}^{t_{l}} e^{-A\left(t_{l}-s\right)}\left(\frac{s-t_{l-1}}{t_{l}-t_{l-1}} b^{l}\right) d s \\
= & h^{l}+B_{l} b^{l}
\end{aligned}
$$

We calculate the Eigenvalues of $B_{l}$ :

$$
\begin{aligned}
B_{l} \phi_{k}= & \int_{t_{l-1}}^{t_{l}} e^{-A\left(t_{l}-s\right)}\left(\frac{s-t_{l-1}}{t_{l}-t_{l-1}} \phi_{k}\right) d s \\
= & \frac{1}{a_{k}\left(t_{l}-t_{l-1}\right)}\left(t_{l}-t_{l-1} e^{-a_{k}\left(t_{l}-t_{l-1}\right)}-\frac{1-e^{a_{k}\left(t_{l}-t_{l-1}\right)}}{a_{k}}\right. \\
& \left.-t_{l-1}\left(1-e^{-a_{k}\left(t_{l}-t_{l-1}\right)}\right)\right) \phi_{k} \\
= & \lambda_{k} \phi_{k} .
\end{aligned}
$$

If we wish to find $b^{l}=-B_{l}^{-1} h^{l} \in H$, then we need to show that $\sum_{k}\left(\frac{h_{k}^{l}}{\lambda_{k}}\right)^{2}<\infty$, where

$$
h^{l}=\sum_{i=1}^{l} \int_{t_{i-1}}^{t_{i}} e^{-A\left(t_{l}-s\right)}\left(\frac{t_{i}-s}{t_{i}-t_{i-1}}\right) b^{i-1} d s+\sum_{i=1}^{l-1} \int_{t_{i-1}}^{t_{i}} e^{-A\left(t_{l}-s\right)}\left(\frac{s-t_{i-1}}{t_{i}-t_{i-1}}\right) b^{i} d s
$$

Therefore,

$$
\begin{aligned}
b_{k}^{l}= & \left(t_{l}-t_{l-1}\right)\left(a_{k}^{-1}\left(1-e^{-a_{k}\left(t_{l}-t_{l-1}\right)}\right)-\left(t_{l}-t_{l-1}\right)\right)^{-1} \\
& \left(\sum_{i=1}^{l} \frac{b_{k}^{i-1}}{t_{i}-t_{i-1}}\left(a_{k}^{-1}\left(e^{-a_{k}\left(t_{l}-t_{i}\right)}-e^{-a_{k}\left(t_{l}-t_{i-1}\right)}\right)-\left(t_{i}-t_{i-1}\right) e^{-a_{k}\left(t_{l}-t_{i-1}\right)}\right)\right. \\
& \left.+\sum_{i=1}^{l-1} \frac{b_{k}^{i}}{t_{i}-t_{i-1}}\left(\left(t_{i}-t_{i-1}\right) e^{-a_{k}\left(t_{l}-t_{i}\right)}-a_{k}^{-1}\left(e^{-a_{k}\left(t_{l}-t_{i}\right)}-e^{-a_{k}\left(t_{l}-t_{i-1}\right)}\right)\right)\right)
\end{aligned}
$$

It now only remains to show that $b^{l} \in H$, where

$$
\begin{aligned}
b^{l}= & \sum_{k \in \mathbb{K}} b_{k}^{l} \phi_{k} \\
= & \left(t_{l}-t_{l-1}\right)\left(A^{-1}\left(I-e^{-A\left(t_{l}-t_{l-1}\right)}\right)-\left(t_{l}-t_{l-1}\right)\right)^{-1} \\
& \left(\sum_{i=1}^{l} \frac{b^{i-1}}{t_{i}-t_{i-1}}\left(A^{-1}\left(e^{-A\left(t_{l}-t_{i}\right)}-e^{-A\left(t_{l}-t_{i-1}\right)}\right)-\left(t_{i}-t_{i-1}\right) e^{-A\left(t_{l}-t_{i-1}\right)}\right)\right. \\
& \left.+\sum_{i=1}^{l-1} \frac{b^{i}}{t_{i}-t_{i-1}}\left(\left(t_{i}-t_{i-1}\right) e^{-A\left(t_{l}-t_{i}\right)}-A^{-1}\left(e^{-A\left(t_{l}-t_{i}\right)}-e^{-A\left(t_{l}-t_{i-1}\right)}\right)\right)\right) .
\end{aligned}
$$

We do this by bounding $\left|b_{k}^{l}\right|$,

$$
\begin{aligned}
\left|b_{k}^{l}\right| \leq & \max _{i \in\{1, \ldots, l\}}\left\{\frac{t_{l}-t_{l-1}}{t_{i}-t_{i-1}}\right\} \\
& \left(\sum_{i=1}^{l}\left|b_{k}^{i-1}\right|\left|\frac{a_{k}^{-1}\left(e^{-a_{k}\left(t_{l}-t_{i}\right)}-e^{-a_{k}\left(t_{l}-t_{i-1}\right)}\right)-\left(t_{i}-t_{i-1}\right) e^{-a_{k}\left(t_{l}-t_{i-1}\right)}}{a_{k}^{-1}\left(1-e^{-a_{k}\left(t_{l}-t_{l-1}\right)}\right)-\left(t_{l}-t_{l-1}\right)}\right|\right. \\
& \left.+\sum_{i=1}^{l-1}\left|b_{k}^{i}\right|\left|\frac{\left(t_{i}-t_{i-1}\right) e^{-a_{k}\left(t_{l}-t_{i}\right)}-a_{k}^{-1}\left(e^{-a_{k}\left(t_{l}-t_{i}\right)}-e^{-a_{k}\left(t_{l}-t_{i-1}\right)}\right)}{a_{k}^{-1}\left(1-e^{-a_{k}\left(t_{l}-t_{l-1}\right)}\right)-\left(t_{l}-t_{l-1}\right)}\right|\right) \\
= & C\left(t_{1}, \ldots, t_{l}\right)\left(\left|g_{1}^{l}\left(a_{k}\right)\right|\left|b_{k}^{l-1}\right|+\sum_{i=1}^{l-1}\left|g_{1}^{i}\left(a_{k}\right)\right|\left|b_{k}^{i-1}\right|+\left|g_{2}^{i}\left(a_{k}\right)\right|\left|b_{k}^{i}\right|\right)
\end{aligned}
$$

where

$$
\begin{aligned}
g_{1}^{i}(x) & =\frac{\left(e^{-x\left(t_{l}-t_{i}\right)}-e^{-x\left(t_{l}-t_{i-1}\right)}\right)-x\left(t_{i}-t_{i-1}\right) e^{-x\left(t_{l}-t_{i-1}\right)}}{\left(1-e^{-x\left(t_{l}-t_{l-1}\right)}\right)-x\left(t_{l}-t_{l-1}\right)} \\
g_{2}^{i}(x) & =\frac{\left(e^{-x\left(t_{l}-t_{i}\right)}-e^{-x\left(t_{l}-t_{i-1}\right)}\right)-x\left(t_{i}-t_{i-1}\right) e^{-x\left(t_{l}-t_{i}\right)}}{\left(1-e^{-x\left(t_{l}-t_{l-1}\right)}\right)-x\left(t_{l}-t_{l-1}\right)}
\end{aligned}
$$

The transcendental equation $f(x)=\left(1-e^{-x\left(t_{l}-t_{l-1}\right)}\right)-x\left(t_{l}-t_{l-1}\right)=0$ has only one solution on $\mathbb{R}^{+}$since $f^{\prime}(x)>0$ on $\mathbb{R}^{+}$, and $f(0)=0$. Therefore $g_{1}^{i}$ and $g_{2}^{i}$ are
continuous on $\left[\min _{k \in \mathbb{K}} a_{k}, \infty\right)$, and

$$
\begin{aligned}
\lim _{x \rightarrow \infty} g_{1}^{i}(x)=0 & \forall i \in\{1, \ldots, l\} \\
\lim _{x \rightarrow \infty} g_{2}^{i}(x)=0 & \forall i \in\{1, \ldots, l-1\} .
\end{aligned}
$$

Therefore $g_{1}^{i}$ and $g_{2}^{i}$ are bounded on $\left[\min _{k \in \mathbb{K}} a_{k}, \infty\right)$ by a constant independent of $k$, but dependent on choice of $\left\{t_{1}, \ldots, t_{l}\right\}$, for the required range of $i$. Therefore

$$
\left|b_{k}^{l}\right| \leq C\left(t_{1}, \ldots, t_{l}\right) \sum_{i=0}^{l-1}\left|b_{k}^{i}\right| .
$$

Since $\left\{b^{i}\right\}_{i=0}^{l-1} \subset H$ and the constant $C$ is independent of $k$, we have shown that $b^{l}=-B_{l}^{-1} h^{l} \in H$.

Corollary 4.7.1. If we discretize $\eta$ in time and truncate the number of Fourier modes $k$ that we assume are non-zero, and assume that $\left\{\eta_{n}\right\}$ is linear between time steps, and that observations of the vector field are only made at time steps, then there are an infinite number of discrete $\left\{\eta_{n}^{\prime}\right\}$ with $\eta_{n} \neq \eta_{n}^{\prime}$ almost everywhere, such that

$$
\mathcal{G}_{E}\left(u_{0},\left\{\eta_{n}\right\}\right)=\mathcal{G}_{E}\left(u_{0},\left\{\eta_{n}^{\prime}\right\}\right) .
$$

Proof. The construction of the proof is identical to that of the previous theorem.

Therefore, Eulerian data is far more informative about the value of

$$
F\left(t_{k}\right)=\int_{0}^{t_{k}} \exp \left(-\nu A\left(t_{k}-s\right)\right) \eta(s) d s
$$

where $\left\{t_{k}\right\}_{k=1}^{K}$ are the observation times. This is simply the contribution of the forcing to the solution map of Stokes' flow. We will demonstrate this in the numerics that follow.

Note that this theorem is not applicable in the Lagrangian case. A Lagrangian particle's position at any observation time $t_{k}$ is dependent on it's initial position, and the value of the vector field at it's current position at each point in time for all time. So any difference in the value of any of the Fourier modes of $\eta$ for any period of time with non-zero measure will result in different final position of most Lagrangian particles, except those particles that happen to lie at the points in the field which are zeros of the Fourier mode which has been changed. In other words, Lagrangian data is informative about the value of the model error function for all $0 \leq t \leq t_{K}$, where $t_{K}$ is the final observation time. This does, however, make a random walk proposal for the model error function much harder to accept, as we will go on to explain. We now look to numerical simulations of Eulerian and Lagrangian data assimilation with model error.

### 4.8 Properties of the Posterior Measure

### 4.8.1 Eulerian Data Assimilation With Model Error

In the previous section we showed that In the Eulerian case, parts of the forcing are unobservable and cannot be determined by the data. To state this precisely, define

$$
F\left(t_{2}\right)=\int_{0}^{t_{2}} \exp \left(-\nu A\left(t_{2}-t_{1}\right) \eta(t) d t\right.
$$

This is the term in the solution map of the Stokes equations which is dependent on the forcing term $\eta$. If the observations are made at times $\left\{t_{i}\right\}_{i=1}^{K}$ then the data is informative about $\underline{F}:=\left\{F\left(t_{j}\right)\right\}_{j=0}^{K}$, rather than about $\eta$ itself since infinitely many functions $\eta$ can give rise to the same vector $\underline{F}$, and therefore to the same velocity field at each of the observation times. Because of this, we expect that the posterior measure will give much greater certainty to estimates of $\underline{F}$ than $\eta$.

This basic analytical fact is manifest in the numerical results which we now describe. Firstly there are infinitely many functions compatible with the observed data so that obtaining a converged posterior on $\eta$ is a computationally arduous task; secondly the prior measure plays a significant role in weighting the many possible forcing functions which can explain the data.

As in the non-model error case of chapter 2 , we look at a set of experiments where the observations are made at one hundred evenly spaced times, with the observation stations evenly spaced on a grid with an increasing number of points. Once again, as can be seen in Figure 4.10, as we increase the number of observation stations, we are able to recover the value of any given Fourier mode in the initial condition with increasing accuracy and certainty as we increase the amount of data. As for all of the results that follow, a step size of $\Delta t=0.01$ is used for the model error problems, where we approximate functions in $L^{2}([0, T], H)$ by piecewise linear functions, which are linear on each interval $[n \Delta t,(n+1) \Delta t)$.


Figure 4.10: $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case.


Figure 4.11: $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case.

Figures 4.11 and 4.12 show the marginal posterior distributions on $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ and on $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ given an increasing number of observation stations. The first figure shows that, even with a large amount of data the standard deviation about the posterior mean for $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ is comparable in magnitude to the posterior mean itself. In contrast, for $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ and for a large amount of data, the standard deviation around the posterior mean is an order of magnitude smaller than the mean value itself. The data is hence much more informative about $\underline{F}$ than it is about $\eta$.

Figure 4.13 shows an example trace of the value of $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ in the Markov chain. Although the chain gives us a decent ballpark estimate of the point-wise value of the forcing, the trace shows the hallmarks of random walk behaviour, and has certainly not converged in distribution after the $10^{7}$ samples that have been calculated here. In contrast to Figure 4.13, Figure 4.14 shows how well the value of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ converges in distribution in the Markov chain. The whole trace stays within a reasonably tight


Figure 4.12: $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case.
band around the ergodic average, and does not display the random walk behaviour of Figure 4.13.

Figure 4.15 shows the expectation of the entire function $\operatorname{Re}\left(F_{0,1}(t)\right)$, given an increasing number of points in the vector field to be observed. As the number of observations increases, the expectation of $\operatorname{Re}\left(F_{0,1}(t)\right)$ nears the true value, given by the solid line.

If we now look at the posterior mean of the initial condition with a varying number of observation stations, and compare this to the true initial condition, we get an error curve as shown in Figure 4.16. This shows that, as the number of observation stations increases in a sensible way, the posterior mean of the initial condition converges to the true initial condition.

Similarly, if we look at the $L^{2}(0, T ; H)$-norm of the difference between the posterior mean of the time-dependent forcing, and the true forcing that created the data, we


Figure 4.13: Value of $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ in the Markov chain
get an error curve as shown in Figure 4.17. This shows that as the number of observation stations increases in a sensible way, that the posterior mean of the time-dependent forcing converges to the true initial condition. Notice however, that the convergence is slower in this quantity than in that given in Figure 4.18.

Figure 4.18 shows the $L^{2}(0, T ; H)$-norm of the difference between the posterior mean of $F=\int_{0}^{t} \exp \left(-a_{k}(t-s)\right) \eta(s) d s$ and the $F$ created from the forcing present in the creation of the data. Again, this shows that as we increase the number of observation stations, the posterior mean converges to the true answer. The convergence of this quantity is much quicker than that of $\eta$, but slower than that of $u_{0}$.

We now consider a similar experiment in which we keep the number of paths or observation stations constant, on a 20 by 20 grid, and increase the number of evenly spaced observation times, the last of which being $T=1$. Figure 4.19 shows the marginal distributions on one particular Fourier mode of the initial condition converging to a Dirac measure on the true value as the number of observation times increases.


Figure 4.14: Value of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ in the Markov chain

Similarly to Figure 4.11, Figure 4.20 again shows that due to theorem 4.7.1, our Markov chain will not be able to ascertain the value of the forcing at any one point in time. Due to this, in Figure 4.21 we once again look to the marginal distribution of $F=\int_{0}^{t_{k}} e^{-A\left(t_{k}-s\right)} \eta(s) d s$ in the Markov chain. As the number of observation times increases, the marginal distribution of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ appears to be converging to an increasingly peaked distribution on the true value.

