

Lecture Notes in Artificial Intelligence 7263

Subseries of Lecture Notes in Computer Science

LNAI Series Editors

Randy Goebel

University of Alberta, Edmonton, Canada

Yuzuru Tanaka

Hokkaido University, Sapporo, Japan

Wolfgang Wahlster

DFKI and Saarland University, Saarbrücken, Germany

LNAI Founding Series Editor

Joerg Siekmann

DFKI and Saarland University, Saarbrücken, Germany

Georg Langs Irina Rish
Moritz Grosse-Wentrup Brian Murphy (Eds.)

Machine Learning and Interpretation in Neuroimaging

International Workshop, MLINI 2011

Held at NIPS 2011

Sierra Nevada, Spain, December 16-17, 2011

Revised Selected and Invited Contributions



Springer

Series Editors

Randy Goebel, University of Alberta, Edmonton, Canada
Jörg Siekmann, University of Saarland, Saarbrücken, Germany
Wolfgang Wahlster, DFKI and University of Saarland, Saarbrücken, Germany

Volume Editors

Georg Langs
Medical University of Vienna, Department of Radiology, CIR Lab
Währinger Gürtel 18-20, 1090 Wien, Austria
E-mail: georg.langs@meduniwien.ac.at

Irina Rish
IBM T.J. Watson Research Center, Computational Biology Center
1101 Kitchawan Road, Yorktown Heights, NY 10598, USA
E-mail: rish@us.ibm.com

Moritz Grosse-Wentrup
Max Planck Institute for Intelligent Systems
Spemannstraße 38, 72076 Tübingen, Germany
E-mail: moritzgw@tuebingen.mpg.de

Brian Murphy
Carnegie Mellon University, Machine Learning Department
5000 Forbes Avenue, Pittsburgh, PA 15213-3891, USA
E-mail: brianmurphy@cmu.edu

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-34712-2 e-ISBN 978-3-642-34713-9
DOI 10.1007/978-3-642-34713-9
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012951409

CR Subject Classification (1998): I.5.1-4, I.2.6, H.2.8, G.3, I.4.6-7, I.4.9, J.3,
I.2.1, F.2.2

LNCS Sublibrary: SL 7 – Artificial Intelligence

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

Brain imaging brings together the technology, methodology, research questions, and approaches of a wide range of scientific fields including physics, statistics, computer science, neuroscience, biology, and engineering. Thus, methodological and technological advances that enable us to obtain measurements, examine relationships across observations, and link these data to neuroscientific hypotheses happen in a highly interdisciplinary environment. Open questions in neuroscience often trigger methodological development, yet original methods can also spur novel perspectives for posing and answering questions when studying the brain. We believe the dynamic field of machine learning with its modern approach to data mining provides many relevant approaches for neuroscience, and enables the exploration of open questions.

In December 2011, we organized a workshop to explore the interface between machine learning and neuroimaging, and how this relationship affects the progress of research, the formulation of novel questions, and the recognition and tackling of big open issues in the field of neuroscience. In order to start a discussion among the involved communities, we invited experts from machine learning, biology, neuroscience, and neuroimaging, to share their views on questions they considered most exciting and important. Before the workshop, we asked all participants to contribute questions, in order to assess the spectrum and relevance of topics. Many replied, and we set out to explore the most pressing issues during two panel discussions that involved all invited speakers and a vocal audience. There were two general themes of discussion. The first focused on the following question: *how can we interpret findings that are obtained with multivariate pattern analysis (MVPA) approaches, in the context of neuroscientific questions we seek to answer*. The second general theme focused on *the shift and divergence of paradigms*, which have emerged while the field has moved on from exclusively univariate approaches. As an introduction to this volume we briefly summarize these two discussions.

The Interpretation of MVPA Findings

How can sophisticated methods be made more relevant and accessible?

Multivariate models are, by construction, difficult to study and visualize since they are based on patterns that span the image and are not localized. Non-linear models, such as those used in kernel-based methods, are even harder to characterize since they cannot be represented with a single discriminative map. Thus, interpreting findings made with multivariate-, or other machine-learning approaches, is not straightforward. In studying multivariate models, our goal is to seek answers to the questions such as the following, whether implicitly or explicitly: What is the link between measurements and physiology? What can we

say about cognitive processes, and their relationships among each other? What is the relationship between observations and experiment conditions?

