THE UNIVERSITY OF
SYDNEY

COPYRIGHT AND USE OF THIS THESIS

This thesis must be used in accordance with the provisions of the Copyright Act 1968.

Reproduction of material protected by copyright may be an infringement of copyright and copyright owners may be entitled to take legal action against persons who infringe their copyright.

Section 51 (2) of the Copyright Act permits an authorized officer of a university library or archives to provide a copy (by communication or otherwise) of an unpublished thesis kept in the library or archives, to a person who satisfies the authorized officer that he or she requires the reproduction for the purposes of research or study.

The Copyright Act grants the creator of a work a number of moral rights, specifically the right of attribution, the right against false attribution and the right of integrity.

You may infringe the author's moral rights if you:

- fail to acknowledge the author of this thesis if you quote sections from the work
- attribute this thesis to another author
- subject this thesis to derogatory treatment which may prejudice the author's reputation

For further information contact the University's Director of Copyright Services

sydney.edu.au/copyright

Kalman-Filter-Based EEG Source Localization

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

by

Matthew James Barton

School of Physics
The University of Sydney
Australia

November 2014

Declaration of originality

To the best of my knowledge, this thesis contains no copy or paraphrase of work published by another person, except where duly acknowledged in the text. This thesis contains no material which has been previously presented for a degree at The University of Sydney or any other university.

Matthew James Barton

Included papers and attribution

Listed below are the papers on which Chapters 2 and 3 are based.

Chapter 2 **Evaluating the Performance of Kalman-Filter-Based EEG Source Localization**
M. J. Barton, P. A. Robinson, S. Kumar, A. Galka, H. F. Durrant-Whyte, J. Guivant, and T. Ozaki
Published in *IEEE Transactions on Biomedical Engineering*, **56** 122–136 (2009)

I was chiefly responsible for this work, with an overall contribution of about 80%.

Chapter 3 **Kalman-Filter-Based EEG Source Localization with a Spatially-Varying Process Model: 1D Simulations**
M. J. Barton, P. A. Robinson, J. Guivant, P. M. Drysdale, A. Galka, and H. F. Durrant-Whyte
To be submitted to *Journal of Neuroscience Methods*

I was chiefly responsible for this work, with an overall contribution of about 90%.

Acknowledgements

I wish to thank my supervisor Prof. Peter Robinson for his guidance, enthusiasm, and unstinting support. He was always ready with excellent advice, constructive insight, and made sure I stayed on track despite the occasional threat of derailment. I thank him for teaching me a great deal about neuroscience, physics, and research, and for giving me confidence in my scientific abilities.

I would like to thank my current and former associate supervisors, Dr Jose Guivant and Dr Suresh Kumar for their assistance, particularly with the Kalman filtering aspects of this thesis. Their interest and support is greatly appreciated.

I would also like to thank Prof. Hugh Durrant-Whyte and Dr Andreas Galka for expressing an interest in my work and taking the time to answer my questions, no matter how elementary, about Kalman filtering and EEG source localization.

Thanks to all the people in the brain modeling group, it has been a great pleasure to work with you all. In particular, I wish to thank Dr Chris Rennie and Dr Peter Drysdale for their assistance with all manner of technical questions. Thanks also goes to my office mates for making the PhD experience all the more entertaining.

Finally, I would like to thank my wife Renee, my parents, my family, and my in-laws. This thesis would not have been completed without their constant love, support, and patience.

As long as our brain is a mystery, the universe, the reflection of the structure of the brain will also be a mystery.

Santiago Ramón y Cajal

Summary

Functional neuroimaging aims to noninvasively characterize the dynamics of the distributed neural networks that mediate brain function in healthy and pathological states. Over recent decades a number of imaging modalities have emerged that allow brain dynamics to be probed on different spatiotemporal scales. This thesis focuses on one such technique, electroencephalographic (EEG) source localization, which solves the EEG inverse problem to estimate the location, magnitude, and time course of the neuronal sources that produce the observed scalp voltages. In particular, this work investigates using the Kalman filter (KF) — a flexible and robust model-based estimation algorithm — to solve this inverse problem.

Chapter 1 begins by introducing EEG source localization and its core components. This is followed by an overview of the KF and subsequently a discussion of how this algorithm can naturally combine these components to solve the EEG inverse problem.

A recently developed spatially whitened Kalman filter (SWKF) that solves the EEG inverse problem for a realistic head model is introduced in Chapter 2. This filter employs a spatial whitening transformation to reduce its computational burden. A telegrapher's equation describes the dynamics of the current dipoles, or more precisely that of their spatial Laplacian. This equation represents a suitable starting point for modeling large-scale brain activity. Likelihood maximization is used to fit spatially uniform model parameters and noise covariances to simulated and clinical EEGs. The resulting inverse solutions are found to accurately reconstruct the underlying source dynamics. This study also applied standard diagnostic tests to objectively evaluate KF performance. These tests compute the statistical properties of the innovation sequences and subsequently identify spatial variation in filter performance which could potentially be improved by spatially-varying model parameters.

Chapter 3 investigates the SWKF using one-dimensional (1D) simulations to

reduce the complexity of the inverse problem. Simulated EEG is generated using a telegrapher's equation and a simplified volume conductor model. Motivated by the findings of Chapter 2, two model parameters in both the simulated data and SWKF are given spatial profiles of a simple functional form (Gaussian) which better reflects true brain dynamics. Unlike Chapter 2, the optimization step here explicitly constrains model parameters to ranges consistent with the process model being a telegrapher's equation. These constraints ensure the estimated model parameters maintain a clear biophysical interpretation. The Courant condition, which places an upper bound on observable wave velocity values for a given spatiotemporal grid, is also introduced. For the purposes of comparison, inverse solutions are also computed using the optimal linear Kalman filter (LKF). This study finds that both filters produce accurate state estimates, with the SWKF and LKF exhibiting similar performance. Standard diagnostic tests show that both filters are well-tuned for all scenarios in the 1D study. Spatially varying parameter profiles are able to be correctly identified from the datasets with transient dynamics, while parameter estimates for the driven datasets are less reliable, because they are degraded by the unmodeled drive term. Temporal undersampling, which occurs when the Courant condition imposes an upper bound on wave velocity estimates that is below the wave velocity's true value, is also found to distort parameter estimates.

Chapter 4 returns to the whole-brain EEG inverse problem and applies several features of the 1D simulation study to the SWKF investigated in Chapter 2. Following the method of Chapter 3, spatially varying model parameters of a known functional form (sinusoidal) are introduced into the simulated EEG and SWKF to reproduce the typical anterior to posterior variation of the alpha rhythm. Compared to Chapter 2, more realistic simulated EEG is generated which exhibits wave-like patterns and spatially varying dynamics. As in Chapter 3, the optimization step constrains model parameters to ranges consistent with the telegrapher's equation definition, and also imposes the Courant condition on wave velocity estimates. State estimation was again found to be reliable for both simulated and clinical EEG. However, the introduction of spatially varying parameters did not improve state estimation for any dataset. In fact, parameter estimation was unreliable for both simulated and clinical EEG. A consistent finding across all datasets is the underestimation of the wave velocity. Several likely contributing factors are identified, including: the use of low density EEG recordings; and the effect of the Courant condition.

Approaches to overcome this limitation are discussed.

Finally, Chapter 5 summarizes the main findings of the thesis and outlines potential directions for future work, including emerging applications for Kalman filtering in computational neuroscience.

Contents

List of Abbreviations

1D One-dimensional

2D Two-dimensional

3D Three-dimensional

AC Autocorrelation

AIC Akaike information criterion

AR Autoregressive

BC Boundary condition

BOLD Blood oxygen level dependent

CI Confidence intervals

CU Simulated EEG with driven dynamics and spatially uniform model parameters used specifically to investigate temporal undersampling

DOT Diffuse Optical Tomography

DU Simulated EEG with driven dynamics and spatially uniform model parameters

DV Simulated EEG with driven dynamics and spatially varying model parameters

ECG Electrocardiography

ECoG Electrocorticography

EEG Electroencephalography

EIT Electrical impedance tomography

EKF Extended Kalman filter

fMRI Functional magnetic resonance imaging

FOCUSS Focal underdetermined system solution

FWHM Full-width half-maximum

GARCH Generalized autoregressive conditional heteroskedasticity

IED Interelectrode distance

KF Kalman filter

KS Kolmogorov-Smirnov

LFM Lead field matrix

LKF Linear Kalman filter

LORETA Low-resolution electromagnetic tomography

MEG Magnetoencephalography

MRI Magnetic resonance imaging

NLS Noise level statistic

OECD Organisation for Economic Co-operation and Development

PET Positron emission tomography

PSD Power spectral density

rCBF Regional cerebral blood flow

RMSE Root mean squared error

SE Spectral entropy

SNR Signal to noise ratio

SPECT Single photon emission computed tomography

SQP Sequential quadratic programming

SQUID Superconducting quantum interference device

SWKF Spatially whitened Kalman filter

TU Simulated EEG with transient dynamics and spatially uniform model parameters

TV Simulated EEG with transient dynamics and spatially varying model parameters

UKF Unscented Kalman filter

VES Voxel-to-electrode separation

Chapter 1

Overview

The human brain is a truly remarkable structure. Approximately 10^{11} neurons form up to 10^{15} synapses [63, 74, 130] to produce a multiscale dynamic network which performs a multitude of complex and diverse functions ranging from maintaining homeostasis (e.g., circadian rhythms and thermoregulation) through to conscious experience and its associated cognitive abilities — many of which are unique to humans.

The primary goal of neuroscientific research is to determine how the human brain performs this constellation of roles. An important outcome of an increasing understanding of brain function is the ability to: (i) potentially prevent; (ii) more reliably diagnose; and (iii) better treat disorders of the central nervous system — both neurological and psychiatric. This will help reduce the considerable burden of disease on both the individual and society associated with these conditions. For example, major depression is estimated to have the second highest disease burden as measured using disability-adjusted life years in Organisation for Economic Co-operation and Development (OECD) countries, being responsible for approximately 7% of the total burden [100]. The ability to record large-scale brain activity has been recognized as a vital component of any effective attempt to understand, and subsequently combat, complex maladies like mental illness. The increasing importance of whole-brain functional data is supported by the growing evidence that the brain's cognitive functions arise from the dynamic interactions of distributed brain areas operating in large-scale networks [17]. Dysfunction of these networks has already been identified in psychopathological states such as depressive illness [111]. These large-scale interactions can be probed through functional neuroimaging, which is a

group of techniques able to noninvasively record neural activity, either directly or indirectly, across the entire brain.

This thesis focuses on the functional imaging modality known as electroencephalographic (EEG) source localization, which through solving the EEG inverse problem is able to estimate the underlying neuronal activity responsible for a given EEG recording. Over the past two decades numerous techniques have been applied to the EEG inverse problem. The work presented here investigates using the Kalman Filter (KF), which is a model-based estimation algorithm successfully applied to similar inverse problems in other fields. Throughout this thesis, the spatially whitened KF (SWKF), which is a KF implementation specifically designed for the EEG inverse problem, is applied to simulated and clinical EEG for both realistic head models and a simplified one-dimensional (1D) cortex. The performance of the filter is then studied and recommendations for improvements are made and subsequently investigated.

This overview chapter provides the background information upon which the work described in this thesis is built. Section 1.1 introduces functional neuroimaging and summarizes the key modalities, including electroencephalography. An overview of EEG source localization is provided in Sec. 1.2. In Sec. 1.3, the KF is discussed. Section 1.4 brings together the two previous topics to introduce KF-based EEG source localization. In Sec. 1.5 the work undertaken in the remainder of this thesis is outlined.

1.1 Functional Neuroimaging

Up until the 1970s investigation of human brain function was largely restricted to clinical studies of patients with brain lesions (secondary to disease or injury), or recording or stimulating brain activity during neurosurgical procedures. However, the field completely changed when advances in physics, mathematics and computing led to the development of modern three-dimensional (3D) imaging techniques, and allowed brain structure and function to be probed with unprecedented fidelity [52]. These advances gave rise to the field of functional neuroimaging which uses various neuroimaging modalities to noninvasively characterize the dynamics of the large-scale distributed neural networks that mediate brain function in healthy and pathological states. Over recent decades a number of imaging modalities have

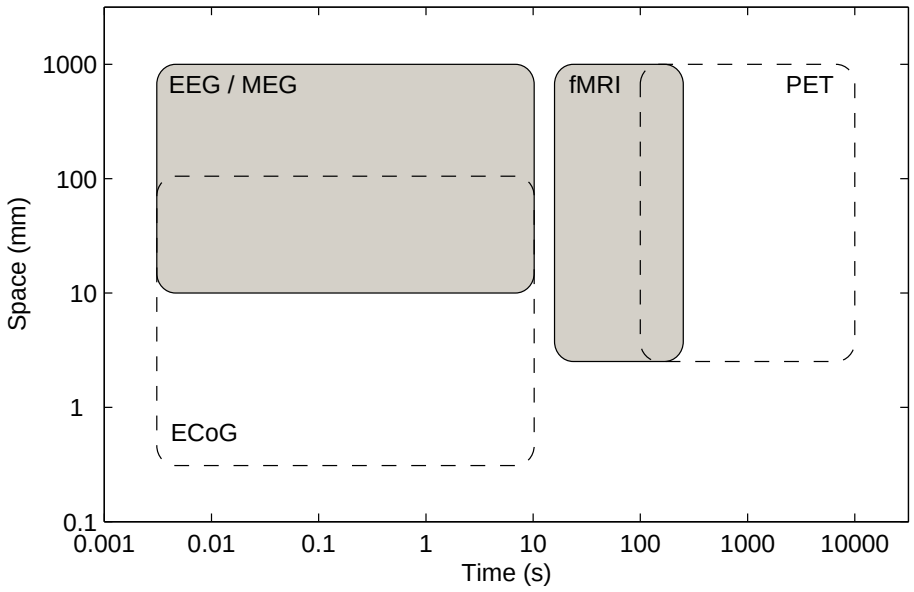

Figure 1.1: Spatiotemporal resolution of some widely used functional neuroimaging modalities. The resolution of electrocorticography (ECoG) — an invasive technique that provides better spatial resolution by placing electrodes directly onto the cerebral cortex — is shown here for comparison.

emerged that allow brain dynamics to be probed on different spatiotemporal scales. Some commonly used techniques include functional magnetic resonance imaging (fMRI), positron emission tomography (PET), and source localization techniques using EEG and magnetoencephalographic (MEG) data. A short description of the these modalities is now provided and their spatiotemporal resolutions are summarized in Fig. 1.1.

1.1.1 Functional Magnetic Resonance Imaging

Functional magnetic resonance imaging is arguably the most well known and widely used functional neuroimaging modality. It is based upon the observation that blood flow within brain tissue is correlated with neural activity [86]. In the early 1990s it was discovered that these changes in blood flow could be detected using MRI, through a phenomenon known as the blood oxygen level dependent (BOLD) effect. This effect arises from changes in blood flow to a particular brain region altering the ratio of oxyhemoglobin to deoxyhemoglobin; since deoxyhemoglobin is paramagnetic it induces local distortions in the MRI signal which can be imaged with

a spatial resolution of several millimeters. The time course of the BOLD signal, which has a temporal resolution of approximately 2 s, can be sampled for the entire brain volume by taking a series of images at short intervals (≈ 2 s). The resulting volumetric time series is then analyzed to identify brain regions whose activity is correlated with the experimental conditions. The ability of fMRI to image the entire brain in vivo, means it can observe the distributed networks involved in cognitive processes. However, fMRI's relatively poor temporal resolution limits its ability to resolve neural activity at cognitive timescales (see Fig. 1.1).

1.1.2 Positron Emission Tomography

Positron emission tomography belongs to the group of imaging modalities that use radioisotopes [52]. This category also includes single photon emission computed tomography (SPECT) and regional cerebral blood flow (rCBF). Positron emission tomography uses a positron-emitting isotope to label a biochemical substance that is used by the brain (e.g., water, glucose, neurotransmitter or a drug). When injected into the body the emitted positrons collide with electrons and a pair of gamma rays are given off in opposite directions. These rays are detected by rings of gamma detectors surrounding the head. This information can then be used to construct images which reflect the brain's activity as a function of the labeled substance. PET is able to image the entire brain volume with a spatial resolution of approximately 5 mm, and a temporal resolution in the order of minutes (see Fig. 1.1). It is the imaging modality of choice for in vivo investigation of brain chemistry and drug receptor activity.

1.1.3 Electroencephalography

The EEG uses electrodes placed on the scalp to record the voltages that result from the electric fields associated with the currents induced by neural activity. These voltages are measured relative to some reference electrodes. EEG has a long history, it was first recorded in rabbits and monkeys by Richard Caton in 1875 [19], and in humans by Hans Berger, approximately 50 years later [14]. The reader is referred to [101] for a detailed history of EEG.

The EEG arises from the extracellular currents that are induced by synaptic activity [102]. When activated, ions flow in and out of synapses, altering a neuron's

transmembrane potential and subsequently its probability of firing an action potential. These ion flows create either a current sink or current source, which causes widespread currents in the extracellular space, linking the source and sink regions on each neuron. The contribution of these post-synaptic currents to the EEG also depends on the alignment of the neurons. Many cortical neurons are randomly aligned, so their contribution to the EEG is small. This is because scalp voltages are due to the linear superposition of the fields from the individual current sources, and if these sources are randomly aligned they will largely cancel each other out. However, cortical pyramidal cells have large, well-separated source and sink regions, and are aligned perpendicular to the cortical surface, which means their current dipole moments can summate to produce a measurable scalp EEG.

The apical dendrites of the pyramidal cells, which are closer to the scalp possess mainly excitatory synapses, which means excitatory post-synaptic potentials are a major contributor to the EEG. In contrast, the inhibitory post-synaptic potentials, resulting from inhibitory interneurons that predominately synapse onto the basal dendrite near the soma, are much less effective at generating an extracellular current. Hence their contribution to the EEG is minimal. Action potentials also induce extracellular currents, but since they are further away from the scalp, are short-lived (1-2 ms) and spatially narrow, their synchronicity is low. Hence their contribution to the EEG is minimal [101].

A similar arguments holds for the relative contribution of synchronous and asynchronous sources. Synchronous sources make a much larger contribution to the EEG, while asynchronous sources will tend to cancel each other out. A single scalp electrode measures the net voltage signal from neural masses containing 10 million to 1 billion neurons. It is estimated that 1 to 10 percent of neurons under a given electrode need to be synchronously active to produce a measurable scalp voltage.

EEG has high temporal resolution (milliseconds), but lower spatial resolution (centimeters) as shown in Fig. 1.1. The poorer spatial resolution results from the extracellular currents being volume conducted from their source and smeared by the different electrical resistances of the brain tissue, skull and scalp.

1.1.4 Magnetoencephalography

Magnetoencephalography is the magnetic counterpart of EEG. This modality records the magnetic fields associated with the currents generated by neural activity [59]. Compared to electrical currents, magnetic fields are not volume conducted or blurred by biological tissue. Instead, they fall off rapidly and predictably from their source, and therefore emerge from the skull largely undistorted. Hence MEG can provide slightly better spatial resolution than EEG, with similar temporal resolution (see Fig. 1.1). However, recording MEG data is much more technically demanding than EEG. In order to detect the minuscule magnetic fields, a superconducting quantum interference device (SQUID) coupled to a sensor coil (gradiometer) must be placed just above the scalp. This has restricted its use to mainly research applications.

1.2 EEG Source Localization

EEG source localization solves the EEG inverse problem to estimate the location, magnitude, and time course of the neuronal sources that produce the observed scalp voltages. This task presents a considerable challenge since, unlike the forward problem (prediction of scalp voltages for a given source configuration), which has a unique solution, the inverse problem is non-unique due to the relatively small number of spatial measurements ($\lessapprox 256$) and volume conduction effects. To make the problem tractable, a priori assumptions (mathematical, anatomical, and physiological) are imposed on the sources and head model. The variety of methodologies being employed has seen a proliferation of source localization techniques in recent years. For a comprehensive review of these see [9, 32, 54, 62, 97, 143].

In general, all EEG source localization techniques consist of several core components. These are shown in Fig. 1.2 and listed below:

- An EEG recording. Electrode montages containing at least 100 electrodes have been recommended for source localization [97], although, as seen in the work presented in this thesis, lower density EEG can produce both reliable and informative results.

- A solution space, which is usually a gray matter mask extracted from a structural MRI scan. If solving the inverse problem using a dynamic technique, a model describing the dynamics of the current sources within this space is also

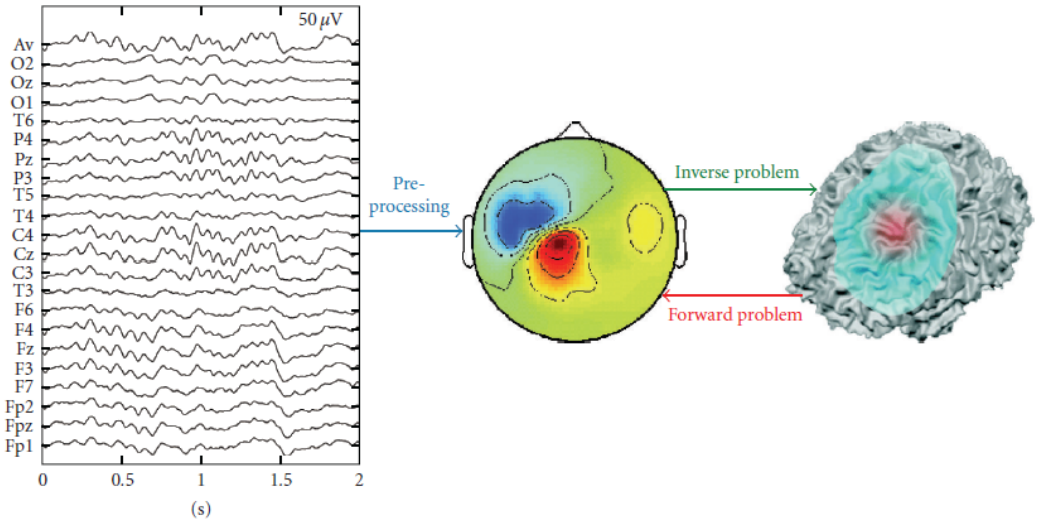

Figure 1.2: The major components of EEG source localization. As outlined in Sec. 1.2, solving the dynamic EEG inverse problem requires: (i) an EEG recording; (ii) a solution space, and, if using a dynamic inverse solution, a mathematical model describing the current dynamics within this space; (iii) a forward model linking the sources to the scalp voltages; and (iv) an inverse solution. This figure has been taken from [143].

required.

- A forward model linking the sources to the scalp voltages. This model, which is often referred to as the lead field matrix, incorporates the geometric and conductive information of the various tissue compartments in the human head (i.e., the brain tissue, cerebrospinal fluid, skull, and scalp), and is computed using quasi-static approximations of Maxwell's equations [90, 102]. See [58] for a detailed review of forward modeling for EEG source localization.

- An inverse solution, which is an algorithm that estimates the underlying current sources by applying some conditions on possible solutions. If using a dynamic inverse solution, an algorithm that combines the EEG measurements with dynamic model predictions is used.

Solutions to the EEG inverse problem can be classified based on a number of key features. One of the most important is the type of source model an inverse solution employs; these typically fall into two main categories [9, 97]. The first type are dipole-fitting approaches (also known as 'parametric' or 'equivalent current dipole' methods), in which the activity is modeled by a relatively small number of focal

sources at locations assumed a priori or estimated from the data. Examples include least-squares source estimation [123], beamformer techniques [140], and multiple signal classification methods [87, 99]. However, a drawback is that the equivalent sources can misrepresent actual activity, especially when spatially extended [104]. The second group of techniques, which this thesis is concerned with, are 'linear distributed' approaches (also known as 'imaging' methods), in which the sources are modeled by a 3D grid of dipoles throughout the head volume.

Distributed source models present a highly ill-posed inverse problem, particularly due to the mismatch between the small number of measurements ($\approx 10^2$) and the number of states to be estimated ($\approx 10^4$). This necessitates the use of constraints to identify an 'optimal' inverse solution. Numerous classes of constraints have been applied to the EEG imaging problem, and the type of constraints used represents another important feature by which inverse solutions can be categorized. Of particular importance to the work undertaken in this thesis is the distinction between instantaneous and dynamic inverse solutions. Instantaneous inverse solutions calculate each source estimate using only the data available at the current instant of time, independent of all other estimates except for the regularization parameter required in these solutions, which is usually computed by optimization over the entire dataset. Examples of instantaneous inverse solutions include: minimizing the norm of the current distribution [60]; variations of weighted minimum norm constraints as implemented in the low-resolution electromagnetic tomography (LORETA) [107]; and focal underdetermined system solution (FOCUSS) [53] algorithms. Since EEG data has temporal structure and is produced by physical processes, this assumption of temporal independence is certainly false and instantaneous techniques ignore much additional information which could further constrain the inverse solution. Incorporating information from previous times into the estimation process yields dynamical EEG inverse solutions, which is the focus of this thesis.

A variety of approaches for solving the dynamical EEG inverse problem have been investigated [4,8,11,23,24,26,29,30,42,45–47,49–51,78,88,116,126,129,137, 147]. One commonly used technique is the introduction of a temporal smoothness term, which has been successfully applied to regularization [23, 126] and Bayesian estimation [8, 26] methods. Another strategy is to use dynamic models for describing source behavior, which can then be used in various estimation schemes. Recent examples include; a particle filter using a random walk model for inverting

magnetoencephalographic (MEG) data [129], a modified LORETA algorithm which generalizes the temporal smoothness constraint into the form of an autoregressive (AR) model, allowing more complex dynamics to be modeled [147], and an inverse solution for evoked responses which uses a neural mass model within a dynamic causal modeling framework [78]. The following section will introduce the Kalman filter, which is an algorithm that can solve dynamic inverse problems, such as the one faced in EEG source localization.

1.3 Kalman Filtering

The Kalman filter is an estimation algorithm proposed by Rudolf Kalman in 1960 [73]. It was originally developed for spacecraft navigation (e.g., Apollo and Voyager [18]) but over the past 50 years has been widely used in the autonomous systems and control engineering communities [33,131]. For a comprehensive overview of Kalman filtering the reader is referred to the following standard texts [10, 55, 56, 92].

A Kalman filter recursively estimates the states and/or parameters of a dynamical system from indirect and uncertain measurements. The state variables are modeled as Gaussian random variables, and their estimates are optimal in the sense that they minimize the variance of the error between the estimated and true state variable values (i.e., the KF minimizes the mean squared error). While many variants of the Kalman filter have been developed, they all share several common features:

- A mathematical model describing the dynamics of the system (states) being estimated in terms of first-order differential equations in the state variables with one or more parameters. This is known as the process model.
- Observations (measurements) of the system, which are related to the states being estimated via a 'measurement function', which is often referred to as the observation model.
- Two noise parameters, both with a zero-mean Gaussian distribution, known as the process and measurement noise. The process noise models random inputs into the dynamic model, as well as compensating for any modeling errors and/or unmodeled dynamics by injecting uncertainty into the state covariance matrix. The effect of this is to make the Kalman filter place greater weight on

the measurements relative to the model predictions. The measurement noise models any noise in the measurement process (e.g. from the sensor(s) being used), as well as compensating for any errors in the observation model by injecting uncertainty into the measurement covariance matrix.

The original KF formulation, known as the linear Kalman filter (LKF), assumes that both the process and observation models are linear. All KF variants operate by performing a series of mathematical operations each time a new measurement (or set of measurements) becomes available. Broadly speaking these operations fall under two headings, namely the time-update (prediction) phase and the measurement-update (correction) phase, which are performed sequentially at each time instant. In the prediction phase, the vector of state variables at the current time instant is predicted from its value at the previous time instant. In the correction phase, the predicted state vector is corrected using the measurement vector at the current time instant. The reader is referred to Fig. 1.3 which shows the equations and data flow for one cycle of the discrete-time LKF. The predictor-corrector structure of this algorithm is clearly visible in this figure.

There are now a large number of KF variants and related algorithms that relax some of the assumptions of the LKF to varying degrees. A snapshot of these is now provided. For instance, filters such as the extended and unscented KFs [61, 141] and the local linearization filter [69], have been developed to handle non-linear process models and perform simultaneous state and parameter estimation. Square root filters are available that provide improved numerical stability, while particle filters relax the assumption of Gaussian uncertainties [5, 34]. For offline applications, the Kalman smoother [55] becomes available. This algorithm uses all available data, past and future, to compute each estimate. Other KFs, like the ensemble KF [36], have been designed to handle high dimensional state spaces, such as those seen in weather models. Another attractive feature of the Kalman filter, which is applicable to all implementations of this algorithm, is its ability to perform multisensor data fusion by combining asynchronous measurements from multiple sensors to estimate a single set of states.

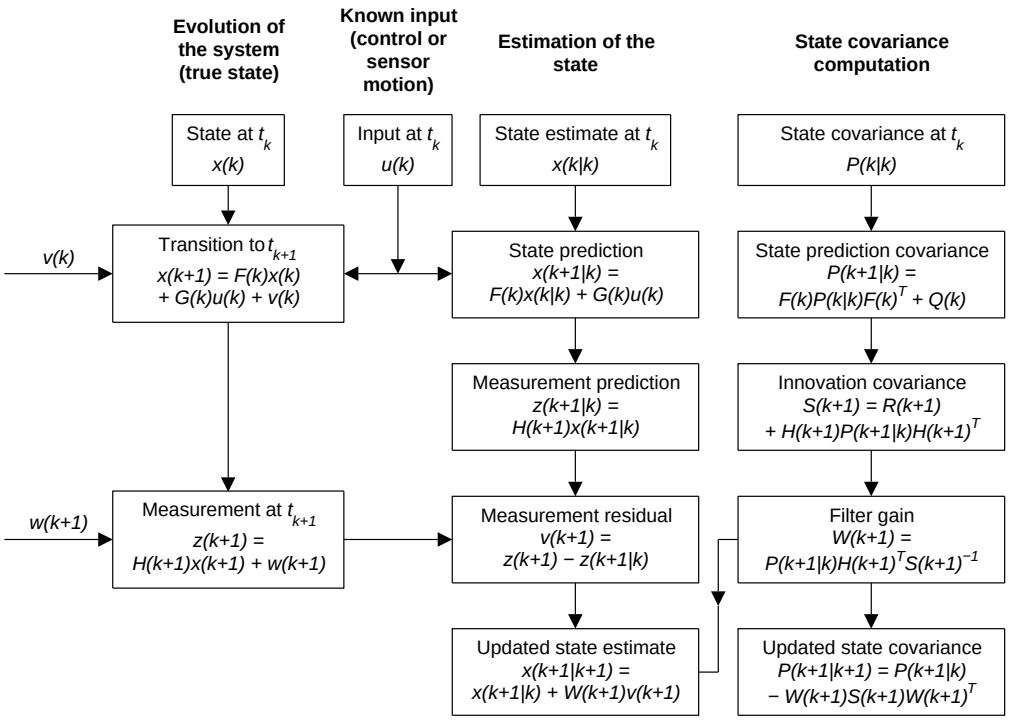

Figure 1.3: One cycle of the discrete-time linear Kalman filter performing state estimation on a linear system. This figure has been adapted from [10]. The process and observation models are denoted by $F(k)$ and $H(k+1)$, respectively. The process noise term is $v(k)$ and the observation noise term is $w(k+1)$. The process and observation noise covariance matrices are denoted by $Q(k)$ and $R(k+1)$, respectively. The input gain matrix is $G(k)$. All other terms are defined in the diagram itself.

1.4 Kalman-Filter-Based EEG Source Localization

There is considerable overlap between the core components of EEG source localization techniques and the features of the Kalman filter (see Secs 1.2 and 1.3). Firstly, the EEG is an indirect and noisy measurement of a dynamic state of interest, namely the underlying neuronal sources generating the observed scalp voltages. Secondly, the forward model (or lead field matrix) performs the role of the KF's observation model. Finally, the KF's process model provides an avenue for dynamic models describing neural activity to be introduced into EEG source localization. The KF is becoming an increasingly important tool for solving inverse problems in spatiotemporal systems. For instance, it has already been used to solve dynamic inverse problems in a variety of biomedical imaging areas including: electrical impedance tomography (EIT) [72], inverse electrocardiography (ECG) [15],

single photon emission computed tomography (SPECT) [76], diffuse optical tomography (DOT) [31], fMRI [64], diffusion MRI [112], and neural tractography [89]. Recently, a number of studies using KF-based EEG source localization have appeared [11, 29, 30, 45–47, 50, 88], and related particle [129] and local linearization [116] filters have also been investigated.

1.5 Thesis Outline

This opening chapter has introduced the concept of EEG source localization and identified the Kalman filter as a promising technique for solving the EEG inverse problem. The remainder of this thesis aims to build upon the initial forays into Kalman-filter-based EEG source localization described in Sec. 1.4. This will be done through a detailed investigation of the SWKF, which is an existing algorithm from this family. This filter's performance is studied and recommendations for improvements are made and subsequently investigated.

Chapter 2 introduces the recently developed SWKF that solves the EEG inverse problem for a realistic head model. This algorithm employs a telegrapher's equation to describe the spatiotemporal dynamics of the current dipoles. Likelihood maximization is used to fit spatially uniform model parameters and noise covariances to simulated and clinical EEGs. Performance of the SWKF is objectively evaluated using standard diagnostic tests that compute the statistical properties of the innovation. The findings of the filter performance evaluation are used to guide the work undertaken in subsequent chapters.

Chapter 3 investigates the SWKF using one-dimensional (1D) simulations to reduce the complexity of the inverse problem. Simulated EEG is generated using a telegrapher's equation, for both transient and driven dynamics, and a simplified volume conductor model. Motivated by the findings of Chapter 2, two parameters of the telegrapher's equation, in both the simulated data and SWKF, are given spatial variation of a simple functional form which better reflects true brain dynamics. Unlike Chapter 2, constrained optimization is used to restrict model parameters to ranges consistent with the process model being a telegrapher's equation. These constraints ensure the estimated model parameters maintain a clear biophysical interpretation. The Courant condition, which places an upper bound on observable wave velocity values for a given spatiotemporal grid, is also introduced and its impact

on inverse solution performance is investigated. For the purposes of comparison, inverse solutions are also computed using the optimal linear Kalman filter (LKF). Detailed assessments of filter performance, for a variety of simulated EEG datasets, are undertaken by comparing estimated and simulated parameter values, and using the aforementioned standard diagnostic tests. The impact of adding spatial variation of model parameters to the KF process model, in the cases of both transient and driven dynamics, is also discussed.

Building on the approaches used and the results obtained in the preceding chapters, Chapter 4 returns to the whole-brain EEG inverse problem and applies several features of the 1D simulation study to the SWKF investigated in Chapter 2. Following the method described in Chapter 3, spatially varying model parameters of a simple functional form are introduced into the simulated EEG and SWKF to reproduce typical spatial variation of the alpha rhythm. Compared to Chapter 2, more realistic simulated EEG is generated which exhibits wave-like patterns and spatially varying dynamics. As done in Chapter 3, constrained optimization is used again to restrict model parameters to ranges consistent with the telegrapher's equation definition, and to also impose the Courant condition on wave velocity estimates. Inverse solution performance is evaluated through a detailed assessment of state and parameter estimates, and the application of standard diagnostic tests. Comments regarding the interpretation of these results are made, and potential options for improving the performance of the SWKF are discussed.

Besides summarizing the main findings of this thesis, Chapter 5 outlines potential directions for future work, focusing on: improvements to the SWKF algorithm; more general extensions to Kalman-filter-based EEG source localization; and emerging applications for Kalman filtering in computational neuroscience.

Chapter 2

Evaluating the Performance of Kalman-Filter-Based EEG Source Localization

Abstract

Electroencephalographic (EEG) source localization is an important tool for noninvasive study of brain dynamics, due to its ability to probe neural activity more directly, with better temporal resolution than other imaging modalities. One promising technique for solving the EEG inverse problem is Kalman filtering, because it provides a natural framework for incorporating dynamic EEG generation models in source localization. Here a recently developed inverse solution is introduced which uses spatiotemporal Kalman filtering tuned through likelihood maximization. Standard diagnostic tests for objectively evaluating Kalman filter performance are then described and applied to inverse solutions for simulated and clinical EEG data. These tests, employed for the first time in Kalman-filter-based source localization, check the statistical properties of the innovation, and validate the use of likelihood maximization for filter tuning. However, this analysis also reveals that the filter's existing space- and time-invariant process model, which contains a single fixed-frequency resonance, is unable to completely model the complex spatiotemporal dynamics of EEG data. This finding indicates that the algorithm could be improved by allowing the process model parameters to vary in space.

2.1 Introduction

Functional neuroimaging aims to noninvasively characterize the dynamics of the distributed neural networks that mediate brain function in healthy and pathological states. A number of imaging techniques have emerged over the past 20 years, providing insights into brain dynamics on different spatiotemporal scales. Functional magnetic resonance imaging (fMRI) and positron emission tomography (PET) use hemodynamic and metabolic fluctuations induced by neural activity to probe brain dynamics with high spatial (millimeters), but only low temporal (seconds to minutes), resolution [133]. Electroencephalographic (EEG) source localization is a complementary imaging technique which accesses, through scalp voltages, a more direct, albeit spatially blurred, measure of the brain's electrical (neural) activity. Typically, the images generated by EEG inverse solutions have a lower spatial resolution (centimeters), but possess a much higher temporal resolution (milliseconds), and are thus important for studying brain dynamics, as they probe neural processes on cognitive timescales [9].

Solving the EEG inverse problem to estimate the location, magnitude, and time course of the neuronal sources that produce the observed scalp voltages presents a considerable challenge. Unlike the forward problem (prediction of scalp voltages for a given source configuration), which has a unique solution, the inverse problem is non-unique due to the relatively small number of spatial measurements ($\lessapprox 256$) and volume conduction effects. To make the problem tractable, a priori assumptions (mathematical, anatomical, and physiological) are imposed on the sources and head model. The variety of methodologies being employed has seen a proliferation of source localization techniques in recent years. For a comprehensive review of these see [9, 97].

Solutions to the EEG inverse problem fall into two main categories. The first type are dipole-fitting approaches (also known as 'parametric' or 'equivalent current dipole' methods), in which the activity is modeled by a relatively small number of focal sources at locations assumed a priori or estimated from the data. Examples include least-squares source estimation [123], beamformer techniques [140], and multiple signal classification methods [87, 99]. However, a drawback is that the equivalent sources can misrepresent actual activity, especially when spatially extended [104]. The second group of techniques, which this chapter is concerned with, are 'linear distributed' approaches (also known as 'imaging' methods), in which the

sources are modeled by a three-dimensional (3D) grid of dipoles throughout the head volume.

Distributed source models present a highly ill-posed inverse problem, particularly due to the mismatch between the small number of measurements ($\approx 10^2$) and the number of states to be estimated ($\approx 10^4$). This necessitates the use of constraints to identify an 'optimal' inverse solution. Numerous classes of constraints have been applied to the EEG imaging problem, such as minimizing the norm of the current distribution [60] and variations of weighted minimum norm constraints as implemented in the low-resolution electromagnetic tomography (LORETA) [107], and focal underdetermined system solution (FOCUSS) [53] algorithms. It is important to note that these inverse solutions [53, 60, 107] are instantaneous; i.e., each source estimate is calculated using only the data available at the current instant of time, independent of all other estimates except that the regularization parameter required in these solutions is usually computed by optimization over the entire dataset. Since EEG data has temporal structure and is produced by physical processes, this assumption of temporal independence is certainly false and instantaneous techniques ignore much additional information which could further constrain the inverse solution. Incorporating information from previous times into the estimation process yields dynamical EEG inverse solutions, which is the focus of this chapter.

Several approaches for solving the dynamical EEG inverse problem have been investigated. One commonly used technique is the introduction of a temporal smoothness term, which has been successfully applied to regularization [23,126] and Bayesian estimation [8, 26] methods. Another strategy is to use dynamic models for describing source behavior, which can then be used in various estimation schemes. Recent examples include; a particle filter using a random walk model for inverting magnetoencephalographic (MEG) data [129], a modified LORETA algorithm which generalizes the temporal smoothness constraint into the form of an autoregressive (AR) model, allowing more complex dynamics to be modeled [147], and an inverse solution for evoked responses which uses a neural mass model within a dynamic causal modeling framework [78].

This chapter investigates another model-based approach, the application of Kalman filtering to solving the dynamical EEG inverse problem. The Kalman filter (KF) is a widely used technique for the estimation of unobservable states in dynamical systems [10, 55, 92]. It has been used to solve dynamic inverse problems

in several biomedical imaging areas including, electrical impedance tomography (EIT) [72], single photon emission computed tomography (SPECT) [76], and diffuse optical tomography (DOT) [31, 114]. An attractive feature of the KF is that it provides a natural framework for introducing predictive models for EEG generation into source localization techniques. These models can be inferred from signal analysis, as done here, or derived from physiology (e.g., [80, 117, 119]). Despite these attributes, Kalman filtering has not been widely explored in the EEG inverse mapping field and its potential remains largely untapped, although a few studies have appeared [29, 46, 47, 88], and the related particle [129], and local linearization [116] filters have also been used. A major reason for this is that the high dimensionality of the underlying state-space makes the application of a standard KF challenging, due to the inability to accurately model the spatiotemporal interactions between all voxels and the high computational costs of running such an algorithm. However, a recently developed KF-based inverse solution [47] avoids these problems and shows considerable promise. It proposes a modified KF algorithm which reduces the high dimensionality of this problem by reformulating it as a coupled set of low-dimensional KFs running in parallel. Using a single telegrapher's equation to model the global source dynamics and likelihood maximization to estimate a small number of model and noise parameters, this technique offers improved source localization over existing instantaneous solutions (e.g., LORETA).

In this chapter the application of Kalman filtering to source localization is examined through a detailed study of the KF-based inverse solution described above [47]. The study aims to characterize, validate, and identify ways to improve the algorithm's performance. To achieve this, several new contributions are made; (i) standard diagnostic tests for objectively evaluating KF performance are introduced to EEG source localization, (ii) the application of these tests is demonstrated, (iii) results for this particular filter whose performance has not been previously evaluated formally are shown and discussed for both simulated and clinical EEG data, and (iv) the outcomes are used to direct future work. These tests have not been discussed in the growing literature on KF-based EEG source localization, despite their proven utility and widespread use in other fields where Kalman filtering is employed [10, 12, 144]. All such tests check the statistical properties required of the innovation sequence, which is the only indicator of KF performance available for real data [10]. Numerous tests, for both off- and on-line applications, have been

developed for this purpose [10, 12, 43, 57, 92, 95, 144]. Using these tests we can determine objectively whether the filter tuning step results in a well-tuned filter, as defined in Sec. 2.5. This analysis is repeated for several process models, so the relative contributions of spatial and temporal components to the inverse solution can be ascertained. Resonant behavior in the process model is then examined to provide the basis for discussing potential improvements to the algorithm.

In Sec. 2.2 the linear distributed EEG inverse problem is described. Section 2.3 outlines the KF-based inverse solution and the likelihood maximization technique used for parameter estimation. In Sec. 2.4 inverse solutions for both simulated and clinical EEG data are presented. Section 2.5 describes the tests for evaluating KF performance and discusses their results. Process model resonance is explored and discussed in Sec. 2.6, which closes by outlining ways to improve both the process model and the filter itself.

2.2 EEG Inverse Problem Formulation

To set up the EEG inverse problem, we define a continuous current vector field $\boldsymbol{j}(\boldsymbol{r}, t)$, where $\boldsymbol{r}$ and t denote space and time, respectively. The solution space is discretized into N_v grid points (voxels) $\boldsymbol{r}_v$, $v = 1, ..., N_v$, restricted to the cortical gray matter of the brain, where the majority of the EEG signal is generated [104]. Time is discretized into N_k points t_k, $k = 1, ..., N_k$. Discretized points are indicated by v and k here, rather than $\boldsymbol{r}_v$ and t_k. At each voxel the state vector is

$$\boldsymbol{j}(v, k) = [j_x(v, k), j_y(v, k), j_z(v, k)]^T. \tag{2.1}$$

The global state vector for the entire system has dimension $N_J = 3N_v$ and is written

$$\boldsymbol{J}(k) = [\boldsymbol{j}(1, k)^T, \ldots, \boldsymbol{j}(N_v, k)^T]^T. \tag{2.2}$$

The currents $\boldsymbol{j}(v, k)$ produce the EEG signal which is recorded on the scalp at N_c electrode sites. If the EEG voltage at a single electrode is denoted by $y(c, k)$, where c is an electrode label, the observation vector containing the scalp voltages at all EEG channels is

$$\boldsymbol{Y}(k) = [y(1, k), \ldots, y(N_c, k)]^T. \tag{2.3}$$

Here voltages refer to average reference (the average voltage is subtracted from each channel).

The observation equation that relates the current vectors to be estimated to the EEG signal is

$$\boldsymbol{Y}(k) = \boldsymbol{K}\boldsymbol{J}(k) + \boldsymbol{\epsilon}(k), \tag{2.4}$$

where the $N_c \times N_J$ matrix $\boldsymbol{K}$, often referred to as the lead field matrix (LFM) or the observation model, maps the current vectors to voltages at the scalp electrodes. In this study the LFM is approximated for the International 10-20 System [67] by solving the vector Laplace equation for a 3-shell spherical head model via the boundary element method [115]. The term $\boldsymbol{\epsilon}(k)$ is a N_c-dimensional vector of observational noise, which is assumed to be white, Gaussian, and unbiased, with covariance matrix $\boldsymbol{C}_{\boldsymbol{\epsilon}}$, and uncorrelated between all pairs of sensors, with equal variance $\sigma_{\boldsymbol{\epsilon}}^2$ at every electrode, so

$$\boldsymbol{C}_{\boldsymbol{\epsilon}} = \sigma_{\boldsymbol{\epsilon}}^2 \boldsymbol{I}_{N_c}, \tag{2.5}$$

where $\boldsymbol{I}_{N_c}$ is the $N_c \times N_c$ identity matrix.

