CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

LONDON MATHEMATICAL SOCIETY STUDENT TEXTS

Managing Editor: Ian J. Leary,
Mathematical Sciences, University of Southampton, UK

67 Elliptic functions, J. V. ARMITAGE & W. F. EBERLEIN
68 Hyperbolic geometry from a local viewpoint, LINDA KEEN & NIKOLA LAKIC
69 Lectures on Kähler geometry, ANDREI MOROIANU
70 Dependence logic, JOUKU VÄÄNÄNEN
71 Elements of the representation theory of associative algebras II, DANIEL SIMSON & ANDRZEJ SKOWROŃSKI
72 Elements of the representation theory of associative algebras III, DANIEL SIMSON & ANDRZEJ SKOWROŃSKI
73 Groups, graphs and trees, JOHN MEIER
74 Representation theorems in Hardy spaces, JAVAD MASHREGHI
75 An introduction to the theory of graph spectra, DRAGOŠ CVETKOVIĆ, PETER ROWLINSON & SLOBODAN SIMIĆ
76 Number theory in the spirit of Liouville, KENNETH S. WILLIAMS
77 Lectures on profinite topics in group theory, BENJAMIN KLOPSCH, NIKOLAY NIKOLOV & CHRISTOPHER VOLL
78 Clifford algebras: An introduction, D. J. H. GARLING
79 Introduction to compact Riemann surfaces and dessins d’enfants, ERNESTO GIRONDO & GABINO GONZÁLEZ-DIEZ
80 The Riemann hypothesis for function fields, MACHIEL VAN FRANKENHUIJSEN
81 Number theory, Fourier analysis and geometric discrepancy, GIANCARLO TRAVAGLINI
82 Finite geometry and combinatorial applications, SIMEON BALL
83 The geometry of celestial mechanics, HANSJÖRG GEIGES
84 Random graphs, geometry and asymptotic structure, MICHAEL KRIVELEVICH *et al*
85 Fourier analysis: Part I - Theory, ADRIAN CONSTANTIN
86 Dispersive partial differential equations, M. BURAK ERDOĞAN & NIKOLAOS TZIRAKIS
87 Riemann surfaces and algebraic curves, R. CAVALIERI & E. MILES
88 Groups, languages and automata, DEREK F. HOLT, SARAH REES & CLAAS E. RÖVER
89 Analysis on Polish spaces and an introduction to optimal transportation, D. J. H. GARLING
90 The homotopy theory of $(\infty,1)$-categories, JULIA E. BERGNER
91 The block theory of finite group algebras I, MARKUS LINCKELMANN
92 The block theory of finite group algebras II, MARKUS LINCKELMANN
93 Semigroups of linear operators, DAVID APPLEBAUM
94 Introduction to approximate groups, MATTHEW C. H. TOINTON
95 Representations of finite groups of Lie type (2nd Edition), FRANÇOIS DIGNE & JEAN MICHEL
96 Tensor products of C^*-algebras and operator spaces, GILLES PISIER
97 Topics in cyclic theory, DANIEL G. QUILLEN & GORDON BLOWER
98 Fast track to forcing, MIRNA DŽAMONJA
99 A gentle introduction to homological mirror symmetry, RAF BOCKLANDT
100 The calculus of braids, PATRICK DEHORNOY
101 Classical and discrete functional analysis with measure theory, MARTIN BUNTINAS
102 Notes on Hamiltonian dynamical systems, ANTONIO GIORGILLI
103 A course in stochastic game theory, EILON SOLAN
104 Differential and low-dimensional topology, ANDRÁS JUHÁSZ
105 Lectures on Lagrangian torus fibrations, JONNY EVANS
106 Compact matrix quantum groups and their combinatorics, AMAURY FRESLON

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

London Mathematical Society Student Texts 107

Inverse Problems and Data Assimilation

DANIEL SANZ-ALONSO
University of Chicago

ANDREW STUART
California Institute of Technology

ARMEEN TAEB
University of Washington

CAMBRIDGE
UNIVERSITY PRESS

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781009414326

DOI: 10.1017/9781009414319

© Daniel Sanz-Alonso, Andrew Stuart and Armeen Taeb 2023

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press & Assessment.