Building on massive univariate approaches, such as those based on the General Linear Model, where every voxel is independently probed for its relationship to a task, many early MVPA methods probed local patterns in a search light style in order to check the data's capability to differentiate between experimental conditions or stimuli. This *encoding-decoding* approach enables the observer to ask not only for the relationship of neuroimaging data to experimental conditions, but also for the relationship *among* experimental conditions, and their shared functional structures. This represents a fundamental breakthrough, since we can now study the internal structure and relationships among tasks, and move toward an understanding of how this functional structure is formed and embedded in the space of anatomy.

Moving beyond local neighborhoods, approaches such as ensemble learning, multivariate regression, or manifold learning typically view the brain as a global pattern or connectivity structure. While this makes physiological interpretation more complex, it enables us to capture distributed processes.

For many methods our understanding of their statistical properties is limited. The common approach to quantify the model fit in MVPA methods is via metrics such as area under the curve, average accuracy, and mean square error obtained from cross-validation. However, we might also be interested in other statistical quantities: how can we assign confidence intervals and statistical significance to the boundaries of the regions we detect, to our estimates of prediction accuracy, and the relationship of both to the experimental conditions; and how confident can we be that our results generalize beyond study populations. What are the methods that achieve statistical interpretability? A rigorous statistical framework to draw neuroscientific or clinical conclusions from observations is essential for their proper adoption. This responsibility is particularly pressing once published results are picked up and form bases for clinical decisions.

Instead of pushing for a small unified and well-understood set of tools that can be used by the neuroscience community, participants suggested that there is a constant dialog among practitioners and method developers. By adopting this dynamic approach, neuroscientists can focus on asking relevant questions, while receiving help on choosing the right tool. Furthermore, they can understand what kind of new questions they can ask if the machine learning community provides novel approaches. Thus, instead of a gap between machine learning and neuroscience there should be a relaxed and focused communication. Instead of consolidation, the process of methodology development and scientific inquiry should progress as a feedback loop, in which one fosters the other.

Divergence of Paradigms

Can MVPA methods help us move beyond simple contrast-based studies?

Multivariate encoding-decoding schemes were originally developed as alternative strategies to analyze neuroimaging data within the boundaries of

traditional experimental paradigms. Yet, MVPA methodology inspired by these early effort has come to free us from the constraints of simple experimental designs by enabling us to ask new and different questions from neuroimaging data. The divergence of methodology, workshop participants observed, helped us move beyond simple contrast-based studies. Today, researchers can choose from a wide array of supervised, unsupervised, or semisupervised multivariate methods to analyze their imaging data, in order to identify structure in neuroimaging data such as resting state fMRI or characterize the space of stimuli, for example, by identifying semantic structure among visual or auditory stimuli. The discussion did not lead to an ultimate consensus regarding a consolidated set of paradigms. Yet the participants agreed that the richness in methodology would continue to feed the divergence of paradigms in neuroscientific research.

Remaining Questions

Many important questions remained unaddressed during the discussion. These include, but are not limited to, the following. Similar to the mass-univariate GLM-based approach, can we develop general MVPA methods that might be specialized for specific situations? Can the machine learning community agree on a few established problems to work on, knowing that they will stay relevant even if particular neuroscientific questions change? How can we choose between alternative models? What are the advantages and disadvantages of generative versus discriminative models? Is there a unified framework for performing brain mapping based on MVPA methodology? We hope that these and many other questions will be explored in future incarnations of this workshop.

In this volume we collect contributions from the MLINI Workshop at the Conference on Neural Information Processing (NIPS 2011). These works aim to shed some light on the state of the art of this interdisciplinary field that involves both the machine learning and neuroimaging communities. The papers underwent a thorough review process, and from an initial 48 submissions, 32 papers were selected for inclusion in the proceedings. Additionally, invited speakers agreed to contribute reviews on various aspects of the field, adding breadth and perspective to the volume.

December 2011

Georg Langs
 Irina Rish
 Moritz Grosse-Wentrup
 Brain Murphy
 Bjoern Menze
 Mert Sabuncu

Organization

MLINI 2011 was organized during the NIPS 2011 Conference at Granada.