Equation (2.4) cannot be inverted directly, due to the large ratio of solution points to measurements. Hence, the inverse problem can only be solved by introducing additional constraints.

2.3 Spatiotemporal Kalman Filtering

In this section, we summarize a recently developed KF-based source localization technique [47] which provides the motivation and basis for the present work. We begin by introducing a model to describe the source dynamics, and a state-space transformation which reduces the filter's computational costs. The spatiotemporal KF algorithm is then outlined. Parameter estimation is then discussed and a method to tune the filter by likelihood maximization is proposed.

2.3.1 Spatiotemporal Models

A key component of any dynamical inverse solution is a model of the system dynamics (i.e., a process model), in this case one which describes the spatiotemporal evolution of the current vectors. For this task we propose a telegrapher's equation [110] of the form

$$\left(\frac{\partial^2}{\partial t^2} + 2\zeta\omega_n\frac{\partial}{\partial t} + \omega_n^2 - b^2\nabla^2\right)\boldsymbol{j}(\boldsymbol{r},t) = \boldsymbol{\eta}(\boldsymbol{r},t), \tag{2.6}$$

where ω_n is the natural frequency, ζ the fractional damping coefficient, b the wave velocity, and $\boldsymbol{\eta}(\boldsymbol{r}, t)$ is a dynamical (process) noise term. This equation is selected for several reasons: (i) it is the continuous form of the discrete model used here and in [47], (ii) it contains an explicit temporal resonance, which is a key feature of EEG data, (iii) it allows physically meaningful parameters to be determined through the estimation step, and (iv) in previous work using mean-field modeling [119], an equivalent equation successfully described the spatiotemporal propagation of neuronal activity. To implement a KF, Eq. (2.6) is discretized with respect to space and time to give

$$\begin{aligned}\boldsymbol{j}(v,k) = \boldsymbol{A}_{L1}\boldsymbol{j}(v,k-1) + \boldsymbol{A}_{L2}\boldsymbol{j}(v,k-2) \\ + \boldsymbol{B}_{L1}[\boldsymbol{L}\boldsymbol{J}(k-1)]_v + \boldsymbol{\eta}_L(v,k),\end{aligned} \tag{2.7}$$

at each voxel where $\boldsymbol{L}$ is a discrete 3D spatial Laplacian operator of dimensions $N_J \times N_J$ which arises from the discretization of the second spatial derivative in Eq. (2.6) and is defined

$$\boldsymbol{L} = \left(\boldsymbol{I}_{N_v} - \frac{\boldsymbol{N}}{6}\right) \otimes \boldsymbol{I}_3, \tag{2.8}$$

where $\otimes$ indicates Kronecker multiplication and $\boldsymbol{N}$ is a $N_v \times N_v$ matrix with element $\boldsymbol{N}(v, v') = 1$ if v' is immediately adjacent to v (maximum of 6 neighbors per voxel in a 3D rectangular grid) and $\boldsymbol{N}(v, v') = 0$ otherwise. Boundary conditions are applied which restrict sources to the grey matter mask, i.e., if a voxel location lies outside of this mask its value is set to 0. The construction of $\boldsymbol{N}$ takes this boundary condition into account by setting to 0 any matrix elements corresponding to an immediately adjacent voxel that lies outside of the grey matter mask. The $[\boldsymbol{J}]_v$ operator selects the column vector composed of the three elements of $\boldsymbol{J}$ that correspond to grid point v. Restricting attention to classes of process models (e.g., [47]) in which the local current components in each voxel are approximated as behaving independently of each other and only interacting with the corresponding current vectors in neighboring voxels, gives the following local parameter matrices in Eq. (2.7)

$$\boldsymbol{A}_{L1} = a_1\boldsymbol{I}_3, \quad \boldsymbol{A}_{L2} = a_2\boldsymbol{I}_3, \quad \boldsymbol{B}_{L1} = b_1\boldsymbol{I}_3. \tag{2.9}$$

From discretization of Eq. (2.6), the model parameters in Eq. (2.9), assumed to be space and time invariant, are

$$a_1 = \frac{2 - (\omega_n \Delta t)^2}{1 + \zeta\omega_n\Delta t}, \tag{2.10}$$

$$a_2 = \frac{\zeta\omega_n\Delta t - 1}{1 + \zeta\omega_n\Delta t}, \tag{2.11}$$

$$b_1 = -\frac{6(b\Delta t)^2}{(\Delta x)^2(1 + \zeta\omega_n\Delta t)}. \tag{2.12}$$

where Δt and Δx are the time step and voxel size (assuming cubic voxels). From Eq. (2.7) we write the global process model as a second-order multivariate AR model:

$$\boldsymbol{J}(k) = \boldsymbol{A}_{G1}\boldsymbol{J}(k-1) + \boldsymbol{A}_{G2}\boldsymbol{J}(k-2) + \boldsymbol{\eta}_G(k), \tag{2.13}$$

where the $N_J \times N_J$ global parameter matrices are

$$\boldsymbol{A}_{G1} = a_1\boldsymbol{I}_{N_J} + b_1\boldsymbol{L}, \quad \boldsymbol{A}_{G2} = a_2\boldsymbol{I}_{N_J}. \tag{2.14}$$

The N_J-dimensional vector $\boldsymbol{\eta}_G(k)$ is a dynamical noise term which is assumed white, Gaussian, and unbiased, with covariance matrix $\boldsymbol{C}_{\boldsymbol{\eta}_G}$. To decompose this high-dimensional problem into a set of coupled low-dimensional, voxel-centered, local filtering problems, as described in the next section, requires this dynamical noise covariance matrix to be diagonal. However, for the process noise, assumption of a diagonal covariance matrix is typically not justified, due to nonvanishing instantaneous correlations between neighboring voxels. So to diagonalize this matrix, a switch to a transformed (Laplacianized) state-space $\tilde{\boldsymbol{J}}(k)$ was proposed [47] where

$$\tilde{\boldsymbol{J}}(k) = \boldsymbol{L}\boldsymbol{J}(k). \tag{2.15}$$

Assuming the same form of dynamics govern the Laplacian of $\boldsymbol{J}$, the process model is

$$\tilde{\boldsymbol{J}}(k) = \boldsymbol{A}_{G1}\tilde{\boldsymbol{J}}(k-1) + \boldsymbol{A}_{G2}\tilde{\boldsymbol{J}}(k-2) + \tilde{\boldsymbol{\eta}}_G(k). \tag{2.16}$$

As a result of this transformation, the dynamical noise covariance matrix $\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_G}$ is closer to diagonal since applying the Laplacian operator $\boldsymbol{L}$ to the state vector $\boldsymbol{J}$ reduces spatial correlations between neighboring voxels through (second-order) spatial differentiation. Assuming the process noise covariance $\sigma^2_{\tilde{\boldsymbol{\eta}}}$ to be fixed in space and time, the covariance matrix is

$$\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_G} = \sigma^2_{\tilde{\boldsymbol{\eta}}}\boldsymbol{I}_{N_J}. \tag{2.17}$$

We can substitute Eq. (2.15) into Eq. (2.16) to obtain

$$\boldsymbol{J}(k) = \boldsymbol{L}^{-1}\boldsymbol{A}_{G1}\boldsymbol{L}\boldsymbol{J}(k-1) + \boldsymbol{L}^{-1}\boldsymbol{A}_{G2}\boldsymbol{L}\boldsymbol{J}(k-2) + \boldsymbol{L}^{-1}\tilde{\boldsymbol{\eta}}_G(k). \tag{2.18}$$

By equating the process noise term in Eq. (2.18) with the one in Eq. (2.13), we find $\boldsymbol{\eta}_G(k) = \boldsymbol{L}^{-1}\tilde{\boldsymbol{\eta}}_G(k)$, which yields the process noise covariance matrix in the original space

$$\boldsymbol{C}_{\boldsymbol{\eta}_G} = \boldsymbol{L}^{-1}E(\tilde{\boldsymbol{\eta}}\tilde{\boldsymbol{\eta}}^T)(\boldsymbol{L}^{-1})^T = \sigma_{\tilde{\boldsymbol{\eta}}}^2(\boldsymbol{L}^T\boldsymbol{L})^{-1}. \tag{2.19}$$

This state-space transformation is called 'spatial whitening', and allows decomposition of the filtering problem, as described in the next section. From now on we will operate in this Laplacianized state-space by replacing $\boldsymbol{J}(k)$ with $\tilde{\boldsymbol{J}}(k)$ and $\boldsymbol{C}_{\boldsymbol{\eta}_G}$ with $\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_G}$. To obtain actual current densities and covariances we simply apply the inverse of the spatial whitening transformation; as seen shortly, this step requires one-off inversion of a very large ($\approx 10^4 \times 10^4$) matrix.

2.3.2 Kalman Filter Algorithm

At this point we could apply standard Kalman filtering to this problem in the original or Laplacianized state-space. However, given the high dimension N_J of the state-space usually seen in EEG inverse problems, the computational time and memory usage for such a filter is large enough to make the numerical estimation of model parameters performed in Sec. 2.4 impractical. To overcome this problem, a modified KF was introduced in [47] which reduces this N_J-dimensional filtering problem to a set of N_v coupled 6-dimensional KFs, one at each voxel in Laplacianized state-space, governed by the local process model Eq. (2.7). This modification requires that $\boldsymbol{C}_{\boldsymbol{\eta}_G}$ be diagonal, which explains the need for spatial whitening.

We now outline the modified KF used here. The reader is referred to [47] for further details regarding its development. Before describing the algorithm, a notational convention is defined. The term $x(k_1|k_2)$ will indicate an estimate of some quantity x computed at time k_1, based on all observations available at time k_2, where $k_1 \geqslant k_2$. Also due to the application of spatial whitening, we will replace the LFM $\boldsymbol{K}$ with $\tilde{\boldsymbol{K}} = \boldsymbol{K}\boldsymbol{L}^{-1}$ henceforth. We start by augmenting the local state vector, as the KF requires the local process model Eq. (2.7) be in the form of a first-order AR model. To achieve this, we define a new 6-dimensional local state vector $\tilde{\boldsymbol{j}}_{KF}(v,k)$ as follows

$$\tilde{\boldsymbol{j}}_{KF}(v,k) = [\tilde{\boldsymbol{j}}(v,k)^T, \tilde{\boldsymbol{j}}(v,k-1)^T]^T, \tag{2.20}$$

so the new local parameter matrices become

$$\boldsymbol{A}_{KF} = \begin{pmatrix} \boldsymbol{A}_{L1} & \boldsymbol{A}_{L2} \\ \boldsymbol{I}_3 & \boldsymbol{0} \end{pmatrix}, \quad \boldsymbol{B}_{KF} = \begin{pmatrix} \boldsymbol{B}_{L1} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix}. \tag{2.21}$$

Rewriting Eq. (2.7), we obtain the local state prediction equation

$$\begin{aligned} \tilde{\tilde{\boldsymbol{j}}}_{KF}(v, k|k-1) = \boldsymbol{A}_{KF}\tilde{\tilde{\boldsymbol{j}}}_{KF}(v, k-1|k-1) \\ + \boldsymbol{B}_{KF}\begin{pmatrix} [\boldsymbol{L}\tilde{\boldsymbol{J}}(v, k-1|k-1)]_v \\ \boldsymbol{0} \end{pmatrix}. \end{aligned} \tag{2.22}$$

The local predicted state covariance is approximated as

$$\tilde{\boldsymbol{P}}(v, k|k-1) = \boldsymbol{A}_{KF}\tilde{\boldsymbol{P}}(v, k-1|k-1)\boldsymbol{A}_{KF}^T + \boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L}, \tag{2.23}$$

where $\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L}$ is the 6×6 local dynamical noise covariance matrix given by

$$\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L} = \begin{pmatrix} \sigma_{\tilde{\boldsymbol{\eta}}}^2 \boldsymbol{I}_3 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix}. \tag{2.24}$$

The contribution of the second (neighborhood) term in Eq. (2.22) to the predicted state covariance Eq. (2.23) is ignored [47], since it is expected to be small, relative to the first (local) term, and therefore will not contribute significantly to the state covariance.

Once the local prediction equations (2.22) and (2.23) have been applied at all voxels, we predict observed scalp voltages from the global state vector:

$$\boldsymbol{Y}(k|k-1) = \tilde{\boldsymbol{K}}\tilde{\boldsymbol{J}}(k|k-1). \tag{2.25}$$

The innovation sequence is the difference between observed and predicted EEG measurements:

$$\boldsymbol{\Delta Y}(k) = \boldsymbol{Y}(k) - \boldsymbol{Y}(k|k-1). \tag{2.26}$$

The associated innovation covariance is approximated by

$$\boldsymbol{R}(k|k-1) = \sum_{v=1}^{N_v} \boldsymbol{Q}(v)\tilde{\boldsymbol{P}}(v, k|k-1)\boldsymbol{Q}(v)^T + \boldsymbol{C}_{\boldsymbol{\epsilon}}, \tag{2.27}$$

where the $N_c \times 6$ matrix $\boldsymbol{Q}(v)$ is defined as

$$\boldsymbol{Q}(v) = ([\tilde{\boldsymbol{K}}]_v \quad \boldsymbol{0}). \tag{2.28}$$

The $[\tilde{\boldsymbol{K}}]_v$ term denotes the 3 columns from $\tilde{\boldsymbol{K}}$ that correspond to the vth voxel, and the $\mathbf{0}$ matrix on the right has dimensions $N_c \times 3$. The $6 \times N_c$ Kalman gain matrix for voxel v is then

$$\boldsymbol{G}(v,k) = \tilde{\boldsymbol{P}}(v,k|k-1)\boldsymbol{Q}(v)^T\boldsymbol{R}(k|k-1)^{-1}. \tag{2.29}$$

The filtering cycle is then completed by calculating the local state estimate and the corresponding local state estimate covariance matrix,

$$\tilde{\boldsymbol{j}}_{KF}(v,k|k) = \tilde{\boldsymbol{j}}_{KF}(v,k|k-1) + \boldsymbol{G}(v,k)\boldsymbol{\Delta Y}(k), \tag{2.30}$$

$$\tilde{\boldsymbol{P}}(v,k|k) = [\boldsymbol{I}_6 - \boldsymbol{G}(v,k)\boldsymbol{Q}(v)]\tilde{\boldsymbol{P}}(v,k|k-1), \tag{2.31}$$

respectively. Applying Eqs (2.30) and (2.31) to all voxels generates the inverse solution for time point k. To obtain actual current density estimates, we then undo the spatial whitening transformation Eq. (2.15) via:

$$\boldsymbol{J}(k|k) = \boldsymbol{L}^{-1}\tilde{\boldsymbol{J}}(k|k). \tag{2.32}$$

The associated $N_J \times N_J$ covariance matrix for the actual current densities at every voxel is given by

$$\boldsymbol{P}(k|k) = \boldsymbol{L}^{-1}\tilde{\boldsymbol{P}}(k|k)(\boldsymbol{L}^T)^{-1}, \tag{2.33}$$

where $\tilde{\boldsymbol{P}}(k|k)$ denotes the Laplacianized $N_J \times N_J$ covariance matrix for all voxels, the diagonal of which consists of the $\tilde{\boldsymbol{P}}(v,k|k)$ matrices (only the first 3 columns of the first 3 rows) given by Eq. (2.31) for each voxel. The remaining elements of $\tilde{\boldsymbol{P}}(k|k)$ are filled with zeros as a result of spatial whitening, which removes off-diagonal covariances.

2.3.3 Parameter Estimation

Since no detailed prior knowledge of parameter values is usually available, a strategy for selecting optimal values for the process model parameters (a_1, a_2, and b_1), and noise covariances ($\sigma_{\boldsymbol{\epsilon}}^2$ and $\sigma_{\tilde{\boldsymbol{\eta}}}^2$) is required. In [47] it was proposed that the filter parameters should be estimated directly from the data using the widely applied technique of likelihood maximization [2, 6, 94, 127]. Following [2], filter parameters are selected by numerically minimizing the Akaike Information Criterion (AIC). The AIC, closely related to the logarithm of the likelihood, estimates the distance between the process model and the unknown true model. Its calculation

allows different model structures and parameter values to be compared objectively, so the best combination can be identified. We begin by defining a parameter vector $\boldsymbol{\vartheta} = (a_1, a_2, b_1)$. The log-likelihood ($\log L$) for the entire EEG time series is

$$\log L(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) = -\frac{1}{2}\sum_{k=1}^{N_k}[\log_e |\boldsymbol{R}(k|k-1)| + \boldsymbol{\Delta Y}(k)^T \boldsymbol{R}(k|k-1)^{-1}\boldsymbol{\Delta Y}(k) + N_c \log_e(2\pi)], \tag{2.34}$$

where $|.|$ denotes absolute value of the matrix determinant. The AIC is then

$$\text{AIC}(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) = -2\log L(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) + 2[\dim(\boldsymbol{\vartheta}) + 2], \tag{2.35}$$

where $\dim(\boldsymbol{\vartheta})$ indicates the number of parameters in $\boldsymbol{\vartheta}$, which is increased by two as we need to fit the noise covariances from the data. Every component of this algorithm is implemented in Matlab [91] and run on an IBM ThinkPad R51 (Intel Pentium 1.6 GHz, 1 GB RAM). The AIC minimization is performed by Matlab's 'fminunc' function which finds the minimum of an unconstrained nonlinear multivariable function using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) quasi-Newton method with a mixed quadratic and cubic line search procedure. For further details regarding the BFGS algorithm, the reader is referred to the Matlab documentation [91] for the 'fminunc' function, and the following standard optimization texts [38, 48]. The convergence criterion used by the optimization algorithm considers an AIC minimum to be reached, and therefore terminates the optimization routine, when the change in the AIC at the next step is less than 1×10^{-10}. The maximum number of filter runs per optimization is set to 500; with each filter iteration taking approximately 2 minutes to complete, this means optimized parameters can be computed within 42 hours. However, in practice, convergence is typically achieved much sooner.

2.4 Results of Inverse Solution

The spatiotemporal KF and parameter estimation technique are now applied to both simulated and clinical EEG data. For each dataset we computed inverse solutions for three process models: (i) full model (discretized telegrapher's equation), (ii) no spatial coupling ($b_1 = 0$), and (iii) random walk ($a_1 = 1$, $a_2 = 0$, $b_1 = 0$). This allows the relative contribution to filter performance of different parts of the process model to be assessed.

The following applies to all simulated and clinical EEG studies in this chapter. Prior to computing an inverse solution, we define a discretized solution space consisting of 3564 (N_v) $7 \times 7 \times 7$ mm gray matter voxels. These voxels cover the cortex and basal ganglia and were taken from the Probabilistic MRI Atlas produced by the Montreal Neurological Institute [93]. At each voxel the 3D local current vector is mapped to the 19 electrode sites for the 10-20 system through the LFM introduced in Sec. 2.2. However, due to the choice of a reference out of the set of electrodes, we exclude one of the electrode sites from the analysis [47], so the number of channels is $N_c = 18$; in this case Pz is chosen. After referencing, both datasets were normalized to unit variance.

The filter requires initial values $\tilde{\boldsymbol{j}}_{KF}(v, 1|1)$ and $\tilde{\boldsymbol{P}}(v, 1|1)$ be given, although the value for $\tilde{\boldsymbol{P}}(v, 1|1)$ is not critical [47]. Here the filter is initialized by setting $\tilde{\boldsymbol{j}}_{KF}(v, 1|1)$ to a $\mathbf{0}$ column vector, and $\tilde{\boldsymbol{P}}(v, 1|1)$ to an identity matrix for all voxels. If it converges, the filter is sensitive to initialization only in the short term, up to 0.5 s. The initial parameter values for the optimization algorithm, when run on both simulated and clinical EEG, were selected by the author. A number of parameter initializations were chosen for each EEG dataset, which allowed the solution space to be explored. However, for almost all reasonable starting points, the optimization algorithm consistently converged on the same parameter sets for each of the simulated and clinical EEG recordings.

2.4.1 Simulated EEG Recording

A major problem with all inverse solutions is obtaining meaningful evaluations of the algorithm's results and performance, because true sources are not available for comparison when working with real data. One solution is to use simulated EEG data, where underlying sources are known. To generate a simulated EEG dataset for this purpose requires us to select a model for the brain dynamics, which displays complex spatiotemporal behavior. Here we propose a highly simplified approximation, similar to the one used in [134], based on the observation that oscillations can be widely distributed but are often strongest in a local region; e.g., alpha activity in the visual cortex.

The temporal dynamics are modeled using a linear combination of sine functions whose components are evenly spaced in the alpha band (8 - 12 Hz). The alpha

band was selected since the clinical data used in the following section displays prominent alpha activity. The amplitude of the oscillations follows a Gaussian centered at $f_0 = 10$ Hz, so the simulated current density is

$$j(k) = \sum_{i=1}^{N_f} A(i) \sin[2\pi f(i) k\Delta t + \psi(i)], \tag{2.36}$$

where N_f is the number of frequency components, $f(i)$ is the frequency of oscillation [8 Hz $\leqslant f(i) \leqslant$ 12 Hz], $\psi(i)$ is a random phase offset [$-\pi \leqslant \psi(i) \leqslant \pi$], and $A(i)$ is the Gaussian scaling coefficient with variance σ_f^2, and

$$A(i) = \frac{1}{\sigma_f \sqrt{2\pi}} \exp\left(-\frac{[f(i) - f_0]^2}{2\sigma_f^2}\right). \tag{2.37}$$

The spatial distribution of the simulated source is modeled by the following 3D Gaussian function,

$$B(v_a) = \frac{1}{(2\pi)^{3/2}|\boldsymbol{\Omega}|^{1/2}} \times \exp\left[-\frac{(\boldsymbol{V}_a - \boldsymbol{V}_c)^T \boldsymbol{\Omega}^{-1} (\boldsymbol{V}_a - \boldsymbol{V}_c)}{2}\right], \tag{2.38}$$

centered at voxel v_c, with coordinates $\boldsymbol{V}_c = (x_c, y_c, z_c)^T$ and evaluated at each voxel v_a in the activation zone, with coordinates $\boldsymbol{V}_a = (x_a, y_a, z_a)^T$. The activation zone comprises the gray matter voxels within a certain radius of v_c; elsewhere $B = 0$. The spatial Gaussian's covariance matrix is $\boldsymbol{\Omega} = \sigma_s^2 \boldsymbol{I}_3$, where σ_s^2 is the variance. Finally, to produce the simulated current densities, the current density (2.36) is multiplied by the spatial coefficient mask (2.38).

For our simulated data, we selected an active region centered in the right occipital pole. The full-width half-maximum (FWHM) values for the frequency and spatial Gaussian distributions were 2 Hz (component spacing 0.25 Hz) and 75 mm, respectively (activation zone radius 100 mm). In this simulation, all current vectors were oriented in the *z*-direction (coronal axis) to maximize the scalp voltages at the occipital electrodes (i.e., O1 and O2). The simulated brain dynamics were then generated for 512 (N_k) time points, assuming a sampling rate of 256 Hz. Two seconds of synthetic EEG data, according to the 10-20 system, was generated from the simulated current densities by multiplication with the LFM, and average reference was applied. Next, white Gaussian observation noise was added to the data to give a signal to noise ratio (SNR) of 20:1 in terms of standard deviations. The resulting EEG data is shown in Fig. 2.1, and displays high amplitude alpha oscillations in the occipital electrodes, which are largest at O2.

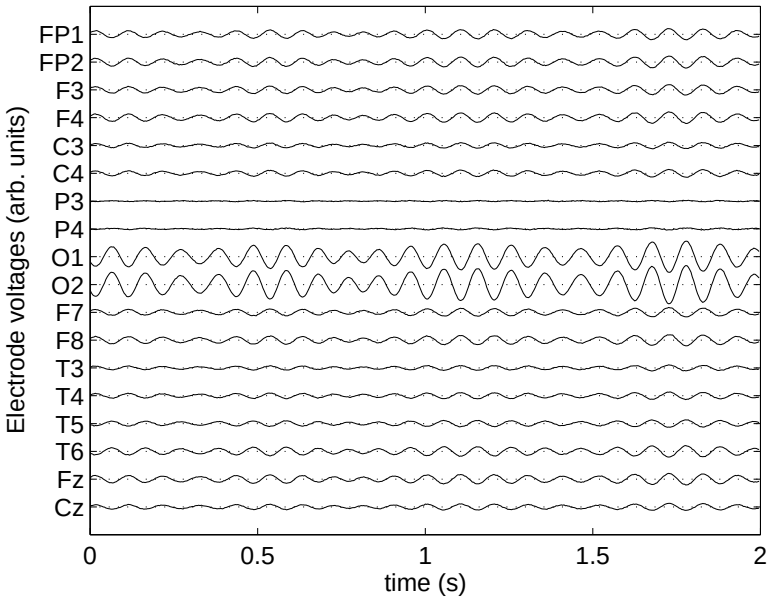

Figure 2.1: Two seconds of simulated EEG data generated for the 10-20 system (Pz is omitted) sampled at 256 Hz with an activation center in the right occipital lobe. Electrode abbreviations are on the vertical axis. The EEG potential uses average reference of all 19 electrodes (including Pz).

We begin by applying the full model inverse solution to the simulated data. Using likelihood maximization to estimate the unknown filter parameters yields $a_1 = 1.85$, $a_2 = -0.91$, $b_1 = -1.88 \times 10^{-3}$, $\sigma_{\boldsymbol{\epsilon}}^2 = 1.94 \times 10^{-3}$, $\sigma_{\boldsymbol{\eta}}^2 = 1.71 \times 10^{-8}$, and a minimized AIC = -13960. The AIC is computed from the 130$^{\text{th}}$ time point onwards (after ≈ 500 ms) for all simulations to allow transients to pass.

By looking at the parameters whose values we know from the simulated data, we immediately gain insight into the optimization's performance. The estimated value of the spatial coupling term is very small, which agrees with the simulated data which contained no spatial interactions (i.e., $b_1 = 0$). The estimated measurement noise covariance $\sigma_{\boldsymbol{\epsilon}}^2$ is also close to the actual value of 2.5×10^{-3}. These findings provide preliminary support for using AIC minimization to tune the filter.

Figures 2.2 and 2.3 illustrate the inverse solution for the simulated data. Figure 2.2 shows the spatial distribution of the current's coronal component when the activation center is maximal. We have displayed the coronal component in Figs 2.2 and 2.3 as the simulated current vectors were restricted to this direction. We see that the algorithm correctly locates the region of alpha activity and its approximate spatial

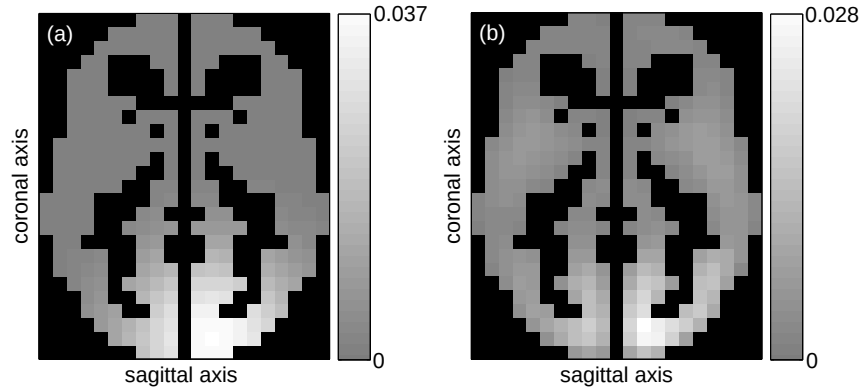

Figure 2.2: Axial slice from the gray matter mask showing the spatial distribution of the coronal component of the local current vectors at a fixed point in time for the simulated data. (a) Original current vectors used in the simulation. (b) Estimated current vectors from the inverse solution.

extent, but slightly underestimates the current densities.

Figure 2.3 shows the time series of the coronal current density component for the simulated data and the inverse solution at two voxels; one in the right occipital pole, at the center of the alpha activity, and the other in the right medial frontal lobe where no simulated activity was present. At the occipital voxel, the simulated current exhibits a large alpha oscillation. This is accurately reconstructed by the inverse solution but, as observed in the spatial data, the current amplitude is marginally underestimated. The frontal voxel is inactive during the simulation (current density is zero throughout). This lack of activity is also identified by the algorithm, where only a very low amplitude oscillation, which lies inside the error interval, is present in the estimated time series.

In Fig. 2.4 the innovation sequence for each electrode is plotted; these should be white in a properly tuned KF (see Sec. 2.5). Here we see that, once filter transients pass, most innovation sequences are near-white, aside from a small alpha oscillation present in some channels. Even in the occipital electrodes, where remaining alpha activity is more pronounced, it is significantly reduced in magnitude relative to the data (by a factor ≈ 10).

We then computed the inverse solution for the case of no spatial coupling, which gave almost identical parameters ($a_1 = 1.85$, $a_2 = -0.91$, $\sigma^2_{\boldsymbol{\epsilon}} = 1.94 \times 10^{-3}$, $\sigma^2_{\tilde{\boldsymbol{\eta}}} = 1.72 \times 10^{-8}$), AIC = -13962, estimated current density, and innovation values. This

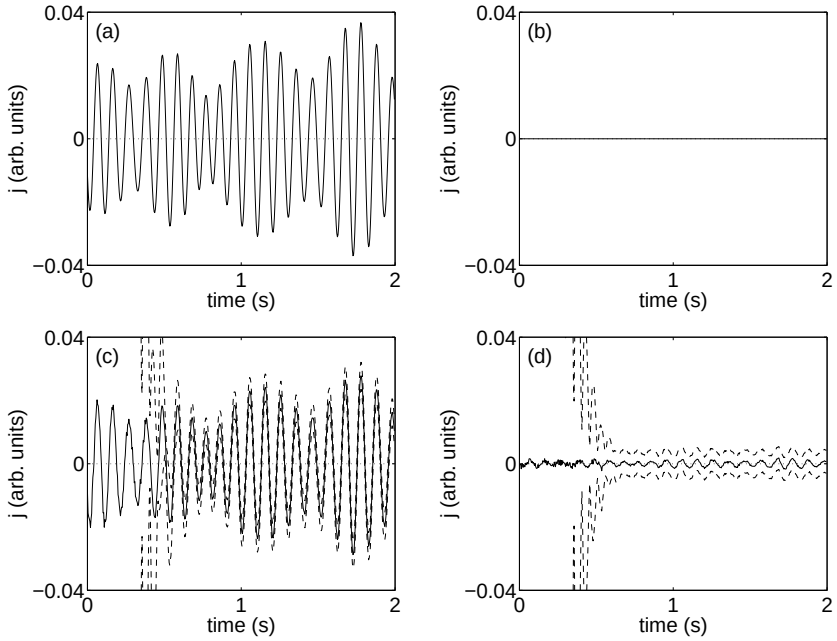

Figure 2.3: Coronal current density component for a voxel in the right occipital pole [(a) and (c)] and the right medial frontal lobe [(b) and (d)] vs. time for simulated data. Frames (a) and (b) display the simulated current vectors, while (c) and (d) show the estimated currents from the inverse solution. Solid lines represent the simulated/estimated currents, while dashed lines indicate 95% confidence intervals.

is expected, as the simulated data assumed no spatial interaction between voxels. These results imply that in both cases a well-tuned filter and accurate process model have been selected to satisfactorily describe the simulated alpha resonance.

Finally, we examined the effect of setting the dynamical model to a random walk, which reduces the process model to a temporal smoothness constraint, and forces the filter to rely largely on the observations and estimation step. However, the optimization step was unable to find a minimized AIC with the following necessary features: realistic noise covariance values; corresponds to a well-tuned filter; reaches steady state by the end of the dataset; and produces an accurate inverse solution. Therefore, the only comments we make about this inverse solution are: (i) the random walk process model functions very poorly for this dataset, resulting in a parameter space where no optimal AIC value exists that corresponds to a well-tuned KF; and (ii) this finding implies the temporal component of the process model is required for the filter to operate soundly. For these reasons, the random walk case is excluded from further analysis, here and for modeling the clinical data.

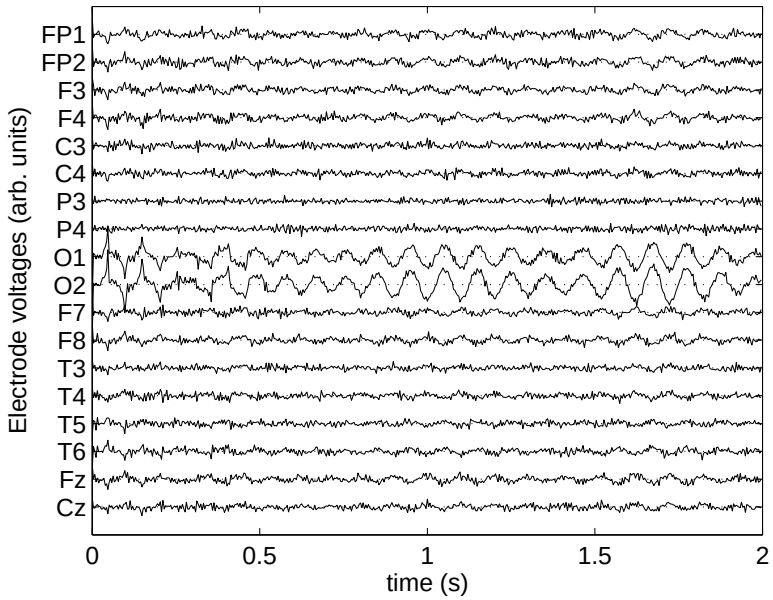

Figure 2.4: Innovation sequences for the simulated data shown in Fig. 2.1. The vertical voltage scale of the traces has been expanded by a factor of 9 relative to Fig. 2.1.

2.4.2 Clinical EEG Recording

We now estimate inverse solutions for a time series chosen from a normal EEG recording collected during routine clinical practice (Neuropediatric Clinic, University of Kiel). The data was recorded from a healthy male child aged 8.5 years, in awake resting state with eyes closed. Electrodes were placed according to the 10-20 system and the data was collected at a sampling rate of 256 Hz. The resolution of the AD conversion was 12 bit. A 2 s time series was selected from the recording for analysis and is shown, using average reference, in Fig. 2.5, which shows characteristic alpha oscillations which are most prominent in the occipital electrodes, and attenuate posterior to anterior.

The full model inverse solution for this EEG recording was computed using likelihood maximization, selecting the following filter parameter values ($a_1 = 1.60$, $a_2 = -0.65$, $b_1 = 3.08 \times 10^{-2}$, $\sigma^2_{\boldsymbol{\epsilon}} = 2.17 \times 10^{-11}$, $\sigma^2_{\tilde{\boldsymbol{\eta}}} = 7.00 \times 10^{-7}$). The minimized AIC = -2922, and is calculated from the 20^{th} time step onward (after ≈ 75 ms) for all three process models.

Figure 2.6 shows the spatial distribution of the inverse solution at a fixed moment in time. This figure shows an area of activity at the right occipital pole, as

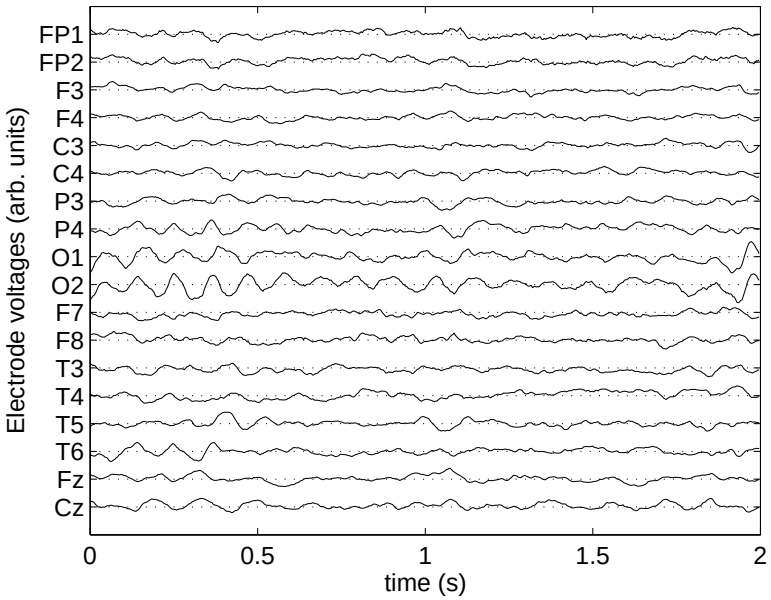

Figure 2.5: Two seconds of awake, eyes-closed EEG recorded from a child aged 8.5 years in the same format as Fig. 2.1.

expected for prominent occipital alpha activity. Figure 2.7 displays the time series for the coronal component of the inverse solution at two voxels, one in the right occipital pole and the other in the right medial frontal lobe. Once again, as expected, we see a large amplitude alpha oscillation in the occipital voxel's time series and very little activity in the frontal voxel. These observations are consistent with an eyes-closed EEG recording.

Innovation sequences are shown in Fig. 2.8. This shows that while the existing dynamical model explains much of the data structure, some of the dynamics still remain uncaptured, especially alpha activity in the occipital electrodes. Relative to the data, these oscillations are approximately five times smaller in magnitude, but twice the size of the corresponding ones in the simulated data.

The inverse solution was then recomputed with spatial interactions removed; the parameters ($a_1 = 1.61$, $a_2 = -0.66$, $\sigma^2_{\boldsymbol{\epsilon}} = 4.77 \times 10^{-11}$, $\sigma^2_{\tilde{\boldsymbol{\eta}}} = 6.49 \times 10^{-7}$), AIC $= -2645$, reconstructed currents, and innovations remained essentially unchanged. This indicates that the spatial coupling term makes only a small contribution to the inverse solution.

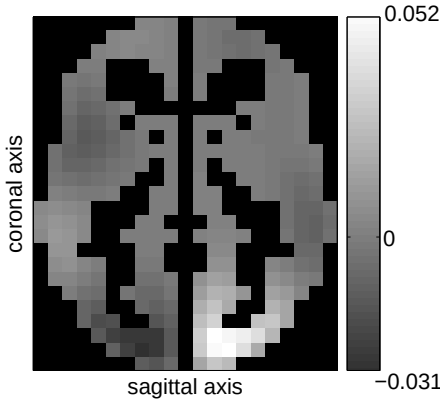

Figure 2.6: Axial slice from the gray matter mask showing the spatial distribution of the coronal component of the local current vectors at a fixed point in time for the clinical recording as estimated by the inverse solution.

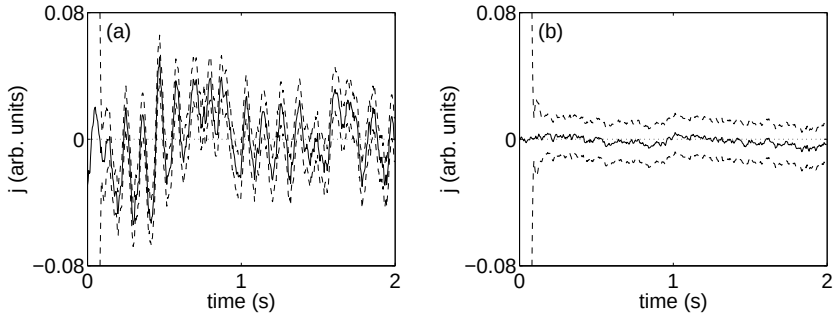

Figure 2.7: Estimated coronal current density component for a voxel in the right occipital pole (a) and the right medial frontal lobe (b) vs. time for the clinical recording. Solid lines represent estimated current vectors from the inverse solution, while dashed lines indicate 95% confidence intervals.

2.5 Analysis of Filter Performance

This section focuses on evaluating the performance of the KF itself. In any real application, validating and optimizing filter performance is difficult because, unlike a simulation study, the true states are unknown and the only information available is contained in the observations of the states. As a result, analysis of the innovation is the principal means of evaluating KF performance [10]. The AIC, which is a function of the innovation and its covariance, has already been used as a relative measure, but we now apply a series of standard diagnostic tests widely used for objectively evaluating and tuning KF performance [10]. These are applied after the

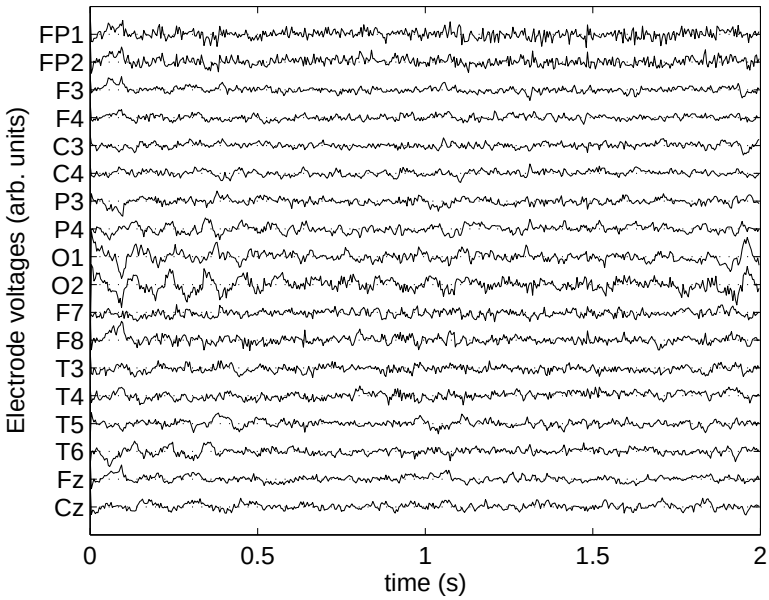

Figure 2.8: Innovation sequences for the clinical data shown in Fig. 2.5. The vertical voltage scale of the traces has been expanded by a factor of 5 relative to Fig. 2.5.

likelihood maximization step, and allow us to determine the overall (rather than relative) quality of the filter, something which is difficult to ascertain from the AIC alone. The tests focus on the properties of the innovation sequence, which should be normally distributed, unbiased (zero mean), uncorrelated (white), and have the correct magnitude (i.e., actual and filter-predicted innovation covariances should be the same). The testing procedure uses the recommendations of the standard reference [10], which are similar to many diagnostic tests described in the literature, and consists of the following five steps:

(i) Using a single sample Kolmogorov-Smirnov (KS) goodness-of-fit hypothesis test [113], we determine if the innovation sequence is normally distributed (innovation non-Gaussian if $P < 0.05$). The tests in Steps (ii) and (iii) assume the innovation is Gaussian.

(ii) To determine if the innovation is unbiased, a one-sample t-test [108] is used (innovation has non-zero mean if $P < 0.05$). The innovation must be unbiased for Steps (iii) and (iv).

(iii) We determine whether the actual and filter-predicted innovation covariances are the same. A mismatch indicates that overall filter noise levels (i.e., the pro-

cess and/or observation noise covariances) have been set incorrectly, which can degrade filter performance, and requires further analysis to ascertain its cause. An inaccurate process model can also result in discordance between the innovation covariances. Assuming (i) and (ii) hold, noise levels can be investigated by checking that approximately 95% of innovation values lie within two standard deviations of zero.

A more precise means of assessing filter noise levels, again requiring that (i) and (ii) hold, is to carry out a χ^2 test on the normalized square of the innovation

$$\boldsymbol{\Delta Y}_N(c,k) = \frac{[\boldsymbol{\Delta Y}(c,k)]^2}{\boldsymbol{R}_c(k|k-1)}, \tag{2.39}$$

where $\boldsymbol{R}_c(k|k-1)$ denotes the innovation covariance for channel c at time k. If (i) and (ii) hold, $\boldsymbol{\Delta Y}_N$ will be a χ^2 random variable (resulting from squaring a Gaussian random variable), with a mean of 1 if the actual and filter-predicted covariances match. To compare these, a test statistic

$$\overline{\boldsymbol{\Delta Y}}_N(c) = \frac{1}{N_k}\sum_{k=1}^{N_k} \boldsymbol{\Delta Y}_N(c,k), \tag{2.40}$$

for each channel is used, whence we obtain 95% confidence intervals from which we can determine whether the noise levels are correct. If the average normalized innovation lies below these bounds, the assumed noise levels are too high, and vice versa.

(iv) The innovation sequence should be white. Any correlations are due to the presence of unmodeled process dynamics or the ratio of observation to process noise being too high. We computed the innovation's time-averaged (biased) autocorrelation r for each channel c by calculating the inverse Fourier transform (FT) of the power spectral density (PSD) [92]. The resulting autocorrelation is equivalent to that obtained via the following time-domain equation [92]

$$r(c,\tau) = \frac{1}{N_k - \tau}\sum_{i=1}^{N_k-\tau} \boldsymbol{\Delta Y}(c,i)\boldsymbol{\Delta Y}(c,i+\tau), \tag{2.41}$$

where τ indicates a discrete time shift ranging from 0 to $N_k - 1$. The autocorrelation was normalized so $r(c,0) = 1$. The number of points in the autocorrelation is halved due to the FT, and is denoted by N_a. The autocorrelation for each channel can be used as a test statistic, which should be

approximately normally distributed with a mean of 0 and standard deviation $1/\sqrt{N_a}$ if the innovation is unbiased and white. Therefore, we can approximate the 95% confidence intervals by $\pm 2/\sqrt{N_a}$. If approximately 95% of the autocorrelation does not lie within these bounds, then the innovation is not white.

Finally we calculated the innovation PSD $\boldsymbol{S}$ for each channel c using Welch's method [142]. The PSD allows the innovation's frequency content to be examined, and this should be flat for an uncorrelated signal [55]. From these PSDs we computed the spectral entropy H_S [65] for each channel c

$$H_S(c) = \frac{-\sum_{f=1}^{N_f} S(c,f) \log_e S(c,f)}{\log_e N_f}, \tag{2.42}$$

where f is the frequency bin number and N_f is the number of bins in the PSD. Spectral entropy is a compact measure of a power spectrum's 'peakedness' (or conversely 'flatness'), ranging from 0 for a monochromatic signal, to 1 for a completely random one. This is useful for comparing the overall 'whiteness' of the innovation sequences between channels and different filter configurations.

