Cambridge University Press
978-0-521-79298-1 - Independent Component Analysis: Principles and Practice
Edited by Stephen Roberts and Richard Everson
Frontmatter
More information

Independent Component Analysis: Principles and Practice

Cambridge University Press
978-0-521-79298-1 - Independent Component Analysis: Principles and Practice
Edited by Stephen Roberts and Richard Everson
Frontmatter
More information

Independent Component Analysis: Principles and Practice

Edited by
Stephen Roberts
Richard Everson

CAMBRIDGE UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521792981

First published 2001

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

ISBN 978-0-521-79298-1 Hardback

Contents

Preface

In recent years there has been an explosion of interest in the application and theory of independent component analysis (ICA). This book is aimed to provide a self-contained introduction to the subject as well as offering a set of invited contributions which we see as lying at the cutting edge of ICA research.

ICA is intimately linked with the problem of blind source separation—attempting to recover a set of underlying *sources* when only a noisy mapping from these sources, the *observations*, is given—and we regard this as the canonical form of ICA. Until recently this mapping was taken to be linear (but see Chapter 4) and "traditionally" (if tradition is allowed in a field of such recent developments) noiseless with the number of observations being equal to the number of hypothesised sources. It is surprising that even the simplest of ICA models can be invaluable and offer new insights into data analysis and interpretation. This, at first sight unreasonable, claim may be supported by noting that many observations of physical systems *are* produced by a linear combination of underlying sources. Furthermore, in many applications, it is an end in itself to produce a set of "sources" which are statistically *independent* rather than just decorrelated (see Chapter 1) and for this ICA would appear an ideal tool.

This book was born from discussions with researchers in the ICA community and aims to provide a snapshot of some current trends in ICA research. Wherever possible, the contributors use a common nomenclature and symbol set, especially for the most frequently used terms. The book has a single *global* set of references and index items which may be found at the end of the book.

A web site dedicated to the book and containing links to author web pages and other useful sites along with code and data related to ICA may be found at `http://www.dcs.ex.ac.uk/ica`

Structure of the book

Chapter 1 offers an introduction to independent component analysis. This chapter aims to give the reader an accessible way into the techniques, issues and jargon of ICA. The field is an extensive one and we have attempted to keep to ideas which we regard as instructive in the key issues of ICA, rather than give a complete description of every sub-method and algorithm modification.

Chapter 2 details one of the most popular approaches to ICA based on polynomial approximations to the mutual information, that of *Fast ICA*. Hyvärinen details in this chapter the theoretical development of the approximations, their justification and the rapid fixed-point algorithm by which the sources are recovered. Background material on the relationship between learning algorithms is also presented along with results on a number of datasets.

Chapter 3 pitches ICA into the important context of graphical models, whereby the relationships between model parameters are represented by a directed acyclic graph. Attias, in this chapter, considers flexible ICA methods (sometimes known as Independent Factor Analysis, IFA) in which the source densities are modelled by mixtures of Gaussians and an explicit additive (sensor) noise term exists. Inference is performed using a *variational learning* approach (see also Chapter 8).

Chapter 4 extends ICA from a general *linear* model of source mixing to the nonlinear case. Karhunen explores in detail the issues involved with forming learning paradigms for such nonlinear ICA and develops promising algorithms to deal with nonlinear mixing. The chapter is illustrated with comparative examples.

Chapter 5 extends ICA to consider the issue of source non-stationarity. Parra and Spence show how the higher order statistics used to locate independent components arise naturally in non-stationary signals. They examine whether linear mixing is a good model for acoustic signals and natural images. Exploiting the property of non-stationarity they show how good blind separation may be achieved using multiple linear decorrelation.

Chapter 6 offers a different perspective on the separation of non-stationary sources. Cardoso and Pham derive an elegant methodology in which non-stationarity may be handled and indeed utilised to aid in the unmixing process.

Chapter 7 considers source separation in the case when the sources are represented by a *sparse* mixture from a signal dictionary (such as wavelet packets). Under these circumstances ICA naturally seeks sources which are as sparse in their representation as possible. This extra information enables Zibulevsky, Pearlmutter, Bofill and Kisilev to obtain impressive results in situations when there are *more* sources than observations.

Chapter 8 pitches ICA as a graphical model with densities over variables being inferred using a variational learning framework. As both *parameters* and *hyper-parameters* of the model are inferred as part of a single learning strategy, this approach is referred to as *ensemble* learning. Miskin and MacKay consider mixture of Gaussian source models and show results from model-selection on real-world problems. They also show that a positivity constraint on the hypothesised mixing process gives rise to ICA solutions which are more local in their support.

Chapter 9 applies ICA to the domain of image decomposition and processing. By assuming that natural images are generated by a linear combination of independent sources (textures and edges for example) ICA may be used to estimate, given an image, a basis for decomposition. Lee and Lewicki show that this basis has an intuitively appealing form (typically that of local filters) and show how ICA may thence be utilised to perform image denoising.