First published 2023

A catalogue record for this publication is available from the British Library

A Cataloging-in-Publication data record for this book is available from the Library of Congress

ISBN 978-1-009-41432-6 Hardback
ISBN 978-1-009-41429-6 Paperback

Cambridge University Press & Assessment has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

CAMBRIDGE
Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

To my wife, parents and brother
DS-A

To the much-loved Crazy 8
AMS

To my parents and my brother and sister
AT

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

Contents

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

© in this web service Cambridge University Press & Assessment www.cambridge.org

Preface

These notes developed as a result of courses taught at Caltech and at the University of Chicago, and aim at providing a clear and concise introduction to the subjects of inverse problems and data assimilation. To cater to students with diverse backgrounds and interests, we complement the material covered in our courses with hands-on assignments; these notes contain several exercises that we have used for this purpose. Additionally, we have found it pedagogically beneficial to ask students to complete an independent project, implementing the methods studied in class to solve an applied problem of their choice. Students can use the bibliographic comments included at the end of each chapter to help them choose and formulate their own projects. The notes are intended to be self-contained, and thus to be useful not only as a teaching resource, but also for independent self-guided learning.

Acknowledgements

A first draft of these notes was created in LaTeX by the students in the Caltech course ACM159 (now ACM/IDS 154), based on lectures presented by the instructor Andrew Stuart, and on input from the course teaching assistant Armeen Taeb. The authors are very grateful to these students, without whom these notes would not exist. The individuals responsible for typesetting the notes, listed in alphabetical order, are: Blancquart, Paul; Cai, Karena; Chen, Jiajie; Cheng, Richard; Cheng, Rui; Feldstein, Jonathan; Huang, De; Idíni, Benjamin; Kovachki, Nikola; Lee, Marcus; Levy, Gabriel; Li, Liuchi; Muir, Jack; Ren, Cindy; Seylabi, Elnaz; Schäfer, Florian; Singhal, Vipul; Stephenson, Oliver; Song, Yichuan; Su, Yu; Teke, Oguzhan; Williams, Ethan; Wray, Parker; Xiao, Fangzhou; Zhan, Eric; Zhang, Shumao. Furthermore, the following students added content to the notes, beyond the materials presented by the instructors: Parker Wray – created

© in this web service Cambridge University Press & Assessment www.cambridge.org

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

an early draft of the Overview; Jiajie Chen – found an alternative proof of early presentations of under-determined inverse problems and smoothing in Gaussian data assimilation; Fangzhou Xiao – providing numerical illustrations of prior, likelihood, and posterior; Elnaz Seylabi and Fangzhou Xiao – catching many typographical errors in an early draft of these notes; Cindy Ren – numerical simulations to enhance understanding of importance sampling; Cindy Ren and De Huang – improving the constants in initial presentations of the approximation error of importance sampling; Richard Cheng and Florian Schäfer – illustrations to enhance understanding of the coupling argument used to study convergence of MCMC algorithms by presenting the finite state-space case; and Ethan Williams and Jack Muir – numerical simulations and illustrations of ensemble Kalman filter and extended Kalman filter that appeared in an early version of these notes. The authors are also grateful to Tapio Helin (LUT University) who used the notes in his own course and provided very helpful feedback on an early draft. The authors are also thankful to Yuming Chen, Andrew Dennehy, Ruoxi Jiang, Phillip Lo, and Walter Zhang (University of Chicago) and Eitan Levin (Caltech) for their generous feedback; they are also grateful to Hwanwoo Kim (University of Chicago) for making substantial improvements to the figures initially provided by the individuals listed above. Finally, the authors are thankful to David Tranah (CUP) whose advice helped to shape this book.

The work of Daniel Sanz-Alonso has been funded by DOE, NGIA, and NSF (USA), and by FBBVA (Spain). The work of Andrew Stuart has been funded by AFOSR, ARL, DoD, NIH, NSF, and ONR (USA), by EPSRC (UK), and by ERC (EU). The work of Armeen Taeb has been funded by the Resnick Fellowship (USA) and by the ETH Foundations of Data Science (Switzerland). All of this funded research has helped to shape the presentation of the material in these notes and is gratefully acknowledged.