(v) It can be difficult to distinguish the relative contributions of process and measurement errors to the innovation, so it is important to look at the error between the state estimate and prediction when evaluating KF performance, because it relates only to process errors, and should be approximately uncorrelated and bounded by its covariance.

We stress that the measurement of KF performance is best thought of as a continuum, spanning from filters that are optimal (i.e., pass the above tests by wide margins) to ones that cannot be tuned to provide any useful estimates. The innovation tests here provide a very strict measure of KF performance, and filters which pass these are close to optimal. However, as stressed in [10] many KFs in real applications lie between these two extremes, and while they may not satisfy all of these rigorous tests all of the time, they do provide useful results, as we will see with the filter here. We term such filters 'well-tuned'. In this situation, the aim of filter tuning is to make the KF as close to optimal as possible, and evaluation of the filter's performance will require a degree of 'engineering judgement'.

We now apply this procedure to the simulated and clinical EEG inverse mappings described in Sec. 2.4. Steps (i) to (iv) are applied to all 18 channels used by the inverse solution, and Step (v) is performed by looking at the error between the state estimate and prediction for voxels of interest in each analysis. Due to some intrinsic short-term filter behavior, and the filter receiving inaccurate initial values, the performance analysis begins, like the AIC calculation, at the 130^{th} and 20^{th} time steps for the simulated and clinical data, respectively, to allow filter transients to pass.

2.5.1 Simulated EEG Recording

When the filter is applied to the simulated data using the full process model, the KS test finds all 18 innovation sequences are normally distributed, while the t-test shows that the 18 innovations are unbiased. However, when we apply Steps (iii) and (iv), a χ^2 test found that the actual innovation covariances did not match the filter-predicted ones at any of the 18 channels. The results at 16 electrode sites indicated that overall filter noise levels were set too high (on average by a factor of 3.0), while the remaining channels (O1 and O2) suggested noise levels were too low. However, the two occipital channels are discounted since the actual innovation covariance value is inflated by a small amount of residual alpha activity (see Fig. 2.4). These tests imply that AIC minimization has selected conservative noise values that will not adversely affect the filter's performance.

Turning our attention to the detection of correlations in the innovation sequences, we computed the PSD for both the recorded data and innovation at each channel, as well as the innovation's autocorrelation. Figure 2.9 shows the PSD and autocorrelation for two typical channels, O2 and P3, and demonstrates that the process model selected through optimization describes the alpha activity present in the simulated data quite well, although a small alpha peak remains in all innovations, particularly O1 and O2, due to the inverse solution underestimating the source magnitudes. This residual alpha activity means that 17 channels have greater than 10% of their autocorrelation lying outside the 95% confidence bounds. From the data and innovation PSDs, the spectral entropy was calculated for each channel and plotted in Fig. 2.11(a). This figure shows that the innovation sequences are significantly whiter than the simulated data across all channels. The dip in the innovation curve at O1 and O2 is due to the alpha activity still present at the occipital electrodes.

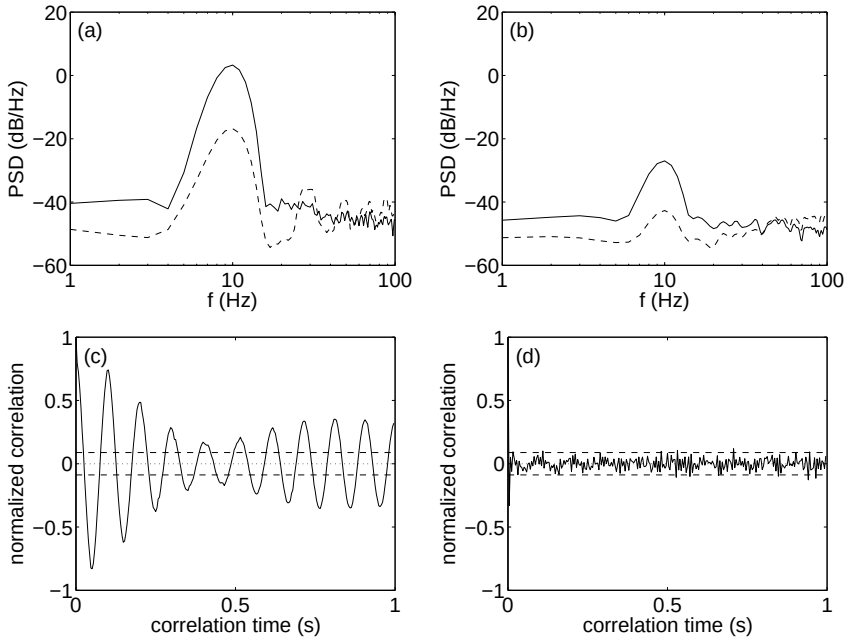

Figure 2.9: Temporal properties of the innovation sequence at channel O2 [(a) and (c)] and P3 [(b) and (d)] for the simulated data. Frames (a) and (b) display the PSD of the data (solid) and innovation (dashed), while (c) and (d) show the autocorrelation of the innovation (solid) and its 95% confidence bounds (dashed).

We completed our filter evaluation by looking at the error between the state estimate and the state prediction for selected voxels. It was found that some alpha activity was present in this error, mainly in the z-component of voxels at or near the center of the simulated activation. However, this appears to result from the inverse solution spatially blurring the reconstructed activation and thus underestimating the current densities, rather than a process model deficiency, and is most likely due to the small number of electrodes. The predicted and estimated current densities being in phase, along with the data and predicted observations, support this explanation.

The performance analysis was repeated for the inverse solution computed with no spatial interaction between voxels. Very similar results were obtained, as expected since the simulated data assumed no spatial coupling.

2.5.2 Clinical EEG Recording

The filter performance analysis is now repeated for the clinical data, starting with the full process model. The KS test found that all innovation sequences are Gaus-

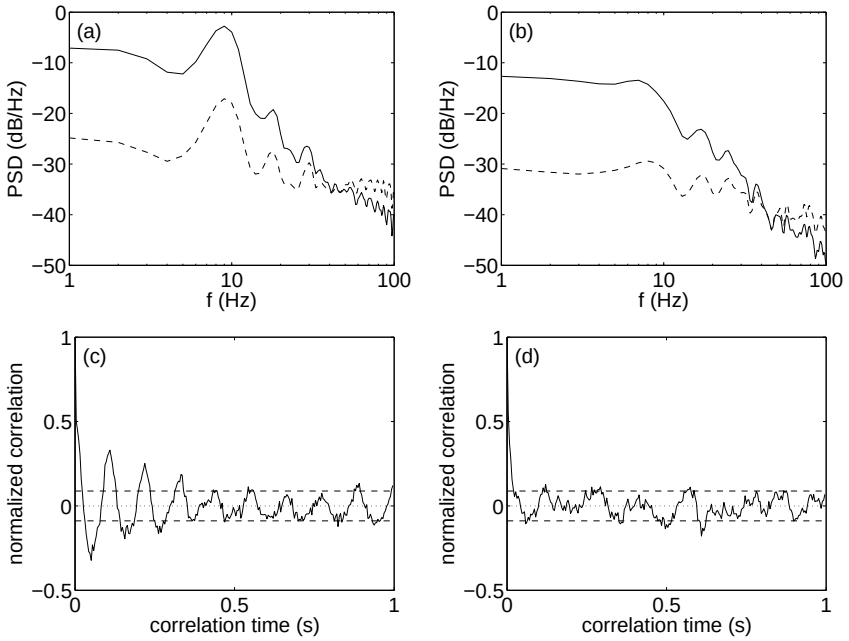

Figure 2.10: Temporal properties of the innovation sequence at channel O2 [(a) and (c)] and F4 [(b) and (d)] for the clinical data. Frames (a) and (b) display the PSD of the data (solid) and innovation (dashed), while (c) and (d) show the autocorrelation of the innovation (solid) and its 95% confidence bounds (dashed).

sian, and the t-test identified 11 out of 18 channels to be unbiased. We then applied the χ^2 test which found that the actual innovation covariances were smaller than the filter-predicted covariances for all 18 channels. This indicates that filter noise levels, selected via AIC minimization, were set too high (on average by a factor of 4.1). However, as the results for the clinical data demonstrate, these somewhat conservative noise values are compatible with correct filter operation.

We then checked whether the innovations were white by computing the PSD and autocorrelation at each channel. Results are shown for illustrative channels, O2 and F4, in Fig. 2.10. This shows that the filter can handle the low-pass characteristic of the EEG recording, as the flatter PSD indicates, but what does remain in the innovation sequence of some channels is a considerable alpha resonance, which is most prominent in the occipital electrodes. The autocorrelations confirm this, as they are clearly not white, and all channels have greater than 10% of the autocorrelation outside the confidence intervals. Despite the unmodeled alpha activity, the filter's ability to significantly whiten the innovation relative to the signal is clear when the spectral entropy for each channel is plotted in Fig. 2.11(b). The small decrease in

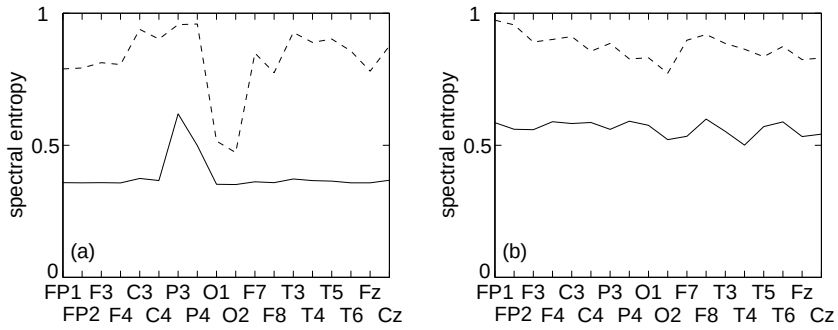

Figure 2.11: Spectral entropy of the EEG recording (solid) and innovation sequence (dashed) at each channel for the simulated (a) and the clinical (b) data.

the innovation spectral entropy is from the alpha activity remaining at the occipital electrodes.

Finally, we looked at the error between the state estimate and prediction for a number of voxels. As with the simulated data, we found alpha oscillations present in this time series, particularly in voxels around the occipital poles. However, unlike the simulated data, this appears to result from a deficiency in the process model, as the predicted current density lags the estimated one. More precisely, the process model is unable to reproduce the alpha activity accurately.

The performance analysis was repeated for the process model without spatial coupling. As seen with the inverse solutions themselves in Sec. 2.4, the performance of this filter is nearly identical to the one using the full model. This again indicates that, while the filter inverts the clinical data quite well, the spatial part of the process model does little to enhance its performance.

2.5.3 Preliminary Overview of Filter Performance

The inverse solutions and the validation tests have shown this tuning technique produces well-tuned filters, although some potential improvements have also become apparent. We found for both simulated and clinical data that the innovations were generally Gaussian and unbiased, while the optimization step selected slightly conservative values for the noise parameters. Similarly, the process models selected for the two datasets modeled the EEG data satisfactorily, as demonstrated by the spectral entropy increasing by ≈ 0.3 (between the data and the innovations) across nearly all channels in both datasets (see Fig. 2.11). However, the correlation analysis and the error between the state estimate and prediction revealed residual al-

pha activity in the innovation sequences of the simulated and clinical data, but for different reasons. The low-amplitude alpha oscillations present in the simulation innovations (most prominent at O1 and O2) appear to result from the filter underestimating the current density's magnitude, which was caused by spatial blurring due to the small number of electrodes used. While this issue is still present in the clinical study, the major reason for the (larger) alpha waveforms in the innovations is the process model being unable to fully capture the alpha resonance; this issue is examined in the next section.

We also found that dropping the spatial coupling from the process model had surprisingly little effect on the filter's performance for the clinical data (although expected for the simulation study). This discovery further indicates that, in its current implementation, the filter's major deficiency most likely lies in the temporal part of the process model, which overshadows any influence the spatial term might have. Alternatively, it is possible that the model's spatial component is inaccurate, and the optimization step seeks to remove it from the inverse solution by selecting $b_1 \approx 0$, or its impact is nullified by spatial whitening. Further analysis is required to resolve this issue.

Finally, a random walk process model was investigated. This model so significantly degraded performance that an optimal, well-tuned, filter could not be found for either dataset. These results imply that the temporal component of the process model, which performs better for the simulated data, is necessary for the filter to function properly. However, the findings of the filter performance analysis, especially for the clinical data, indicate that the modeling part of this estimation technique could be substantially improved. Possible modeling improvements, along with resonant behavior in the process model, are the subject of the next section.

2.6 Resonant Behavior of the Process Model

Here we investigate the process model's resonant behavior, particularly for the clinical data where model deficiencies have been identified. We focus on the model's temporal aspects, as Sec. 2.5 revealed the spatial term has minimal impact on the inverse solution. We begin by obtaining expressions for the parameters describing the model's resonant behavior, which provides additional insight into the resonant properties of the inverse solutions generated in Sec. 2.4 and why the alpha resonance

was modeled better in the simulated data. Then the inverse solution is computed for a series of process models, each containing an explicit alpha resonance to capture the posterior alpha activity present in the clinical data. Suggestions for potential future improvements to both the dynamical model and the filter algorithm are then made.

2.6.1 Resonant Process Model

The equations for the temporal AR parameters, Eqs (2.10) and (2.11), can be manipulated to give expressions for ω_n and ζ as functions of a_1 and a_2:

$$\omega_n = \sqrt{\frac{2(1-a_1-a_2)}{(\Delta t)^2(1-a_2)}}, \tag{2.43}$$

$$\zeta = \left(\frac{1+a_2}{1-a_2}\right)\sqrt{\frac{1-a_2}{2(1-a_1-a_2)}}. \tag{2.44}$$

These expressions convert a_1 and a_2 into parameters from the original telegrapher's equation that have a clear physical interpretation, and allow us to better characterize the resonant behavior of the process model selected by likelihood maximization.

From the simulated data model parameters, we find $\omega_n = 62.2$ s^{-1} (9.9 Hz) and $\zeta = 0.20$, while $\omega_n = 63.5$ s^{-1} (10.1 Hz) and $\zeta = 0.86$ for the clinical data. To further illustrate the model's temporal characteristics, the frequency response for the model's temporal component is shown in Fig. 2.12 for both datasets. The process model selected for the simulation has resonant properties which closely match the data itself, as the estimated natural frequency of 9.9 Hz lies near the center of the Gaussian frequency envelope and ζ is also close to its actual value of zero for undamped sine waves. As expected, the model has a sharp resonance at $\approx$ 10 Hz, which can be seen in Fig. 2.12. However, the process model for the clinical data has $\zeta \approx 1$, so it displays no discernible resonant behavior, as is clear from the absence of any peak in the frequency response in Fig. 2.12.

If we recall that the AIC minimization step identifies the global (space- and time-invariant) process model which best explains the data, these results are to be expected. In the case of the simulated data, which was generated by a single source centered in the right occipital lobe, the dynamics can be described sufficiently by a single, globally resonant, process model. This is not the case for the clinical data,

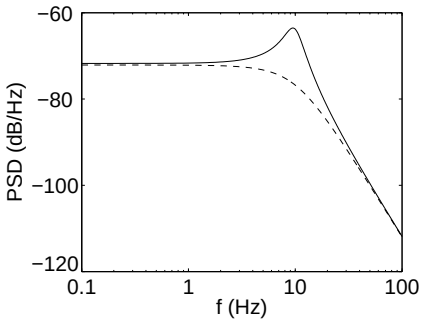

Figure 2.12: PSD of the process model's temporal component as selected by likelihood maximization for the simulated (solid) and clinical (dashed) datasets.

where the recording's dominant alpha resonance shows considerable spatial dependence (diminishing in amplitude frontally) which cannot be accurately captured by a process model with spatial and temporal uniformity. Instead, the optimization step identified the one global feature of the clinical EEG, its low-pass filter characteristic. As a result, the alpha activity seen in the innovation of some channels (e.g., O1 and O2) is expected, as the alpha resonance is unmodeled.

2.6.2 Inverse Solution With an Explicit Resonance

With the selection of appropriate values for ω_n and ζ, the process model can describe resonant features of the EEG, as seen for the simulated data. However, due to the space- and time-invariant process model, the parameter estimation step selected a non-resonant model for the clinical data, despite the presence of posterior alpha activity. So we now examine the effect on filter performance of applying a process model with an alpha resonance to the clinical data, to provide further insight into how the filter could be improved.

We began by fixing the process model's natural frequency $\omega_n = 56.5$ s^{-1} (9 Hz) to match the alpha frequency at O2. The strength of the resonance was varied across four filter runs by setting $\zeta = 0.2$, 0.3, 0.4, 0.5. To find the optimal filter for each run AIC minimization was used. Due to the spatial term's minimal impact on the inverse solutions for the clinical data it was ignored (i.e., $b_1 = 0$), so only the two noise covariances needed to be estimated. To allow transient filter behavior to pass, the AIC was calculated from the 130$^{\text{th}}$ time step onwards for each case. For the purposes of comparison the full model filter was reoptimized over this segment of clinical data, giving $\omega_n = 36.4$ s^{-1} (5.8 Hz), $\zeta = 0.75$, $b_1 = 1.11 \times 10^{-2}$,

$\sigma_{\epsilon}^2 = 1.27 \times 10^{-9}$, $\sigma_{\tilde{\eta}}^2 = 1.25 \times 10^{-7}$, and AIC $= -3307$.

We find that introducing an explicit resonance degrades the filter's overall performance; for instance, AIC values increased from -2091 to 2408 as ζ was decreased from 0.5 to 0.2. The spectral entropy of the innovations, which are plotted for each run in Fig. 2.13, tell a similar story: they generally decrease (innovations become less white) as the model is made more resonant, although resonant behavior, up to a point, does marginally whiten the innovation of channel O2; the site with the most significant alpha activity. This is shown in Fig. 2.13, where the spectral entropy increases slightly in channel O2 for $\zeta = 0.3$ and 0.4. However, this small improvement comes at a considerable cost, as the global resonance distorts (with decreasing ζ) the modeling of the low (sub-alpha) frequencies, which can be seen in the innovations of the other channels, and results in poorer filter performance. This is not surprising since the alpha resonance is not present at all sites in the clinical data.

From this analysis we conclude that, overall, the 'optimal' non-resonant process model outperforms the model with an explicit uniform alpha resonance, as expected, although neither of these space- and time-invariant models can accurately describe the spatiotemporal complexity of clinical EEG data. Therefore, improved dynamical models are required.

2.6.3 Future Directions

Guided by our filter analysis, we now discuss options for improving the inverse solution, focusing particularly on the process model. The first issue is what form the dynamical model should take. Given that resonant behavior is a key feature of EEG data, the existing telegrapher's equation, which contains a single resonance, would be a reasonable choice, although physiology-based models of brain dynamics, such as those presented in [80,117,119], are also attractive, because estimated parameters are physiologically meaningful, increasing the information provided by the inverse solution. Furthermore, these models could better describe the spatial interactions between voxels, an issue that warrants further investigation.

As noted, a uniform global model of brain dynamics is unrealistic, so regardless of what form of model is selected, its behavior will require spatiotemporal variation, e.g., to model the spatial properties of the alpha rhythm seen in the clinical data. This issue was previously investigated for this filter using generalized au-

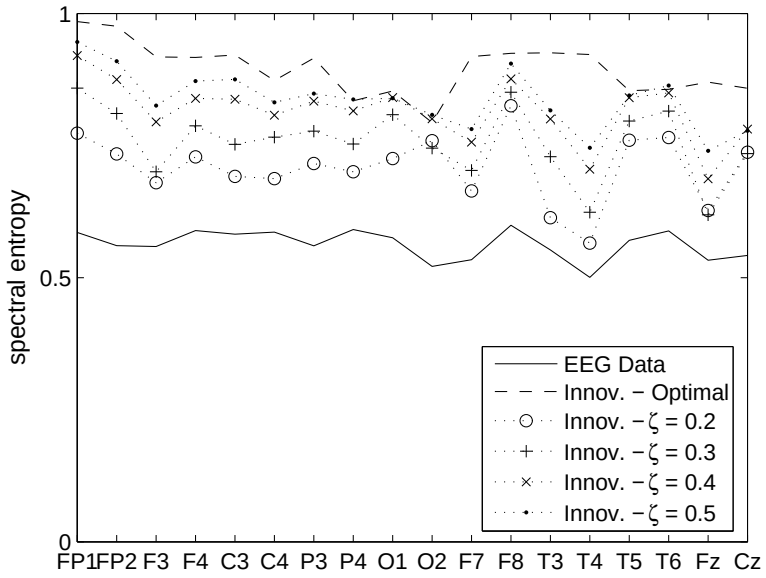

Figure 2.13: Innovation spectral entropy as the damping coefficient in the process model is varied ($\zeta = 0.2, 0.3, 0.4, 0.5$) at a fixed natural frequency (9 Hz) for the clinical data. The spectral entropy for the optimized filter innovations and the clinical data itself are also shown, as labeled.

toregressive conditional heteroskedasticity (GARCH) modeling of covariance [46], which was found to enhance performance. In that study, the same homogeneous model was used, but the process noise (which measures our confidence in the dynamical model) was allowed to vary in space and time, as a function of how well the process model was performing at a particular voxel. We propose the alternative approach of letting the process model parameters (e.g., ζ and ω_n) have spatiotemporal variation, which opens the possibility of parametric imaging, but poses a more difficult parameter estimation task than in this chapter, although for systems with spatial variation described by a relatively small number of parameters it may be possible to apply the likelihood maximization technique used here. Another strategy is to estimate the parameters within the KF itself, with the added benefit of providing the quality of each parameter estimate via the state covariance matrix. This produces a nonlinear filtering problem which can be solved using algorithms such as the extended Kalman filter (EKF) [10], or the unscented Kalman filter (UKF) [71]. In other fields, both the EKF [120, 148] and UKF [128] have been successfully applied to system identification for spatiotemporal systems modeled by partial dif-

ferential equations. Another option, which also permits estimation inside a KF, is to extend the state-space and model each parameter as a Gaussian random field [22]. However, as discussed in [47], regardless of what technique is used, the estimation problem must be observable [10] (i.e., the states must be estimable from the measurements), which is of increasing concern as the number of quantities estimated from the data rises.

A difficulty inherent to the EEG inverse problem, and of particular significance to the KF, is its high dimensionality. Typically, this necessitates simplification of the filter algorithm to reduce memory consumption and achieve practical run times. Here spatial whitening is used for this purpose, reducing the algorithm to a set of low-dimensional KFs. A filter that operates in the full untransformed state-space would offer two key advantages: (i) removal of any distortions introduced by the 'strong' whitening transformation, which will allow its effect on the inverse solution to be properly assessed, and (ii) the process model will describe the state of interest; i.e., the current density, rather than its second spatial derivative, which is especially important for physiology-derived models. The gap between the full filter and the single voxel centered, spatially whitened version could potentially be bridged by partitioned filters [128]. These filters divide the state-space into local filtering neighborhoods, and allow the trade-off between computation time/memory usage and filter performance to be examined.

Finally, it is worth mentioning that the inverse solution could be further constrained by introducing additional information. High density EEG recordings could be used to provide extra observations (up to 256 channels), which will improve the spatial resolution of the inverse mappings. When the inverse solution is computed off-line, the Kalman smoother [55] becomes available. This algorithm uses all available data, past and future, to compute each estimate, and has recently been applied to the EEG inverse problem [29, 88]. Using the KF to fuse EEG data with other imaging modalities, particularly fMRI, is a natural extension of this work [10, 29, 116] that could improve spatial resolution beyond what is possible with EEG alone.

2.7 Summary and Conclusion

We have investigated the application of dynamical inverse solutions to EEG source localization. Dynamical techniques are of particular interest because they provide a natural framework for introducing the growing number of models describing brain dynamics [80, 117, 119] into inverse solutions. The KF is an example of a model-based estimation technique that is well suited to solving inverse problems, but has only lately been applied in this field. Motivated by its potential, we introduced a recently proposed KF-based source localization technique [47]. Key features of this algorithm are: the process model is a space- and time-invariant telegrapher's equation, a spatial whitening transformation is used to reduce its computational burden, and the filter is tuned using likelihood maximization.

Inverse solutions for simulated and clinical data, both containing alpha activity in the occipital lobe, were computed and presented for three process models. The optimized filters were then analyzed in detail using standard diagnostic tests for evaluating KF performance. Following this, the resonant properties of the process model were examined and the effect of introducing an explicit alpha resonance into the filter was explored for clinical data. The major findings are:

(i) The AIC minimization step selects appropriate model parameters, and noise covariances, which result in a well-tuned filter as indicated by the reconstructed current densities and diagnostics. This shows that likelihood maximization is effective for filter tuning, but tuning still requires an appropriate process model to be chosen. For instance, the simulated and clinical data could be modeled by either the full model, or the model without spatial coupling, although these models performed better for the simulated data, and the spatial term made only a small contribution to the clinical data's inverse solution. In contrast, a random walk model could not be optimized for either dataset.

(ii) The process model is a telegrapher's equation, which contains a resonance whose properties were examined. It was found that AIC minimization, which finds the space- and time- invariant model that best describes the data, selected a process model containing an alpha resonance for the simulated, but not the clinical, EEG. This makes sense as the model chosen should capture any spatially and temporally uniform features of the time series, which for the simulated data is the alpha activity (the only salient feature), and for the clini-

cal recording is the low-pass characteristic. Thus, these findings explain why (a) the innovations for the clinical data (especially the occipital electrodes) contain unmodeled alpha activity of a higher magnitude than in the simulated EEG, and (b) the predicted and estimated current densities are out of phase only in the clinical study.

(iii) The introduction of an explicit alpha resonance into the process model for the clinical data degraded filter performance. This is due to a mismatch between the globally resonant dynamical model and the data, where the alpha activity is confined predominantly to the occipital electrodes. However, introduction of a resonance did improve the modeling of the posterior alpha rhythm.

(iv) We demonstrated the utility of applying a battery of diagnostic tests to this KF, as they provide numerous insights into filter performance, and a means of validating the parameters selected by likelihood maximization. This step is very important because a minimized AIC does not necessarily correspond to a well-tuned filter.

From these results, a number of potentially rewarding future directions were identified, which focused on selecting an appropriate process model, the need for spatiotemporal variation of model parameters, handling the problem's high dimensionality, and introducing additional information to further constrain the inverse solution.

Chapter 3

Kalman-Filter-Based EEG Source Localization with a Spatially-Varying Process Model: 1D Simulations

Abstract

Electroencephalographic (EEG) source localization enables the brain's distributed dynamics to be imaged with high temporal resolution. However, estimating the actual current sources requires the solution of an ill-posed dynamic inverse problem. Recently, there has been growing interest in dynamic inverse solutions for source reconstruction, as they can incorporate dynamical models that describe brain activity. The Kalman filter (KF) is one such algorithm currently receiving attention, and is the focus of this chapter. Motivated by the need to better understand KF-based EEG source localization, this study applies two KF algorithms — one a standard linear filter, the other spatially whitened — to simulated EEG data generated using a telegrapher's equation and a simplified volume conductor model. To better reflect true brain dynamics, the simulated EEG was given spatial variation by allowing two parameters of the telegrapher's equation to be functions of position. Estimation of the spatial variations was incorporated into the filtering algorithm by allowing the same parameters in the process model of both KFs to have spatial variations of the same functional form. Both filters were found to reliably reconstruct the simulated current vectors, while the accuracy of the estimated model parameter spatial profiles was more variable. This is due to unmodeled inputs which can distort pa-

rameter estimates, meaning that their values should be interpreted with caution. Future algorithms will benefit from explicit handling of any inputs. Despite this finding, this study demonstrates that spatially varying process models can improve filter performance, and provides further validation of the spatially whitened Kalman filter (SWKF) algorithm used in Chapters 2 and 4 (and elsewhere) to solve the EEG inverse problem for a realistic head model. The importance of adequate temporal sampling when using discretized damped wave process models was also identified. These findings are directly applicable to other spatiotemporal systems where KFs are used to solve inverse problems.

3.1 Introduction

Electroencephalography (EEG) measures the brain's electrical activity through the recording of scalp voltages. These measurements result from large-scale synchronized activation of cortical pyramidal cells [104], and probe brain dynamics with higher temporal resolution (milliseconds), but lower spatial resolution (centimeters), than other imaging techniques including functional magnetic resonance imaging (fMRI) and positron emission tomography (PET) [133]. The spatiotemporal resolution of various neuroimaging modalities was displayed in Fig. 1.1 in Chapter 1. A popular approach for processing EEG data is to apply source localization algorithms. These techniques solve the EEG inverse problem to estimate the location, magnitude, and time course of the neuronal sources that produce the observed scalp voltages. However, solving this inverse problem is nontrivial as it is both dynamic (the states to be estimated vary in time) and ill-posed (any given dataset could be explained by an infinite number of source configurations). The non-uniqueness is due to volume conduction effects and the mismatch between the small number of spatial measurements ($\approx 10^2$) and the number of states to be estimated ($\approx 10^4$). To obtain unique inverse solutions, additional constraints must be imposed on the sources and head model. A variety of approaches have been used to solve the EEG inverse problem. For a comprehensive review of these see [9, 32, 54, 62, 97, 143].

Solutions to the EEG inverse problem fall into two main categories. The first type are 'equivalent current dipole' methods, in which the activity is modeled by a relatively small number of focal sources at locations assumed a priori or estimated from the data (see [39, 77, 99, 123] for examples). The second group of tech-

niques, with which this chapter is concerned, are 'linear distributed' approaches (also known as 'imaging' methods), in which the sources are modeled by a three-dimensional grid of dipoles throughout the head volume. The linear distributed inverse solutions can be further categorized into two broad groups, namely instantaneous and dynamic techniques. Instantaneous methods calculate each estimate using only the data available at the current instant of time. Examples include the minimum norm algorithm and its variants (e.g., LORETA) [53, 60, 107]. More recently, dynamic inverse solutions have been developed which incorporate information from multiple times into the estimation process [4, 8, 11, 23, 24, 26, 29, 30, 42, 45–47, 49–51, 78, 88, 116, 126, 129, 137, 147]. These techniques are the focus of this chapter, and have two important advantages over instantaneous methods: (i) the additional information further constrains the inverse solution, and (ii) they provide a natural framework for introducing the growing number of dynamic models describing brain activity (e.g., [7, 27, 66, 119]) into EEG source localization.

Several approaches for introducing dynamic constraints have already been investigated. One commonly used technique is the introduction of a temporal smoothness term, which has been successfully applied to regularization [23,126] and Bayesian estimation [8, 26] methods. Another strategy is to use dynamic models, either inferred from signal analysis or derived from physiology, for describing source behavior. These can then be used in various estimation schemes. A number of dynamic models of varying complexity have been employed including; random walk [24, 129] and autoregressive [29, 30, 147] models, damped wave equations [11, 45–47, 50, 88], as used in this thesis, and neural mass models [78, 116].

To incorporate dynamic models into EEG source localization, we employ the Kalman filter (KF), which is a widely used model-based technique for recursively estimating states in dynamical systems from indirect and uncertain measurements [10, 55, 92]. The KF is becoming an increasingly important tool for solving inverse problems in spatiotemporal systems; for instance, it has already been used to solve dynamic inverse problems in a variety of biomedical imaging areas including: electrical impedance tomography (EIT) [72], inverse electrocardiography (ECG) [15], single photon emission computed tomography (SPECT) [76], diffuse optical tomography (DOT) [31], fMRI [64], diffusion MRI [112], and neural tractography [89]. Recently, a number of studies using KF-based EEG source localization have appeared [4, 11, 29, 30, 45–47, 49–51, 88], and related particle [129] and

local linearization [116] filters have also been investigated.

The work presented here is motivated by the findings of the study described in Chapter 2 (published in [11]) which evaluated the performance of an existing KF-based EEG source localization technique [47] using standard diagnostic tests. A key issue identified in Chapter 2 stems from the fact that brain activity, and consequently the EEG, exhibits spatiotemporal variation; e.g., the alpha rhythm typically increases in strength and frequency from anterior to posterior, and diminishes when the eyes are opened [101]. However, many dynamic inverse solutions, including [47], do not model such features of brain dynamics, using only space- and time-invariant process models. In Chapter 2 this deficiency manifested itself as spatial variation in the KF's performance. We argue that allowing spatially varying model parameters will improve both the accuracy of the inverse solution and measures of filter performance. The filter's current method for system identification, using likelihood maximization, could be expanded to accommodate more model parameters, as demonstrated with a similar technique in [116].

A difficulty associated with KF-based inverse solutions operating in high dimensional state spaces is their large computational cost. To overcome this, the algorithm in [47] used a state-space transformation to reformulate the filtering problem as a coupled set of low-dimensional KFs running in parallel. Methods like this are important for reducing computational costs, but a thorough understanding of their effects is needed. For instance, the inverse solution's performance has not been verified for spatially varying brain dynamics, for either simulated or clinical EEG data. Furthermore, the accuracy of the parameter estimates needs clarification as the transformation means the filter is not modeling the sources directly. This clearly has implications for interpreting physically meaningful parameters, such as the model's spatial coupling term (which describes interactions between neighboring points), which was consistently found to have a very small value in Chapter 2. Also the effect of the temporal sampling rate, which sets an upper bound on wave velocity estimates via the Courant condition [21], needs to be investigated.

In this chapter, we address the issues above using a one-dimensional (1D) simulation of the linear distributed inverse problem. The source dynamics are modeled by a telegrapher's equation with parameters whose spatial profile is based on physiology [101], and of a known functional form. This reduces the complexity of the problem, and offers several advantages over a whole-brain inverse solution, includ-

ing: a modest-sized state-space, meaning that for the purposes of comparison, both a standard linear KF (LKF) and the spatially whitened KF (SWKF) can be applied to the data, simplified geometry and volume conductor model, and straightforward introduction of spatially varying model parameters.

The remainder of this chapter is structured as follows. Section 3.2 introduces a 1D simulation of the linear distributed EEG inverse problem. In Sec. 3.3, the model describing the current source dynamics is outlined. Section 3.4 presents the two KF-based source localization algorithms used in this chapter. In Sec. 3.5, the simulation studies are described, along with the likelihood maximization technique for filter tuning, and the tests which evaluate filter performance. The results are presented in Sec. 3.6. Section 3.7 discusses these results and makes some concluding remarks.

3.2 EEG Inverse Problem Formulation

We begin by outlining an EEG simulation, similar to those in [23, 126], to which our KF-based inverse solutions will be applied. The simulation is set up by defining a continuous current vector field $\boldsymbol{j}(\boldsymbol{r}, t)$, where $\boldsymbol{r}$ and t denote space and time, respectively. The solution space, which simulates the cortical surface, is modeled as a 1D line of current dipoles, which is discretized into N_v grid points (voxels) $\boldsymbol{r}_v$, $v = 1, ..., N_v$, with a spatial separation Δx, as shown in Fig. 3.1. Time is discretized into N_k points t_k, $k = 1, ..., N_k$, with time step Δt. Discretized points are indicated by v and k here, rather than $\boldsymbol{r}_v$ and t_k. Each voxel contains a single current dipole vector $\boldsymbol{j}(v, k)$ which is perpendicular to the voxel array. The state vector for the entire system has dimension N_v and is written

$$\boldsymbol{J}(k) = [\boldsymbol{j}(1, k), \ldots, \boldsymbol{j}(N_v, k)]^T. \tag{3.1}$$

The currents $\boldsymbol{j}(v, k)$ produce the EEG signal which is recorded at N_c electrode sites. These electrodes are also arranged in a 1D array parallel to the voxels with an interelectrode distance (IED) and voxel-to-electrode separation (VES) as shown in Fig. 3.1. If the EEG voltage at a single electrode is denoted by $y(c, k)$, where c is an electrode label, the observation vector containing the voltages in all EEG channels is

$$\boldsymbol{Y}(k) = [y(1, k), \ldots, y(N_c, k)]^T. \tag{3.2}$$

Here voltages refer to average reference (the average voltage over all channels at k is subtracted from each channel).

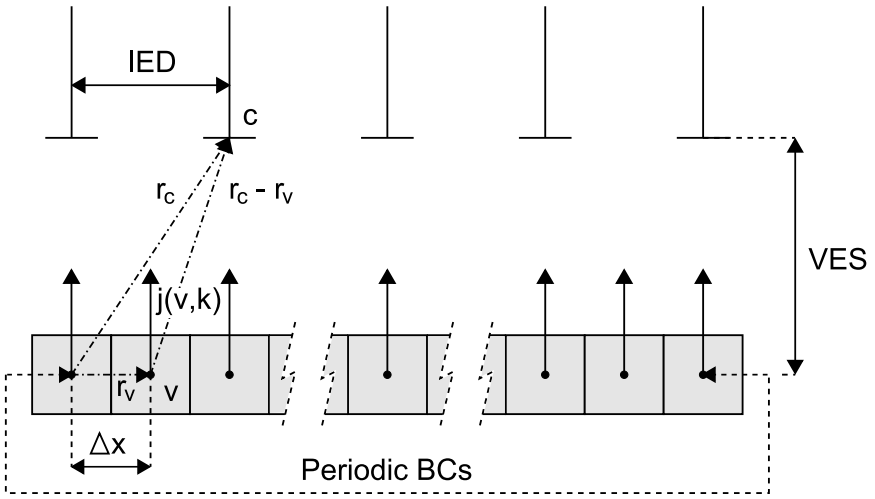

Figure 3.1: Schematic diagram of 1D EEG simulation. An infinite, homogeneous volume conductor is assumed, which maps the single current vector $j(v,k)$ at each voxel v to a voltage at each electrode c. The origin is at the leftmost voxel's center, and vectors r_v, r_c, and $r_c - r_v$, which are used to construct the Green function Eq. (3.4), are shown for an arbitrary voxel and electrode pair. Periodic boundary conditions (BCs) connect the terminal voxels.

For the spatially discretized model used here, the forward model for a single electrode is [104]

$$y(c,k) = \sum_{v=1}^{N_v} \boldsymbol{G}(c,v) \cdot \boldsymbol{j}(v,k), \tag{3.3}$$

where $\boldsymbol{G}$ is the Green function that contains all the geometric and conductive information about the volume conductor, and weights the contribution of the current vector at v to the voltage at electrode c. In this study, we assume the idealized case of sources in an infinite, homogeneous medium of scalar conductivity γ, which has the following Green function [58, 103]

$$\boldsymbol{G}(c,v) = \frac{\boldsymbol{r}_c - \boldsymbol{r}_v}{4\pi\gamma|\boldsymbol{r}_c - \boldsymbol{r}_v|^3}, \tag{3.4}$$

where $\boldsymbol{r}_c$ and $\boldsymbol{r}_v$ are the position vectors from the origin to electrode c and voxel v, respectively.

The observation equation that relates all the current vectors to be estimated to all the measured EEG signals is

$$\boldsymbol{Y}(k) = \boldsymbol{K}\boldsymbol{J}(k) + \boldsymbol{\epsilon}(k), \tag{3.5}$$

where the $N_c \times N_v$ matrix $\boldsymbol{K}$, often referred to as the lead field matrix (LFM)

or the observation model, maps the current vectors to voltages at the electrodes and can be obtained from Eq. (3.4). The term $\boldsymbol{\epsilon}(k)$ is a N_c-dimensional vector of observational noise, which is assumed to be white, Gaussian, and unbiased, with covariance matrix $\boldsymbol{C}_{\boldsymbol{\epsilon}}$, and uncorrelated between all pairs of sensors, with equal variance $\sigma_{\boldsymbol{\epsilon}}^2$ at every electrode, so

$$\boldsymbol{C}_{\boldsymbol{\epsilon}} = \sigma_{\boldsymbol{\epsilon}}^2 \boldsymbol{I}_{N_c}, \tag{3.6}$$

where $\boldsymbol{I}_{N_c}$ is the $N_c \times N_c$ identity matrix.

Equation (3.5) cannot be inverted directly, due to the large ratio of solution points to measurements. Hence, the inverse problem can only be solved by introducing additional constraints.

3.3 Dynamic Model

The process model describes the spatiotemporal evolution of the current vectors, and takes the form of a telegrapher's equation [110]. This equation was used in Chapter 2 and [47], and is chosen again because: (i) it is physiologically plausible; (ii) contains an explicit temporal resonance; (iii) is consistent with experiments demonstrating smooth waves of cortical activity [3,13,35,37,146]; and (iv) has successfully described the spatiotemporal propagation of neuronal activity in a variety of systems [119]. The process model is given by

$$\left(\frac{\partial^2}{\partial t^2} + 2\zeta(\boldsymbol{r})\omega_n(\boldsymbol{r})\frac{\partial}{\partial t} + [\omega_n(\boldsymbol{r})]^2 - b^2\nabla^2\right)\boldsymbol{j}(\boldsymbol{r},t) = \boldsymbol{u}(\boldsymbol{r},t) + \boldsymbol{\eta}(\boldsymbol{r},t), \tag{3.7}$$

where $\omega_n = 2\pi f_n$ is the natural frequency, ζ the fractional damping coefficient, b the wave velocity, $\boldsymbol{u}(\boldsymbol{r},t)$ is a deterministic input and $\boldsymbol{\eta}(\boldsymbol{r},t)$ is a dynamical (process) noise term. The parameters f_n, ζ, and b are restricted to values greater than zero as per the telegrapher's equation definition [110]. Note that in Chapter 2, which followed the approach used in [47], no bounds were imposed on the model parameters. To allow the investigation of spatially varying source dynamics, the model parameters f_n and ζ are given spatial dependences, whose forms are discussed in Sec. 3.5. Note that the deterministic input $\boldsymbol{u}(\boldsymbol{r},t)$ is only used for driving the simulated cortex and is not known to, or explicitly modeled by, either KF in this study. Conversely, the process noise $\boldsymbol{\eta}(\boldsymbol{r},t)$ is not used when generating simulated data, but instead provides the KFs with a means of compensating for any process model

inaccuracies and the unmodeled deterministic drive. These features are consistent with the KFs used in Chapter 2 and [47]. To implement the KFs, we discretize Eq. (3.7) with respect to space and time using finite difference approximations for the derivatives, giving

$$\begin{aligned}\boldsymbol{j}(v,k) = a_1(v)\boldsymbol{j}(v,k-1) + a_2(v)\boldsymbol{j}(v,k-2) \\ + a_3(v)[\boldsymbol{L}\boldsymbol{J}(k-1)]_v + \boldsymbol{u}_L(v,k) + \boldsymbol{\eta}_L(v,k),\end{aligned} \tag{3.8}$$

at each voxel, where $\boldsymbol{L}$ is a 5-point discrete 1D Laplacian operator of dimensions $N_v \times N_v$ defined as [1]

$$\boldsymbol{L} = \left(\boldsymbol{I}_{N_v} - \frac{\boldsymbol{N}}{\kappa}\right), \tag{3.9}$$

where $\kappa = 1.25$, and $\boldsymbol{N}$ is a $N_v \times N_v$ matrix with element $\boldsymbol{N}(v,v') = 0.5$ if v' is immediately adjacent to v, or 0.125 if v' is located two voxels from v. All other entries in $\boldsymbol{N}$ are 0. The $[\boldsymbol{J}]_v$ operator selects the element of $\boldsymbol{J}$ that corresponds to grid point v. The terms $\boldsymbol{u}_L(v,k)$ and $\boldsymbol{\eta}_L(v,k)$ denote the deterministic input and process noise at v, respectively. The spatially dependent parameters in Eq. (3.8) are

$$a_1(v) = \frac{2 - [\omega_n(v)\Delta t]^2}{1 + \zeta(v)\omega_n(v)\Delta t}, \tag{3.10}$$

$$a_2(v) = \frac{\zeta(v)\omega_n(v)\Delta t - 1}{1 + \zeta(v)\omega_n(v)\Delta t}, \tag{3.11}$$

$$a_3(v) = -\frac{5(b\Delta t)^2}{4(\Delta x)^2[1 + \zeta(v)\omega_n(v)\Delta t]}. \tag{3.12}$$

Periodic boundary conditions are imposed; i.e., voxels $v = 1$ and $v = N_v$ are considered to be adjacent to each other.

We also mention that for the discretized process model Eq. (3.8) to be stable, the Courant number ν [21], which for a 1D hyperbolic equation is given by

$$\nu = \frac{b\Delta t}{\Delta x}, \tag{3.13}$$

should be less than $\nu_{\text{max}} = \sqrt{2}$ for a 5-point approximation of the Laplacian term [98]. This constraint places an upper limit on the wave velocity that can be used in the simulations or estimated from the data for a given spatiotemporal grid. To ensure numerical stability in practice, ν must be below this value by an appreciable margin — typically $10 - 20\%$ in our experience. This issue is discussed and explored in Sec. 3.5.

3.4 Kalman-Filter-Based Inverse Solution

In this section, the two KF-based source localization algorithms used in this study are introduced; one a standard LKF [10, 55, 92], and the other a SWKF [47]. The inverse solutions from the optimal LKF serve as a reference when analyzing those from the less accurate, but computationally more efficient, SWKF.

3.4.1 Linear Kalman Filter

We now outline the standard LKF equations. For more details, the reader should consult [10, 55, 92]. Before describing the algorithm a notational convention is defined: $x(k_1|k_2)$ indicates an estimate of some quantity x computed at time k_1, based on all observations available at time k_2, where $k_1 \geqslant k_2$. Due to the discretized process model Eq. (3.8), we introduce an expanded state vector

$$\boldsymbol{J}_2(k) = [\boldsymbol{j}(1,k), \boldsymbol{j}(1,k-1), \ldots, \boldsymbol{j}(N_v,k), \boldsymbol{j}(N_v,k-1)]^T, \tag{3.14}$$

which has dimension $2N_v$ and is used by both filters in this chapter. As a result, a new LFM with dimensions $N_c \times 2N_v$ is required for the LKF and defined as follows:

$$\boldsymbol{K}_2 = \boldsymbol{K} \otimes (1 \quad 0), \tag{3.15}$$

where $\otimes$ indicates Kronecker multiplication.