Figure 4.22 shows the expectation of the entire function $\operatorname{Re}\left(F_{0,1}(t)\right)$ with an increasing number of observation times. As the number of observations in time increases, the expectation comes closer to the true function that created the data.


Figure 4.15: $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case.


Figure 4.16: $\left\|\mathbb{E}(u)-u_{\mathrm{Act}}\right\|_{L^{2}}$ : Increasing numbers of observations in space, Eulerian Model Error Case. $u_{\text {Act }}$ is actual initial condition that created the data.


Figure 4.17: $\left\|\mathbb{E}(\eta)-\eta_{\text {Act }}\right\|_{L^{2}(0, T ; H)}$ : Increasing numbers of observations in space, Eulerian Model Error Case. $\eta_{\text {Act }}$ is the actual forcing functions that created the data.


Figure 4.18: $\left\|\mathbb{E}(F)-F_{\text {Act }}\right\|_{L^{2}(0, T ; H)}$ : Increasing numbers of observations in space, Eulerian Model Error Case. $F_{\text {Act }}$ is the actual forcing functions that created the data.


Figure 4.19: $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case.


Figure 4.20: $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case.


Figure 4.21: $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case.


Figure 4.22: $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in time, Eulerian Model Error Case.

We may also be interested in understanding how well we are able to characterise high frequency (in space) forcing from Eulerian data. In the following experiment, all Fourier modes in the forcing that created the data were set to zero, apart from two high frequency modes for $k=(5,5)$ and $k=(4,5)$. An increasing number of observation stations were placed on a grid, with observations made at 100 evenly spaced times up to $T=1$. In the following graphs, the top graph in each pair (Figures 4.23, 4.25 and 4.27) shows the mean forcing function (an average over all the realisations in the Markov chain) for particular Fourier modes, as a function of time. The actual values that were present in the forcing that created the data are indicated by the solid line. The bottom graph in each pair (Figures 4.24, 4.26 and 4.28) shows the absolute value of the difference between the mean function and the true value that created the data.


Figure 4.23: $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing

In each Fourier mode, the more observations in space that are assimilated, the better the estimate. Moreover, the variance in these estimates (omitted here) reduces


Figure 4.24: $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing


Figure 4.25: $\operatorname{Re}\left(\eta_{4,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing


Figure 4.26: $\operatorname{Re}\left(\eta_{4,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing


Figure 4.27: $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, high frequency forcing


Figure 4.28: $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, high frequency forcing
as the amount of information increases, leading to peaked distributions on the forcing function which created the data.

Note that high frequency modes require more spacial observations to be able to make accurate estimates than the low frequency modes. This is simply due to the fact that with few spacial observations, no matter how many observations we have in time, our information about the high frequency Fourier modes is under-determined, and aliasing leads to a great deal of uncertainty about this modes.

So far we have only considered examples where the model error forcing that created the data is constant in time. In the following experiment we take a draw from the model error prior as the forcing function that is used in the creation of the data. This means that we have non-zero forcing in all of the Fourier modes, and this value is constantly changing at each time step also.

Similarly to the last experiment we present pairs of graphs where the first (Figures
4.29 and 4.33) shows the estimates (with varying numbers of spatial observations with a fixed number of 100 observations in time) of the forcing function for a particular Fourier mode, along with the actual forcing function that created the data. In the second graph in each pair (Figures 4.30 and 4.30) the absolute value of the difference between the estimates and the function that created the data are plotted.


Figure 4.29: $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior

These graphs show that as we increase the number of spatial observations, the estimates of the forcing functions converge to the function which created the data. Similarly more spatial observations are required before we get good estimates of the values of the high frequency Fourier modes.

### 4.8.2 Lagrangian Data Assimilation With Model Error

The Lagrangian equivalent of the model error problem is much harder to sample from in comparison with the Eulerian case. This is due to the fact that the position of a passive


Figure 4.30: $\operatorname{Re}\left(\eta_{0,1}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior


Figure 4.31: $\operatorname{Re}\left(\eta_{2,2}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior


Figure 4.32: $\operatorname{Re}\left(\eta_{2,2}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior


Figure 4.33: $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Increasing numbers of observations in space, Eulerian Model Error Case, forcing function taken from prior


Figure 4.34: $\operatorname{Re}\left(\eta_{5,5}(t)\right)$ : Absolute value of the difference between the mean and truth, Eulerian Model Error Case, forcing function taken from prior
tracer is dependent upon the entire history of the vector field up to the observation time. Moreover, a small change in the forcing function can result in the small displacement of hyperbolic points in the flow, which in turn can drastically alter the trajectory of one or more tracers. This makes it hard to move in the Lagrangian model error state space. As a consequence a simple random walk proposal is highly unlikely to be accepted. Indeed, this was borne out in our initial results, which showed that even after very large amounts of samples, the Markov chains were far from converged, as the value of $\beta$ in the proposal (6.9) that is required to give reasonable acceptance probabilities was simply too small to allow efficient exploration of the state space in a reasonable time.

One way to tackle this is to alter the likelihood to allow freer exploration of the state space. The data was created with observational noise with variance $\sigma^{2} I$, but if we were to increase the size of variance $\sigma^{2}$ used in the likelihood, then the acceptance probabilities increase, allowing larger steps in the proposal to be accepted more of the
time.
The results that follow in this section use data that was created with $\sigma^{2}=10^{-4}$, but use a covariance operator of $\hat{\sigma}^{2} I$ in the likelihood, with $\hat{\sigma}^{2}=25$ so that

$$
\mathbb{P}\left(y \mid u_{0}, f\right) \propto \exp \left(-\frac{1}{50}\left\|\mathcal{G}_{L}\left(u_{0}, f\right)-y\right\|\right)
$$

We will show that, despite the huge size of $\frac{\hat{\sigma}}{\sigma}$, it is possible to obtain reasonable estimates of the true forcing and initial condition. In particular, for large data sets, the posterior mean of these functions is close to the true values which generated the data. However, because $\frac{\hat{\sigma}}{\sigma}$ is large, the variance around the mean is much larger than in previous sections where $\hat{\sigma}=\sigma$.

Equivalently to the Eulerian case, we first consider the scenario where we have 100 equally spaced observation times up to time $T=1$, at which we observe the passive tracers, whose initial positions are given on a grid with an increasing number of points. Figure 4.35 shows how the marginal distributions on $\operatorname{Re}\left(u_{0,1}(0)\right)$ change as we increase the number of tracers to be observed. This figure indicates that as the number of spatial observations is increased in a sensible way, the marginal distribution on this particular Fourier mode is converging to an increasingly peaked distribution on the true value that created the data.

Figure 4.36 shows how the marginal distributions on $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ changes as we increase the number of paths to be observed. In comparison to the Eulerian case, we are able to determine much more about the point-wise value of the forcing function, here using the idea of inflated observational noise variance in the statistical model. Notice that, as in the Eulerian case, the uncertainty in the point-wise value of the forcing is far greater than that for the initial condition of the dynamical system.

Figure 4.37 shows distributions of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ and demonstrates convergence to a sharply peaked distribution on the true value that created the data in the limit of


Figure 4.35: $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case.
large data sets. The uncertainty in this quantity is less than for the point-wise value of the forcing, as in the Eulerian case, but the discrepancy is considerably less than in the Eulerian case.

Figure 4.38 shows the expectation of the entire function $\operatorname{Re}\left(F_{0,1}(t)\right)$, given an increasing number of points in the vector field to be observed. As the number of paths to be observed increases, the approximation does improve in places. However, since we altered the likelihood to make it possible to sample from this distribution, this also vastly increased the variance in each of the marginal distributions as there is relatively more influence from the prior. Therefore the expectation of this function does not tell the whole story. This picture does show us however that it is certainly possible to get ballpark estimates for these functions given Lagrangian data.

We now consider a scenario in which we have a fixed number of 400 paths whose initial conditions are on a 20 by 20 grid. We observe these paths on an increasing number


Figure 4.36: $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case.
of evenly spaced observation times, the last of which is made at $T=1$. The marginal distributions of $\operatorname{Re}\left(u_{0,1}\right)$ as shown in Figure 4.39 show that as we increase the number of observation times, we converge to an increasingly peaked distribution on the true value that created the data.

Figure 4.40 shows the marginal distributions on $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ as we increase the number of temporal observations. As we accrue more information, we can be more certain about the value of this quantity which created the data.

Figure 4.41 shows that as we increase the number of observation times, we are able to recover better information about the true value of $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ that created the data.

And finally, Figure 4.42 shows the expectation of the entire function $\operatorname{Re}\left(F_{0,1}(t)\right)$, given an increasing number of observation times. Once again, the fits are not fantastic to the true values that created the data, but the variances of the marginal distributions


Figure 4.37: $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case.
of the value of this function for each time are very large. It does certainly gives us a decent indicator as to what the forcing present in the system may be however.

### 4.9 Conclusions

By careful analysis of the forward problem, we have been able to formulate a well-posed Bayesian inverse problem regarding Eulerian and Lagrangian data of the Stokes' flow dynamical system, with model error. We have shown that the likelihood function is continuous with respect to a space that has full measure with respect to a specified choice of Gaussian prior measure. Using this, we have shown how to draw samples from well defined posterior distributions on function space using the RWMH MCMC sampler. We have then implemented this algorithm in $C$, and gained insight into what kind of information is available in Eulerian and Lagrangian data regarding the forcing in the system. We have also verified the algorithm behaves as we would expect in a range of


Figure 4.38: $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in space, Lagrangian Model Error Case.
situations in the limit of a large amount of data.
This chapter demonstrates again our philosophy underlying these methods; the belief that formulating numerical methods on infinite dimensional spaces and only discretizing once we choose to implement such a method gives us better algorithms that are robust under different discretisations and refinements.

The future directions suggested in sections 2.13 and 3.11 also hold here, with the addition of some others. Since we now have a framework with which to quantify the difference between the dynamical system of interest and the model on our computer, we should in theory be able to attempt to assimilate data which comes from fluid dynamical systems in the real world, for instance the oceans. Lagrangian float data is readily available on the internet, and by naively placing this data in the framework we have described, we should be able to recovery the nonlinear behaviour of the system through the model error term. The viscosity of the ocean is, in comparison with the


Figure 4.39: $\operatorname{Re}\left(u_{0,1}\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case.
contribution from the advection term in the full Navier-Stokes equations, very small. Assimilating this data in this way should mean that the model error term we recover will mainly describe that advection, which would be an interesting task to undertake.

To do this accurately, however, we would first need to tackle the problem of being able to explore the Lagrangian model error state-space without altering the covariance in the likelihood function. Since one of the main advantages of MCMC, other than being able to characterise non-Gaussian distributions, is to be able to accurately quantify uncertainty, it is important not to distort the distribution in this way. Otherwise one might as well implement other variational methods instead.

In the next section, we will briefly consider how the methods we have consider in chapters 3 and 4 and particularly chapter 2, can be built into a filtering framework.


Figure 4.40: $\operatorname{Re}\left(\eta_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case.


Figure 4.41: $\operatorname{Re}\left(F_{0,1}(0.5)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case.


Figure 4.42: $\operatorname{Re}\left(F_{0,1}(t)\right)$ : Increasing numbers of observations in time, Lagrangian Model Error Case.

## Chapter 5

## Filtering in Data Assimilation

### 5.1 Motivation

In many applications, particularly in weather forecasting, our data does not arrive at one time. We may wish to take our current best guess as to what the state of the system is, and then as the data comes in, we may wish to be able to improve that guess via Bayesian inference using that data. This then becomes our next best guess until more data comes in. This online analysis of data in real time is crucial to making regular accurate forecasts. This is referred to as filtering, as opposed to the method of using all of the data in one go, which is termed smoothing, of which we have demonstrated one approach in chapters 2-4. The approach that we have taken so far in these chapters is termed offline data assimilation or reanalysis.

In this chapter, we will briefly explain how we might go about applying the concepts we have discussed to a filtering algorithm in the scenario of Eulerian data assimilation with no model error. We will also look at filtering using the fact that in the Eulerian case of data assimilation in Stokes' flow we are able to explicitly determine
the Gaussian posterior at each step of the filtering. For the purposes of this chapter, we will assume that there is no stochastic forcing of the system, and that $f \equiv 0$ in the governing Stokes' flow equations (2.4)-(2.5).

### 5.2 Sequential Sampling

Firstly, we must think about the data scenario that we are intending to perform filtering on. We divide the time line into intervals, with data about each time interval being made available as we reach the end of that time interval. Each interval starts at a time $T_{k}$ and ends at $T_{k+1}$. For each iteration of the filtering algorithm, we will have a prior distribution (which is our best guess) on the value of the velocity field at the time $T_{k}$. We will then perform data assimilation as described before in chapter 2 or 3 , with the only difference being that we now introduce a non-zero mean for our prior distribution on $u_{k}$, the state of the velocity field at time $T_{k}$, which we shall denote $m_{k}$.

In the Eulerian case, we know that the posterior will be Gaussian, so we can characterise this distribution completely from the mean and covariance operator. In this case, given mean and covariance at a time $T_{k}$, we can calculate how this distribution would behave when pushed forward to time $T_{k+1}$ in Stokes' flow. This then gives us the prior distribution for the next iteration of the filtering process.

In the case that we have a non-linear observation operator, the posterior will not be Gaussian. However, in certain scenarios, the distribution will be close to Gaussian, and in others it is simply the best approximation that we can easily use. We make the approximation to a Gaussian, and then similarly evolve this distribution up to time $T_{k+1}$.