Workshop Chairs

Georg Langs	CIR Dep. Radiology, MUW, Austria; CSAIL, MIT, USA
Irina Rish	IBM, USA
Moritz Grosse-Wentrup	MPI for Intelligent Systems, Germany
Brian Murphy	Carnegie Mellon University, USA

Organizers

Melissa Carroll	Google, New York, USA
Guillermo Cecchi	IBM T.J. Watson Research Center, USA
Kai-min Kevin Chang	LTI and CNBC, Carnegie Mellon University, USA
Moritz Grosse-Wentrup	Max Planck Institute for Intelligent Systems, Tübingen, Germany
James V. Haxby	Dartmouth College, USA
Anna Korhonen	Computer Laboratory and RCEAL, University of Cambridge, UK
Georg Langs	CIR Dep. Radiology, MUW, Austria; CSAIL, MIT, USA
Bjoern Menze	ETH Zürich, Switzerland; CSAIL, MIT, USA
Brian Murphy	Machine Learning Dept., Carnegie Mellon University, USA
Janaina Mourao-Miranda	University College London, UK
Vittorio Murino	University of Verona/Istituto Italiano di Tecnologia, Italy
Francisco Pereira	Princeton University, USA
Irina Rish	IBM T.J. Watson Research Center, USA
Mert Sabuncu	Harvard Medical School, USA
Irina Simanova	Max Planck Psycholinguistics and Donders Institute, Nijmegen, The Netherlands
Bertrand Thirion	INRIA, Neurospin, France

Program Committee

Yi Chen	Bernstein Center and Charité Universitätsmedizin, Berlin, Germany
Andy Connolly	HaxbyLab, Dartmouth University, USA
Scott Fairhall	Centre for Mind/Brain Sciences, University of Trento, Italy
Swaroop Guntupalli	HaxbyLab, Dartmouth University, USA
Yaroslav Halchenko	HaxbyLab, Dartmouth University, USA
Michael Hanke	Experimental Psychology, University of Magdeburg, Germany
Marius Peelen	Centre for Mind/Brain Sciences, University of Trento, Italy
Diego Sona	FBK/CIMeC Neuroinformatics Lab, Trento, Italy
Marcel van Gerven	Donders Institute and Radboud University, Nijmegen, The Netherlands
Gaël Varoquaux	INRIA, Neurospin, France
John Anderson	Carnegie Mellon University, USA
Mark Cohen	University of California, Los Angeles, USA
Kevyn Collins-Thompson	Microsoft Research, USA
Jack Gallant	University of California, Berkeley, USA
Tom Heskes	Radboud University Nijmegen, The Netherlands
Mark Johnson	Macquarie University, Australia
Russ Poldrack	University of Texas, Austin, USA
Dean Pomerleau	Intel Labs, Pittsburgh, USA

Table of Contents

Coding and Decoding

A Comparative Study of Algorithms for Intra- and Inter-subjects fMRI Decoding	1
<i>Vincent Michel, Alexandre Gramfort, Evelyn Eger, Gaël Varoquaux, and Bertrand Thirion</i>	
Beyond Brain Reading: Randomized Sparsity and Clustering to Simultaneously Predict and Identify	9
<i>Alexandre Gramfort, Gaël Varoquaux, and Bertrand Thirion</i>	
Searchlight Based Feature Extraction	17
<i>Shahar Jamshy, Omri Perez, Yehezkel Yeshurun, Talma Hendler, and Nathan Intrator</i>	
Looking Outside the Searchlight	26
<i>Josef A. Etzel, Michael W. Cole, and Todd S. Braver</i>	
Population Codes Representing Musical Timbre for High-Level fMRI Categorization of Music Genres	34
<i>Michael Casey, Jessica Thompson, Olivia Kang, Rajeev Raizada, and Thalia Wheatley</i>	
Induction in Neuroscience with Classification: Issues and Solutions	42
<i>Emanuele Olivetti, Susanne Greiner, and Paolo Avesani</i>	
A New Feature Selection Method Based on Stability Theory – Exploring Parameters Space to Evaluate Classification Accuracy in Neuroimaging Data	51
<i>Jane M. Rondina, John Shawe-Taylor, and Janaina Mourão-Miranda</i>	
Identification of OCD-Relevant Brain Areas through Multivariate Feature Selection	60
<i>Emilio Parrado-Hernández, Vanessa Gómez-Verdejo, Manel Martínez-Ramón, Pino Alonso, Jesús Pujol, José M. Menchón, Narcís Cardoner, and Carles Soriano-Mas</i>	
Deformation-Invariant Sparse Coding for Modeling Spatial Variability of Functional Patterns in the Brain	68
<i>George H. Chen, Evelina G. Fedorenko, Nancy G. Kanwisher, and Polina Golland</i>	

Decoding Complex Cognitive States Online by Manifold Regularization in Real-Time fMRI	76
<i>Toke Jansen Hansen, Lars Kai Hansen, and Kristoffer Hougaard Madsen</i>	