The filtering cycle begins with the prediction equations

$$\boldsymbol{J}_2(k|k-1) = \boldsymbol{F}\boldsymbol{J}_2(k-1|k-1), \tag{3.16}$$

$$\boldsymbol{P}_2(k|k-1) = \boldsymbol{F}\boldsymbol{P}_2(k-1|k-1)\boldsymbol{F}^T + \boldsymbol{C}_{\boldsymbol{\eta}}, \tag{3.17}$$

where $\boldsymbol{F}$ is the $2N_v \times 2N_v$ state transition matrix, which can be obtained from Eq. (3.8) when the deterministic and stochastic input terms are ignored, and $\boldsymbol{P}_2$ is the $2N_v \times 2N_v$ state covariance matrix. The process noise is assumed to be white, Gaussian, and unbiased, and fixed in space and time, with covariance $\sigma^2_{\boldsymbol{\eta}}$, and $2N_v \times 2N_v$ covariance matrix $\boldsymbol{C}_{\boldsymbol{\eta}}$ of the form suggested in [47]

$$\boldsymbol{C}_{\boldsymbol{\eta}} = \sigma^2_{\boldsymbol{\eta}}(\boldsymbol{L}^T\boldsymbol{L})^{-1} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{3.18}$$

where $\boldsymbol{L}$ is given by Eq. (3.9), with $\kappa = 1.26$ to ensure $\boldsymbol{L}$ is not close to singular [44]. The matrix $(\boldsymbol{L}^T\boldsymbol{L})^{-1}$ in Eq. (3.18) was normalized so all elements on its main diagonal are equal to unity. This covariance matrix was selected because: (i) correlations between the activity of neighboring voxels are expected, and need to be modeled, and (ii) it produces the diagonal process noise covariance matrix required when the spatial whitening transformation is applied in Sec. 3.4.2.

The LKF's next steps involve computing the predicted EEG $\boldsymbol{Y}$, the innovation $\boldsymbol{\Delta Y}$, its covariance $\boldsymbol{R}$, and the Kalman gain $\boldsymbol{W}$ as follows:

$$\boldsymbol{Y}(k|k-1) = \boldsymbol{K}_2\boldsymbol{J}_2(k|k-1), \tag{3.19}$$

$$\boldsymbol{\Delta Y}(k) = \boldsymbol{Y}(k) - \boldsymbol{Y}(k|k-1), \tag{3.20}$$

$$\boldsymbol{R}(k|k-1) = \boldsymbol{K}_2\boldsymbol{P}_2(k|k-1)\boldsymbol{K}_2^T + \boldsymbol{C}_{\boldsymbol{\epsilon}}, \tag{3.21}$$

$$\boldsymbol{W}(k) = \boldsymbol{P}_2(k|k-1)\boldsymbol{K}_2^T\boldsymbol{R}(k|k-1)^{-1}. \tag{3.22}$$

The filtering cycle is then completed by calculating the state estimate and its associated covariance matrix:

$$\boldsymbol{J}_2(k|k) = \boldsymbol{J}_2(k|k-1) + \boldsymbol{W}(k)\boldsymbol{\Delta Y}(k), \tag{3.23}$$

$$\boldsymbol{P}_2(k|k) = \boldsymbol{P}_2(k|k-1) + \boldsymbol{W}(k)\boldsymbol{R}(k|k-1)\boldsymbol{W}(k)^T. \tag{3.24}$$

3.4.2 Spatially Whitened Kalman Filter

The SWKF is now briefly described. Further details regarding its development can be found in Chapter 2 and [47]. The SWKF decomposes the $2N_v$-dimensional filtering problem solved by the LKF into a computationally more efficient set of N_v coupled two-dimensional (2D) KFs. However, this requires $\boldsymbol{C}_{\boldsymbol{\eta}}$ to be diagonal; an assumption that is typically not justified. Therefore, the SWKF operates in a transformed (Laplacianized) state-space, $\tilde{\boldsymbol{J}}(k) = \boldsymbol{L}\boldsymbol{J}(k)$, where $\kappa = 1.26$ for all $\boldsymbol{L}$ in this filter. This transformation, known as 'spatial whitening', was selected since

spatial differentiation reduces spatial correlations between neighboring voxels. It also produces the required diagonal $2N_v \times 2N_v$ process noise covariance matrix

$$\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}} = \boldsymbol{L}_2 \boldsymbol{C}_{\boldsymbol{\eta}} \boldsymbol{L}_2^T = \sigma_{\tilde{\boldsymbol{\eta}}}^2 \boldsymbol{I}_{N_v} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{3.25}$$

when $\boldsymbol{C}_{\boldsymbol{\eta}}$, as defined in Eq. (3.18), is switched to the Laplacianized state-space. In Eq. (3.25), $\boldsymbol{L}_2 = \boldsymbol{L} \otimes \boldsymbol{I}_2$ is the $2N_v \times 2N_v$ Laplacian operator for the expanded state vector Eq. (3.14), and $\sigma_{\tilde{\boldsymbol{\eta}}}^2$ denotes the process noise covariance in the transformed state-space. Now the filtering problem can be reduced to a 2D KF running at every voxel, with each local filter maintaining the state vector

$$\tilde{\boldsymbol{j}}_{SW}(v,k) = [\tilde{\boldsymbol{j}}(v,k), \tilde{\boldsymbol{j}}(v,k-1)]^T, \tag{3.26}$$

and 2×2 state covariance matrix $\tilde{\boldsymbol{p}}(v,k)$ in the Laplacianized state-space. The same form of dynamic model Eq. (3.8) is assumed to govern $\tilde{\boldsymbol{J}}$, which is a reasonable assumption since $\boldsymbol{J}$ and $\tilde{\boldsymbol{J}}$ have similar spatiotemporal properties, although $\tilde{\boldsymbol{J}}$ has less spatial structure.

Again, the filtering cycle starts with the prediction equations. Rather than applying the process model at each voxel individually, the smaller state-space, compared to the whole-brain inverse solutions in Chapters 2 and 4, allows the predicted current vectors to be computed at all voxels using $\boldsymbol{F}$ as follows:

$$\tilde{\boldsymbol{J}}_2(k|k-1) = \boldsymbol{F} \tilde{\boldsymbol{J}}_2(k-1|k-1), \tag{3.27}$$

where $\tilde{\boldsymbol{J}}_2(k) = \boldsymbol{L}_2 \boldsymbol{J}_2(k)$ and contains $\tilde{\boldsymbol{j}}_{SW}(v,k)$ for all voxels.

The local predicted state covariance is approximated as

$$\tilde{\boldsymbol{p}}(v,k|k-1) = \boldsymbol{A}_{SW}(v) \tilde{\boldsymbol{p}}(v,k-1|k-1) \boldsymbol{A}_{SW}(v)^T + \boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L}, \tag{3.28}$$

where the local temporal state transition matrix is

$$\boldsymbol{A}_{SW}(v) = \begin{pmatrix} a_1(v) & a_2(v) \\ 1 & 0 \end{pmatrix}, \tag{3.29}$$

and $\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L}$ is the local dynamical noise covariance matrix given by

$$\boldsymbol{C}_{\tilde{\boldsymbol{\eta}}_L} = \begin{pmatrix} \sigma_{\tilde{\boldsymbol{\eta}}}^2 & 0 \\ 0 & 0 \end{pmatrix}. \tag{3.30}$$

Once the predicted current dipoles and their covariances have been computed at every voxel, we predict the observed scalp voltages from the global state vector:

$$\boldsymbol{Y}(k|k-1) = \tilde{\boldsymbol{K}}_{\boldsymbol{2}}\tilde{\boldsymbol{J}}_2(k|k-1), \tag{3.31}$$

where $\tilde{\boldsymbol{K}}_2$ is the LFM for the SWKF, which undoes the spatial whitening transformation

$$\tilde{\boldsymbol{K}}_2 = \boldsymbol{K}\boldsymbol{L}^{-1} \otimes (1 \quad 0). \tag{3.32}$$

Next, the innovation $\boldsymbol{\Delta Y}$ is computed using Eq. (3.20) and its covariance is approximated by

$$\boldsymbol{R}(k|k-1) = \sum_{v=1}^{N_v} \boldsymbol{Q}(v)\tilde{\boldsymbol{p}}(v,k|k-1)\boldsymbol{Q}(v)^T + \boldsymbol{C}_{\boldsymbol{\epsilon}}, \tag{3.33}$$

where $\boldsymbol{Q}(v)$ is the $N_c \times 2$ LFM for v, consisting of the two columns from $\tilde{\boldsymbol{K}}_2$ that correspond to the vth voxel. The $2 \times N_c$ Kalman gain matrix for voxel v is then

$$\boldsymbol{w}(v,k) = \tilde{\boldsymbol{p}}(v,k|k-1)\boldsymbol{Q}(v)^T\boldsymbol{R}(k|k-1)^{-1}. \tag{3.34}$$

The filtering cycle is then completed by calculating the local state estimate and its corresponding covariance matrix,

$$\tilde{\tilde{\boldsymbol{j}}}_{SW}(v,k|k) = \tilde{\tilde{\boldsymbol{j}}}_{SW}(v,k|k-1) + \boldsymbol{w}(v,k)\boldsymbol{\Delta Y}(k), \tag{3.35}$$

$$\tilde{\boldsymbol{p}}(v,k|k) = [\boldsymbol{I}_2 - \boldsymbol{w}(v,k)\boldsymbol{Q}(v)]\tilde{\boldsymbol{p}}(v,k|k-1), \tag{3.36}$$

respectively. Applying Eqs (3.35) and (3.36) to all voxels generates the inverse solution for time point k. To obtain the actual current dipole estimates, we undo the spatial whitening transformation via

$$\boldsymbol{J}(k|k) = \boldsymbol{L}^{-1}\tilde{\boldsymbol{J}}(k|k). \tag{3.37}$$

The associated $N_v \times N_v$ covariance matrix for the actual current dipoles at every voxel is given by

$$\boldsymbol{P}(k|k) = \boldsymbol{L}^{-1}\tilde{\boldsymbol{P}}(k|k)(\boldsymbol{L}^T)^{-1}, \tag{3.38}$$

where $\tilde{\boldsymbol{P}}(k|k)$ denotes the diagonal $N_v \times N_v$ covariance matrix for all voxels in the transformed state space. The diagonal of $\tilde{\boldsymbol{P}}(k|k)$ is a vector containing the upper-left element of $\tilde{\boldsymbol{p}}(v,k|k)$ for each voxel. The remaining entries in $\tilde{\boldsymbol{P}}(k|k)$ are zeros as a result of spatial whitening, which removes off-diagonal covariances.

3.5 Details of Simulation Study

This section begins by outlining the simulated EEG datasets used in this study. Then the maximum likelihood-based method for filter tuning is described and tests for evaluating filter performance are summarized.

3.5.1 Simulated EEG Data

The evaluation of the whole-brain inverse solution presented in Chapter 2 identified a number of factors affecting filter performance. The simulated datasets used in this study were designed to investigate four of these factors believed to be particularly important, namely: (i) unmodeled inputs; (ii) spatially varying dynamics and model parameters; (iii) temporal undersampling; and (iv) the spatial whitening transformation. To perform this investigation six groups of simulated EEG data were generated using the process model Eq. (3.8), with each group having its own source dynamics. The first two groups (TU and TV) are produced by the system's transient dynamics, while the next two (DU and DV) are driven by a spatiotemporal input, and the final groups (CU1 and CU2), which also have driven dynamics, are used to investigate temporal undersampling. The letter U in a dataset's name indicates the simulated current dipoles were produced using spatially uniform model parameters, while the letter V indicates spatially varying parameters.

Before introducing the simulated datasets individually, the parameters common to all are discussed and summarized in Table 3.1. Firstly, the spatiotemporal discretization (Δx and Δt) was selected to match typical values from clinical source localization studies, while N_v was chosen so the length of the simulated cortex is similar to human anatomy. The model parameters f_n and ζ, which are specified for the simulations and estimated by the KFs, are given either a uniform or Gaussian spatial profile. The uniform profile mirrors the one used in Chapter 2 and allows for comparison with the Gaussian profile, which is selected for its ability to produce source dynamics with realistic spatial variation. Furthermore, using a spatial profile of a known functional form permits the investigation of spatially varying model parameters using the existing filtering framework, with the addition of only a small number of parameters. In contrast, the wave velocity is assumed to be spatially uniform for all datasets. This spatial invariance is a reasonable approximation [104], and simplifies the investigation of temporal undersampling. The magnitude of b —

Table 3.1: Simulated current dipole parameters for transient (TU and TV) and driven (DU and DV) source dynamics, and temporal undersampling (CU1 and CU2). The parameters N_c and $\sigma_{\boldsymbol{\epsilon}}^2$ are shown for simulated EEG where IED $= 0.02$ m and SNR $= 10$. Due to the introduction of an arbitrary scaling (see Sec. 3.5.1), the values of $\sigma_{\boldsymbol{\epsilon}}^2$ are much larger than expected from physiology.

Dataset	Δt (s)	N_k	Δx (m)	N_v	N_c	b (m s^{-1})	ν	c_{f_n} (Hz)	d_{f_n} (Hz)	μ_{f_n} (m)	σ_{f_n} (m)	c_ζ	d_ζ	μ_ζ (m)	σ_ζ (m)	$\sigma_{\boldsymbol{\epsilon}}^2$ (V^2)
TU	0.004	251	0.005	101	26	1	0.8	10	—	—	—	0.01	—	—	—	3.77×10^6
TV	0.002	251	0.005	51	13	1	0.4	5	5	0.125	0.04	0.1	-0.09	0.125	0.04	1.84×10^7
DU	0.004	251	0.005	101	26	1	0.8	10	—	—	—	0.05	—	—	—	6.43×10^6
DV	0.004	251	0.005	101	26	1	0.8	5	5	0.370	0.04	0.3	-0.25	0.370	0.04	2.28×10^7
CU1	0.004	1001	0.005	51	13	1	0.8	10	—	—	—	0.1	—	—	—	1.33×10^7
CU2	0.001	1001	0.005	51	13	4	0.8	10	—	—	—	0.1	—	—	—	1.37×10^9

which is 1 m s^{-1} for all datasets except CU2 where $b = 4$ m s^{-1} — was selected so the Courant condition is satisfied for the spatiotemporal grids used. This is shown in Table 3.1 where $\nu \leqslant 0.8$ for all six datasets, which is comfortably below its maximum allowed value of $\sqrt{2}$. However, cortical activity is known to propagate at up to 10 m s^{-1} [105, 119]. This discrepancy was present in Chapter 2, where the spatiotemporal discretization restricted b to values well below 10 m s^{-1}, and motivates the temporal undersampling study outlined in Sec. 3.5.1.3.

We now introduce a generalized spatial Gaussian $g(v)$ which is used throughout the simulation study, where $g = f_n$ and $g = \zeta$ for the natural frequency and damping coefficient spatial profiles, respectively, while $g = i$ denotes the initial condition and $g = u_s$ for the spatial component of the deterministic drive. The Gaussian is given by

$$g(v) = c_g + \frac{d_g n_g}{\sigma_g \sqrt{2\pi}} \exp\left(-\frac{[(v-1)\Delta x - \mu_g]^2}{2\sigma_g^2} \right), \tag{3.39}$$

where c_g is a constant offset, n_g is a scaling coefficient introduced for f_n and ζ so the Gaussian term's amplitude is d_g, while μ_g and σ_g are the mean and standard deviation, respectively. Note that the periodic boundary conditions must be taken into account when computing the distance between v and μ_g.

A final general comment concerns the magnitude of the simulated current dipoles and scalp voltages in this study, which are much larger than physiological values [104]. This results from the initial conditions in Sec. 3.5.1.1, the drive terms in Secs 3.5.1.2 and 3.5.1.3, and the conductivity in Sec. 3.5.1.4 not being scaled to match physiology. However, the net effect of these features of the simulation on both the current dipoles and scalp voltages is simply to introduce multiplication by a constant, which does not alter the results or their interpretation.

3.5.1.1 Transient Source Dynamics

The first two sets of simulated current dipoles (TU and TV) are generated by the source model's transient response to spatial Gaussian initial conditions, with the drive term set to zero. These datasets provide algorithm validation as both KFs should perform well when the unmodeled drive term is absent. The first transient dataset (TU) assumes f_n and ζ are spatially uniform, so only c_{f_n} and c_ζ are required [see Table 3.1 and Figs 3.2(g) and 3.2(h)], while the initial conditions are given by a spatial Gaussian with parameters $c_i = 0$, $d_i = n_i = 1$, $\mu_i = 0.25$ m, and $\sigma_i = 0.01$

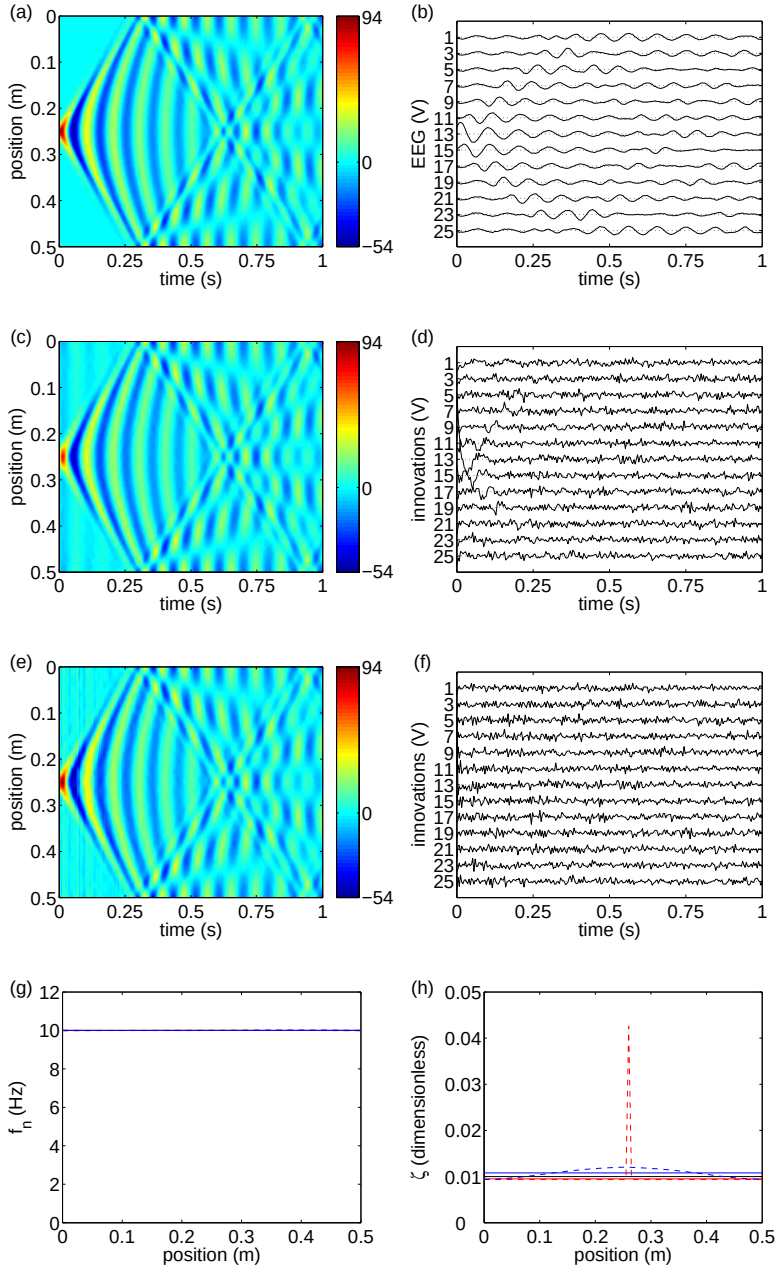

Figure 3.2: Results of inverse solution for transient source dynamics with uniform parameters (TU), where IED $= 0.02$ m and SNR $= 10$. The simulated source dynamics and associated EEG are shown in (a) and (b), respectively. The LKF's inverse solution when fitting a Gaussian profile for f_n and ζ is shown in (c), and the innovations in (d). Frames (e) and (f) show the same results for the SWKF. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 9.6 and 9.1 relative to (b), respectively. Frame (g) shows the simulated (solid black line) and estimated spatial profiles for f_n. In (g) the red and blue lines indicate the LKF- and SWKF-selected profiles, respectively, which are solid when uniform and dashed when Gaussian. The spatial profile for ζ is shown in (h) using the same key as (g).

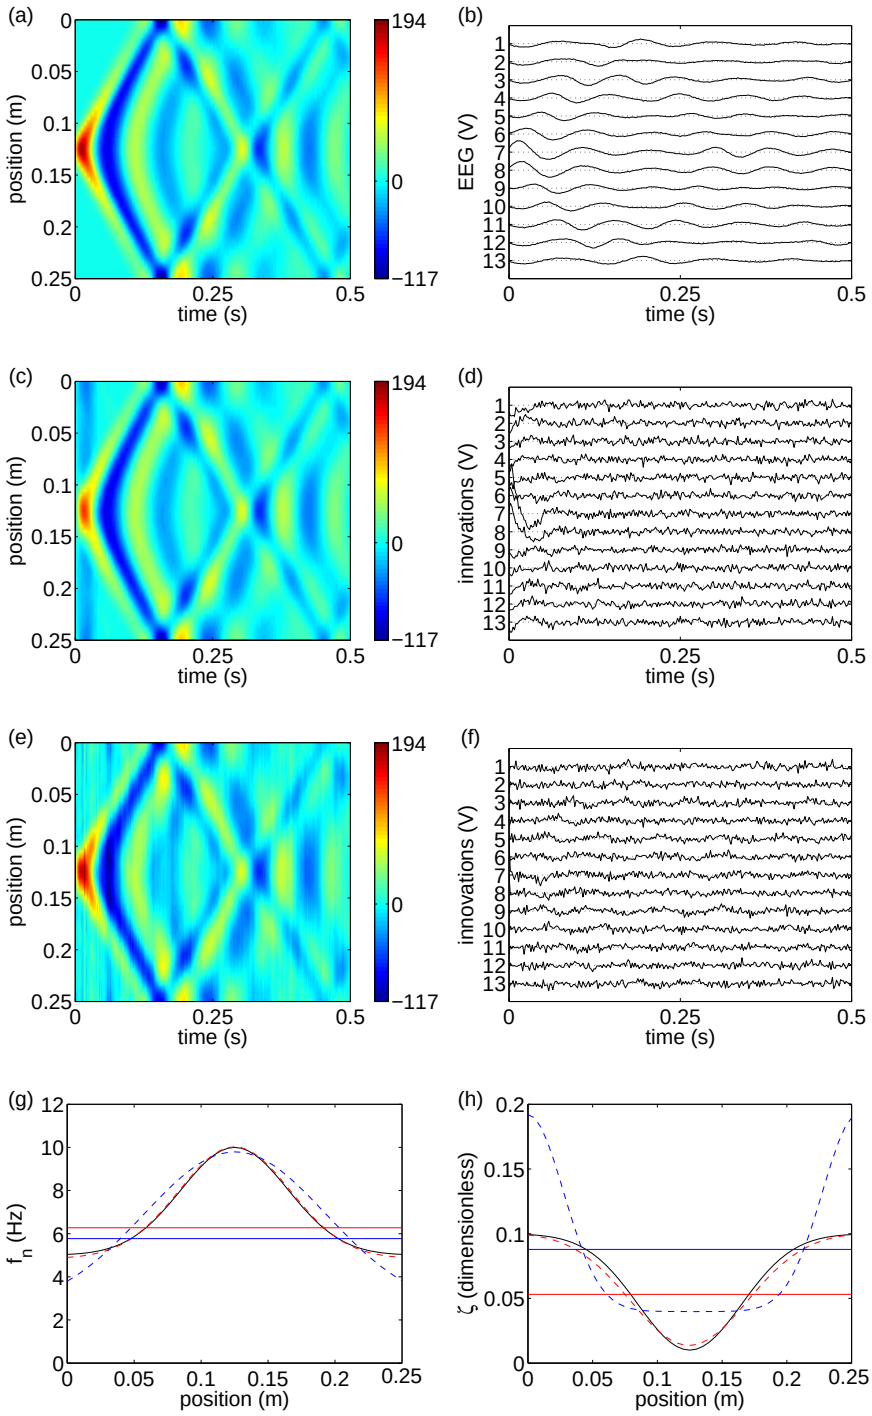

Figure 3.3: Results of inverse solution for transient source dynamics with spatially varying parameters (TV), where IED $= 0.02$ m and SNR $= 10$. This figure uses the same format as Fig. 3.2. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 9.8 and 7.9 relative to (b), respectively.

m. For the second transient dataset (TV), f_n and ζ have a Gaussian profile with the parameters listed in Table 3.1 and plotted in Figs 3.3(g) and 3.3(h), while the initial conditions are the same as TU, except $\mu_i = 0.125$ m. The resulting current dipoles for TU and TV are shown in Figs 3.2(a) and 3.3(a), respectively.

3.5.1.2 Driven Source Dynamics

The next pair of simulated current dipoles (DU and DV) are produced by driving the dynamic model with a spatiotemporal input. This term simulates thalamic input to the cortex and is similar to the one used in Chapter 2. The temporal component of the drive $u_t(k)$ is modeled by a sum of randomly-phased sine functions evenly spaced between 0 and 20 Hz, and is defined as

$$u_t(k) = \sum_{m=1}^{N_f} \sin[2\pi f(m) k\Delta t + \psi(m)], \tag{3.40}$$

where N_f is the number of frequency components, $f(m)$ is the frequency of oscillation [0 Hz $\leqslant f(m) \leqslant 20$ Hz], $\psi(m)$ is a random phase offset [$-\pi \leqslant \psi(m) \leqslant \pi$], and the frequency spacing is 0.1 Hz. The drive term's spatial component u_s is modeled by a Gaussian, and multiplied by u_t to obtain the spatiotemporal drive from Eq. (3.8)

$$\boldsymbol{u}_L(v,k) = u_s(v)u_t(k). \tag{3.41}$$

Note that all drive terms in this chapter have parameter values for u_s of $c_{u_s} = 0$, $d_{u_s} = n_{u_s} = 1$, and are multiplied by 0.01 so the simulated current dipoles for both the transient and driven datasets are the same order of magnitude.

The first dataset with driven dynamics (DU) assumes f_n and ζ are spatially uniform [see Table 3.1 and Figs 3.4(g) and 3.4(h)], with one drive, where u_s has parameters $\mu_{u_s} = 0.25$ m, and $\sigma_{u_s} = 0.01$ m. This produces a highly resonant source of 10 Hz oscillations which spread over the cortex as shown in Fig. 3.4(a). The other dataset (DV) uses Gaussian profiles for f_n and ζ [see Table 3.1 and Figs 3.5(g) and 3.5(h)] to reproduce some typical features of the alpha rhythm, namely the spatial variation of its frequency and strength [101]. In order to generate these dynamics, the simulated cortex is driven by two inputs which excite regions with different parameter values. This results in two alpha sources with different spatiotemporal characteristics, which more closely models what is observed in real EEG recordings. To see an example of the spatiotemporal variation of clinical EEG, the reader

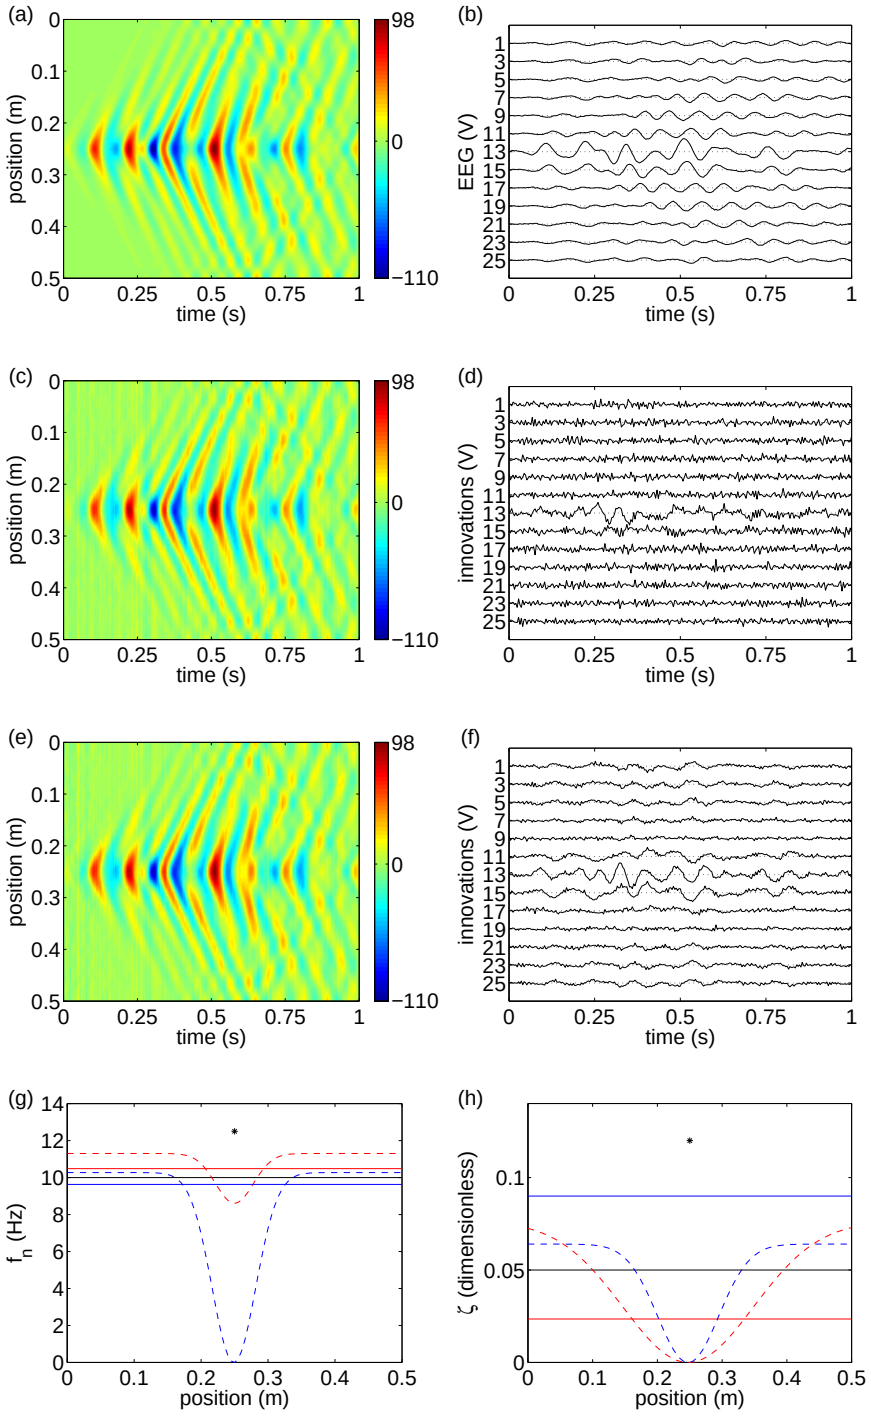

Figure 3.4: Results of inverse solution for driven source dynamics with uniform parameters (DU), where IED $= 0.02$ m and SNR $= 10$. This figure uses the same format as Fig. 3.2. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 6.2 and 3.5 relative to (b), respectively. In frames (g) and (h) the spatial position of the drive term's center is indicated by an asterisk.

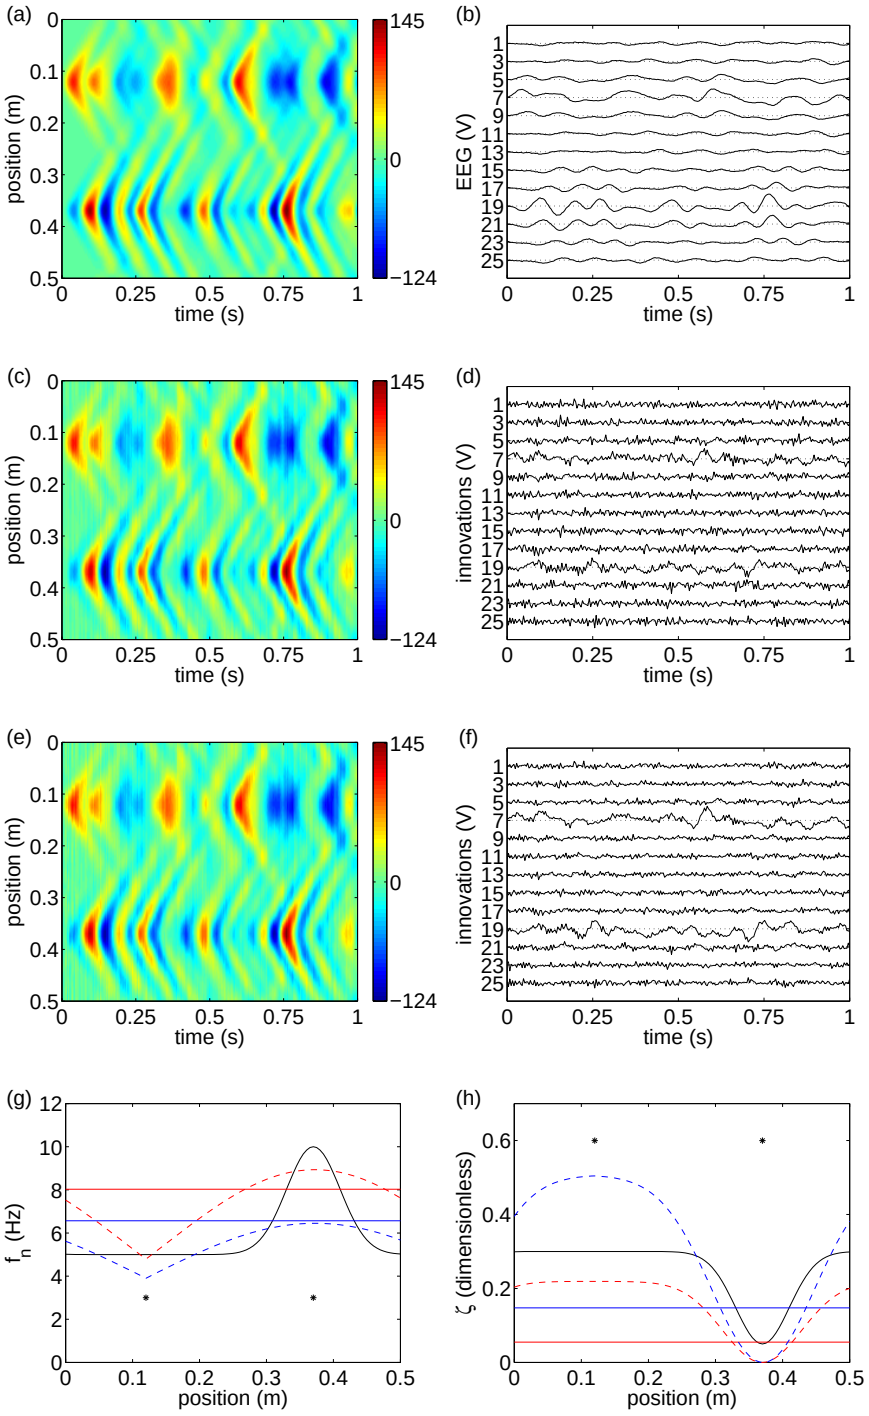

Figure 3.5: Results of inverse solution for driven source dynamics with spatially varying parameters (DV), where IED $= 0.02$ m and SNR $= 10$. This figure uses the same format as Fig. 3.2. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 6.8 and 5.5 relative to (b), respectively. In frames (g) and (h) the spatial position of each drive term's center is indicated by an asterisk.

is referred to the recording used in Chapters 2 and 4. The first drive term is centered on the Gaussian peaks for f_n and ζ, which generates a strong 10 Hz resonance, with u_s specified by $\mu_{u_s} = 0.37$ m and $\sigma_{u_s} = 0.01$ m. The second drive is located where the parameter profiles are flat, with $\mu_{u_s} = 0.12$ m and $\sigma_{u_s} = 0.015$ m, and produces a more damped 5 Hz oscillation. Figure 3.5(a) displays the spatially varying dynamics present in DV.

3.5.1.3 Temporal Undersampling of EEG Data

As stated in Sec. 3.3, the Courant condition must be satisfied for the discretized process model Eq. (3.8) to be stable. This means the maximum wave velocity b_{max} for the spatiotemporal grid used in this 1D simulation study is

$$b_{\text{max}} = \frac{\nu_{\text{max}} \Delta x}{\Delta t}. \tag{3.42}$$

For the 3D whole-brain inverse solution used in Chapters 2 and 4, $\nu_{\text{max}} = \sqrt{3}$ which sets $b_{\text{max}} \approx 1$ m s^{-1} for the spatiotemporal grid employed (see Chapter 4 for further details regarding this calculation). This upper bound is well below the previously stated 10 m s^{-1} for the propagation of cortical activity and therefore prevents b from being estimated correctly, and may even explain the small spatial terms seen in Chapter 2. To facilitate further investigation, this simulation study attempts to recreate this important scenario, which we referred to as temporal undersampling — a name derived from the fact that b, and to a large extent Δx, are fixed, and hence one possible solution is to increase the temporal sampling rate to satisfy the Courant condition (see Sec. 3.7 for a more detailed discussion about potential solutions). To examine the impact of temporal undersampling on inverse solution performance, two simulated datasets (CU1 and CU2) were produced. These datasets have driven dynamics, and assume f_n and ζ are spatially uniform (see Table 3.1). CU1 and CU2 each possess one drive term of the form described in Sec. 3.5.1.2, with the spatial parameters $\mu_{u_s} = 0.125$ m and $\sigma_{u_s} = 0.01$ m being identical for both datasets. While $\nu = 0.8$ for both CU1 and CU2, the datasets display quite different dynamics due to the choice of b. Like datasets DU and DV, the dynamics of the simulated current densities in CU1 result primarily from the transient response of the telegrapher's equation to the drive term's input. Contrastingly, the dynamics in CU2 — owing to its higher b value — are more heavily dominated by the drive term, and less so by the intrinsic dynamics of the telegrapher's equation itself. Designing the

datasets in this way allows the effect of temporal undersampling on inverse solution performance to be assessed for two realistic scenarios: (i) where the dynamics are well approximated by a telegrapher's equation whose parameter values are similar to those used in the simulation (CU1), and (ii) the more complicated situation where significantly different model parameters provide the best-fit to the source dynamics (CU2). The latter scenario provides important insight into inverse solution performance when fitting a telegrapher's equation to activity displaying dynamics of a different form. Furthermore, these findings will help interpret the whole-brain inverse solutions described in Chapters 2 and 4, since temporal undersampling and a mismatch between the form of the proposed and actual dynamics will both affect algorithm performance in this setting.

Finally, temporal undersampling was reproduced by creating six new datasets which use only every second, third, or fourth time point from both CU1 and CU2. This undersampling prevents correct estimation of b, as $b_{\max}$ falls below its true value. Figures 3.6(a) and 3.7(a) display the current dipoles for CU1 and CU2, respectively, when every time step is used.

3.5.1.4 Computing EEG Data

To produce the simulated EEG data, we simply multiply $\boldsymbol{J}$ by the LFM $\boldsymbol{K}$ and apply average reference. For the volume conductor's Green function Eq. (3.4) we assumed unit scalar conductivity; i.e., $\gamma = 1$ S m^{-1}, VES $= 0.015$ m [104], and IED $= 0.02$ m, to match the high density electrode arrays recommended for source localization [121]. We also produced EEG data with an IED $= 0.05$ m, which is equivalent to the commonly used 10-20 system [121]. Next, white Gaussian observation noise was added to the EEG to give signal to noise ratios (SNRs) of 2, 5, 10, and 20 in terms of standard deviations of the voltages. This means 8 EEG datasets were constructed for each of the simulated current dipoles, TU, TV, DU, and DV. However, only one EEG dataset (IED $= 0.02$ m, SNR $= 10$) was produced from CU1 and CU2. In this chapter we focus largely on the high density data (IED $= 0.02$ m), with a SNR $= 10$, which is typical for EEG recordings [121]. For the IED $= 0.02$ m and SNR $= 10$ datasets, the number of electrodes N_c along with the simulated observation noise covariance σ_ϵ^2 are listed in Table 3.1, and the EEG for TU, TV, DU, DV, CU1, and CU2 is plotted in Figs 3.2(b), 3.3(b), 3.4(b), 3.5(b), 3.6(b), and 3.7(b), respectively.

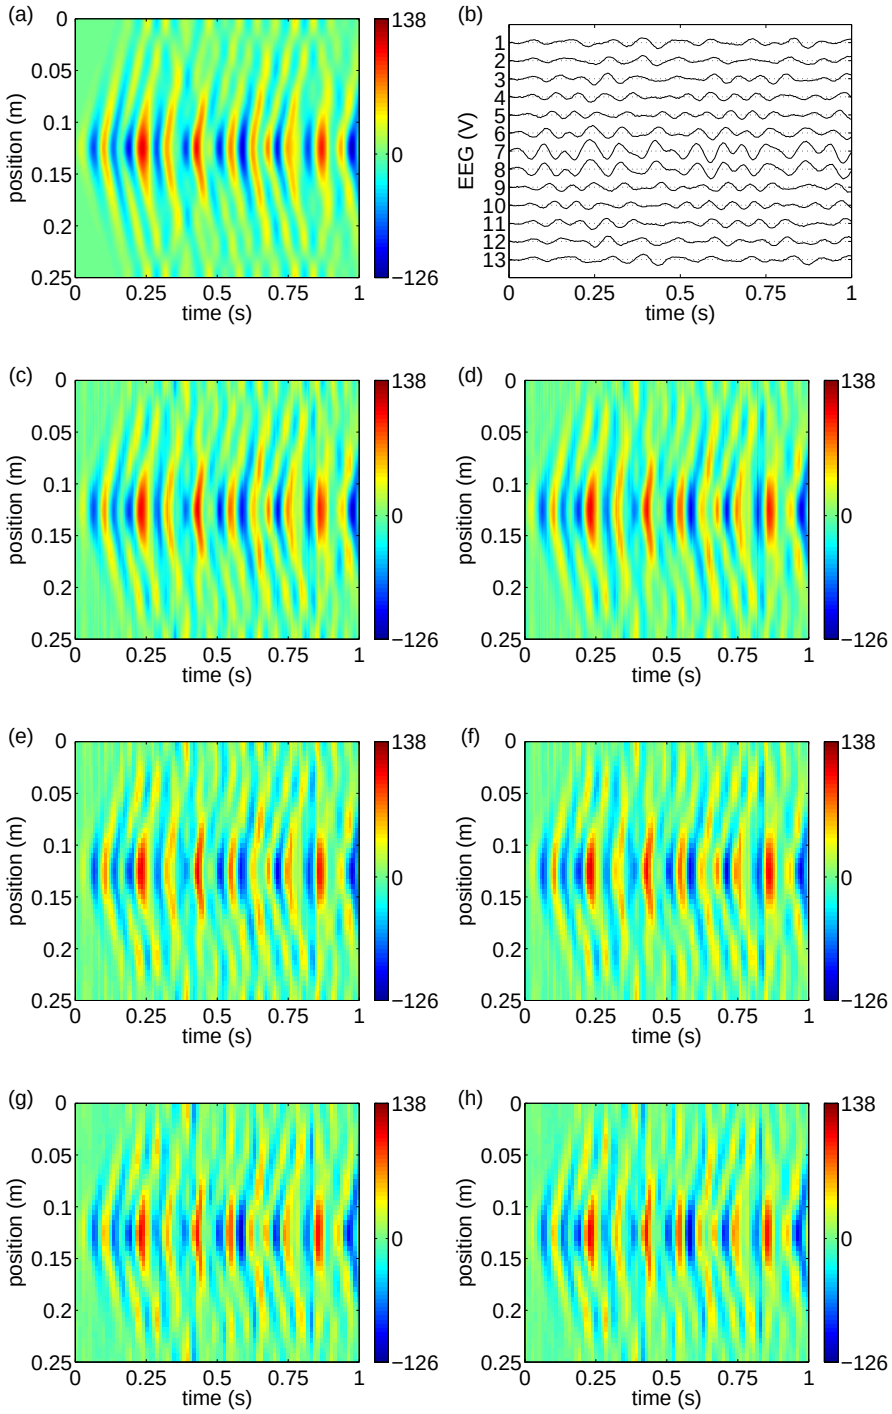

Figure 3.6: Results of inverse solution for the temporal undersampling study using dataset CU1. To improve visualization all frames display the first second of data only. The simulated source dynamics and associated EEG are shown in (a) and (b), respectively. The inverse solutions for the LKF and SWKF when using all time points are shown in (c) and (d), respectively. The remaining two rows follow the same format as (c) and (d), except that (e) and (f) display the results when every second time point is used, and (g) and (h) the results when every third point is supplied.