The main thing that has changed in this setup, apart from it's iterative nature, is the fact that we now have a non-zero mean for our prior distribution. We need to consider how we might go about sampling from this new distribution, which has

Radon-Nikodym derivative given by

$$
\frac{d \mu(u)}{d \mu_{0}}=\exp (-\Phi(u)), \quad \mu_{0} \sim \mathcal{N}(m, \mathcal{C})
$$

where $\Phi(\cdot)=\frac{1}{2}\|\mathcal{G}(\cdot)-y\|$. If we are to use the same proposal scheme as we have previously, so that

$$
v=\left(1-\beta^{2}\right)^{1 / 2} u+\beta \xi, \quad \xi \sim \mathcal{N}(0, \mathcal{C})
$$

then the acceptance probabilities are given by

$$
a(v, u)=\exp \left(\Phi(u)-\Phi(v)+\left\langle\mathcal{C}^{-1 / 2} m, \mathcal{C}^{-1 / 2}(v-u)\right\rangle\right.
$$

This follows from the following informal calculation.
Proposal state: $v=\left(1+\beta^{2}\right)^{\frac{1}{2}} u+\beta \mathcal{C}^{\frac{1}{2}} \xi, \xi \sim \mathcal{N}(0, I), \beta \in(0,1]$.
Target measure density:

$$
\begin{aligned}
\pi(u) & =\exp (-\Phi(u)) \exp \left(\frac{1}{2}\langle(u-m), \mathcal{L}(u-m)\rangle\right) \\
& =\exp (-\Phi(u)) \exp \left(-\frac{1}{2}\left|\mathcal{C}^{-\frac{1}{2}}(u-m)\right|^{2}\right)
\end{aligned}
$$

Transition kernel: $q(u, v)=\exp \left(-\frac{1}{2 \beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}}\left(v-\left(1-\beta^{2}\right)^{\frac{1}{2}} u\right)\right|^{2}\right)$.
Let $\pi(u) q(u, v)=\exp (-\Phi(u)) A(u, v)$. Then,

$$
\begin{aligned}
-2 \log (A)= & \left|\mathcal{C}^{-1 / 2}(u-m)\right|^{2}+\frac{1}{\beta^{2}}\left|\mathcal{C}^{-1 / 2}\left(v-\left(1-\beta^{2}\right)^{1 / 2} u\right)\right|^{2} \\
= & \left|\mathcal{C}^{-1 / 2} u\right|^{2}+\frac{1}{\beta^{2}}\left|\mathcal{C}^{-1 / 2}\left(v-\left(1-\beta^{2}\right)^{1 / 2} u\right)\right|^{2}+\left|\mathcal{C}^{-1 / 2} m\right|^{2} \\
& -2\left\langle\mathcal{C}^{-1 / 2} m, \mathcal{C}^{-1 / 2} u\right\rangle \\
= & \frac{1}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} v\right|^{2}+\frac{1}{\beta^{2}}\left|\mathcal{C}^{-\frac{1}{2}} u\right|^{2}-\frac{2\left(1-\beta^{2}\right)^{\frac{1}{2}}}{\beta^{2}}\left\langle\mathcal{C}^{-\frac{1}{2}} u, \mathcal{C}^{-\frac{1}{2}} v\right\rangle \\
& +\left|\mathcal{C}^{-1 / 2} m\right|^{2}-2\left\langle\mathcal{C}^{-1 / 2} m, \mathcal{C}^{-1 / 2} u\right\rangle
\end{aligned}
$$

Therefore the acceptance probabilities are given by

$$
a(v, u)=\exp \left(\Phi(u)-\Phi(v)+\left\langle\mathcal{C}^{-1 / 2} m, \mathcal{C}^{-1 / 2}(v-u)\right\rangle\right)
$$

Since the acceptance probability contains terms which would be infinite on function space, this method is not well-defined on function spaces and would lose efficiency as the mesh is refined.

Alternatively, we could sample $w=u_{0}-m$. Therefore the Radon-Nikodym derivative is given by

$$
\frac{d \mu(w)}{d \mu_{0}}=\exp (-\Phi(w+m)), \quad \mu_{0} \sim \mathcal{N}(0, \mathcal{C})
$$

where the prior is now on $w$, not on $u_{0}=m+w$. If we let our proposal for $w_{n+1}$, given currently accepted state $w_{n}$, be equal to

$$
w_{n+1}=\left(1-\beta^{2}\right)^{1 / 2} w_{n}+\beta \xi, \quad \xi \sim \mathcal{N}(0, \mathcal{C}),
$$

then the acceptance probabilities are given by

$$
a\left(w_{n+1}, w_{n}\right)=\exp \left(\Phi\left(w_{n}+m\right)-\Phi\left(w_{n+1}+m\right)\right),
$$

by a very similar argument to those shown in section 1.12.1. This method is far superior to the one which we previously described, as the acceptance probabilities are independent of the discretisation used, and the method is well defined on function space. This is the much more efficient choice for our sampling method.

The idea of this sequential sampler is, on each time interval, to approximate the posterior measure on the initial condition of that section by a Gaussian measure, and then push this measure forwards in time to create a prior measure on the value of the velocity field at the beginning of the next time interval. So far we have only considered updating the mean of the Gaussian measure. We will now set out how we could also push forwards the covariance operator. Consider the following system:

$$
\frac{d u}{d t}+\mathcal{A} u=0, \quad u(0) \sim \mathcal{C}(m, \Sigma)
$$

Since $u(t)=e^{-\mathcal{A t}} u(0)$, therefore

$$
\mathbb{E} u(t)=e^{-\mathcal{A} t} \mathbb{E} u(0)=e^{-\mathcal{A} t} m .
$$

Similarly,

$$
\begin{aligned}
\mathbb{E}\left(u(t)-e^{-\mathcal{A} t} m\right)\left(u(t)-e^{-\mathcal{A} t} m\right)^{\mathrm{T}} & = \\
\mathbb{E} e^{-\mathcal{A} t} u(0) u(0)^{\mathrm{T}} e^{-\mathcal{A} t}-e^{-\mathcal{A} t} m m^{\mathrm{T}} e^{-\mathcal{A} t} & = \\
\mathbb{E} e^{-\mathcal{A} t}\left(u(0) u(0)^{\mathrm{T}}-m m^{\mathrm{T}}\right) e^{-\mathcal{A} t} & = \\
\mathbb{E} e^{-\mathcal{A} t} \Sigma e^{-\mathcal{A} t} . &
\end{aligned}
$$

So if we run the sampler on a given section that starts at time $T_{i}$ and finishes at time $T_{i+1}$, we will keep a note of the mean and variance of the Fourier modes, which we denote by $m_{i}$ and $\Sigma_{i}$ respectively. Then the Gaussian approximation of $u_{i}=u\left(\cdot, T_{i}\right)$ is $u_{i} \sim \mathcal{N}\left(m_{i}, \Sigma_{i}\right)$. We then push this forward to give us a prior distribution on $u_{i+1}$ given by $\mu_{i+1} \sim \mathcal{N}\left(e^{-\mathcal{A}\left(T_{i+1}-T_{i}\right)} m_{i}, e^{-\mathcal{A}\left(T_{i+1}-T_{i}\right)} \Sigma_{i} e^{-\mathcal{A}\left(T_{i+1}-T_{i}\right)}\right)$.

We shall now compare three different sequential samplers. With the first (SS1), we shall apply exactly the same prior on each section, namely $\mu_{i}=\mathcal{N}\left(0, \mathcal{A}^{-\alpha}\right)$, so that no information is passed forwards to the next section. We would expect this sampler not to work too well.

The second sampler (SS2) will update the mean of the prior, but not the covariance operator. Therefore, the prior for each section is given by

$$
\mu_{i}=\mathcal{N}\left(e^{-\mathcal{A}\left(T_{i}-T_{i-1}\right)} m_{i-1}, \mathcal{A}^{-\alpha}\right)
$$

The third (SS3) will update the mean in the same fashion, and approximate the covariance operator as follows. We will approximate the covariance operator by a
diagonal matrix. That is, in each section, we will calculate the variance of each Fourier mode, and put these into a diagonal matrix, $\Sigma_{i-1}$. Then our prior distribution on $u_{i}$ is given by $\mu_{i}=\mathcal{N}\left(e^{-\mathcal{A}\left(T_{i}-T_{i-1}\right)} m_{i-1}, e^{-\mathcal{A}\left(T_{i}-T_{i-1}\right)} \Sigma_{i} e^{-\mathcal{A}\left(T_{i}-T_{i-1}\right)}\right)$.

The following figures show both the values of Fourier coefficients that we are trying to recover (blue crosses), and the error which we made in their estimation (red crosses).

Figures 5.1(a), 5.1(c) and 5.1(e) below are virtually identical, as with only one section there is no update of the prior, and so the experiments are all the same. Similarly, Figures 5.1(b), 5.1(d) and 5.1(f) are also the same, being the first section of a two section run. Each of these figures is made with the same prior distribution and with the same data, meaning their results differ only because of the use of a different set of random numbers.

Figures 5.2(a), 5.2(c) and 5.2(e) show the results of the 2 nd section, with varying updates on the prior distribution of the initial condition of the section, $u_{1}=u(\cdot, 0.5)$. For 5.2(a) and 5.2(c), which do not update the covariance matrix as 5.2(e) does, we see that although the red crosses lie below the blue crosses for the first two values of $|k|$, the errors quickly become far larger as $|k|$ increases, comparatively to the value of the Fourier coefficients we are trying to recover.

However, in 5.2(e), the errors now decay in the same way as the Fourier coefficients that we are trying to recover. The red crosses are also slightly lower than the blue crosses for at least the first 5 values of $|k|$, which is significant on this logarithmic scale. Figures 5.2(b), 5.2(d) and 5.2(f) are once again very similar as they are the first section of a 5 section sequential sampler.

Figures 5.2(a), 5.2(c) and 5.2(e) show the errors in the final section of the 5 section sequential sampler. As was shown in Figures 5.2(a), 5.2(c) and 5.2(e), updating


Figure 5.1: Results of different filtering approaches

(b) SS1, 5 Sections, 1st Section, Fourier Mode Errors

(c) SS2, 2 Sections, 2nd Section, Fourier ( Mode Errors Mode Errors

(e) SS3, 2 Sections, 2nd Section, Fourier (f) SS3, 5 Sections, 1st Section, Fourier Mode Errors

Figure 5.2: Results of different filtering approaches
the covariance matrix and the mean makes a huge difference to the quality of the results.
Figures $5.2(\mathrm{~b}), 5.2(\mathrm{~d})$ and $5.2(\mathrm{f})$ show plots of the errors for each section in a 20 section sequential sampler. In each plot, the red line shows the $L^{2}$-error. The blue line denotes the relative error, that is

$$
\frac{\| \text { Error } \|_{L^{2}}}{\left\|u_{i}\right\|_{L^{2}}}
$$

There is no consistency in the way the errors in Figures 5.2(b) and 5.2(d) change through time. However, interestingly the relative error appears to converge to a single value in Figure $5.2(\mathrm{f})$. This is not altogether surprising, as earlier plots showed that the errors appear to decay in the same way as the Fourier modes themselves.

In summary we can conclude that updating the mean and approximating the covariance operator as a diagonal matrix is well worthwhile for very little computational effort. There is no doubt that approximating the full covariance matrix would give further improvements.

### 5.3 Analytical Posteriors in Eulerian Sequential Sampling

We can also apply the concepts of subsection 2.7 to the case of Eulerian sequential sampling. To recap, in each section, starting at time $T_{k}$ and ending at $T_{k+1}$, we consider the observations in that time interval only. Incorporating the prior distribution $\mu_{0}^{k}$ on $u_{k}$, the state of the velocity field at time $T_{k}$, and the data, we arrive at a Gaussian distribution $\mu^{k}$ on $u_{k} \sim \mathcal{N}\left(m_{k}, \Sigma_{k}\right)$. Since in the case of Eulerian data, our observation operator is linear and therefore the posterior is Gaussian, there is no need to approximate this posterior as we have done previously in the Lagrangian case.

We can then "push-forward" this posterior distribution in time using the dynamics of the model, to arrive at a prior distribution $\pi_{0}^{k+1}$ on $u_{k+1}$, the state of the

(a) SS1, 5 Sections, 5th Section, Fourier Mode Errors

(c) SS2, 5 Sections, 5th Section, Fourier Mode Errors

(e) SS3, 5 Sections, 5th Section, Fourier Mode Errors

(b) SS1, 20 Sections, $L^{2}$ Error in Time

(d) SS2, 20 Sections, $L^{2}$ Error in Time

(f) SS3, 20 Sections, $L^{2}$ Error in Time

Figure 5.3: Results of different filtering approaches
vector field at time $T_{k+1}$. Using (2.7-2.8) we can calculate analytically the posterior distributions explicitly for each section in the sequence.

In the following results, each of these posterior distributions has been "pushed back" to become distributions on the initial condition of the velocity field. Let us consider a situation where we have 25 observation stations in space, making 1000 observations over 10 units of time. Suppose now we break these data sets into 100 equally sized sections, and apply a sequential algorithm as previously described. We can now look at the analytical posteriors on the initial condition as a function of the amount of data that has been considered. Let $m_{n}, \Sigma_{n}$ denote the $n^{\text {th }}$ analytical mean and covariance respectively. That is, the posterior distribution on the initial condition given the information up to $t=\frac{n}{10}$.


Figure 5.4: Convergence of $m_{n}$ to the analytical mean of the posterior using the whole data set

Figures 5.4 and 5.5 show how as more information is included, the analytical


Figure 5.5: Convergence of $\Sigma_{n}$ to the analytical covariance of the posterior using the whole data set
posterior distributions converge to the analytical posterior distribution of the whole data set. Also included here are exactly the same plots for the algorithm implemented previously with the covariance matrix approximated by a diagonal matrix. As the covariance matrices in this particular example are very strongly diagonally dominated, the red and blue plots are almost indistinguishable. Both of these error curves hit the floor of machine epsilon. Note here that these results, even if were able to calculate the analytic posteriors, would not be mirrored in the Lagrangian case. The non-linear nature of the problem means that with each Gaussian approximation we would create further errors.

### 5.4 Conclusions

In this chapter, we have very briefly introduced a way in which the smoothing methods described in chapters 2-4 could be adapted into a filtering algorithm.

However, current computational limits dictate that these methods cannot be currently used in the context which these methods are often used - atmospheric forecasting. The forward models in these cases are far more complex than the simple linear example we have been looking at so far, and to run this forward model enough times to have a converged chain in a short enough time that the forecast is still relevant is simply not possible yet. This is not to say that this may not become possible with more powerful computers in the future. When/if this ever happens, then this approach may become relevant to people who currently are interested in filtering methods, and may be worth addressing again.

This chapter concludes our analysis of data assimilation problems involving noisy observations of Stokes' flow. In the next chapter we introduce another data assimilation problem with applications in the biomedical sciences.

## Chapter 6

## A Data Assimilation Problem in

## Shape Registration

### 6.1 Motivation

There are many applications of shape registration, the process of comparing two shapes, particularly in the biomedical sciences. For example, after completing an antenatal ultrasound scan, you might wish to check whether the shape of the foetus' organs show traits which are common in foeti with certain congenital conditions. This would first involve segmentation of the image to isolate that organ's boundary, and then matching that shape against a library of organ shapes of prime examples of those congenital conditions. At this point, we need a measure on how "far apart" two shapes are.

In this chapter, we will first define this problem in the framework of a minimisation problem over solutions of a PDE. We will then show how this minimisation problem can be translated into a data assimilation problem on function space. At this point, we will be able to utilise the data assimilation framework that we have introduced in previous
chapters to compute numerical results which incorporate estimates of uncertainty.
The use of flow fields to describe deformation of shapes is a commonly used method [8]. However other techniques, such as elastic matching [29, 67, 85], and spline methods[62] also exist.

The method that we will describe herein relates closely to the Lagrangian data assimilation application addressed in chapter 3 , in the sense that we are trying to find a flow field which has advected one curve into another, guide by noisy observations of the final curve. Thus we construct a very similar approach, once again framed on function space. We go on to prove results regarding the forward problem which then allow us to frame the ill-posed inverse problem as a well-posed Bayesian inverse problem with well-defined posterior measures for which we have already demonstrated algorithms with which we are able to draw samples. We will then go on to present numerical results from the algorithm.