Chapter 10 regards ICA as a general linear transform of the same form as the linear discriminant of pattern classification. Using a nested hierarchy of ICA decompositions of real data, Girolami shows that excellent results may be obtained in difficult unsupervised classification problems. He then proceeds to consider the process of visualisation of high-dimensional data (mapping to a two-dimensional space, for example) as an ICA-like procedure for which learning rules may be obtained.

Chapter 11 allows a model in which the mixing matrix of a linear ICA model is considered to be *non-stationary*. This matrix may thence be tracked using a particle filter. The approach is shown to be very effective when tracking non-stationary mixing of temporally uncorrelated sources. Some solutions to the more difficult problem of tracking mixtures of temporally correlated sources are presented.

Chapter 12 extends the standard ICA model by allowing the source density models to be *dynamic* rather than static. This is achieved by the use of linear dynamic models within the ICA framework. The authors also consider the important issue of model order selection to determine the most probable number of underlying sources. Results are presented for a variety of datasets.

Acknowledgements

There are many people we would like to thank for their contributions to this book. In particular we thank Iead Rezek, Dirk Husmeier, Peter Sykacek, Will Penny and Riz Choudrey for interesting and enlightening conversations. We are especially grateful to Will Penny for a careful reading of the entire book as it neared completion. We would also like to thank David Tranah, our editor at CUP, for his unfailing support for the project, even as the deadline approached and went.

The editors would like to express thanks to the other contributors for their hard work and their patience and help during the editing process.

SR would like to say a big thanks, with much love, to Anna, Rebecca and Elise.

RE thanks Frances, Gwyneth and Juliet for their love and endurance.

Stephen Roberts
Richard Everson

Contributors

Hagai Attias
hagaia@microsoft.com
Microsoft Research 113/3359,
One Microsoft Way,
Redmond, WA 98052-6399,
USA.

Pau Bofill
pau@ac.upc.es
Departament d'Arquitectura de
Computadors,
Universitat Politècnica de Catalunya,
Campus Nord,
Sergi Girona s/n,
08071 Barcelona,
Spain.

Jean-Francois Cardoso
cardoso@tsi.enst.fr
CNRS and ENST,
ENST, Dept TSI,
46, rue Barrault,
75634 Paris CEDEX 13,
France.

Richard Everson
reverson@exeter.ac.uk
Department of Computer Science,
University of Exeter,
Exeter EX4 4PT,
UK.

Mark Girolami
mark.girolami@paisley.ac.uk
Department of Computing and
Information Systems,
University of Paisley,
High Street,
Paisley, PA1 2BE,
Scotland, UK.

Aapo Hyvärinen
Aapo.Hyvarinen@hut.fi
Neural Networks Research Centre,
P.O. Box 5400,
Helsinki University of Technology,
FIN-02015 TKK,
Finland.

Juha Karhunen
Juha.Karhunen@hut.fi
Helsinki University of Technology,
Laboratory of Computer and
Information Science,
P.O. Box 5400,
FIN-02015 HUT,
Finland.

Pavel Kisilev
paulk@tx.technion.ac.il
Electrical Engineering Department,
Technion Israel Institute of
Technology,
Haifa 32000, Israel.

Te-Won Lee
tewon@salk.edu
INC, UC San Diego,
L a Jolla,
CA 9093-0523,
USA.

Michael S. Lewicki
lewicki@cnbc.cmu.edu
CMU/CNBC,
Mellon Inst. 115,
4400 Fifth Avenue,
Pittsburgh, PA 15213,
USA.

David MacKay
mackay@mrao.cam.ac.uk
Cavendish Laboratory,
Madingley Road,
Cambridge, CB3 0HE,
UK.

James Miskin
jwm1003@mrao.cam.ac.uk
Cavendish Laboratory,
Madingley Road,
Cambridge, CB3 0HE,
UK.

Lucas Parra
lparra@sarnoff.com
Sarnoff Corporation, CN5300,
Princeton, NJ, 08543-5300,
USA.

Barak Pearlmutter
bap@cs.unm.edu
Computer Science Department,
FEC 313,
University of New Mexico,
Albuquerque, NM 87131,
USA.

William Penny
wpenny@robots.ox.ac.uk
Robotics Research Group,
Department of Engineering Science,
University of Oxford,
Oxford OX1 3PJ,
UK.

Dinh-Tuan Pham
Dinh-Tuan.Pham@imag.fr
CNRS IMAG-LMC,
51 rue des Mathématiques,
B. P. 53,
38041 Grenoble Cedex 9,
France.

Stephen Roberts
sjrob@robots.ox.ac.uk
Robotics Research Group,
Department of Engineering Science,
University of Oxford,
Oxford OX1 3PJ,
UK.

Clay Spence
cspence@sarnoff.com
Sarnoff Corporation, CN5300,
Princeton, NJ, 08543-5300,
USA.

Michael Zibulevsky
michael@cs.unm.edu
Electrical Engineering Department,
Technion Israel Institute of
Technology,
Haifa 32000,
Israel.