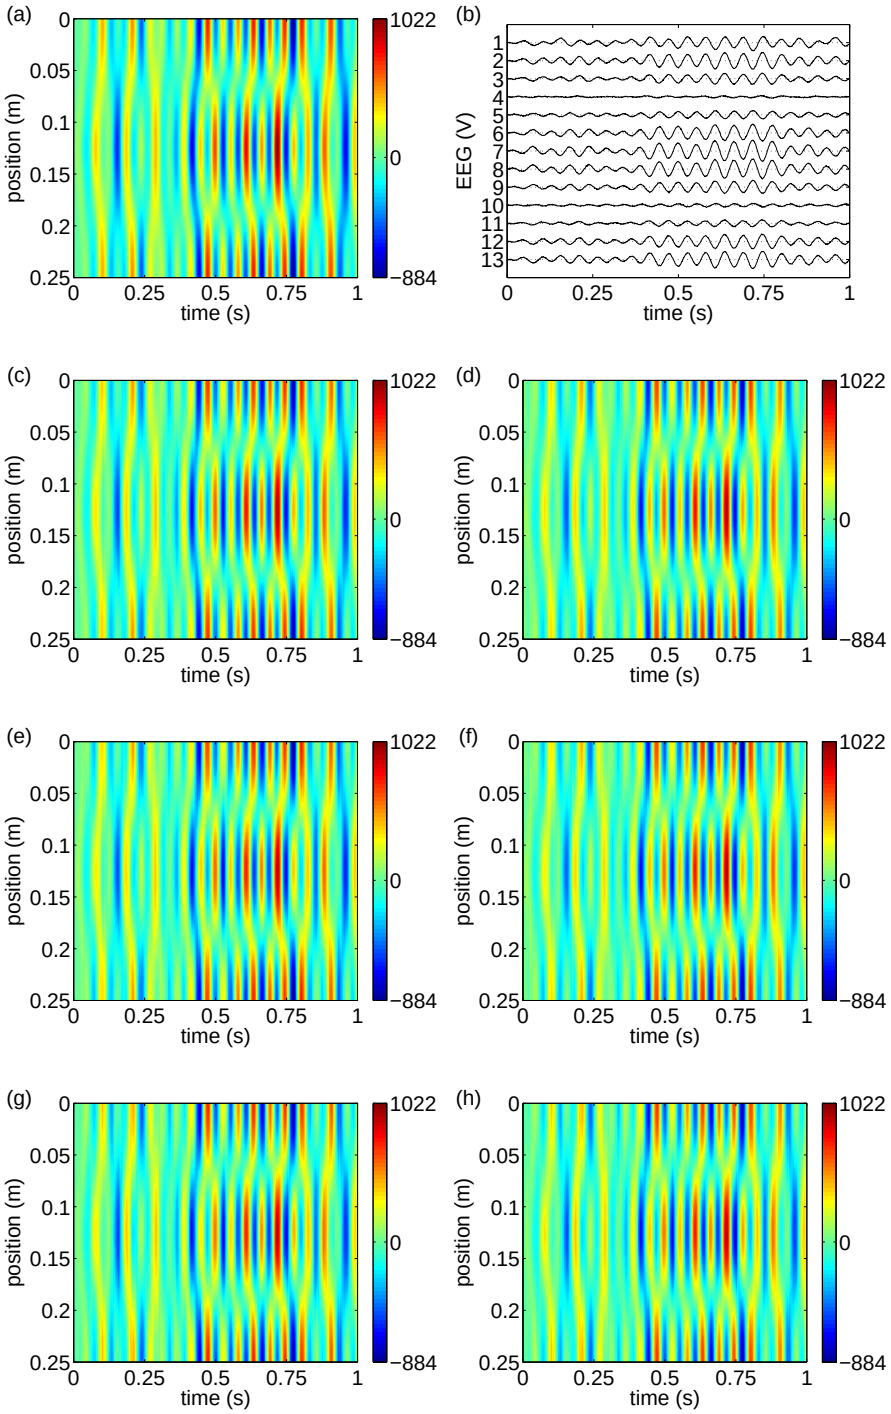

Figure 3.7: Results of inverse solution for the temporal undersampling study using dataset CU2. The simulated source dynamics and associated EEG are shown in (a) and (b), respectively. The inverse solutions for the LKF and SWKF when using all time points are shown in (c) and (d), respectively. The remaining two rows follow the same format as (c) and (d), except that (e) and (f) display the results when every second time point is used, and (g) and (h) the results when every third point is supplied.

3.5.2 Parameter Estimation

To select optimal values for the process model parameters and noise covariances, we used numerical minimization of the Akaike Information Criterion (AIC) [2]. The AIC, which is closely related to the log-likelihood, estimates the distance between the process model and the true one, and has been successfully employed for filter tuning in Chapter 2 and [47]. To compute the AIC we begin by defining two parameter vectors; one for fitting spatially uniform parameters $\boldsymbol{\vartheta}_u = (c_{f_n}, c_\zeta, b)$, the other for parameters with a Gaussian profile $\boldsymbol{\vartheta}_g = (c_{f_n}, d_{f_n}, \mu_{f_n}, \sigma_{f_n}, c_\zeta, d_\zeta, \mu_\zeta, \sigma_\zeta, b)$. The log-likelihood ($\log L$) for the entire EEG time series is

$$\log L(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) = -\frac{1}{2}\sum_{k=1}^{N_k}[\log_e |\boldsymbol{R}(k|k-1)| + \boldsymbol{\Delta Y}(k)^T \boldsymbol{R}(k|k-1)^{-1}\boldsymbol{\Delta Y}(k) + N_c \log_e(2\pi)], \tag{3.43}$$

where $|.|$ denotes absolute value of the matrix determinant. The AIC is then

$$\text{AIC}(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) = -2\log L(\boldsymbol{\vartheta}, \sigma_{\boldsymbol{\epsilon}}^2, \sigma_{\boldsymbol{\eta}}^2) + 2[\dim(\boldsymbol{\vartheta}) + 2], \tag{3.44}$$

where $\dim(\boldsymbol{\vartheta})$ indicates the number of parameters in $\boldsymbol{\vartheta}$, which is increased by two as we need to fit the noise covariances from the data. Unlike Chapter 2 and [47], constrained optimization is used here. This takes advantage of our knowledge of the parameters, which are given bounds where possible, as displayed in Table 3.2. To ensure numerical stability of the process model, the upper bound for b is restricted to 80% of its absolute maximum b_{max}.

For each EEG dataset, the optimal parameters are estimated assuming spatially uniform and spatial Gaussian model parameters for both the linear and spatially whitened KFs. The filters' state estimates are initialized by setting $\boldsymbol{J}(0|0)$ and $\tilde{\boldsymbol{j}}_{SW}(v, 0|0)$ to be $\mathbf{0}$ column vectors, and $\boldsymbol{P}(0|0)$ and $\tilde{\boldsymbol{p}}(v, 0|0)$ to be identity matrices. Each model parameter's starting value is randomly selected from an interval spanning $\pm 50\%$ of the parameter's simulated value. This interval is modified by the parameter bounds in Table 3.2 where appropriate. To allow filter transients to pass, the AIC is computed from the 50^{th} time point onwards for all datasets. Every component of this algorithm is implemented in Matlab [91] and run on an IBM ThinkPad R51 (Intel Pentium 1.6 GHz, 1 GB RAM). The AIC minimization is performed by Matlab's 'fmincon' function which finds the minimum of a constrained nonlinear multivariable function using a sequential quadratic programming (SQP)

Table 3.2: Bounds for the KF parameters. These bounds are imposed on the parameter estimates by the optimization step (see Sec. 3.5.2).

Parameter	Lower Bound	Upper Bound
c_{f_n} (Hz)	0	—
$c_{f_n} + d_{f_n}$ (Hz)	0	—
μ_{f_n} (m)	0	$(N_v - 1)\Delta x$
σ_{f_n} (m)	0	—
c_ζ	0	—
$c_\zeta + d_\zeta$	0	—
μ_ζ (m)	0	$(N_v - 1)\Delta x$
σ_ζ (m)	0	—
b (m s^{-1})	0	$0.8b_{\text{max}}$
$\sigma^2_{\boldsymbol{\eta}}$ (A^2 m^2)	0	—
$\sigma^2_{\boldsymbol{\epsilon}}$ (V^2)	0	—

method. For further details regarding SQP, the reader is referred to the Matlab documentation [91] for the 'fmincon' function, and the following standard optimization texts [38, 48]. The convergence criterion used by the optimization algorithm considers an AIC minimum to be reached, and therefore terminates the optimization routine, when the change in the AIC at the next step is less than 1×10^{-10}. The maximum number of filter runs per optimization is set to 2000; with each filter iteration taking approximately 2 minutes to complete, this means optimized parameters can be computed within 67 hours. However, in practice, convergence is typically achieved much sooner.

3.5.3 Evaluating Kalman Filter Performance

Using simulated data means KF performance can be judged by directly comparing simulated and estimated values for the current dipoles and model parameters. This is done in Sec. 3.6. For the current vectors, this comparison can also be performed by calculating the root-mean-squared error (RMSE) as follows:

$$\text{RMSE} = \sqrt{\frac{1}{N_v N_k} \sum_{v=1}^{N_v} \sum_{k=1}^{N_k} [\boldsymbol{j}(v,k|k) - \boldsymbol{j}(v,k)]^2}. \tag{3.45}$$

However, as discussed in Chapter 2, evaluating KF performance in real applications involves testing the statistical properties of the innovation sequence, which should be Gaussian, unbiased, uncorrelated, and have the correct magnitude (i.e., actual and filter-predicted innovation covariances should be the same). In this study the diagnostic tests from Chapter 2 are applied to the innovations of the optimized filters. An additional test statistic $\overline{\boldsymbol{\Delta Y}}_N$ is also introduced, which replaces Step 3 in Sec. 2.5 and encapsulates overall filter noise levels into a single number [10]. To calculate $\overline{\boldsymbol{\Delta Y}}_N$ we begin by computing the normalized square of the innovation

$$\boldsymbol{\Delta Y}_N(k) = \boldsymbol{\Delta Y}(k)^T \boldsymbol{R}(k|k-1)^{-1} \boldsymbol{\Delta Y}(k), \tag{3.46}$$

which will be a χ^2 variable (resulting from squaring a Gaussian random variable), with N_c degrees of freedom and $E[\boldsymbol{\Delta Y}_N(k)] = N_c$ if $\boldsymbol{\Delta Y}$ is Gaussian, unbiased, and the actual and filter-predicted covariances match [10]. The test statistic $\overline{\boldsymbol{\Delta Y}}_N$ is then computed

$$\overline{\boldsymbol{\Delta Y}}_N = \frac{1}{N_k} \sum_{k=1}^{N_k} \boldsymbol{\Delta Y}_N(k), \tag{3.47}$$

and 95% confidence intervals obtained for its expected value, from which we can determine whether the filter's noise levels are correct. If $\overline{\boldsymbol{\Delta Y}}_N$ is above these bounds, noise levels are too low and the filter is termed 'optimistic', while if $\overline{\boldsymbol{\Delta Y}}_N$ is below the lower bound, the noise levels are too high and the filter is known as 'pessimistic'. Such conservative noise covariances can help ensure the filter is robust [70].

3.6 Results

In this section the optimized inverse solutions for the transient (TU and TV), driven (DU and DV), and temporal undersampling (CU1 and CU2) datasets are presented for the LKF and SWKF, fitting both uniform and Gaussian spatial profiles for f_n and ζ. Note that only uniform spatial profiles are fitted to datasets CU1 and CU2. While the inverse solutions were computed for each filter for two IEDs across a range of SNRs (for TU, TV, DU, and DV, but not CU1 and CU2), we focus on the results for the simulated EEG with IED $= 0.02$ m and SNR $= 10$, which are summarized in Tables 3.3 – 3.9 and Figs 3.2 – 3.8. More general comments regarding the algorithms and their performance are made in Sec. 3.7.

3.6.1 Transient Dynamics

We begin by describing the performance of the filters when applied to the simulated EEG resulting from TU [see Fig. 3.2(b) and Sec. 3.5.1.1]. As expected, the LKF and the SWKF perform very well when fitting either uniform or spatially varying parameters. Hence the results of all four filters will be discussed as a single group. Firstly, the inverse solutions estimated the underlying current dipoles very accurately. This is demonstrated by the close agreement between the simulated and estimated current dipoles, which can be observed by: (i) comparing Fig. 3.2(a) with Figs 3.2(c) and 3.2(e), and (ii) looking at the time series shown in Fig. 3.8(a). The small covariance and RMSE values for TU shown in Fig. 3.8(a) and Table 3.3, respectively, also indicate an accurate inverse solution. Estimates of the spatial profiles for f_n and ζ match the uniform parameters used in the simulation, with even the Gaussian parameters selected resulting in uniform spatial profiles [see Figs 3.2(g) and 3.2(h)]. The correct wave velocity was also identified by all four filters. As the process model parameters were estimated very accurately, the process noise covariances have very small values (≈ 0), which indicates an extremely high level of confidence in the selected process model. Each filter also estimated the true value for the observation noise covariance.

To further assess the filters' performance, the innovation sequences were analyzed. An initial assessment can be made from the appearance of the innovation time series. For TU, the innovations shown in Figs 3.2(d) and 3.2(f) appear to be white. To provide a quantitative evaluation, the statistical properties of the innovation sequences were computed. As anticipated, these indicated very well-tuned KFs, since the innovations were Gaussian, unbiased, white, and have the correct magnitude (see Table 3.4). In conclusion, these results clearly demonstrate that the LKF and SWKF perform well under favorable conditions. The results presented from this point forward describe the performance of these algorithms as increasingly more challenging scenarios are presented.

The second dataset investigating transient dynamics is TV, which introduces spatial variation of the model parameters f_n and ζ for the first time. Details of the spatial profiles were provided in Sec. 3.5.1.1, and the simulated current dipoles and EEG are shown in Figs 3.3(a) and 3.3(b), respectively. The results for dataset TV are described in two parts; the LKF first, followed by the SWKF.

Overall the LKF performed well when fitting either the uniform or the Gaussian

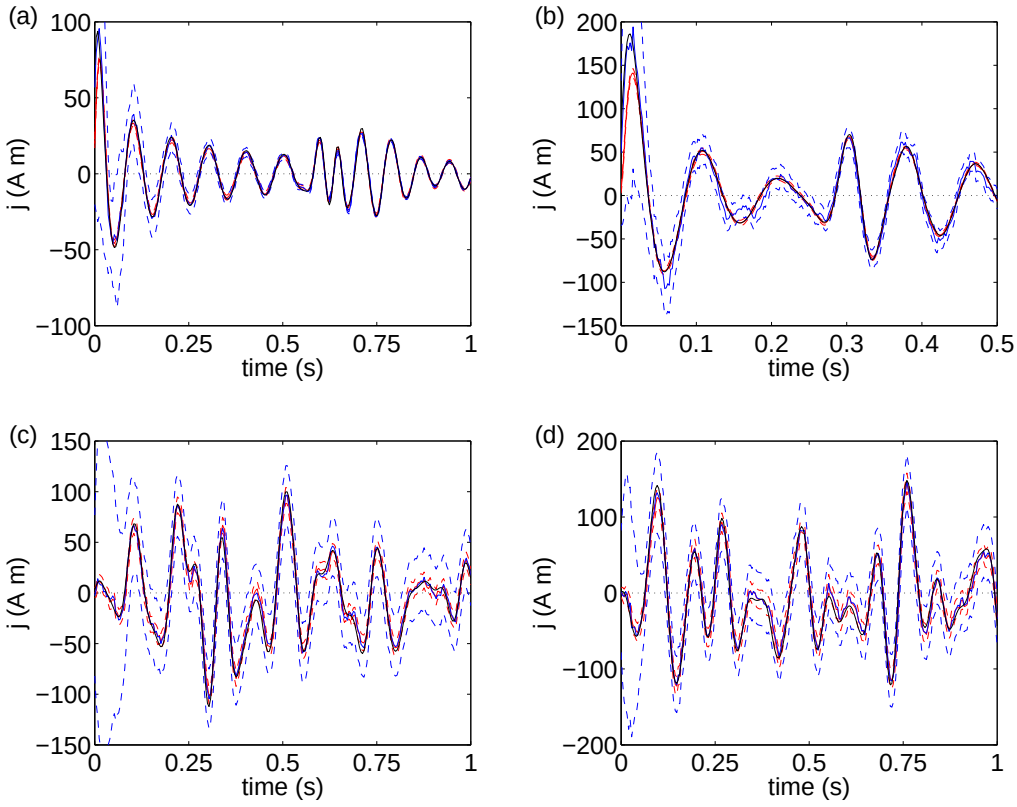

Figure 3.8: Current dipole time series for a voxel v in the datasets TU, TV, DU, and DV with IED $= 0.02$ m and SNR $= 10$. Frame (a) displays the time series for TU at $v = 51$. The solid black line indicates the simulated current vector, while the solid red and blue curves represent the currents estimated by the LKF and SWKF, respectively, when fitting the Gaussian parameter profiles. The dashed lines indicate the 95% CI for the estimates using the same color coding. The CIs were computed from the estimated state covariance matrix at each time step. Frames (b), (c), and (d) show the same information for datasets TV ($v = 26$), DU ($v = 51$), and DV ($v = 75$), respectively.

Table 3.3: Estimated parameter values and performance metrics for the two EEG datasets displaying transient dynamics with IED $= 0.02$ m and SNR $= 10$. For the purposes of comparison, the simulated parameter values used to generate datasets TU and TV are displayed in rows TUsim and TVsim, respectively. Note that all parameters estimated for the SWKF (except $\sigma^2_{\boldsymbol{\epsilon}}$) describe the spatially whitened current dipole $\tilde{\boldsymbol{J}}$, rather than $\boldsymbol{J}$.

Dataset	Filter	c_{f_n} (Hz)	d_{f_n} (Hz)	μ_{f_n} (m)	σ_{f_n} (m)	c_ζ	d_ζ	μ_ζ (m)	σ_ζ (m)	b (m s^{-1})	$\sigma^2_{\boldsymbol{\eta}}$ (A^2 m^2)	$\sigma^2_{\boldsymbol{\epsilon}}$ (V^2)	AIC	RMSE
TUsim	—	10.00	—	—	—	0.0100	—	—	—	1.00	—	3.77×10^6	—	—
TU	LKF	10.00	—	—	—	0.0095	—	—	—	1.00	2.68×10^{-4}	3.77×10^6	94934	1.08
TU	LKF	6.89	3.12	0.26	1.73	0.0094	0.033	0.26	7.93×10^{-4}	1.00	1.74×10^{-4}	3.77×10^6	94910	1.14
TU	SWKF	10.00	—	—	—	0.0108	—	—	—	1.00	0.00	3.77×10^6	95875	1.39
TU	SWKF	10.00	0.026	0.41	0.078	0.0122	-0.0028	0.00	0.106	1.00	0.00	3.77×10^6	95881	1.39
TVsim	—	5.00	5.00	0.125	0.040	0.0100	-0.09	0.125	0.040	1.00	—	1.84×10^7	—	—
TV	LKF	6.27	—	—	—	0.053	—	—	—	1.01	1.54	1.84×10^7	52413	8.86
TV	LKF	4.85	5.18	0.125	0.042	0.101	-0.087	0.124	0.045	1.00	1.19×10^{-4}	1.84×10^7	51553	1.67
TV	SWKF	5.77	—	—	—	0.088	—	—	—	1.01	1.97×10^{-4}	1.84×10^7	53548	9.59
TV	SWKF	2.45	7.34	0.125	0.068	0.0398	0.152	0.00	0.0276	0.992	8.66×10^{-5}	1.84×10^7	53014	8.57

Table 3.4: Results of KF performance evaluation tests for the two EEG datasets displaying transient dynamics with IED $= 0.02$ m and SNR $= 10$. The KF's spatial profile for f_n and ζ is either uniform (U) or Gaussian (G) as indicated. The number of channels whose innovation is Gaussian, and unbiased, according to Kolmogorov-Smirnov (KS) and t-tests, respectively, are shown (i.e., $P > 0.05$). Also displayed is the overall noise level statistic (NLS) Eq. (3.47) and the 95% confidence intervals (CI) for its expected value. Innovation whiteness is characterized by: (i) counting the channels with $< 90\%$ of their autocorrelation (AC) within the 2 standard deviations for a white signal; and (ii) showing the range of the innovation spectral entropy ($\boldsymbol{\Delta Y}$ SE) across all channels.

Dataset	Filter	Profile	KS	t-test	NLS (95% CI)	AC	$\boldsymbol{\Delta Y}$ SE
TU	LKF	U	26	26	27.7 (25.0-27.0)	1	0.974-0.994
TU	LKF	G	26	26	27.6 (25.0-27.0)	1	0.975-0.994
TU	SWKF	U	26	26	25.2 (25.0-27.0)	0	0.974-0.991
TU	SWKF	G	26	26	25.1 (25.0-27.0)	0	0.975-0.991
TV	LKF	U	13	13	13.0 (12.4-13.7)	1	0.958-0.992
TV	LKF	G	13	11	13.6 (12.4-13.7)	1	0.974-0.990
TV	SWKF	U	13	12	13.0 (12.4-13.7)	3	0.849-0.988
TV	SWKF	G	13	13	13.3 (12.4-13.7)	8	0.904-0.988

parameter profiles, although as expected the filter with spatially varying parameters performed slightly better. The LKF reconstructed the current dipoles accurately, as shown by the concordance between the simulated current dipoles shown in Fig. 3.3(a), and the estimated activity which is plotted for the Gaussian parameter fit in Fig. 3.3(c). Looking at the LKF's Gaussian fit further, the time series for a single voxel shown in Fig. 3.8(b) confirms the agreement between the simulated (black line) and estimated activity (red line), and also reveals small error covariances (dashed red lines). When estimating the spatial profiles for f_n and ζ using a uniform profile, the optimization step would select a value between that parameter's maximum and minimum values, as shown in Figs 3.3(g) and 3.3(h). Pleasingly, when fitting f_n and ζ with Gaussian profiles, likelihood maximization chose parameters with an almost identical spatial profile to those from the simulation [again see Figs 3.3(g) and 3.3(h)]. The exact parameter values can be read from Table 3.3. The performance advantage offered by fitting Gaussian parameter profiles over uniform parameters is further evidenced by the reduction of the RMSE, which dropped from 8.86 to 1.67. Both of the LKFs correctly estimated the simulated wave velocity. Turning to the noise covariances; the observation noise was estimated accurately, while the process noise covariance decreased by over four orders of magnitude when fitting the Gaussian profiles rather than the uniform parameters. This finding indicates a large increase in the filter's confidence in the process model, and was expected since this filter implementation selected parameter profiles matching those in the simulated data. The innovations also revealed very well-tuned filters for both the uniform and spatial Gaussian LKFs. The innovations for the LKF fitting Gaussian parameters are shown in Fig. 3.3(d) and appear to be white. This is confirmed by the statistical properties of the innovations for both LKFs which are found to be Gaussian, unbiased, white, and have the correct magnitude (see Table 3.4).

When operating on the TV dataset, the SWKF also produced accurate inverse solutions and reasonable parameter estimates. However, its overall performance was not as good as the LKF; a finding which is plausible given the SWKF is an approximation of the superior LKF algorithm. The inverse solution computed by the SWKF when fitting the Gaussian parameter profile is plotted in Fig. 3.3(e) and shows a fairly accurate reconstruction of the simulated current dipoles [see Fig. 3.3(a)], although the LKF's inverse solution which is displayed in Fig. 3.3(c)

has better fidelity. This observation is supported by Fig. 3.8(b), where the SWKF's estimated time series (blue line) is slightly less accurate and the error covariances larger than the LKF. As already seen with the LKF when estimating uniform spatial parameters, the tuning step for the SWKF also selects values for f_n and ζ that lie between the parameter's maximum and minimum values, albeit different values to those chosen for the LKF [see Figs 3.3(g) and 3.3(h), and Table 3.3]. When fitting Gaussian parameter profiles, the AIC minimization chose a spatial profile for f_n that closely approximates its true profile [see Fig. 3.3(g)]. However, the Gaussian profile proposed for ζ is less accurate, and only captures the parameter's spatial trend, which is shown in Fig. 3.3(h). It may be possible to improve this parameter estimate by running the optimization with a number of initial values, which will generate a probability distribution for that parameter, rather than relying on a single optimization, as done in this study. Such an approach was employed successfully in [116], and could be used in future work. The RMSEs for the SWKF show only a modest improvement when Gaussian parameter profiles are added, which is due to the inaccurate ζ profile. Nonetheless, the RMSEs for the SWKF are similar to that of the LKF when fitting spatially uniform parameters. Once more, the SWKFs correctly estimated the simulated wave velocity and the observation noise covariance. Unlike the LKF, there was only a small decrease in the process noise covariance when spatially varying parameters were introduced; a finding which is again due to the inexact ζ profile. The SWKF fitting uniform parameters produced innovations that were Gaussian, unbiased, of the correct magnitude, and largely uncorrelated, although a decrease in the minimum spectral entropy to ≈ 0.85 was observed (see Table 3.4). Looking at the SWKF with Gaussian parameter profiles, the discrepant ζ profile appears to introduce correlations into some of the innovations plotted in Fig. 3.3(f). This is supported by the quantitative analysis which reveals the innovations to be Gaussian, unbiased, and to have the correct magnitude, although eight channels exhibited significant temporal correlations (see Table 3.4). Overall, these issues are relatively minor and we conclude that the two SWKFs are well-tuned, and produce inverse solutions of a similar quality to the LKF.

In summary, the results presented in this section provide validation of both KF algorithms, as they were found to be well-tuned and produced accurate estimates of both states and parameters for a 1D simulation of the EEG inverse problem displaying transient dynamics. In particular, spatial variation of model parameters was able

to be correctly estimated from the simulated EEG. These same KFs are now applied to simulated EEG produced by driven dynamics. Their performance is described in the following section.

3.6.2 Driven Dynamics

The first dataset investigating driven dynamics (DU) contains a single source and uniform spatial parameters, as outlined in Sec. 3.5.1.2. The simulated current dipoles and EEG for DU are shown in Figs 3.4(a) and 3.4(b), respectively. When the two KFs were applied to the simulated EEG for DU, they were found to produce quite similar results, regardless of the parameter profile employed. Therefore, the results for all four filters will be discussed together. Firstly, the filters estimated the current dipoles very accurately. This is clearly illustrated for the LKF and SWKF — when fitting Gaussian parameter profiles — by the close agreement between the simulated and estimated current dipoles, which is evident when comparing Fig. 3.4(a) with Figs 3.4(c) and 3.4(e). The single voxel time series plotted in Fig. 3.8(c) also demonstrates the accuracy of both algorithms, with the LKF's error covariances being smaller than the SWKF's. Compared to the state estimates, the parameter estimates of all four filters were less accurate (see Table 3.5). This stems from the drive term(s) introduced into the simulated data, which are not explicitly modeled by either filter, and lead to distorted parameter values as the optimization step attempts to compensate for these inputs by maladjustment of the parameter values. This finding is not confined to DU alone, but is seen throughout this study when investigating driven dynamics (see also the results for datasets DV, CU1, CU2). The DU estimates for f_n and ζ are plotted in Figs 3.4(g) and 3.4(h), respectively, and show that both filters select uniform parameters close to the true simulation values, with f_n being estimated more accurately than ζ. As foreshadowed above, when fitting Gaussians for f_n and ζ, the optimization step selects parameter profiles that differ significantly from the simulated ones, particularly around the drive term's location. This is a clear illustration of how unmodeled inputs can distort parameter estimates.

All four filters estimated the wave velocity to within 10% of its actual value (see Table 3.5 for all parameter estimates). Looking again at the accuracy of the inverse solutions, the RMSEs show similar values for all four filters with only a small reduction conferred by spatially varying parameters. This finding further reinforces the

Table 3.5: Estimated parameter values and performance metrics for the two EEG datasets with driven dynamics, IED $= 0.02$ m and SNR $= 10$. For the purposes of comparison, the simulated parameter values used to generate datasets DU and DV are displayed in rows DUsim and DVsim, respectively. Note that all parameters estimated for the SWKF (except $\sigma^2_{\boldsymbol{\epsilon}}$) describe the spatially whitened current dipole $\tilde{\boldsymbol{J}}$, rather than $\boldsymbol{J}$.

Dataset	Filter	c_{f_n} (Hz)	d_{f_n} (Hz)	μ_{f_n} (m)	σ_{f_n} (m)	c_ζ	d_ζ	μ_ζ (m)	σ_ζ (m)	b (m s^{-1})	$\sigma^2_{\boldsymbol{\eta}}$ (A^2 m^2)	$\sigma^2_{\boldsymbol{\epsilon}}$ (V^2)	AIC	RMSE
DUsim	—	10.00	—	—	—	0.0500	—	—	—	1.00	—	6.43×10^6	—	—
DU	LKF	10.48	—	—	—	0.0235	—	—	—	0.932	11.74	6.43×10^6	102591	3.76
DU	LKF	11.30	-2.69	0.250	0.028	0.076	-0.076	0.248	0.101	0.921	10.31	6.43×10^6	102274	3.30
DU	SWKF	9.63	—	—	—	0.0900	—	—	—	0.964	5.93×10^{-3}	6.43×10^6	106411	3.90
DU	SWKF	10.27	-10.27	0.249	0.033	0.064	-0.064	0.247	0.047	1.00	1.06×10^{-3}	6.43×10^6	103956	3.36
DVsim	—	5.00	5.00	0.370	0.040	0.300	-0.250	0.370	0.040	1.00	—	2.28×10^7	—	—
DV	LKF	8.03	—	—	—	0.055	—	—	—	0.854	21.1	2.28×10^7	107971	6.04
DV	LKF	0.00	8.93	0.372	0.226	0.219	-0.219	0.370	0.059	0.874	17.7	2.28×10^7	107696	5.53
DV	SWKF	6.56	—	—	—	0.147	—	—	—	0.897	5.74×10^{-3}	2.28×10^7	111074	5.89
DV	SWKF	0.00	6.45	0.372	0.252	0.507	-0.507	0.371	0.077	0.944	3.63×10^{-3}	2.28×10^7	110343	5.63

Table 3.6: Results of KF performance evaluation tests for the two EEG datasets with driven dynamics, where IED $= 0.02$ m and SNR $= 10$. This table uses the same format as Table 3.4.

Dataset	Filter	Profile	KS	t-test	NLS (95% CI)	AC	$\boldsymbol{\Delta Y}$ SE
DU	LKF	U	26	25	26.2 (25.0-27.0)	5	0.808-0.984
DU	LKF	G	26	25	25.7 (25.0-27.0)	9	0.834-0.983
DU	SWKF	U	26	16	26.4 (25.0-27.0)	18	0.638-0.988
DU	SWKF	G	26	20	26.6 (25.0-27.0)	23	0.658-0.982
DV	LKF	U	26	26	24.5 (25.0-27.0)	8	0.844-0.983
DV	LKF	G	26	24	24.6 (25.0-27.0)	7	0.858-0.989
DV	SWKF	U	26	18	24.4 (25.0-27.0)	17	0.611-0.988
DV	SWKF	G	26	11	24.9 (25.0-27.0)	12	0.699-0.990

view that the major source of inaccuracy is the unmodeled drive term, which cannot be significantly overcome through spatially varying model parameters alone. The process noise covariances confirm this conclusion, with only small decreases seen when spatial variation of f_n and ζ was added. Each filter estimated the observation noise covariance correctly.

Turning to the innovation sequences plotted in Figs 3.4(d) and 3.4(f), there are clearly temporal correlations in some channels, particularly those closest to the unmodeled drive term. These correlations are larger and more widespread for the SWKF. As shown in Table 3.6, the innovations for both filters were found to be Gaussian, unbiased, and to have the correct magnitude. However, as indicated by the time series plots, there are significant correlations present in 9 (out of 26) channels for the LKF, and 23 channels for the less optimal SWKF. This is supported by a concomitant increase in the range of the innovation spectral entropies. Furthermore, employing spatially varying parameters does not improve these measures of filter performance. Therefore, these results demonstrate how the unmodeled drive term can affect algorithm performance, with the greatest impact felt by parameter estimation, rather than state estimation.

The second dataset containing driven dynamics is DV. This dataset poses an even greater challenge for the KFs, because it was produced using Gaussian parameter profiles and two drive terms (see Sec. 3.5.1.2). Figures 3.5(a) and 3.5(b) show

the simulated current dipoles and EEG, respectively. The results for DV follow a similar pattern to DU, with the two KF algorithms producing comparable results for either parameter profile. For this reason, the performance of the four filters are reviewed together once more. Since the interpretation of the results for DV is similar to DU, a detailed discussion is omitted here, and the reader is referred to the earlier exposition for DU.

Starting with the inverse solutions for DV, all four filters were found to reconstruct the current dipoles accurately. This is demonstrated by the concordance between the simulated current dipoles in Fig. 3.5(a) and their estimated counterparts shown in Figs 3.5(c) and 3.5(e). The single voxel time series in Fig. 3.8(d) also shows this close agreement between the simulated and estimated activity for both filters. However, as seen previously for DU, the LKF's error covariances are smaller than the SWKF's. As already seen with DU, parameter estimation is again found to be less reliable than state estimation, with the optimization step choosing inexact parameter profiles for the two filters fitting spatial Gaussians. These discrepant estimates are again due to the unmodeled drive terms. The DV estimates for the f_n and ζ profiles are shown in Figs 3.5(g) and 3.5(h), respectively. These figures show that for both filters employing spatially uniform parameters, the optimization step selects a value between that parameter's maximum and minimum. As intimated above, the estimated Gaussian profiles for f_n and ζ differ from their simulated values for both the LKF and SWKF. However, these profiles do retain some qualitatively similar features such as the location and approximate size of the Gaussian's peak or nadir. The wave velocity was estimated to within 15% of its actual value by all four filters, with the discrepancy again due to the unmodeled inputs. Once more, the RMSEs for all KFs have comparable values, with Gaussian parameters providing only a slight improvement. This result is supported by the limited reduction in process noise covariance when spatially varying parameters are used. These findings reconfirm the unmodeled drive terms as the main source of error. The observation noise covariance was estimated correctly by all four filters.

Looking at the innovations shown in Figs 3.5(d) and 3.5(f), some temporal correlations are clearly visible, especially in the channels nearest the unmodeled inputs. Again, these correlations are more prominent for the SWKF than the LKF. Calculating the statistical properties of the innovations reveals some interesting features (see Table 3.6). The innovations of all four KFs are found to be Gaussian and pos-

sess noise levels larger than anticipated. A filter operating in this regime is termed pessimistic (see Sec. 3.5.3), and can exhibit more robust performance [70]. The innovations for both LKFs are unbiased, although the SWKF innovations are losing this property, in particular the one fitting Gaussian parameters. Confirming the observations from the innovation time series, significant correlations are found in 8 LKF channels, and 17 SWKF channels. The innovation spectral entropies show a spread similar to dataset DU, with lower values for the SWKF yet again.

In conclusion, the results presented here have demonstrated the ability of two KF algorithms to solve the EEG inverse problem for the more realistic case of driven dynamics. For both driven datasets the filters produced accurate state estimates, and generally positive results for the diagnostic tests; indicating that the filters are indeed well-tuned. These findings notwithstanding, the less reliable parameter estimates — regardless of the parameter profile employed — and the considerable correlations present in the innovations for DU and DV indicate the need to model system inputs explicitly. Such a task will be the focus of future work, leading to not only more precise parameter estimates, but further improvement of overall filter performance.

3.6.3 Temporal Undersampling of EEG Data

The first study of temporal undersampling investigates the performance of the filters when applied to CU1. The design of this dataset, which contains a single drive term and uniform parameters, was described in Sec. 3.5.1.3, and the simulated current dipoles and EEG are plotted in Figs 3.6(a) and 3.6(b), respectively. To begin, the inverse solution with uniform parameters was computed for the EEG using all available time points. The results for both filters were similar to one another, and exhibit a pattern much like those described for DU and DV. Therefore, the results for the two algorithms will be discussed together. When using all observations, the filters estimated the current dipoles accurately. This is demonstrated by the close agreement between the simulated and estimated current dipoles, which is evident when comparing Fig. 3.6(a) with Figs 3.6(c) and 3.6(d). As seen previously with DU and DV, both filters selected slightly inaccurate values for the model parameters as a result of the unmodeled drive term (see Table 3.7). However, this discrepancy in the parameter values is relatively small, which indicates that the telegrapher's equation selected by each of the KFs display similar dynamics to the one used to generate the

simulated data. Dataset CU1 — whose dynamics result largely from the transient response of the telegrapher's equation to the drive term — was designed to produce this scenario. The reader is referred back to Sec. 3.6.2 for more detailed discussion about the effect of unmodeled inputs on parameter estimates. The innovations for the LKF and SWKF were found to be Gaussian, unbiased, and to have the correct magnitude (see Table 3.8). Temporal correlations were present in the innovations of both filters, although they were more widespread for the SWKF, where significant correlations were found in 9 (out of 13) channels, compared to only 3 for the LKF.

The effect of temporal undersampling on KF performance will now be described. When the data was undersampled, the quality of the state estimation decreased only marginally for both filters; even when up to three-quarters of the data was discarded. This is demonstrated by the close agreement between the simulated and estimated current densities [Figs 3.6(e) - 3.6(h)] when temporal undersampling was applied. The RMSEs for both filters support this observation, since they only displayed a small increase as less data was used. Conversely, parameter estimation became significantly less accurate when the time series was undersampled.

One important source of this inaccuracy arises from the Courant condition which imposes an upper bound (b_{max}) on the wave velocity estimable from the data. As Δt is increased, b_{max} decreases (see Eq. 3.42) until it falls below b's true value, whereafter b is unable to be estimated correctly through the optimization step — this was the case for CU1 when every second, third, or fourth time point was used. In this situation, we found both filters would select wave velocities equal or very close to b_{max} (see Table 3.7).

In the setting of temporal undersampling, we would like to know how the filters maintain the accuracy of their state estimates. One possibility is to adjust its estimates of ζ and f_n to maintain the system's effective damping ζ_{eff} near its value for the true model parameters. Effective damping can be approximated as follows:

$$\zeta_{\text{eff}} \approx \omega_n \zeta + \frac{b}{l}, \tag{3.48}$$

where the spatial width of the drive term is $l = 2\sigma_{u_s}$. As b_{max} decreases, a compensatory rise in the $\omega_n \zeta$ term of Eq. (3.48) could reasonably be expected. However, no clear trend was observed for either filter, with the value of $\omega_n \zeta$ remaining approximately constant as less data was used. This finding suggests the involvement of another compensatory mechanism. The only parameter showing a reliable trend was the process noise covariance, which consistently increased for both filters as

Table 3.7: Estimated parameter values and performance metrics for temporal undersampling study with IED $= 0.02$ m and SNR $= 10$. The number in parentheses in the leftmost column indicates that either all (1), or every second (2), third (3) or fourth (4) time points from the original dataset are used in a particular filter run. An asterisk next to a wave velocity estimate denotes its value is either at, or within 5%, of its upper bound, as determined by the Courant condition with a 20% margin of safety. The caret adjacent to the estimate of f_n for CU1 (4) indicates its value is at the upper bound used by the optimizer (see Sec. 3.6.3 for more details). For the purposes of comparison, the simulated parameter values used to generate datasets CU1 and CU2 are displayed in rows CU1sim and CU2sim, respectively. Note that all parameters estimated for the SWKF (except $\sigma^2_{\boldsymbol{\epsilon}}$) describe the spatially whitened current dipole $\tilde{\boldsymbol{J}}$, rather than $\boldsymbol{J}$.

Dataset	Filter	c_{f_n} (Hz)	c_ζ	b (m s^{-1})	$\sigma^2_{\boldsymbol{\eta}}$ (A^2 m^2)	$\sigma^2_{\boldsymbol{\epsilon}}$ (V^2)	AIC	RMSE
CU1sim	—	10.00	0.100	1.00	—	1.33×10^7	—	—
CU1 (1)	LKF	11.07	0.070	0.829	47.4	1.33×10^7	229053	7.37
CU1 (2)	LKF	11.69	0.058	0.707*	260.0	1.33×10^7	112402	9.30
CU1 (3)	LKF	12.46	0.070	0.471*	835.8	1.33×10^7	72852	9.76
CU1 (4)	LKF	12.00^	0.108	0.351*	1714.3	1.33×10^7	52665	10.05
CU1 (1)	SWKF	10.15	0.152	0.860	0.019	1.33×10^7	257668	7.88
CU1 (2)	SWKF	11.32	0.154	0.687*	0.294	1.33×10^7	125189	9.25
CU1 (3)	SWKF	12.28	0.136	0.471*	1.25	1.33×10^7	80091	9.44
CU1 (4)	SWKF	12.18	0.143	0.350*	2.86	1.33×10^7	57350	9.93
CU2sim	—	10.00	0.100	4.00	—	1.37×10^9	—	—
CU2 (1)	LKF	16.71	0.041	1.368	10.23	1.37×10^9	297492	38.29
CU2 (2)	LKF	16.82	0.041	1.324	90.09	1.37×10^9	141904	48.11
CU2 (3)	LKF	16.16	0.039	1.773	302.97	1.37×10^9	89864	46.62
CU2 (4)	LKF	16.90	0.037	1.238	691.16	1.37×10^9	64028	56.49
CU2 (1)	SWKF	15.44	0.116	0.904	2.94×10^{-3}	1.37×10^9	299114	46.79
CU2 (2)	SWKF	16.21	0.107	0.631	0.03844	1.37×10^9	142965	54.79
CU2 (3)	SWKF	16.30	0.101	0.782	0.164	1.37×10^9	90652	55.68
CU2 (4)	SWKF	16.76	0.086	0.483	0.463	1.37×10^9	64570	60.53

Table 3.8: Results of KF performance evaluation tests for the temporal undersampling study. This table uses the same format as Table 3.4.

Dataset	Filter	Profile	KS	t-test	NLS (95% CI)	AC	$\boldsymbol{\Delta Y}$ SE
CU1 (1)	LKF	U	13	13	11.8 (12.7-13.3)	3	0.812-0.963
CU1 (2)	LKF	U	13	13	12.9 (12.5-13.5)	2	0.796-0.977
CU1 (3)	LKF	U	13	13	13.2 (12.4-13.6)	13	0.816-0.950
CU1 (4)	LKF	U	13	13	13.3 (12.4-13.7)	11	0.853-0.941
CU1 (1)	SWKF	U	13	10	12.8 (12.7-13.3)	9	0.649-0.943
CU1 (2)	SWKF	U	13	13	13.0 (12.5-13.5)	13	0.731-0.925
CU1 (3)	SWKF	U	13	13	13.3 (12.4-13.6)	13	0.793-0.905
CU1 (4)	SWKF	U	13	13	13.3 (12.4-13.7)	13	0.841-0.920
CU2 (1)	LKF	U	13	13	12.7 (12.7-13.3)	2	0.937-0.984
CU2 (2)	LKF	U	13	13	12.5 (12.5-13.5)	4	0.922-0.987
CU2 (3)	LKF	U	13	13	12.6 (12.4-13.6)	7	0.895-0.987
CU2 (4)	LKF	U	13	12	12.9 (12.4-13.7)	5	0.897-0.985
CU2 (1)	SWKF	U	13	13	12.3 (12.7-13.3)	10	0.812-0.981
CU2 (2)	SWKF	U	13	13	12.2 (12.5-13.5)	12	0.771-0.984
CU2 (3)	SWKF	U	13	13	12.5 (12.4-13.6)	13	0.763-0.979
CU2 (4)	SWKF	U	13	13	12.8 (12.4-13.7)	12	0.781-0.985

more data points were discarded (see Table 3.7). This is typical behavior of a KF when confronted with an increasingly less accurate process model, in this case due to $b_{\max}$ decreasing and Δt increasing — the filter's solution is to trust the data more and the model less. The effectiveness of this strategy is clearly dependent on the quality and quantity of the observations. The accurate and robust inverse solutions seen for CU1 — and other driven datasets throughout this chapter — indicate that the observations in these simulations are sufficient for this strategy to be effective. In Sec. 3.7, some general comments regarding the impact of changing two observation parameters, namely the IED and SNR, will be made.

The innovations for both filters when temporal undersampling was applied were Gaussian, unbiased and had the correct magnitude (see Table 3.8). This demonstrates the ability of the LKF and SWKF to operate robustly under suboptimal conditions. Significant correlations were present in nearly all innovations of both filters. However, this is not surprising given the deficiencies of the process model.

A final comment regarding CU1 concerns the optimization step. We found that as less time points from CU1 were made available to the LKF, the optimization algorithm would increasingly select parameter sets that resulted in the covariance matrices becoming singular, causing the filter to crash. This issue was associated with the optimizer's performance becoming increasingly sensitive to the parameters' initial values and their associated bounds, and was particularly prominent when only every fourth time point was used. To remedy this problem, parameters were initialized closer to their likely estimated values and tighter bounds were imposed as well. Evidence of this issue can be seen in Table 3.7 when the LKF is applied to dataset CU1 (4); here c_{f_n} is at its maximum value of 12. Interestingly, no such problem was observed for the SWKF.

We will now discuss the second temporal undersampling study which investigates filter performance for dataset CU2. This dataset, which also contains a single drive term and uniform parameters, was described in Sec. 3.5.1.3. The simulated current dipoles and EEG are displayed in Figs 3.7(a) and 3.7(b), respectively. Once more, we begin by computing the inverse solution with uniform parameters using all available EEG time points. Again, both filters produced similar results; hence they will be discussed together. When using all observations, the filters reconstructed the current dipoles accurately. This is demonstrated by the concurrence between the simulated [Fig. 3.7(a)] and estimated [Figs 3.7(c) and 3.7(d)] current densities.

Unlike datasets DU, DV, and CU1, both filters selected parameter values that were considerably different to those used in the simulation (see Table 3.7). For example, both filters selected wave velocities which were well below the simulated value of 4 m s^{-1} despite the optimizer's upper bound for b being above this value when all time points were used. The mismatch indicates that the telegrapher's equation selected by each of the KFs exhibit significantly different dynamics compared to the one used to generate the simulated data. This arises from the fact that the dynamics of CU2 are dominated by the unmodeled drive term and less so by the transient dynamics of the telegrapher's equation itself. Designing CU2 in this way allowed us to investigate the situation where the filters must fit their process model (i.e., a telegrapher's equation with no drive term) to current densities generated by a different dynamic model. This scenario is important because it provides a simple simulation of what happens for the clinical data (see Chapters 2 and 4), where a simplified dynamic model — potentially affected by temporal undersampling — is fitted to complex spatiotemporal data. Despite the mismatched dynamic models, the innovations largely maintained their desired properties for the LKF and the SWKF, as they were found to be Gaussian, unbiased, and to have the correct magnitude (see Table 3.8). Again, temporal correlations were observed in some of the innovations, and were more pronounced for the SWKF, where significant correlations were found in 10 (out of 13) channels, compared to only 2 for the LKF.