### 6.2 Equations of motion for the curve matching problem

We derive the equations of motion for curves in the plane acted on by geodesics in the diffeomorphism group. We parametrise the curve as a continuous function $q$ from a space $S$ (such as the circle, $S^{1}$ ) into $\mathbb{R}^{2}$ (all of this can be easily extended to surfaces in $\mathbb{R}^{3}$ which is a problem which provides many applications in medicine, for example) i.e., $q \in C^{0}\left(S^{1}, \mathbb{R}^{2}\right)$. The motion of the curve is written as $q(s, t)$, where $s \in S$ is the parameter around the curve, and $t \in[0,1]$ is the time parameter. We wish to find the geodesic that takes the "template" curve $\Gamma^{A}$ (parametrised by $\left.q^{A}(s)\right)$ to the "target" curve $\Gamma^{B}$ (parametrised by $q^{B}(s)$ ). However, we do not wish to enforce that any specific
point $q^{A}(s)$ gets mapped to any specific point on $\Gamma^{B}$, so the boundary conditions are

$$
\begin{equation*}
q(s, 0)=q^{A}(\eta(s)), \quad q(s, 1)=q^{B}(s), \tag{6.1}
\end{equation*}
$$

where $\eta \in \operatorname{Diff}_{+}(S)$, the orientation-preserving subgroup of the diffeomorphism group $\operatorname{Diff}(S)$ on $S$. We shall minimise over all reparametrisations $\eta$ to obtain a parameterindependent description of the curves. This is equivalent to finding geodesics on closed components of $C^{0}\left(S, \mathbb{R}^{2}\right) /$ Diff $_{+}(S)$, in particular the component $\operatorname{Emb}(S)^{1} / \operatorname{Diff}(S)$, often referred to as "shape space". However, it is more convenient computationally to work with the unreduced functions $q$, as we shall do here.

Following the methodology of $[57,37,79]$ we constrain the motion of the curve $q(s, t)$ to the action of diffeomorphisms by requiring that

$$
\begin{equation*}
\frac{\partial}{\partial t} q(s, t)=u(q(s, t), t) \tag{6.2}
\end{equation*}
$$

where $u(x, t)$ is a time-parametrised family of vector fields on $\mathbb{R}^{2}$. This guarantees that the topology of the curve is preserved (i.e. there are no overlaps or cavitations). If the boundary conditions 6.1 are satisfied, we say that $u$ describes a path between $\Gamma^{A}$ and $\Gamma^{B}$. We select a function space $B$ for vector fields, and define the distance along the path as

$$
\begin{equation*}
\int_{0}^{1} \frac{1}{2}\|u\|_{B}^{2} \mathrm{~d} t \tag{6.3}
\end{equation*}
$$

For simplicity we assume that $B$ is a Hilbert space and that there exists an operator $A$ such that

$$
\|u\|_{B}^{2}=\langle u, A u\rangle_{L^{2}} .
$$

The shortest path between $\Gamma^{A}$ and $\Gamma^{B}$ is defined by minimising (6.3) over $u$ and $\eta$ subject to (6.2) and the boundary conditions (6.1). See [40] chapter 2 for other

[^1]examples of how to tackle optimisation problems with PDE constraints. To obtain the equations of motion, we introduce Lagrange multipliers $p(s, t)$ (which we call the "momentum") to enforce (6.2), and seek extrema of the action
$$
S=\int_{0}^{1} \frac{1}{2}\|u\|_{B}^{2}+\langle p, \dot{q}-u(q)\rangle \mathrm{d} t
$$

We obtain the weak form of the equations of motion by variational calculus. Define $\delta$ to be such that for a functional $L$,

$$
\delta L[\phi]=\lim _{\epsilon \rightarrow 0}=\frac{L[\phi+\epsilon \delta \phi]-L[\phi]}{\epsilon}
$$

where $\delta \phi$ is an arbitrary test function from the domain of $L$. Then $\delta L$ is the Gâteux derivative of $L$. Applying this definition to $S$ we get that

$$
\begin{aligned}
\delta S= & \delta \int_{0}^{1} \frac{1}{2}\|u\|_{B}^{2}+\langle p, \dot{q}-u(q)\rangle \mathrm{d} t \\
= & \int_{0}^{1}\langle\delta u, A u\rangle_{L^{2}}+\langle\delta p, \dot{q}-u(q)\rangle+\langle p, \delta \dot{q}-\delta u(q)-\nabla u(q) \cdot \delta q\rangle \mathrm{d} t \\
= & \int_{0}^{1}\langle\delta u, A u\rangle_{L^{2}}-\langle p, \delta u(q)\rangle_{L^{2}}+\langle\delta p, \dot{q}-u(q)\rangle \\
& +\left\langle-\dot{p}-(\nabla u(q))^{T} p, \delta q\right\rangle \mathrm{d} t+[\langle p, \delta q\rangle]_{0}^{1}
\end{aligned}
$$

At time $t=1, \delta q=0$, since the boundary condition for $q$ is fixed. At time $t=0$,

$$
\delta q=\delta q^{A} \circ \eta=\frac{\partial q^{A}}{\partial s} \circ \eta \delta \eta=\left.\frac{\partial q}{\partial s}\right|_{t=0} \frac{\nu \circ \eta}{\frac{\partial \eta}{\partial s}}
$$

Here $\delta \eta=\nu \circ \eta$, where $\nu$ is a vector field on $S$ (see [41] for a review of how to perform variational calculus on Lie groups such as $\operatorname{Diff}\left(S^{1}\right)$ ). The boundary term becomes

$$
\begin{equation*}
0=\left\langle\left. p\right|_{t=1}, \frac{\partial q}{\partial s} \frac{\nu \circ \eta}{\frac{\partial \eta}{\partial s}}\right\rangle=\left\langle\left.\left. p\right|_{t=1} \circ \eta^{-1} \cdot \frac{\partial q}{\partial s} \circ \eta^{-1}\right|_{t=0}, \nu(s)\right\rangle \tag{6.4}
\end{equation*}
$$

for an arbitrary test function $\nu(s)$. As discussed in [21], the quantity $p \cdot \partial q / \partial s$ is a conserved quantity of the dynamics (since the action principle is invariant under reparametrisations), and hence we obtain

$$
p \cdot \frac{\partial q}{\partial s}=0, \quad \forall t \in[0,1]
$$

The condition states that the momentum $p$ is normal to the shape, and guarantees that we have obtained the minimum over all reparametrisations.

This all leads to the equations of motion in weak form:

$$
\begin{align*}
\langle\delta u, A u\rangle_{L^{2}}-\langle p, \delta u(q)\rangle_{L^{2}} & =0  \tag{6.5}\\
\int_{0}^{1}\left\langle\delta p, \frac{\partial q}{\partial t}-u(q)\right\rangle \mathrm{d} t & =0  \tag{6.6}\\
\int_{0}^{1}\left\langle p, \frac{\partial \delta q}{\partial t}-(\nabla u(q)) \delta q\right\rangle \mathrm{d} t & =0 \tag{6.7}
\end{align*}
$$

where $\delta p$ and $\delta q$ are space-time test functions, with

$$
w, u \in B, \quad p, \delta p \in L^{2}, \quad q, \delta q \in H^{1}
$$

These equations can be solved numerically using the method described in [21]. We define the mapping $\Psi$ as

$$
\left.q\right|_{t=1}=\Psi\left(\left.p\right|_{t=0},\left.q\right|_{t=0}\right)
$$

where $p$ and $q$ satisfy equations (6.5-6.7), and we seek the solution $\left(p_{0}, \eta\right)$ of

$$
q^{B}(\eta(s))=\Psi\left(p_{0} n, q^{A}\right)
$$

where $p_{0}$ is the normal component of $p$ and $n$ is the normal to the curve, which guarantees the condition

$$
p(s, 0) \cdot \frac{\partial q^{A}}{\partial s}(s)=0
$$

which in turn guarantees that we have obtained the minimum over all reparametrisations. The problem of finding geodesics between $\Gamma^{A}$ and $\Gamma^{B}$ can be solved as a shooting problem for $q$ in which one has to find initial conditions for the normal component of $p$ such that the end condition for $q$ at time $t=1$ is satisfied.

In the case in which one has only observed a finite number of points on $q^{B}$ (e.g. when $q^{B}$ is obtained from a segmentation of a medical image), it is not convenient to
compute with this reparametrised condition. However, since the reparametrisation operation commutes with the forward model (see [22]), we can apply the reparametrisation first, choosing

$$
q(s, 0)=q^{A}\left(\eta^{-1}(s)\right)
$$

We construct $\eta$ from a vector field $\nu(s)$ on $S$, so that

$$
\frac{\partial}{\partial t} \chi(s, t)=\nu(\chi(s, t))=\nu \circ \chi, \quad \chi(s, 0)=s, \quad \eta(s)=\chi(s, 1)
$$

This is to preserve the ordering of the points on $S$ under the reparametrisation. It is then necessary to choose a transformation of $p$ such that (6.5) holds for the reparametrisation $(p, q) \mapsto(\bar{p}, \bar{q})$ where $\bar{q}=q \circ \eta$.

This transformation can easily be obtained by writing the formula for $\bar{u}$ :

$$
\begin{aligned}
\langle w, A u\rangle & =\langle p, w(q)\rangle \\
& =\langle p, w(\bar{q} \circ \eta)\rangle \\
& =\left\langle p \circ \eta \frac{\partial \eta}{\partial s} \circ \eta, w(\bar{q})\right\rangle \\
& =\langle\bar{p}, w(\bar{q})\rangle \\
& =\langle w, A \bar{u}\rangle
\end{aligned}
$$

where $w \in B$ is an arbitrary test function. This gives us the weak form of $\bar{p}$, given by

$$
\langle\bar{p}, \gamma\rangle_{L^{2}}=\left\langle p, \gamma \circ \eta^{-1}\right\rangle_{L^{2}}, \quad \forall w \in B
$$

We write the reparametrisation operator (which does not change the velocity fields and hence does not change the flow map) as

$$
(\bar{p}, \bar{q})=\mathcal{R}(p, q, \nu)
$$

Hence to find the optimum path over all reparametrisations, we solve

$$
q^{B}=\Psi \circ \mathcal{R}\left(p_{0} n, q^{A}, \nu\right)
$$

The normal component of $p_{0}$ then characterises the shape of the target curve $\Gamma^{B}$ relative to the curve $\Gamma^{A}$, whilst the generator variable $\nu$ merely describes the reparametrisation of the target curve which is obtained at the minimum.

### 6.3 Prior Distribution on $\left(p_{0}, \nu\right)$

As in the previous chapters, we will pick prior distributions on the unknowns to enforce regularity conditions. In these other examples, a mean zero Gaussian prior was chosen with covariance operator $\delta A^{-\alpha}=\delta(-\mathcal{P} \Delta)^{-\alpha}$ where $\mathcal{P}$ is the Leray projector onto divergence-free spaces. This was adequate in this context, as we defined our fields to be (among other things) mean zero, meaning that we did not have to consider the variance for the constant Fourier mode.

This is not the case in the scenario presented in this chapter, where our unknown functions are also only one dimensional, and where both the momentum and the reparametrisation can have a constant term in their Fourier expansions. Therefore we need to consider another choice of covariance operator, which is defined for frequency $k=0$. With this in mind, let us consider the Helmholtz operator

$$
\mathcal{H}=\ell I-\Delta
$$

where $\ell \in \mathbb{R}$ defines the length scale. This is a positive definite operator with the same smoothing properties as the Stokes operator. However, unlike the Stokes operator it's inverse is well-defined on constant functions.

Lemma 6.3.1. Functions drawn from $\mathcal{N}\left(0,(-\Delta)^{-\alpha}\right)$ and $\mathcal{N}\left(0, \mathcal{H}^{-\alpha}\right)$ have the same regularity properties. In particular $u \in H^{s}$ if $s<\alpha-\frac{1}{2}$ in one dimension.

Proof. Proof by measure equivalence, given by corollary 1.9.1.

The length scale parameter $\ell \in \mathbb{R}^{+}$allows us to control at which scale the smoothing properties of the Laplacian take effect. With a larger value of $\ell$, a larger value of $|k|$ is required before the effect of the Laplacian becomes dominant. Likewise, as $\ell \rightarrow 0, \mathcal{H} \rightarrow-\Delta$, meaning that the Laplacian is dominant on all scales. The choice of this value, however, does not affect the overall regularity of samples drawn from the distribution $\mathcal{N}\left(0, \mathcal{H}^{-\alpha}\right)$.

### 6.4 The Observational Noise Model

As we have previously shown in the thesis, we assume that observations $y$ of the quantity of interest are noisy in nature, satisfying:

$$
y=\mathcal{G}\left(p_{0}, \nu\right)+\xi, \quad \xi \sim \mathcal{N}(0, \Sigma)
$$

where $\Sigma$ is assumed to be known, and where $\mathcal{G}$ is defined to be the observation operator given by

$$
\mathcal{G}\left(p_{0}, \nu\right)=\left\{\Psi \circ \mathcal{R}\left(p_{0} n, q^{A}\left(s_{i}\right), \nu\right)\right\}_{i=1}^{n}
$$

This allows us to compute the likelihood that $y$ was observed with a given $\left(p_{0}, \nu\right)$ :

$$
\mathbb{P}\left(y \mid p_{0}, \nu\right) \propto \exp \left(-\frac{1}{2}\left\|\mathcal{G}\left(p_{0}, \nu\right)-y\right\|_{\Sigma}^{2}\right)
$$

where $\|x\|_{\Sigma}^{2}:=x^{T} \Sigma^{-1} x$ is the covariance weighted norm. We now have all the components that we require to define the posterior distribution on $\left(p_{0}, \nu\right)$.

### 6.5 The Posterior Distribution

By Bayes law the Radon-Nikodym derivative of the posterior with respect to the prior, assuming $\mu$ is $\mu_{0}$-measurable, is given by a function proportional to the likelihood.

That is, in our case, the measure with Radon-Nikodym derivative, assuming $\mu$ is $\mu_{0^{-}}$ measurable, given by

$$
\begin{equation*}
\frac{d \mu\left(p_{0}, \nu\right)}{d \mu_{0}} \propto \exp \left(-\frac{1}{2}\left\|y-\mathcal{G}\left(p_{0}, \nu\right)\right\|_{\Sigma}^{2}\right) . \tag{6.8}
\end{equation*}
$$

We now need to consider for what values of $\alpha$ do samples from the prior Gaussian distribution $\mu_{0}$ with covariance operator $\mathcal{C}=\mathcal{H}^{-\alpha}$ have sufficient regularity to ensure that $\mu=\mathbb{P}\left(p_{0}, \nu \mid y\right)$ is measurable with respect to $\mu_{0}=\mathbb{P}\left(p_{0}, \nu\right)$. This require analysis of the forward problem, which we will undertake in the next section.

### 6.5.1 Properties of the Observation Operator

In this section we define the observation operator for $\Psi \circ \mathcal{R}\left(p_{0} n, q^{A} \nu\right)$, and prove that it is global Lipschitz in $p_{0}$ and $\nu$.