Neuroscience

Modality Neutral Techniques for Brain Image Understanding	84
<i>David B. Keator</i>	
How Does the Brain Represent Visual Scenes? A Neuromagnetic Scene Categorization Study	93
<i>Pavan Ramkumar, Sebastian Pannasch, Bruce C. Hansen, Adam M. Larson, and Lester C. Loschky</i>	
Finding Consistencies in MEG Responses to Repeated Natural Speech.....	101
<i>Miika Koskinen</i>	
Categorized EEG Neurofeedback Performance Unveils Simultaneous fMRI Deep Brain Activation	108
<i>Sivan Kinreich, Ilana Podlipsky, Nathan Intrator, and Talma Hendler</i>	
Predicting Clinically Definite Multiple Sclerosis from Onset Using SVM	116
<i>Philip P. Kwok, Olga Ciccarelli, Declan T. Chard, David H. Miller, and Daniel C. Alexander</i>	
MKL-Based Sample Enrichment and Customized Outcomes Enable Smaller AD Clinical Trials	124
<i>Chris Hinrichs, N. Maritza Dowling, Sterling C. Johnson, and Vikas Singh</i>	
Pairwise Analysis for Longitudinal fMRI Studies	132
<i>Diego Sona, Paolo Avesani, Stefano Magon, Gianpaolo Basso, and Gabriele Miceli</i>	

Dynamics

Non-separable Spatiotemporal Brain Hemodynamics Contain Neural Information	140
<i>Felix Bießmann, Yusuke Murayama, Nikos K. Logothetis, Klaus-Robert Müller, and Frank C. Meinecke</i>	
The Dynamic Beamformer	148
<i>Ali Bahramisharif, Marcel A.J. van Gerven, Jan-Mathijs Schoffelen, Zoubin Ghahramani, and Tom Heskes</i>	

Covert Attention as a Paradigm for Subject-Independent Brain-Computer Interfacing	156
<i>Hans J.P. Wouters, Marcel A.J. van Gerven, Matthias S. Treder, Tom Heskes, and Ali Bahramisharif</i>	
The Neural Dynamics of Visual Processing in Monkey Extrastriate Cortex: A Comparison between Univariate and Multivariate Techniques	164
<i>Marime Cauchoix, Ali Bilgin Arslan, Denis Fize, and Thomas Serre</i>	
Statistical Learning for Resting-State fMRI: Successes and Challenges	172
<i>Gaël Varoquaux and Bertrand Thirion</i>	
Relating Brain Functional Connectivity to Anatomical Connections: Model Selection	178
<i>Fani Deligianni, Gaël Varoquaux, Bertrand Thirion, Emma Robinson, David J. Sharp, A. David Edwards, and Daniel Rueckert</i>	
Information-Theoretic Connectivity-Based Cortex Parcellation	186
<i>Nico S. Gorbach, Silvan Step, Jenia Jitsev, Corina Melzer, and Marc Tittgemeyer</i>	
Inferring Brain Networks through Graphical Models with Hidden Variables	194
<i>Justin Dauwels, Hang Yu, Xueou Wang, Francois Vialatte, Charles Latchoumane, Jaeseung Jeong, and Andrzej Cichocki</i>	
Pitfalls in EEG-Based Brain Effective Connectivity Analysis	202
<i>Stefan Haufe, Vadim V. Nikulin, Guido Nolte, and Klaus-Robert Müller</i>	

Probabilistic Models and Machine Learning

Data-Driven Modeling of BOLD Drug Response Curves Using Gaussian Process Learning	210
<i>Orla M. Doyle, Mitul A. Mehta, Michael J. Brammer, Adam J. Schwarz, Sara De Simoni, and Andre F. Marquand</i>	
Variational Bayesian Learning of Sparse Representations and Its Application in Functional Neuroimaging	218
<i>Evangelos Roussos, Steven Roberts, and Ingrid Daubechies</i>	
Identification of Functional Clusters in the Striatum Using Infinite Relational Modeling	226
<i>Kasper Winther Andersen, Kristoffer Hougaard Madsen, Hartwig Siebner, Lars Kai Hansen, and Morten Mørup</i>	

A Latent Feature Analysis of the Neural Representation of Conceptual Knowledge 234
 Kai-min Chang, Brian Murphy, and Marcel Just

Real-Time Functional MRI Classification of Brain States Using Markov-SVM Hybrid Models: Peering Inside the rt-fMRI Black Box 242
 Ariana Anderson, Dianna Han, Pamela K. Douglas, Jennifer Bramen, and Mark S. Cohen

Restoring the Generalizability of SVM Based Decoding in High Dimensional Neuroimage Data 256
 Trine Julie Abrahamsen and Lars Kai Hansen

Author Index 265