The effect of temporal undersampling on KF performance for dataset CU2 is now described. As seen previously with CU1, temporal undersampling again results in only a small decrease in state estimation accuracy for both filters. This is illustrated by the concordance between the simulated and estimated current densities [Figs 3.7(e) - 3.7(h)] when temporal undersampling was applied. The RMSEs for both filters support this observation, since they increased only marginally as less data was made available. As expected, the model parameters selected are still significantly different to their simulated values (see Table 3.7). Unlike CU1, the Courant condition does not appear to restrict wave velocity estimation, as the optimizer always chose values for b well below its upper bound for the datasets using every second, third, or fourth time point. This is observed despite $b_{\max}$ falling below b's simulated value for these three datasets.

As done previously for CU1, we again searched for any mechanisms compensating for temporal undersampling to maintain the accuracy of the state estimates.

The process model's effective damping ζ_{eff} was computed for both filters for each of the datasets derived from CU2. Once again, no clear pattern could be observed for either filter, with $\omega_n\zeta$ remaining roughly constant as more data was discarded. However, the process noise covariance displayed the same trend seen in CU1 whereby it increased for both filters as more data points were discarded (see Table 3.7). This increases the filters' confidence in the observations relative to the model predictions and compensates for the growing discordance between the estimated and the simulated dynamics as Δt grows. Notwithstanding the simultaneous difficulties posed by mismatched dynamic models and temporal undersampling, the innovations for both filters were found to be Gaussian, unbiased, and to have the correct magnitude (see Table 3.8). This is another demonstration of how these algorithms can maintain their statistical integrity despite suboptimal conditions. As seen throughout this study, temporal correlations were again observed in many of the innovations of both filters — the SWKF more so than the LKF. However, this is expected given the shortcomings of the dynamic model.

In summary, the temporal undersampling study has shown that both filters can produce accurate inverse solutions when the Courant condition is violated for source dynamics dominated by either: (i) the dynamic model's transient response (CU1); or (ii) the drive term (CU2). The diagnostic tests revealed well-tuned filters for both datasets, although temporal correlations were observed in the innovations. These were due to the presence of unmodeled inputs, and were worsened by temporal undersampling. As temporal undersampling was applied (i.e., Δt increased), the process noise covariance rose to compensate for the growing inaccuracy of the process model. Parameter estimation was much less accurate for both filters. In the case of CU1, the Courant condition prevented the wave velocity from being estimated accurately. In contrast, this condition was less important for CU2 where differences between the simulated and estimated model parameters were greater. These findings demonstrate the robust nature of these algorithms, but also reaffirm the need to model the drive terms better in order to improve the accuracy of parameter estimates. Finally, it is also important to be cognisant of the impact of temporal undersampling on filter performance.

3.7 Discussion and Conclusions

In this chapter, a detailed study of KF-based EEG source localization was undertaken using 1D simulations. Two KF algorithms — one a standard linear filter, the other spatially whitened — were applied to simulated EEG data generated using a telegrapher's equation and a simplified volume conductor model. To better reflect true brain dynamics, the simulated EEG was given spatial variation by allowing two parameters of the telegrapher's equation to be functions of position. Estimation of the spatial variations was incorporated into the filtering algorithm by allowing the same parameters in the process model of both KFs to have spatial variations of the same functional form. Model and noise parameters were selected through AIC minimization. The KFs were applied to simulated data displaying both transient and driven dynamics. Temporal undersampling was also investigated. The inverse solutions for each scenario were analyzed in detail. This analysis included the application of standard diagnostic tests for evaluating KF performance. The key findings of this study are discussed below:

(i) The combination of Kalman filtering with a likelihood maximization tuning step can correctly estimate the spatial variation of model parameters from simulated EEG with transient dynamics. This leads to more accurate inverse solutions and improved filter performance. It also provides a simple, preliminary demonstration of parametric imaging.

(ii) The estimation of spatially varying model parameters is both less reliable and less advantageous for the more realistic scenario of driven EEG dynamics. This is because neither filter models the drive term(s) explicitly. Therefore, estimated parameter values need to be interpreted with caution, and the modeling of exogenous inputs introduced in future filter implementations.

(iii) Despite the difficulties experienced with parameter estimation, both filter algorithms can accurately reconstruct the simulated current dipoles, with the suboptimal, but computationally less burdensome, SWKF performing similarly to the more optimal LKF, especially in terms of RMSE values, and in the most part for the various tests evaluating KF performance. However, the estimated state covariances for the SWKF are found to be consistently larger for driven dynamics (see Table 3.9). This expected finding is due to the spatial whitening transformation which removes correlation information between

Table 3.9: Estimated $[\sigma_e(v, N_k)]$ and predicted $[\sigma_p(v, N_k)]$ standard deviations for the current dipole at voxel v at time step N_k for the datasets TU, TV, DU, and DV, with IED $= 0.02$ m and SNR $= 10$. These values are obtained from the estimated and predicted state covariance matrices, $P(k|k)$ and $P(k|k-1)$, respectively. The distance of v from the origin is indicated by $x(v) = (v-1)\Delta x$.

Dataset	Filter	Profile	v	$x(v)$ (m)	$\sigma_e(v, N_k)$ (A m)	$\sigma_p(v, N_k)$ (A m)
TU	LKF	U	51	0.25	0.608	0.611
TU	LKF	G	51	0.25	0.574	0.576
TU	SWKF	U	51	0.25	0.355	0.357
TU	SWKF	G	51	0.25	0.345	0.347
TV	LKF	U	26	0.125	5.65	6.97
TV	LKF	G	26	0.125	0.746	0.747
TV	SWKF	U	26	0.125	9.23	9.67
TV	SWKF	G	26	0.125	5.64	5.82
DU	LKF	U	51	0.25	4.10	6.44
DU	LKF	G	51	0.25	3.90	6.17
DU	SWKF	U	51	0.25	12.30	13.22
DU	SWKF	G	51	0.25	14.58	15.86
DV	LKF	U	75	0.37	6.75	10.04
DV	LKF	G	75	0.37	6.12	9.16
DV	SWKF	U	75	0.37	17.20	18.32
DV	SWKF	G	75	0.37	17.30	18.47

neighboring voxels, and can be observed in the smaller reductions of the state covariance achieved through the update step of the SWKF compared to the LKF (see Table 3.9). Importantly, these larger state covariance values are not significant enough to affect the interpretation of the inverse solutions.

(iv) Hyperbolic partial differential equations, such as the telegrapher's equation used in this work, are an appropriate choice for modeling large-scale brain activity. However, via the Courant condition, the spatiotemporal discretization places an upper bound on the wave velocity, and therefore the spatial coupling term. If this condition is ignored, filter performance can be degraded,

although reliable state — but not parameter — estimates are still produced even when b cannot be correctly estimated. This comes at the price of larger state covariance estimates and less reliable optimization performance.

(v) The performance of both filters was found to be robust across SNRs ranging from 2–20.

(vi) As expected the electrode spacing was found to affect inverse solution performance. For instance, when less data was available to constrain the inverse solution (IED = 10 voxels), the AIC minimum found through optimization became dependent on parameter starting values. In this situation, there exists a number of AIC minima with values similar to the 'correct' one. Such ambiguity can lead to state and parameter estimates which are inaccurate, yet still describe the observed data.

(vii) More broadly, these findings are also applicable to other spatiotemporal systems where KFs are used to solve joint state and parameter estimation.

The above findings suggest several potential improvements and future directions for this work:

(i) Perhaps the most significant finding from this study is the ability of both KFs to reliably perform state estimation, even when challenged by driven dynamics and temporal undersampling. Parameter estimation was found to be less reliable, and was due to the unmodeled drive term, which needs to be handled by the KFs. Several potential strategies are now proposed. One approach is to estimate the drive term using basis functions (e.g., [116]). Another option is to employ a process model that contains an explicit cortical input, such as [80, 117, 119] through its corticothalamic loop. An alternative approach is to use a GARCH formulation of these algorithms, which could compensate for the unmodeled drive term.

(ii) To accurately capture the spatial variation of brain dynamics, the algorithms should ideally be able to estimate arbitrary spatial profiles for model parameters.

(iii) Temporal undersampling issues must be overcome in order to accurately estimate the wave velocity. If available, an obvious workaround is to collect the

EEG data at a higher (temporal) frequency. A more practical solution is to run the process model (i.e., the filter's prediction step) at a faster rate, and assimilate the observations as they become available. This approach has already been used when applying filtering algorithms to fMRI time series [30, 64, 69, 109]. Here, this strategy is needed because the time steps over which models describing BOLD signal dynamics can be reliably integrated are significantly smaller than the time between consecutive scans.

(iv) Parameter estimates could be improved and local minima avoided if the optimization step was run from multiple starting points. This would produce a probability distribution for the parameter values, as done in [116].

Chapter 4

Kalman-Filter-Based EEG Source Localization with a Spatially-Varying Process Model: Whole-Brain EEG

4.1 Introduction

In this chapter the concepts and findings from the 1D simulation study described in Chapter 3 are applied — through a pilot study — to the whole-brain inverse solution analyzed in Chapter 2. This pilot study expands the analysis of the SWKF algorithm to address the following issues:

- The accuracy and biophysical interpretability of process model parameters estimated from both simulated and clinical EEG data.
- The effect on inverse solution performance of introducing simple spatial variation of model parameters.
- The impact of temporal undersampling — via the Courant condition — on parameter estimation.

Focusing on these issues helps to achieve this study's overall aim of providing further insights into KF-based EEG source localization on two levels. Firstly, specific comments regarding the application, performance, and enhancement of the SWKF algorithm are made. Secondly, we provide a more general discussion about using Kalman filters to solve the joint state and parameter estimation problem for functional neuroimaging data. This is currently a very active and technically demanding

area of research, and we outline some possible future directions that have emerged from the work described in this thesis. The significant task of implementing these recommendations will be the subject of future work. Since the two previous chapters, both of which are self-contained papers, have already introduced much of the background material and mathematics relevant to this chapter, we do not repeat it here, and include only what is necessary to describe the motivation for this work, along with any new material. Where necessary, the reader is referred to the relevant parts of the previous chapters. We now briefly review Chapters 2 and 3 to provide the context and motivation for the pilot study.

In Chapter 2, we introduced and evaluated a recently developed SWKF that solved the EEG inverse problem for a realistic head model. The filter employed a discrete spatiotemporal model, originally chosen as an ansatz, to describe the dynamics of the current dipoles — or more precisely their spatial Laplacian. Operating in a Laplacianized state space significantly reduced the filter's computational burden. In this work it was recognized that for certain parameter ranges this dynamic model's continuous analog was the telegrapher's equation, which is an appropriate starting point for modeling large-scale brain activity for reasons that are outlined in Secs 2.3.1 and 3.3. However, in Chapter 2 the coefficients a_1, a_2, and b_1 [see Eqs (2.10) – (2.12)] were estimated using unconstrained optimization, which meant they could take on values that were inconsistent with the telegrapher's equation definition, which requires the natural frequency f_n, fractional damping coefficient ζ, and wave velocity b to be $\geqslant 0$ (see Sec. 3.3). We also note that when estimated directly — as was done in Chapter 2 — b_1 should be $\leqslant 0$. The optimized values for the temporal parameters (f_n and ζ) were found to agree with the telegrapher's equation definition for both simulated and clinical data (see Sec. 2.6.1). In contrast, the spatial term b_1 [Eq. (2.12)] had the wrong sign for the clinical data (see Sec. 2.4.2). We aim to remedy that here by estimating the wave velocity directly while imposing appropriate constraints. Additionally, this study also applied standard diagnostic tests to objectively evaluate KF performance. These tests identified spatial variation in filter performance which could be improved by spatially-varying model parameters.

In Chapter 3, the SWKF algorithm was examined in detail using a 1D simulation study. The simulated EEG data was generated using a telegrapher's equation and a simplified volume conductor model. Two of the model parameters (f_n and ζ) were also given spatial variation of a simple functional form to better reflect true

brain dynamics (see Sec. 3.5.1). In the optimization step, model parameters were explicitly constrained to ranges consistent with the process model being a telegrapher's equation. For the purposes of comparison, inverse solutions were also computed using the optimal LKF. This study found both filters produced accurate state estimates, with the SWKF exhibiting performance similar to the LKF. Standard diagnostic tests also showed both filters to be well-tuned for all scenarios investigated in the 1D study. However, parameter estimates for the SWKF and LKF were less reliable, as they were degraded by both the unmodeled drive term and temporal undersampling.

In this chapter, we return to solving the EEG inverse problem for both simulated and clinical whole-brain EEG recordings. Following the method described in Chapter 2, the SWKF is again applied to a realistic head model. Using the strategy from Chapter 3, spatially varying model parameters (f_n and ζ) are introduced into the simulated EEG and optimization step using a known functional form. Unlike Chapter 2, simulated whole-brain EEG is generated with non-zero wave velocities, which means wave-like patterns of brain activity can be produced, and the SWKF's performance in this more realistic scenario can be studied in detail. Finally, at the optimization step, model parameters are constrained to ranges consistent with the process model being a telegrapher's equation. The Courant condition, which places an upper bound on stable estimates of the wave velocity, is also imposed by the optimizer. These constraints were previously applied to parameter estimates in Chapter 3, but not Chapter 2.

The remainder of this chapter is structured as follows. Section 4.2 outlines the pilot study. Here the whole-brain inverse problem is recapitulated, and the introduction of spatially varying model parameters into the gray matter mask is described. Details of the simulated and clinical EEG datasets are also provided in this section, along with a brief overview of the SWKF algorithm and the methods used to tune and evaluate its performance. The results for the simulated and clinical EEG are presented in Sec. 4.3. Section 4.4 discusses these results and makes some concluding remarks.

4.2 Details of Pilot Study

The pilot study reinvestigates the EEG inverse problem described in Chapter 2. Here the inverse solution's core components — including the problem formulation, the dynamic model, the SWKF, the gray matter mask, and the lead field matrix — remain unchanged from Chapter 2. To avoid repetition, the reader is referred to Secs 2.2 – 2.4 which contain the original introduction of the whole-brain inverse problem.

One important modification in this chapter is the introduction of spatially varying model parameters to the whole-brain EEG inverse problem using the method described in Chapter 3. This approach restricts each parameter's spatial profile to a known functional form, which allows the existing filtering framework to be used with the addition of only a small number of parameters. For the pilot study we have chosen to model the spatial variation of the alpha rhythm, which typically increases in strength and frequency from anterior to posterior [101]. The spatial dependence of alpha activity was selected for several reasons: (i) it is a robust and widely studied feature of EEG data; (ii) it can be modeled by making the parameters that describe the temporal resonance of the telegrapher's equation (f_n and ζ) a simple function of space; and (iii) it allows the eyes-closed EEG recording from Chapter 2, which displays the aforementioned spatial characteristics of the alpha rhythm, to be reinvestigated with spatially varying model parameters.

To achieve the desired anterior to posterior variation of the alpha rhythm, we propose a sinusoidal function $s(v)$ to describe the spatial profiles for the natural frequency ($s = f_n$) and damping coefficient ($s = \zeta$) in this study. The sinusoid at voxel v, with coordinates $\boldsymbol{V} = (x, y, z)^T$ is given by

$$s(v) = c_s + d_s \cos\theta, \tag{4.1}$$

where c_s and d_s are the sinusoid's mean and amplitude respectively, and the cosine of the angle θ measured in the sagittal plane between the origin v_o — whose coordinates are $\boldsymbol{V_o} = (x_o, y_o, z_o)^T$ — and voxel v is

$$\cos\theta = \frac{z_o - z}{\sqrt{(y_o - y)^2 + (z_o - z)^2}}. \tag{4.2}$$

Figure 4.1(a) shows how θ is constructed for the origin coordinates used in this study [$\boldsymbol{V_o} = (11, 0, 13)^T$]. We emphasize that θ is only a function of each voxel's axial (y) and coronal (z) — but not sagittal (x) — position, which means the choice

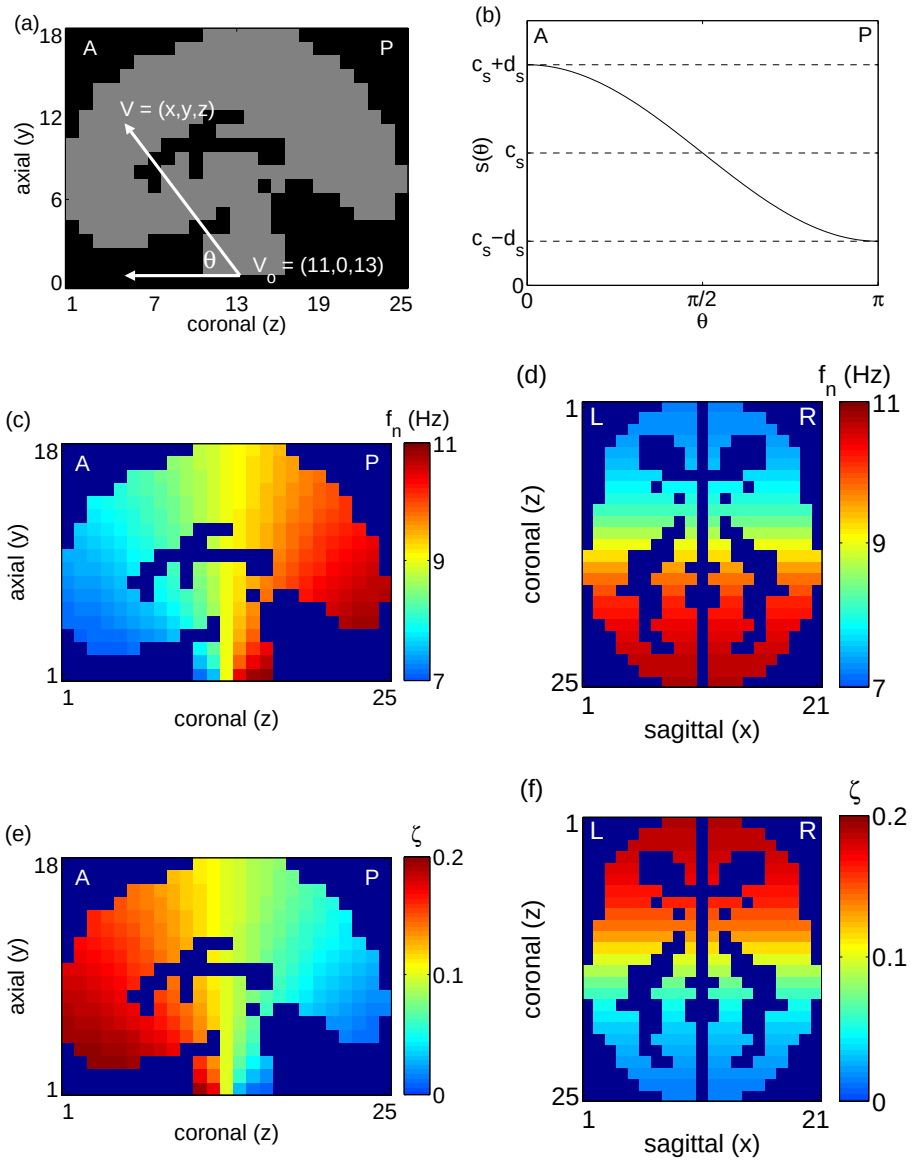

Figure 4.1: Spatially varying model parameters for whole-brain gray matter mask. The parameters f_n and ζ are a function of θ [see Eq. (4.1)]. Frame (a) shows how θ is constructed for the origin coordinates used in this study. Note that θ is independent of sagittal (x) position. Frame (b) plots the general sinusoidal profile between $\theta = 0$ (anterior) and $\theta = \pi$ (posterior). Frames (c) – (f) show the spatial profiles for dataset DV (see Sec. 4.2.1.2). Frames (c) and (d) show the spatial profile for f_n in the sagittal and axial planes, respectively. The spatial profile for ζ is shown for the same planes in frames (e) and (f). Note that for the sagittal slices, the letters A and P denote anterior and posterior, respectively. Axial slices are shown using neurological convention; i.e., left is left and right is right.

of x_o is arbitrary. Therefore, f_n and ζ exhibit no left-to-right variation, which is a reasonable approximation for an initial model of the alpha rhythm's spatial properties. More complex models could be introduced in the future. The general spatial profile produced by this sinusoidal function is plotted in Fig. 4.1(b). A specific example of the resulting spatial profiles for f_n and ζ are shown for dataset DV in Figs 4.1(c)-4.1(f). The spatial variation of f_n within the gray matter mask is displayed in the sagittal and axial planes, in Figs 4.1(c) and 4.1(d), respectively, while the spatial profile for ζ using the same planes is shown in Figs 4.1(e) and 4.1(f). Giving the parameters f_n and ζ spatial dependence means the coefficients for the discretized 3D telegrapher's equation a_1, a_2, and b_1 [Eqs (2.10) – (2.12)] must now be evaluated at each voxel. This topic is discussed in Sec. 3.3 for the 1D simulation study.

4.2.1 Simulated EEG Data

To investigate various aspects of filter performance, four simulated EEG datasets were generated using the process model Eq. (2.7), with each dataset having its own source dynamics. The first two datasets (TU and TV) are produced by the system's transient dynamics, while the next two (DU and DV) are driven by a spatiotemporal input. Following the convention from Chapter 3, the letter U indicates the simulated current dipoles were produced using spatially uniform model parameters, while the letter V indicates spatially varying parameters.

Before introducing the simulated datasets individually, the features common to all are discussed. Firstly, the spatiotemporal discretization (Δx and Δt) matches the one used in Chapter 2. The spatial discretization $\Delta x = 7$ mm is determined by the gray matter mask (see Sec. 2.4). To mirror the clinical data, three of the four simulated datasets — TU, DU, and DV — comprise 2 s of EEG generated at a sampling rate of 256 Hz. Due to the stronger damping present in TV, this EEG dataset is 1 s in length and sampled at 512 Hz.

The model parameters f_n and ζ, which are specified for the simulations and estimated by the KFs, are given either a uniform or sinusoidal spatial profile. The model parameters are summarized in Table 4.1. The uniform profile was used in Chapter 2 and allows for comparison with the sinusoidal profile, which is selected for its ability to produce source dynamics with realistic spatial variation. Furthermore, using a spatial profile of a known functional form permits the investigation

Table 4.1: Simulated current dipole parameters for transient (TU and TV) and driven (DU and DV) datasets.

Dataset	b (m s^{-1})	ν	c_{f_n} (Hz)	d_{f_n} (Hz)	c_ζ	d_ζ	σ_ϵ^2
TU	0.8	1.34	10	—	0.01	—	9.76×10^{-3}
TV	1.6	1.34	9	-2	0.05	0.05	9.80×10^{-3}
DU	0.8	1.34	10	—	0.1	—	9.58×10^{-3}
DV	0.8	1.34	9	-2	0.1	0.1	9.81×10^{-3}

of spatially varying model parameters using the existing filtering framework, with the addition of only a small number of parameters. In contrast, the wave velocity is again assumed to be spatially uniform for all datasets. This spatial invariance is a reasonable approximation [104]. The magnitude of b — which is 0.8 m s^{-1} for all datasets except TV where $b = 1.6$ m s^{-1} — was selected so the Courant condition is satisfied for the spatiotemporal grids used. This is shown in Table 4.1 where $\nu = 1.34$ for all four datasets, which is comfortably below its maximum allowed value of $\sqrt{3}$. However, cortical activity is known to propagate at up to 10 m s^{-1} [105, 119]. This discrepancy was present in Chapter 2, where the same spatiotemporal discretization was used and restricted b to values well below 10 m s^{-1}. Following the approach from Chapter 2, all current vectors in the simulated datasets are oriented in the *z*-direction (coronal axis) to maximize the scalp voltages at the occipital electrodes (i.e., O1 and O2).

4.2.1.1 Transient Source Dynamics

The first two sets of simulated current dipoles (TU and TV) are generated by the source model's transient response to spatial Gaussian initial conditions, with the drive term set to zero. These datasets provide a baseline for assessing the SWKF's performance as the drive term was found to distort parameter estimates in Chapter 3. The first transient dataset (TU), assumes f_n and ζ are spatially uniform, so only c_{f_n} and c_ζ are required [see Table 4.1 and Figs 4.2(g) and 4.2(h)]. The initial conditions are defined by the 3D spatial Gaussian Eq. (2.38) introduced in Sec. 2.4.1. Here the Gaussian has FWHM $= 30$ mm (activation zone radius 30 mm) and is centered in the right occipital pole $\boldsymbol{V}_c = (13, 7, 24)^T$. For the second transient dataset (TV), f_n and ζ have sinusoidal profiles with the parameters listed in Table 4.1 and plotted in

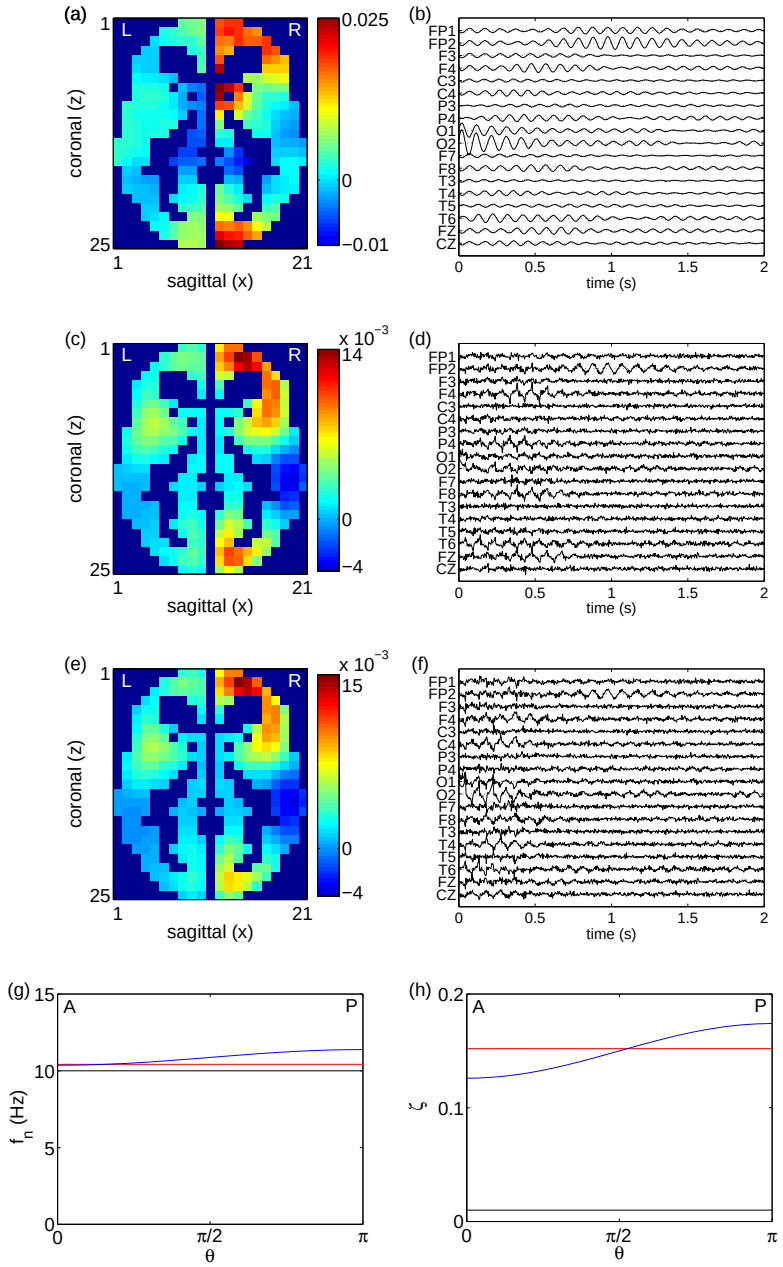

Figure 4.2: Results of inverse solution for transient source dynamics with uniform parameters (TU). A snapshot of the simulated source dynamics is shown on an axial slice in frame (a). Frame (b) shows the simulated EEG. The SWKF's inverse solution when fitting a uniform profile for f_n and ζ is shown in (c), and the innovations in (d). Frames (e) and (f) show the same results when fitting sinusoidal profiles. Frames (a), (c), and (e) are all shown at 1.21 s. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 6.1 and 6.0 relative to (b), respectively. Frame (g) shows the simulated (black line) and estimated spatial profiles for f_n. In (g) the red and blue lines indicate the uniform and sinusoidal profiles, respectively. The spatial profile for ζ is shown in (h) using the same key as (g).

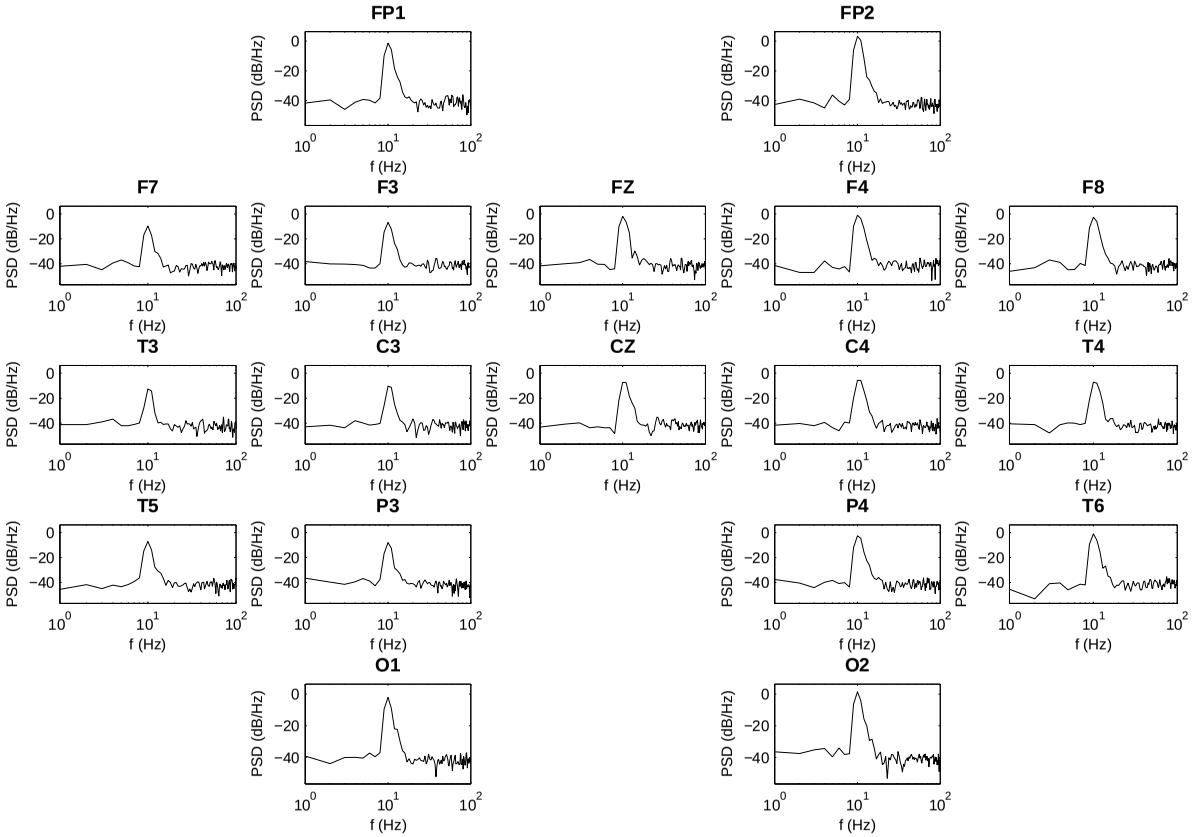

Figure 4.3: PSD at each electrode of the 10-20 system for the simulated EEG with transient dynamics and uniform parameters (TU). Pz is not shown since it is excluded from the analysis (see Sec. 4.2.4 for more details).

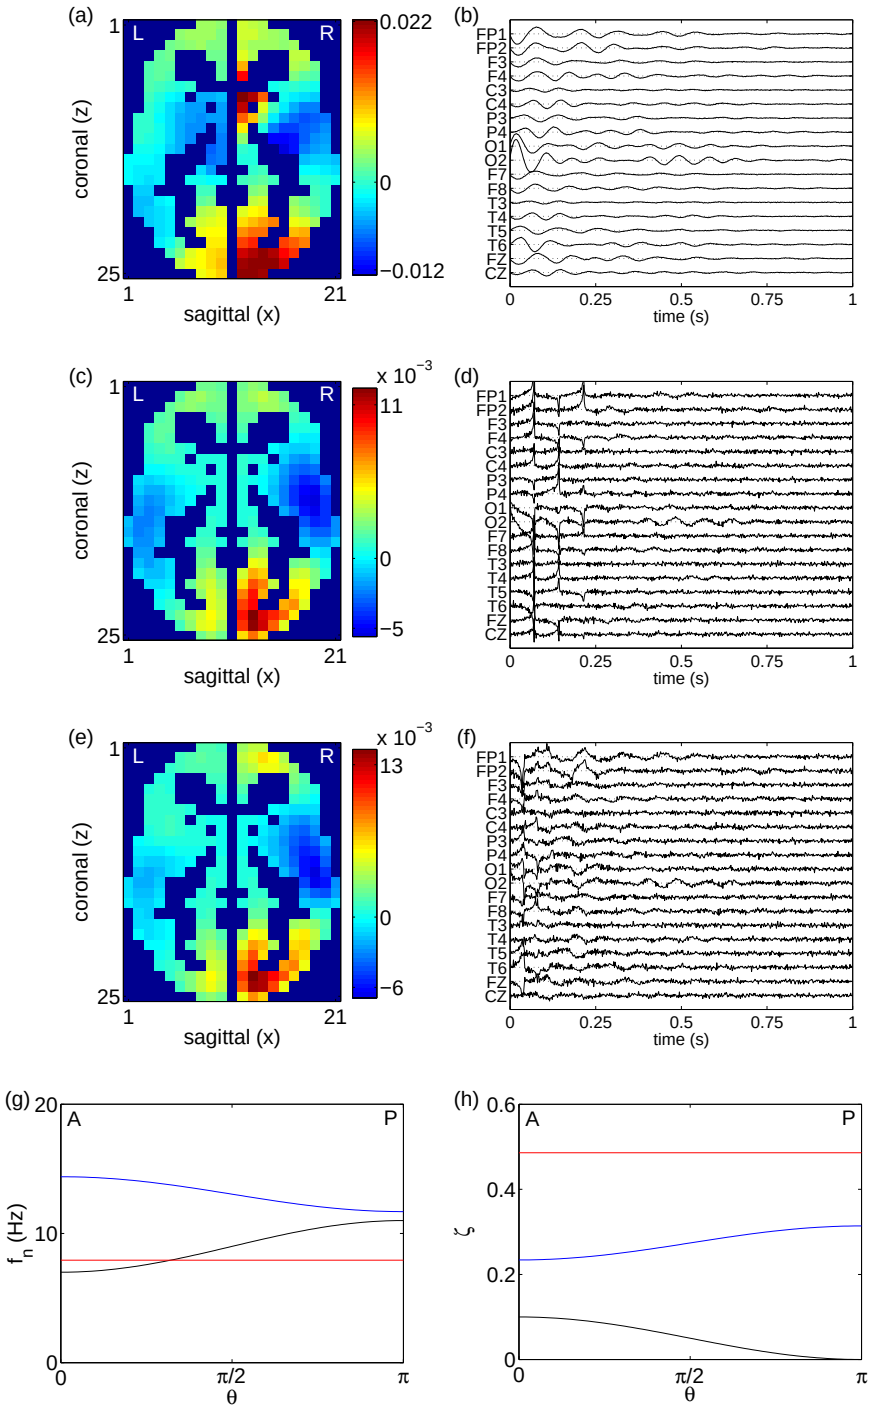

Figure 4.4: Results of inverse solution for transient source dynamics with spatially varying parameters (TV). This figure uses the same format as Fig. 4.2. Frames (a), (c), and (e) are all shown at 0.41 s. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 5.5 and 6.1 relative to (b), respectively.

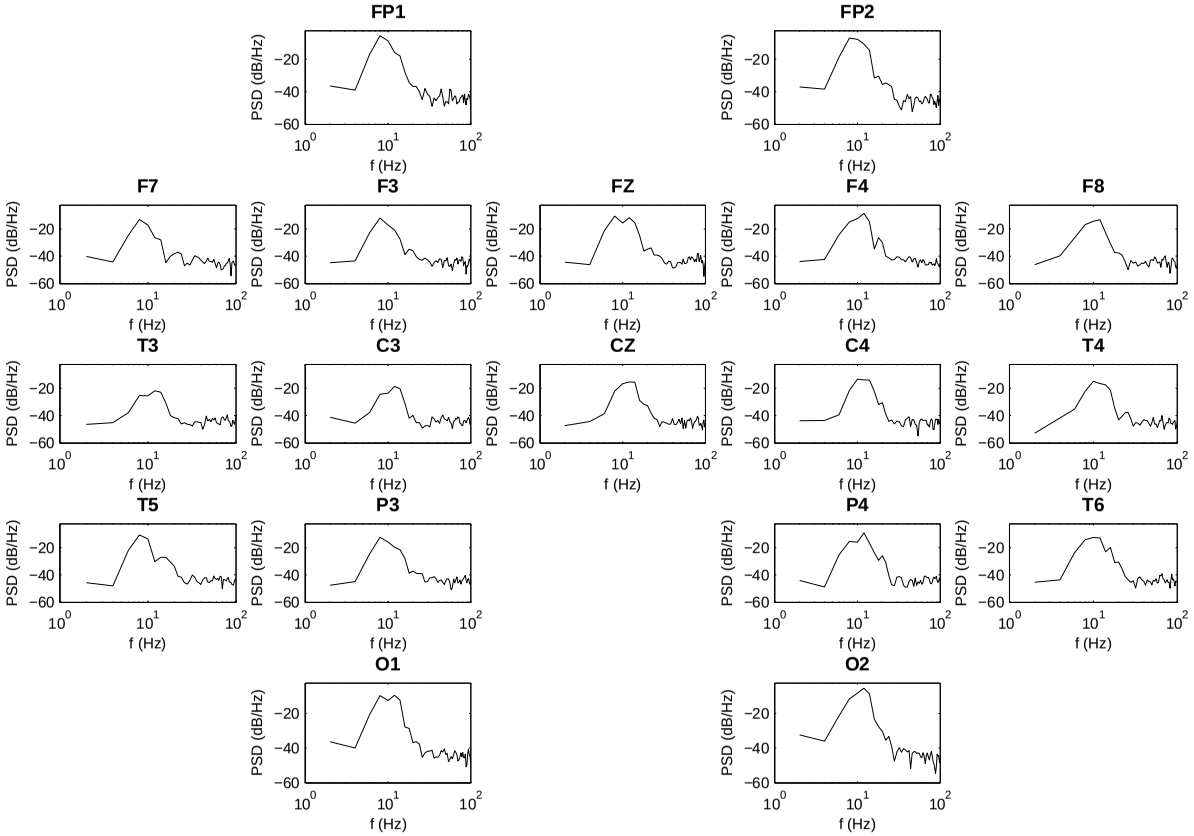

Figure 4.5: PSDs for the simulated EEG with transient dynamics and spatially varying parameters (TV). This figure uses the same format as Fig. 4.3.

Figs 4.4(g) and 4.4(h). The initial conditions for TV comprise two 3D spatial Gaussians. Two are required to produce transient dynamics which clearly demonstrate spatial dependence. One Gaussian — located in the right occipital lobe — is the same as used in TU, and the other is centered in the left frontal lobe $\boldsymbol{V}_c = (6, 6, 4)^T$ with FWHM $= 30$ mm (activation zone radius 30 mm). A snapshot of the resulting current dipoles for TU and TV are shown in Figs 4.2(a) and 4.4(a), respectively.

4.2.1.2 Driven Source Dynamics

The second pair of simulated current dipoles (DU and DV) is produced by driving the dynamic model with a spatiotemporal input. This term simulates thalamic input to the cortex and is similar to those used in Chapters 2 and 3. Following the approach outlined in Sec. 3.5.1.2, the temporal component of the drive $u_t(k)$ is again modeled by the sum of randomly-phased sine functions defined by Eq. (3.40). The sine functions have an even spacing of 0.1 Hz between 0 and 20 Hz. The drive term's spatial component u_s is modeled by the same 3D spatial Gaussian, Eq. (2.38), that provided the initial conditions for the transient source dynamics in this study (see Sec. 4.2.1.1). To compute the spatiotemporal drive term $\boldsymbol{u}_L$, the temporal u_t and spatial u_s components are multiplied together, as in Eq. (3.41). As done in the 1D study, the deterministic drive $\boldsymbol{u}_L$ is then added to the right hand side of the discretized 3D telegrapher's equation Eq. (2.7).

The first dataset with driven dynamics (DU) assumes f_n and ζ are spatially uniform [see Table 4.1 and Figs 4.6(g) and 4.6(h)], with one drive, where u_s is the same spatial Gaussian used for dataset TU. This produces a highly resonant source of 10 Hz oscillations in the right occipital lobe which spread over the cortex as shown in Fig. 4.6(a). The other dataset (DV) uses sinusoidal profiles for f_n and ζ [see Table 4.1 and Figs 4.8(g) and 4.8(h)] to reproduce some typical features of the alpha rhythm, namely the spatial variation of its frequency and strength [101]. In order to generate these dynamics, the gray matter mask is driven by two inputs which excite regions with different parameter values. This results in two alpha sources with distinct spatiotemporal characteristics, which more closely models what is observed in real EEG recordings. The first drive term is the same as the one used for dataset DU. Since it is centered in the right occipital lobe where $f_n \approx 11$ and $\zeta \approx 0$, it generates a strong 11 Hz resonance. The second drive term is comprised of the spatial Gaussian from TV which was located in the left frontal lobe, and the same temporal

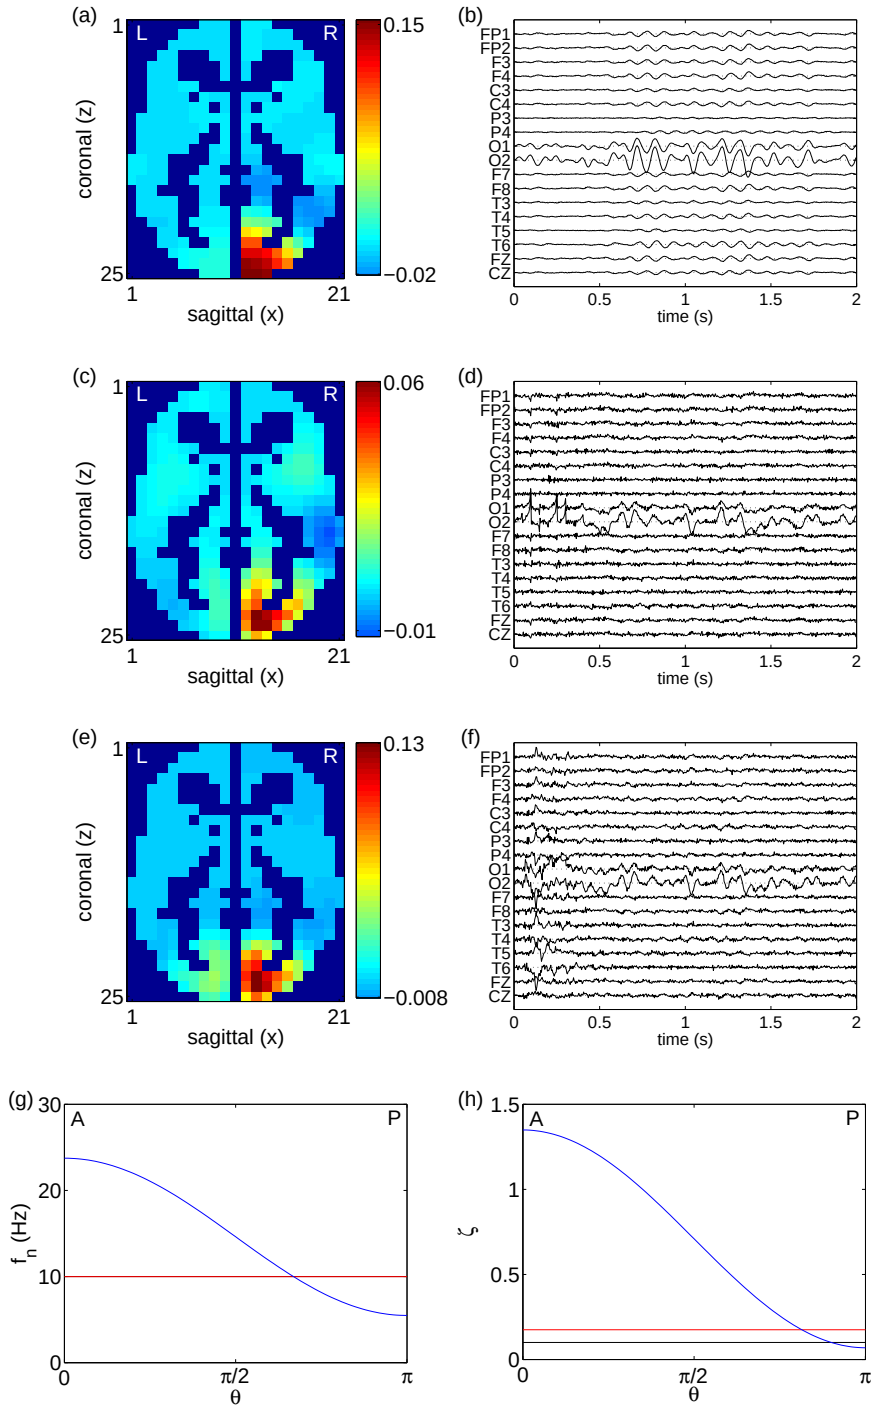

Figure 4.6: Results of inverse solution for driven source dynamics with uniform parameters (DU). This figure uses the same format as Fig. 4.2. Frames (a), (c), and (e) are all shown at 0.71 s. The vertical voltage scale of the traces in (d) and (f) have both been expanded by a factor 4.4 relative to (b), respectively.