© in this web service Cambridge University Press & Assessment www.cambridge.org

Introduction

Aim and Overview of the Notes

The aim of these notes is to provide a clear and concise mathematical introduction to the subjects of Inverse Problems and Data Assimilation, and their interrelations, together with bibliographic pointers to literature in this area that goes into greater depth. The target audiences are advanced undergraduates and beginning graduate students in the mathematical sciences, together with researchers in the sciences and engineering who are interested in the systematic underpinnings of methodologies widely used in their disciplines.

In its most basic form, inverse problem theory is the study of how to estimate model parameters from data. Often the data provide indirect information about these parameters, corrupted by noise. The theory of inverse problems, however, is much richer than just parameter estimation. For example, the underlying theory can be used to determine the effects of noisy data on the accuracy of the solution; it can be used to determine what kind of observations are needed to accurately determine a parameter; and it can be used to study the uncertainty in a parameter estimate and, relatedly, is useful, for example, in the design of strategies for control or optimization under uncertainty, and for risk analysis. The theory thus has applications in many fields of science and engineering.

To apply the ideas in these notes, the starting point is a mathematical model mapping the unknown parameters to the observations: termed the "forward" or "direct" problem, and often a subject of research in its own right. A good forward model will not only identify how the data is dependent on parameters, but also what sources of noise or model uncertainty are present in the postulated relationship between unknown parameters and data. For example, if the desired forward problem cannot be solved analytically, then the forward model may be approximated by a numerical simulation; in this case, discretization may be considered as a source of error. Once a relationship between model parameters,

© in this web service Cambridge University Press & Assessment www.cambridge.org

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

sources of error, and data is clearly defined, the inverse problem of estimating parameters from data can be addressed. The theory of inverse problems can be separated into two cases: (1) the ideal case where data is not corrupted by noise and is derived from a known perfect model; and (2) the practical case where data is incomplete and imprecise. The first case is useful for classifying inverse problems and determining if a given set of observations can, in principle, allow to fully reconstruct the model parameters; this provides insight into conditions needed for existence, uniqueness, and stability of a solution to the inverse problem. The second case is useful for the formulation of practical algorithms to learn about parameters, and uncertainties in their estimates, and will be the focus of these notes.

A model for which a solution exists, is unique, and changes continuously with input (stability) is termed "well-posed." Conversely, a model lacking any of these properties is termed "ill-posed." Ill-posedness is present in many inverse problems, and mitigating it is an extensive part of the subject. Out of the different approaches to formulating an inverse problem, our notes emphasize the Bayesian framework. Nonetheless, practical algorithms in this area include a variety of related optimization approaches, and these are also discussed in detail.

The goal of the Bayesian framework is to find a probability measure that assigns a probability to each possible solution for a parameter u, given the data y. Bayes' formula states that

$$\mathbb{P}(u \mid y) = \frac{1}{\mathbb{P}(y)} \mathbb{P}(y \mid u)\, \mathbb{P}(u).$$

This formula enables calculation of the posterior probability on $u \mid y$, $\mathbb{P}(u \mid y)$, in terms of the product of the data likelihood $\mathbb{P}(y \mid u)$ and the prior information on the parameter encoded in $\mathbb{P}(u)$. The likelihood describes the probability of the observed data y if the input parameter were set to be u; it is determined by the forward model, and the structure of the noise. The normalization constant $\mathbb{P}(y)$ ensures that $\mathbb{P}(u \mid y)$ is a probability measure. There are four primary benefits to this framework: (1) it provides a clear theoretical setting in which the forward model choice, the description of how noise enters the data and the forward model, and *a priori* information on the unknown parameter are all explicit; (2) it provides information about the entire solution space for possible input parameter choices; (3) it naturally leads to quantification of uncertainty and risk in parameter estimates; (4) it is generalizable to a wide class of inverse problems, in finite and infinite dimension, and comes with a well-posedness theory mitigating the ill-posedness of a naive deterministic approach.

The first part of the notes is dedicated to studying the Bayesian framework for inverse problems. Techniques such as importance sampling and Markov Chain

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE
Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

Topic	Inverse Problems	Data Assimilation
Bayesian Formulation	Chapter 1	Chapter 7
Linear Setting	Chapter 2	Chapter 8
Optimization Perspective	Chapter 3	Chapter 9
Gaussian Approximation	Chapter 4	Chapter 10
Sampling	Chapters 5 and 6	Chapters 11 and 12
Kalman Inversion	Chapter 13	

Table 1 *Structure of the notes: the organization of the material emphasizes the unity between the subjects of inverse problems and data assimilation.*

Monte Carlo (MCMC) methods are introduced; these methods have the desirable property that in the limit of an infinite number of samples they reproduce the full posterior distribution. Since it is often computationally intensive to implement these methods, especially in high-dimensional problems, techniques to approximate the posterior by a Dirac or a Gaussian distribution are also discussed, along with related optimization algorithms to determine the best approximation.