When the curve is observed from a segmentation from a grey scale image (for example), the observed data is an ordered list of points $\left\{q_{i}\right\}_{i=1}^{n}$. We seek $\left(p_{0}, \nu\right)$ such that

$$
\Psi \circ \mathcal{R}\left(p_{0} n, q^{A}\left(s_{i}\right), \nu\right)=q_{i}, \quad i=1, \ldots, n
$$

with $\left\{s_{i}\right\}_{i=1}^{n}$ is a monotonic and distinct sequence of points in $S$. This guarantees to preserve the ordering since the curve is mapped by a diffeomorphism. Hence we define the observation operator $\mathcal{G}: L^{2} \times H^{3}$ by

$$
\mathcal{G}\left(p_{0}, \nu\right)=\left(\begin{array}{c}
\Psi \circ \mathcal{R}\left(p_{0} n, q^{A}\left(s_{1}\right), \nu\right) \\
\vdots \\
\Psi \circ \mathcal{R}\left(p_{0} n, q^{A}\left(s_{n}\right), \nu\right)
\end{array}\right)
$$

To show that $\mathcal{G}$ is Lipschitz, we first prove existence, uniqueness and Lipschitz continuity for the maps $\Psi$ and $\mathcal{R}$. This is then used to show that the observation operator is Lipschitz continuous with respect to the normal component $p_{0}$ of the initial conditions
$p$ and the generator variable $\nu$. For an arbitrary $u \in B$, we define the time- $t$ flow map $\phi_{t}: \mathbb{R}^{2} \rightarrow \mathbb{R}^{2}$ of $u$, as the solution of the differential equation

$$
\frac{\partial}{\partial t} \phi_{t}(x)=u\left(\phi_{t}(x), t\right), \quad \phi_{0}(x)=x
$$

First we state two lemmas which show that if $B$ is sufficiently smooth, then the flow map is a diffeomorphism with smooth inverse, and that the flow map is Lipschitz. We define the norm

$$
\|u\|_{1, T}=\left(\int_{0}^{T}\|u\|_{B} \mathrm{~d} t\right)
$$

Lemma 6.5.1. If $k \geq 1$ and $B$ is embedded in $C_{0}^{k}\left(\Omega, \mathbb{R}^{k}\right)$, then, for all flows $u \in$ $L^{1}([0, T], \Omega), \Phi$, the corresponding flow map, is $k$-times differentiable. Moreover, there exists constants $C$ such that for all $u \in L^{1}([0, T], \Omega)$,

$$
\sup _{t \in[0, T]}\left\|\phi_{t}\right\|_{C^{k}} \leq C e^{C\|u\|_{1, T}}, \quad \sup _{t \in[0, T]}\left\|\phi_{t}^{-1}\right\|_{C^{k}} \leq C e^{C\|u\|_{1, T}}
$$

Proof. See lemmas 7 and 9 in appendix C of [77].

Lemma 6.5.2. Assume that $B$ is continuously embedded in $C_{0}^{k}(\Omega)$. Let $\phi_{t}$ and $\phi_{t}^{\prime}$ be the time-t flow maps for $u$ and $u^{\prime}$ respectively, where $u, u^{\prime} \in L^{2}([0, T], B)$. Then, for $t \leq T$ we have

$$
\left\|\left(\phi_{t}-\phi_{t}^{\prime}\right)\right\|_{C^{k-1}} \leq C\left\|u-u^{\prime}\right\|_{1, t} e^{C\left(\|u\|_{1, t}+\left\|u^{\prime}\right\|_{1, t}\right)}
$$

Proof. See lemma 11 in Appendix C of [77].

First we prove existence and uniqueness for the equations (6.5-6.7).

Theorem 6.5.1 (Existence and uniqueness of solutions). Let $B$ be continuously embedded into $C_{0}^{1}$. Given initial conditions $p_{0} \in L^{2}, q_{0} \in H^{1}$, equations (6.5-6.7) have a unique solution for all times $t>0$.

Proof. The proof follows the techniques developed in [77].
For any $u \in L^{2}([0, T], B)$ (not necessarily satisfying equation (6.5)), and chosen initial conditions $\left.p\right|_{t=0}=p_{0},\left.q\right|_{t=0}=q_{0}$, we can write the solutions of equations (6.66.7) as

$$
q=\phi_{t} \circ q_{0}, \quad p=\left(\left(\nabla \phi_{t}^{-1}\right) \circ q\right) \cdot p_{0}
$$

This means that estimates of $p$ and $q$ can be written entirely in terms of estimates of $\phi_{t}$.

For $u \in L^{2}([0, T], B)$, we define the function $\Omega: B \rightarrow B$ by

$$
\Omega(u)=u-\int_{S} p G(x-q) \mathrm{d} s
$$

When $\Omega(u)=0$ for all $t<T$ then we have a solution of equations (6.5-6.7).
First we show that $\Omega$ maps $B$ into $B$.

$$
\begin{aligned}
\|\Omega(u)\|_{B} & \leq \sup _{\|w\|_{B}=1}\left\langle u-\int p G(x-q) \mathrm{d} s, w\right\rangle_{B} \\
& =\sup _{\|w\|_{B}=1}\left(\langle u, w\rangle_{B}-\langle p, w(q)\rangle_{L^{2}}\right) \\
& \leq \sup _{\|w\|_{B}=1}\left(\|u\|_{B}\|w\|_{B}+\|p\|_{L^{2}} C\|w\|_{\infty}\right) \\
& \leq \sup _{\|w\|_{B}=1}\left(\|u\|_{B}\|w\|_{B}+\|p\|_{L^{2}} C\|w\|_{B}\right) \\
& =\|u\|_{B}+C\|p\|_{L^{2}} \\
& =\|u\|_{B}+C\left\|\left(\nabla \phi_{t}^{-1}\right) \circ q \cdot p_{0}\right\|_{L^{2}} \\
& \leq\|u\|_{B}+C\left\|\nabla \phi_{t}^{-1}\right\|_{\infty}\left\|p_{0}\right\|_{L^{2}} \\
& \leq\|u\|_{B}+C e^{C\|u\|_{1, T}}\left\|p_{0}\right\|_{L^{2}}
\end{aligned}
$$

where we have used lemma 6.5.1. Hence, the $B$-norm of $\Omega$ is bounded provided that $p_{0} \in L^{2}$.

To show that the solution exists and is unique, we show that $\Omega(u)$ is a contraction mapping in $B$. For $u$ and $u^{\prime}$ in $B$, we define $\phi_{t}$ and $\phi_{t}^{\prime}$ as the time- $t$ flow maps of $u$ and $u^{\prime}$ respectively. Furthermore, we obtain $p, q$ and $p^{\prime}, q^{\prime}$ from $\phi_{t}$ and $\phi_{t}^{\prime}$ respectively. We write

$$
\Omega(u)-\Omega\left(u^{\prime}\right)=u-u^{\prime}-\int_{S}\left(p-p^{\prime}\right) G(x-q) \mathrm{d} s+\int_{S} p^{\prime}\left(G(x-q)-G\left(x-q^{\prime}\right)\right) \mathrm{d} s
$$

Then

$$
\begin{aligned}
\left\|\Omega(u)-\Omega\left(u^{\prime}\right)\right\|_{B} \leq & \sup _{\|w\|_{B}=1}\left\langle w, u-u^{\prime}-\int_{S}\left(p-p^{\prime}\right) G(x-q) \mathrm{d} s\right. \\
& \left.+\int_{S} p^{\prime}\left(G(x-q)-G\left(x-q^{\prime}\right)\right) \mathrm{d} s\right\rangle_{B} \\
= & \sup _{\|w\|_{B}=1}\left(\left\langle w, u-u^{\prime}\right\rangle_{B}+\left\langle A w, \int_{S}\left(p-p^{\prime}\right) G(x-q) \mathrm{d} s\right\rangle_{L^{2}}\right. \\
& \left.+\left\langle A w, \int_{S} p^{\prime}\left(G(x-q)-G\left(x-q^{\prime}\right)\right) \mathrm{d} s\right\rangle_{L^{2}}\right) \\
= & \sup _{\|w\|_{B}=1}\left(\left\langle w, u-u^{\prime}\right\rangle_{B}+\left\langle p-p^{\prime}, w(q)\right\rangle_{L^{2}}\right. \\
& \left.+\left\langle p^{\prime}, w(q)-w\left(q^{\prime}\right)\right\rangle_{L^{2}}\right) \\
\leq & \sup _{\|w\|_{B}=1}\left(\|w\|_{B}\left\|u-u^{\prime}\right\|_{B}+\left\|p-p^{\prime}\right\|_{L^{2}} C\|w\|_{\infty}\right. \\
& \left.+\left\|p^{\prime}\right\|_{L^{2}} C\|\nabla w\|_{\infty}\left\|q-q^{\prime}\right\|_{L^{2}}\right) \\
\leq & \sup _{\|w\|_{B}=1}\left(\|w\|_{B}\left\|u-u^{\prime}\right\|_{B}+\left\|p-p^{\prime}\right\|_{L^{2}} C\|w\|_{B}\right. \\
& \left.+\left\|p^{\prime}\right\|_{L^{2}} C\|w\|_{B}\left\|q-q^{\prime}\right\|_{L^{2}}\right) \\
\leq & \left\|u-u^{\prime}\right\|_{B}+C\left\|p-p^{\prime}\right\|_{L^{2}}+C\left\|p^{\prime}\right\|_{L^{2}}\left\|q-q^{\prime}\right\|_{L^{2}} .
\end{aligned}
$$

It remains to find bounds for $\left\|p-p^{\prime}\right\|_{L^{2}}$ and $\left\|q-q^{\prime}\right\|_{L^{2}}$.
For the bounds on $q-q^{\prime}$ we have

$$
\begin{aligned}
\left\|q-q^{\prime}\right\|_{L^{2}} & =\left\|\phi_{t}\left(q_{0}\right)-\phi_{t}^{\prime}\left(q_{0}\right)\right\|_{L^{2}} \\
& \leq C\left\|u-u^{\prime}\right\|_{1, t} e^{C\left(\|u\|_{1, t}+\left\|u^{\prime}\right\|_{1, t}\right) T}
\end{aligned}
$$

where we have used lemma 6.5.2. To bound $p-p^{\prime}$ we have

$$
\begin{aligned}
\left\|p-p^{\prime}\right\|_{L^{2}} & =\left\|\left(\left(\nabla \phi_{t}^{-1}\right) \circ q-\left(\nabla \phi_{t}^{\prime-1}\right) \circ q^{\prime}\right) \cdot p_{0}\right\|_{L^{2}} \\
& \leq C\left\|u-u^{\prime}\right\|_{1, t} e^{C\left(\|u\|_{1, t}+\left\|u^{\prime}\right\|_{1, t}\right) T}\left\|p_{0}\right\|_{L_{2}} .
\end{aligned}
$$

Hence we have

$$
\left\|\Omega(u)-\Omega\left(u^{\prime}\right)\right\|_{1, T} \leq C\left\|u-u^{\prime}\right\|_{1, T} e^{C T\left(\|u\|_{1, t}+\left\|u^{\prime}\right\|_{1, t}\right)},
$$

where the constants $C$ only depend on $\left\|q_{0}\right\|_{L^{2}}$ and $\left\|p_{0}\right\|_{L^{2}}$. Hence, $\Omega$ is a contraction mapping in $B$ for all $T$ sufficiently small.

Next we prove that there is in fact a unique solution for all $T<\infty$. To do this, we define $T_{0}$ to be the largest value such that $\Omega_{T}$ has a unique fixed point $u_{T}$ for all $T<T_{0}$. Aiming for a contradiction, we assume that $T_{0}<\infty$.

Choose $T<T_{0}$, and $\epsilon>0$. Given $u^{\prime} \in L^{2}([0, \epsilon], B)$, we define $v \in L^{2}([0, T+$ $\epsilon], B$ ) by

$$
v(x, t)=\left\{\begin{array}{rr}
u_{T}(x, t) & t<T \\
v(x, t-T) & t>T
\end{array}\right.
$$

The proof of a unique solution for $t<T+\epsilon$ can be obtained using the same fixed point argument, provided that $\epsilon<\epsilon_{0}$ where $\epsilon_{0}$ is a function of $\|q\|_{L^{2}}$ and $\|p\|_{L^{2}}$ at time $T$. These themselves can be bounded by a function of $\|u\|_{1, T}$ multiplied by the initial values $\left\|q_{0}\right\|_{L^{2}}$ and $\left\|p_{0}\right\|_{L^{2}}$ respectively. However, since $\|u\|_{B}^{2}$ is conserved along solution trajectories (it is the conserved Hamiltonian), this means that $\epsilon_{0}$ can be uniformly bounded for $T<T_{0}$, hence a contradiction. This means that a unique solution exists for all $t$.

Theorem 6.5.2. We fix the initial conditions (template) for $q^{A} \in C^{1}$. Let $B$ be continuously embedded into $C_{0}^{2}$. Let $\left(\bar{p}_{0}, \bar{q}_{0}\right)=\mathcal{R}\left(p_{0}, q^{A}, \nu\right)$ and $\left(\bar{p}_{0}^{\prime}, \bar{q}_{0}^{\prime}\right)=\mathcal{R}\left(p_{0}^{\prime}, q^{A}, \nu^{\prime}\right)$.

Then let ( $u, q, p$ ) and ( $u^{\prime}, q^{\prime}, p^{\prime}$ ) be solutions to equations (6.5-6.7) with initial conditions $\bar{p}_{0} \in L^{2}$ and $\bar{p}_{0}^{\prime} \in L^{2}$ respectively. Then, at time $t=1$, there exists a constant $C$ which depends on $\left\|\bar{q}_{0}\right\|_{\infty},\left\|\bar{q}_{0}^{\prime}\right\|_{\infty},\left\|\bar{p}_{0}\right\|_{\infty},\left\|\bar{p}_{0}^{\prime}\right\|_{\infty}$ only, such that

$$
\left\|q-q^{\prime}\right\|_{\infty} \leq C\left(\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty}\right)
$$

and hence $\Psi$ is Lipschitz continuous.

Proof. To show that the solution depends continuously on the initial conditions for $p$, consider two solutions ( $q, p, u$ ) and ( $q^{\prime}, p^{\prime}, u^{\prime}$ ) with initial conditions $\bar{p}_{0}$ and $\bar{p}_{0}^{\prime}$ for $p$ respectively, with initial conditions $\bar{q}_{0}$ and $\bar{q}_{0}^{\prime}$ for $q$. At time $t$ the difference between $u$ and $u^{\prime}$ in $B$ is

$$
\begin{aligned}
\left\|u-u^{\prime}\right\|_{B}= & \left\|\int p(s, t) G(x-q(s, t)) \mathrm{d} s-\int p^{\prime}(s, t) G\left(x-q^{\prime}(s, t)\right) \mathrm{d} s\right\|_{B} \\
\leq & \sup _{|w|=1}\left\langle\int p(s, t) G(x-q(s, t)) \mathrm{d} s-\int p^{\prime}(s, t) G\left(x-q^{\prime}(s, t)\right) \mathrm{d} s, w\right\rangle \\
\leq & \sup _{|w|=1}\left(\langle p, w(q)\rangle-\left\langle p^{\prime}, w\left(q^{\prime}\right)\right\rangle\right) \\
\leq & \sup _{|w|=1}\left(\left\langle p-p^{\prime}, w(q)\right\rangle+\left\langle p^{\prime}, w(q)-w\left(q^{\prime}\right)\right\rangle\right) \\
\leq & \sup _{|w|=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}}\|w(q)\|_{L_{2}}+\left\|p^{\prime}\right\|_{L^{2}}\left\|w(q)-w\left(q^{\prime}\right)\right\|_{L_{2}}\right) \\
\leq & \sup _{|w|=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} C\|w\|_{\infty}+\left\|p^{\prime}\right\|_{L^{2}} C\|\nabla w\|_{\infty}\left\|q-q^{\prime}\right\|_{L^{2}}\right) \\
\leq & \sup _{|w|=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} C\|w\|_{B}+\left\|p^{\prime}\right\|_{L^{2}} C\|w\|_{B}\left\|q-q^{\prime}\right\|_{L^{2}}\right) \\
= & C\left\|p-p^{\prime}\right\|_{L^{2}}+C\left\|p^{\prime}\right\|_{L^{2}}\left\|q-q^{\prime}\right\|_{L^{2}} \\
= & C\left\|\left(\nabla \phi_{t}^{-1} \circ q\right) \cdot \bar{p}_{0}-\left(\nabla \phi_{t}^{\prime-1} \circ q^{\prime}\right) \cdot \bar{p}_{0}^{\prime}\right\|_{L^{2}} \\
& +C\left\|\nabla \phi_{t}^{\prime-1} \circ q^{\prime} \cdot \bar{p}_{0}^{\prime}\right\|_{L^{2}}\left\|\phi_{t} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}^{\prime}\right\|_{L^{2}}
\end{aligned}
$$