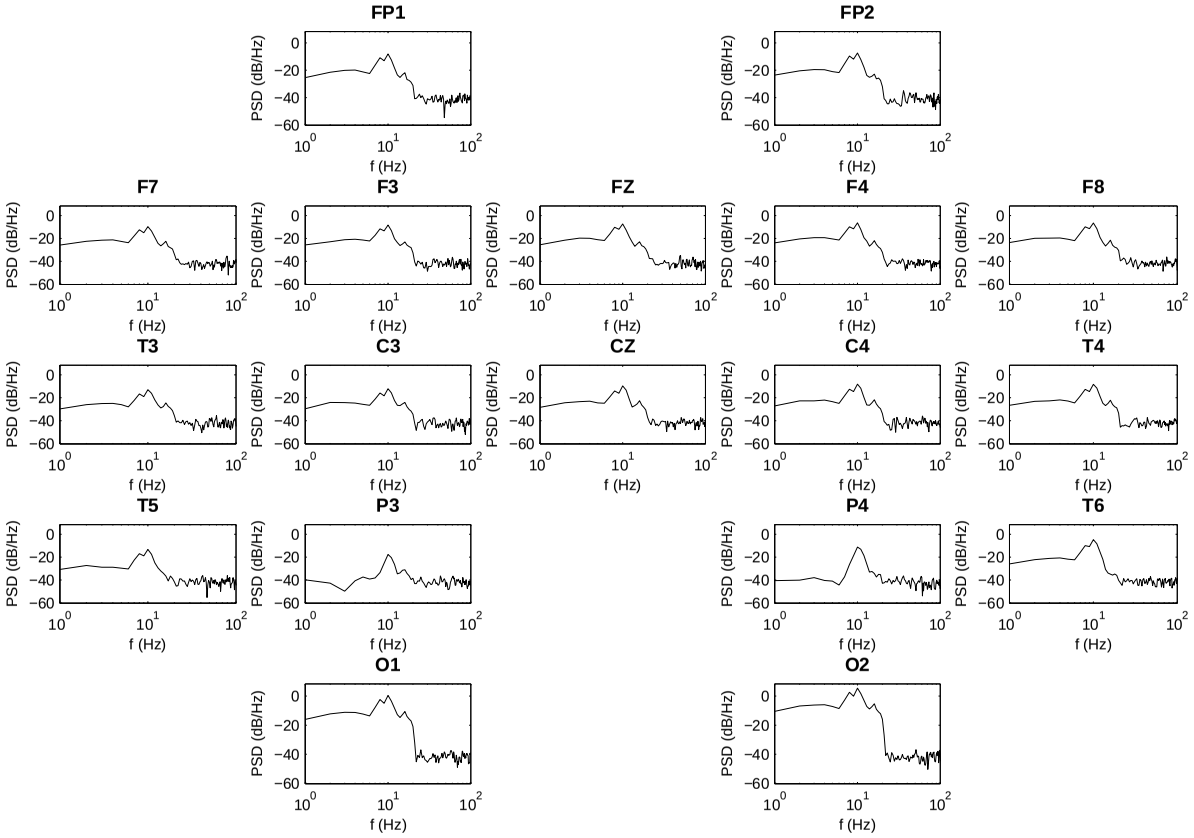

Figure 4.7: PSDs for the simulated EEG with driven dynamics and uniform parameters (DU). This figure uses the same format as Fig. 4.3.

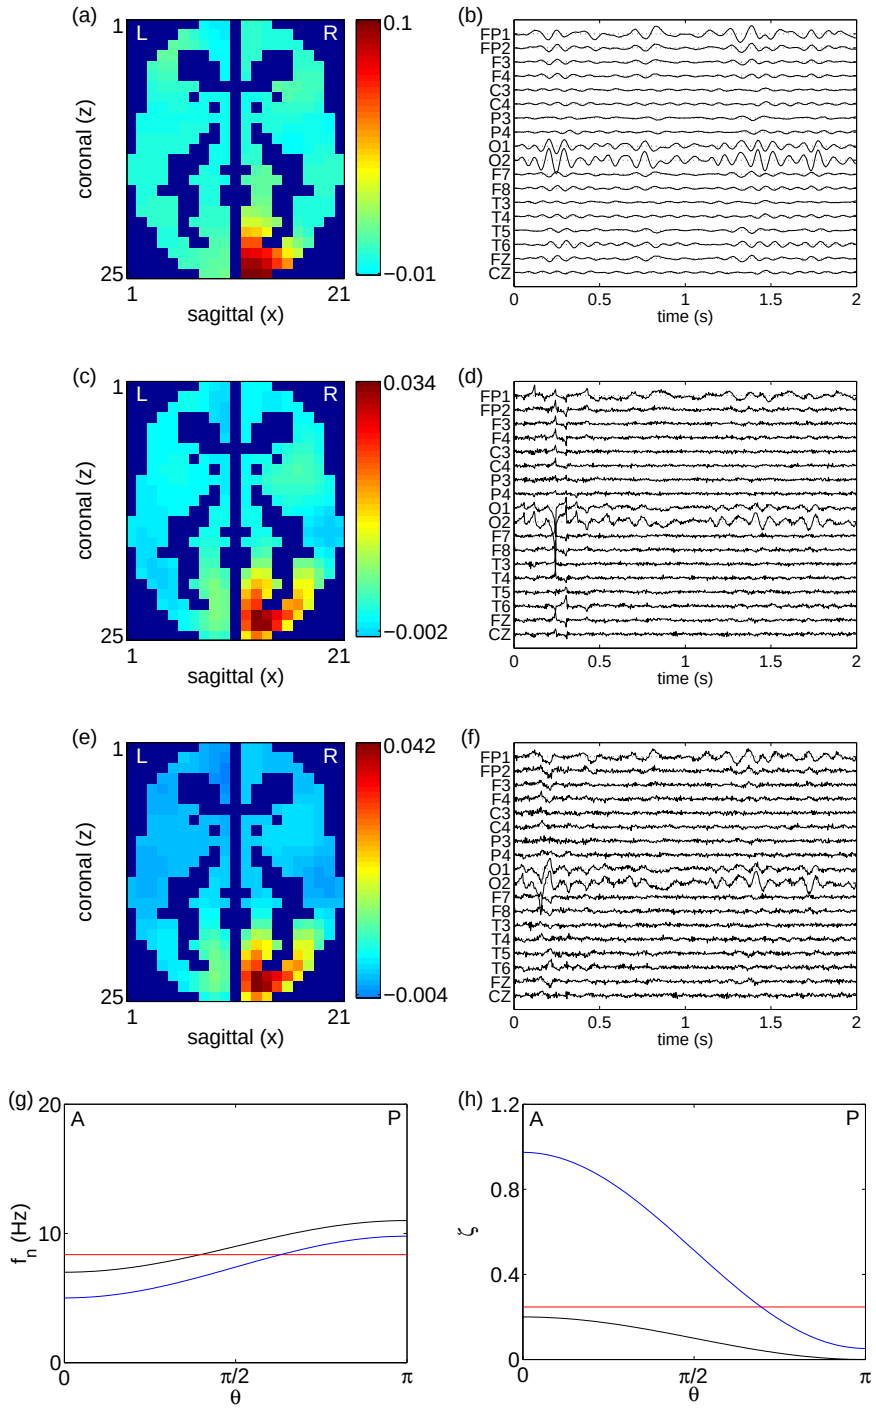

Figure 4.8: Results of inverse solution for driven source dynamics with spatially varying parameters (DV). This figure uses the same format as Fig. 4.2. Frames (a), (c), and (e) are all shown at 1.41 s. The vertical voltage scale of the traces in (d) and (f) have been expanded by factors of 3.8 and 4.8 relative to (b), respectively.

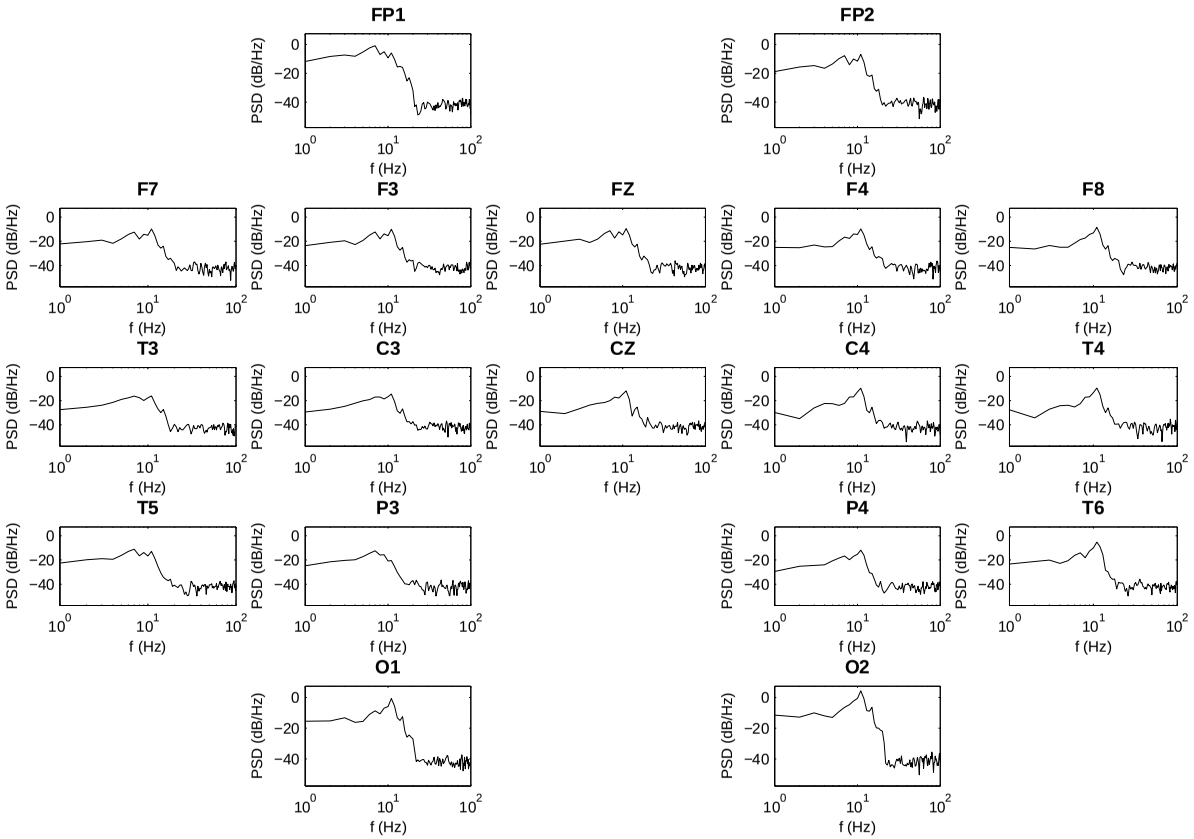

Figure 4.9: PSDs for the simulated EEG with driven dynamics and spatially varying parameters (DV). This figure uses the same format as Fig. 4.3.

input used for the first drive term. Given this source is located where $f_n \approx 7$ Hz and $\zeta \approx 0.2$, it produces a weaker 7 Hz resonance. Figure 4.8(a) shows a snapshot of the resulting current densities for dataset DV.

4.2.1.3 Computing EEG Data

To produce the simulated EEG data, we multiply $\boldsymbol{J}$ by the LFM $\boldsymbol{K}$ and apply average reference. The LFM from Chapter 2 is used again. This matrix is computed for the International 10-20 System ($N_c = 19$) by solving the vector Laplace equation for a 3-shell spherical head model via the boundary element method [115]. White Gaussian observation noise was added to the EEG to give a signal to noise ratio (SNR) of 10 in terms of standard deviations of the voltages. This SNR is typical for EEG recordings [121]. The simulated observation noise covariances $\sigma_{\boldsymbol{\epsilon}}^2$ are listed in Table 4.1, and the EEG for TU, TV, DU, and DV, is plotted in Figs 4.2(b), 4.4(b), 4.6(b), and 4.8(b) respectively. Figures 4.3, 4.5, 4.7, and 4.9, show the PSD of the EEG time series at each electrode. The spatial dependence of the EEG — which reflects the spatial dependence of the underlying source dynamics — can be seen for DV in Fig. 4.9 where the frequency and strength of the alpha rhythm displays anterior to posterior variation.

4.2.2 Clinical EEG Data

In this study the clinical recording from Chapter 2 is used again (see Sec. 2.4.2). This data was recorded from a healthy male child aged 8.5 years, in awake resting state with eyes closed. Electrodes were placed according to the 10-20 system and the data was collected at a sampling rate of 256 Hz. A 2 s time series was selected from the recording for analysis and is shown, using average reference, in Fig. 4.10(a). This data displays characteristic alpha oscillations which are most prominent in the occipital electrodes, and attenuate posterior to anterior. Figure 4.11, which shows the time series PSD at each electrode, clearly demonstrates the spatial dependence of brain activity present within this dataset.

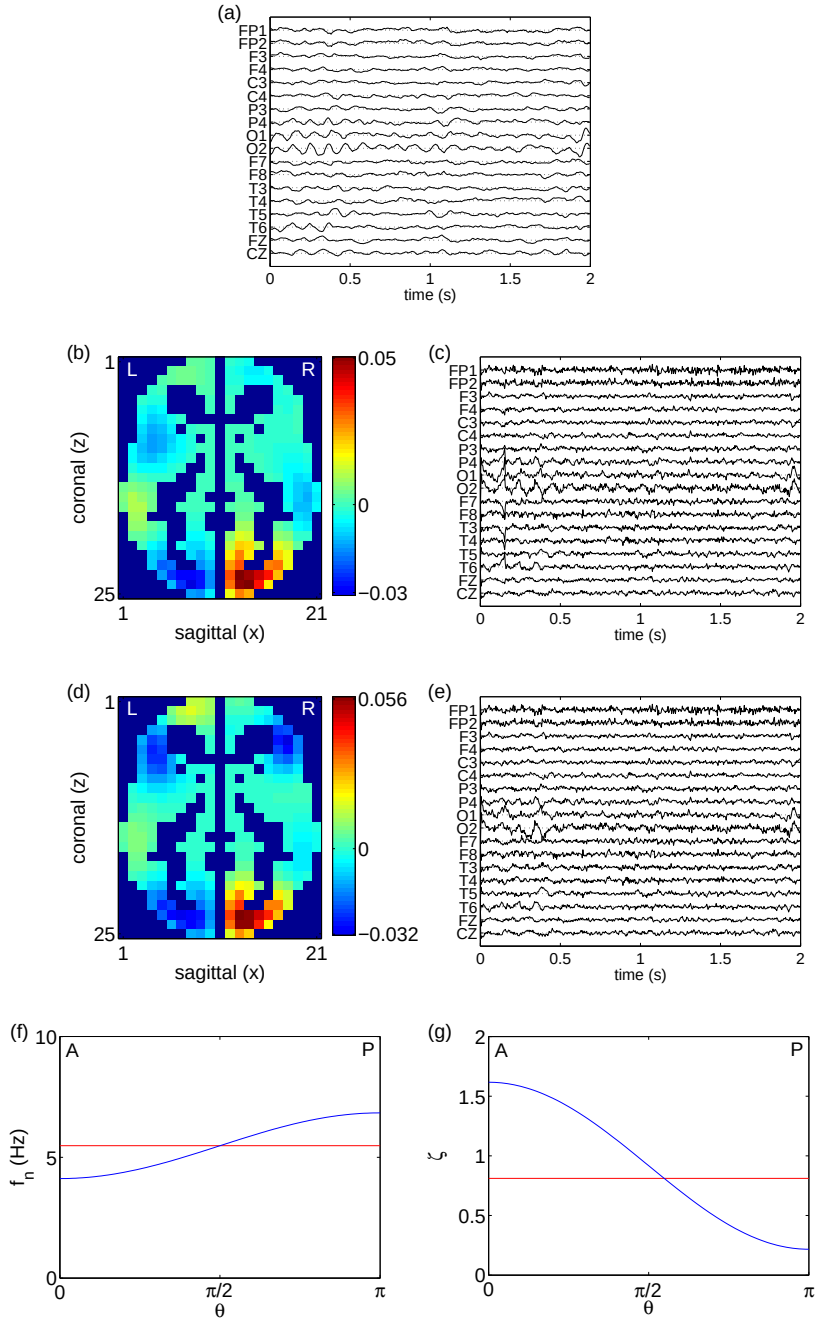

Figure 4.10: Results of inverse solution for clinical EEG. Frame (a) shows the clinical EEG. A snapshot of the SWKF's inverse solution when fitting a uniform profile for f_n and ζ is shown on an axial slice in (b), and the innovations in (c). Frames (d) and (e) show the same results when fitting sinusoidal profiles. Frames (b) and (d) are both shown at 0.47 s. The vertical voltage scale of the traces in (c) and (e) have been expanded by factors of 4.7 and 4.6 relative to (a), respectively. Frame (f) shows the estimated spatial profiles for f_n. In (f) the red and blue lines indicate the uniform and sinusoidal profiles, respectively. The spatial profile for ζ is shown in (g) using the same key as (f).

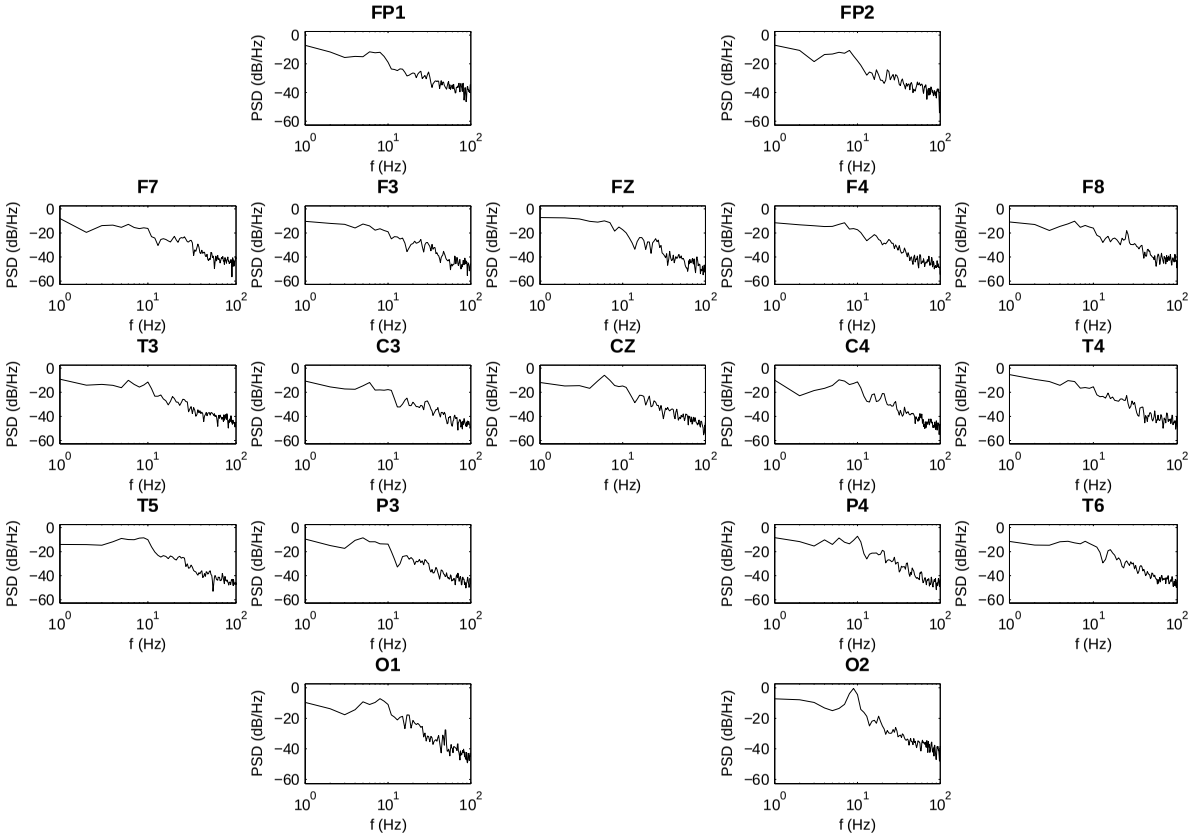

Figure 4.11: PSDs for the clinical EEG recording. This figure uses the same format as Fig. 4.3.

4.2.3 Courant Condition

The Courant condition was introduced in Chapter 3 for the 1D telegrapher's equation, where it was shown to impose an upper bound on the wave velocity that can be used in simulations or estimated from the data for a given spatiotemporal grid (see Secs 3.3 and 3.5.1.3). The 3D telegrapher's equation used here, and previously in Chapter 2, is also subject to this condition. To ensure stability of the 3D discretized process model Eq. (2.7), the Courant number ν [21], which for a 3D hyperbolic equation is given by

$$\nu = \frac{b_x \Delta t}{\Delta x} + \frac{b_y \Delta t}{\Delta y} + \frac{b_z \Delta t}{\Delta z}, \tag{4.3}$$

should be less than $\nu_{\text{max}} = \sqrt{3}$ for a nearest neighbor approximation of the Laplacian term [132]. Here (b_x, b_y, b_z) and (Δx, Δy, Δz) are the wave velocities and the spatial grid size in each of the three Cartesian coordinates, respectively. Given the voxels are cubic (i.e., $\Delta x = \Delta y = \Delta z$) and the wave velocity is assumed to be isotropic (i.e., $b_x = b_y = b_z = b$), the Courant condition becomes

$$\frac{b \Delta t}{\Delta x} \leqslant \frac{1}{\sqrt{3}}. \tag{4.4}$$

Compared to the 1D case [see Eq. (3.13)], the Courant condition in 3D more severely restricts the maximum wave velocity b_{max} for a given spatiotemporal grid. For the values of Δx and Δt used by the whole-brain inverse solution in Chapters 2 and 4, $b_{\text{max}} = 1.03$ m s^{-1}, which is well below the 10 m s^{-1} cortical activity is known to propagate at [105, 119]. Furthermore, to ensure numerical stability in practice, b_{max} must be below this value by approximately $10 - 20\%$. Therefore, the Courant condition can prevent the inverse solution from correctly estimating b — and as a result other model parameters — from clinical EEG recordings. In Chapter 3 the temporal undersampling study investigated this issue on a simulated 1D cortex and found parameter estimation to be inaccurate, while state estimation remained robust. In this chapter, we extend this work into 3D to study the impact of temporal undersampling on a whole-brain inverse solution.

4.2.4 Parameter Estimation

Owing to its success in Chapters 2 and 3, optimal values for the process model parameters and noise covariances are once again selected through numerical minimization of the AIC. The reader is referred to Secs 2.3.3 and 3.5.2 for further

information regarding this technique. Similarly to Chapter 3, we define two process model parameter vectors; one for fitting spatially uniform parameters $\boldsymbol{\vartheta}_u = (c_{f_n}, c_\zeta, b)$, and the other for parameters with a sinusoidal profile $\boldsymbol{\vartheta}_s = (c_{f_n}, d_{f_n}, c_\zeta, d_\zeta, b)$. Recall that the process and observation noise covariances ($\sigma^2_{\boldsymbol{\eta}}$ and $\sigma^2_{\boldsymbol{\epsilon}}$) are also estimated in this step. The AIC is calculated via Eqs (3.43) and (3.44). As done in Chapter 3, constrained optimization is used here. The parameters are given bounds consistent with the telegrapher's equation definition; these are displayed in Table 4.2. To ensure numerical stability of the process model, the upper bound for b is restricted to 90% of its absolute maximum b_{max}. Since a reference was selected from the set of electrodes, we again exclude one of the electrode sites from the analysis [47], so the number of channels is $N_c = 18$. As in Chapter 2, Pz is chosen to be omitted. All datasets were normalized to unit variance after referencing.

For each of the five EEG datasets used in this study, the SWKF's optimal parameters are estimated assuming the spatial profile for the model parameters is either uniform or sinusoidal. The filter initialization follows the procedure employed in Chapter 2 where $\tilde{\boldsymbol{j}}_{KF}(v, 1|1)$ is set to a $\mathbf{0}$ column vector, and $\tilde{\boldsymbol{P}}(v, 1|1)$ to an identity matrix for all voxels. For the simulated data, each model parameter's starting value is randomly selected from an interval spanning $\pm 50\%$ of the parameter's simulated value. This interval is modified by the parameter bounds in Table 4.2 where appropriate. The starting values for the clinical data are randomly chosen from an interval covering the broadest possible range of permitted values. To allow filter transients to pass, the AIC is computed from the 150^{th} time point onwards. Once again, all components of this algorithm are implemented in Matlab [91] and run on an IBM ThinkPad R51 (Intel Pentium 1.6 GHz, 1 GB RAM). The AIC minimization is performed by Matlab's 'fmincon' function which finds the minimum of a constrained nonlinear multivariable function using sequential quadratic programming (see Sec. 3.5.2 for relevant references). The same convergence criterion used in Chapters 2 and 3 is employed again; i.e, the optimization algorithm terminates when the change in the AIC at the next step is less than 1×10^{-10}. The maximum number of filter runs per optimization is set to 500; with each filter iteration taking approximately 5 minutes to complete, this means optimized parameters can be computed within 42 hours. However, in practice, convergence is typically achieved much sooner.

Table 4.2: Bounds for the KF parameters. These bounds are imposed on the parameter estimates by the optimization step (see Sec. 4.2.4).

Parameter	lower bound	upper bound
c_{f_n} (Hz)	0	—
$c_{f_n} - \lvert d_{f_n} \rvert$ (Hz)	0	—
c_ζ	0	—
$c_\zeta - \lvert d_\zeta \rvert$	0	—
b (m s^{-1})	0	$0.9 b_{\text{max}}$
$\log_{10}(\sigma^2_{\boldsymbol{\eta}})$	—	—
$\log_{10}(\sigma^2_{\boldsymbol{\epsilon}})$	—	—

4.2.5 Evaluating Kalman Filter Performance

The methods described in Chapters 2 and 3 for evaluating KF performance are again employed in this study. For a detailed overview the reader is referred to Secs 2.1, 2.5, 3.5.3, and the references cited therein. As discussed previously, simulation studies allow KF performance to be assessed through the direct comparison of simulated and estimated values for the current dipoles and model parameters. These comparisons, including calculation of the RMSE, are performed for the datasets with transient (TU and TV) and driven (DU and DV) dynamics in Secs 4.3.1 and 4.3.2, respectively. Once again, the diagnostic tests which evaluate the statistical properties of the innovation sequence are applied to each of the optimized filters in this study. The results of these tests are shown in Table 4.4 and discussed throughout Sec. 4.3.

4.3 Results

In this section, the optimized inverse solutions for the transient (TU and TV), driven (DU and DV) and clinical EEG are presented for the SWKF, fitting both uniform and sinusoidal spatial profiles for f_n and ζ. The results are summarized in Tables 4.3 and 4.4, and Figs 4.2 – 4.12. More general comments regarding the inverse solution and its performance are made in Sec. 4.4.

4.3.1 Simulated EEG Data: Transient Dynamics

We begin by describing the performance of the inverse solution when applied to the EEG for dataset TU [see Fig. 4.2(b) and Sec. 4.2.1.1]. Since the SWKF exhibits similar performance when fitting either uniform or spatially varying parameters, the results for both filters will be discussed together. Firstly, the inverse solutions estimated the underlying current dipoles quite accurately. This is demonstrated by the agreement between the simulated and estimated current dipoles, which can be observed by: (i) comparing Fig. 4.2(a) with Figs 4.2(c) and 4.2(e), and (ii) looking at the time series shown in Fig. 4.12(a). The small covariance and RMSE values for TU shown in Fig. 4.12(a) and Table 4.3, respectively, also indicate an accurate inverse solution. As may be expected, the introduction of spatially varying parameters does not significantly alter inverse solution accuracy. Estimates of the spatial profile for f_n are reasonably accurate [see Fig. 4.2(g)], while ζ is overestimated [see Fig. 4.2(h)]. The wave velocity was underestimated for both uniform and spatially varying parameters (see Table 4.3). Each filter estimated the observation noise covariance reliably. To further assess the filters' performance, the innovation sequences were analyzed. On inspection both sets of innovations appear quite similar [see Figs 4.2(d) and 4.2(f)]. To provide a quantitative evaluation, the statistical properties of the innovation sequences were computed. These showed the innovations to be Gaussian, unbiased, and have the correct magnitude (see Table 4.4). However, the autocorrelations and the spectral entropy values indicate that significant correlations are present in a majority of the innovations. This means some of the source dynamics remain unmodeled, which is expected since the optimization step 'explained' the data with a parameter set that differs from the true parameter values.

The second dataset investigating transient dynamics is TV, which introduces spatial variation of the model parameters f_n and ζ for the first time. Details of the spatial profiles were provided in Sec. 4.2.1.1, and the simulated current dipoles and EEG are shown in Figs 4.4(a) and 4.4(b), respectively. The results for dataset TV are now described.

Once again the filter displayed comparable performance when fitting either uniform or sinusoidal parameter profiles. The filter reconstructed the current dipoles accurately, as shown by the concordance between the simulated current dipoles shown in Fig. 4.4(a), and the estimated activity plotted in Figs 4.4(c) and 4.4(e).

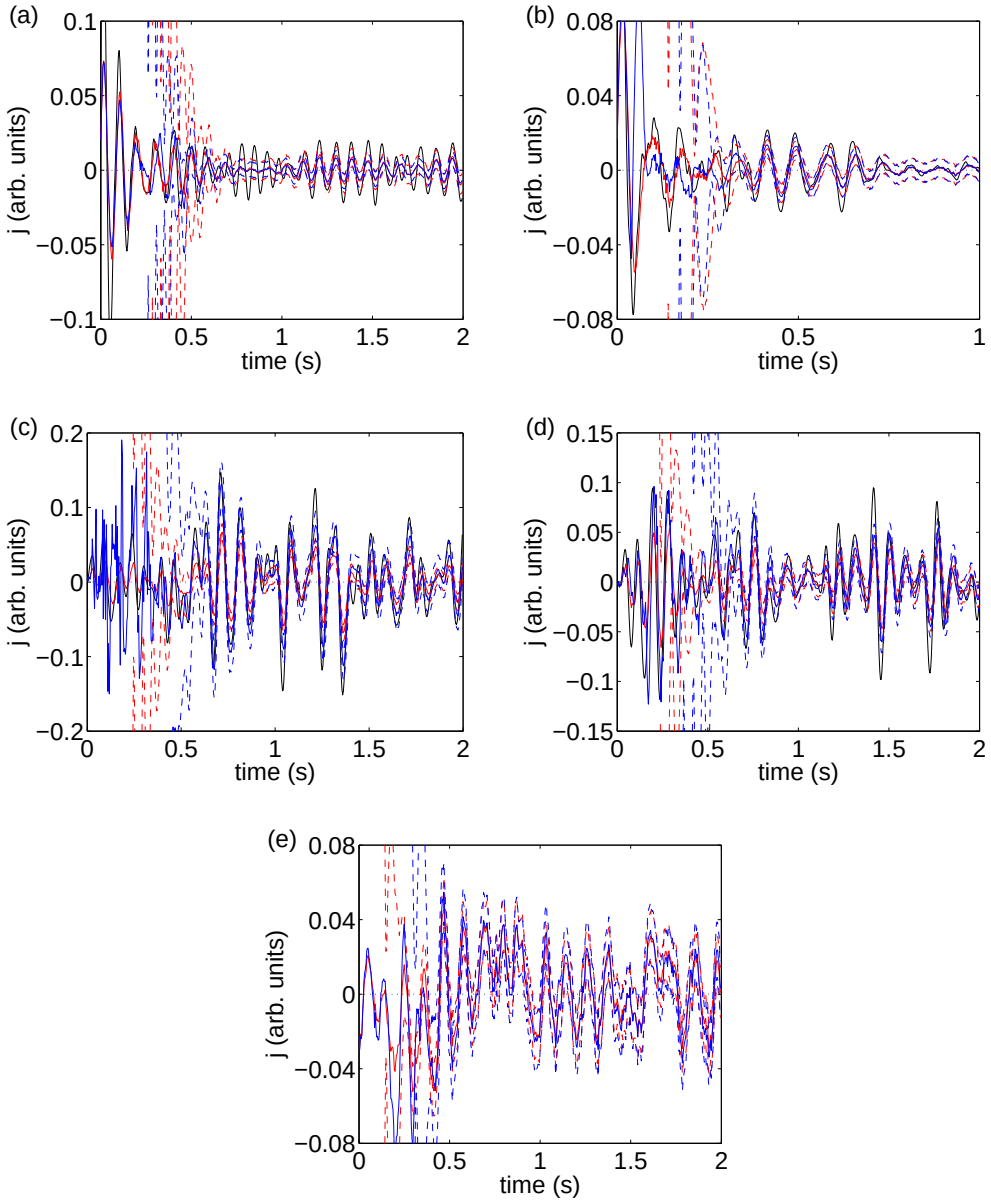

Figure 4.12: Current dipole time series for a voxel in the right occipital pole for the simulated data (TU, TV, DU, and DV) and the clinical EEG recording. Frame (a) displays the time series at this particular voxel for TU. The solid black line indicates the simulated current vector (available for TU, TV, DU, and DV datasets only), while the solid red and blue curves represent the currents estimated by the SWKF when fitting uniform or sinusoidal parameter profiles, respectively. The dashed lines indicate the 95% confidence intervals for the estimates using the same color coding. Frames (b), (c), (d), and (e) show the same information for datasets TV, DU, DV, and the clinical EEG, respectively.

Table 4.3: Estimated parameters and performance metrics for simulated and clinical EEG. The asterisk next to the estimate of b for the clinical EEG when fitting spatially uniform parameters indicates its value is at the optimizer's lower bound (see Sec. 4.3.3 for more details). For the purposes of comparison, the simulated parameter values used to generate datasets TU, TV, DU, and DV are displayed in rows TUsim, TVsim, DUsim, and DVsim, respectively. Note that all estimated parameters (except σ_{ϵ}^2) describe the spatially whitened current dipole $\tilde{\boldsymbol{J}}$, rather than $\boldsymbol{J}$.

Dataset	c_{f_n} (Hz)	d_{f_n} (Hz)	c_ζ	d_ζ	b (m s^{-1})	$\sigma_{\boldsymbol{\eta}}^2$	$\sigma_{\boldsymbol{\epsilon}}^2$	AIC	RMSE
TUsim	10.00	—	0.010	—	0.80	—	9.76×10^{-3}	—	—
TU	10.42	—	0.152	—	0.373	3.29×10^{-8}	8.30×10^{-3}	-5991	7.54×10^{-3}
TU	10.87	-0.509	0.150	-0.0240	0.320	2.94×10^{-8}	8.55×10^{-3}	-6414	1.04×10^{-2}
TVsim	9.00	-2.00	0.050	0.050	1.60	—	9.80×10^{-3}	—	—
TV	7.93	—	0.486	—	1.43	3.82×10^{-9}	8.85×10^{-3}	-6979	7.01×10^{-3}
TV	13.04	1.35	0.274	-0.0400	0.278	7.71×10^{-9}	9.12×10^{-3}	-7315	1.16×10^{-2}
DUsim	10.00	—	0.100	—	0.80	—	9.58×10^{-3}	—	—
DU	9.98	—	0.175	—	0.0983	1.07×10^{-7}	7.01×10^{-3}	-2387	6.46×10^{-3}
DU	14.61	9.13	0.709	0.640	0.752	1.16×10^{-7}	8.83×10^{-3}	-5980	2.30×10^{-2}
DVsim	9.00	-2.00	0.100	0.100	0.80	—	9.81×10^{-3}	—	—
DV	8.36	—	0.247	—	0.581	8.60×10^{-8}	7.15×10^{-3}	-2585	7.09×10^{-3}
DV	7.40	-2.39	0.513	0.461	0.701	7.16×10^{-8}	7.58×10^{-3}	-3480	1.01×10^{-2}
Clin.	5.48	—	0.811	—	1.00×10^{-10}*	1.15×10^{-7}	3.48×10^{-10}	-3113	—
Clin.	5.45	-1.36	0.917	0.700	5.52×10^{-7}	1.08×10^{-7}	6.08×10^{-15}	-3317	—

Table 4.4: Results of KF performance evaluation tests for simulated and clinical EEG data. The KF's spatial profile for f_n and ζ is either uniform (U) or sinusoidal (S) as indicated. The number of channels whose innovation is Gaussian and unbiased according to Kolmogorov-Smirnov (KS) and t-tests, respectively, are shown (i.e., $P > 0.05$). Also displayed is the overall noise level statistic (NLS) Eq. (3.47) and its 95% confidence intervals (CI). Innovation whiteness is characterized by: (1) counting the channels $< 90\%$ of its autocorrelation (AC) within the 2 standard deviations for a white signal; and (2) showing the range of the innovation spectral entropy ($\boldsymbol{\Delta Y}$ SE) across all channels.

Dataset	Profile	KS	t-test	NLS (95% CI)	AC	$\boldsymbol{\Delta Y}$ SE
TU	U	18	17	17.6 (17.4-18.6)	16	0.647-0.974
TU	S	18	18	17.6 (17.4-18.6)	15	0.682-0.974
TV	U	17	17	17.8 (17.4-18.6)	14	0.840-0.974
TV	S	18	18	17.8 (17.4-18.6)	16	0.647-0.979
DU	U	18	17	18.0 (17.4-18.6)	12	0.724-0.974
DU	S	18	18	17.9 (17.4-18.6)	18	0.689-0.930
DV	U	18	17	18.0 (17.4-18.6)	17	0.683-0.960
DV	S	18	18	17.8 (17.4-18.6)	16	0.667-0.964
Clin.	U	18	11	18.0 (17.4-18.6)	15	0.777-0.984
Clin.	S	18	12	18.0 (17.4-18.6)	15	0.793-0.981

The time series for a single voxel is shown in Fig. 4.12(b) and confirms the agreement between the simulated and estimated activity. Small error covariances are also noted in Fig. 4.12(b). The RMSE values are small for both parameter profiles. When fitting a uniform profile for f_n the optimization step selected a value between its simulated minimum and maximum values, as shown in Fig. 4.4(g). However, when fitting the sinusoidal profile, we find that f_n is slightly overestimated and the spatial trend is not correctly identified. For ζ, both spatial profiles significantly overestimate its value; the uniform profile more than the sinusoidal one [see Fig. 4.4(h)]. The exact parameter values can be read from Table 4.3. The wave velocity for the uniform profile was only 10% below its true value, while the sinusoidal profile's b value was underestimated by nearly a factor of six. Turning to the noise covariances; the observation noise was estimated accurately, while the process noise covariance had similar values for both spatial profiles.

The innovations also revealed well-tuned filters for the inverse solutions fitting uniform and sinusoidal parameters. The diagnostic tests showed the innovations to be Gaussian, unbiased and have the correct magnitude (see Table 4.4). As observed in Figs 4.4(d) and 4.4(f), some structure is still present in the innovation sequences. Formal statistical testing confirms the presence of correlations in most of the innovations. Spectral entropy values also vary greatly between channels.

In summary, when the SWKF was applied to simulated EEG recordings with transient dynamics, it was found to be well-tuned and produce accurate state estimates. Parameter estimation was less reliable, particularly for ζ and b, these being consistently over- and under-estimated, respectively. Interestingly, this phenomenon of fitting the EEG with parameter values significantly different from their simulated ones was also observed in the 1D simulation study in Chapter 3 when the electrode density was reduced (IED $= 0.05$ m), although the results were not reported in detail (see Sec. 3.7). Returning to the whole-brain inverse solution, we note that the parameters' spatial trends were also not correctly identified. The relatively poor performance of the parameter estimation seen here is in contrast to the accurate parameter estimates produced by the SWKF when applied to EEG displaying transient dynamics in the 1D simulation study (see Sec. 3.6.1). This issue will be discussed in Sec. 4.4. The SWKF is now applied to simulated EEG produced by driven dynamics. Its performance is described in the following section.

4.3.2 Simulated EEG Data: Driven Dynamics

The first dataset investigating driven dynamics is DU, which contains a single source and uniform spatial parameters, as outlined in Sec. 4.2.1.2. The simulated current dipoles and EEG for DU are shown in Figs 4.6(a) and 4.6(b), respectively. When the two spatial profiles were applied to the simulated EEG for DU, they were found to produce quite similar results. Therefore, the results for both spatial profiles are discussed together. Firstly, the inverse solution estimated the current dipoles quite accurately. This is illustrated for the two profiles by the close agreement between the simulated and estimated current dipoles, which is evident when comparing Fig. 4.6(a) with Figs 4.6(c) and 4.6(e). The single voxel time series plotted in Fig. 4.12(c) also demonstrates the accuracy of both profiles. We also note that the RMSE values are small for both parameter profiles. In terms of parameter estimation, the same inaccuracies observed for the simulated data with transient dynamics are present here. Estimation of f_n was again more reliable, with the uniform profile almost exactly matching its simulated value [see Fig. 4.6(g)]. While c_{f_n} for the sinusoidal profile was similar to its true value, this profile introduced spatial variation of f_n that was not present in the simulated data. As already observed for datasets TU and TV, ζ is again inflated by both spatial profiles [see Fig. 4.4(h)]. Furthermore, the estimated sinusoidal profile displays significant, and incorrect, spatial variation of ζ. Wave velocity estimation continues to be unpredictable. It was greatly underestimated for the uniform profile, while for the sinusoidal profile it was within 6% of its true value. The process noise covariances have similar values for both spatial profiles, and the observation noise was estimated accurately.

Turning to the innovation sequences plotted in Figs 4.6(d) and 4.6(f), there are clearly temporal correlations in some channels, particularly O1 and O2, these being closest to the unmodeled drive term. As shown in Table 4.4, the innovations for both profiles were found to be Gaussian, unbiased, and to have the correct magnitude. However, as indicated by the time series plots, there are significant small amplitude correlations present in nearly all channels for both the uniform and sinusoidal profiles. This finding is consistent with the considerable variation seen in the innovation spectral entropies.

The second dataset containing driven dynamics is DV. This dataset is the most complex in the simulation study and was produced using sinusoidal parameter profiles and two drive terms (see Sec. 4.2.1.2). Figure 4.8(a) shows a snapshot of the

simulated current dipoles while the EEG time series is displayed in Fig. 4.8(b). The results for DV follow a similar pattern to DU, with the two parameter profiles producing comparable results. Therefore, the performance of the two filters are reviewed together once more.

Starting with the inverse solutions for DV, both filters were found to reconstruct the current dipoles accurately. This is demonstrated by the concordance between the simulated current dipoles in Fig. 4.8(a) and their estimated counterparts shown in Figs 4.8(c) and 4.8(e). The single voxel time series in Fig. 4.12(d) also confirms the close agreement between the simulated and estimated activity for both spatial profiles. The RMSEs for both profiles have similar values, with the sinusoidal parameters not improving the inverse solution's accuracy. As already seen with TU, TV, and DU, parameter estimation is again found to be less reliable than state estimation, with the optimization step choosing inexact parameter profiles. The DV estimates for f_n are displayed in Fig. 4.8(g). This figure shows when fitting the uniform profile for f_n the optimization step selects a value between the parameter's true maximum and minimum value. For the sinusoidal case, f_n is slightly overestimated and the sign of the amplitude is incorrect. Figure 4.8(h) shows that ζ is overestimated again for both profiles, and the sign of amplitude for the sinusoidal profile is reversed. The wave velocity was underestimated by 27.5% and 12.5% when fitting uniform and sinusoidal parameters, respectively. The process noise covariance values are small and similar to each other. The observation noise covariance is accurately estimated for both profiles.

Looking at the innovations shown in Figs 4.8(d) and 4.8(f), some temporal correlations are clearly visible, especially in the channels nearest the drive terms (FP1, FP2, O1, and O2). The diagnostic tests show the innovations for both spatial profiles to be Gaussian, unbiased and have the correct magnitude. Confirming the observations from the innovation time series, significant correlations are found in nearly all channels for both profiles. The innovation spectral entropies show a spread similar to those seen for TU, TV, and DU.

In conclusion, the results presented here have demonstrated the ability of SWKF to solve the whole-brain EEG inverse problem for the more realistic case of driven dynamics. For both driven datasets, the filters produced accurate state estimates, and generally positive results for the diagnostic tests; indicating that the filters are indeed well-tuned. As seen previously in the transient dynamics study, parameter

estimates were found to be less reliable, particularly b and ζ, who are again under- and overestimated, respectively. Once more the parameters' spatial trends were not correctly identified, and the addition of spatially varying parameters did not improve filter performance measurably. Similarly to the transient study, widespread correlations were consistently seen in a majority of the innovation sequences. Unlike Chapter 3, where the introduction of unmodeled drive terms significantly affected both state and parameter estimation, we find that the performance of the SWKF appears stable regardless of whether the EEG data displays transient or driven dynamics, and irrespective of the spatial profile being fitted. This is a somewhat surprising result and will be discussed further in Sec. 4.4.