The second part of the notes covers data assimilation. This refers to a particular class of inverse problems in which the unknown parameter is the initial condition of a dynamical system or, in the case of stochastic dynamics, the entire sequence of subsequent states of the system, and the data comprises partial and noisy observations of the (possibly stochastic) dynamical system. A primary use of data assimilation is in forecasting, where the purpose is to provide better future estimates than can be obtained using either the data or the model alone. All the methods from the first part of the course may be applied directly, but there are other new methods which exploit the Markovian structure to update the state of the system sequentially, rather than to learn about the initial condition. (But, of course, knowledge of the initial condition may be used to inform the state of the system at later times.)

The third and final part of the notes describes methods for generic inverse problems that build on data assimilation ideas, thus bringing together the material in the first two parts. The structure of the notes, as well as the presentation, emphasizes the inter-relations between inverse problems and data assimilation.

© in this web service Cambridge University Press & Assessment www.cambridge.org

CAMBRIDGE

Cambridge University Press & Assessment
978-1-009-41432-6 — Inverse Problems and Data Assimilation
Daniel Sanz-Alonso , Andrew Stuart , Armeen Taeb
Frontmatter
More Information

As summarized in Table 1, each chapter in the first part (inverse problems) has its counterpart in the second part (data assimilation).

Notation

Throughout the notes we use $\mathbb{N}$ to denote the positive integers $\{1, 2, 3, \ldots\}$, and $\mathbb{Z}^+$ to denote the non-negative integers $\mathbb{N} \cup \{0\} = \{0, 1, 2, 3, \ldots\}$. The symbol I_d denotes the identity matrix on $\mathbb{R}^d$, and Id denotes the identity mapping. We use $|\cdot|$ to denote the Euclidean norm corresponding to the inner-product $\langle a, b\rangle = a^\top b$; we also use the notation $|\cdot|$ to denote the induced norm on matrices.

A symmetric matrix A is positive definite (resp. positive semi-definite) if $\langle u, Au\rangle$ is positive (resp. non-negative) for all $u \neq 0$. This will sometimes be denoted by $A > 0$ (resp. $A \geq 0$). For $A > 0$, we denote by $|\cdot|_A$ the weighted norm defined by $|v|_A^2 = v^\top A^{-1} v$. The corresponding weighted Euclidean inner-product is given by $\langle \cdot\,,\ \cdot\rangle_A := \langle \cdot\,, A^{-1}\cdot\rangle$. We use $\otimes$ to denote the outer product between two vectors: $(a \otimes b)c = \langle b, c\rangle a$. We let $B(u, \delta)$ denote the open ball of radius δ at u, in the Euclidean norm. We also use det and Tr to denote the determinant and trace functions on matrices.

Throughout, we denote by $\mathbb{P}(\cdot), \mathbb{P}(\cdot \mid \cdot)$ the probability density function (pdf) of a random variable and its conditional pdf, respectively. We write

$$\rho(f) = \mathbb{E}^\rho[f] = \int_{\mathbb{R}^d} f(u)\rho(u)du$$

to denote expectation of $f : \mathbb{R}^d \mapsto \mathbb{R}$ with respect to pdf ρ on $\mathbb{R}^d$. The distribution of the random variables in these notes will often have density with respect to Lebesgue measure, but occasional use of Dirac masses will be required; we will use the notational convention that Dirac mass at point v has “density” $\delta(\cdot - v)$, also denoted by $\delta_v(\cdot)$. When a random variable u has pdf ρ we will write $u \sim \rho$. We use $\Rightarrow$ to denote weak convergence of probability measures; that is, $\rho_n \Rightarrow \rho$ if $\rho_n(f) \to \rho(f)$ for all bounded and continuous $f : \mathbb{R}^d \mapsto \mathbb{R}$.

© in this web service Cambridge University Press & Assessment www.cambridge.org