$$
\begin{aligned}
= & C\left\|\left(\nabla \phi_{t}^{-1} \circ q\right) \cdot\left(\bar{p}_{0}-\bar{p}_{0}^{\prime}\right)+\left(\left(\nabla \phi_{t}^{-1} \circ q\right)-\left(\nabla \phi_{t}^{\prime-1} \circ q^{\prime}\right)\right) \cdot \bar{p}_{0}^{\prime}\right\|_{L^{2}} \\
& +C\left\|\nabla \phi_{t}^{\prime-1} \circ q^{\prime} \cdot \bar{p}_{0}^{\prime}\right\|_{L^{2}}\left\|\phi_{t} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}+\phi_{t}^{\prime} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}^{\prime}\right\|_{L^{2}} \\
= & C \|\left(\nabla \phi_{t}^{-1} \circ q\right) \cdot\left(\bar{p}_{0}-\bar{p}_{0}^{\prime}\right) \\
& +\left(\left(\nabla \phi_{t}^{-1} \circ q\right)-\left(\nabla \phi_{t}^{-1} \circ q^{\prime}\right)+\left(\nabla \phi_{t}^{-1} \circ q^{\prime}\right)-\left(\nabla \phi_{t}^{\prime-1} \circ q^{\prime}\right)\right) \cdot \bar{p}_{0}^{\prime} \|_{L^{2}} \\
& +C\left\|\nabla \phi_{t}^{\prime-1} \circ q^{\prime} \cdot \bar{p}_{0}^{\prime}\right\|_{L^{2}}\left\|\phi_{t} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}+\phi_{t}^{\prime} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}^{\prime}\right\|_{L^{2}} \\
\leq & C\left\|\nabla \phi_{t}^{-1}\right\|_{\infty}\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left(\left\|\phi_{t}^{-1}\right\|_{C^{2}}\left\|q-q^{\prime}\right\|_{L^{2}}+\left\|\phi_{t}^{-1}-\phi_{t}^{\prime-1}\right\|_{C^{1}}\right)\left\|\bar{p}_{0}^{\prime}\right\|_{L^{2}} \\
& +C\left\|\nabla \phi_{t}^{\prime-1}\right\|_{\infty}\left\|\bar{p}_{0}^{\prime}\right\|_{L^{2}}\left(\left\|\phi_{t}-\phi_{t}^{\prime}\right\|_{\infty}+\left\|\phi_{t}^{\prime}\right\|_{C^{1}}\left\|_{q_{0}}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}\right) \\
\leq & C\left\|\nabla \phi_{t}^{-1}\right\|_{\infty}\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left(\left\|\phi_{t}^{-1}\right\|_{C^{2}}\left(\left\|\phi_{t}-\phi_{t}^{\prime}\right\|_{\infty}+\left\|\phi_{t}^{\prime}\right\|_{C^{1}}\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}\right)\right. \\
& \left.+\left\|\phi_{t}^{-1}-\phi_{t}^{\prime-1}\right\|_{C^{1}}\right)\left\|\bar{p}_{0}^{\prime}\right\|_{L^{2}}+C\left\|\nabla \phi_{t}^{\prime-1}\right\|_{\infty}\left\|\bar{p}_{0}^{\prime}\right\|_{L^{2}}\left(\left\|\phi_{t}-\phi_{t}^{\prime}\right\|_{\infty}\right. \\
& \left.+\left\|\phi_{t}^{\prime}\right\|_{C^{1}}\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}\right) \\
\leq & C\left(\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}+\left\|u-u^{\prime}\right\|_{1, T}\right)
\end{aligned}
$$

by using lemmas 6.5 .1 and 6.5 .2 with $k=2$. Grönwall's lemma implies that

$$
\left\|u-u^{\prime}\right\|_{B} \leq C\left(\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}\right)
$$

Then the $\infty$-norm distance between $q$ and $q^{\prime}$ at time $T$ is

$$
\begin{aligned}
\left\|q-q^{\prime}\right\|_{\infty} & =\left\|\phi_{t} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}^{\prime}\right\|_{\infty} \\
& =\left\|\phi_{t} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}+\phi_{t}^{\prime} \circ \bar{q}_{0}-\phi_{t}^{\prime} \circ \bar{q}_{0}^{\prime}\right\|_{\infty} \\
& \leq\left\|\phi_{t}-\phi_{t}^{\prime}\right\|_{\infty}+\left\|\phi_{t}^{\prime}\right\|_{C^{1}}\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty} \\
& \leq C\left(\left\|u-u^{\prime}\right\|_{1, T}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty}\right) \\
& \leq C\left(\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty}\right) \\
& \leq C\left(\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty}\right) .
\end{aligned}
$$

Lemma 6.5.3. Let $D$ be continuously embedded in $C^{1}(S, \mathbb{R})$, and the initial template parametrisation be $q^{A} \in C^{1}\left(S, \mathbb{R}^{2}\right)$. Then for $\nu, \nu^{\prime} \in D$ and $p_{0}, p_{0}^{\prime} \in L^{2}(S, \mathbb{R})$, there exists a constant $C$ dependent only on $\|\nu\|_{D},\left\|\nu^{\prime}\right\|_{D},\left\|p_{0}\right\|_{L^{2}},\left\|p_{0}^{\prime}\right\|_{L^{2}}$ and $\left\|q^{A}\right\|_{C^{1}}$, and where $\left(\overline{p_{0}}, \overline{q_{0}}\right)=\mathcal{R}\left(p_{0}, q^{A}, \nu\right)$ and $\left({\overline{p_{0}}}^{\prime}, \overline{q_{0}}\right)=\mathcal{R}\left(p_{0}^{\prime}, q^{A}, \nu^{\prime}\right)$,

$$
\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}+\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty} \leq C\left(\left\|p_{0}-p_{0}^{\prime}\right\|_{L^{2}}+\left\|\nu-\nu^{\prime}\right\|_{D}\right) .
$$

Hence $\mathcal{R}$ is Lipschitz continuous.

Proof. We first attempt to find a bound for $\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty}$. The reparametrisation formula for $q$ is:

$$
\bar{q}=q \circ \eta,
$$

where $\bar{q}$ is the reparametrised $q$. Using lemma 6.5 .2 with $k=1$, we get that

$$
\begin{aligned}
\left\|\bar{q}_{0}-\bar{q}_{0}^{\prime}\right\|_{\infty} & =\left\|q^{A} \circ \eta-q^{A} \circ \eta^{\prime}\right\|_{\infty} \\
& \leq\left\|q^{A}\right\|_{C^{1}}\left\|\eta-\eta^{\prime}\right\|_{\infty} \\
& \leq C\left\|q^{A}\right\|_{C^{1}}\left\|\nu-\nu^{\prime}\right\|_{1,1} \\
& =C\left\|q^{A}\right\|_{C^{1}}\left\|\nu-\nu^{\prime}\right\|_{D},
\end{aligned}
$$

since $\nu, \nu^{\prime}$ are constant in time. We now attempt to bound $\left\|\bar{p}_{0}-\bar{p}_{0}^{\prime}\right\|_{L^{2}}$. The reparametrisation formula for $p$ is defined weakly:

$$
\langle\bar{p}, \gamma\rangle_{L^{2}}=\left\langle p, \gamma \circ \eta^{-1}\right\rangle_{L^{2}},
$$

for all vector-valued functions $\gamma: S^{1} \mapsto \mathbb{R}^{2}$.

Hence, to bound $\left|\bar{p}-\bar{p}^{\prime}\right|$ :

$$
\begin{aligned}
\left\|\bar{p}-\bar{p}^{\prime}\right\|_{L^{2}}= & \sup _{\|\gamma\|_{L^{2}}=1}\left\langle\bar{p}-\bar{p}^{\prime}, \gamma\right\rangle_{L^{2}} \\
= & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\langle p, \gamma \circ \eta^{-1}\right\rangle_{L^{2}}-\left\langle p^{\prime}, \gamma \circ\left(\eta^{\prime}\right)^{-1}\right\rangle_{L^{2}}\right) \\
= & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\langle p-p^{\prime}, \gamma \circ \eta^{-1}\right\rangle_{L^{2}}+\left\langle p^{\prime}, \gamma \circ \eta^{-1}-\gamma \circ\left(\eta^{\prime}\right)^{-1}\right\rangle_{L^{2}}\right) \\
\leq & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}}\left\|\gamma \circ \eta^{-1}\right\|_{L^{2}}+\left\|p^{\prime}\right\|_{L^{2}}\left\|\gamma \circ \eta^{-1}-\gamma \circ\left(\eta^{\prime}\right)^{-1}\right\|_{L^{2}}\right) \\
= & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\int_{S^{1}}\left(\gamma \circ \eta^{-1}\right)^{2} \mathrm{~d} s}\right. \\
& +\left\|p^{\prime}\right\|_{L^{2}} \sqrt{\left.\int_{S^{1}}\left(\gamma \circ \eta^{-1}-\gamma \circ\left(\eta^{\prime}\right)^{-1}\right)^{2} \mathrm{~d} s\right)} \\
= & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\int_{S^{1}}\left(\gamma\left(s^{\prime}\right)\right)^{2} \frac{\partial \eta}{\partial s^{\prime}}\left(s^{\prime}\right) \mathrm{d} s^{\prime}}\right. \\
& +\left\|p^{\prime}\right\|_{L^{2}} \sqrt{\left.\int_{S^{1}}\left(\gamma \circ \eta^{-1}\right)^{2}-2 \gamma \circ \eta^{-1} \gamma \circ\left(\eta^{\prime}\right)^{-1}+\left(\gamma \circ\left(\eta^{\prime}\right)^{-1}\right)^{2} \mathrm{~d} s\right)} \\
= & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\int_{S^{1}}\left(\gamma\left(s^{\prime}\right)\right)^{2} \frac{\partial \eta}{\partial s^{\prime}}\left(s^{\prime}\right) \mathrm{d} s^{\prime}}\right. \\
& +\left\|p^{\prime}\right\|_{L^{2}}\left(\int_{S^{1}}(\gamma)^{2} \frac{\partial \eta}{\partial s^{\prime}} \mathrm{d} s^{\prime}-2 \int_{S^{1}} \gamma\left(\gamma \circ\left(\eta^{\prime}\right)^{-1} \circ \eta\right) \frac{\partial \eta}{\partial s^{\prime}} \mathrm{d} s^{\prime}\right. \\
& \left.\left.+(\gamma)^{2} \frac{\partial \eta^{\prime}}{\partial s^{\prime}} \mathrm{d} s^{\prime}\right)^{1 / 2}\right) \\
\leq & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\int_{S^{1}}\left(\gamma\left(s^{\prime}\right)\right)^{2} \frac{\partial \eta}{\partial s^{\prime}}\left(s^{\prime}\right) \mathrm{d} s^{\prime}}\right. \\
& +\left\|p^{\prime}\right\|_{L^{2}}\left(\|\gamma\|_{L^{2}}^{2}\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}+2\|\gamma\|_{L^{2}}\left\|\gamma \circ\left(\left(\eta^{\prime}\right)^{-1} \circ \eta\right)\right\|_{L^{2}}\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}\right. \\
& \left.\left.\left.+\|\gamma\|_{L^{2}}^{2}\left\|\frac{\partial \eta^{\prime} \|}{\partial s}\right\|_{L^{2}}\right)\right)^{1 / 2}\right)
\end{aligned}
$$

$$
\begin{aligned}
\leq & \sup _{\|\gamma\|_{L^{2}}=1}\left(\left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\int_{S^{1}}\left(\gamma\left(s^{\prime}\right)\right)^{2} \frac{\partial \eta}{\partial s^{\prime}}\left(s^{\prime}\right) \mathrm{d} s^{\prime}}\right. \\
& +\left\|p^{\prime}\right\|_{L^{2}}\left(\|\gamma\|_{L^{2}}^{2}\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}+2\|\gamma\|_{L^{2}}^{2}\left\|\frac{\partial}{\partial s}\left(\left(\eta^{\prime}\right)^{-1} \circ \eta\right)\right\|_{L^{2}}\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}\right. \\
& \left.\left.+\|\gamma\|_{L^{2}}^{2}\left\|\frac{\partial \eta^{\prime}}{\partial s}\right\|_{L^{2}}\right)^{1 / 2}\right) \\
\leq & \left\|p-p^{\prime}\right\|_{L^{2}} \sqrt{\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}} \\
& +\left\|p^{\prime}\right\|_{L^{2}}\left(\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}+2\left\|\frac{\partial}{\partial s}\left(\left(\eta^{\prime}\right)^{-1} \circ \eta\right)\right\|_{L^{2}}\left\|\frac{\partial \eta}{\partial s}\right\|_{L^{2}}+\left\|\frac{\partial \eta^{\prime}}{\partial s}\right\|_{L^{2}}\right)^{1 / 2} .
\end{aligned}
$$

To bound $\left(\left(\eta^{\prime}\right)^{-1} \circ \eta\right)$, we note that this map can be obtained by first applying the time-1 flow map of $\nu$, and then the time-1 flow map of $-\nu^{\prime}$. An application of lemma 6.5.1 gives

$$
\left\|\left(\left(\eta^{\prime}\right)^{-1} \circ \eta\right)\right\|_{C^{k}} \leq C\left\|\nu-\nu^{\prime}\right\|_{D},
$$

where $C$ depends only on $\|\nu\|_{D}$. Moreover, since

$$
\|\eta\|_{C^{1}} \leq C\|\nu\|_{D}
$$

the result follows.

Theorem 6.5.3. Let $B$ be continuously embedded in $C_{0}^{2}$, and $D$ be continuously embedded in $C_{0}^{1}$. Then for $p_{0} \in L^{2}$ and $\nu \in D$, the observation operator $\mathcal{G}$ is Lipschitz continuous.

Proof. The observation operator is the composition of three Lipschitz continuous functions, $\mathcal{R}$ (by lemma 6.5.3, $\Psi$ (by lemma 6.5.2, and the projection

$$
q \rightarrow\left(\begin{array}{c}
q\left(s_{1}\right) \\
\vdots \\
q\left(s_{n}\right)
\end{array}\right)
$$

Corollary 6.5.1. Let $\mu_{0}\left(p_{0}, \nu\right)=\mathcal{N}\left(0, \delta_{1} \mathcal{H}_{1}^{-\alpha_{1}}\right) \times \mathcal{N}\left(0, \delta_{2} \mathcal{H}_{2}^{-\alpha_{2}}\right)$ for $\alpha_{1}>\frac{1}{2}, \alpha_{2}>2$, $\delta_{1}, \delta_{2} \neq 0$, and where $\mathcal{H}_{i}=\left(\ell_{i} I-\Delta\right)$ where $\ell_{i} \in \mathbb{R}^{+}$for $i \in\{1,2\}$. Then $\mathcal{G}$ is measurable with respect to $\mu_{0}$, and the posterior measure $\mu$ is absolutely continuous with respect to $\mu_{0}$, with Radon-Nikodym derivative given by (6.8).