4.3.3 Clinical EEG Data

We now review the performance of the inverse solution when applied to the clinical EEG recording displayed in Fig. 4.10(a). The reader is referred to Sec. 4.2.2 for further details regarding this dataset. Figures 4.10(b) and 4.10(d) show the spatial distribution of the inverse solution at the same time step for the uniform and sinusoidal parameter profiles, respectively. These figures both show an area of activity at the right occipital pole, which is consistent with an EEG recording displaying prominent occipital alpha activity that is strongest on the right side. Figure 4.12(e) displays the time series for the coronal component of the inverse solution at a voxel in the right occipital pole. As expected, a large amplitude alpha oscillation is observed in the time series for both parameter profiles. These findings are consistent with an eyes-closed EEG recording. As seen throughout the simulation study, the state estimates are again similar regardless of the parameter profile employed. Turning to the parameter estimates, we find that when fitting the uniform profile $f_n \approx 5.5$ Hz and $\zeta \approx 0.8$ [see Figs 4.10(f) and 4.10(g)]. The high damping coefficient is similar to the value selected in the KF evaluation study in Chapter 2. Section 2.6 contains a discussion of this result. When fitting the sinusoidal parameter profile the optimizer selected similar mean values to those chosen for the uniform profile (see Table 4.3). Interestingly, the amplitude parameters are chosen so that f_n increases anterior to posterior and ζ does the opposite. These spatial trends match those typically seen in an eyes-closed EEG recording [see Sec. 4.2 and Fig. 4.11]. However, given the unpredictability of parameter estimates seen throughout the simulation study, we would advise caution when interpreting the significance of these purported trends.

In this chapter, the wave velocity b has been estimated directly to ensure parameter values are consistent with the telegrapher's equation definition. The Courant condition has also imposed an upper limit on all wave velocity estimates which is well below the speed at which cortical activity is known to propagate (see Sec. 4.2.3). Here we find $b \approx 0$ for both spatial profiles, which means the spatial component of the process model does not contribute to the inverse solution. This finding was also encountered and discussed in Chapter 2 (see Sec. 2.5.3). The process noise covariance had similar values to those seen in the simulation study, while the observation noise covariances were very small like the values selected for the clinical EEG in Chapter 2 (see Sec. 2.4.2). The innovations shown in Figs 4.8(d) and 4.8(f) appear quite similar, with some temporal correlations visible, particularly in channels O1 and O2. The diagnostic tests show the innovations for both spatial profiles to be Gaussian and have the correct magnitude. Most of the innovations are unbiased. Significant correlations are found in 15 out of the 18 channels for both profiles. The innovation spectral entropies show a range of values consistent with a majority of the innovation sequences being non-white.

4.4 Discussion and Conclusions

In this chapter, the whole-brain KF-based EEG source localization algorithm from Chapter 2 has been reinvestigated with the introduction of spatially varying model parameters. This scenario better reflects true brain dynamics and follows the method from Chapter 3 where f_n and ζ — which control the strength and frequency of the telegrapher's equation's resonance — are a simple function of position. Estimation of this spatial variation was incorporated into the filtering algorithm by allowing the same parameters in the process model of the KF to have spatial variations of the same functional form. Model and noise parameters were selected through AIC minimization. The KFs were applied to simulated data, with uniform and spatially varying parameters for both transient and driven dynamics, and a clinical EEG recording. The inverse solutions for each scenario were analyzed in detail. This analysis included the application of standard diagnostic tests for evaluating KF performance. The key findings and future directions arising from this study are:

(i) The whole-brain simulation studies demonstrated that the SWKF reliably solves the state estimation problem for wave-like patterns of brain activity generated

by either transient or driven dynamics. This result is consistent with the findings described in Chapter 3. The state estimates for the clinical EEG also appeared to be accurate. The introduction of spatially varying model parameters did not significantly improve the accuracy of state estimates for any dataset.

(ii) Parameter estimation was much less accurate, even for the transient datasets where the drive term — which is known to distort parameter estimates — was not present. This is unlike the 1D study where the parameter estimates for the data with transient dynamics were very accurate (see Sec. 3.6.1). This observation suggests the involvement of other factors. A possible explanation may be provided by the tendency for the inverse solutions for both simulated and clinical data to underestimate b and overestimate ζ. This behavior was consistently observed for both parameter profiles and may be due to the sparse spatial sampling provided by the 10-20 system, which leads to — under AIC minimization — the EEG data being best explained by a model where each electrode time series is generated by its own source ('region of influence'), rather than each electrode observing a single spatiotemporal process (i.e., the brain's activity) from a different vantage point. Further investigation with higher density EEG recordings will be required to adjudicate on this possibility. As mentioned in Sec. 4.3.1, this phenomenon was also observed in the 1D simulation study in Chapter 3 when the electrode density was reduced. Estimates of f_n were more accurate, but still occasionally introduced spatial variations that were not present in the data. Finally, no clear trend was identified in the process noise covariance values, with similar values being selected irrespective of parameter profile or EEG dataset (see Table 4.3). These results mirror those seen for the driven datasets in Chapter 3 (see Sec. 3.6.2), and show that the filter places a similar level of confidence in the predictions of the process model regardless of the parameter profile employed.

(iii) An investigation of temporal undersampling was not undertaken for the simulated data in this pilot study. However, the inverse solutions for the clinical EEG were temporally undersampled because the Courant condition restricted the wave velocity in the gray matter to $\lessapprox 1$ m s^{-1}, which is up to an order of magnitude below its true value. For both parameter profiles $b \approx 0$ (see Table 4.3). This result is similar to that seen in Chapter 2 where the estimated spatial interaction term b_1 also had limited impact on the quality of the inverse solu-

tion due to its small value (see Sec. 2.4.2). Several factors that are likely to contribute to the low b estimates are: low density EEG, the Courant condition, spatial whitening, and the form of the process model. Further investigation is required to determine the relative contributions of these factors.

(iv) During this study we again found that the AIC minimization can be dependent upon initialization. Hence in future work, we recommend running the optimization step from multiple starting points, which would generate a probability distribution for the parameters, from which the most likely value could be determined. As noted previously, this strategy has been implemented in other inverse solutions such as [116].

In light of the above findings, any parameter estimates produced by this inverse solution should be interpreted with caution. The issue of inaccurate parameter estimates arising in KF schemes performing simultaneous state and parameter estimation has recently been discussed in [20, 124]. Here the authors point out that this is expected and acceptable in the setting of imperfect process models, where *parameters* may lose some biophysical meaning/interpretability in the course of providing more accurate *state* estimates; as we have witnessed in this study. Whether this is acceptable depends upon the goals of the inverse solution. Regardless, future work should aim to introduce more realistic dynamic models into estimation algorithms such as the KF. The value of physiology-based models is evident in the model fitting study presented in [118], which demonstrates how meaningful estimates of parameters that describe the spatiotemporal features of brain activity can be extracted from EEG data. Future inverse solutions should incorporate dynamic models such as these and introduce spatial and/or temporal variation of the relevant model parameters; something which is yet to be extensively investigated. However, work in this field is already underway. For example, a recent study [122] investigated the feasibility of using average parameter sets to represent spatially varying model parameters in an experiment using a voltage-sensitive dye to probe brain dynamics. Another promising study [75] fitted parameters from a physiology-based model to ERPs recorded using an extended 10-20 system to produce spatial parameter maps to compare depression and control groups.

Chapter 5

Concluding Remarks and Future Directions

This chapter provides a brief overview of the outcomes of the thesis. It begins with an outline of the key findings and outstanding issues arising from the research, and closes with a discussion of possible directions for future work. Since the individual studies have already been discussed in detail in Chapters 2 – 4, this chapter presents a summary of the major results and future directions.

This thesis has investigated EEG source localization using Kalman filtering, which is a widely used and robust technique for estimating states in dynamical systems from indirect and uncertain measurements. The KF is an obvious candidate for solving the dynamic EEG inverse problem since it naturally brings together the four core components of the inverse problem that were described in Chapter 1 (see Sec. 1.2). In Chapter 2 we began by introducing a Kalman filter algorithm (SWKF) that solved the EEG inverse problem for a realistic head model. This filter employed a discretized telegrapher's equation to describe the dynamics of the current dipoles, or more precisely that of their spatial Laplacian. Such an equation represents an appropriate starting point for modeling large-scale brain activity. Furthermore, the filter was operated in a Laplacianized state space to reduce its computational burden. Spatially uniform model parameters were fitted to both simulated and clinical EEG using AIC minimization. The resulting inverse solutions were found to accurately reconstruct the underlying source dynamics. This study also applied standard diagnostic tests to objectively evaluate KF performance. These tests compute the statistical properties of the innovation sequences and subsequently identified spatial

variation in filter performance which could potentially be improved by spatially-varying model parameters.

In Chapter 3 the SWKF algorithm introduced in Chapter 2 was studied in detail using 1D simulations which reduced the complexity of the inverse problem. The simulated EEG data was generated using a telegrapher's equation and a simplified volume conductor model. Two of the model parameters were given spatial profiles of a simple functional form (Gaussian) to better reflect true brain dynamics. To ensure the estimated model parameters maintained a clear biophysical interpretation, the optimization step explicitly constrained their values to ranges consistent with the process model being a telegrapher's equation. This is in contrast to Chapter 2 where unconstrained optimization was used, and parameter estimates could take on values that were discordant with the equation's definition, and therefore potentially lose their interpretability. For the purposes of comparison, inverse solutions were also computed using the optimal LKF. This study found both filters produced accurate state estimates, with the SWKF exhibiting performance similar to the LKF. Standard diagnostic tests also showed both filters to be well-tuned for all scenarios investigated in the 1D study. Spatially varying parameter profiles were able to be correctly identified from the datasets with transient dynamics, while parameter estimates for the driven datasets were less reliable because they were degraded by the unmodeled drive term. Temporal undersampling was also found to distort parameter estimates. This occurred when the Courant condition — for a given spatiotemporal grid — imposed an upper bound on wave velocity estimates that was below the wave velocity's true value.

Chapter 4 returned to the whole-brain EEG inverse problem using the SWKF described in Chapter 2. As in Chapter 3 spatially varying model parameters of a known functional form (sinusoidal) were introduced into the simulated EEG and optimization step to reproduce the typical anterior to posterior variation of the alpha rhythm. Compared to Chapter 2, more complex simulated whole-brain EEG was generated which displayed wave-like patterns and spatially varying dynamics. Like Chapter 3, the optimization step constrained model parameters to ranges consistent with the telegrapher's equation definition, and also imposed the Courant condition on estimates of b. State estimation was again found to be reliable for both simulated and clinical EEG. However, the introduction of spatially varying parameters did not improve state estimation for any dataset. In fact, parameter estimation

was unreliable for both the transient and driven simulated EEG, with b and ζ being consistently under- and overestimated, respectively. For the clinical EEG, the estimated wave velocity was very small, which means the spatial coupling term had a negligible influence on the inverse solution, as was the case in Chapter 2. Several factors that are likely to contribute to the low b estimates are: low density EEG, the Courant condition, spatial whitening, and the form of the process model. Approaches to overcome this limitation were described throughout this thesis and are summarized below. In light of these results, parameter estimates produced by this inverse solution should be interpreted with caution.

The work presented in this thesis offers many potential directions for future investigation. To conclude this thesis a selection of these are outlined under the following three categories: (i) improvements to the SWKF algorithm; (ii) more general extensions to Kalman-filter-based EEG source localization; and (iii) emerging applications for Kalman filtering in computational neuroscience.

This thesis has focused on the SWKF algorithm as an example of Kalman-filter-based EEG source localization. To improve its performance some simple steps can be taken. Firstly, higher density EEG recordings — up to ≈ 100 electrodes — should be used as these are known to improve inverse solution performance [97], and will also help provide more accurate parameter estimates. To raise the upper bound imposed on estimates of b by the Courant condition, we suggest running the process model with a finer time step, and assimilating the observations as they become available. Since AIC minimization is used to obtain optimal parameter values, the reliability of these estimates could be improved, and local minimums avoided if the optimization algorithm was run from multiple starting points and the most likely parameter values selected from the resulting probability distributions. To allow broader comments about the SWKF's performance to be made, the algorithm also needs to be applied to multiple clinical EEG recordings.

The telegrapher's equation process model employed by the SWKF does not explicitly model exogenous inputs, which was found to distort parameter estimates in Chapter 3. Possible strategies for handling these inputs vary in complexity and include: using a GARCH formulation of the SWKF algorithm to compensate for the unmodeled drive term [45, 46]; estimating the drive term using basis functions [116]; and/or employing a physiology-based process model that contains an explicit cortical input [80, 117, 119]. In order to facilitate the introduction of the

physiology-based models described in [80, 117, 119] we recommend that future implementations of the SWKF use a two-dimensional (2D) surface model of the cortex, rather than the 3D volumetric model employed currently. An example of an EEG source localization algorithm employing a cortical surface model is described in [139]. Recent advances in modeling and imaging, coupled with increases in computing power has seen the emergence of large scale models of neural activity, that utilize the connectivity information provided by diffusion tensor imaging [136] and provide links to observed signals such as EEG and fMRI [66, 82, 138]. Such models display complex nonlinear dynamics and represent the future of process models for inverse solutions in functional neuroimaging. The need to perform joint state and parameter estimation with increasingly large and complex dynamic models — not just in neuroimaging — is driving the parallel development of specialized algorithms such as the UKF [49, 61, 71], ensemble KF [36], variational filtering schemes [25, 40, 42], and those that apply multiple penalty terms [139]. Hence, there is great potential for a cross-pollination of ideas and techniques between the neuroimaging community and other fields where large-scale nonlinear spatiotemporal inverse problems are solved.

The introduction of model parameters with spatiotemporal variations is a further extension of EEG inverse solutions that is motivated by observations of real brain dynamics and structure, but is yet to be extensively investigated. The ability to image the spatiotemporal variation of model parameters will provide more accurate inverse solutions and additional insights into brain dynamics. This thesis undertook a preliminary investigation of the topic through the introduction of spatially varying model parameters. While this study identified several problems that still require resolution; a combination of the improvements suggested above, coupled with the ongoing development of neural models and estimation techniques, should unlock the rich potential of this approach. Several recent studies have already demonstrated that the introduction of either time-varying [49, 50] or spatially-varying [42] model parameters can improve the performance of a Kalman-filter-based EEG inverse solution. In [4], a novel technique for solving the EEG inverse problem was introduced which uses the EKF framework to integrate multiple source models. Another future direction that is already under investigation is combining the aforementioned neural models with multiple observation models — most often EEG and fMRI — to realize multimodal neuroimaging which further constrains the inverse problem and

provides spatiotemporal resolution beyond what is possible with a single imaging modality [7, 30, 138].

There is growing interest in using model-based closed-loop control of neural dynamics to treat conditions such as epilepsy and Parkinson's disease [68, 79, 81]. Such an approach would offer improved performance over current brain stimulation methods which are all open-loop. The KF, with its ability to provide optimal control of dynamic systems (see Sec. 1.3), is an obvious candidate for any attempt at controlling brain activity. The algorithm's potential has already been recognized and investigated in several simulation studies. For instance, in [125] a KF was used on a simulated cortex to control features of its spatiotemporal dynamics. The same research group has recently expanded upon this work to investigate model-based control of epilepsy [135] and Parkinson's disease [124].

Another interesting and ongoing avenue of study arises from the observation that the human brain and a Kalman filter must both assimilate internally generated predictions, with uncertain and incomplete observations of the external environment in order to generate an optimal estimate of the state of the world. Perhaps unsurprisingly, a number of studies have found that KFs accurately model how humans integrate information, learn, and make decisions [28, 83–85, 96, 145]. These findings are further supported by the growing evidence that animal nervous systems encode and manipulate probabilities [41, 83]. Kalman filters have also been used to model brain regions that represent space-time, such as the hippocampus [16] and the entorhinal cortex [106].

In conclusion the Kalman filter and its variants provide a flexible approach to model-based estimation whose scope extends well beyond EEG source localization. The algorithm's ability to combine dynamic models and measurements within a sound statistical framework to provide reliable state and parameter estimates means it will continue to have an important role in the fields of functional neuroimaging and computational neuroscience for the foreseeable future.

References

[1] ABRAMOWITZ, M., AND STEGUN, I. A., Eds. *Handbook of Mathematical Functions*. Dover, New York, 1965.

[2] AKAIKE, H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control AC-19* (1974), 716–723.

[3] ALEXANDER, D. M., JURICA, P., TRENGOVE, C., NIKOLAEV, A. R., GEPSHTEIN, S., ZVYAGINTSEV, M., MATHIAK, K., SCHULZE-BONHAGE, A., RUESCHER, J., BALL, T., AND VAN LEEUWEN, C. Traveling waves and trial averaging: The nature of single-trial and averaged brain responses in large-scale cortical signals. *NeuroImage 73* (2013), 95–112.

[4] ANTELIS, J. M., AND MINGUEZ, J. DYNAMO: Concurrent dynamic multi-model source localization method for EEG and/or MEG. *Journal of Neuroscience Methods 212* (2013), 28–42.

[5] ARULAMPALAM, M. S., MASKELL, S., GORDON, N., AND CLAPP, T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing 50* (2002), 174–188.

[6] ÅSTRÖM, K. J. Maximum likelihood and prediction error methods. *Automatica 16* (1980), 551–574.

[7] BABAJANI-FEREMI, A., AND SOLTANIAN-ZADEH, H. Multi-area neural mass modeling of EEG and MEG signals. *NeuroImage 52* (2010), 793–811.

[8] BAILLET, S., AND GARNERO, L. A Bayesian approach to introducing anatamo-functional priors in the EEG/MEG inverse problem. *IEEE Transactions on Biomedical Engineering 44* (1997), 374–385.

[9] BAILLET, S., MOSHER, J. C., AND LEAHY, R. M. Electromagnetic brain mapping. *IEEE Signal Processing Magazine 18* (2001), 14–30.

[10] BAR-SHALOM, Y., RONG LI, X., AND KIRUBARAJAN, T. *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. John Wiley & Sons, New York, 2001.

[11] BARTON, M. J., ROBINSON, P. A., KUMAR, S., GALKA, A., DURRANT-WHYTE, H. F., GUIVANT, J., AND OZAKI, T. Evaluating the performance of Kalman-filter-based EEG source localization. *IEEE Transactions on Biomedical Engineering 56* (2009), 122–136.

[12] BASSEVILLE, M. Detecting changes in signals and systems - A survey. *Automatica 24* (1988), 309–326.

[13] BENUCCI, A., FRAZOR, R. A., AND CARANDINI, M. Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron 55* (2007), 103–117.

[14] BERGER, H. Über das elektrenkephalogramm des menschen (On the human electroencephalogram) (in German). *Archiv für Psychiatrie und Nervenkrankheiten 87* (1929), 527–570.

[15] BERRIER, K. L., SORENSEN, D. C., AND KHOURY, D. S. Solving the inverse problem of electrocardiography using a Duncan and Horn formulation of the Kalman filter. *IEEE Transactions on Biomedical Engineering 51* (2004), 507–515.

[16] BOUSQUET, O., BALAKRISHNAN, K., AND HONAVAR, V. Is the hippocampus a Kalman filter? In *Proceedings of the Pacific Symposium on Biocomputing* (1998), pp. 655–666.

[17] BRESSLER, S. L., AND MENON, V. Large-scale brain networks in cognition: emerging methods and principles. *Trends in Cognitive Sciences 14* (2010), 277–290.

[18] CAMPBELL, J. K., SYNNOTT, S. P., AND BIERMAN, G. J. Voyager orbit determination at Jupiter. *IEEE Transactions on Automatic Control AC-28* (1983), 256–268.

[19] CATON, R. The electric currents of the brain. *British Medical Journal 2*, 765 (1875), 278.

[20] CORNICK, M., HUNT, B., OTT, E., KURTULDU, H., AND SCHATZ, M. F. State and parameter estimation of spatiotemporally chaotic systems illustrated by an application to Rayleigh-Bénard convection. *Chaos 19* (2009), 013108/1–10.

[21] COURANT, R., FRIEDRICHS, K., AND LEWY, H. On the partial difference equations of mathematical physics. *IBM Journal of Research and Development 11* (1967), 215–234.

[22] CRESSIE, N. A. C. *Statistics for Spatial Data*. John Wiley & Sons, New York, 1991.

[23] DARVAS, F., SCHMITT, U., LOUIS, A. K., FUCHS, M., KNOLL, G., AND BUCHNER, H. Spatio-temporal current density reconstruction (stCDR) from EEG/MEG-data. *Brain Topography 13* (2001), 195–207.

[24] DAUNIZEAU, J., AND FRISTON, K. J. A mesostate-space model for EEG and MEG. *NeuroImage 38* (2007), 67–81.

[25] DAUNIZEAU, J., FRISTON, K. J., AND KIEBEL, S. J. Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D 238* (2009), 2089–2118.

[26] DAUNIZEAU, J., MATTOUT, J., CLONDA, D., GOULARD, B., BENALI, H., AND LINA, J.-M. Bayesian spatio-temporal approach for EEG source reconstruction: Conciliating ECD and distributed models. *IEEE Transactions on Biomedical Engineering 53* (2006), 503–516.

[27] DAVID, O., HARRISON, L., AND FRISTON, K. J. Modelling event-related responses in the brain. *NeuroImage 25* (2005), 756–770.

[28] DAYAN, P., AND YU, A. J. Uncertainty and learning. *IETE Journal of Research 49* (2003), 171–181.

[29] DENEUX, T., AND FAUGERAS, O. EEG-fMRI fusion of non-triggered data using Kalman filtering. In *Proceedings of the 3rd IEEE International Symposium on Biomedical Imaging* (2006), pp. 1068–1071.

[30] DENEUX, T., AND FAUGERAS, O. EEG-fMRI fusion of paradigm-free activity using Kalman filtering. *Neural Computation 22* (2010), 906–948.

[31] DIAMOND, S. G., HUPPERT, T. J., KOLEHMAINEN, V., FRANCESCHINI, M. A., KAIPIO, J. P., ARRIDGE, S. R., AND BOAS, D. A. Dynamic physiological modeling for functional diffuse optical tomography. *NeuroImage 30* (2006), 88–101.

[32] DING, L., AND YUAN, H. Inverse source imaging methods in recovering distributed brain sources. *Biomedical Engineering Letters 2* (2012), 2–7.

[33] DISSANAYAKE, M. W. M. G., NEWMAN, P., CLARK, S., DURRANT-WHYTE, H. F., AND CSORBA, M. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation 17* (2001), 229–241.

[34] DOUCET, A., DE FREITAS, N., AND GORDON, N., Eds. *Sequential Monte Carlo Methods in Practice*. Statistics for Engineering and Information Science. Springer, New York, 2001.

[35] ERMENTROUT, G. B., AND KLEINFELD, D. Traveling electrical waves in cortex: Insights from phase dynamics and speculation on a computational role. *Neuron 29* (2001), 33–44.

[36] EVENSEN, G. The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics 53* (2003), 343–367.

[37] FEREZOU, I., BOLEA, S., AND PETERSEN, C. C. H. Visualizing the cortical representation of whisker touch: Voltage-sensitive dye imaging in freely moving mice. *Neuron 50* (2006), 617–629.

[38] FLETCHER, R. *Practical Methods of Optimization*, second ed. John Wiley and Sons, Chichester, United Kingdom, 1987.

[39] FOXE, J. J., MCCOURT, M. E., AND JAVITT, D. C. Right hemisphere control of visuospatial attention: line-bisection judgments evaluated with high-density electrical mapping and source analysis. *NeuroImage 19* (2003), 710–726.

[40] FRISTON, K. J. Variational filtering. *NeuroImage 41* (2008), 747–766.

[41] FRISTON, K. J., AND STEPHAN, K. E. Free energy and the brain. *Synthese 159* (2007), 417–458.

[42] FUKUSHIMA, M., YAMASHITA, O., KANEMURA, A., ISHII, S., KAWATO, M., AND SATO, M. A state-space modeling approach for localization of focal current sources from MEG. *IEEE Transactions on Biomedical Engineering 59* (2012), 1561–1571.

[43] GADZHIEV, C. M. New method for checking the statistical characteristics of the innovation sequence of the Kalman filter. *Engineering Simulation 14* (1997), 83–91.

[44] GALKA, A., OZAKI, T., BAYARD, J. B., AND YAMASHITA, O. Whitening as a tool for estimating mutual information in spatiotemporal data sets. *Journal of Statistical Physics 124* (2006), 1275–1325.

[45] GALKA, A., OZAKI, T., MUHLE, H., STEPHANI, U., AND SINIATCHKIN, M. A data-driven model of the generation of human EEG based on a spatially distributed stochastic wave equation. *Cognitive Neurodynamics 2* (2008), 101–113.

[46] GALKA, A., YAMASHITA, O., AND OZAKI, T. GARCH modelling of covariance in dynamic estimation of inverse solutions. *Physics Letters A 333* (2004), 261–268.

[47] GALKA, A., YAMASHITA, O., OZAKI, T., BISCAY, R., AND VALDES-SOSA, P. A solution to the dynamical inverse problem of EEG generation using spatiotemporal Kalman filtering. *NeuroImage 23* (2004), 435–453.

[48] GILL, P. E., MURRAY, W., AND WRIGHT, M. H. *Practical Optimization*. Academic Press, London, 1981.

[49] GIRALDO, E., AND CASTELLANOS-DOMINGUEZ, G. Current density reconstruction from EEG based on a time varying nonlinear physiological model. In *Proceedings of the 5th International IEEE EMBS Conference on Neural Engineering* (2011), pp. 225–228.

[50] GIRALDO, E., DEN DEKKER, A. J., AND CASTELLANOS-DOMINGUEZ, G. Estimation of dynamic neural activity using a Kalman filter approach

based on physiological models. In *Proceedings of the 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (2010), pp. 2914–2917.

[51] GIRALDO-SUÁREZ, E., PADILLA-BURITICÁ, J. I., AND CASTELLANOS-DOMÍNGUEZ, C. G. Dynamic inverse problem solution using a Kalman filter smoother for neuronal activity estimation. *Revista Tecno Lógicas 27* (2011), 33–51.

[52] GORDON, E., Ed. *Integrative Neuroscience: Bringing Together Biological, Psychological and Clinical Models of the Human Brain*. Francis & Taylor Ltd, London, 2000.

[53] GORODNITSKY, I. F., GEORGE, J. S., AND RAO, B. D. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. *Electroencephalography and Clinical Neurophysiology 95* (1995), 231–251.

[54] GRECH, R., CASSAR, T., MUSCAT, J., CAMILLERI, K., FABRI, S., ZERVAKIS, M., XANTHOPOULOS, P., SAKKALIS, V., AND VANRUMSTE, B. Review on solving the inverse problem in EEG source analysis. *Journal of NeuroEngineering and Rehabilitation 5* (2008), Article 25.

[55] GREWAL, M. S., AND ANDREWS, A. P. *Kalman Filtering: Theory and Practice*. Prentice Hall, New Jersey, 1993.

[56] GYLYS, V. B. Kalman filters and nonlinear filters. In *Encyclopedia of Physical Science and Technology*, R. A. Myers, Ed., third ed., vol. 8. Academic Press, London, 2002, pp. 153–175.

[57] HAJIYEV, C. Innovation approach based measurement error self-correction in dynamic systems. *Measurement 39* (2006), 585–593.

[58] HALLEZ, H., VANRUMSTE, B., GRECH, R., MUSCAT, J., DE CLERCQ, W., VERGULT, A., D'ASSELER, Y., CAMILLERI, K., FABRI, S., VAN HUFFEL, S., AND LEMAHIEU, I. Review on solving the forward problem in EEG source analysis. *Journal of NeuroEngineering and Rehabilitation 4* (2007), Article 46.

[59] HÄMÄLÄINEN, M., HARI, R., ILMONIEMI, R. J., KNUUTILA, J., AND LOUNASMAA, O. V. Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics 65* (1993), 413–497.

[60] HÄMÄLÄINEN, M. S., AND ILMONIEMI, R. J. Interpreting magnetic fields of the brain - minimum norm estimates. *Medical and Biological Engineering and Computing 32* (1994), 35–42.

[61] HAYKIN, S., Ed. *Kalman Filtering and Neural Networks*. John Wiley & Sons, New York, 2001.

[62] HE, B., AND LIAN, J. High-resolution spatio-temporal functional neuroimaging of brain activity. *Critical Reviews in Biomedical Engineering 30* (2002), 283–306.

[63] HERCULANO-HOUZEL, S. The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in Human Neuroscience 3* (2009), Article 31.

[64] HU, Z., ZHAO, X., LIU, H., AND SHI, P. Nonlinear analysis of the BOLD signal. *EURASIP Journal on Advances in Signal Processing* (2009), 215409/1–13.

[65] INOUYE, T., SHINOSAKI, K., SAKAMOTO, H., TOI, S., UKAI, S., IYAMA, A., KATSUDA, Y., AND HIRANO, M. Quantification of EEG irregularity by use of the entropy of the power spectrum. *Electroencephalography and Clinical Neurophysiology 79* (1991), 204–210.

[66] IZHIKEVICH, E. M., AND EDELMAN, G. M. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences (USA) 105* (2008), 3593–3598.

[67] JASPER, H. H. The ten twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology 10* (1958), 371–375.

[68] JOHNSON, M. D., LIM, H. H., NETOFF, T. I., CONNOLLY, A. T., JOHNSON, N., ROY, A., HOLT, A., LIM, K. O., CAREY, J. R., VITEK, J. L., AND HE, B. Neuromodulation for brain disorders: Challenges and opportunities. *IEEE Transactions on Biomedical Engineering 60* (2013), 610–624.

[69] JOHNSTON, L. A., DUFF, E., MAREELS, I., AND EGAN, G. F. Nonlinear estimation of the BOLD signal. *NeuroImage 40* (2008), 504–514.

[70] JULIER, S. J., AND DURRANT-WHYTE, H. F. On the role of process models in autonomous land vehicle navigation systems. *IEEE Transactions on Robotics and Automation 19* (2003), 1–14.

[71] JULIER, S. J., AND UHLMANN, J. K. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE 92* (2004), 401–422.

[72] KAIPIO, J. P., KARJALAINEN, P. A., SOMERSALO, E., AND VAUHKONEN, M. State estimation in time-varying electrical impedance tomography. *Annals of the New York Academy of Sciences 873* (1999), 430–439.

[73] KALMAN, R. E. A new approach to linear filtering and prediction problems. *Transactions of the ASME: Journal of Basic Engineering 82* (1960), 35–45.

[74] KANDEL, E. R., SCHWARTZ, J. H., AND JESSELL, T. M. *Principles of Neural Science*, fourth ed. McGraw-Hill, New York, 2000.

[75] KERR, C. C., KEMP, A. H., RENNIE, C. J., AND ROBINSON, P. A. Thalamocortical changes in major depression probed by deconvolution and physiology-based modeling. *NeuroImage 54* (2011), 2672–2682.

[76] KERVINEN, M., VAUHKONEN, M., KAIPIO, J. P., AND KARJALAINEN, P. A. Time-varying reconstruction in single photon emission computed tomography. *International Journal of Imaging Systems and Technology 14* (2004), 186–197.

[77] KIEBEL, S. J., DAUNIZEAU, J., PHILLIPS, C., AND FRISTON, K. J. Variational Bayesian inversion of the equivalent current dipole model in EEG/MEG. *NeuroImage 39* (2008), 728–741.

[78] KIEBEL, S. J., DAVID, O., AND FRISTON, K. J. Dynamic causal modelling of evoked responses in EEG/MEG with lead field parameterization. *NeuroImage 30* (2006), 1273–1284.

[79] KIM, J. W., ROBERTS, J. A., AND ROBINSON, P. A. Dynamics of epileptic seizures: Evolution, spreading, and suppression. *Journal of Theoretical Biology 257* (2009), 527–532.

[80] KIM, J. W., AND ROBINSON, P. A. Compact dynamical model of brain activity. *Physical Review E 75* (2007), 031907/1–10.

[81] KIM, J. W., AND ROBINSON, P. A. Controlling limit-cycle behaviors of brain activity. *Physical Review E 77* (2008), 051914/1–4.

[82] KNOCK, S. A., MCINTOSH, A. R., SPORNS, O., KÖTTER, R., HAGMANN, P., AND JIRSA, V. K. The effects of physiologically plausible connectivity structure on local and global dynamics in large scale brain models. *Journal of Neuroscience Methods 183* (2009), 86–94.

[83] KÖRDING, K. Decision theory: What "should" the nervous system do? *Science 318* (2007), 606–610.

[84] KORDING, K. P., TENENBAUM, J. B., AND SHADMEHR, R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience 10* (2007), 779–786.

[85] KRAKAUER, J. W., MAZZONI, P., GHAZIZADEH, A., RAVINDRAN, R., AND SHADMEHR, R. Generalization of motor learning depends on the history of prior action. *PLoS Biology 4* (2006), 1798–1808.

[86] LAGOPOULOS, J. Functional MRI: an overview. *Acta Neuropsychiatrica 19* (2007), 64–65.

[87] LIU, H., AND SCHIMPF, P. H. Efficient localization of synchronous EEG source activities using a modified RAP-MUSIC algorithm. *IEEE Transactions on Biomedical Engineering 53* (2006), 652–661.

[88] LONG, C. J., DESAI, N. U., HÄMÄLÄINEN, M., TEMEREANCA, S., PURDON, P. P., AND BROWN, E. N. A dynamic solution to the ill-conditioned magnetoencephalography (MEG) source localization problem. In *Proceedings of the 3rd IEEE International Symposium on Biomedical Imaging* (2006), pp. 225–228.

[89] MALCOLM, J. G., SHENTON, M. E., AND RATHI, Y. Neural tractography using an unscented Kalman filter. In *Information Processing in Medical Imaging*, J. L. Prince, D. L. Pham, and K. J. Myers, Eds., vol. 5636 of *Lecture Notes in Computer Science*. Springer, Berlin, 2009, pp. 126–138.

[90] MALMIVUO, J., AND PLONSEY, R. *Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields*. Oxford University Press, New York, 1995.

[91] MATLAB VERSION 7.0.0 (R14). `http://www.mathworks.com`. The MathWorks Inc., Natick, Massachusetts, 2004.

[92] MAYBECK, P. S. *Stochastic models, estimation, and control*, vol. 141 of *Mathematics in Science and Engineering*. Academic Press, New York, 1979.

[93] MAZZIOTTA, J. C., TOGA, A. W., EVANS, A., FOX, P., AND LANCASTER, J. A probabalistic atlas of the human brain: Theory and rationale for its development. *NeuroImage 2* (1995), 89–101.

[94] MEHRA, R. K. Identification of stochastic linear dynamic systems using Kalman filter representation. *AIAA Journal 9* (1971), 28–31.

[95] MEHRA, R. K., AND PESCHON, J. An innovations approach to fault detection and diagnosis in dynamic systems. *Automatica 7* (1971), 637–640.

[96] MEHTA, B., AND SCHAAL, S. Forward models in visuomotor control. *Journal of Neurophysiology 88* (2002), 942–953.

[97] MICHEL, C. M., MURRAY, M. M., LANTZ, G., GONZALEZ, S., SPINELLI, L., AND GRAVE DE PERALTA, R. EEG source imaging. *Clinical Neurophysiology 115* (2004), 2195–2222.

[98] MITCHELL, A. R. *Computational Methods in Partial Differential Equations*. John Wiley & Sons, London, 1969.

[99] MOSHER, J. C., LEWIS, P. S., AND LEAHY, R. M. Multiple dipole modeling and localization from spatio-temporal MEG data. *IEEE Transactions on Biomedical Engineering 39* (1992), 541–557.

[100] MURRAY, C. J. L., AND LOPEZ, A. D., Eds. *The global burden of disease: a comprehensive assessment of mortality and disability from diseases, injuries and risk factors in 1990 and projected to 2020*. Harvard University Press, Cambridge, 1996.

[101] NIEDERMEYER, E., AND LOPES DA SILVA, F., Eds. *Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*, fourth ed. Lippincott Williams & Wilkins, Baltimore, 1999.

[102] NUNEZ, P. L. *Neocortical Dynamics and Human EEG Rhythms*. Oxford University Press, New York, 1995.

[103] NUNEZ, P. L., AND SILBERSTEIN, R. B. On the relationship of synaptic activity to macroscopic measurements: Does co-registration of EEG with fMRI make sense? *Brain Topography 13* (2000), 79–96.

[104] NUNEZ, P. L., AND SRINIVASAN, R. *Electric Fields of the Brain: The Neurophysics of EEG*, second ed. Oxford University Press, New York, 2006.

[105] NUNEZ, P. L., AND SRINIVASAN, R. A theoretical basis for standing and traveling brain waves measured with human EEG with implications for an integrated consciousness. *Clinical Neurophysiology 117* (2006), 2424–2435.

[106] OSBORN, G. W. A Kalman filtering approach to the representation of kinematic quantities by the hippocampal-entorhinal complex. *Cognitive Neurodynamics 4* (2010), 315–335.

[107] PASCUAL-MARQUI, R. D., MICHEL, C. M., AND LEHMANN, D. Low resolution electromagnetic tomography: A new method for localizing electrical activity in the brain. *International Journal of Psychophysiology 18* (1994), 49–65.

[108] PHIPPS, M. C., AND QUINE, M. P. *A primer of statistics: data analysis, probability, inference*, third ed. Prentice Hall, Sydney, 1998.

[109] PLIS, S. M., LANE, T., WEISEND, M. P., AND CALHOUN, V. D. MEG and fMRI for nonlinear estimation of neural activity. In *Proceedings of the Asilomar Conference on Signals, Systems, and Computers* (2009), pp. 1598–1602.

[110] POLYANIN, A. D. *Handbook of Linear Partial Differential Equations for Engineers and Scientists*. Chapman & Hall/CRC, United States of America, 2002.

[111] POSNER, J., HELLERSTEIN, D. J., GAT, I., MECHLING, A., KLAHR, K., WANG, Z., MCGRATH, P. J., STEWART, J. W., AND PETERSON, B. S. Antidepressants normalize the default mode network in patients with dysthymia. *JAMA Psychiatry 70* (2013), 373–382.

[112] POUPON, C., ROCHE, A., DUBOIS, J., MANGIN, J.-F., AND POUPON, F. Real-time MR diffusion tensor and Q-ball imaging using Kalman filtering. *Medical Image Analysis 12* (2008), 527–534.

[113] PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A., AND VETTERLING, W. T. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, New York, 1986.

[114] PRINCE, S., KOLEHMAINEN, V., KAIPIO, J. P., FRANCESCHINI, M. A., BOAS, D., AND ARRIDGE, S. R. Times-series estimation of biological factors in optical diffusion tomography. *Physics in Medicine and Biology 48* (2003), 1491–1504.

[115] RIERA, J. J., FUENTES, M. E., VALDÉS, P. A., AND OHÁRRIZ, Y. EEG-distributed inverse solutions for a spherical head model. *Inverse Problems 14* (1998), 1009–1019.

[116] RIERA, J. J., JIMENEZ, J. C., WAN, X., KAWASHIMA, R., AND OZAKI, T. Nonlinear local electrovascular coupling. II: From data to neuronal masses. *Human Brain Mapping 28* (2007), 335–354.

[117] ROBINSON, P. A. Propagator theory of brain dynamics. *Physical Review E 72* (2005), 011904/1–13.

[118] ROBINSON, P. A., RENNIE, C. J., ROWE, D. L., AND O'CONNOR, S. C. Estimation of multiscale neurophysiologic parameters by electroencephalographic means. *Human Brain Mapping 23* (2004), 53–72.

[119] ROBINSON, P. A., RENNIE, C. J., ROWE, D. L., O'CONNOR, S. C., AND GORDON, E. Multiscale brain modelling. *Philosophical Transactions of the Royal Society of London, B 360* (2005), 1043–1050.

[120] ROSSI, L. A., KRISHNAMACHARI, B., AND KUO, C.-C. J. Distributed parameter estimation for monitoring diffusion phenomena using physical mod-

els. In *Proceedings of the 1st IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks* (2004), pp. 460–469.

[121] RYYNÄNEN, O. R. M., HYTTINEN, J. A. K., LAARNE, P. H., AND MALMIVUO, J. A. Effect of electrode density and measurement noise on the spatial resolution of cortical potential distribution. *IEEE Transactions on Biomedical Engineering 51* (2004), 1547–1554.

[122] SAUER, T. D., AND SCHIFF, S. J. Data assimilation for heterogeneous networks: The consensus set. *Physical Review E 79* (2009), 051909/1–4.

[123] SCHERG, M., AND VON CRAMON, D. Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. *Electroencephalography and Clinical Neurophysiology 62* (1985), 32–44.

[124] SCHIFF, S. J. Towards model-based control of Parkinson's disease. *Philosophical Transactions of the Royal Society of London, A 368* (2010), 2269–2308.

[125] SCHIFF, S. J., AND SAUER, T. Kalman filter control of a model of spatiotemporal cortical dynamics. *Journal of Neural Engineering 5* (2008), 1–8.

[126] SCHMITT, U., LOUIS, A. K., WOLTERS, C., AND VAUKHONEN, M. Efficient algorithms for the regularization of dynamic inverse problems: II. Applications. *Inverse Problems 18* (2002), 659–676.

[127] SCHWEPPE, F. C. Evaluation of likelihood functions for Gaussian signals. *IEEE Transactions on Information Theory 11* (1965), 61–70.

[128] SITZ, A., KURTHS, J., AND VOSS, H. U. Identification of nonlinear spatiotemporal systems via partitioned filtering. *Physical Review E 68* (2003), 016202/1–9.

[129] SOMERSALO, E., VOUTILAINEN, A., AND KAIPIO, J. P. Non-stationary magnetoencephalography by Bayesian filtering of dipole models. *Inverse Problems 19* (2003), 1047–1063.

[130] SPORNS, O., TONONI, G., AND KÖTTER, R. The human connectome: A structural description of the human brain. *PLoS Computational Biology 1* (2005), 245–251.

[131] SUKKARIEH, S., NEBOT, E. M., AND DURRANT-WHYTE, H. F. A high integrity IMU/GPS navigation loop for autonomous land vehicle applications. *IEEE Transactions on Robotics and Automation 15* (1999), 572–578.

[132] SULLIVAN, D. M. Exceeding the Courant condition with the FDTD method. *IEEE Microwave and Guided Wave Letters 6* (1996), 289–291.

[133] TOGA, A. W., AND MAZZIOTTA, J. C., Eds. *Brain Mapping: The Methods*, second ed. Academic Press, London, 2002.

[134] TRUJILLO-BARRETO, N. J., AUBERT-VAZQUEZ, E., AND VALDES-SOSA, P. A. Bayesian model averaging in EEG/MEG imaging. *NeuroImage 21* (2004), 1300–1319.

[135] ULLAH, G., AND SCHIFF, S. J. Tracking and control of neuronal Hodgkin-Huxley dynamics. *Physical Review E 79* (2009), 040901/1–4.

[136] VALDÉS-HERNÁNDEZ, P. A., OJEDA-GONZÁLEZ, A., MARTÍNEZ-MONTES, E., LAGE-CASTELLANOS, A., VIRUÉS-ALBA, T., VALDÉS-URRUTIA, L., AND VALDES-SOSA, P. A. White matter architecture rather than cortical surface area correlates with the EEG alpha rhythm. *NeuroImage 49* (2010), 2328–2339.

[137] VALDES-SOSA, P., RIERA, J., AND CASANOVA, R. Spatio temporal distributed inverse solutions. In *Biomag 96: Proceedings of the Tenth International Conference on Biomagnetism (Volume 1)* (1996), pp. 377–380.

[138] VALDES-SOSA, P. A., SANCHEZ-BORNOT, J. M., SOTERO, R. C., ITURRIA-MEDINA, Y., ALEMAN-GOMEZ, Y., BOSCH-BAYARD, J., CARBONELL, F., AND OZAKI, T. Model driven EEG/fMRI fusion of brain oscillations. *Human Brain Mapping 30* (2009), 2701–2721.

[139] VALDÉS-SOSA, P. A., VEGA-HERNÁNDEZ, M., SÁNCHEZ-BORNOT, J. M., MARTÍNEZ-MONTES, E., AND BOBES, M. A. EEG source imaging with spatio-temporal tomographic nonnegative independent component analysis. *Human Brain Mapping 30* (2009), 1898–1910.

[140] VAN VEEN, B. D., VAN DRONGELEN, W., YUCHTMAN, M., AND SUZUKI, A. Localization of brain electrical activity via linearly constrained

minimum variance spatial filtering. *IEEE Transactions on Biomedical Engineering 44* (1997), 867–880.

[141] VOSS, H. U., TIMMER, J., AND KURTHS, J. Nonlinear dynamical system identification from uncertain and indirect measurements. *International Journal of Bifurcation and Chaos 14* (2004), 1905–1933.

[142] WELCH, P. D. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics AU-15* (1967), 70–73.

[143] WENDEL, K., VÄISÄNEN, O., MALMIVUO, J., GENCER, N. G., VANRUMSTE, B., DURKA, P., MAGJAREVIĆ, R., SUPEK, S., PASCU, M. L., FONTENELLE, H., AND GRAVE DE PERALTA MENENDEZ, R. EEG/MEG source imaging: Methods, challenges, and open issues. *Computational Intelligence and Neuroscience* (2009), 656092/1–12.

[144] WILLSKY, A. S. A survey of design methods for failure detection in dynamic systems. *Automatica 12* (1976), 601–611.

[145] WOLPERT, D. M., GHAHRAMANI, Z., AND JORDAN, M. I. An internal model for sensorimotor integration. *Science 269* (1995), 1880–1882.

[146] XU, W., HUANG, X., TAKAGAKI, K., AND WU, J. Compression and reflection of visually evoked cortical waves. *Neuron 55* (2007), 119–129.

[147] YAMASHITA, O., GALKA, A., OZAKI, T., BISCAY, R., AND VALDES-SOSA, P. Recursive penalized least squares solution for dynamical inverse problems of EEG generation. *Human Brain Mapping 21* (2004), 221–235.

[148] YIN, J., SYRMOS, V. L., AND YUN, D. Y. Y. System identification using nonlinear filtering methods with applications to medical imaging. In *Proceedings of the 39th IEEE Conference on Decision and Control* (2000), pp. 3313–3318.