Proof. Result follows by theorems 1.10.2 and 1.8.1, lemmas 1.8.1, 6.3.1, 6.5.1 and theorem 6.5.3.

### 6.6 RWMH with Deterministic Burn-In

Using the results in the previous section, we are now able to use our RWMH method as previously described in section 1.12 .2 to sample from well-defined posterior measures on function spaces, conditioned on our observations of a shape.

One problem with this particular application, however, is that the forward model is more expensive computationally than the simple linear PDE model of Stokes flow. This means that, in particular, the usual burn-in period where we start the chain at zero (which will usually be in the tales of the distribution of interest) and iterate until we feel the chain has entered stationarity, is also rather drawn out and expensive.

To combat this, deterministic methods can be utilised to start the chain closer to the region of state space which has higher probability density. In particular, we might want to solve the Tikhonov regularisation problem which relates to maximising the probability density of the posterior measure. That is finding the solution to:

$$
\min _{p_{0} \in L^{2}, \nu \in H^{2}} L\left(p_{0}, \nu\right)=\min _{p_{0} \in L^{2}, \nu \in H^{2}} \frac{1}{2}\left\|\mathcal{G}\left(p_{0}, \nu\right)-y\right\|_{\Sigma}^{2}+\frac{1}{2}\left\|\left(p_{0}, \nu\right)\right\|_{\mu_{0}}^{2}
$$

where $\|(\cdot, \cdot)\|_{\mu_{0}}$ is the equivalent penalty term to that induced in the posterior measure by the choice of prior measure $\mu_{0}$, or the Cameron-Martin norm corresponding to
$\mu_{0}$. For more details on the relationship between the Bayesian approach and Tikhonov regularisation, see section 1.11. In particular, if (as is the case for our purposes) $\mu_{0}\left(p_{0}, \nu\right)=\mathcal{N}\left(0, \delta_{1} \mathcal{H}^{-\alpha_{1}}\right) \times \mathcal{N}\left(0, \delta_{2} \mathcal{H}^{-\alpha_{2}}\right)$, then

$$
\left\|\left(p_{0}, \nu\right)\right\|_{\mu_{0}}^{2}=\sum_{k} \delta_{1}\left|p_{k}\right|^{2}\left(\ell_{1}-|k|^{2}\right)^{\alpha_{1}}+\delta_{2}\left|\nu_{k}\right|^{2}\left(\ell_{2}-|k|^{2}\right)^{\alpha_{2}}
$$

Various hill-climbing methods can be used to try to find local minima of this quantity. Most of these methods incorporate gradient information of some type to attempt to search for the minima in appropriate directions. These methods include steepest descent and conjugate gradient among others.

The method that we will utilise in our numerics is called the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method[16, 71]. This method uses the gradient information from the last two states in the chain to approximate a Hessian matrix, which is then used to choose an appropriate direction in which to search for the local minimum. When the value of $\nabla L$ falls below a given threshold the algorithm is assumed to have converged and terminates. In more detail (adapted from [81]):

From an initial guess $x_{0}$ and an approximate Hessian matrix $B_{0}$ the following steps are repeated until $x$ converges to the solution.
while $\left|\nabla L\left(x_{k}\right)\right|>$ tol do
3: Obtain a direction $p_{k}$ by solving: $B_{k} p_{k}=-\nabla L\left(x_{k}\right)$
4: Perform a line search to find an acceptable step size $\alpha_{k}$ in the direction found in the first step, then update $x_{k+1}=x_{k}+\alpha_{k} p_{k}$

Set $s_{k}=\alpha_{k} p_{k}$
$y_{k}=\nabla L\left(x_{k+1}\right)-\nabla L\left(x_{k}\right)$
$B_{k+1}=B_{k}+\frac{y_{k} y_{k}^{T}}{y_{k}^{T} s_{k}}-\frac{B_{k} s_{k}\left(B_{k} s_{k}\right)^{T}}{s_{k}^{T} B_{k} s_{k}}$
end while

By running this deterministic solver before the adaptive burn-in period (where different proposal step sizes are tried to attempt to optimise the MCMC algorithm), we can hope that we are not in the tails of the posterior distribution, and therefore that the burn-in process can be completed with far fewer evaluations of the forward model.

From this point onwards the same RWMH MCMC algorithm can be implemented. That is, given a currently accepted state in the chain $x_{n}=\left(p_{0}, \nu\right)_{n}$, we propose a new state $z$ where

$$
\begin{equation*}
z=\left(1-\beta^{2}\right)^{1 / 2} x_{n}+\beta w, \quad w \sim \mu_{0} \tag{6.9}
\end{equation*}
$$

Recall that the acceptance probability in this algorithm is given by

$$
\alpha(x, z)=\exp (\Phi(x)-\Phi(z)),
$$

where $\Phi(\cdot)=\frac{1}{2}\|\mathcal{G}(\cdot, \cdot)-y\|_{\Sigma}^{2}$.

### 6.7 Numerical Approximation of $\mathcal{G}$

For the purposes of this thesis, we will not go into the details of how the forward model is approximated numerically, but will instead treat the process as a "black box" (particularly as the software for calculating this forward model was provided for the author). This particle mesh method is drawn directly from [21].

### 6.8 General Setup

In this section we describe the scenarios in which we will numerically test the algorithm in the next section. The first thing to address is our choice of template shape $\Gamma^{A}$ and the parametrisation that we use for this shape $q^{A}(s)$. Since we are only considering closed curves, it seems natural to choose a circle for this to keep things as simple as
possible. We also wish to choose a nice smooth parametrisation for this shape, centred in the middle of our domain $\mathbb{T}^{2}=[0,2 \pi)^{2}$ so we pick

$$
\begin{equation*}
q^{A}(s)=(\cos (s)+\pi, \sin (s)+\pi), \quad s \in[0,2 \pi) \tag{6.10}
\end{equation*}
$$

We will engage in simulation studies in which the data is itself produced by employing the numerical simulation of a forward PDE model. This will involve picking a $\left(p_{0}, \nu\right)$ from the prior and using the numerical approximation of the observation operator to create our data. We will also use data from recognisable shapes to show that when there are enough observations the algorithm is able to explore the probability measure whose mean can recreate a shape close to that from which the data was taken.

In all the numerics, we assume that the noise $\xi$ through which we make the observations

$$
y=\mathcal{G}\left(p_{0}, \nu\right)+\xi, \quad \xi \sim \mathcal{N}(0, \Sigma)
$$

has a diagonal covariance matrix $\Sigma=\sigma^{2} I$ for some $\sigma>0$.
The prior distributions on $p_{0}$ and $\nu$ are $\mathcal{N}\left(0, \delta_{1} \mathcal{H}^{-\alpha_{1}}\right)$ and $\mathcal{N}\left(0, \delta_{1} \mathcal{H}^{-\alpha_{1}}\right)$ respectively, with $\alpha_{1}=0.55, \alpha_{2}=3.05, \delta_{1}=40$ and $\delta_{2}=0.3$. Note that these parameters are sufficient to ensure that corollary 6.5 .1 holds.

In terms of the approximation of the forward model in the algorithm, 50 time steps are used. The curves themselves are approximated by 100 points, and with a $64 \times 64$ grid approximating the underlying velocity field. The values of $s_{j}$ at which observations are made are given by $\left\{\frac{2 \pi s_{j}}{N}\right\}_{j=0}^{N-1}$. These parameters are used in the model to create the data, and in the implementation of the statistical algorithm.

### 6.9 Numerical Results

We now present some numerical results to show that the algorithm is successfully drawing samples from the posterior distribution.

### 6.9.1 Posterior Consistency

It seems reasonable to expect that as we increase the amount of informative data that we are using in our inference, the closer our posterior mean will be to the functions that created the data, and that at the same time the uncertainty in that estimation will decrease. In this set of numerical experiments, we take a draw $\left(p_{0}, \nu\right)$ from the prior measure, and using our approximation of the forward model, create data $y$ such that

$$
y=\left\{q^{B}\left(\frac{2 \pi n}{N}\right)\right\}_{n=0}^{N-1}+\xi, \quad \xi \sim \mathcal{N}(0, \Sigma)
$$

with increasing $N$.
In the following graphs we look at the marginal distributions for the two lowest frequency Fourier modes of both $p_{0}$ and $\nu$ as estimated by our MCMC method, for $N=10,50,100$. We choose the low frequencies as these are the Fourier modes which are most informed by our observations.

Figures 6.1 and 6.2 show how as we increase the number of observations, the marginal distributions on these two particular Fourier modes become increasingly peaked and centred near the value of that Fourier which was present in the $p_{0}$ which created the data.

Similarly Figures 6.3 and 6.4 show that as we increase the number of observations, the marginal distributions on these two particular Fourier modes of the reparametrisation function $\nu$ also become increasingly peaked on the values which were present in the function that created the data.


Figure 6.1: Marginal distributions on $p_{0}(0)$ with increasing numbers of observations


Figure 6.2: Marginal distributions on $p_{1}(0)$ with increasing numbers of observations


Figure 6.3: Marginal distributions on $\nu_{0}$ with increasing numbers of observations


Figure 6.4: Marginal distributions on $\nu_{1}$ with increasing numbers of observations


Figure 6.5: Distributions of acceptance probabilities in the MCMC method with increasing numbers of observations, average acceptance probability tuned to $\approx 25 \%$

Figure 6.5 shows the distribution of acceptance probabilities for the different Markov chains with varying numbers of observations. Since the step size $\beta$ has been tuned is such a way that the average acceptance probability is approximately $25 \%$, we would expect these distributions to be similar, and this is borne out in this figure.

Figure 6.6 shows the distribution of $\Phi(\cdot)=\frac{1}{2}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$ in each of the Markov chains with varying numbers of observations. These distribution differ in this way since the dimension of the observation spaces are different.

Figure 6.7 shows the same plot, but with the norm divided by the number of observations made. This shows that as the number of observations increase, the contribution to the norm from each observation decreases.


Figure 6.6: Distribution of $\Phi(\cdot)=\frac{1}{2}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$ in the MCMC method with increasing numbers of observations


Figure 6.7: Distribution of $\frac{1}{N} \Phi(\cdot)=\frac{1}{2 N}\|\mathcal{G}(\cdot)-y\|_{\Sigma}^{2}$ in the MCMC method with increasing numbers of observations

### 6.9.2 "Real Life" Data

In this section we attempt to use the algorithm on a recognisable shape for which we do not know the value of $\left(p_{0}, \nu\right)$ which morphs our initial template parametrisation $q^{A}(\cdot)$ given by (6.10) into the shape for which we have data.

## The Christmas Tree

In our first experiment, we attempt to find which pair of functions $\left(p_{0}, \nu\right)$ are required to morph our initial circle into a shape close to the observations, plotted in Figure 6.8.


Figure 6.8: Christmas tree shape data

The main issue with this data is that it contains several sharp edges. Since the initial shape is completely smooth, it would require a discontinuity in the flow fields to cause a sharp edge in the final shape. Since we have enforced that the deformation is continuous, this cannot occur. This complexity was verified in the deterministic burn-in process. The solver found a local minima of the functional, such that when the MCMC process was invoked, the chain quickly moved away from this state.


Figure 6.9: Marginal distributions on $p_{0}(0)$ with increasing numbers of observations


Figure 6.10: Marginal distributions on $p_{1}(0)$ with increasing numbers of observations

Figures 6.9 and 6.10 show the marginal distributions of the first two Fourier modes of $p_{0}$ in the posterior. It can clearly be seen that these distributions are nonGaussian. This is due in no small part to the spiky nature of the data. Since the result of the deformations is always smooth, only one edge of a sharp corner in the data can be satisfied close to that corner at a time, giving us a multi-modal distribution. Given many sharp corners, as in this example, we get quite a complex distribution, which took a great deal of samples to converge.


Figure 6.11: Marginal distributions on $\nu_{0}$ with increasing numbers of observations

In comparison Figures 6.11 and 6.12 show very smooth marginal distributions on the Fourier modes of $\nu$. The distribution of the observation points is not affected in the same way by these singularities as the flow field itself. The next two plots represent the forward model of the mean values in that distribution.

As can be seen particularly in Figure 6.14, the effect of the assimilation process has very much been to smooth off the harsh edges in the data.


Figure 6.12: Marginal distributions on $\nu_{1}$ with increasing numbers of observations


Figure 6.13: Representations of the mean functions


Figure 6.14: Mesh deformation for the forward model of the mean functions. Original data included for comparison in the left hand frame

## The Face

In an attempt to use some data that is a little closer in shape to a circle than the previous example, we now consider some data of the outline of a face.

There are still some sharp corners in this data, around the ears and at the hair parting, so we would expect to see some significant smoothing of these features again, due to our choice of $\Gamma^{A}$.

Figures 6.16 and 6.17 show the marginal distributions of the first two Fourier modes of $p_{0}$ in the posterior. These distributions would appear to be slightly nonsymmetrical, with Figure 6.16 having a heavier tail to the right than the left, and Figure 6.17 having the opposite.

As in the previous example, in comparison Figures 6.18 and 6.19 show marginal distributions on the Fourier modes of $\nu$ which on first glance would appear to be symmetric and approximately Gaussian. The following two plots represent the forward model


Figure 6.15: Face shaped data


Figure 6.16: Marginal distributions on $p_{0}(0)$ with increasing numbers of observations


Figure 6.17: Marginal distributions on $p_{1}(0)$ with increasing numbers of observations


Figure 6.18: Marginal distributions on $p_{0}(0)$ with increasing numbers of observations


Figure 6.19: Marginal distributions on $p_{1}(0)$ with increasing numbers of observations of the mean values in that distribution.

As can be seen particularly in Figure 6.21, the effect of the assimilation process has very much been to smooth off the harsh edges in the data, due to the mismatch in the template shape and the data. This would be an issue if it wasn't for the fact that in the application that we are aspiring to, we would be attempting to match images of two organs, which would have the same kind of topology and smoothness as each other on the whole.

### 6.10 Conclusions

By careful analysis of the forward problem, we have been able to formulate a wellposed Bayesian inverse problem regarding shape matching of a curve with a set of noisy observations of another. We have shown that the likelihood function is continuous with respect to a space that has full measure with respect to a specified choice of Gaussian prior measure. Using this, we have shown how to draw samples from well defined


Figure 6.20: Representations of the mean functions


Figure 6.21: Mesh deformation for the forward model of the mean functions. Original data included for comparison in the left hand frame
posterior distributions on function space using the RWMH MCMC sampler. We have then implemented this algorithm, and presented briefly some initial numerics. In these numerics we have shown that, using a draw from the prior to create our target data, that as we increase the number of informative observations, the marginal distributions on both $p_{0}$ and $\nu$ become increasingly peaked on the functions that created the data.

Moreover, we have presented two examples of matching a circle with some "real life" data. We have discussed the effect that a mismatch in the differentiability of the template shape and target data can cause on the posterior.

Although only still in the early stages of development, this chapter presents a Bayesian approach to shape registration which takes into account that observations of organs' shapes are not exact, and gives a distribution on the shortest distance in shape space between the template shape to the observed shape. The sampling method is framed on function space in line with the philosophy of this whole thesis, and so is robust under refinements of discretization.

Since we already have at our disposal an implementation of the adjoint problem so that we can calculate the gradient of the observation operator (which we use for the deterministic burn-in), we can very simply adapt the MCMC method from RWMH to MALA, which includes gradient information in the proposal distribution. This would increase the efficiency of the algorithm markedly. Further analytical results would be needed, however, to ensure that the gradient of the observation operator is continuous on a space which has full measure with respect to an appropriately chosen prior distribution.

Certainly another excellent test of this algorithm would be to attempt to use it on some genuine images taken and segmented from scans of bodily organs.

Another idea would be to also make the problem translation, rotation and scale invariant, so that any misalignment of the imaging equipment has a negligible effect on
the results. This would involve adding parameters into the state space to allow for these types of operations.

One might also consider the case where the template shapes themselves are only noisily observed, so that we are trying to find a distribution on the length of geodesic paths in shape space between two noisily observed shapes. This might be a little closer to the scenario in the application, where the template shapes themselves will simply have come from scans of organs with a particular condition.

## Bibliography

[1] Gnu scientific library reference manual.
http://www.gnu.org/software/gsl/manual/html_node/
Random-number-generator-algorithms.html, February 2007.
[2] T. Albrecht, A. Dedner, M. Luethi, and T. Vetter. Curvature guided level set registration using adaptive finite elements. Submitted to Medical Image Analysis.
[3] J. Anderson. An ensemble adjustment kalman filter for data assimilation. Monthly Weather Review, 129:2884-2903, 2001.
[4] A. Apte, M. Hairer, A.M. Stuart, and J. Voss. Sampling the posterior: An approach to non-gaussian data assimilation. Physica D: Nonlinear Phenomena, 230:50-64, 2006.
[5] A. Apte, C.K.R.T. Jones, and A.M. Stuart. A bayesian approach to lagrangian data assimilation. Tellus, 60:336-347, 2008.
[6] A. Apte, C.K.R.T. Jones, A.M. Stuart, and J. Voss. Data assimilation: Mathematical and statistical perspectives. International Journal of Numerical Methods in Fluids, 56:1033-1046, 2008.
[7] H. Aref and S. Balachandar. Chaotic advection in a stokes flow. Physics of Fluids, 29, 1986.
[8] R. Azencott, R. Glowinski, and A.M. Ramos. A controllability approach to shape identification. Mathematics Letters, 21:861-865, 2008.
[9] T. Bengtsson, P. Bickel, and B. Li. Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. IMS Collections, 2, 2008. Probability and Statistics: Essays in Honor of David A. Freedman.
[10] A.F. Bennett. Inverse Modelling of the Ocean and Atmosphere. Cambridge University Press, 2002.
[11] A. Beskos, G. Roberts, A. Stuart, and J. Voss. Mcmc methods for diffusion bridges. Stochastics and Dynamics, 8(3):319-350, 2008.
[12] A. Beskos and A. Stuart. Mcmc methods for sampling function space. ICIAM plenary lecture, 2007.
[13] P. Bickel, B. Li, and T. Bengtsson. Sharp failure rates for the bootstrap particle filter in high dimensions. IMS Collections, 3:318-329, 2008. Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh.
[14] M. Bocquet. Grid resolution dependence in the reconstruction of an atmospheric tracer source. Nonlinear Processes in Geophysics, 12:219-234, 2005.
[15] M. Bocquet. Inverse modelling of atmospheric tracers: non-gaussian methods and second-order sensitivity analysis. Nonlinear Processes in Geophysics, 15:127-143, 2008.
[16] C.G. Broyden. The convergence of a class of double-rank minimization algorithms. J. Inst. Maths Applics, 6:76-90, 1970.
[17] R.E. Carlson and F.N. Fritsch. Monotone piecewise bicubic interpolation. Siam J. Numer. Anal, 22(2), April 1985.
[18] S. Castellari, A. Griffa, T.M. Ögökmen, and P. Poulain. Prediction of particle trajectories in the adriatic sea using lagrangian data assimilation. Journal of Marine Systems, 29:33-50, 2001.
[19] S. Chib and E. Greenberg. Understanding the metropolis-hastings algorithm. The American Statistician, 49(4):327-335, 1995.
[20] S.C. Chu and F.T. Metcalf. On grönwall's inequality. Proceedings of the American Mathematical Society, 18:439-440, 1967.
[21] C. J. Cotter. The variational particle-mesh method for matching curves. J. Phys. A: Math. Theor., 41:344003, 2008.
[22] C. J. Cotter and D. D. Holm. Geodesic boundary value problems with symmetry. submitted, 2009.
[23] S. Cotter, M. Dashti, J. Robinson, and A. Stuart. Mcmc methods on function space and applications to fluid mechanics. In progress.
[24] S. Cotter, M. Dashti, J. Robinson, and A. Stuart. Bayesian inverse problems for functions and applications to fluid mechanics. Inverse Problems, 25, 2009.
[25] S. Cotter, M. Dashti, and A. Stuart. Approximation of bayesian inverse problems for pdes. SIAM Journal on Numerical Analysis, 2010. To appear.
[26] M. Dashti and J.C. Robinson. A simple proof of uniqueness of the particle trajectories for solutions of the Navier-Stokes equations. Nonlinearity, 22:735-746, 2009.
[27] K.W.Morton D.F.Mayers. Numerical Solution of Partial Differential Equations. Cambridge University Press, 1994.
[28] Y. Dong, Y. Gu, and D.S. Oliver. Sequential assimilation of 4d seismic data for reservoir description using the ensemble kalman filter. Journal of Petroleum Science and Engineering, 53:83-99, 2006.
[29] M. Droske, W. Ring, and M. Rumpf. Mumford-shah based registration: a comparison of a level set and a phase field approach. Computing and Visualization in Science, 12:101-114, 2009.
[30] C.L. Farmer. Algorithms for Approximation, chapter Bayesian field theory applied to scattered data interpolation and inverse problems, pages 147-166. Springer-Verlag Heidelberg, 2007.
[31] B. Farrell and P. Ioannou. State estimation using a reduced-order kalman filter. Journal of the Atmospheric Sciences, 58:3666-3680, 2001.
[32] R. Farwig, H. Kozono, and H. Sohr. The helmholtz decomposition in arbitrary unbounded domains - a theory beyone $l^{2}$. Proceedings of Equadiff-11, pages 7785, 2005.
[33] M. Fisher and D.J. Lary. Lagrangian four dimensional variational data assimilation of chemical species. Q.J.R. Meteorol. Soc., 121(527):1681-1704, 1995.
[34] M. Fisher, J. Nocedal, Y. Trémolet, and S.J. Wright. Data assimilation in weather forecasting: a case study in pde-constrained optimization. Optimization and Engineering, 10, 2009.
[35] D. Fujiwara and H. Morimoto. An $l_{r}$ theory of the helmholtz decomposition of vector fields. J. Fac. Sci. Univ. of Tokyo, 1977.
[36] J. Glaunes, A. Trouvé, and L. Younes. Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub manifolds matching. Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition.
[37] J. Glaunes, A. Trouvé, and L. Younes. Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition, volume 2, pages 712-718, 2004.
[38] Y. Gu and D.S. Oliver. The ensemble kalman filter for continuous updating of reservoir simulation models. Journal of Energy Resources Technology, 128(1):7987, 2006.
[39] Y. Hino, S. Murakami, T. Naito, and N. Van Minh. A variation-of-constants formula ofr abstract functional differential equations in the phase space. Journal of Differential Equations, 179:336-355, 2002.
[40] M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich. Optimization with PDE constraints. Springer, 2008.
[41] D. D. Holm, J. E. Marsden, and T. S. Ratiu. The Euler-Poincaré equations and semidirect products with applications to continuum theories. Adv. in Math., 137:181, 1998. http://xxx.lanl.gov/abs/chao-dyn/9801015.
[42] M.Y. Hussaini and T.A. Zang. Spectral methods in fluid dynamics. Ann. Rev. Flui Mech., 19:339-367, 1987.
[43] K. Ide, L. Kuznetsov, and C.K.R.T. Jones. Lagrangian data assimilation for point vortex systems. Journal of Turbulence, 3, 2002.
[44] Armin Iske and Trygve Randen, editors. Mathematical Methods and Modelling in Hydrocarbon Exploration and Production, volume 7 of Mathematics in Industry. Springer-Verlag, Berlin, 2005.
[45] E. Somersalo J. Kaipio. Statistical and Computational Inverse Problems. Springer, 2004.
[46] R.E. Kalman. A new approach to linear filtering and prediction problems. Journal of basic Engineering, 1960.
[47] E. Kalnay. Atmospheric Modelling, Data Assimilation and Predictability. Cambridge University Press, 2003.
[48] P. Krause. The diffusion kernel filter. J. Stat. Phys., 134:365-380, 2009.
[49] A.S. Lawless, S. Gratton, and N.K. Nichols. An investigation of incremental 4d-var using non-tangent linear models. Q. J. R. Meteorol. Soc., 131:459-476, 2005.
[50] L.C.Evans. Partial Differential Equations, volume 19 of Graduate Studies in Mathematics. American Mathematical Society, 1998.
[51] J. Li and D. Xiu. A generalized polynomial chaos based ensemble kalman filter with high accuracy. Journal of computational physics, 228(15):5454-5469, 2009.
[52] D.M. Livings, S.L. Dance, and N.K. Nichols. Unbiased ensemble square root filters. Physica D, 237:1021-1028, 2008.
[53] R. Lumpkin and Pierre Flament. Lagrangian statistics in the central north pacific. Journal of Marine Systems 29, 2001.
[54] A.J. Majda, J. Harlim, and B. Gershgorin. Mathematical strategies for filtering turbulent dynamical systems. Discrete and Continuous Dynamical Systems, 27, 2010.
[55] G. Marsaglia. The Marsaglia Random Number CDROM including the Diehard Battery of Tests of Randomness. http://www.csis.hku.hk/ diehard/cdrom/, 1995.
[56] Y. Marzouk and D. Xiu. A stochastic collocation approach to inference in inverse problems. Communications in computational physics, 6(4):826-847, 2009.
[57] M. T. Miller and L. Younes. Group actions, homeomorphisms, and matching: A general framework. International Journal of Computer Vision, 41:61-84, 2001.
[58] G. Nævdal, L.M. Johnsen, S.I. Aanonsen, and E.H. Vefring. Reservoir monitoring and continuous model updating using ensemble kalman filter. SPE Journal, 10(1):66-74, 2005.
[59] G. Nævdal, T. Mannseth, and E.H. Vefring. Near-well reservoir monitoring through ensemble kalman filter. SPE/DOE Improved Oil Recovery Symposium, 2002.
[60] H.J. Nussbaumer. Fast Fourier transform and convolution algorithms. SpringerVerlag, 2 edition, 1982.
[61] G. Da Prato and J. Zabczyk. Stochastic Equations in Infinite Dimensions. Cambridge University Press, 1992.
[62] A. Rao, R. Chandrashekara, G.I. Sanchez-Ortiz, R. Mohiaddin, P. Aljabar, J.V. Hajnal, B.K. Puri, and D. Rueckert. Spatial transformation of motion and deformation fields using nonrigid registration. IEEE Transactions on Medical Imaging, 23, 2004.
[63] G. O. Roberts, A. Gelman, and W. R. Gil. Weak convergence and optimal scaling of random walk metropolis algorithms. Annals of Applied Probability, 1997.
[64] G.O. Roberts. Optimal scaling of discrete approximations to langevin diffusions. J. R. Statist. Soc. B, 60:255-268, 1998.
[65] G.O. Roberts and J. Rosenthal. Optimal scaling for various metropolis-hastings algorithms. Statistical Science, 2001.
[66] James C. Robinson. Infinite-Dimensional Dynamical Systems. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2001.
[67] M. Rumpf. chapter Variational methods in image matching and motion extraction in level set and PDE based reconstruction methods: applications to inverse problems and image processing. 2009.
[68] H. Salman. A hybrid grid/particle filter for lagrangian data assimilation. i: Formulation the passive scalar approximation. Quarterly Journal of the Royal Meteorological Society, 134:1539-1550, 2008.
[69] H. Salman. A hybrid grid/particle filter for lagrangian data assimilation. ii: Application to a model vortex flow. Quarterly Journal of the Royal Meteorological Society, 134:1551-1565, 2008.
[70] M.Frigo S.G.Johnson. Fftw homepage. http://www.fftw.org/ , 2010.
[71] D.F. Shanno. On broyden-fletcher-goldfarb-shanno method. J. Optimiz. Theory Appl., 46, 1985.
[72] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson. Obstacles to high-dimensional particle filtering. Monthly Weather Review, 136:46294640, 2008.
[73] A.M. Stuart. Inverse Problems: A Bayesian Perspective. Acta Numerica, 2010.
[74] R. Temam. Navier-Stokes equations and nonlinear functional analysis. Siam, 1983.
[75] W.J. Thistleton, J.A. Marsh, K. Nelson, and C. Tsallis. Generalized box-müller method for generating q-gaussian random deviates. IEEE transactions on information theory, 53(12):4805-4810, 2007.
[76] E. Thönnes. Monte Carlo Methods, Lecture notes for CY904/ST407. The University of Warwick, 2009.
[77] A. Trouve and L. Younes. Local geometry of deformable templates. SIAM Journal on Mathematical Analysis, 37(1):17-59, 2005.
[78] M. Vaillant and J. Glaunes. Surface matching via currents. IPMI, pages 381-392, 2005.
[79] Marc Vaillant and Joan Glaunes. Surface matching via currents. In IPMI, pages 381-392, 2005.
[80] E. Vanmarcke. Random Fields: Analysis and Synthesis. The MIT Press, 1983.
[81] Wikipedia. BFGS method.
http://en.wikipedia.org/wiki/BFGS.method, 2010.
[82] Wikipedia. Bicubic interpolation.
http://en.wikipedia.org/wiki/Bicubic_interpolation, 2010.
[83] Wikipedia. Bilinear interpolation.
http://en.wikipedia.org/wiki/Bilinear_interpolation, 2010.
[84] Wikipedia. Grönwall's inequality.
http://en.wikipedia.org/wiki/Gronwall's_inequality, 2010.
[85] B. Wirth, L. Bar, M. Rumpf, and G. Sapiro. Geodesics in shape space via variational time discretization. volume 5681, pages 288-302. Proceedings of the 7 th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR'09), 2009.
[86] Benedikt Wirth, Leah Bar, Martin Rumpf, and Guillermo Sapiro. Geodesics in shape space via variational time discretization. In EMMCVPR '09: Proceedings of the 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition, pages 288-302. Springer-Verlag, 2009.
[87] L. Wu, V. Mallet, M. Bocquet, and B. Sportisse. A comparison study of data assimilation algorithms for ozone forecasts. J. Geophys. Res., 113, 2008.
[88] D. Xiu. Fast numerical method for stochastic computations: A review. Communications in computational physics, 5(2-4):242-272, 2008.
[89] D. Xiu and G.E. Karniadakis. Modeling uncertainty in flow simulations via generalized polynomial chaos. Journal of computational physics, 187(1):137-167, 2003.
[90] M. Zupanski, M. Navon, and D. Zupanski. The maximum likelihood ensemble filter as a non-differentiable minimization algorithm. Quarterly Journal of the Royal Meteorological Society, 134:1039-1050, 2008.


[^0]:    ${ }^{1}$ Here by number of Fourier modes, we mean the dimension of the Fourier space approximation, ie

[^1]:    ${ }^{1}$ Space of smooth maps between $S$ and $R^{2}$ which are homeomorphisms onto their images

